

論文 / 著書情報
Article / Book Information

題目(和文)	コーパスに基づく多義性解消
Title(English)	Corpus-Based Word Sense Disambiguation
著者(和文)	藤井 敦
Author(English)	Atsushi Fujii
出典(和文)	学位:博士(工学), 学位授与機関:東京工業大学, 報告番号:甲第3877号, 授与年月日:1998年3月26日, 学位の種別:課程博士, 審査員:
Citation(English)	Degree:Doctor (Engineering), Conferring organization: Tokyo Institute of Technology, Report number:甲第3877号, Conferred date:1998/3/26, Degree Type:Course doctor, Examiner:
学位種別(和文)	博士論文
Type(English)	Doctoral Thesis

Corpus-Based Word Sense Disambiguation

Atsushi Fujii

Ph.D. thesis

Department of Computer Science

Tokyo Institute of Technology

March 1998

Supervisor : Hozumi Tanaka

Preface

Resolution of lexical ambiguity, commonly termed “word sense disambiguation”, is expected to improve the analytical accuracy for tasks which are sensitive to lexical semantics. Such tasks include machine translation, information retrieval, parsing, natural language understanding and lexicography. Reflecting the growth in utilization of machine readable texts, word sense disambiguation techniques have been explored variously in the context of corpus-based approaches. Within one corpus-based framework, that is the similarity-based method, systems use a database, in which example sentences are manually annotated with correct word senses. Given an input, systems search the database for the most similar example to the input. The lexical ambiguity of a word contained in the input is resolved by selecting the sense annotation of the retrieved example.

In this research, we apply this method of resolution of verbal polysemy, in which the similarity between two examples is computed as the weighted average of the similarity between complements governed by a target polysemous verb. We explore similarity-based verb sense disambiguation focusing on the following three methods. First, we propose a weighting schema for each verb complement in the similarity computation. Second, in similarity-based techniques, the overhead for manual supervision and searching the large-sized database can be prohibitive. To resolve this problem, we propose a method to select a small number of effective examples, for system usage. Finally, the efficiency of our system is highly dependent on the similarity computation used. To maximize efficiency, we propose a method which integrates the advantages of previous methods for similarity computation.

Acknowledgments

First and foremost, I would like to thank my supervisor, Hozumi Tanaka (TITECH, Japan) for his guidance, support and encouragement throughout the years of my studentship. I would also like to thank Takenobu Tokunaga (TITECH, Japan) for his enormous support with this research, and other members of my thesis committee: Sadaoki Furui, Hiroshi Maruyama and Masayuki Numao (TITECH, Japan).

Additionally, I would like to express my gratitude to members of the Tanaka & Tokunaga laboratories (TITECH, Japan) for their contributions to this research, amongst whom, I can never thank Kentaro Inui and Timothy Baldwin enough.

Manabu Okumura (JAIST, Japan), Michael Zock (LIMSI, France) and Dan Tufis (Romanian Academy, Romania) provided invaluable comments on this research, for which I am greatly indebted.

Last but not least, I would also like to thank Masayuki Kameda (RICOH Co., Ltd., Japan) for his support with the QJP parser, and Wakako Kashino and Minako Hashimoto (Information-technology Promotion Agency, Japan) for their support with the corpus used in this research.

This research was partially supported by a Research Fellowship of the Japan Society for the Promotion of Science for Young Scientists.

Contents

Preface	i
Acknowledgments	iii
1 Introduction	1
1.1 Background	1
1.2 Focus of this Research	3
1.3 Outline of the Proceeding Chapters	5
2 Past Research on Word Sense Disambiguation	7
2.1 Overview: Terminology and Task Description	7
2.2 Different Methodologies for Word Sense Disambiguation	8
2.2.1 General classification of methodologies	9
2.2.2 Restricting supervision	22
2.3 Evaluation Methodology	26
2.3.1 Common evaluation criteria	27
2.3.2 Multiple human judgement	29
2.3.3 Empirical comparison	30
2.3.4 Further considerations	32
2.4 Related NLP Research	33
2.5 Discussion	35
2.6 Summary	36
3 Verb Sense Disambiguation System	39
3.1 Overall Design	39
3.2 Basic Idea	41
3.3 Methodology	45
3.3.1 Computation of similarity between case filler nouns	46

3.3.2	Computation of case contribution factor	48
3.4	Experimentation	49
3.5	Further Enhancement	55
3.5.1	Computation of interpretation certainty	55
3.5.2	Incorporation of contextual constraints	56
3.5.3	Remaining problems	59
3.6	Summary	60
4	Selective Sampling of Effective Example Sets	63
4.1	Motivation	63
4.2	Methodology	66
4.2.1	Overview	66
4.2.2	Computation of training utility	68
4.2.3	Enhancement of computation	70
4.3	Experimentation	71
4.4	Related Work	73
4.4.1	Uncertainty sampling	73
4.4.2	Committee-based sampling	75
4.5	Discussion	77
4.5.1	Sense ambiguity of case fillers in selective sampling	77
4.5.2	Applicability of our selective sampling method	77
4.6	Summary	79
5	Exploration of Similarity Computation	81
5.1	Motivation	81
5.2	Methodology	84
5.2.1	Overview	84
5.2.2	Resolution of the simultaneous equations	84
5.2.3	Word similarity using SBL	86
5.3	Experimentation	86
5.4	Application to verb sense disambiguation system	87
5.5	Related Work	88
5.6	Summary	89
6	Conclusion	91
6.1	Contribution	91

<i>CONTENTS</i>	vii
6.2 Outstanding Issues	92
Bibliography	95

List of Tables

- 3.1 The relation between the length of the path between two nouns n_1 and n_2 in the *Bunruigoihyo* thesaurus ($len(n_1, n_2)$), and their relative similarity ($sim(n_1, n_2)$) . . . 47
- 3.2 The verbs contained in the corpus used, and the accuracy of the different verb sense disambiguation methods (LB: lower bound, RB: rule-based method, NB: Naive-Bayes method, VSM: vector space model, BGH: the *Bunruigoihyo* thesaurus) 51
- 3.3 The frequency of nouns and resultant accuracy of verb sense disambiguation . . . 54

List of Figures

2.1	Classification of different methodologies of word sense disambiguation	9
2.2	A fragment of the decision tree for the English verb <i>take</i>	12
2.3	A fragment of the decision list for the word <i>plant</i>	13
2.4	An example rule set for CN2	14
2.5	Example salient words for two categories	16
2.6	Example indicator nouns for the adjective <i>old</i>	16
2.7	The word and context vectors for two contexts	20
2.8	Contingency table	29
2.9	A fragment of the EDR thesaurus including senses related to the Japanese verb <i>tsukau</i>	32
3.1	The overall design of the verb sense disambiguation system	40
3.2	A fragment of the database, and the entry associated with the Japanese verb <i>toru</i>	41
3.3	An example of a subentry for the Japanese verb <i>kakeru</i>	42
3.4	The hierarchy of IPAL semantic features	43
3.5	The semantic ranges of the nominative and accusative for the verb <i>toru</i>	44
3.6	An input and the database	45
3.7	A fragment of the <i>Bunruigoihyo</i> thesaurus	47
3.8	The relation between the number of training data used and accuracy of the dif- ferent methods	53
3.9-a	The case where the interpretation certainty of the enclosed x 's is great	56
3.9-b	The case where the interpretation certainty of the x 's contained in the intersection of senses 1 and 2 is small	56
3.9	The concept of interpretation certainty	56
3.10	The relation between coverage and accuracy with different λ 's	57
4.1	Flow of control of the selective sampling	66

4.2	Example of a given corpus associated with the verb <i>yameru</i>	68
4.3-a	The case where the training utility of <i>a</i> is greater than that of <i>b</i> because <i>a</i> has more unsupervised neighbors	69
4.3-b	The case where the training utility of <i>a</i> is greater than that of <i>b</i> because <i>b</i> closely neighbors <i>e</i> , contained in the database	69
4.3	The concept of training utility	69
4.4	A fragment of the thesaurus including neighbors of <i>x</i> associated with case <i>c</i> . . .	71
4.5	The relation between the number of training data sampled and accuracy of the system	73
4.6	The relation between the number of training data sampled and accuracy of the system without using examples from IPAL	74
4.7	A case where either <i>x</i> or <i>y</i> can be selected in committee-based sampling	76
4.8-a	Interpretation certainty of <i>x</i> is small because <i>x</i> lies in the intersection of distinct verb senses	77
4.8-b	Interpretation certainty of <i>y</i> is small because <i>y</i> is semantically ambiguous	77
4.8	Two separate scenarios in which the interpretation certainty of <i>x</i> is small	77
4.9	The case where informative example <i>x</i> is not selected	78
5.1	A fragment of the thesaurus associated with x_1 to x_6	83
5.2	A fragment of the set of simultaneous equations associated with Figure 5.1	84
5.3	Approximation of the statistics-based length x_i	86
5.4	The relation between the number of equations used and the accuracy	87

Chapter 1

Introduction

1.1 Background

Natural language processing (NLP) involves resolution of various types of *ambiguity*. Lexical ambiguity is one of these ambiguity types, and occurs when a single word (lexical form) is associated with multiple senses or meanings. For applications which are sensitive to semantic denotation, or more precisely lexical semantics, this ambiguity type can pose a major obstacle. Resolution of lexical ambiguity, which is commonly termed “word sense disambiguation” (WSD), is expected to improve the quality of the following research fields.

- Machine translation (MT) can safely be identified as one of the major beneficiaries of word sense disambiguation, because a single word in a source language is frequently associated with multiple translations in a target language, each of which is often associated with a different sense to the source word. For example, the “tax” and “obligation” senses of the English word *duty* correspond to the French translations of *droit* and *devoir*, respectively. Sense disambiguation of *duty* is expected to allow this word to be translated appropriately for a given context. In fact, a number of MT-oriented word sense disambiguation methods have been explored based on this notion [10, 28].
- Information retrieval (IR) and text categorization (TC) suffer from the effects of noisy words associated with multiple senses, and IR/TC systems can easily end up relating documents containing the same words but in different senses (usages). For example, documents containing the word *AIDS* can easily be associated with those containing the word *aids*. Conventional systems have tentatively avoided this problem through usage of inverse document frequency (IDF) [127]. The rationale behind IDF is that words which rarely occur over document collections are valuable, or that in other words, the IDF of

a word is inversely proportional to the number of documents containing that word. As Krovets [74] identified, this can be recast in the context of lexical ambiguity. That is, words of higher frequency also tend to be associated with a greater number of senses, and therefore the degree to which these noisy words affect the system output can be minimized by introducing IDF. However, as a number of experimental results have shown/suggested, word sense disambiguation is a crucial task for the further improvement of IR/TC systems [40, 74, 75, 128, 129, 131].

- Syntactic analysis (or parsing) often fails to identify the correct syntactic structure for an input sentence when syntactic relations are associated with semantic content. Prepositional phrase (PP) attachment problems¹ and predicate-argument structures associated with selectional restrictions are immediate examples of this problem type, in that they require the intervention of the semantic content of lexical entries for knowledge representation. Given the fact that syntactic and semantic (lexical) ambiguity are not independent of each other, a number of methods have proposed the mutual resolution of these two ambiguity types [84, 87, 98].
- So called class-based NLP approaches [66, 120] are also potential beneficiaries of word sense disambiguation². These approaches involve mapping of each word entry to a semantic class (usually taken from a thesaurus taxonomy). Consequently, disambiguation of word senses is poignant, because each word is often associated with multiple class candidates, which are closely related to the senses of that word.
- In natural language understanding (NLU), semantic structures are constructed by considering the meaning of each word [17]. Kilgarriff [68] argued that current practical NLU systems, such as dialogue and information extraction (IE) systems, have commonly employed domain-specific knowledge representation rather than word sense disambiguation, in order to counter lexical ambiguity. However, we would like to note that his observation does not immediately reject the possibility of the potential contribution of word sense disambiguation to NLU systems.
- Kilgarriff [68] also points to the advantages of word sense disambiguation in lexicography. By this is meant that sense-annotated linguistic data reduces the considerable overhead imposed on lexicographers in sorting large-scaled corpora according to word usage for

¹Ravin [118], for example, focused on the resolution of PP-attachment problem through word sense disambiguation.

²A number of methods described in the previous item (“parsing”) can also be seen as instances of this category.

different senses. Not only lexicography, but the general process of compiling linguistic resources is expected to improve through the interaction between lexicographers and computers [5, 146]. In addition, word sense disambiguation techniques can also allow language learners to access example sentences containing a certain word usage from large corpora, without excessive overhead.

It should be noted that past word sense disambiguation methods have not contributed to all the research given above. In fact, no quantitative data has been forthcoming documenting empirical improvement through word sense disambiguation, except for in the fields of MT [10] and IR/TC [40]. One reason for this past failure to enhance existing methods is presumably due to the immaturity of word sense disambiguation research, which stimulates us to further explore this exciting research area.

Reflecting the growth in utilization of machine readable texts, word sense disambiguation techniques have been explored variously in the context of “corpus-based NLP approaches”³. These methods generally use a corpus in which component words of each example sentence are annotated (either manually or automatically) with their correct word sense, to automatically induce rules or probabilistic models for disambiguation. Unlike conventional rule-based approaches relying on *hand-crafted* selectional rules (some of which are reviewed, for example, by Hirst [54] and Small et al. [135]), corpus-based approaches release us from the task of generalizing observed phenomena through a set of rules. While certain methods require manual annotation of the given corpora (namely “supervised methods”)⁴, other methods exclude or minimize the overhead for supervision (namely “unsupervised methods”). However, we observe that the applicability of unsupervised methods has so far been limited to relatively specific applications, and that the overhead for supervision still remains as a major drawback of corpus-based word sense disambiguation (we will elaborate on this issue in Chapter 2).

1.2 Focus of this Research

First, let us precisely state the focus of this research, i.e. which subcategory of word sense disambiguation we are targetting, given the considerable variation in types and associated methods of lexical disambiguation (between noun and adjective senses, for example). At the same time, we

³Corpus-based approaches have been explored in terms of other types of NLP research, a sample of which are reviewed, for example, by Church and Mercer [21].

⁴One may argue that supervised corpus-based methods have not released us from hand-encoding tasks. However, we would like to note that manual annotation is still easier than describing rule sets based on human introspection.

consider the difficulty of the task, which can range between coarse-grained (totally distinct) and fine-grained (closely related) word sense distinctions. In this research, we explore disambiguation of verb senses (verbal polysemy⁵) based on an existing machine readable dictionary. Note that the dictionary we use provides relatively fine-grained verbal polysemy. Among the research fields described in Section 1.1, our research focus is expected to improve the quality of machine translation, parsing, natural language understanding and lexicography. With regard to current information retrieval systems, the disambiguation of noun senses seems to be more crucial a task than that for verb senses (for one thing, keywords and user queries usually comprised of noun phrases). However, the potential contribution of verb sense disambiguation extends to a significant proportion of the tremendous range of information retrieval systems.

Second, let us describe the approach we use to tackle the verb sense disambiguation task. Obviously, our research methodology must be contextualized in terms of past research literature associated with corpus-based word sense disambiguation, for which the reader should refer to Chapter 2. In this process, it is important not to limit our focus only to *verb* sense disambiguation, because methods employed in other types of disambiguation may also be applicable. In brief, our system employs a similarity-based method, in which disambiguation is performed based on the similarity between a given input sentence and example sentences associated with each verb sense. The similarity is computed by averaging the similarity for each case (or case filler noun) syntactically governed by the target verb. Through preliminary experimentation on Japanese verbs, we identified a number of problems associated with our research focus. In the following, we describe our approaches to these problems.

- The degree to which each case contributes to verb sense disambiguation is not consistent. One may intuitively understand that in the case of English, for example, object case is more closely related to verb senses than subject case. The same observation can also be made in the case of Japanese. We explore a method of introducing this notion *computationally* into similarity-based methods, and demonstrate the effectivity of our proposal through comparison with a number of different methods (see Chapter 3 for details).
- As with most corpus-based word sense disambiguation systems, our system uses a large-scaled corpus annotated with correct verb senses. However, a considerable overhead is required when one tries to manually perform the annotation process. One possible solution would involve automatic annotation, that is, an unsupervised method. However, our experiments show that at least for the particular unsupervised method targetted in our

⁵We will define “polysemy” in Section 2.1.

research, unsupervised methods still find it difficult to match the performance achieved for supervised methods. In view of this result, we explore a *semi*-supervised method in which we *selectively* sample a small number of effective example sentences from a given corpus in the annotation process. In other words, our sampling method aims at minimizing the overhead for manual supervision, without degrading the system performance (see Chapter 4 for details).

- We also identified that the performance of our system is highly dependent on the similarity computation between the input and example sentences. Roughly speaking, past approaches for similarity computation can be subdivided into statistically-driven and thesaurus-driven methods. In addition, integration methods combined these two approaches have recently been proposed. We also explore the similarity computation in the context of the integration of different methods, and demonstrate the effectivity of our proposal for our verb sense disambiguation system (see Chapter 5 for details).

1.3 Outline of the Proceeding Chapters

Chapter 2 surveys past research on word sense disambiguation, and identifies open questions in the field. Chapter 3 describes the overall architecture of our similarity-based verb sense disambiguation system, in which we newly introduce the notion of degree of case contribution to verb sense. Chapter 4 proposes a selective sampling method, which samples a smaller-sized, effective example set for use with our verb sense disambiguation system, so as to minimize the overhead for supervision and that to search a large-sized corpus. Chapter 5 explores and discusses the similarity computation used in our system, through integration of thesaurus taxonomy and co-occurrence statistics. Finally, Chapter 6 summarizes our contributions and discusses outstanding issues.

Chapter 2

Past Research on Word Sense Disambiguation

Reflecting the rapid growth in utilization of machine readable texts, word sense disambiguation techniques have been explored variously in the context of “corpus-based NLP approaches”. This chapter surveys past research associated with corpus-based word sense disambiguation, focusing mainly on methodology and evaluation criteria.

2.1 Overview: Terminology and Task Description

The task of a word sense disambiguation system (“system” or “WSD system”, hereafter) is to resolve the lexical ambiguity of a word¹ in a given context. To put it more precisely, the term “lexical ambiguity” refers to two different concepts: “homonymy” and “polysemy”. The former is the case where two different words happen to have the same lexical form, and the latter is the case where one word has several (related) meanings. Conventionally, the distinction between *bank* (“river edge”) and *bank* (“financial institution”) has been used as an example of homonymy, and *rust* (verb) and *rust* (noun) for polysemy. In this dissertation, we will generally use the term “polysemy” to refer to both lexical ambiguity types, because (a) the difference between these two ambiguity types has been relatively less controversial in word sense disambiguation tasks (although from a linguistic point of view, the two ambiguity types should be rigorously defined), and (b) the focus of this research is on disambiguating verbal polysemy rather than homonymy. To derive plausible word senses (polysemy), most past WSD systems have used lexical resources, such as machine readable dictionaries (MRDs) or thesauri. Thus, the task

¹In most cases, past methods disambiguate only “content word”, and “functional words” such as prepositions are beyond the scope of word sense disambiguation.

of WSD systems can be termed as “categorization” because a plausible word sense is selected from *predefined* candidates. Note that this task should not to be confused with automatic identification of word senses [41, 144, 158] and word clustering (grouping) [49, 53, 113, 139]. One may notice that the distinction between “ambiguity” and “vagueness” is also a controversial linguistics issue [44, 67, 140]. A comprehensible example of vagueness would be *aunt*, which is unspecific between “father’s sister” and “mother’s sister” without explicit context. However, this type of under-specification is beyond the scope of word sense disambiguation unless dictionaries list father/mother’s sister as different senses of *aunt*.

Let us leave the definition of lexical ambiguity and turn to the process of word sense disambiguation. Given an input sentence containing polysemous word(s), most WSD systems first preprocess the input to extract a set of features (clues) used for the disambiguation. This preprocessing typically involves morphological/syntactic analysis because the parts-of-speech of words appearing in the input and syntactic relations involving polysemous words, can be informative features. Macroy [90], for example, identified syntactic tags, morphology, collocations, and word associations as the most important sources of information for word sense disambiguation. Thereafter the system interprets polysemous word(s) by selecting a single plausible word sense. While certain systems interpret only one polysemous word in the input, other systems simultaneously interpret all polysemous words appearing in the input. Wilks and Stevenson [152] call this second task type “word sense tagging” (based on the analogy of part-of-speech tagging). However, the difference between the two task types is relatively unimportant for this research.

Section 2.2 classifies past word sense disambiguation methods. Section 2.3 then describes the way past research has evaluated word sense disambiguation methods. Section 2.4 describes related NLP research which is expected to enhance word sense disambiguation. Finally, discussion associated with word sense disambiguation is added in Section 2.5.

2.2 Different Methodologies for Word Sense Disambiguation

This section surveys different past methods for corpus-based word sense disambiguation, to clarify problems tackled in this research. First, Section 2.2.1 classifies past methods, according to their induction mechanism (rule-based method or probabilistic model, say). Second, Section 2.2.2 focuses on a different viewpoint, that is, supervised vs. unsupervised learning methods. While supervised methods require manual annotation of correct sense to each polysemous word contained in a corpus, unsupervised methods automatically acquire corpora annotated with (presumably correct) word senses. These automatic methods are what Section 2.2.2 focuses on principally. It should be noted that criteria for the classification of past methods can

vary depending on the viewpoint or interest of the researcher².

2.2.1 General classification of methodologies

This section classifies past methodologies as shown in Figure 2.1, which first divides methods into two different approaches. The first approach can be called “qualitative approach”³ and uses selectional rules associated with each word sense candidate. Given an input containing a polysemous word, the system deterministically selects the sense(s) for which the rules are satisfied. Generally speaking, the granularity of selectional rules can be a problem for this approach. That is, specified rules often fail to select a word sense for exceptional inputs. On the other hand, generalized rules run the risk of selecting incorrect word senses. To counter this problem, the second approach – what we call the “quantitative” approach – computes scalable values for each word sense candidate, and selects the sense with maximal value as the interpretation of the polysemous word. Compared with the rule-based approach, this approach is more robust for exceptional inputs⁴. Both approaches generally use a corpus, in which examples are associated with each word sense candidate. We shall call this corpus the “training data”. Note that in this section, examples in the training data can be annotated either *manually* or *automatically*, without loss of generality. The qualitative approach can be further subdivided in terms of rule format. Subcategories described in this section are “selectional restrictions”, “decision trees” and “decision lists”. The quantitative approach can be divided in terms of the method used to compute scalable values, as “probabilistic models” and “similarity-based methods”. All subcategories of these two approaches correspond to the different titles of the following items.




Figure 2.1: Classification of different methodologies of word sense disambiguation

²For example, Wilks and Stevenson [152] classified past word sense disambiguation methods differently.

³One may notice that the terms “symbolic/rule-based approach” can be interchangeably used for qualitative approach. However, we used the term “qualitative approach” to make the direct contrast to “quantitative approach”.

⁴A number of *compromised* methods can be found in past research. Wilks [150] used selectional restrictions as *preferences* rather than constraints. Uramoto [142] combined qualitative and quantitative approaches.

Selectional restrictions Selectional restrictions [64], which impose constraints on arguments for a given word (sense), have commonly been used in word sense disambiguation relying on hand-crafted rules [15, 54, 84]. Let us take the following example sentences (a) and (b), which contain different usages of the verb *employ*, that is, “to hire” and “to accept”:

- (1) a. The facility will *employ* new employees. (“to hire”)
 b. The committee *employed* his proposal. (“to accept”)

One can intuitively differentiate the senses of *employ* in sentences (1-a) and (1-b) with the complements of each *employ*. To be more precise, *employ* in (1-a) restricts its subject and object nouns to those associated with the semantic features HUMAN/ORGANIZATION and HUMAN, respectively. On the other hand, *employ* in (1-b) restricts its subject and object nouns to those associated with the semantic features HUMAN/ORGANIZATION and IDEA, respectively. Consequently, given *employees* as the object, the sense “to hire” is selected as the interpretation of *employ* in (1-a), and the sense “to accept” is ruled out. The same reasoning can be used to select the sense “to accept” as the interpretation of *employ* in (1-b). One may notice that selectional restriction can also disambiguate polysemy of verb complements (the subject and object). For example, *facility* in (1-a) has multiple senses, a sample of which are “installation”, “proficiency” and “readiness”. However, the selectional restriction imposed for the subject of *employ* (“to hire”) can correctly select the sense “installation” as the interpretation of *facility*. It should be noted that the polysemy of both *facility* and *employ* are theoretically disambiguated simultaneously. In fact, the disambiguation process can be seen as mutually propagating semantic constraints to each polysemous word through selectional restriction⁵.

However, considerable human effort is required to describe large-scaled selectional restrictions. Manual construction is also associated with human bias and inconsistency of granularity. Besides this, revision requires additional human overhead, that is, lexicographers have to laboriously identify and revise associated entries. Resnik [120] proposed an information-theoretic method to automatically identify selectional restrictions, which is expected to counter this problem to some degree. Resnik identified selectional restrictions as semantic classes defined in the taxonomy of the well-known English semantic network WordNet [93]⁶, and only those nouns dominated in the taxonomy by the identified class can satisfy the restriction. While Resnik used this method to resolve syntactic ambiguity, Ribas [124] applied it to disambiguate senses of case filler nouns (for example, disambiguation of the polysemy of *facility* in (1-a) as performed above). The basis of this method is to estimate the information-theoretic association degree

⁵Lytinen [84] and Nagao [98] used constraint propagation method to simultaneously resolve syntactic and semantic ambiguity.

⁶WordNet terms semantic classes as “synsets”.

between case fillers and semantic classes, for each verb (or verb sense). Intuitively speaking, the association degree gives a greater values for semantic classes that are likely to appear as complement of a given verb sense. Formally speaking, the association degree between verb sense s and class r (restriction candidate) with respect to case c is computed by Equation 2.1 [120, 124].

$$A(s, c, r) = P(r|s, c) \cdot \log \frac{P(r|s, c)}{P(r|c)} \quad (2.1)$$

Here, $P(r|s, c)$ is the conditional probability that a case filler marked with case c of sense s is dominated by class r in the WordNet taxonomy. $P(r|c)$ is the conditional probability that a case filler marked with case c (disregarding verb sense) is dominated by class r . The distribution obtained from training data is usually used to estimate each probability.

Decision trees In spite of its long history of applications in AI research, “decision trees” have rarely been applied to word sense disambiguation⁷. Among a number of proposed decision tree algorithms “C4.5” [117] has been used relatively commonly as a benchmark comparison. Mooney [96] and Pedersen and Bruce [111] (individually) compared the performance of various word sense disambiguation methods with the C4.5 algorithm. Tanaka [137] used the C4.5 algorithm to acquire English-Japanese verbal translation rules. One may notice that this task can also be seen as a type of rule induction for word sense disambiguation, because a unique English verb can be interpreted as different Japanese verbs. Figure 2.2 shows a fragment of the decision tree for the English verb *take*, which corresponds to multiple Japanese translations such as *erabu* (“to choose”), *tsureteiku* (“to take along”) and *motteiku* (“to take away”). Given an input containing *take* followed by the object noun *him* and preposition *to*, the decision tree selects *tsureteiku* as the interpretation, by traversing the branches corresponding to “object=him” and “prep=to”.

Here, let us devote a little space to explaining the basis of the C4.5 algorithm. In this decision tree algorithm, classification rules are formulated by recursively partitioning the training data. Each nested partition is based on the feature value that provides the greatest increase in the information gain ratio for the current partition. The final partitions correspond to a set of classification rules where the antecedent of each rule is a conjunction of the feature values used to form the corresponding partition. Let $\{C_1, C_2, \dots, C_k\}$ denote class candidates, one of which is assigned to the input. Suppose we already have a partition, which divides the set T of training

⁷Okumura and Tanaka [109] proposed generalized discrimination networks (GDNs), which can be seen a variant of decision trees, for word sense disambiguation. However, we will not discuss their method further here, because they focus on exploration of an “incremental disambiguation model” [91] and not automatic construction of the networks the main topic of this research.




Figure 2.2: A fragment of the decision tree for the English verb *take*

data into subsets $\{T_1, T_2, \dots, T_n\}$. Given that $freq(C_i, T)$ and $|T|$ denote the number of training data in T that belong to class C_i and the number of training data in set T , respectively, the “entropy” (uncertainty) of set T can be estimated by Equation 2.2.

$$info(T) = - \sum_i \frac{freq(C_i, T)}{|T|} \cdot \log \frac{freq(C_i, T)}{|T|} \quad (2.2)$$

Suppose T is partitioned into n subsets using feature value X as a classification rule, the expected entropy over partitioned subsets, $info_X(T)$, is estimated by Equation 2.3.

$$info_X(T) = \sum_i \frac{|T_i|}{|T|} \cdot info(T_i) \quad (2.3)$$

The decrease in entropy resulting from the partition, which represents the information gain for X , is then estimated by Equation 2.4.

$$gain(X) = info(T) - info_X(T) \quad (2.4)$$

Decision lists Decision lists are a form of rule representation as proposed by Rivest [126], and consist of tuples of the form “(condition, value)”. As Rivest observes, decision lists can be seen as “if-then-else” rules, in other words, exceptional conditions appear earlier while general conditions appear later in the list⁸. Given a query, each condition in the decision list is applied sequentially until a condition which is satisfied by the query is found. Thereupon, the value which corresponds to that condition is selected as the answer. Yarowsky applied decision lists to the task of accent restoration [156] (this is one type of lexical disambiguation, in which a single word is associated with multiple pronunciations), and word sense disambiguation [157].

⁸The last condition accepts all cases (namely “true”), otherwise the system could potentially fail to make any decision for certain input types.

In Yarowsky’s case [156, 157], each condition corresponds to a word collocation which can be used as evidence to resolve lexical ambiguity, and each value corresponds to a correct word sense (or pronunciation). Since manual identification of effective conditions is expensive and inconsistent, Yarowsky used word collocation (within a fixed window size) obtained from a large corpus to automatically identify effective evidence types. The effectivity degree of a given piece of evidence is estimated as the likelihood that it supports a given sense candidate more strongly than another⁹. Formally speaking, this notion is represented as the log-likelihood, that is, the ratio between the conditional probability that sense s_1 and sense s_2 occur, respectively, given evidence E (Equation 2.5).

$$\log \frac{P(s_1|E)}{P(s_2|E)} \quad (2.5)$$

In decision lists, evidence (along with their supporting word senses) is sorted according to log-likelihood, in descending order. Figure 2.3 shows a fragment of the decision list trained for the disambiguation of the word *plant* (“organism”/“factory”) [157], where each piece of evidence denotes a specific collocational pattern or word collocating with a certain distance associated with *plant*. For example, given the input containing the pattern *plant height*, the interpretation for *plant* is “organism”. Note that Yarowsky [156, 157] used a method binary sense distinction method, in that the number of sense candidates was limited to two. To apply this method to the disambiguation of words with multiple ambiguity, the denominator in Equation 2.5 is presumably computed based on the probability of (a) all the other sense candidates or (b) only the most competitive sense candidate (that is, the sense candidate with the second highest probability given evidence E).

evidence	sense
<i>plant</i> growth	organism
car (within $\pm k$ words)	factory
<i>plant</i> height	organism
union (within $\pm k$ words)	factory
equipment (within $\pm k$ words)	factory
⋮	⋮

Figure 2.3: A fragment of the decision list for the word *plant*

Pedersen and Bruce [111] used the “CN2” rule induction algorithm [23] as one form of comparison in their “decomposable model”¹⁰. The CN2 algorithm generates decision lists based on a given set of training data. The rule induction algorithm consists of finding rules, and

⁹In Yarowsky’s case [156, 157], the number of sense candidates for each word was limited to two.

¹⁰See the following “probabilistic models” item for details of the decomposable model.

measuring the significance of those rules based on entropy and their coverage, over the range of the training data. The quality of decision lists is highly dependent on the sequence of rules, and thus naive rule sequences end up degrading the disambiguation performance. To overcome this problem, the latest version of CN2 [22] optionally employs an “unordered rule” technique. Let us explain, in passing, the basic algorithm here. The training procedure is the same as in the original version. However, given an input and decision rule set, rather than applying the rules in a fixed order, the algorithm collects *all* the rules which the input satisfies, and their associated conditions. Note that by disregarding the rule order (in other words the “context”), different rules can potentially support different values and consequently ambiguity often remains unresolved. In such a case, additional consideration is given to the distribution of the training data covered by the collected rules, in choosing a unique value. Let us take Figure 2.4 as an example rule set [22], where each line corresponds to a rule for the two values of “bird” and “elephant”. In this figure, the “coverage” column denotes the number of training examples which satisfy the corresponding rule, for the values “bird” and “elephant”, respectively. Let us consider the following example input.

`<size=large, beaked=yes, legs=two, feather=yes, flies=no>`

The rule sets in Figure 2.4 cannot uniquely decide the answer for this input, which satisfies all three rules. However, the summation of the coverage for these rules “(35,10)” puts the preference for the final answer on “bird” over “elephant”.

rule	coverage
if legs=two and feather=yes then class=bird	(13,0)
if size=large and flies=no then class=elephant	(2,10)
if beaked=yes then class=bird	(20,0)

Figure 2.4: An example rule set for CN2

Probabilistic models From the viewpoint of probability theory, the task of word sense disambiguation is to select the sense with maximal probability for a given input¹¹. The probability for word sense s , $P(s|x)$, is commonly transformed into Equation 2.6 through use of the Bayesian theorem.

$$P(s|x) = \frac{P(s) \cdot P(x|s)}{P(x)} \quad (2.6)$$

¹¹Some statistical methods are reviewed, for example, by Charniak [16].

In practice, $P(x)$ can be omitted because this factor is constant for all the sense candidates, and therefore does not affect the relative probability for different senses (Equation 2.7).

$$\arg \max_s P(s|x) = \arg \max_s P(s) \cdot P(x|s) \quad (2.7)$$

The probability of sense s , $P(s)$, is usually estimated based on its distribution, obtained from training data. Thus, the performance of the method fundamentally depends on the method for approximating $P(x|s)$. For the purposes of general discussion, let the input be represented by a vector comprised of word sense disambiguation features, as given below.

$$\langle F_1 = f_1, F_2 = f_2, \dots, F_N = f_N \rangle$$

Here, F_i and f_i are the i -th feature type and its value, respectively. Let us summarize the diverse range of past methods based on the following two principles.

The first principle is to identify an informative feature set, and ideally, different feature sets for each different target polysemous word. Typically, words which *saliently* collocate with a word sense are used as features. In this case, each feature takes a binary value, that is, 1 for existence and 0 for absence in the input. A number of methods have been proposed to automatically identify informative collocating words (termed “salient words” [154] or “indicators” [60]). Yarowsky [154] used mutual information between w and sense s (in Yarowsky’s case, word senses are semantic categories defined in Roget’s thesaurus [14]) to estimate the degree of salience of word w to sense s . Intuitively speaking, the mutual information between two phenomena gives a greater value when these phenomena are more likely to co-occur. The mutual information of w and s , $I(w; s)$, is computed as shown in Equation 2.8.

$$I(w; s) = \log \frac{P(w|s)}{P(w)} \quad (2.8)$$

Here, $P(w|s)$ is the probability that w appears given s , and $P(w)$ is the probability that w appears in the context¹². These factors are estimated based on the relative distributions of w and s in the training data. Figure 2.5 shows examples of salient words related to the categories ANIMAL and TOOLS [154]. Intuitively speaking, when words like *species* and *family* appear in the input, the probability for ANIMAL tends to be greater than that for TOOLS. Justeson and Katz [60] select word w as indicator of sense s such that w appears more frequently with

¹²Strictly speaking, $P(w|s)$ and $P(w)$ should be denoted as $P(F_w = 1|s)$ and $P(F_w = 1)$, respectively, where F_w is the feature representing the existence of word w . However, we use a simplified notation here.

sense s than with other sense candidates. This is equivalent to selecting a word w that satisfies Equation 2.9, where t is any other sense candidate of the polysemous word under evaluation¹³.

$$P(s|w) > P(t|w) \quad (2.9)$$

Their objective is to disambiguate adjective senses, or in their case, polysemy of an adjective as defined by its antonyms. For example, the word *old* has two senses, i.e. “not new” and “not young”. Figure 2.6 shows examples of indicator nouns identified in their paper [60]. Note that in Justeson and Katz’s case, the disambiguation process in itself does not rely on a probabilistic model, or more precisely, indicators are used simply as restriction rules. Ng and Lee [105] and Pedersen et al. [112] use multiple feature types along with collocating words. The following additional features are also usually used¹⁴:

- the morphological properties of polysemous words (for example, singular/plural in the case of a polysemous noun, and tense in the case of a polysemous verb),
- parts-of-speech of collocating words,
- syntactic relations associated with polysemous words.

category	salient words
ANIMAL	species, family, bird, fish
TOOLS	tool, machine, engine, blade

Figure 2.5: Example salient words for two categories

sense	indicators
<i>old</i> (“not new”)	world, thing, car, way
<i>old</i> (“not young”)	man, people, woman, wine

Figure 2.6: Example indicator nouns for the adjective *old*

The second principle is to compute (approximate) $P(x|s)$, based on the conditional probability of feature(s) given sense s . The most simple model, the “Naive-Bayes method”, assumes that features are conditionally independent of each other given sense s . That is, $P(x|s)$ is

¹³In practice, $P(s|w)$ and $P(t|w)$ are transformed using the Bayesian theorem as for Equation 2.6.

¹⁴Ng and Lee [105] use these features for exemplar-based word sense disambiguation, which will be described in the next paragraph.

approximated simply by the product of each $P(F_i = f_i | s)$, as shown in Equation 2.10.

$$P(x|s) \simeq \prod_i P(F_i = f_i | s) \quad (2.10)$$

A number of past research applied the Naive-Bayes method to word sense disambiguation [43, 96, 78, 103, 110, 111].

A more complex model, namely the “decomposable model” [13, 112], considers interdependency between different features. Here, let features (F_1, F_2, F_3) represent the input. Suppose F_1 and (F_2, F_3) are interdependent given s , such that $P(x|s)$ can be expressed by Equation 2.11.

$$P(x|s) = P(F_1 = f_1 | s) \cdot P(F_2 = f_2, F_3 = f_3 | s) \quad (2.11)$$

In the decomposable model, the number of parameters to be estimated, which is proportional to the number of combinations of values for interdependent features, can be enormous. This often leads to the data sparseness problem. In addition, it is difficult to identify which features are interdependent given a particular sense. To overcome this problem, Pedersen et al. [112] proposed a method of automatically identifying the optimal model (higher performance and fewer parameter estimations), by iteratively altering the complexity level of the model. However, their experimental results show that identified models did not generally outperform the Naive-Bayes method.

Similarity-based methods Given problems along with their solution (the way they have been solved), humans are able to solve new problems based on the analogy of the previously observed cases. This analogy-based process has been explored in many AI applications, under the headings of case-based reasoning [72], exemplar-based reasoning [4], memory-based reasoning [149] and instance-based learning [1]¹⁵. A number of word sense disambiguation systems have variously applied these methods¹⁶. One critical issue in this has been the computation of *similarity* between an input (a new problem) and examples in the training data (previous problems), and thus, we term these methods “similarity-based methods”. In the k -nearest neighbor method (k -NN), one similarity-based method, processing proceeds as follows. First, k examples similar to the input are retrieved from the training data. Thereafter, retrieved examples *vote* on the sense of the polysemous word in the input, or in other words, the sense receiving the highest frequency

¹⁵Nagao [99], for example, explored the analogy principle in NLP applications.

¹⁶In the case of word sense disambiguation, the reasoning mechanism is relatively simple when compared with tasks focused on in AI research (for example, resolution of a “political dispute” focused on by Kolodner [72]), because previous problems are merely example sentences containing polysemous words annotated with senses.

of annotation with the k examples is selected as the interpretation of the input word. In the case of $k = 1$, this method is termed the nearest neighbor method, in which the input word is disambiguated simply by superimposing the word sense associated with the example of highest similarity. Generally speaking, the nearest neighbor method has been used more commonly than k -NN in past word sense disambiguation research¹⁷. In the following, we classify past similarity-based methods in terms of the method of computing the similarity (or distance).

Suppose each training example and the input (“examples”, hereafter) are represented by a feature vector defined by Equation 2.8, in which one may notice that each example is positioned in an N -dimensional space, where feature F_i corresponds to the i -axis. One common such implementation termed the “vector space model” (VSM) computes the similarity between two examples by the angle between the two vectors representing the examples. Note that VSM has a long history of application in information retrieval (IR) and text categorization (TC) tasks [127]. However, in the case of IR/TC, VSM is used to compute the similarity between documents, which is represented by a vector comprising statistical factors of content words in each document. Formally speaking, the similarity between two examples x and y is computed as the cosine of the angle between their associated vectors. This can be expressed by Equation 2.12, where \vec{x} and \vec{y} are vectors representing examples x and y , respectively.

$$\text{sim}(x, y) = \frac{\vec{x} \cdot \vec{y}}{|\vec{x}| |\vec{y}|} \quad (2.12)$$

Schütze [130] applies the vector space model to word sense disambiguation, although vectors are used for each word sense not individual examples¹⁸. First, Schütze represents each word by a “word vector”, that is, a vector comprising the statistics (e.g. frequency) of its collocating words. The collocational statistics are usually collected within a fixed proximity. Then, each context is represented by a “context vector”, which is the sum (or “centroid”) of word vectors related to words appearing in the context. Figure 2.7 illustrates the idea of the context and word vectors, in which each w denotes a word vector corresponding to a context vector V . It should be noted that unlike the vector space model in IR/TC, Schütze’s method returns a positive similarity value even when two given context vectors have no words in common (such as with contexts 1 and 2 in Figure 2.7), that is, two context vectors can be similar when they comprise similar word vectors. Then, automatic clustering algorithms [18, 27] are used to cluster each polysemous

¹⁷Ng [103] automatically identified the optimal value of k over the range of a given training data, which reportedly improved the performance of the nearest neighbor method.

¹⁸Within the IR/TC community, this approach is called “text-to-category comparison”, and contrasted with “text-to-text comparison” which computes the similarity between the input and individual examples.

word into word senses, which are also represented by “sense vectors”¹⁹. In practice, Schütze used the “singular value decomposition” (SVD) technique [9, 45], which finds the major axes and reduces the dimension of the vector space. Finally, Equation 2.12 computes the similarity between the input and each word sense cluster, to select the word sense with maximal similarity. Leacock et al. [78] compared the vector space model, the Naive-Bayes method (see the previous paragraph) and neural networks, and reported that the vector space model and neural networks slightly outperformed the Naive-Bayes method. Niwa and Nitta [106] also explored the vector space model, in an implementation resembling Schütze’s method. The vector representation is expressed by Equation 2.13, where $V(s)$ and $V(w)$ are the context vector for sense s and the word vector for word w within a fixed proximity from s , respectively.

$$V(s) = \sum_{w \in \text{context}} V(w) \quad (2.13)$$

In Niwa and Nitta’s case, the word vector for w , $V(w)$, consists of the mutual information between w and “origin words” (commonly used words), as in Equation 2.14.

$$V(w) = \langle I(w; o_1), I(w; o_2), \dots, I(w; o_i), \dots \rangle \quad (2.14)$$

Here, $I(w; o_i)$ denotes the mutual information between w and origin word o_i . In practice, origin words are made up of the 51st to 1050th most frequently used words in Collins English Dictionary definitions [133]. Equation 2.12 is used to compute the similarity between two context vectors. However, unlike Schütze’s method, which averages context vectors into sense vectors, Niwa and Nitta prefer to use the “text-to-text comparison” method: they compute the similarity between the input and *individual* examples.

One may argue that the “Euclidean distance” between two examples in N -dimensional space can also be applied to the similarity computation (as performed by Aha et al. [1] for instance-based learning). However, this similarity measure seems less effective for word sense disambiguation, because Euclidean distance is sensitive to vector length, which is usually proportional to the frequency of collocating words. In other words, Euclidean distance usually assigns lower similarities to frequently appearing word senses.

Ng and Lee [105], following Cost and Salzberg [25], used a different similarity measure. In their method, the distance between two examples is computed by summing the distances between the feature values associated with those examples. In other words, two examples are similar when they have feature values that roughly correspond in distribution as obtained from

¹⁹Schütze used two different clustering algorithms, depending on the data size.

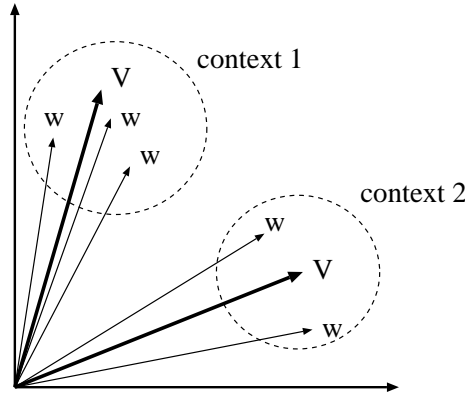


Figure 2.7: The word and context vectors for two contexts

the training data. The distance between two feature values f_1 and f_2 of a feature F is computed as shown in Equation 2.15.

$$dist(f_1, f_2) = \sum_i |P(s|F=f_1) - P(s|F=f_2)| \quad (2.15)$$

Here, $P(s|F=f_1)$ is the conditional probability of sense s given feature F takes value f_1 . $P(s|F=f_2)$ denotes a similar probability for value f_2 of feature F . Ng and Lee [105] used multiple features described in the previous item (see “probabilistic models”), and showed its effectivity over the use of single type feature. Their results suggest that feature selection is also a crucial task in the similarity-based methods.

Cho and Kim [19] used “relative entropy”, which estimates the degree to which two probability distributions differ, for disambiguation of verb sense. They used only one feature, in that each example is represented simply by the object case noun for the target polysemous verb. The similarity between examples is computed based on the distribution of their associated nouns. Cho and Kim represented the distribution of noun n by Equation 2.16, where $P(v_i|n)$ is the conditional probability that verb v_i appears given noun n .

$$d(n) = \langle P(v_1|n), P(v_2|n), \dots, P(v_m|n) \rangle \quad (2.16)$$

Given two examples associated with distributions $d(n_1)$ and $d(n_2)$, the similarity between them

is computed by Equation 2.17²⁰.

$$\text{sim}(d(n_1), d(n_2)) = - \sum_i P(v_i|n_1) \log \frac{P(v_i|n_1)}{P(v_i|n_2)} \quad (2.17)$$

All the methods reviewed so far compute the similarity based on the statistical distribution obtained from training data. The remainder of this section is devoted to reviewing methods relying on hand-crafted resources (mostly thesauri) for the similarity computation. Here, let us assume that sentences (1) (see page 10) are examples incorporated into training data, and that we have the following sentence as input.

(2) The company *employed* the graduate.

One may notice that the sense of *employ* in the input would be “to hire”, because the subject and object of *employ* in the input are semantically closer to those in (1-a), respectively, than to those in (1-b). This is where hand-crafted thesauri can be applied, based on the intuitively feasible assumption that words located near each other within the structure of a thesaurus have similar meaning. In practice, most methods heuristically predefine the relation between the similarity between nouns, and the length of the path between them in the thesaurus structure. Kurohashi and Nagao [77] used the *Bunruigoihyo* thesaurus [102] for disambiguating senses of Japanese verbs. Uramoto [142] used LDOCE [114] for disambiguating senses of English verbs. Note that complements of verbs in the training data can be considered as an extensional description of selectional restriction. In other words, instead of merely ruling out inappropriate case filler nouns, example complements estimate the degree that input complements satisfy the restriction imposed on them. It should also be noted that lexical ambiguity of complement nouns can be resolved in the same manner as previously demonstrated for the disambiguation of *facility* in (1-a). One may notice that *company* in the input can be interpreted as “installation” because *facility* in (1-a) is also related to “installation”, and therefore these two words are close vicinity within located in the thesaurus²¹. Li et al. [81] and Lin [82] (independently) explored this notion. In both cases, the similarity between two words is computed based on the taxonomy defined in WordNet [93]. Li [81] heuristically predefined the relation between the length of the path in the taxonomy and the similarity. Lin [82] used a more formal measure resembling the information-theoretic taxonomy-based similarity measure proposed by Resnik [121, 122] (see Section 5.5 for details).

²⁰In practice, Cho and Kim [19] used a smoothing technique in the case of an occurrence of $P(v_i|n) = 0$ in the given vectors. However, we do not further describe the technique utilized.

²¹Regrettably, Kurohashi and Nagao [77] and Uramoto [142] did not make significant comment on this point.

2.2.2 Restricting supervision

Corpus-based approaches have recently pointed the way to a promising trend in word sense disambiguation. A number of experiments have shown that the performance of word sense disambiguation can be significantly improved by enhancing the volume of supervised examples [96, 104]. However, to do this, supervised methods require considerable manual annotation in supervising large-sized training data sets. To resolve this problem, unsupervised and semi-supervised methods, which (semi-)automatically acquire annotated training data, have been variously explored. On the other hand, one may wonder if systems without human supervision can perform to a reasonable level. Below, we survey different approaches for restricting reliance on supervision. In addition, despite the promising features of each method, we discuss problems and open questions related to the various approach types.

Bilingual corpora Based on the observation that different word senses in a given language can correspond to distinct words in other languages, Dagan and Itai [28] used bilingual corpora for word sense disambiguation (they used Hebrew-English and German-English language-pair corpora, respectively). In their case, word polysemy in the source language is defined based on the existence of separate translations in the target language. In other words, the objective of this research can be seen as translating a source language sentence to a target language sentence on a word-to-word basis. The method proceeds as follows:

1. the source language corpus is parsed to extract syntactic tuples (source syntactic tuples),
2. a bilingual lexicon is used to generate alternative target syntactic tuples for each source syntactic tuple,
3. the target language corpus is parsed to extract target syntactic tuples, which are used to evaluate the plausibility of each alternative target tuple generated in the previous step,
4. the target sentence maximizing the combined plausibility of target syntactic tuples is generated.

For example, the Hebrew tuple “*higdil sikkuy*” (verb-obj) containing the polysemous word *higdil* is associated with three English tuples: “increase chance”, “enlarge chance” and “magnify chance”. However, the polysemy can be resolved by selecting the tuple which is most likely to occur, base on the target corpus, i.e. “increase chance”. It should be noted that Dagan and Itai’s method uses information extracted *independently* from the source and target language corpora, which means that the two corpora may not necessarily be translations of one another,

and that manual supervision is avoided. This property of potential non-correspondence leads to a salient contrast with the method proposed by Brown et al. [10], which uses bilingual *aligned* corpora. The applicability of bilingual corpus-based methods is relatively limited, that is, they are restricted to translation-oriented applications. Apparently, word sense distinctions differ depending on application, and therefore (as identified by Krovets [75], for example) word sense disambiguation aimed at machine translation is not necessarily useful for information retrieval.

Bootstrapping The basis of “bootstrapping” is, given an initial training data set (usually consisting of a small number of annotated examples), to progressively enhance the training data by iteratively acquiring presumably correctly annotated examples from previous disambiguation results. As can be imagined, this method often ends up incorporating noise (incorrectly annotated examples) into the training data. To avoid this problem, Hearst [50] used a relatively small number of *supervised* examples as initial training data. Yarowsky [157] completely excluded manual supervision by automatically acquiring the initial training data set from a dictionary. In addition, Yarowsky used discourse constraints to exclude noise from the decision list. To put it more precisely, when significantly large number of examples associated with a given discourse are annotated with a common sense in the training data, all examples associated with that discourse are standardized to the same sense annotation. Yarowsky’s experimental results show that the performance of this method is equivalent to that achieved by supervised learning.

However, one controversial issue would be the applicability of bootstrapping to the disambiguation of *finer-grained* word sense distinction [105], in that this method has generally been applied only to *coarse-grained* binary word distinctions, such as *sake* (“benefit”/“Japanese liquor”). Karov and Edelman [63] used bootstrapping to automatically enhance word sense classifiers using dictionary definitions and a corpus. As with similar research, their method also focused on coarse-grained binary sense distinction, such as *suit* (“court”/“garment”). To sum up, the effectivity of fully-automatic bootstrapping for the disambiguation of fine-grained sense distinction remains an open question.

Spreading noise through monosemous words Yarowsky [154] projected polysemy onto word categories defined in Roget’s thesaurus [14], and trained statistical classifiers for *categories* rather than individual words. Consequently, monosemous words associated with each category provide reliable co-occurrence statistics, and noisy statistics arising from polysemous words should be relatively dispersed (tolerable), without any supervision in sense tagging. For example, the classifier for the category ANIMAL is expected to be reliably trained by monosemous words like *sparrow*, although polysemous words like *crane* potentially introduce a certain proportion

of noise. Luk [83] adopted a similar approach to obtain statistics about word senses defined in LDOCE [114]. While these methods completely remove any overhead for supervision, the applicability of the methods is limited to word senses mapped onto thesauri, where word entries are associated with categories. This property does not apply to many dictionaries.

Automatic clustering Schütze [130] reduced manual supervision by using automatic clustering algorithms (proposed by Cheeseman et al. [18] and Cutting et al. [27], respectively). First, clustering algorithms are used to divide the training data into a certain number of clusters. Thereafter, a human expert examines a small number of examples (from 10 to 20) contained in each cluster, which are applied in determining an appropriate word sense for each cluster. Thus, strictly speaking, this method does not constitute supervised learning. Given an input, the cluster (word sense) with the maximum similarity to the input is selected as the interpretation (see Equation 2.12 for this computation). Regrettably, since reported experiments have not compared this semi-supervised method with fully-supervised learning methods, it is not possible to ascertain whether automatic clustering algorithms are expected to reduce manual supervision without degrading the performance demonstrated by supervised methods.

Pedersen and Bruce [110] used automatic clustering algorithms relying on McQuitty's similarity analysis [89] and Ward's minimum-variance method [148], respectively. Their training/test data includes polysemous nouns, verbs and adjectives collected from the ACL/DCI Wall Street Journal corpus [86], in which each word is annotated with a single sense defined in LDOCE [114] or WordNet [93]. Their comparative experiments showed that automatic methods still find it difficult to match the accuracy achievable with fully-supervised learning methods in the disambiguation of relatively fine-grained polysemy (respective accuracies of roughly 66% vs. 84%). They also tested the expectation maximization (EM) algorithm [31] for unsupervised learning, which resulted in an accuracy of about 63%.

Linguistic behavior Justeson and Katz [60] proposed a method to automatically acquire training examples for the disambiguation of adjective senses. As described in Section 2.2.1, in this case, polysemy is projected onto adjective antonyms. Their acquisition method uses the following three principles:

- (a) antonyms often co-occur during direct comparison, in contrastive opposition,
- (b) antonyms are frequently joined by *and* or *or*,
- (c) antonyms frequently appear in noun phrases joined by prepositions and with the same head noun.

Each principle can be understood through the following example sentences presented in their paper [60], which contain the adjectives *old* (“not new”/“not young”) and *hard* (“not easy”/“not soft”).

- (a) They indicated that no *new errors* were being made and that all *old errors* would be corrected within 60 days.
- (b) It was pitiful to see the thin ranks of warriors, *old* and *young*...
- (c) But there is no sudden transition from *hard* to *soft*.

The corresponding principles annotate the instances of *old* and *hard* in sentences (a), (b) and (c) with the senses “not new”, “not young” and “not soft”, respectively. These training examples are used to identify indicator nouns (see Section 2.2.1). While this method successfully performs on these examples, the following issues must be considered. First, the above principles can be applied only to adjective sense disambiguation in which senses are mapped onto concepts associated with adjective antonyms. Second, the coverage of this method is relatively small: Justeson and Katz used a corpus consisting of 1.5 million sentences, and could obtain only about 1500 co-occurrences of a target adjective and antonym. Besides this, indicator nouns identified from these co-occurrences could disambiguate only about 27% of (open) test inputs.

Machine readable dictionaries Dictionaries provide definitions (and in some cases example sentences) for each word sense, which contain a number of “clue words”. Let us take the following sentences as example definitions for different senses of the word *bank* (“river edge”/“financial institution”):

- (3) a. land along side of river/lake. (“river edge”)
- b. place money kept. (“financial institution”)

As can be seen, these definitions contain clue words, such as *river* or *money*, which are intuitively associated with the respective senses of *bank* given above. Supposing a given input contains *bank* and *river*, one can easily select the former sense for the interpretation of the input *bank*. A number of methods follow this intuition [26, 30, 47, 63, 79, 108, 151]. Given an input, these methods generally compute the number of (clue) words appearing in the input and definitions, as the score for each sense candidate. Thereafter, they select the sense with the maximum score. A variation of these methods is to normalize the score by the length of the input.

Methods of this type can be categorized as instances of similarity-based methods (see Section 2.2.1), because the score represents the similarity between the input and each sense candidate²². Unlike supervised methods, these methods do not involve an excessive overhead for

²²Other methods, such as the Naive-Bayes method, can also use clue words as a feature set.

large-scaled sense annotation²³.

However, the quality and quantity of definitions, which highly influence performance, are problematic for these approaches. Definitions are relatively arbitrary, and thus different dictionaries often provide different definitions for the same sense. This problem is less crucial for human users because they can rely on external knowledge. However, lacking sufficient knowledge, computers can potentially fail to perform correctly given undesirable definitions. As for the quantity problem, it is hard to obtain adequate statistical information because definitions usually contain small numbers of words, or in some cases simply comprise synonym words. The method proposed by Luk [83], in which an outside corpus (the Brown corpus) was combined to enhance the statistics, would be one solution to this problem. Co-occurrence statistics were obtained in terms of a “control vocabulary” obtained from LDOCE [114]. However, this method requires additional manual compilation of the control vocabulary set (pruning rarely occurring words, for example). Okumura and Matsunaga [108] also overcame the quantity problem by expanding definitions through the EDR thesaurus [58]. However, these general methods do not directly address the quality problem. To sum up, we claim that the use of dictionary definitions is not a stand-alone approach: definitions can be used as an initial resource, but must be combined with additional (supervised) training data. This claim was experimentally validated, as described in Chapter 3.

2.3 Evaluation Methodology

From a scientific point of view, performance evaluation is invaluable. The procedure to evaluate a dedicated method is fundamentally to simulate run-time usage of the method by providing a corpus as training/test data. However, the performance of corpus-based methods is generally strongly biased by the training data provided, as well as test data. To minimize this bias, most experiments iteratively carry out the same trials, and average the results derived from each trial. In each trial, a fixed number of examples are *randomly* selected from the corpus, as the training and test data, respectively. This evaluation methodology can be called “cross validation”. Strictly speaking, the performance of methods under evaluation is inherently biased by the given corpus. However, since collecting broader coverage sense-annotated corpora requires tremendous human labor, most researchers have conducted experiments using relatively smaller-sized corpora at hand at the time. Another problem is that since interpretation of “performance” seems

²³Strictly speaking, human lexicographers need to provide sense definitions, a process which is associated with considerable overhead. However, given that a number of machine readable dictionaries are currently available, word sense disambiguation system developers are virtually released from the task of establishing sense definitions.

to vary according to the individual, a number of separate evaluation criteria exist. The following sections detail evaluation methodologies used in past word sense disambiguation methods.

2.3.1 Common evaluation criteria

One common evaluation criterion is the degree to which a method can be applied to run-time inputs, that is, the “coverage”²⁴. This criterion is crucial especially for qualitative approaches (see Section 2.2.1), because naive rules often reject all the sense candidates and therefore no decision can be made. Additionally, the ratio between the number of cases where the correct decision was made and the total number of decisions made, termed as “accuracy”, is commonly used. These two criteria can be summarized as in Equations 2.18.

$$\begin{aligned} \text{coverage} &= \frac{\# \text{ of decisions made}}{\text{total } \# \text{ of inputs}} \\ \text{accuracy} &= \frac{\# \text{ of correct decisions made}}{\text{total } \# \text{ of decisions made}} \end{aligned} \tag{2.18}$$

It should be noted that these criteria can also be applied to quantitative approaches. Given the degree of confidence about the decision, the method can *purposefully* refuse to make a risky decision. A number of instances of this principle can be found in the evaluation of quantitative word sense disambiguation methods. Dagan and Itai [28] estimated the degree of confidence by the statistical significance of the given training data [153]²⁵. Li et al. [81] empirically defined the relation between their heuristic rules and the confidence degree. In these cases, coverage and accuracy are contradictory criteria. However, most recent experiments have commonly employed a “backing-off” strategy (for example, random choice) when no decision can be made, resulting in a coverage of virtually 100% and performance being evaluated simply through accuracy. One plausible reason for this approach is that a single criterion is more intuitively understandable than multiple criteria. In fact, in the case of Li et al. [81] mentioned above, the trade-off between coverage and accuracy can be interpreted as the validation of the definition of confidence degree, rather than the performance of the method itself.

One may notice that coverage and accuracy resemble “recall” and “precision”, respectively, which have been commonly used in performance evaluation of information retrieval (IR) and text categorization (TC) tasks. In fact, the definition of recall/precision in experiments carried by Ribas [124], for example, is almost the same as that of coverage/accuracy, as given above.

²⁴The term “coverage” can have different interpretations. One example is the degree to which a system restricts the type of input (speech/on-line text). However, we confine interpretation of coverage to the narrower meaning given in the text.

²⁵Dagan and Itai [28] used the criteria of “applicability” and “precision”, which correspond to “coverage” and “accuracy”, respectively, in this section.

Let us devote a little more space to explaining the notion of recall and precision. In the case of IR, recall favors systems that retrieve as many documents salient to a user query as possible (disregarding noise contained in the retrieved data), while precision favors systems that retrieve as few irrelevant documents as possible. As can be seen, when all the documents are retrieved, recall is always 100%, potentially sacrificing precision. On the other hand, in the case of TC, recall favors systems that assign as many correct categories to each document as possible, while precision favors systems that assign as few incorrect categories to each document as possible. Note that since a single document can usually be associated with multiple categories according to different points of view, evaluation relying on a single criterion, i.e. accuracy, is less effective. The notion of these two contradictory criteria can be generalized as follows. Let us assume a situation in which a subject has to answer “yes” or “no” to N given questions (correct answers are not given to the subject, of course). The results can be classified into four cases as shown in Table 2.8, which is called a “contingency table”. In this table, capital letters (from A to D) denote the number of instances associated with each case. Correspondingly, recall and precision are defined as in Equation 2.19.

$$\begin{aligned} \text{recall} &= \frac{A}{A+C} \\ \text{precision} &= \frac{A}{A+B} \end{aligned} \tag{2.19}$$

Note that in the case of IR, a question corresponds to a document, and “yes” means that the document is retrieved. Therefore, $A+B$ and $A+C$ denote the number of retrieved documents and salient documents, respectively. On the other hand, in the case of TC, a question corresponds to each combination of document and category. Thus, $A+B$ and $A+C$ denote the number of categories assigned to documents and correct categories assigned to documents, respectively. We may note, in passing, that to integrate recall and precision, “F-measure”, which is expressed by Equation 2.20, is often used as an evaluation criterion.

$$\text{F-measure} = \frac{(\beta^2 + 1) \cdot \text{recall} \cdot \text{precision}}{\beta^2 \cdot \text{precision} + \text{recall}} \tag{2.20}$$

Here, β is a parametric constant used to control the preference between recall and precision. One may notice that as a type of categorization task, word sense disambiguation can equally be evaluated as performed for text categorization. However, most researchers coarsely assign a *single* sense to each word, and seem to prefer accuracy as the evaluation criterion.

Finally, let us describe several entropy-driven criteria, which have not been used for the past evaluation of word sense disambiguation. As Resnik and Yarowsky [123] identified, “cross-

	correct	yes	no
subject			
yes		A	B
no		C	D

Figure 2.8: Contingency table

entropy”, which has commonly been used for the evaluation of speech recognition systems, can also be used to evaluate probabilistic word sense disambiguation systems. Intuitively speaking, this criterion allocates a higher score when the system assigns a higher probability to a correct answer. Kononenko and Bratko [73] consider the prior distribution of candidate classes as well as the posterior distribution, that is, a higher score is allocated when the system improves on the prior probability for the correct answer. This criterion is expected to make it easier to compare results derived from different test data collections. The same motivation can be found in the evaluation methodology proposed by Gale et al. [42], in which the system performance was compared with the lower bound performance gained through the prior distribution of word senses.

2.3.2 Multiple human judgement

In using any of the evaluation criterion described in Section 2.3.1, the system decision must be compared with the correct answer. In most cases, human judgement provides the “correct” answer. However, there are cases where even human experts cannot correctly perform word sense disambiguation, because of differing interpretation of word senses between individuals [42, 118]. Consequently, system evaluation is biased by human judgement. One can expect to minimize this bias through the use of multiple human experts, that is, by validating the judgement of one human through evaluation by other humans. Although the associated cost is greater than for a single human judge, there have been a number of past cases of multiple judges being used in word sense disambiguation research²⁶.

Ravin [118] asked multiple humans (including both linguists and non-linguists) to disambiguate the test data which was intended for subsequent system evaluation, and used only those inputs where humans agreed on the answer. Gale et al. [42] also used multiple human judges to estimate the upper bound on their system’s performance. By “upper bound”, they meant human performance on disambiguation, which can be evaluated by way of comparison with

²⁶Note that multiple human judges have also been used in the evaluation of other applications, such as speech recognition systems.

other human judgements. To estimate the upper bound more reliably, they conducted the same trials by iteratively changing the human under evaluation, to average the performance of different humans. Hatzivassiloglou and McKeown [49] used multiple human judges to *directly* score individual decisions made by their system (although they used this methodology to evaluate the performance of adjective clustering). Intuitively speaking, their system was attributed a higher score when it made a correct decision which more humans agreed on the answer to.

2.3.3 Empirical comparison

Empirical (quantitative) comparison of a proposed method with various existing methods is a common evaluation methodology, assuming some well-defined evaluation criterion. In most cases, proposed (new) methods are compared with lower bound techniques, that is, in the case of word sense disambiguation, to systematically choose the most frequently appearing sense in the training data [42]. One aspect lacking in the empirical comparison of word sense disambiguation systems is that there is no contest for systems, as exists for information extraction systems in the Message Understanding Conferences (MUCs). However, a number of extensive comparisons involving a relatively large number of different methods can be found in past experiments. Let us devote a little space to describe these experiments, which fundamentally used the accuracy as a evaluation criterion (some of the below disambiguation methods are described in Section 2.2.1).

In the case of word sense disambiguation for English, the ACL/DCI Wall Street Journal corpus [86], in which each word is annotated with a single sense defined in LDOCE [114] or WordNet [93], has commonly been used. Leacock et al. [78] compared the vector space model, the Naive-Bayes method and neural networks, and reported that the Naive-Bayes method is marginally outperformed by the two other methods. Mooney [96] reported that the Naive-Bayes method outperformed, variously, neural networks, C4.5, k -NN ($k = 3$), and variations of the “FOIL” rule induction algorithm [116]. An experiment conducted by Ng [103] showed that “PEBLS” [25] and the Naive-Bayes method are comparable in performance. Pedersen and Bruce [110] compared three unsupervised methods: two automatic clustering algorithms (namely, McQuitty’s similarity analysis [89] and Ward’s minimum-variance method [148]) and the expectation maximization (EM) algorithm [31]. In the experimental results, McQuitty’s similarity clustering outperformed the other two methods. Pedersen and Bruce [111] reported that C4.5 slightly outperformed all of the Naive-Bayes method, PEBLS, CN2 and variations of their decomposable model [13, 112].

The relative results derived for these above experiments differ because each method involves separate parameters, which can vary between experiments. For example, a feature set, which

would be inconsistent between most of the methods above, is highly influential in performance.

Fujii et al. [39] compared a number of different methods using the EDR Japanese corpus [58]. According to their results, the performance of the compared methods can be sorted in descending order as follows: the thesaurus-driven similarity-based method²⁷, the vector space model, the Naive-Bayes method and the qualitative method using selectional restriction. To derive more general results, Fujii et al. [39] used a different type of criterion along with the accuracy. Instead of the conventional binary (correct/incorrect) judgement (i.e. accuracy), they applied an evaluation criterion based on the semantic similarity between word sense candidates, as advocated by a number of researchers [82, 123]. To exemplify this notion, let us take the polysemous Japanese verb *tsukau*, which has multiple senses in EDR, such as “to employ”, “to operate”, “to spend” and “to use MATERIAL”. These senses are associated with the EDR thesaurus [58]. Figure 2.9 shows a fragment of the thesaurus, in which an oval denotes a word sense. As with most thesauri, the length of the path between two word senses can be seen as the relative semantic distance between them. As one can see, the verb sense “to spend” is physically closer to “to use MATERIAL” than to “to employ” or “to operate”, in structure. In fact, these two proximal verb senses can be merged into one common category, that is, “to use up”. Furthermore, they can be merged with “to operate” to form “to use PHYSICAL OBJECT”, as distinct from the remaining sense of “to employ (HUMAN/CONCEPT)”. Let the correct sense of an input *tsukau* be “to use MATERIAL”, and assume the system incorrectly outputs “to spend”. In this case, the error would be more acceptable than outputting “to employ”, which is totally different to the correct interpretation of “to use MATERIAL”. Therefore, we should allot the system a scaled acceptability factor instead of a score of zero. Note that the binary judgment simply scores the system zero, irrespective of the extent to which the error is acceptable given a particular practical application. Fujii et al. [39] formalized the scaled acceptability factor (“acceptability”, hereafter) as in Equation 2.21.

$$A(x, s) = \left(\frac{MAXLEN - EDR(x, s)}{MAXLEN} \right)^\alpha \quad (2.21)$$

Here, x and s are the system’s interpretation and the correct answer, respectively, and $A(x, s)$ is the acceptability of the given x and s . $EDR(x, s)$ represents the length of the path between x and s in the EDR thesaurus. $MAXLEN$ is the maximum length of the path between senses associated with an input verb. For example, $MAXLEN = 7$ in Figure 2.9 (the length of the path between “to operate” and “to employ”). α is a parametric constant, which acts to control

²⁷Fujii et al. [39] used the *Bunruigoihyo* thesaurus [102] for the similarity computation.

the reduction factor of $A(x, s)$ for incorrect interpretations. With a larger α , $A(x, s)$ becomes smaller for incorrect interpretations, and becomes closer to the binary judgment. One can notice that the acceptability ranges from 0 (where x and s are the most dissimilar verb senses) to 1 (x and s are identical). The result derived through the acceptability with different α values showed the same tendency as the accuracy.




Figure 2.9: A fragment of the EDR thesaurus including senses related to the Japanese verb *tsukau*

2.3.4 Further considerations

This section discusses several further issues which should be noted for the evaluation of word sense disambiguation systems.

The first issue is that the granularity of word senses differs depending on the viewpoint. This problem is closely related to the problem that different viewpoints result in different distinctions of word senses (as discussed in Section 2.3.2). Consequently, the performance of the system is likely to be degraded when finer-grained sense definitions are used for the evaluation. Conversely, the system would generally perform better for applications that require coarser-grained sense definitions. The best way to minimize this bias is to use multiple word sense definition paradigms, with different levels of granularity. Lin [82] used multiple granularity levels by gradually relaxing sense distinction based on the taxonomy of WordNet [93]. To put it plainly, one can simulate coarse-grained sense definition by regarding distinct senses dominated by the same parent node (synset) in the taxonomy as the same sense. The experimental results showed that Lin’s method outperformed the lower bound method as the granularity of sense distinction was lessened. However, at the same time, multiple evaluation criteria can complicate the interpretation of the performance.

Second, one may notice that the most straightforward evaluation type is “task-based evaluation”, where the criterion is the degree to which the system improved on the performance

of another application. However, a surprisingly small number of cases of task-based evaluation can be found in past word sense disambiguation research. Brown et al. [10] improved on the performance of French-English machine translation by introducing their word sense disambiguation method. Fukumoto and Suzuki [40] reported that application of word sense disambiguation improved on the performance of text categorization²⁸. Voorhees [147] used word sense disambiguation (with respect to WordNet senses) to enhance query terms for information retrieval (IR), which unfortunately did not improve the performance of a conventional IR system. One reason for this result would be the rudimentary nature of existing word sense disambiguation methods, especially for fine-grained sense definitions. This poses the question as to what degree word sense disambiguation should be correctly performed for operational applications. Experiments conducted by Sanderson [128] aimed at an IR system application provide a pointer to the answer to this question. Sanderson carried out experiments as follows. First, due to the lack of large-scaled sense-annotated data, artificial polysemous words, namely “pseudo-words” [130, 155] are created, based on words appearing in a given document collection. For example, distinct words like *guerrilla* and *reptile* are considered to be two distinct senses of a pseudo-word. In other words, pseudo-words simulate extremely coarse-grained lexical ambiguity. Based on this, one can artificially alter the performance of (virtual) word sense disambiguation systems. For example, a system with 50% accuracy can be simulated by *knowingly* selecting incorrect senses for half of the pseudo-words contained in the document collection. The results show that (a) a theoretical perfect WSD system improved on the baseline performance of the IR system (for which polysemy of pseudo-words remains unresolved), and (b) WSD system with 90% accuracy rate did not improved on the baseline performance, which means, at a bare minimum, more than 90% accuracy is required to enhance IR systems. At the same time, we note that the applicability to practical issues of this conclusion derived from the artificial evaluation, is arguable.

2.4 Related NLP Research

Let us discuss NLP research which is expected to enhance word sense disambiguation.

Morphological analysis As reviewed in Section 2.2.1, past word sense disambiguation methods generally rely on the morphological content of words in the input (or training data), such as the root form or part-of-speech. In addition, in the case of agglutinative languages, such as

²⁸Schütze [129] improved the quality of information retrieval through word sense *discrimination*, which differs from word sense disambiguation in that the former does not use predefined sense candidates.

Chinese and Korean, lexical segmentation is also poignant²⁹. Note that numerous types of morphological analyzers (in various languages) with reasonable performance have been established for easy access.

Syntactic analysis Syntactic analysis can enhance the performance of word sense disambiguation methods which rely on syntactic relations (associated with polysemous words). This analysis is especially poignant when the input comprises a complex sentence. Note that given the fact that full parsing is still expensive, a number of methods for partial syntactic analysis have also been applied to identify syntactic relations associated with polysemous words [19, 39, 159]. For Japanese verb sense disambiguation, Fujii et al. [39] used two simple heuristics for extracting verb complements³⁰, rather than syntactic analysis on the whole input sentence. Note that verb complements are useful features for verb sense disambiguation. The heuristics utilized are given below:

- each complement is associated to the verb of highest proximity,
- complements containing the genitive case marker *no* are not considered because they can constitute either possessive or nominative case markers, and are thus confusing.

They reported that the performance of verb sense disambiguation combined with these heuristics is comparable to that combined with a full syntactic analysis. In other words, the overhead required for syntactic analysis can be reduced without degrading the system performance.

Identification of idiomatic expressions Idiomatic (fixed or freezed) expressions, in which a specific word collocation pattern stands for a certain meaning, is one obstacle to word sense disambiguation. For example, similarity-based systems generally fail to interpret inputs where polysemous words comprise elements of idiomatic expressions, by semantic overgeneralization through the use of a thesaurus. Possible solutions would include one proposed by Uramoto [142], in which idiomatic expressions are described separately in the database so that the system can control their overgeneralization. At the same time, given the fact that there is no universal consensus as to idiomatic expression types, automatic methods to identify idiomatic expressions [20, 134] may also prove useful.

Discourse analysis The discourse or domain dependency of each word sense, as been used in a small proportion of past word sense disambiguation techniques [47, 101, 157], can provide

²⁹Continuous speech inputs pose the same problem for non-agglutinative languages.

³⁰In Japanese, a verb complement consists of noun phrase (case filler) and postposition (case marker suffix). See Section 3.2 for details.

strong external constraints. Given a single discourse segment, polysemous words tend to coincide in interpretation. Anaphora/ellipsis resolution [48, 65, 97] is also expected to enhance the contextual information. Let us take the following example, in which the pronoun *it* refers to *taxi* in the previous sentence:

- (4) A taxi is coming. Let's take *it*.

One may notice that correct anaphora resolution would make it easier to disambiguate the inherent polysemy of *take* in the second sentence.

Establishment of linguistic resources As Kilgariff [68] identified, lexicography is both a benefactor and beneficiary for word sense disambiguation. That is, dictionaries provide sense candidates for polysemous words and clue words for each sense. In addition, automatic identification of word senses [41, 144, 158] (perhaps combined with lexicography) is expected to avoid human bias in sense distinction.

Establishment of thesauri is also useful for a number of thesaurus-driven word sense disambiguation methods [77, 81, 82, 124, 125, 142, 154]. Conventional methods for automatic thesaurus construction have utilized dictionary definitions to extract IS-A or hyper/hypo relations [51, 92, 100]. However, given the number of electronic thesauri currently available, the recent trend seems to be to focus on enhancement of existing thesauri (not development *from scratch*). Besides this, given the fact that most thesauri aim at general purpose applications, they need to be adjusted for a particular usage [7, 8]. Tokunaga et al. [138] and Uramoto [143] (independently) extended existing thesauri by appropriately positioning unregistered words in the taxonomy. Hearst and Schütze [52] proposed a method to adjust the general taxonomy defined in WordNet [93] for particular usages (although they targetted the application toward text categorization systems).

Finally, the establishment of a sense-annotated corpus as a benchmark collection is inevitable for the standardized evaluation of word sense disambiguation methods [94, 95].

2.5 Discussion

Pustejovsky and Boguraev [115] identified one limitation of current word sense disambiguation methods in their dependence on fixed numbers of predefined word senses, and that therefore a *novel* word usage cannot be considered. This limitation can be especially problematic when one targets an *all-purpose* system, because exhaustive sense enumeration seems practically impossible. However, focusing on a particular domain (sublanguage), where vocabulary size is

relatively limited, this problem is expected to be resolvable. One may argue that compilation of word senses for different domains poses a considerable overhead. However, automatic word sense identification [41, 144, 158] is expected to reduce this overhead.

In addition, Basili et al. [6, 7] proposed a method to *tune* word senses taken from a general purpose lexicon (in their case, WordNet [93]) to a specific sublanguage. They first (manually) identified some of higher level word classes in WordNet [93] as “kernel senses” (25 for nouns and 15 for verbs), for each of which representative verbs are automatically identified. Then, a sublanguage corpus is used to train statistical classifiers for each kernel sense³¹. Finally, statistical classifiers evaluate the membership of each (prospective) polysemous word to the kernel senses. Consequently, novel senses are acquired and irrelevant senses specified in the sublanguage are discarded.

Thesaurus extending methods [138, 143] can also potentially resolve the novel sense problem. Tokunaga et al. [138] used the “SVMV model” [57] to train probabilistic classifiers of thesaurus nodes (they use the *Bunruigoihyo* thesaurus [102] as the core thesaurus). Each node is represented by the co-occurrence vector associated with words belonging to the node. Given a new word unlisted in the thesaurus (which is also represented by a vector), classifiers compute the probability that the word belongs to each node, and the word is positioned in the node with maximum probability. One may notice that this technique is equal to positioning a novel usage of a word in a conceptual hierarchy (such as WordNet [93] or EDR [58])³².

At the same time, we concede that this issue still needs to be further explored, and do not pretend to draw any premature conclusions in this research.

2.6 Summary

In this chapter, we surveyed past research associated with word sense disambiguation. First, we elaborated on different classes of corpus-based word sense disambiguation methods, which consists of qualitative approaches (selectional restrictions, decision trees and decision lists) and quantitative approaches (probabilistic models and statistically/thesaurus-driven similarity-based methods). We identified that one critical content is selection of useful features (clues) for word sense disambiguation. In addition, the performance of the similarity methods strongly depends on the way to compute similarity between given examples. We also identified that past methods

³¹Basili et al. [6, 7] used a statistical method proposed by Yarowsky [154] for the classifier training.

³²One may argue that, strictly speaking, the number of word classes (senses) defined in a thesaurus is already limited, that is, novel senses undefined in thesauri cannot be considered. The answer to this problem remains as an open question.

generally use corpus, where each polysemous word is associated with correct sense, as source knowledge. Past experiments showed that large-scaled corpora are required for operational systems. In other words, a considerable overhead is required when one tries to supervise, i.e. manually annotate word senses to large corpora. In view of this problem, we second investigated how past research has tried to minimize the overhead required for human supervision. While past methods reportedly successfully minimized (or completely excluded) the overhead for supervision, we identified that applicability of past methods seems to be limited to a certain specified usage, such as machine translation oriented system [28], coarse-grained word sense distinction [63, 130, 157], and disambiguation of adjective senses [60]. Third, from a scientific point of view, we described evaluation methodologies for word sense disambiguation, which turned out that so far there have been no standardized evaluation criterion. This issue needs to be further explored. Fourth, we described a numerous NLP research which is expected to enhance word sense disambiguation. Finally, we discussed a limitation of current word sense disambiguation and a possible solution for this problem.

In the following chapters, we will tackle some of identified problems. Chapter 3 describes overall design of our verb sense disambiguation system, in which we propose a method to weigh the degree to which each feature contributes to disambiguation. Chapter 4 proposes novel methods to minimize the overhead required for our system, i.e. the overhead for supervision on large corpora and the overhead for searching large corpora. Finally, Chapter 5 explores the similarity computation integrating hand-crafted thesaurus and statistical information, which is expected to enhance the similarity-based word sense disambiguation.

Chapter 3

Verb Sense Disambiguation System

This chapter describes a similarity-based method verb sense disambiguation system. The disambiguation process follows the nearest neighbor method. Given an input consisting of a polysemous verb and its governing case complements (or cases), the system searches a database for the example most similar to the input. Thereafter, the polysemy of the verb is disambiguated by superimposing the sense of the verb associated with the retrieved example. The system uses a hand-crafted thesaurus or alternatively co-occurrence statistics, for the similarity computation. The main interest in this chapter lies in the introduction of the notion of ‘case contribution to disambiguation’ (CCD), which is a weighting schema for each case in similarity computation. Intuitively speaking, greater diversity of semantic range of case filler examples will lead to that case contributing more highly to verb sense disambiguation. We also report the results of comparative experiments, in which the performance of disambiguation is improved by considering the CCD factor.

3.1 Overall Design

This chapter elaborates on our proposed similarity-based verb sense disambiguation system¹ and its evaluation through experimentation. Although our system is currently implemented for Japanese word sense disambiguation, the methodology can theoretically be applied to other languages. Let us briefly explain the disambiguation process based on Figure 3.1. In this figure, “input” denotes a (complex/simple) sentence containing polysemous verb(s). Then, “morpho-

¹We originally used the term “example-based” method [34, 35, 36, 37, 38]. However, we shall use the more general “similarity-based” terminology in this research, because no explanatory distinction between similar approaches, such as the case-based and exemplar-based reasoning, has been identified.

logical and syntactic analyzer” extracts predicate-argument structures (“pred-args”)² from the input. It should be noted that in the case of Japanese, morphological analysis entails lexical segmentation as well as part-of-speech tagging. Thereafter, given the pred-args, the core of the verb sense disambiguation system (“VSD”) outputs plausible sense(s) for each verb in the pred-args. In this process, the VSD uses supervised examples (in the “database”) and “thesaurus”/“co-occurrence” data. The VSD additionally outputs the degree on certainty of its decision (“certainty”). Currently, we rely on existing NLP tools for the morph/syntax analyzer, and this component of the system is beyond the scope of this research. In fact, for Japanese, a number of existing tools such as JUMAN (a morphological analyzer) [88] and QJP (a morphological and syntactic analyzer) [62] have performed promisingly. Beside these tools, a number of simple heuristics to identify predicate-argument structures have been proposed for verb sense disambiguation [19, 39], in an attempt to minimize the overhead for syntactic analysis. To sum up, the focus of this chapter can be identified by the portion of the overall system enclosed within the dashed region of Figure 3.1. Hereafter, “system” refers to the VSD, and we shall use the term “inputs” to refer to the inputs which are morphologically and syntactically analyzed (in other words, predicate-argument structures).

Section 3.2 overviews the basis of our verb sense disambiguation system, and then Section 3.3 elaborates on the disambiguation methodology. Section 3.4 evaluates our our system by way of experiments. Section 3.5 describes a way to further enhance our system, including the computation of the certainty degree.

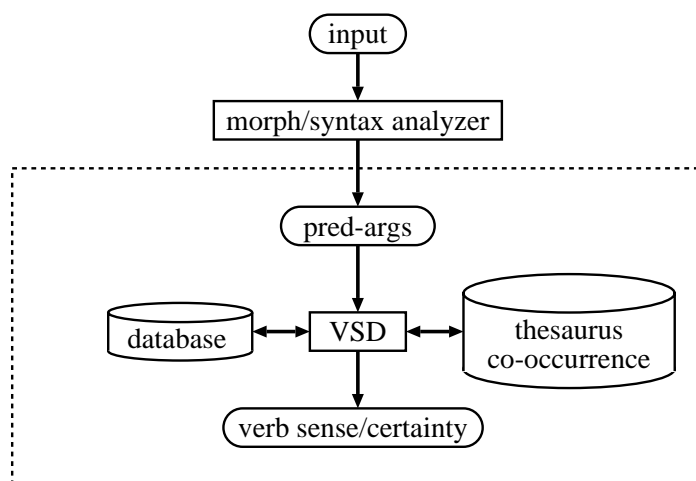


Figure 3.1: The overall design of the verb sense disambiguation system

²We shall interchangeably use the terms “predicate-argument” and “verb-complement” structures.

3.2 Basic Idea

The core mechanism in our verb sense disambiguation system is based on the method proposed by Kurohashi and Nagao [77] and later enhanced by Fujii et al [37, 38]. The system uses an example-database (database, hereafter) containing examples of collocations for each verb sense and its associated case frame(s). It should be noted that while the term “case” has been used variously by different researchers, including surface and deeper case level senses, we shall consistently use “case” with reference to the surface level³. Figure 3.2 shows a fragment of the entry associated with the Japanese verb *toru*. The verb *toru* has multiple senses, a sample of which are “to take/steal”, “to attain”, “to subscribe” and “to reserve”. The database specifies the case frame(s) associated with each verb sense. In Japanese, a complement of a verb consists of a noun phrase (case filler) and its case marker suffix, for example *ga* (nominative) or *wo* (accusative). The database also lists several case filler examples for each case. In practice, the database is usually compiled from machine readable dictionaries (MRDs) and text corpora. The task of the system is “to interpret” the verbs occurring in the input text, i.e. to choose one sense from among a set of candidates⁴.

<i>toru:</i>		
$\left\{ \begin{array}{l} \textit{suri} \quad (\textit{pickpocket}) \\ \textit{kanojo} \quad (\textit{she}) \\ \textit{ani} \quad (\textit{brother}) \end{array} \right\} \textit{ga}$	$\left\{ \begin{array}{l} \textit{kane} \quad (\textit{money}) \\ \textit{saifu} \quad (\textit{wallet}) \\ \textit{otoko} \quad (\textit{man}) \\ \textit{uma} \quad (\textit{horse}) \\ \textit{aidea} \quad (\textit{idea}) \end{array} \right\} \textit{wo}$	<i>toru</i> (to take/steal)
$\left\{ \begin{array}{l} \textit{kare} \quad (\textit{he}) \\ \textit{kanojo} \quad (\textit{she}) \\ \textit{gakusei} \quad (\textit{student}) \end{array} \right\} \textit{ga}$	$\left\{ \begin{array}{l} \textit{menkyoshou} \quad (\textit{license}) \\ \textit{shikaku} \quad (\textit{qualification}) \\ \textit{biza} \quad (\textit{visa}) \end{array} \right\} \textit{wo}$	<i>toru</i> (to attain)
$\left\{ \begin{array}{l} \textit{kare} \quad (\textit{he}) \\ \textit{chichi} \quad (\textit{father}) \\ \textit{kyaku} \quad (\textit{client}) \end{array} \right\} \textit{ga}$	$\left\{ \begin{array}{l} \textit{shinbun} \quad (\textit{newspaper}) \\ \textit{zasshi} \quad (\textit{journal}) \end{array} \right\} \textit{wo}$	<i>toru</i> (to subscribe)
$\left\{ \begin{array}{l} \textit{kare} \quad (\textit{he}) \\ \textit{dantai} \quad (\textit{group}) \\ \textit{ryokoukyaku} \quad (\textit{passenger}) \\ \textit{joshu} \quad (\textit{assistant}) \end{array} \right\} \textit{ga}$	$\left\{ \begin{array}{l} \textit{kippu} \quad (\textit{ticket}) \\ \textit{heya} \quad (\textit{room}) \\ \textit{hikouki} \quad (\textit{airplane}) \end{array} \right\} \textit{wo}$	<i>toru</i> (to reserve)
⋮	⋮	⋮

Figure 3.2: A fragment of the database, and the entry associated with the Japanese verb *toru*

All verb senses we use are defined in the machine readable dictionary “IPAL” [56] (which

³Conventional case systems for natural language are reviewed, for example, by Bruce [12].

⁴Note that unlike the automatic acquisition of word sense definitions [41, 115, 144, 158], the task of the system is to identify the best matched category with a given input, from *predefined* candidates.

parallels that of Kurohashi and Nagao [77]). Let us devote a little space to explaining IPAL. IPAL lists about 900 Japanese verbs and categorizes each verb into multiple “subentries”, based on verbal syntax and semantics. In our case, subentries are equivalent to verb senses. Figure 3.3 shows a fragment of a subentry for the Japanese verb *kakeru*, including the subentry number, definition of the sense, case pattern, semantic features for case fillers, and example case filler nouns. The “case pattern” entry defines the obligatory cases associated with the verb sense, in which N_i represents the i -th case filler⁵. In addition, optional cases are indicated in parentheses (such as N_2 -*ni* in Figure 3.3). IPAL defines 19 semantic feature types, which comprise a hierarchy, as shown in Figure 3.4. In other words, these semantic features can be used as selectional restrictions (see Section 2.2.1). However, according to experiments conducted by Kurohashi and Nagao [77], the IPAL semantic features are not sufficiently fine-grained for rule-based verb sense disambiguation. Kurohashi and Nagao thus implemented their similarity-based method based on examples listed in each “example nouns” entry, and showed its effectivity over the rule-based method. Although Kurohashi and Nagao used only examples listed in IPAL, their method is open for the use of additional examples. In our case, we take from IPAL only the “case pattern” and “example nouns” entries to initialize the database, and later enhance the database with additional supervised examples.

# subentry	008	
definition	attach something to a body	
case pattern	N_1 - <i>ga</i> (N_2 - <i>ni</i>) N_3 - <i>wo</i>	
N_1	semantic feature	human
	example nouns	<i>kare</i> (he), <i>kanojo</i> (she)
N_2	semantic feature	parts
	example nouns	<i>kata</i> (shoulder), <i>te</i> (hand), <i>kubi</i> (neck)
N_3	semantic feature	products
	example nouns	<i>epuron</i> (apron), <i>megane</i> (glasses), <i>pendanto</i> (pendant)

Figure 3.3: An example of a subentry for the Japanese verb *kakeru*

Let us now turn to the process of verb sense disambiguation. Given an input, which consists of a polysemous verb and its governing complements, the system identifies the verb sense on the basis of the scored similarity between the input and the examples given for each verb sense. Let us look at Figure 3.2 again and take the sentence below as an example input:

- (5) *hisho ga shindaisha wo toru.*
 (secretary-NOM) (sleeping car-ACC) (?)

⁵Note that IPAL describes only the *typical* case order, and that word order in Japanese is less strictly restricted than other language such as English.

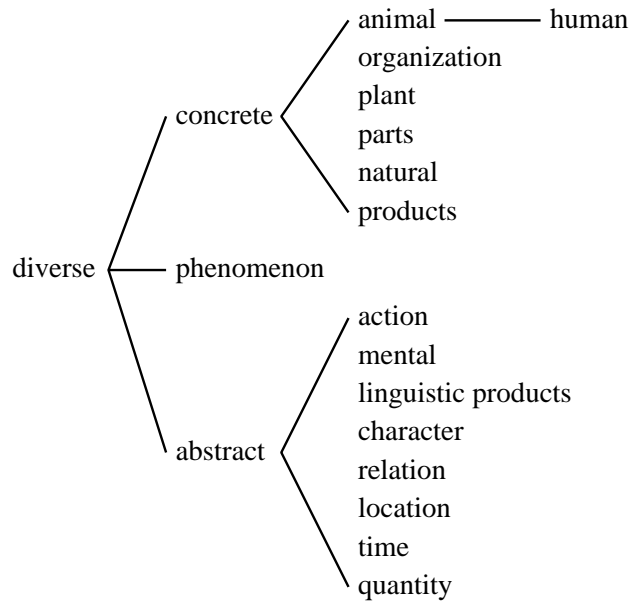


Figure 3.4: The hierarchy of IPAL semantic features

In this example, one may consider *hisho* (“secretary”) and *shindaisha* (“sleeping car”) to be semantically similar to *joshu* (“assistant”) and *hikouki* (“airplane”) respectively, and since both collocate with the “to reserve” sense of *toru*, one could infer that *toru* should be interpreted as “to reserve”. This resolution can be called the nearest neighbor method because the verb in the input is disambiguated by superimposing the sense of the verb appearing in the example of highest similarity⁶. As one can see, the similarity between an input and an example is estimated based on the similarity between case fillers marked with the same case. While the notion of nearest neighbor does not predefine the type of similarity measurement used, we will explain two different types of similarity measurement in the following sections.

Furthermore, since the restrictions imposed by the case fillers in choosing the verb sense are not equally selective, we introduce a weighted ‘case contribution to disambiguation (CCD)’ of the verb senses. Let us consider another example input:

- (6) *gakusei ga shuukanshi wo toru.*
 (student-NOM) (magazine-ACC) (?)

The nominative, *gakusei* (“student”), in sentence (6) is found in the “to attain” case frame of *toru* and there is no other co-occurrence in any other sense of *toru*. Therefore, the nominative supports an interpretation “to attain”. On the other hand, the accusative, *shuukanshi*

⁶Hereafter, “similarity-based” systems basically refers to those which are based on the nearest neighbor method.

(“magazine”), is most similar to the examples included in the accusative of “to subscribe” and therefore the accusative supports another interpretation “to subscribe”. Although the most plausible interpretation here is actually the latter, the former would be chosen if one always relies equally on the similarity in the nominative and the accusative. However, in the case of *toru*, since the semantic range of nouns collocating with the verb in the nominative does not seem to have a strong delinearization in a semantic sense (in Figure 3.2, the nominative of each verb sense displays the same general concept, i.e. HUMAN), it would be difficult, or even risky, to properly interpret the verb sense based on the similarity in the nominative. In contrast, since the semantic ranges are disparate in the accusative, it would be feasible to rely more strongly on the similarity here.

This argument can be illustrated as in Figure 3.5, in which the symbols e_1 and e_2 denote example case fillers of different case frames, and an input sentence includes two case fillers denoted by x and y . The figure shows the distribution of example case fillers for the respective case frames, denoted in a semantic space. The semantic similarity between two given case fillers is represented by the physical distance between the two symbols. In the nominative, since x happens to be much closer to an e_2 than any e_1 , x may be estimated to belong to the range of e_2 's, although x actually belongs to both sets of e_1 's and e_2 's. In the accusative, however, y would be properly estimated to belong to the set of e_1 's due to the disjunction of the two accusative case filler sets, even though examples do not fully cover each of the ranges of e_1 's and e_2 's. Note that this difference would be critical if example data were sparse. One may argue that this property can be generalized as the notion that the system always relies only on the similarity in the accusative for verb sense disambiguation. Although some typical verbs show this general notion, it is not guaranteed for any kind of verb. Our approach, which computes the degree of contribution for each verb respectively, can handle exceptional cases as well as typical ones. We will explain the method used to compute CCD in Section 3.3.




Figure 3.5: The semantic ranges of the nominative and accusative for the verb *toru*

3.3 Methodology

To illustrate the overall algorithm, we will consider an abstract specification of both an input and the database (see Figure 3.6). Let the input be $\{n_{c_1-m_{c_1}}, n_{c_2-m_{c_2}}, n_{c_3-m_{c_3}}, v\}$, where n_{c_i} denotes the case filler for the case c_i , and m_{c_i} denotes the case marker for c_i , and assume that the interpretation candidates for v are derived from the database as s_1 , s_2 and s_3 . The database also contains a set \mathcal{E}_{s_i,c_j} of case filler examples for each case c_j of each sense s_i (“—” indicates that the corresponding case is not allowed).

input	$n_{c_1-m_{c_1}}$	$n_{c_2-m_{c_2}}$	$n_{c_3-m_{c_3}}$		$v(?)$
database	\mathcal{E}_{s_1,c_1}	\mathcal{E}_{s_1,c_2}	\mathcal{E}_{s_1,c_3}	—	$v(s_1)$
	\mathcal{E}_{s_2,c_1}	\mathcal{E}_{s_2,c_2}	\mathcal{E}_{s_2,c_3}	\mathcal{E}_{s_2,c_4}	$v(s_2)$
	—	\mathcal{E}_{s_3,c_2}	\mathcal{E}_{s_3,c_3}	—	$v(s_3)$

Figure 3.6: An input and the database

During the verb sense disambiguation process, the system first discards those candidates whose case frame does not fit the input. In the case of Figure 3.6, s_3 is discarded because the case frame of $v(s_3)$ does not subcategorize for the case c_1 ⁷. In contrast, s_2 will not be rejected at this step. This is based on the fact that in Japanese, cases can be easily omitted if they are inferable from the given context. Note that no omission of case fillers is allowed in the database for the reasonable system coverage. In the case of IPAL, example case fillers cover every slot.

In the next step the system computes the score of the remaining candidates and chooses as the most plausible interpretation the one with the highest score. The score of an interpretation is computed by considering the weighted average of the similarity degrees of the input case fillers with respect to each of the example case fillers (in the corresponding case) listed in the database for the sense under evaluation⁸. Formally, this is expressed by Equation 3.1, where $Score(s)$ is the score of sense s of the input verb, and $SIM(n_c, \mathcal{E}_{s,c})$ is the maximum similarity degree between the input case filler n_c and the corresponding case fillers in the database example set $\mathcal{E}_{s,c}$ (calculated through Equation 3.2). $CCD(c)$ is the weight factor of case c , which we will

⁷Since IPAL does not necessarily enumerate all the possible optional cases, the absence of case c_1 from $v(s_3)$ in the figure may denote that c_1 is optional. If so, the interpretation s_3 should not be discarded in this stage. To avoid this problem, we use the same technique as used in Kurohashi’s method. That is, we define several particular cases beforehand, such as the nominative, the accusative and the dative, to be obligatory, and impose the grammatical case frame constraints above only for these obligatory cases. Optionality of case needs to be further explored.

⁸ \mathcal{E}_{s_2,c_4} is not taken into consideration in the computation since c_4 does not appear in the input.

explain later in this section.

$$Score(s) = \frac{\sum_c SIM(n_c, \mathcal{E}_{s,c}) \cdot CCD(c)}{\sum_c CCD(c)} \quad (3.1)$$

$$SIM(n_c, \mathcal{E}_{s,c}) = \max_{e \in \mathcal{E}_{s,c}} sim(n_c, e) \quad (3.2)$$

It should be noted that while Equation 3.1 proposes one implementation of the score computation, there can be a number of variations. For example, an alternative model could not consider cases with diminished CCD values (for the verb under evaluation) in this computation. In fact, one may notice that introduction of the CCD factor can be characterized as one type of “feature selection”⁹: cases associated with greater CCD values consist of a useful feature set. In the extreme case, one can rely solely on the case with greatest CCD value¹⁰. Optimization of the model still remains as an open question and needs further exploration.

The following two sections detail the computation of similarity between case fillers, and the CCD factor.

3.3.1 Computation of similarity between case filler nouns

With regard to the computation of the similarity between two different case fillers ($sim(n_c, e)$ in Equation 3.1), we experimentally used two alternative approaches. The first approach uses semantic resources, that is, hand-crafted thesauri (such as Roget’s thesaurus [14] or WordNet [93] in the case of English, and *Bunruigoihyo* [102], EDR [58] or *Goi-Taikai* [55]¹¹ in the case of Japanese), based on the intuitively feasible assumption that words located near each other within the structure of a thesaurus have similar meaning. Therefore, the similarity between two given words is represented by the length of the path between them in the thesaurus structure [37, 38, 77, 81, 142]¹². We used the similarity function empirically identified by Kurohashi and Nagao [77], in which the relation between the length of the path in the *Bunruigoihyo* thesaurus and the similarity, is defined as shown in Table 3.1. In this thesaurus, each entry is assigned a 7 digit class code. In other words, this thesaurus can be considered as a tree, 7 levels in depth, with each leaf as a set of words. Figure 3.7 shows a fragment of the *Bunruigoihyo* thesaurus including some of the nouns in both Figure 3.2 and the input sentence above.

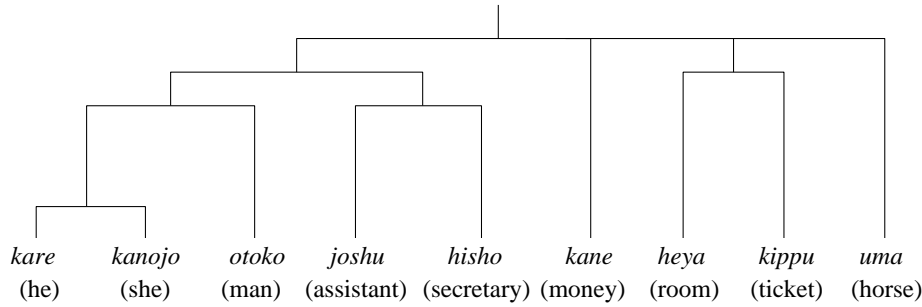
The second approach is based on statistical modeling. We adopted one typical implemen-

⁹See the item “probabilistic models” (p. 14) in Section 2.2.1 for past feature selection methods.

¹⁰Yarowsky [157], for example, advocated a simple implementation using only the single most useful feature.

¹¹*Goi-Taikai* is a relatively newly released Japanese dictionary.

¹²Different types of application of hand-crafted thesauri to word sense disambiguation have been proposed, for example, by Yarowsky [154] (see Section 2.2.2).

Figure 3.7: A fragment of the *Bunruigoihyo* thesaurusTable 3.1: The relation between the length of the path between two nouns n_1 and n_2 in the *Bunruigoihyo* thesaurus ($len(n_1, n_2)$), and their relative similarity ($sim(n_1, n_2)$)

$len(n_1, n_2)$	0	2	4	6	8	10	12
$sim(n_1, n_2)$	11	10	9	8	7	5	0

tation called the “vector space model” (VSM) [33, 78, 106, 127, 130], which has a long history of application in information retrieval (IR) and text categorization (TC) tasks. In the case of IR/TC, VSM is used to compute the similarity between documents, which is represented by a vector comprising statistical factors of content words in a document. Similarly, in our case, each noun is represented by a vector comprising statistical factors, although statistical factors are calculated in terms of the predicate-argument structure in which each noun appears. Predicate-argument structures, which consist of complements (case filler nouns and case markers) and verbs, have also been used in the task of noun classification [53]. This can be expressed by Equation 3.3, where \vec{n} is the vector for the noun in question, and items t_i represent the statistics for predicate-argument structures including n .

$$\vec{n} = \langle t_1, t_2, \dots, t_i, \dots \rangle \quad (3.3)$$

In regard to t_i , we used the notion of TF·IDF [127]. TF (term frequency) gives each context (a case marker/verb pair) importance proportional to the number of times it occurs with a given noun. The rationale behind IDF (inverse document frequency) is that contexts which rarely occur over collections of nouns are valuable, and that therefore, the IDF of a context is inversely proportional to the number of noun types that appear in that context. This notion is expressed by Equation 3.4, where $f(\langle n, c, v \rangle)$ is the frequency of the tuple $\langle n, c, v \rangle$, $nf(\langle c, v \rangle)$ is the number of *noun types* which collocate with verb v in the case c , and N is the number of noun

types within the overall co-occurrence data.

$$t_i = f(\langle n, c, v \rangle) \cdot \log \frac{N}{nf(\langle c, v \rangle)} \quad (3.4)$$

We then compute the similarity between nouns n_1 and n_2 by the cosine of the angle between the two vectors \vec{n}_1 and \vec{n}_2 . This is realized by Equation 3.5.

$$\text{sim}(n_1, n_2) = \frac{\vec{n}_1 \cdot \vec{n}_2}{|\vec{n}_1| |\vec{n}_2|} \quad (3.5)$$

We extracted co-occurrence data from the RWC text base RWC-DB-TEXT-95-1 [119]. This text base consists of 4 years worth of “Mainichi Shimbun” newspaper articles [85], which have been automatically annotated with morphological tags. The total morpheme content is about 100 million. Since full parsing is usually expensive, a simple heuristic rule was used in order to obtain collocations of nouns, case markers and verbs in the form of tuples $\langle n, c, v \rangle$. This rule systematically associates each sequence of noun and case marker to the verb of highest proximity, and produced 419,132 tuples. This co-occurrence data was used in the preliminary experiment described in Section 3.4¹³.

3.3.2 Computation of case contribution factor

In Equation 3.1, $CCD(c)$ expresses the weight factor of the contribution of case c to (current) verb sense disambiguation. Intuitively speaking, preference should be given to cases displaying case fillers which are classified in semantic categories of greater disjunction. As such, c 's contribution to the sense disambiguation of a given verb, $CCD(c)$, is likely to be higher if the example case filler sets $\{\mathcal{E}_{s_i, c} \mid i = 1, \dots, n\}$ share fewer elements as in Equation 3.6.

$$CCD(c) = \left(\frac{1}{C_2^n} \sum_{i=1}^{n-1} \sum_{j=i+1}^n \frac{|\mathcal{E}_{s_i, c}| + |\mathcal{E}_{s_j, c}| - 2|\mathcal{E}_{s_i, c} \cap \mathcal{E}_{s_j, c}|}{|\mathcal{E}_{s_i, c}| + |\mathcal{E}_{s_j, c}|} \right)^\alpha \quad (3.6)$$

Here, α is a constant for parameterizing the extent to which CCD influences verb sense disambiguation. The larger α , the stronger CCD's influence on the system output. To avoid data sparseness, we smooth each element (noun example) in $\mathcal{E}_{s_i, c}$. In practice, this involves generalizing each example noun into a 5 digit class based on the *Bunruigoihyo* thesaurus, as has been commonly used for smoothing.

¹³Note that each verb in co-occurrence data should ideally be annotated with its verb sense. However, there is no existing Japanese text base with sufficient volume of word sense tags.

3.4 Experimentation

We estimated the performance of our verb sense disambiguation method through a comparative experiment with other existing methods, in which we compared the following five methods:

- lower bound (LB), in which the system systematically chooses the most frequently appearing verb sense in the database [42],
- rule-based method (RB), in which the system uses a thesaurus to (automatically) identify appropriate semantic classes as selectional restrictions for each verb complement,
- Naive-Bayes method (NB), in which the system interprets a given verb based on the probability that it takes each verb sense,
- similarity-based method using the vector space model (VSM), in which the system uses the above mentioned co-occurrence data extracted from the RWC text base,
- similarity-based method using the *Bunruigoihyo* thesaurus (BGH), in which the system uses Table 3.1 for the similarity computation.

Note that the last two similarity-based methods consider CCD factor for the similarity computation.

In the rule-based method, the selectional restrictions are represented by thesaurus classes, and allow only those nouns dominated by the given class in the thesaurus structure as verb complements. In order to identify appropriate thesaurus classes, we used the association measure proposed by Resnik [120], which computes the information-theoretic association degree between case fillers and thesaurus classes, for each verb sense¹⁴. Equation 3.7 duplicates Equation 2.1 here for the sake of enhanced readability.

$$A(s, c, r) = P(r|s, c) \cdot \log \frac{P(r|s, c)}{P(r|c)} \quad (3.7)$$

Here, $A(s, c, r)$ is the association degree between verb sense s and class r (restriction candidate) with respect to case c . We used the semantic classes defined in the *Bunruigoihyo* thesaurus. $P(r|s, c)$ is the conditional probability that a case filler example associated with case c of sense s is dominated by class r in the *Bunruigoihyo* thesaurus. $P(r|c)$ is the conditional probability that a case filler example for case c (disregarding verb sense) is dominated by class r . Each probability is estimated based on the distribution obtained from the training data. In practice, every r

¹⁴Note that previous research has applied this technique to tasks other than verb sense disambiguation, such as syntactic disambiguation [120] and disambiguation of case filler noun senses [124].

whose association degree is above a certain threshold is chosen as a selectional restriction [120, 124]. Intuitively speaking, by decreasing the value of the threshold, the system coverage can be broadened while opening the way for irrelevant (noisy) selectional restrictions.

The Naive-Bayes method, which is one type of Bayesian classification, assumes that each feature (i.e. case filler) included in a given input is conditionally independent of other case fillers: the system approximates the probability that an input x takes a verb sense s , $P(s|x)$, simply by computing the product of the probability that each verb sense s takes n_c as a case filler for case c . The verb sense with maximal probability is then selected as the interpretation (Equation 3.8)¹⁵.

$$\begin{aligned} \arg \max_s P(s|x) &= \arg \max_s \frac{P(s) \cdot P(x|s)}{P(x)} \\ &= \arg \max_s P(s) \cdot P(x|s) \\ &\simeq \arg \max_s P(s) \prod_c P(n_c|s) \end{aligned} \tag{3.8}$$

Here, $P(n_c|s)$ is the probability that a case filler associated with sense s for case c in the training data is n_c . We estimated $P(s)$ based on the distribution of the verb senses in the training data. In practice, data sparseness leads to not all case fillers n_c appearing in the database, and as such, we generalize each n_c into semantic class defined in the *Bunruigoihyo* thesaurus.

A number of methods involve a parametric constant: the threshold value for the association degree (RB), a generalization level for case filler nouns (NB), and α in Equation 3.6 (VSM and BGH). For these parameters, we conducted several trials prior to the actual comparative experiment, to determine the optimal parameter values over a range of data sets. For our method, we set α extremely large, which is equivalent to virtually relying solely on the SIM of the case with greatest CCD. However, note that when the SIM of the case with greatest CCD is equal for multiple verb senses, the system computes the SIM of the case with second highest CCD. This process is repeated until only one verb sense remains. When more than one verb sense is selected for any given method (or none of them remains, for the rule-based method), the system simply selects the verb sense that appears most frequently in the database¹⁶.

¹⁵A number of experimental results have shown the effectivity of the Naive-Bayes method for word sense disambiguation [43, 78, 96, 103, 112].

¹⁶One may argue that this goes against the basis of the rule-based method, in that, given a proper threshold value for the association degree, the system could improve on the accuracy (potentially sacrificing the coverage), and that the trade-off between the coverage and the accuracy is therefore a more appropriate evaluation criterion. However, our trials on the rule-based method with different threshold values did not show significant correlation between the improvement of the accuracy and the degeneration of the coverage.

In the experiment, we conducted six-fold cross validation, that is, we divided the training/test data into six equal parts¹⁷, and conducted six trials in which a different part was used as test data each time, and the rest as training data (the database). We evaluated the performance of each method according to its accuracy, that is the ratio of the number of correct outputs, compared to the total number of inputs. The training/test data used in the experiment contained about one thousand simple Japanese sentences collected from news articles¹⁸. Each sentence in the training/test data contained one or more complement(s) followed by one of the eleven verbs described in Table 3.2. In Table 3.2, the column of “English gloss” describes typical English translations of the Japanese verbs. The column of “# of sentences” denotes the number of sentences in the corpus, and “# of senses” denotes the number of verb senses contained in IPAL [56]. The column of “accuracy” shows the accuracy of each method.

Table 3.2: The verbs contained in the corpus used, and the accuracy of the different verb sense disambiguation methods (LB: lower bound, RB: rule-based method, NB: Naive-Bayes method, VSM: vector space model, BGH: the *Bunruigoihyo* thesaurus)

verb	English gloss	# of sentences	# of senses	accuracy (%)				
				LB	RB	NB	VSM	BGH
<i>ataeru</i>	give	136	4	66.9	62.1	75.8	84.1	86.0
<i>kakeru</i>	hang	160	29	25.6	24.6	67.6	73.4	76.2
<i>kuwaeru</i>	add	167	5	53.9	65.6	82.2	84.0	86.8
<i>motomeru</i>	require	204	4	85.3	82.4	87.0	85.5	85.5
<i>noru</i>	ride	126	10	45.2	52.8	81.4	80.5	85.3
<i>osameru</i>	govern	108	8	30.6	45.6	66.0	72.0	74.5
<i>tsukuru</i>	make	126	15	25.4	24.9	59.1	56.5	69.9
<i>toru</i>	take	84	29	26.2	16.2	56.1	71.2	75.9
<i>umu</i>	bear offspring	90	2	83.3	94.7	95.5	92.0	99.4
<i>wakaru</i>	understand	60	5	48.3	40.6	71.4	62.5	70.7
<i>yameru</i>	stop	54	2	59.3	89.9	92.3	96.2	96.3
total	—	1315	—	51.4	54.8	76.6	78.6	82.3

Looking at Table 3.2, one can see that our similarity-based method outperformed the other methods (irrespective of the similarity computation), although the Naive-Bayes method is relatively comparable in performance. Surprisingly, despite the relatively ad-hoc similarity definition utilized (see Table 3.1), the *Bunruigoihyo* thesaurus led to a greater accuracy gain than the vector space model. In order to estimate the upper bound (limitation) of the disambiguation task,

¹⁷Ideally speaking, training and test data should be drawn from different sources, to simulate a *real* application. However, the sentences were already scrambled when provided to us, and therefore we could not identify the original source corresponding to each sentence.

¹⁸Morph/syntax analyses were *manually* conducted on the corpus to avoid errors potentially caused by existing tools

that is, to what extent a human expert makes errors in disambiguation [42], we analyzed incorrect outputs and found that roughly 30% of the system errors using the *Bunruigoihyo* thesaurus fell into this category. It should be noted that while the vector space model requires computational cost (time/memory) of an order proportional to the size of the vector, determination of paths in the *Bunruigoihyo* thesaurus comprises a trivial cost¹⁹.

In order to evaluate the effectivity of the CCD factor, we also investigated the accuracy rate for a similarity-based method which does not consider the CCD factor (as performed by Kurohashi and Nagao [77]). In other words, this method computes the score for verb sense s simply by summing the similarity degrees of the input case fillers with respect to each of the example case fillers, as in Equation 3.9. We used the *Bunruigoihyo* thesaurus for the similarity computation between two case filler nouns.

$$Score(s) = \sum_c SIM(n_c, \mathcal{E}_{s,c}) \quad (3.9)$$

We found that the CCD factor led to an accuracy gain from 76.9% to 82.3% (5.4% gain).

We also investigated how the accuracy of each method improved as the training data was increased, because the performance of corpus-based methods generally depends on the size of training data. For this purpose, we initially used only the examples taken from IPAL, and progressively increased the size of the training data used. Figure 3.8 shows the relation between the number of the training data used and the accuracy of different methods. In this figure, zero on the x-axis represents the system using only the IPAL examples, which imitates a dictionary definition-based method²⁰. We derive from this figure that our method using the *Bunruigoihyo* outperformed the other methods, irrespective of the size of training data, and that enhancing the volume of training data significantly improved on the accuracy using only the IPAL examples. The latter observation justifies the necessity of supervision on large-sized training data, as shown in past experiments [96, 104].

One may argue that given sufficient co-occurrence statistics, the vector space model should outperform hand-crafted thesauri, in other words, human lexicographers' knowledge is no longer needed. We investigate this prediction in Table 3.3, which shows the relation between the frequency of nouns appearing in the co-occurrence data extracted from the RWC text base (see Section 3.3) and the accuracy of verb sense disambiguation, in which the "frequency" entry denotes the threshold of the frequency of nouns. The "coverage" entry denotes the ratio between the number of inputs including at least one noun with frequency over a given threshold, and the

¹⁹We will propose a method to optimize the computational cost for the vector space model in Chapter 5.

²⁰See the item "Machine readable dictionary" (p. 25) in Section 2.2.2.




Figure 3.8: The relation between the number of training data used and accuracy of the different methods

total number of inputs. The last two entries show the accuracy with different similarity measures, for each coverage. Surprisingly, not only the accuracy of VSM but also the accuracy of the *Bunruigoihyo* thesaurus increased as the threshold of the frequency increased, and VSM did not outperform the *Bunruigoihyo* thesaurus for any of the thresholds. We could assume frequently appearing nouns are used so commonly that even human lexicographers can reasonably define the similarity between them in a thesaurus. In addition, nouns which frequently appear in the co-occurrence data also appear in the training data, and therefore they provide the maximal similarity (that is, “exact matching”) independent of which similarity measure is used. We would like to note that human knowledge is useful in the task of word sense disambiguation, as with other NLP research [69].

Table 3.3: The frequency of nouns and resultant accuracy of verb sense disambiguation

frequency	≥ 100	≥ 1000	≥ 10000
coverage	72.4%	57.2%	16.3%
VSM	83.5	87.9	91.4
BGH	88.4	92.0	94.4

We also investigated errors made by the rule-based method to find a rational explanation for its inferiority. We identified that the association measure in Equation 3.7 tends to give a greater value to less frequently appearing verb senses and lower level (more specified) classes, and therefore chosen rules are generally overspecified²¹. Consequently, frequently appearing verb senses are likely to be rejected. On the other hand, when attempting to enhance the rule set by setting a smaller threshold value for the association score, overgeneralization can be a problem. We also note that one of the theoretical differences between the rule-based and similarity-based methods is that the former statically generalizes examples (prior to system usage), while the latter does so dynamically. Static generalization would appear to be relatively risky for sparse training data.

Although comparison of different approaches to word sense disambiguation should be further investigated, this experimental result gives us good motivation to explore similarity-based verb sense disambiguation approaches in the following sections and chapters.

²¹This problem has also been identified by Charniak [16].

3.5 Further Enhancement

3.5.1 Computation of interpretation certainty

Since, as shown in Table 3.2, the system still finds it difficult to achieve a 100% accuracy, it is important to select presumably correct outputs from the overall outputs (potentially sacrificing system coverage), for practical applications. Integrated with other NLP systems, it is desirable that other systems can vary the degree of reliance on the output of the verb sense disambiguation system. To achieve this, it is useful to estimate the degree of certainty as to the interpretation, so that one can gain higher accuracy selecting only outputs with greater certainty degree.

A number of methods have been proposed to compute the interpretation certainty in word sense disambiguation [28, 157] and text categorization [80]. These methods estimate interpretation certainty as the ratio between the probability of the most plausible category (word sense/text category), and the probability of any other category, excluding the most probable one. Similarly, in the verb sense disambiguation system, we introduce the notion of interpretation certainty based on the following preference conditions [36]:

1. the highest interpretation score is greater,
2. the difference between the highest and second highest interpretation scores is greater.

The rationale for these conditions is given below. Consider Figure 3.9, where each symbol denotes an example in a given corpus, with symbols x as unsupervised examples and symbols e as supervised examples. The curved lines delimit the semantic vicinities (extents) of the two verb senses 1 and 2, respectively²². The semantic similarity between two examples is graphically portrayed by the physical distance between the two symbols representing them. In Figure 3.9-a, x 's located inside a semantic vicinity are expected to be interpreted as being similar to the appropriate example e with high certainty, a fact which is in line with condition 1 above. However, in Figure 3.9-b, the degree of certainty for the interpretation of any x which is located inside the intersection of the two semantic vicinities cannot be great. This occurs when the case fillers associated with two or more verb senses are not selective enough to allow for a clear cut delineation between them. This situation is explicitly rejected by condition 2.

Based on the above two conditions, we compute interpretation certainties using Equation 3.10, where $C(x)$ is the interpretation certainty of an example x . $Score_1(x)$ and $Score_2(x)$ are the highest and second highest scores for x , respectively. λ , which ranges from 0 to 1, is a parametric constant used to control the degree to which each condition affects the computation

²²Note that this method can easily be extended for a verb which has more than two senses.




Figure 3.9-a: The case where the interpretation certainty of the enclosed x 's is great




Figure 3.9-b: The case where the interpretation certainty of the x 's contained in the intersection of senses 1 and 2 is small

Figure 3.9: The concept of interpretation certainty

of $C(x)$.

$$C(x) = \lambda \cdot \text{Score}_1(x) + (1 - \lambda) \cdot (\text{Score}_1(x) - \text{Score}_2(x)) \quad (3.10)$$

We estimate the validity of the notion of the interpretation certainty, by the trade-off between the accuracy and coverage of the system. Note that in this experiment, the accuracy is the ratio between the number of correct outputs, and the number of cases where the interpretation certainty of the output is above a certain threshold. The coverage is the ratio between the number of cases where the interpretation certainty of the output is above a certain threshold, and the number of inputs. By raising the value of the threshold, the accuracy also increases (at least theoretically), while the coverage decreases.

The system used the *Bunruigoihyo* thesaurus for the similarity computation, and was evaluated by way of six-fold cross validation using the same corpus as that used for the experiment described in Section 3.4. Figure 3.10 shows the result of the experiment with several values of λ , from which the optimal λ value seems to be in the range around 0.5. It can be seen that, as we assumed, both of the above conditions are essential for the estimation of the interpretation certainty.

3.5.2 Incorporation of contextual constraints

A number of researchers have pointed out that words tend to maintain the same sense within a given context [101, 157]. In other words, when the same verb appears multiply in the same context, we can generally assume that it will take the same sense. The crucial issue then becomes which verb sense to select if each verb occurrence is interpreted with different sense by a dedicated method (for example, the methods compared in Section 3.4). Here, we newly




Figure 3.10: The relation between coverage and accuracy with different λ 's

introduce a method to propagate contextual constraint through the degree of interpretation certainty. Let us assume that a given input includes multiple distinct verbs of common lexical content, and that one of them is interpreted with (unreasonably) smaller certainty than the others. In the following, we list cases where the similarity cannot be reasonably computed:

- insufficient supervised training data,
- input case markers are topicalized or substituted for,
- input cases are omitted,
- input case filler nouns are unlisted in the thesaurus²³ (or unlisted in the co-occurrence data in the case of the vector space model),
- input case fillers comprise compound nouns.

In such cases, we superimpose interpretations with higher certainty onto those with lower certainty. It should be noted that while the four problems above can be individually solved through different approaches (some of which are discussed in Section 3.5.3), our approach can be used as a *complementary* solution in a general formalism. To exemplify our contextual propagation method, let us take the following example (taken from the EDR Japanese corpus [58]):

(7) *satou wo tsuka-wazuni ... parachinousu wo tsukau.*
 (sugar-ACC) (do not “use MATERIAL”) ... (palatinose-ACC) (?)

This example contains two occurrences of the polysemous Japanese verb *tsukau*, which has multiple senses in the EDR dictionary [58], such as “to employ”, “to operate”, “to spend” and “to use MATERIAL”. While the sense of the former *tsukau* can be correctly identified as “to use MATERIAL”, the disambiguation of the latter *tsukau* fails because the *Bunruigoihyo* thesaurus does not list the technical term *parachinousu* (“palatinose”), which is a artificial sweetener. However, we can identify the correct interpretation for the latter *tsukau* by superimposing the interpretation for the former one, i.e. “use MATERIAL”. Computation of the degree of certainty is performed using the method proposed in Equation 3.10 (see Section 3.5.1).

For the evaluation of this method, we used the EDR Japanese corpus [58], which was originally extracted from news articles. Note that the corpus used for experiments in Section 3.4 was not applicable to this evaluation because this corpus was compiled on a simple sentence basis. The EDR corpus provides sense information for each word, based on the EDR dictionary, and we

²³Most thesauri, including WordNet and the *Bunruigoihyo* thesaurus, lack proper nouns and technical terms.

used this as a means of checking the interpretation²⁴. From the EDR corpus, we first collected sentences containing one of ten frequently appearing verbs, producing a total of 10,880 sentences. Thereafter, we collected sentences in which a polysemous verb appears more than once, from the aforementioned collection. We limited the range of context to one sentence because the EDR corpus does not provide wider contextual information, such as paragraph boundaries and sentence genres. The number of derived sentences was 462 (out of 10,880), which means the applicability of this method is relatively small. However, the use of contextual information enhanced the accuracy of the similarity-based method using the *Bunruigoihyo* thesaurus from 60.4% to 64.1%.

3.5.3 Remaining problems

Let us discuss how further enhancements to our verb sense disambiguation system could be made.

First, it should be noted that in Japanese, case markers can be omitted or topicalized (i.e. marked with postposition *wa*), an issue which our framework does not currently consider. In addition, polysemous verbs in relative clauses, where the surface case markers of the head noun is omitted, pose a similar problem. Kurohashi and Nagao [77] proposed a way of modeling such (irregular) verb complements, by matching them to complements followed by *ga*, *ni* or *wo* based on the similarity between respective case fillers. Baldwin et al. [3] and Baldwin [2] proposed a head gapping method for Japanese relative clauses, which identifies appropriate case slots for head nouns in relative clauses. Anaphora/ellipsis resolution is expected to overcome the case omission problem. Given that this processing can be carried out successfully, the similarity between an input and examples is expected to be more reliable.

Second, our system is currently limited to the vocabulary defined in the *Bunruigoihyo* thesaurus. Thesaurus extending methods [138, 143] are expected to counter this problem. The problem of vocabulary size is also problematic when case filler nouns comprise compound nouns, and therefore extraction of the semantic head is a crucial task. Note that in Japanese, compound nouns lack lexical segmentation. It has been empirically shown that in more than 90% of cases of four *kanji* character Japanese compound nouns²⁵, the semantic head consists of the last two characters [70]. However, analysis of longer compound nouns still remains as a challenging task. Kobayashi et al. [70, 71], for example, proposed a method of analyzing the syntactic/semantic

²⁴It should be noted that according to our preliminary observation, the EDR corpus contains a number of sense annotation errors and “nil” verb senses (unanalysed/unanalyseable verb senses). In view of this problem, we did not use this corpus for other experiments in this research.

²⁵*Kanji* characters are basic ideograms in Japanese.

structure of Japanese compound nouns, which can further enhance our system coverage.

Third, some idiomatic expressions represent highly restricted collocations, and overgeneralizing them semantically through the use of a thesaurus can cause further errors. Possible solutions would include one proposed by Uramoto, in which idiomatic expressions are described separately in the database so that the system can control their overgeneralization [142] (as discussed in Section 2.4).

Finally, external information such as the discourse or domain dependency of each word sense [47, 101, 157] is expected to lead to system improvement²⁶.

3.6 Summary

This chapter described our similarity-based verb sense disambiguation system [37, 38]. The basis of the system is as follows. First, provided with sentences containing a polysemous verb, the system searches the database for the most similar example to the input, following the nearest neighbor method. The database consists of supervised examples which are manually annotated with correct verb senses. Thereafter, the verb is disambiguated by superimposing the sense of the verb appearing in the most similar supervised example. The similarity (score) between the input and an example, or more precisely, the weighted average of the similarity between case filler nouns included in them, is computed based on either a statistical measure or an existing thesaurus. For the statistical similarity measure, we use the ubiquitous vector space model [33, 78, 106, 127, 130]. In this, each case filler noun is represented as a vector comprising statistical factors about its collocation, which are taken from a large-scaled text base, and the similarity between two nouns is computed as the cosine of the angle between their associated vectors. As for the thesaurus-driven similarity measure [77, 81, 142], we applied a method proposed by Kurohashi and Nagao [77], which determines the similarity between two nouns as the length of the path between them in the hand-crafted *Bunruigoihyo* thesaurus [102]. In practice, morphological and syntactic analyses are needed prior to the verb sense disambiguation process. We currently assume the use of existing tools for morphological and syntactic analysis on the input, and do not focus on these modules in this research.

Let us summarize the main points that have been made in this chapter.

First, we introduced the weighting factor of case contribution to disambiguation (CCD), which computes the degree to which each case filler contributes to verb sense disambiguation. Intuitively speaking, greater diversity of semantic range of case filler examples will lead to

²⁶Note that we limited the range of context to one sentence in Section 3.5.2, because the EDR corpus does not provide wider contextual information, such as paragraph boundaries and sentence genres.

greater contribution of that case to verb sense disambiguation. We then compute the similarity between the input and each supervised example, placing emphasis on cases with greater CCD values. We also identified that the introduction of the CCD factor can be seen as a feature selection method, that is, cases with greater CCD values consist of a useful feature set. In order to investigate the effectivity of CCD, we compared our system with the rule-based and Naive-Bayes methods. Empirical results showed that the similarity-based system combined with the CCD factor improved on the lower bound performance to a larger degree than the rule-based and Naive-Bayes methods. We also showed that the use of the *Bunruigoihyo* thesaurus is comparable with the vector space model in the similarity computation. While criticism has been made of inherent limitations in hand-crafted resources, our experimental results advocate the effectivity of these resources for NLP applications. In addition, we showed that additional supervised training examples significantly improved on the performance, relying solely on a small number of examples taken from the machine readable dictionary IPAL [56].

Second, in order to achieve higher accuracy, we selected presumably correct outputs from the overall outputs by use of the notion of interpretation certainty. The interpretation certainty is greater when (a) the score associated with the selected verb sense is greater, and (b) the difference between the highest and second highest scores is greater. We showed the effectivity of this computation by way of experiments.

Finally, our prototype method of propagating contextual constraints, further improved on the performance of the similarity-based method.

Chapter 4

Selective Sampling of Effective Example Sets

This chapter proposes an efficient example sampling method for similarity-based word sense disambiguation systems. To construct a database of a practical size, a considerable overhead for manual sense disambiguation (“overhead for supervision”) is required. In addition, the time complexity of searching a large-sized database poses a considerable problem (“overhead for search”). To counter these problems, our method selectively samples a smaller-sized effective subset from a given example set for use in word sense disambiguation. Our method is characterized by its reliance on the notion of ‘training utility’: the degree to which each example is informative for future example sampling when used for the training of the system. The system progressively collects examples by selecting those with greatest utility. This paper reports on the effectivity of our method through experiments on about one thousand sentences. Compared to experiments with other example sampling methods, our method reduced both the overhead for supervision and the overhead for search, without degenerating the performance of the system.

4.1 Motivation

In Chapter 3, we described the similarity-based verb sense disambiguation system, and showed its effectivity through comparative experiments. Following the nearest neighbor method, our system uses the database, which contains example sentences associated with each verb sense. Given an input sentence containing a polysemous verb, the system chooses the most plausible verb sense from predefined candidates. In this process, the system computes a scored similarity

between the input and examples in the database, and chooses the verb sense associated with the example which maximizes the score. To realize this, we have to manually disambiguate polysemous verbs appearing in examples, prior to their use by the system. We shall call these examples “supervised examples”. A preliminary experiment conducted in Section 3.4 on eleven polysemous Japanese verbs showed that (a) the more supervised examples we provided to the system, the better it performed, and (b) in order to achieve a reasonable result (say over 80% accuracy), the system needed a hundred-order supervised example set for each verb. Therefore, in order to build an operational system, the following problems have to be taken into account¹:

- given human resource limitations, it is not reasonable to supervise every example in large corpora (“overhead for supervision”),
- given the fact that similarity-based systems, including our system, search the database for the examples most similar to the input, the computational cost becomes prohibitive if one works with a very large database size (“overhead for search”).

Empirically speaking, we observed that about 84% of the 900 or so verbs defined in IPAL [56] are polysemous, and that some of these verbs are associated with debilitatingly large numbers of senses (the maximum number of senses is 32). This statistics shows that the overhead for supervision in a real application is never trivial. It should be noted that the overhead for supervision is also crucial when one tries to customize a WSD system to several distinct domains, because (a) word sense distinctions are often different depending on domains, and therefore (b) multiple overhead for supervision in each domain is required.

One may argue that unsupervised WSD methods (for example, those described in Section 2.2.2) can overcome the former problem. To investigate this assumption, we preliminary applied the bootstrapping method to the corpus used in Section 3.4. In this experiment, we first provided the similarity-based system (using the *Bunruigoihyo* thesaurus) with an initial database, consisting of examples taken from IPAL [56] (see the “example nouns” entry listed in Figure 3.3).

Thereafter, the system repeats the training process and automatically incorporates the example with maximal interpretation certainty into the database, until no training data remains. We used Equation 3.10 for the computation of interpretation certainty (see Section 3.5.1). Finally, we evaluated the accuracy of verb sense disambiguation on open test data. For this purpose, we divided the corpus into test and training data and conducted six-fold cross validation, as carried

¹Note that these problems are associated with corpus-based approaches in general, and have been identified by a number of researchers [32, 80, 141, 157].

out in Section 3.4. We found that the accuracy was 66.1%, which is less than satisfactory when compared with the accuracy of the supervised methods listed in Table 3.2. In fact, the bootstrapping method did not significantly improve on the accuracy rate achieved only using the initial database (64.7%). This experimental result advocates that we still need supervised methods in certain cases, although we concede that unsupervised methods also need to be explored to counter the overhead for supervision.

Motivated by these above arguments, we propose a different approach, namely to *select* a small number of optimally informative examples from given corpora [35, 36]. Hereafter we will call these examples “samples”. Our example sampling method, based on the utility maximization principle, decides on the preference for including a given example in the database. This decision procedure is usually called “selective sampling” [24]. The overall control flow of selective sampling systems can be depicted as in Figure 4.1, where “system” refers to our verb sense disambiguation system, and “examples” refers to an unsupervised example set². The sampling process basically cycles between the word sense disambiguation (WSD) and training phases. During the WSD phase, the system generates an interpretation for each polysemous verb contained in the input example (“WSD outputs”). This phase is equivalent to normal word sense disambiguation execution. During the training phase, the system selects samples for training from the previously produced outputs. During this phase, a human expert supervises samples, that is, provides the correct interpretation for the verbs appearing in the samples. Thereafter, samples are simply incorporated into the database without any computational overhead (as would be associated with globally re-estimating parameters in statistics-based systems), meaning that the system can be trained on the remaining examples (“residue”) for the next iteration. Iterating between these two phases, the system progressively enhances the database. It should be noted that the selective sampling procedure gives us an optimally informative database of a given size irrespective of the stage at which processing is terminated.

Several researchers have proposed this type of approach for NLP applications. Engelson and Dagan [32] proposed a committee-based sampling method, which is currently applied to HMM training for part-of-speech tagging. This method sets several models (committee) taken from a given supervised data set, and selects samples based on the degree of disagreement among the committee members as to the output. However, as this method is implemented for statistics-based models, there is a need to explore how to formalize and map the concept of selective sampling into similarity-based approaches. Lewis and Gale [80] proposed an uncertainty

²The “system” in Figure 4.1 includes the “VSD”, “thesaurus/co-occurrence” and “morph/syntax analyzer” from Figure 3.1. However, note that the morphological and syntactic analyses for “examples” are needed only at the first stage of the iteration.

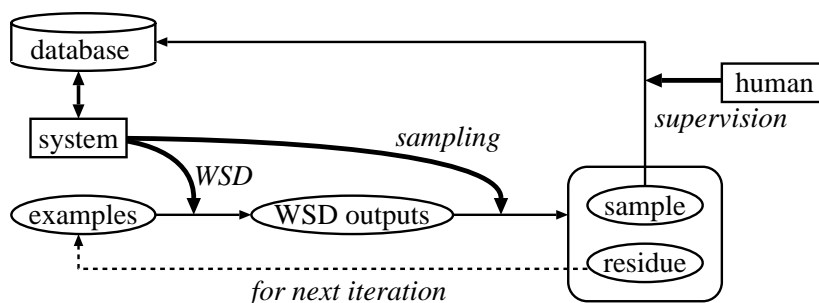


Figure 4.1: Flow of control of the selective sampling

sampling method for statistics-based text classification. In this method, the system always samples outputs with an uncertain level of correctness. However, we should take into account the training effect a given example has on other unsupervised examples, introduced as ‘training utility’ in our method. We devote Section 4.3 to further comparison of our approach and other related works.

With respect to the problem of overhead for search, possible solutions would include the generalization of similar examples through a thesaurus taxonomy [61, 107] or the reconstruction of the database using a small portion of useful instances selected from a given supervised example set [1, 136]. Aha et al. [1] select examples that significantly contribute to the performance gain on test data. However, such approaches imply a significant overhead for supervision of each example prior to the system’s execution. This shortcoming is precisely what our approach aims to avoid: reduction of the overhead for supervision as well as the overhead for search.

In the followings, Section 4.2 first elaborates on the methodology of our selective sampling method. Section 4.3 then evaluates our method by way of experiments. Section 4.4 describes related work. Finally, discussion is added in Section 4.5.

4.2 Methodology

4.2.1 Overview

Let us look again at Figure 4.1, in which “WSD outputs” refers to a corpus in which each sentence is assigned an expected verb interpretation during the WSD phase. In the training phase, the system stores supervised samples (with each interpretation simply checked or appropriately corrected by a human) in the database, to be used in a later WSD phase. In this section, we turn to the problem as to which examples should be selected as samples.

Lewis and Gale [80] proposed the notion of uncertainty sampling for the training of statistics-

based text classifiers. Their method selects those examples that the system classifies with minimum certainty, based on the assumption that there is no need for teaching the system the correct answer when it has answered with sufficiently high certainty. However, we should take into account the training effect a given example has on other remaining (unsupervised) examples. In other words, we would like to select samples such as to be able to correctly disambiguate as many examples as possible in the next iteration. If this is successfully done, the number of examples to be supervised will decrease. We consider maximization of this effect by means of a training utility function aimed at ensuring that the example with the greatest training utility factor, is the most useful example at a given point in time. Intuitively speaking, the training utility of an example is greater when we can expect greater increase in the interpretation certainty of the remaining examples after training using that example (see Section 3.5.1 for the computation of interpretation certainty).

To explain this notion intuitively, let us take Figure 4.2 as an example corpus. In this corpus, all sentences contain the verb *yameru*, which has two senses according to IPAL, s_1 (“to stop (something)”) and s_2 (“to quit (occupation)”). In this figure, sentences e_1 and e_2 are supervised examples associated with the senses s_1 and s_2 , respectively, and x_i ’s are unsupervised examples. For the sake of enhanced readability, the examples x_i ’s are partitioned according to their verb senses, that is, x_1 to x_5 correspond to sense s_1 , and x_6 to x_9 correspond to sense s_2 . In addition, one may notice that examples in the corpus can be readily categorized based on case similarity, that is, into clusters $\{x_1, x_2, x_3, x_4\}$ (“someone/something stops service”), $\{e_2, x_6, x_7\}$ (“someone leaves organization”), $\{x_8, x_9\}$ (“someone quits occupation”), $\{e_1\}$ and $\{x_5\}$. Let us simulate the sampling procedure with this example corpus. In the initial stage with $\{e_1, e_2\}$ in the database, x_6 and x_7 can be interpreted as s_2 with greater certainty than for other x_i ’s, because these two examples are similar to e_2 . Therefore, uncertainty sampling selects any example excepting x_6 and x_7 as the sample. However, it is evident that any of x_1 to x_4 is more desirable because by way of incorporating one of these examples, we can obtain more x_i ’s with greater certainty. Assuming that x_1 is selected as the sample and incorporated into the database with sense s_1 , either of x_8 and x_9 will be more highly desirable than other unsupervised x_i ’s in the next stage.

Let \mathbf{S} be a set of sentences, i.e. a given corpus, and \mathbf{D} be the subset of supervised examples stored in the database. Further, let \mathbf{X} be the set of unsupervised examples, realizing Equation 4.1.

$$\mathbf{S} = \mathbf{D} \cup \mathbf{X} \quad (4.1)$$

The example sampling procedure can be illustrated as:

1. $WSD(\mathbf{D}, \mathbf{X})$

e_1 :	<i>seito ga</i> (student-NOM)	<i>shitsumon wo</i> (question-ACC)	<i>yameru</i> (s_1)
e_2 :	<i>ani ga</i> (brother-NOM)	<i>kaisha wo</i> (company-ACC)	<i>yameru</i> (s_2)
x_1 :	<i>shain ga</i> (employee-NOM)	<i>eigyō wo</i> (sales-ACC)	<i>yameru</i> (?)
x_2 :	<i>shoten ga</i> (store-NOM)	<i>eigyō wo</i> (sales-ACC)	<i>yameru</i> (?)
x_3 :	<i>kōjō ga</i> (factory-NOM)	<i>sōgyō wo</i> (operation-ACC)	<i>yameru</i> (?)
x_4 :	<i>shisetsu ga</i> (facility-NOM)	<i>uten wo</i> (operation-ACC)	<i>yameru</i> (?)
x_5 :	<i>sensyū ga</i> (athlete-NOM)	<i>renshū wo</i> (practice-ACC)	<i>yameru</i> (?)
x_6 :	<i>musuko ga</i> (son-NOM)	<i>kaisha wo</i> (company-ACC)	<i>yameru</i> (?)
x_7 :	<i>kangōfu ga</i> (nurse-NOM)	<i>byōin wo</i> (hospital-ACC)	<i>yameru</i> (?)
x_8 :	<i>hikoku ga</i> (defendant-NOM)	<i>giin wo</i> (congressman-ACC)	<i>yameru</i> (?)
x_9 :	<i>chichi ga</i> (father-NOM)	<i>kyōshi wo</i> (teacher-ACC)	<i>yameru</i> (?)

Figure 4.2: Example of a given corpus associated with the verb *yameru*

2. $e \leftarrow \arg \max_{x \in \mathbf{X}} TU(x)$
3. $\mathbf{D} \leftarrow \mathbf{D} \cup \{e\}$, $\mathbf{X} \leftarrow \mathbf{X} \cap \overline{\{e\}}$
4. goto 1

where $WSD(\mathbf{D}, \mathbf{X})$ is the verb sense disambiguation process on input \mathbf{X} using \mathbf{D} as the database. In this disambiguation process, the system outputs the following for each input: (a) a set of verb sense candidates with interpretation scores, and (b) an interpretation certainty. These factors are used for the computation of $TU(x)$, newly introduced in our method. $TU(x)$ computes the training utility factor for an example x . The sampling algorithm gives preference to examples of maximum utility.

We will explain in the following sections how one can estimate $TU(x)$, based on the estimation of the interpretation certainty.

4.2.2 Computation of training utility

The training utility of an example a is greater than that of another example b when the total interpretation certainty of unsupervised examples increases more after training with example a than with example b . Let us consider Figure 4.3, in which the x-axis mono-dimensionally denotes the semantic similarity between two unsupervised examples, and the y-axis denotes the interpretation certainty of each example. Let us compare the training utility of the examples a and b in Figure 4.3-a. Note that in this figure, whichever example we use for training, the interpretation certainty for each unsupervised example (x) neighboring the chosen example increases based on its similarity to the supervised example. Since the increase in the interpretation certainty of a given x becomes smaller as the similarity to a or b diminishes, the training utility of the two examples can be represented by the shaded areas. It is obvious that the training

utility of a is greater as it has more neighbors than b . On the other hand, in Figure 4.3-b, b has more neighbors than a . However, since b is semantically similar to e , which is already contained in the database, the total *increase* in interpretation certainty of its neighbors, i.e. the training utility of b , is smaller than that of a .




Figure 4.3-a: The case where the training utility of a is greater than that of b because a has more unsupervised neighbors




Figure 4.3-b: The case where the training utility of a is greater than that of b because b closely neighbors e , contained in the database

Figure 4.3: The concept of training utility

Let $\Delta C(x=s, y)$ be the difference in the interpretation certainty of $y \in \mathbf{X}$ after training with $x \in \mathbf{X}$, taken with the sense s . We compute the interpretation certainty of an example using Equation 3.10 in Section 3.5.1. $TU(x=s)$, which is the training utility function for x taken with sense s , can be computed by way of Equation 4.2.

$$TU(x=s) = \sum_{y \in \mathbf{X}} \Delta C(x=s, y) \quad (4.2)$$

It should be noted that in Equation 4.2, we can replace \mathbf{X} with a subset of \mathbf{X} which consists of neighbors of x . However, in order to facilitate this, an efficient algorithm to search for neighbors of an example is required. We will discuss this problem in Section 4.5.

Since there is no guarantee that x will be supervised with any given sense s , it can be risky to rely solely on $TU(x=s)$ for the computation of $TU(x)$. We estimate $TU(x)$ by the *expected* value of x , calculating the average of each $TU(x=s)$, weighted by the probability that x takes sense s . This can be realized by Equation 4.3, where $P(s|x)$ is the probability that x takes the sense s .

$$TU(x) = \sum_s P(s|x) \cdot TU(x=s) \quad (4.3)$$

Given the fact that (a) $P(s|x)$ is difficult to estimate in the current formulation, and (b) the cost of computation for each $TU(x=s)$ is not trivial, we temporarily approximate $TU(x)$ as in Equation 4.4, where \mathbf{K} is a set of the k -best verb sense(s) of x with respect to the interpretation

score in the current state.

$$TU(x) \simeq \frac{1}{k} \sum_{s \in \mathbf{K}} TU(x=s) \quad (4.4)$$

4.2.3 Enhancement of computation

In this section, we discuss how to enhance the computation associated with our example sampling algorithm.

First, we note that computation of $TU(x=s)$ in Equation 4.2 (see Section 4.2.2) becomes time consuming because the system is required to search the whole set of unsupervised examples for examples whose interpretation certainty will increase after x is used for training. To avoid this problem, we could potentially apply a method used in efficient database search techniques, by which the system can search for neighbor examples of x with optimal time complexity [145]. However, in this section, we will explain another efficient algorithm to identify neighbors of x , in which neighbors of case fillers are taken as being given directly by the thesaurus structure³. The basic idea is the following: the system searches for neighbors of each case filler of x instead of x as a whole, and merges them as a set of neighbors of x . Note that by dividing examples along the lines of each case filler, we can retrieve neighbors based on the structure of the *Bunruigoihyo* thesaurus (instead of the conceptual semantic space as in Figure 3.9). Let $\mathbf{N}_{x=s,c}$ be a subset of unsupervised neighbors of x whose interpretation certainty will increase after x is used for training, considering only case c of sense s . The *actual* neighbor set of x with sense s ($\mathbf{N}_{x=s}$) is then defined as in Equation 4.5.

$$\mathbf{N}_{x=s} = \bigcup_c \mathbf{N}_{x=s,c} \quad (4.5)$$

Figure 4.4 shows a fragment of the thesaurus, in which x and the y 's are unsupervised case filler examples. Symbols e_1 and e_2 are case filler examples stored in the database taken as senses s_1 and s_2 , respectively. The triangles represent subtrees of the structure, and the labels n_i represent nodes. In this figure, it can easily be seen that the interpretation score of s_1 never changes for examples other than the children of n_4 , after x is used for training with sense s_1 . In addition, incorporating x into the database with sense s_1 never changes the score of examples y for other sense candidates. Therefore, $\mathbf{N}_{x=s_1,c}$ includes only examples dominated by n_4 , in other words, examples which are more closely located to x than e_1 in the thesaurus structure. Since, during the WSD phase, the system determines e_1 as the supervised neighbor of x for sense s_1 , identifying $\mathbf{N}_{x=s_1,c}$ does not require any extra computational overhead. One may notice that

³Utsuro's method requires the construction of large-scaled similarity templates prior to similarity computation [145], and this is what we would like to avoid.

the technique presented here is not applicable when the vector space model (see Section 3.3) is used for the similarity computation. However, automatic clustering algorithms, which give a set of words hierarchy based on the similarity between them (for example, one proposed by Tokunaga et al. [139]), could potentially facilitate the application of this retrieval method to the vector space model.




Figure 4.4: A fragment of the thesaurus including neighbors of x associated with case c

Second, sample size at each iteration should ideally be one, so as to avoid the supervision of similar examples. On the other hand, a small sampling size generates a considerable computation overhead for each iteration of the sampling procedure. This can be a critical problem for statistics-based approaches, as the reconstruction of statistic classifiers is expensive. However, similarity-based systems fortunately do not require reconstruction, and examples simply have to be stored in the database. Furthermore, in each disambiguation phase, our similarity-based system needs only compute the similarity between each newly stored example and its unsupervised neighbors, rather than between every example in the database and every unsupervised example. Let us reconsider Figure 4.4. As mentioned above, when x is stored in the database with sense s_1 , only the interpretation score of y 's dominated by n_4 , i.e. $\mathbf{N}_{x=s_1,c}$, will be changed with respect to sense s_1 . This algorithm reduces the time complexity of each iteration from $O(N^2)$ to $O(N)$, given that N is the total number of examples in a given corpus.

4.3 Experimentation

In order to investigate the effectivity of our example sampling method, we conducted an experiment, in which we compared the following four sampling methods:

- a control (random), in which a certain proportion of a given corpus is randomly selected for training,
- uncertainty sampling (US), in which examples with minimum interpretation certainty are selected [80],

- committee-based sampling (CBS) [32],
- our method based on the notion of training utility (TU).

We elaborate on uncertainty sampling and committee-based sampling in Section 4.4. We compared these sampling methods by evaluating the relation between the number of training examples sampled and the performance of the system. We conducted six-fold cross validation and carried out sampling on the training set. With regard to the training/test data set, we used the same corpus as that used for the experiment described in Section 3.4. Each sampling method uses examples from IPAL to initialize the system, with the number of example case fillers for each case being an average of about 3.7. For each sampling method, the system uses the *Bunruigoihyo* thesaurus for the similarity computation. In Table 3.2 (in Section 3.4), the column of “accuracy” for “BGH” denotes the accuracy of the system with the entire set of training data contained in the database. Each of the four sampling methods achieved this figure at the conclusion of training.

We evaluated each system performance according to its accuracy, that is the ratio of the number of correct outputs, compared to the total number of inputs. For the purpose of this experiment, we set the sample size to 1 for each iteration, $\lambda = 0.5$ for Equation 3.10, and $k = 1$ for Equation 4.4. Based on a preliminary experiment, increasing the value of k either did not improve the performance over that for $k = 1$, or lowered the overall performance. Figure 4.5 shows the relation between the number of the training data sampled and the accuracy of the system. In Figure 4.5, zero on the x-axis represents the system using only the examples provided by IPAL. Looking at Figure 4.5 one can see that compared with random sampling and committee-based sampling, our sampling method reduced the number of the training data required to achieve any given accuracy. For example, to achieve a accuracy of 80%, the number of the training data required for our method was roughly one-third of that for random sampling. Although the accuracy for our method was surpassed by that for uncertainty sampling for larger sizes of training data, this minimal difference for larger data sizes is overshadowed by the considerable performance gain attained by our method for smaller data sizes.

Since IPAL has, in a sense, been manually selectively sampled in an attempt to model the maximum verb sense coverage, the performance of each method is biased by the initial contents of the database. To counter this effect, we also conducted an experiment involving the construction of the database from scratch, without using examples from IPAL. During the initial phase, the system randomly selected one example for each verb sense from the training

set, and a human expert provided the correct interpretation to initialize the system⁴. Figure 4.6 shows the performance of the various methods, from which the same general tendency as seen in Figure 4.5 is observable. However, in this case, our method was generally superior to other methods. Through these comparative experiments, we can conclude that our example sampling method is able to decrease the number of the training data, i.e. the overhead for both supervision and searching, without degrading the system performance.




Figure 4.5: The relation between the number of training data sampled and accuracy of the system

4.4 Related Work

4.4.1 Uncertainty sampling

The procedure for uncertainty sampling [80] is as follows, where $C(x)$ represents the interpretation certainty for an example x (see our sampling procedure in Section 4.2.1 for comparison):

⁴In order to minimize the potential bias caused by selection of the initial examples, we conducted the same trials with three different initial example sets and averaged the results for each trial.




Figure 4.6: The relation between the number of training data sampled and accuracy of the system without using examples from IPAL

1. $WSD(\mathbf{D}, \mathbf{X})$
2. $e \leftarrow \arg \min_{x \in \mathbf{X}} C(x)$
3. $\mathbf{D} \leftarrow \mathbf{D} \cup \{e\}$, $\mathbf{X} \leftarrow \mathbf{X} \cap \overline{\{e\}}$
4. goto 1

Let us discuss the theoretical difference between this and our method. Considering Figure 4.3 again, one can see that the concept of training utility is supported by the following properties:

- (a) an example which neighbors more unsupervised examples is more informative (Figure 4.3-a),
- (b) an example less similar to one already existing in the database is more informative (Figure 4.3-b).

Uncertainty sampling directly addresses property (b), but ignores property (a). It differs from our method more crucially when more unsupervised examples remain, because these unsupervised examples have a greater influence on the computation of training utility. This assumption can also be seen in the comparative experiments in Section 4.3, in which our method outperformed uncertainty sampling to the highest degree in early stages.

4.4.2 Committee-based sampling

In committee-based sampling [32], which follows the “query by committee” principle [132], the system selects samples based on the degree of disagreement between models randomly taken from a given training set (these models are called “committee members”). This is achieved by iteratively repeating the steps given below, in which the number of committee members is given as two without loss of generality:

1. draw two models randomly,
2. classify unsupervised example x according to each model, producing classifications C_1 and C_2 ,
3. if $C_1 \neq C_2$ (the committee members disagree), select x for the training of the system.

Figure 4.7 shows a typical disparity evident between committee-based sampling and our sampling method. The basic notation in this figure is the same as in Figure 3.9, and both x and y denote unsupervised examples, or more formally $\mathbf{D} = \{e_1, e_2\}$, and $\mathbf{X} = \{x, y\}$. Assume a pair

of committee members $\{e_1\}$ and $\{e_2\}$ have been selected from the database \mathbf{D} . In this case, the committee members disagree as to the interpretations of both x and y , and consequently, both examples can potentially be selected as a sample for the next iteration. In fact, committee-based sampling tends to require a number of similar examples (similar to e_1 and y) in the database, otherwise committee members taken from the database will never agree. This feature provides a salient contrast to our method for which a similar example is less informative, and x is therefore preferred to y as a sample. This contrast can also correlate to the fact that committee-based sampling is currently applied to statistics-based language models (HMM classifiers), in other words, statistical models generally require that the distribution of the training data reflects that of the overall text. Through this argument, one can assume that committee-based sampling is better suited to statistics-based systems, while our method is more suitable for similarity-based systems.




Figure 4.7: A case where either x or y can be selected in committee-based sampling

Engelson and Dagan [32] criticized uncertainty sampling [80], which they call a “single model” approach, as distinct from their “multiple model” approach:

... sufficient statistics may yield an accurate 0.51 probability estimate for a class c in a given example, making it certain that c is the appropriate classification⁵. However, the certainty that c is the correct classification is low, since there is a 0.49 chance that c is the wrong class for the example. A single model can be used to estimate only the second type of uncertainty, which does not correlate directly with the utility of additional training.

We note that this criticism cannot be applied to our sampling method, despite the fact that our method falls into the category of a single model approach. In our sampling method, given sufficient statistics, the increment of the certainty degree for unsupervised examples, i.e. the training

⁵By appropriate classification, Engelson and Dagan mean the classification given by a perfectly-trained model.

utility of additional supervised examples, becomes small (theoretically, for both similarity-based and statistics-based systems). As such, the utility factor can be considered to correlate directly with additional training, for our method.

4.5 Discussion

4.5.1 Sense ambiguity of case fillers in selective sampling

The semantic ambiguity of case fillers (nouns) should be taken into account during selective sampling. Figure 4.8, which uses the same basic notation as Figure 3.9, illustrates one possible problem caused by case filler ambiguity. Let x_1 be a sense of a case filler x , and y_1 and y_2 be different senses of a case filler y . On the basis of Equation 3.10, the interpretation certainty of x and y is small in Figures 4.8-a and 4.8-b, respectively. However, in the situation shown in Figure 4.8-b, since (a) the task of distinguishing between the *verb* senses 1 and 2 is easier, and (b) instances where the sense ambiguity of case fillers corresponds to distinct verb senses will be rare, training using either y_1 or y_2 will be less effective than using a case filler of the type of x . It should also be noted that since *Bunruiyohyo* fundamentally associates each word with a single concept, this problem is not critical in our case. However, given other existing thesauri like the EDR electronic dictionary [58] or WordNet [93], these two situations should be strictly differentiated.




Figure 4.8-a: Interpretation certainty of x is small because x lies in the intersection of distinct verb senses




Figure 4.8-b: Interpretation certainty of y is small because y is semantically ambiguous

Figure 4.8: Two separate scenarios in which the interpretation certainty of x is small

4.5.2 Applicability of our selective sampling method

First, we note that our selective sampling method is expected to be applicable to most similarity-based methods, or more precisely, those which follow the nearest neighbor method, although

in this chapter, we describe the sampling method being applied only to our verb sense disambiguation method (from Chapter 3). However, our sampling method does not seem effective for some similarity-based methods, such as the one proposed by Ng and Lee [105]. Ng and Lee used the similarity measure as shown in Equation 2.15 (see p. 20). It should be noted that their similarity computation uses the statistics obtained from supervised examples. *Selective* sampling for the supervised examples potentially biases the estimation of $P(s|F=f)$, i.e. the conditional probability that sense s occurs, given that feature F takes value f . This issue needs to be further explored. On the other hand, our sampling method can be applied to similarity-based methods which use resources *independent* of the supervised examples, for the similarity computation [19, 77, 106, 142].

Second, Figure 4.9 exemplifies a limitation of our sampling method. The basic notation is the same as in Figure 3.9. In this figure, the only supervised examples contained in the database are e_1 and e_2 , and x represents an unsupervised example belonging to sense 2. Given this scenario, x is informative because (a) it clearly evidences the semantic vicinity of sense 2, and (b) without x as sense 2 in the database, the system may misinterpret other examples neighboring x . However, in our current implementation, the training utility of x would be small because it would be mistakenly interpreted as sense 1 with great certainty due to its relatively close semantic proximity to e_1 . Even if x has a number of unsupervised neighbors, the total increment of their interpretation certainty cannot be expected to be large. This shortcoming often presents itself when the semantic vicinities of different verb senses are closely aligned or their semantic ranges are not disjunctive. Here, let us consider Figure 3.5 again, in which the nominative case would parallel the semantic space shown in Figure 4.9 more closely than the accusative. Relying more on the similarity in the accusative (the case with greater CCD) as is done in our system, we aim to map the semantic space in such a way as to achieve higher semantic disparity and minimize this shortcoming.




Figure 4.9: The case where informative example x is not selected

4.6 Summary

Corpus-based approaches have recently pointed the way to a promising trend in word sense disambiguation. However, these approaches tend to require a considerable overhead for supervision in constructing a large-sized database, additionally resulting in a computational overhead to search the database. To overcome these problems, our method, which is currently applied to an similarity-based verb sense disambiguation system, selectively samples a smaller-sized subset from a given example set [35, 36]. This method is expected to be applicable to other similarity-based systems. Applicability for other types of systems needs to be further explored.

The process basically iterates through two phases: (normal) word sense disambiguation and a training phase. During the disambiguation phase, the system is provided with sentences containing a polysemous verb, and searches the database for the most semantically similar example to the input (nearest neighbor method). Thereafter, the verb is disambiguated by superimposing the sense of the verb appearing in the supervised example. The similarity between the input and an example, or more precisely the similarity between the case fillers included in them, is computed based on an existing thesaurus (or the vector space model is alternatively applicable). In the training phase, a sample is then selected from the system outputs and provided with the correct interpretation by a human expert. Through these two phases, the system iteratively accumulates supervised examples into the database. The critical issue in this process is to decide which example should be selected as a sample in each iteration. To resolve this problem, we considered the following properties: (a) an example which neighbors more unsupervised examples is more influential for subsequent training, and therefore more informative, and (b) since our verb sense disambiguation is based on nearest neighbor resolution, an example similar to one already existing in the database is redundant. Motivated by these properties, we introduced and formalized the concept of training utility as the criterion for example selection. Our sampling method always gives preference to that example which maximizes training utility.

We reported on the performance of our sampling method by way of experiments, in which we compared our method with random sampling, uncertainty sampling [80] and committee-based sampling [32]. The result of the experiments showed that our method reduced both the overhead for supervision and the overhead for searching the database to a larger degree than any of the above three methods, without degrading the performance of verb sense disambiguation. Through the experiment and discussion, we claim that uncertainty sampling considers property (b) mentioned above, but lacks property (a). We also claim that committee-based sampling differs from our sampling method in terms of its suitability to statistics-based systems as compared to similarity-based systems.

Chapter 5

Exploration of Similarity Computation

This chapter explores the computation of word similarity, on which the performance of our similarity-based verb sense disambiguation system is highly dependent. The statistics-based computation of word similarity has been popular in recent research, but is associated with significant computational cost. On the other hand, the use of hand-crafted thesauri as semantic resources is simple to implement, but lacks mathematical rigor. To integrate the advantages of these two approaches, we aim at calculating a statistical weight for each branch of a thesaurus, so that we can compute word similarity based simply on the length of the path between two words in the thesaurus. Our experiment on Japanese nouns shows that this framework upheld the inequality of statistics-based word similarity with an accuracy of more than 70%. We also applied our framework of word similarity computation to our verb sense disambiguation system.

5.1 Motivation

As with other similarity-based systems, the performance of our system (as described in Chapter 3) is highly dependent on the computation of the (relative) similarity between two examples, which is further decomposed into the similarity between two case filler nouns (we shall call this “word similarity”, hereafter). By the performance of the system, we refer to the following different viewpoints:

- the relative number of correct outputs (“accuracy”),

- the cost needed for the similarity computation (“efficiency”).

We described two previous approaches for word similarity computation in Section 3.3. Here, let us summarize these approaches again.

The first approach for word similarity computation is statistics-based methods [11, 16, 29, 46, 53, 113, 139]. From a number of different implementations, we used the vector space model (VSM) [33, 78, 106, 127, 130], for our verb sense disambiguation system. In the vector space model, each word is represented by a vector consisting of co-occurrence statistics, such as relative frequency, with respect to other words. Note that in our case, we used only the co-occurrence between a noun and verbs. The similarity between two given words is then measured computationally using the two vectors representing those words. One typical implementation computes the similarity as the cosine of the angle between the two vectors, a method which is also commonly used in information retrieval and text categorization systems to measure the similarity between documents [127]. Since it is based on mathematical rationale, this type of similarity measurement has been popular. Besides this, since the similarity is computed based on given co-occurrence data, word similarity can easily be adjusted according to the domain. However, data sparseness is an inherent problem. This fact was observed in our preliminary experiment (see Section 3.4), despite using statistical information taken from news articles for as many as *four* years. Furthermore, in this approach, each vector requires $O(N)$ computational cost, given that N is the length of the vector. Note that although, in our case, each vector term is confined to a verb, additional parts-of-speech for vector terms are problematic. In particular, the vector length can potentially be overly great when the vectors comprise *noun* terms, because nouns are generally constitute an open class.

The other category of word similarity approach uses a hand-crafted thesaurus (such as Roget’s thesaurus [14] or WordNet [93] in the case of English, and *Bunruigoihyo* [102] or EDR [58] in the case of Japanese), based on the intuitively feasible assumption that words located near each other within the structure of a thesaurus have similar meaning. Therefore, the similarity between two given words is represented by the length of the path between them in the thesaurus structure [77, 81, 142]. Unlike the former approach, the required computational cost can be restricted to constant order, because only a list of semantic codes for each word is required. For example, the commonly used Japanese *Bunruigoihyo* thesaurus [102] represents each semantic code with only 7 digits¹. However, computationally speaking, the relation between the similarity (namely the semantic length of the path), and the physical length of the path is not clear².

¹The revised version of the *Bunruigoihyo* thesaurus assigns an 8 digit code each word.

²Most researchers *heuristically* define functions between the similarity and physical path length [76, 77, 81, 142].

Furthermore, since most thesauri aim at a *general* word hierarchy, the similarity between words used in specific domains (technical terms) cannot be measured to the desired level of accuracy.

To sum up, the use of hand-crafted thesauri is preferable to statistics-based approaches in terms of “efficiency”. As for the “accuracy”, while our experiments in Section 3.4 quantitatively showed that the use of the *Bunruigoihyo* thesaurus outperformed the vector space model, there has been no qualitative explanation for this general tendency. In addition, new and improved models can potentially be introduced for statistics-based approaches. In view of these arguments, we aim at intergrating the advantages of the two above methodological types, or more precisely, realizing statistics-based word similarity based on the length of the thesaurus path [34]. Consequently, our approach allows us to measure the statistics-based word similarity, while retaining optimal computational cost ($O(1)$). The crucial concern in this process is how to determine the ‘statistics-based length’ (SBL) of each branch in a thesaurus. We tentatively use the *Bunruigoihyo* thesaurus, in which each word corresponds to a leaf in the tree structure. Let us take Figure 5.1, which shows a fragment of the thesaurus. In this figure, w_i ’s denote words and x_i ’s denote the statistics-based length of each branch i . Let the statistics-based (vector space model) word similarity between w_1 and w_2 be $vsm(w_1, w_2)$. We hope to estimate this similarity by the length of the path through branches 3 and 4, and derive an equation “ $x_3 + x_4 = sim(w_1, w_2)$ ”. Intuitively speaking, any combination of x_3 and x_4 which satisfies this equation can constitute the SBLs for branches 3 and 4. Formalizing equations for other pairs of words in the same manner, we can derive the set of simultaneous equations shown in Figure 5.2. That is, we can assign the SBL for each branch by way of finding answers for each x_i .

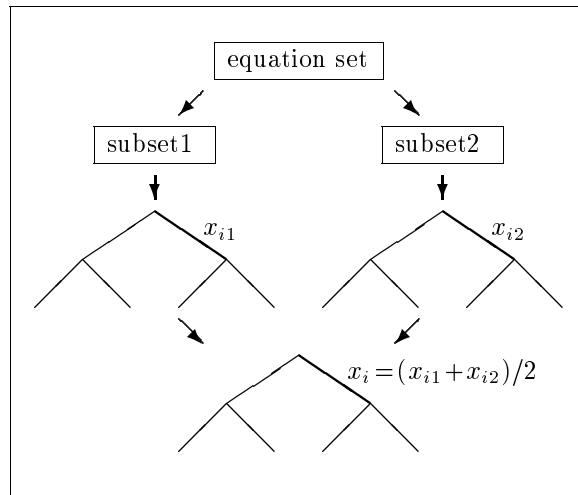


Figure 5.1: A fragment of the thesaurus associated with x_1 to x_6

$$\begin{array}{rcl}
 x_1 + x_2 + x_3 + x_5 & = & vsm(w_1, w_3) \\
 x_1 + x_2 + x_3 + x_6 & = & vsm(w_1, w_4) \\
 x_1 + x_2 + x_4 + x_5 & = & vsm(w_2, w_3) \\
 & \vdots &
 \end{array}$$

Figure 5.2: A fragment of the set of simultaneous equations associated with Figure 5.1

In Section 5.2, we elaborate on the methodology of our word similarity measurement. We then evaluate our method by way of an experiment in Section 5.3 and apply this method to the task of word sense disambiguation in Section 5.4.

5.2 Methodology

5.2.1 Overview

Our word similarity measurement proceeds in the following way:

1. compute the statistics-based similarity of every combination of given words,
2. set up a set of simultaneous equations through use of the thesaurus and previously computed word similarities, and find solutions for the statistics-based length (SBL) of the corresponding thesaurus branch (see Figures 5.1 and 5.2),
3. the similarity between two given words is measured by the sum of SBLs included in the path between those words.

For step 1, we used the vector space model, for which the reader is referred to Section 3.3.1. However, note that our framework is independent of the implementation of the similarity computation. We elaborate on steps 2 and 3 in the following sections.

5.2.2 Resolution of the simultaneous equations

The set of simultaneous equations used in our method is expressed by Equation 5.1, where A is a matrix comprising only the values 0 and 1, and B is the list of vsm 's (i.e. statistics-based similarities) for all possible combinations of given words. X is a list of variables, which represents the statistics-based length (SBL) for the corresponding branch in the thesaurus.

$$AX = B \tag{5.1}$$

Here, let the i -th similarity in B be $vs_m(a, b)$, and let $path(a, b)$ denote the path between words a and b in the thesaurus. Each equation contained in the set of simultaneous equations is represented by Equation 5.2, where x_j is the statistics-based length (SBL) for branch j , and α_{ij} is either 0 or 1, as in Equation 5.3.

$$[\alpha_{i1} \ \alpha_{i2} \ \dots \ \alpha_{ij} \ \dots] \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_j \\ \vdots \end{bmatrix} = vs_m(a, b) \quad (5.2)$$

$$\alpha_{ij} = \begin{cases} 1 & \text{if } j \in path(a, b) \\ 0 & \text{otherwise} \end{cases} \quad (5.3)$$

By finding the solutions for X , we can assign SBLs to branches. However, the set of similarity values outnumbers the variables. For example, the revised version of the *Bunruigoihyo* thesaurus contains about 55,000 noun entries, and therefore, the number of similarity values for those nouns becomes about 1.5×10^9 ($C_2^{55,000}$). On the other hand, the number of the branches is only about 53,000. As such, overly many equations are redundant, and the time complexity to solve the simultaneous equations becomes a crucial problem. To counter this problem, we randomly divide the overall equation set into equal parts, which can be solved reasonably. Thereafter we approximate the solution for x by averaging the solutions for x derived from each subset. Let us take Figure 5.3, in which the number of subsets is given as two without loss of generality. In this figure, x_{i1} and x_{i2} denote the answers for branch i individually derived from subsets 1 and 2, and x_i is approximated by the average of x_{i1} and x_{i2} (that is, $\frac{x_{i1} + x_{i2}}{2}$). To generalize this notion, let x_{ij} denote the solution associated with branch i in subset j . The approximate solution for branch i is given by Equation 5.4, where n is the number of divisions of the equation set.

$$x_i = \frac{1}{n} \sum_{j=1}^n x_{ij} \quad (5.4)$$

With regard to resolving the simultaneous equations, we used the mathematical analysis tool “MATLAB”³.

³Developed by Cybernet System, Inc.

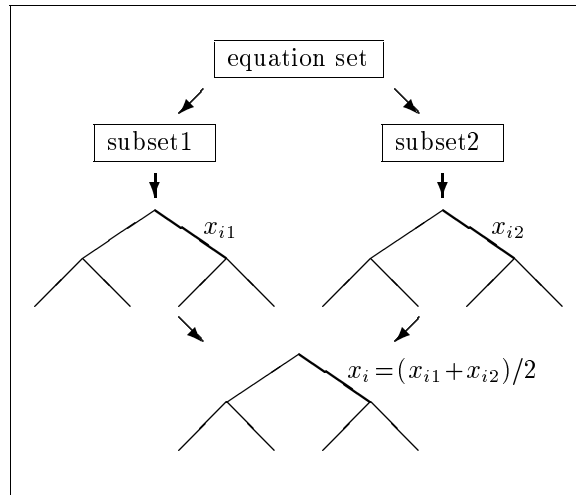


Figure 5.3: Approximation of the statistics-based length x_i

5.2.3 Word similarity using SBL

Let us reconsider Figure 5.1. In this figure, the similarity between w_1 and w_2 , for example, is measured by the sum of x_3 and x_4 . In general, the similarity between words a and b using SBL ($sbl(a, b)$, hereafter) is realized by Equation 5.5, where x_i is the SBL for branch i , and $path(a, b)$ is the path that includes thesaurus branches located between a and b .

$$sbl(a, b) = \sum_{i \in path(a, b)} x_i \quad (5.5)$$

5.3 Experimentation

Here, we evaluate the degree to which simultaneous equations are successfully approximated through the use of the technique described in Section 5.2. In other words, we analyze to what extent the (original) statistics-based word similarity can be realized by our framework. We conducted this evaluation in the following way. Let the statistics-based similarity between words a and b be $vsm(a, b)$, and the similarity based on SBL be $sbl(a, b)$. Here, let us assume the inequality “ $vsm(a, b) > vsm(c, d)$ ” for words a, b, c and d . If this inequality can be maintained for our method, that is, “ $sbl(a, b) > sbl(c, d)$ ”, the similarity measurement is taken to be successful. The accuracy is then estimated by the ratio between the number of successful measurements and the total number of trials. Since resolution of equations is time-consuming, we tentatively generalized 23,223 nouns into 303 semantic classes (represented by the first 4 digits of the semantic code given in the *Bunruigoihyo* thesaurus), reducing the total number of equations to

45,753. Figure 5.4 shows the relation between the number of equations used and the accuracy: we divided the overall equation set into n equal subsets⁴ (see Section 5.2.2), and progressively increased the number of subsets used in the computation. When the whole set of equations was provided, the accuracy became about 72%. We also estimated the lower bound of this evaluation, that is, we also conducted the same trials using the *Bunruigoihyo* thesaurus. In this case, if word a is more closely located to b than c is to d and “ $vsm(a, b) > vsm(c, d)$ ”, that trial measurement is taken to be successful. We found that the lower bound method produced an accuracy of about 56%, and therefore, that our framework outperformed this method.




Figure 5.4: The relation between the number of equations used and the accuracy

5.4 Application to verb sense disambiguation system

We further evaluated our word similarity technique in the task of verb sense disambiguation. The evaluation methodology is simple, that is, the performance of word similarity measurement is evaluated in the context of the improvement on our similarity-based system. In practice,

⁴We arbitrarily set $n = 15$ so as to be able to resolve equations reasonably.

Equation 5.5 is used as a substitution for sim in Equation 3.2 (see Section 3.3 for details). As performed in Section 3.4, this experiment involved six-fold cross validation, that is, we divided the training/test data into six equal parts, and conducted six trials in which a different part was used as test data each time, and the rest as the database. We evaluated the performance of the system according to its accuracy, that is the ratio of the number of correct outputs, compared to the total number of inputs. We used as training/test data the same corpus as in Section 3.4, where sentences contain one of eleven polysemous verbs, producing a total of 1315 sentences (see Table 3.2). We found that the combined performance with SBLs did not differ from that for the *Bunruigoihyo* thesaurus (see the column “BGH” in Table 3.2), which means our method still finds it difficult to enhance existing hand-crafted thesauri. However, it should be noted that given the inferiority of the vector space model as shown in Table 3.2, our word similarity measure cannot theoretically exceed the performance for the *Bunruigoihyo* thesaurus, because the objective of SBL is to *simulate* word similarity based on the vector space model. The use of SBL is expected to improve other applications where the vector space model outperforms hand-crafted thesauri.

5.5 Related Work

In this chapter, we focused solely on the “efficiency” problem related to the word similarity computation, disregarding the “accuracy” problem. Resnik [122] (later enhanced by Jiang and Conrath [59]) integrates a thesaurus taxonomy and information content for similarity computation, focusing on the accuracy problem. Resnik’s method computes similarity between given word classes (synsets) taken from WordNet [93], based on the degree of *commonality* between them. In other words, the more information two synsets share in common, the more similar they are. Formally speaking, the similarity between two synsets c_1 and c_2 is computed by Equation 5.6.

$$sim(c_1, c_2) = \max_{c \in S(c_1, c_2)} [-\log P(c)] \quad (5.6)$$

Here, $S(c_1, c_2)$ is the set of synsets that dominate both c_1 and c_2 . $P(c)$ is estimated based on the distribution of words associated with synset c , obtained from a corpus. Resnik [122] (and Jiang and Conrath [59]) showed the effectivity of this method using 30 pairs of words, which seems a relatively small data size. We applied Resnik’s method to our verb sense disambiguation system, using the *Bunruigoihyo* as a core thesaurus. We conducted the same procedure as performed in Section 3.4, and found that the accuracy of verb sense disambiguation was about 23%. It should be noted that this result represents only one example among the numerous usages of

word similarity, and does not discount the validity of Resnik's method. However, we claim that enhancement of the accuracy problem for similarity computation still remains difficult, and needs to be further explored.

5.6 Summary

The previous methods for word similarity computation can be divided into two approaches. Statistics-based approaches have been popular because of their mathematical rigor. However, these approaches generally require prohibitive computational cost. Remaining approaches, which use a hand-crafted thesaurus, are computationally cheaper. However, the quality of these approaches is highly dependent on a priori human judgement. In this chapter, we proposed a new method to integrate two different approaches for word similarity computation [34]. That is, we realized statistics-based word similarity computation with complexity equivalent to that for thesaurus-based approaches. Our objective is to determine the statistical weight for each thesaurus branch, so that one can measure statistics-based similarity by simply traversing the thesaurus structure. For this purpose, we set up a set of simultaneous equations, in which each variable and value correspond to a thesaurus branch and statistics-based similarity, respectively. By resolving the simultaneous equations, we can expect to assign appropriate statistical weights to each branch. By way of this method, we can expect to optimize the computational efficiency required for our verb sense disambiguation.

Chapter 6

Conclusion

6.1 Contribution

In this research, we targetted various problems associated with corpus-based word sense disambiguation, focusing on verbal polysemy. Let us summarize the main points that have been made in this research.

- First, in Chapter 3, we described our similarity-based verb sense disambiguation system and demonstrated its effectivity. One of the major features of this chapter was that we modeled and computationally implemented the linguistic behavior of ‘case contribution to disambiguation’ (CCD). Through comparative experiments, we confirmed that the performance of our system was improved by considering the CCD factor. Our experiments also showed that a hand-crafted thesaurus (in our case, the *Bunruigoihyo* thesaurus [102]) is an effective resource for word sense disambiguation. Due to the lack of a large-sized sense-annotated test collection for Japanese, we used a relatively small-sized corpus for our experimentation. However, our experimental results are expected to reflect general performance tendencies to a certain degree. We further enhanced our system by way of computation of the degree of interpretation certainty, so as to obtain more reliable outputs. In addition, we proposed a prototype method of propagating contextual constraints when a unique polysemous verb appears in an input sentence. It should be noted that while our experiments were carried out on Japanese, our proposed system can be applied to other languages.
- Although the performance of our similarity-based system proved to be relatively satisfactory, we identified the overhead for manual sense annotation (overhead for supervision) and the overhead searching a database (overhead for search) as being limitations on oper-

ational applications. Given this observation, in Chapter 4, we proposed a novel sampling method, which minimizes the overhead for supervision and the overhead for search. Note that these overheads can be major drawbacks for corpus-based NLP approaches in general. Our sampling method is characterized by its reliance on the notion of ‘training utility’: the degree to which each example is informative for future sampling when incorporated into the database. Comparative experiments showed that our sampling method reduced both overheads to a larger degree than any existing sampling methods, without degrading the performance of verb sense disambiguation. It should be noted that example sets sampled by our method can also be used as initial training data to further improve the quality of unsupervised learning, as occurs in bootstrapping.

- Finally, in Chapter 5, we proposed a new method for word similarity computation, given that the performance of our system is highly dependent on the similarity computation between given words. Our method integrates the advantages of the two different approaches for similarity computation: the mathematical rigor of statistics-based approaches and easy implementation of thesaurus-based approaches. The core process is to determine the ‘statistics-based length’ (SBL) of each branch in a thesaurus, reflecting statistically computed word similarities. In this way, our method realizes statistics-based word similarity by simply traversing the thesaurus structure.

6.2 Outstanding Issues

Out standing issue is the problem of balancing qualitative and quantitative approaches [5, 69, 142] (see Chapter 2.2.1 for a description of these approaches). Among a number of past research attempts to tackle this issue, Uramoto [142], for example, used selectional restrictions and scalable similarity based on a thesaurus taxonomy, as (relatively shallow) constraints and preference conditions, respectively. However, given that Uramoto did not compare his integrated approach with monotonic approaches, the empirical study of balancing different approaches requires further exploration and is expected to improve on current word sense disambiguation methods.

Another direction would be the establishment of standardized evaluation criteria for word sense disambiguation systems, which is expected to streamline system comparison/enhancement¹. However, as discussed in Section 2.3, numerous problems exist associated with evaluation method-

¹ACL SIGLEX is going to hold a workshop (co-ordinated by Adam Kilgarriff) in September 1998, to initiate the standardized evaluation of word sense disambiguation systems for English, French, German, Italian and Spanish.

ologies. Ultimately speaking, we think that task-based (application-based) evaluation is the most straightforward, because the performance of a system is directly related to benefits to real-world applications. In addition, word sense distinction, which can be one major obstacle in standardized evaluation, can be fixed relatively easily for specific applications. However, on the other hand, numerous sense-annotated corpora are required for different applications. Our sampling method (as well as past efforts targetting unsupervised methods) are expected to reduce the human overhead for establishing these types of sense-annotated corpora.

Bibliography

- [1] David W. Aha, Dennis Kibler, and Marc K. Albert. Instance-based learning algorithms. *Machine Learning*, Vol. 6, No. 1, pp. 37–66, 1991.
- [2] Timothy Baldwin. Relative clause coordination and subordination in Japanese. In *Proceedings of the Australian Natural Language Processing Post Graduate Workshop*, pp. 1–10, 1998.
- [3] Timothy Baldwin, Takenobu Tokunaga, and Hozumi Tanaka. Semantic verb classes in the analysis of head gapping in Japanese relative clauses. In *Proceedings of the Natural Language Processing Pacific Rim Symposium*, pp. 409–414, 1997.
- [4] Ray Bareiss. *Exemplar-Based Knowledge Acquisition*. Academic Press, 1990.
- [5] Roberto Basili, Maria Teresa Pazienza, and Paola Velardi. An empirical symbolic approach to natural language processing. *Artificial Intelligence*, Vol. 85, No. 1–2, pp. 59–99, 1996.
- [6] Roberto Basili, Michelangelo Della Rocca, and Maria Teresa Pazienza. Contextual word sense tuning and disambiguation. *Journal of Applied Artificial Intelligence*, Vol. 11, pp. 235–262, 1997.
- [7] Roberto Basili, Michelangelo Della Rocca, and Maria Teresa Pazienza. Towards a bootstrapping framework for corpus semantic tagging. In *Proceedings of ACL SIGLEX Workshop on Tagging Text with Lexical Semantics: Why, What, and How?*, pp. 66–72, 1997.
- [8] Roberto Basili, Paola Velardi, and Maria Teresa Pazienza. Integrating general-purpose and corpus-based verb classification. *Computational Linguistics*, Vol. 22, No. 4, pp. 559–568, 1996.
- [9] Michael W. Berry. Large-scale sparse singular value computations. *The International Journal of Supercomputer Applications*, Vol. 6, No. 1, pp. 13–49, 1992.

- [10] Peter F. Brown, Stephen A. Della Pietra, and Vincent J. Della Pietra. Word-sense disambiguation using statistical methods. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, pp. 264–270, 1991.
- [11] Peter F. Brown, Vincent J. Della Pietra, Peter V. deSouza, Jenifer C. Lai, and Robert L. Mercer. Class-based n -gram models of natural language. *Computational Linguistics*, Vol. 18, No. 4, pp. 467–479, 1992.
- [12] Bertram Bruce. Case systems for natural language. *Artificial Intelligence*, Vol. 6, No. 4, pp. 327–360, 1975.
- [13] Rebecca Bruce and Janyce Wiebe. Word-sense disambiguation using decomposable models. In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics*, pp. 139–146, 1994.
- [14] Robert L. Chapman. *Roget's International Thesaurus (Fourth Edition)*. Harper and Row, 1984.
- [15] Eugene Charniak. A neat theory of marker passing. In *Proceedings of AAAI-86*, pp. 584–588, 1986.
- [16] Eugene Charniak. *Statistical Language Learning*. MIT Press, 1993.
- [17] Eugene Charniak and Robert Goldman. A logic for semantic interpretation. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, pp. 87–94, 1988.
- [18] Peter Cheeseman, James Kelly, Matthew Self, John Stutz, Will Taylor, and Don Freeman. AutoClass: A Bayesian classification system. In *Proceedings of the 5th International Conference on Machine Learning*, pp. 54–64, 1988.
- [19] Jeong-Mi Cho and Gil Chang Kim. Korean verb sense disambiguation using distributional information from corpora. In *Proceedings of the Natural Language Processing Pacific Rim Symposium*, pp. 691–696, 1995.
- [20] Kenneth W. Church and Patrick Hanks. Word association norms, mutual information, and lexicography. *Computational Linguistics*, Vol. 16, No. 1, pp. 22–29, 1990.
- [21] Kenneth W. Church and Robert L. Mercer. Introduction to the special issue on computational linguistics using large corpora. *Computational Linguistics*, Vol. 19, No. 1, pp. 1–24, 1993.

- [22] Peter Clark and Robin Boswell. Rule induction with CN2: Some recent improvements. In *Proceedings of European Working Session on Learning*, pp. 151–163, 1991.
- [23] Peter Clark and Tim Niblett. The CN2 induction algorithm. *Machine Learning*, Vol. 3, No. 4, pp. 261–283, 1989.
- [24] David Cohn, Les Atlas, and Richard Ladner. Improving generalization with active learning. *Machine Learning*, Vol. 15, No. 2, pp. 201–221, 1994.
- [25] Scott Cost and Steven Salzberg. A weighted nearest neighbor algorithm for learning with symbolic features. *Machine Learning*, Vol. 10, No. 1, pp. 57–78, 1993.
- [26] Jim Cowie, Joe Guthrie, and Louise Guthrie. Lexical disambiguation using simulated annealing. In *Proceedings of the 14th International Conference on Computational Linguistics*, pp. 359–365, 1992.
- [27] Douglass R. Cutting, David R. Karger, Jan O. Pedersen, and John W. Tukey. Scatter/Gather: A cluster-based approach to browsing large document collections. In *Proceedings of the 15th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 318–329, 1992.
- [28] Ido Dagan and Alon Itai. Word sense disambiguation using a second language monolingual corpus. *Computational Linguistics*, Vol. 20, No. 4, pp. 563–596, 1994.
- [29] Ido Dagan, Fernando Pereira, and Lillian Lee. Similarity-based estimation of word co-occurrence probabilities. In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics*, pp. 272–278, 1994.
- [30] George C. Demetriou. Lexical disambiguation using constraint handling in prolog (CHIP). In *Proceedings of the 6th Conference of the European Chapter of the Association for Computational Linguistics*, pp. 431–436, 1993.
- [31] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, Vol. 39, pp. 1–38, 1977.
- [32] Sean P. Engelson and Ido Dagan. Minimizing manual annotation cost in supervised training from corpora. In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics*, pp. 319–326, 1996.
- [33] William B. Frakes and Ricardo Baeza-Yates. *Information Retrieval: Data Structure & Algorithms*. PTR Prentice-Hall, 1992.

- [34] Atsushi Fujii, Toshihiro Hasegawa, Takenobu Tokunaga, and Hozumi Tanaka. Integration of hand-crafted and statistical resources in measuring word similarity. In *Proceedings of the ACL-EACL Workshop on Automatic Information Extraction and Building of Lexical Semantic Resources*, pp. 45–51, 1997.
- [35] Atsushi Fujii, Kentaro Inui, Takenobu Tokunaga, and Hozumi Tanaka. Selective sampling for example-based word sense disambiguation. *Computational Linguistics*. (To appear).
- [36] Atsushi Fujii, Kentaro Inui, Takenobu Tokunaga, and Hozumi Tanaka. Selective sampling of effective example sentence sets for word sense disambiguation. In *Proceedings of the 4th Workshop on Very Large Corpora*, pp. 56–69, 1996.
- [37] Atsushi Fujii, Kentaro Inui, Takenobu Tokunaga, and Hozumi Tanaka. To what extent does case contribute to verb sense disambiguation? In *Proceedings of the 16th International Conference on Computational Linguistics*, pp. 59–64, 1996.
- [38] Atsushi Fujii, Kentaro Inui, Takenobu Tokunaga, and Hozumi Tanaka. Case contribution in example-based verb sense disambiguation. *Journal of Natural Language Processing*, Vol. 4, No. 2, pp. 111–123, 1997.
- [39] Atsushi Fujii, Takenobu Tokunaga, and Hozumi Tanaka. An empirical comparison of recent corpus-based word sense disambiguation techniques. In *Proceedings of the Natural Language Processing Pacific Rim Symposium*, pp. 201–206, 1997.
- [40] Fumiyo Fukumoto and Yoshimi Suzuki. An automatic clustering of articles using dictionary definitions. In *Proceedings of the 16th International Conference on Computational Linguistics*, pp. 406–411, 1996.
- [41] Fumiyo Fukumoto and Jun'ichi Tsujii. Automatic recognition of verbal polysemy. In *Proceedings of the 15th International Conference on Computational Linguistics*, pp. 764–768, 1994.
- [42] William Gale, Kenneth Ward Church, and David Yarowsky. Estimating upper and lower bounds on the performance of word-sense disambiguation programs. In *Proceedings of the 30th Annual Meeting of the Association for Computational Linguistics*, pp. 249–256, 1992.
- [43] William Gale, Kenneth Ward Church, and David Yarowsky. A method for disambiguating word senses in a large corpus. *Computers and the Humanities*, Vol. 26, pp. 415–439, 1992.
- [44] Dirk Geeraerts. Vagueness's puzzles, polysemy's vagaries. *Cognitive Linguistics*, Vol. 4, No. 3, pp. 223–272, 1993.

- [45] Gene H. Golub and Charles F. van Loan. *Matrix Computations*. John Hopkins University Press, 1989.
- [46] Ralph Grishman and John Sterling. Generalizing automatically generated selectional patterns. In *Proceedings of the 15th International Conference on Computational Linguistics*, pp. 742–747, 1994.
- [47] Joe A. Guthrie, Louise Guthrie, Yorick Wilks, and Homa Aidinejad. Subject-dependent co-occurrence and word sense disambiguation. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, pp. 146–152, 1991.
- [48] Daniel Hardt. An empirical approach to VP ellipsis. *Computational Linguistics*, Vol. 23, No. 4, pp. 525–541, 1997.
- [49] Vasileios Hatzivassiloglou and Kathleen R. McKeown. Towards the automatic identification of adjectival scales: Clustering adjectives according to meaning. In *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, pp. 172–182, 1993.
- [50] Marti A. Hearst. Noun homograph disambiguation using local context in large text corpora. In *Proceedings of the 7th Annual Conference of the University of Waterloo Centre for the New OED and Text Research*, pp. 1–22, 1991.
- [51] Marti A. Hearst. Automatic acquisition of hyponyms from large text corpora. In *Proceedings of the 14th International Conference on Computational Linguistics*, pp. 539–545, 1992.
- [52] Marti A. Hearst and Hinrich Schütze. Customizing a lexicon to better suit a computational task. Branimir Boguraev and James Pustejovsky (ed.), *Corpus Processing for Lexical Acquisition*, MIT Press.
- [53] Donald Hindle. Noun classification from predicate-argument structures. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*, pp. 268–275, 1990.
- [54] Graeme Hirst. *Semantic Interpretation and the Resolution of Ambiguity*. Cambridge University Press, 1987.
- [55] Satoru Ikehara, et al. *Goi-Taikei — A Japanese Lexicon*. Iwanami Shoten, 1997. (In Japanese).

- [56] Information-technology Promotion Agency. *IPAL Japanese Dictionary for Computers (Basic Verbs)*, 1987. (In Japanese).
- [57] Makoto Iwayama and Takenobu Tokunaga. A probabilistic model for text categorization: Based on a single random variable with multiple values. In *Proceedings of the 4th Conference on Applied Natural Language Processing*, pp. 162–167, 1994.
- [58] Japan Electronic Dictionary Research Institute. EDR electronic dictionary technical guide, 1995. (In Japanese).
- [59] Jay J. Jiang and David W. Conrath. Semantic similarity based on corpus statistics and lexical taxonomy. In *Proceedings of ROCLING X*, pp. 19–33, 1997.
- [60] John S. Justeson and Slava M. Katz. Principled disambiguation: Discriminating adjective senses with modified nouns. *Computational Linguistics*, Vol. 21, No. 1, pp. 1–27, 1995.
- [61] Hiroyuki Kaji, Yuuko Kida, and Yasutsugu Morimoto. Learning translation templates from bilingual text. In *Proceedings of the 14th International Conference on Computational Linguistics*, pp. 672–678, 1992.
- [62] Masayuki Kameda. A portable & quick Japanese parser : QJP. In *Proceedings of the 16th International Conference on Computational Linguistics*, pp. 616–621, 1996.
- [63] Yael Karov and Shimon Edelman. Learning similarity-based word sense disambiguation. In *Proceedings of the 4th Workshop on Very Large Corpora*, pp. 42–55, 1996.
- [64] Jerrold Jacob Katz and Jerry Alan Fodor. The structure of a semantic theory. *Language*, Vol. 39, No. 2, pp. 170–210, 1963.
- [65] Andrew Kehler and Stuart Shieber. Anaphoric dependencies in ellipsis. *Computational Linguistics*, Vol. 23, No. 3, pp. 457–466, 1997.
- [66] Sue J. Ker and Jason S. Chang. A class-based approach to word alignment. *Computational Linguistics*, Vol. 23, No. 2, pp. 313–343, 1997.
- [67] Adam Kilgarriff. *Polysemy*. Ph.D. thesis, University of Sussex, 1992.
- [68] Adam Kilgarriff. What is word sense disambiguation good for? In *Proceedings of the Natural Language Processing Pacific Rim Symposium*, pp. 209–214, 1997.
- [69] Judith Klavans and Philip Resnik, editors. *The Balancing Act: Combining Symbolic and Statistical Approaches to Language*. MIT press, 1996.

- [70] Yoshiyuki Kobayashi, Takenobu Tokunaga, and Hozumi Tanaka. Analysis of Japanese compound nouns using collocational information. In *Proceedings of the 15th International Conference on Computational Linguistics*, pp. 865–869, 1994.
- [71] Yoshiyuki Kobayashi, Takenobu Tokunaga, and Hozumi Tanaka. Analysis of syntactic structure of Japanese compound noun. In *Proceedings of the Natural Language Processing Pacific Rim Symposium*, pp. 326–331, 1995.
- [72] Janet Kolodner. *Case-Based Reasoning*. Morgan Kaufmann, 1993.
- [73] Igor Kononenko and Ivan Bratko. Information-based evaluation criterion for classifier's performance. *Machine Learning*, Vol. 6, No. 1, pp. 67–80, 1991.
- [74] Robert Krovets and W. Bruce Croft. Lexical ambiguity and information retrieval. *ACM Transactions on Information Systems*, Vol. 10, No. 2, pp. 115–141, 1992.
- [75] Robert Krovetz. Homonymy and polysemy in information retrieval. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics and the 8th Conference of the European Chapter of the Association for Computational Linguistics (ACL-EACL'97)*, pp. 72–79, 1997.
- [76] Sadao Kurohashi and Makoto Nagao. Dynamic programming method for analyzing conjunctive structures in Japanese. In *Proceedings of the 14th International Conference on Computational Linguistics*, pp. 170–176, 1992.
- [77] Sadao Kurohashi and Makoto Nagao. A method of case structure analysis for Japanese sentences based on examples in case frame dictionary. *IEICE Transactions on Information and Systems*, Vol. E77-D, No. 2, pp. 227–239, 1994.
- [78] Claudia Leacock, Geoffrey Towell, and Ellen Voorhees. Corpus-based statistical sense resolution. In *Proceedings of ARPA Human Language Technology Workshop*, pp. 260–265, 1993.
- [79] Michael Lesk. Automatic sense disambiguation using machine-readable dictionaries: How to tell a pine cone from an ice cream cone. In *Proceedings of ACM SIGDOC*, pp. 24–26, 1986.
- [80] David D. Lewis and William A. Gale. A sequential algorithm for training text classifiers. In *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 3–12, 1994.

- [81] Xiaobin Li, Stan Szpakowicz, and Stan Matwin. A WordNet-based algorithm for word sense disambiguation. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pp. 1368–1374, 1995.
- [82] Dekang Lin. Using syntactic dependency as local context to resolve word sense ambiguity. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics and the 8th Conference of the European Chapter of the Association for Computational Linguistics (ACL-EACL'97)*, pp. 64–71, 1997.
- [83] Alpha K. Luk. Statistical sense disambiguation with relatively small corpora using dictionary definitions. In *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics*, pp. 181–188, 1995.
- [84] Steven L. Lytinen. Dynamically combining syntax and semantics in natural language processing. In *Proceedings of AAAI-86*, pp. 574–578, 1986.
- [85] Mainichi Shimbun. Mainichi shimbun CD-ROM '91-'94, 1991-1994. (In Japanese).
- [86] Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, Vol. 19, No. 2, pp. 313–330, 1993.
- [87] Hiroshi Maruyama. Structural disambiguation with constraint propagation. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*, pp. 31–38, 1990.
- [88] Yuji Matsumoto, Sadao Kurohashi, Takehito Utsuro, Yutaka Myoki, and Makoto Nagao. *JUMAN Users Manual*. Kyoto University and Nara Institute of Science and Technology, 1993. (In Japanese).
- [89] L. McQuitty. Similarity analysis by reciprocal pairs for discrete and continuous data. *Educational and Psychological Measurement*, Vol. 26, pp. 825–831, 1966.
- [90] Susan W. McRoy. Using multiple knowledge sources for word sense discrimination. *Computational Linguistics*, Vol. 18, No. 1, pp. 1–30, 1992.
- [91] C. S. Mellish. *Computer Interpretation of Natural Language Descriptions*. Ellis Horwood, 1985.
- [92] A. Michiels and J. Noël. Approaches to thesaurus production. In *Proceedings of the 7th International Conference on Computational Linguistics*, pp. 227–232, 1982.

- [93] George A. Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, Katherine Miller, and Randee Teng. Five papers on WordNet. Technical Report CLS-Rep-43, Cognitive Science Laboratory, Princeton University, 1993.
- [94] George A. Miller, Martin Chodorow, Shari Landes, Claudia Leacock, and Robert G. Thomas. Using a semantic concordance for sense identification. In *Proceedings of ARPA Human Language Technology Workshop*, pp. 240–243, 1994.
- [95] George A. Miller, Claudia Leacock, Randee Teng, and Ross T. Bunker. A semantic concordance. In *Proceedings of ARPA Human Language Technology Workshop*, pp. 303–308, 1993.
- [96] Raymond J. Mooney. Comparative experiments on disambiguating word senses: An illustration of the role of bias in machine learning. In *Proceedings of the 1st Conference on Empirical Methods in Natural Language Processing*, pp. 82–91, 1996.
- [97] Masaki Murata and Makoto Nagao. Resolution of verb ellipsis in Japanese sentence using surface expressions and examples. In *Proceedings of the Natural Language Processing Pacific Rim Symposium*, pp. 75–80, 1997.
- [98] Katashi Nagao. A preferential constraint satisfaction technique for natural language analysis. *IEICE Transactions on Information and Systems*, Vol. E77-D, No. 2, pp. 161–170, 1994.
- [99] Makoto Nagao. A framework of a mechanical translation between Japanese and English by analogy principle. *Artificial and Human Intelligence*, pp. 173–180, 1984.
- [100] Jun'ichi Nakamura and Makoto Nagao. Extraction of semantic information from an ordinary English dictionary and its evaluation. In *Proceedings of the 10th International Conference on Computational Linguistics*, pp. 459–464, 1988.
- [101] Tetsuya Nasukawa. Discourse constraint in computer manuals. In *Proceedings of the 5th International Conference on Theoretical and Methodological Issues in Machine Translation*, pp. 183–194, 1993.
- [102] National Language Research Institute. *Bunruigoihyo*. Shuei publisher, 1964. (In Japanese).
- [103] Hwee Tou Ng. Exemplar-based word sense disambiguation: Some recent improvements. In *Proceedings of the 2nd Conference on Empirical Methods in Natural Language Processing*, pp. 208–213, 1997.

- [104] Hwee Tou Ng. Getting serious about word sense disambiguation. In *Proceedings of ACL SIGLEX Workshop on Tagging Text with Lexical Semantics: Why, What, and How?*, pp. 1–7, 1997.
- [105] Hwee Tou Ng and Hian Beng Lee. Integrating multiple knowledge sources to disambiguate word sense: An exemplar-based approach. In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics*, pp. 40–47, 1996.
- [106] Yoshiki Niwa and Yoshihiko Nitta. Co-occurrence vectors from corpora vs. distance vectors from dictionaries. In *Proceedings of the 15th International Conference on Computational Linguistics*, pp. 304–309, 1994.
- [107] Hiroshi Nomiya. Machine translation by case generalization. *Transactions of Information Processing Society of Japan*, Vol. 34, No. 5, pp. 905–912, 1993. (In Japanese).
- [108] Manabu Okumura and Yuukichi Matsunaga. Word sense disambiguation by marker passing on very large semantic networks. In *Proceedings of the Natural Language Processing Pacific Rim Symposium*, pp. 71–76, 1995.
- [109] Manabu Okumura and Hozumi Tanaka. Towards incremental disambiguation with a generalized discrimination network. In *Proceedings of AAAI-90*, pp. 990–995, 1990.
- [110] Ted Pedersen and Rebecca Bruce. Distinguishing word senses in untagged text. In *Proceedings of the 2nd Conference on Empirical Methods in Natural Language Processing*, pp. 197–207, 1997.
- [111] Ted Pedersen and Rebecca Bruce. A new supervised learning algorithm for word sense disambiguation. In *Proceedings of AAAI/IAAI-97*, pp. 604–609, 1997.
- [112] Ted Pedersen, Rebecca Bruce, and Janyce Wiebe. Sequential model selection for word sense disambiguation. In *Proceedings of the 5th Conference on Applied Natural Language Processing*, pp. 388–395, 1997.
- [113] Fernando Pereira, Naftali Tishby, and Lillian Lee. Distributional clustering of English words. In *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, pp. 183–190, 1993.
- [114] Paul Procter. *Longman Dictionary of Contemporary English*. Harlow and London, 1978.
- [115] James Pustejovsky and Branimir Boguraev. Lexical knowledge representation and natural language processing. *Artificial Intelligence*, Vol. 63, No. 1–2, pp. 193–223, 1993.

- [116] J. Ross Quinlan. Learning logical definitions from relations. *Machine Learning*, Vol. 5, No. 3, pp. 239–266, 1990.
- [117] J. Ross Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, 1993.
- [118] Yael Ravin. Disambiguating and interpreting verb definitions. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*, pp. 260–267, 1990.
- [119] Real World Computing Partnership. *RWC text database*, 1995. (In Japanese).
- [120] Philip Resnik. *Selection and Information: A Class-Based Approach to Lexical Relationships*. Ph.D. thesis, University of Pennsylvania, 1993.
- [121] Philip Resnik. Disambiguating noun groupings with respect to WordNet senses. In *Proceedings of the 3rd Workshop on Very Large Corpora*, 1995.
- [122] Philip Resnik. Using information content to evaluate semantic similarity in a taxonomy. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pp. 448–453, 1995.
- [123] Philip Resnik and David Yarowsky. A perspective on word sense disambiguation methods and their evaluation. In *Proceedings of ACL SIGLEX Workshop on Tagging Text with Lexical Semantics: Why, What, and How?*, pp. 79–86, 1997.
- [124] Francesc Ribas. On learning more appropriate selectional restrictions. In *Proceedings of the 7th Conference of the European Chapter of the Association for Computational Linguistics*, pp. 112–118, 1995.
- [125] German Rigau, Jordi Atserias, and Eneko Agirre. Combining unsupervised lexical knowledge methods for word sense disambiguation. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics and the 8th Conference of the European Chapter of the Association for Computational Linguistics (ACL-EACL'97)*, pp. 48–55, 1997.
- [126] Ronald L. Rivest. Learning decision lists. *Machine Learning*, Vol. 2, pp. 229–246, 1987.
- [127] Gerard Salton and Michael J. McGill. *Introduction to Modern Information Retrieval*. McGraw-Hill, 1983.
- [128] Mark Sanderson. Word sense disambiguation and information retrieval. In *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 142–151, 1994.

- [129] Hinrich Schütze. Automatic word sense discrimination. *Computational Linguistics*. (To appear).
- [130] Hinrich Schütze. Dimensions of meaning. In *Proceedings of Supercomputing*, pp. 787–796, 1992.
- [131] Hinrich Schütze and Jan O. Pedersen. Information retrieval based on word senses. In *Proceedings of the 4th Annual Symposium on Document Analysis and Information Retrieval*, pp. 161–175, 1995.
- [132] H. S. Seung, M. Opper, and H. Sompolinsky. Query by committee. In *Proceedings of the 5th Annual ACM Workshop on Computational Learning Theory*, pp. 287–294, 1992.
- [133] John Sinclair, editor. *Collins COBUILD English Language Dictionary*. Collins and the University of Birmingham, 1987.
- [134] Frank A. Smadja and Kathleen R. McKeown. Automatically extracting and representing collocations for language generation. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*, pp. 252–259, 1990.
- [135] Steven L. Small, Garrison W. Cottrell, and Michael K. Tanenhaus, editors. *Lexical Ambiguity Resolution*. Morgan Kaufmann, 1988.
- [136] Barry Smyth and Mark T. Keane. Remembering to forget: A competence-preserving case deletion policy for case-based reasoning systems. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pp. 377–382, 1995.
- [137] Hideki Tanaka. Verbal case frame acquisition from a bilingual corpus: Gradual knowledge acquisition. In *Proceedings of the 15th International Conference on Computational Linguistics*, pp. 727–731, 1994.
- [138] Takanobu Tokunaga, Atsushi Fujii, Makoto Iwayama, Naoyuki Sakurai, and Hozumi Tanaka. Extending a thesaurus by classifying words. In *Proceedings of the ACL-EACL Workshop on Automatic Information Extraction and Building of Lexical Semantic Resources*, pp. 16–21, 1997.
- [139] Takanobu Tokunaga, Makoto Iwayama, and Hozumi Tanaka. Automatic thesaurus construction based on grammatical relations. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pp. 1308–1313, 1995.

- [140] David Tuggy. Ambiguity, polysemy, and vagueness. *Cognitive Linguistics*, Vol. 4, No. 3, pp. 273–290, 1993.
- [141] Naohiko Uramoto. A best-match algorithm for broad-coverage example-based disambiguation. In *Proceedings of the 15th International Conference on Computational Linguistics*, pp. 717–721, 1994.
- [142] Naohiko Uramoto. Example-based word-sense disambiguation. *IEICE Transactions on Information and Systems*, Vol. E77-D, No. 2, pp. 240–246, 1994.
- [143] Naohiko Uramoto. Positioning unknown words in a thesaurus by using information extracted from a corpus. In *Proceedings of the 16th International Conference on Computational Linguistics*, pp. 956–961, 1996.
- [144] Takehito Utsuro. Sense classification of verbal polysemy based-on bilingual class/class association. In *Proceedings of the 16th International Conference on Computational Linguistics*, pp. 968–973, 1996.
- [145] Takehito Utsuro, Kiyotaka Uchimoto, Mitsutaka Matsumoto, and Makoto Nagao. Thesaurus-based efficient example retrieval by generating retrieval queries from similarities. In *Proceedings of the 15th International Conference on Computational Linguistics*, pp. 1044–1048, 1994.
- [146] Paola Velardi, Maria Teresa Pazienza, and Michela Fasolo. How to encode semantic knowledge: A method for meaning representation and computer-aided acquisition. *Computational Linguistics*, Vol. 17, No. 2, pp. 153–170, 1991.
- [147] Ellen M. Voorhees. Using WordNet to disambiguate word senses for text retrieval. In *Proceedings of the 16th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 171–180, 1993.
- [148] J. Ward. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, Vol. 58, pp. 236–244, 1963.
- [149] S. M. Weiss and C. A. Kulikowski. *Computer Systems That Learn*. Morgan Kaufmann, 1991.
- [150] Yorick Wilks. A preferential, pattern-seeking, semantics for natural language inference. *Artificial Intelligence*, Vol. 6, No. 1, pp. 53–74, 1975.

- [151] Yorick Wilks, et al. Providing machine tractable dictionary tools. James Pustejovsky (ed.), *Semantics and the Lexicon*, Kluwer Academic, 1993.
- [152] Yorick Wilks and Mark Stevenson. Sense tagging: Semantic tagging with a lexicon. In *Proceedings of ACL SIGLEX Workshop on Tagging Text with Lexical Semantics: Why, What, and How?*, pp. 47–51, 1997.
- [153] Anthony Woods, Paul Fletcher, and Arthur Hughes. *Statistics in language studies*. Cambridge University Press, 1986.
- [154] David Yarowsky. Word-sense disambiguation using statistical models of Roget’s categories trained on large corpora. In *Proceedings of the 14th International Conference on Computational Linguistics*, pp. 454–460, 1992.
- [155] David Yarowsky. One sense per collocation. In *Proceedings of ARPA Human Language Technology Workshop*, pp. 266–271, 1993.
- [156] David Yarowsky. Decision lists for lexical ambiguity resolution: Application to accent restoration in Spanish and French. In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics*, pp. 88–95, 1994.
- [157] David Yarowsky. Unsupervised word sense disambiguation rivaling supervised methods. In *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics*, pp. 189–196, 1995.
- [158] Uri Zernik. Lexicon acquisition: Learning from corpus by capitalizing on lexical categories. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pp. 1556–1562, 1989.
- [159] Uri Zernik and Paul Jacobs. Tagging for learning: Collecting thematic relations from corpus. In *Proceedings of the 12th International Conference on Computational Linguistics*, Vol. 1, pp. 34–39, 1990.