Word Sense Implantation by Orthographical Conversion

Kazuhide Yamamoto, Yuki Gumizawa and Yuki Mikami
Nagaoka University of Technology
Nagaoka City, Niigata Prefecture, 940-2188 JAPAN
{yamamoto, gumizawa, mikami}@jnlp.org

Abstract

We present a word sense disambiguation (WSD) tool of Japanese Hiragana words. Unlike other WSD tasks which output something like "sense #3" as a result, our WSD task rewrites the target Hiragana word into a Kanji word, which is a different orthography. This means that the task is also a kind of orthographical normalization as well as WSD. We use pointwise mutual information (PMI) for disambiguating word. Our proposed method shows 90 percent of Hiragana words can be accurately converted to Kanji words. In this paper, we present the detail of task, our proposed method, and the performance.

Keywords

Word sense disambiguation; Lexical substitution, Pointwise mutual information, Orthographical normalization

1. Introduction

Word sense disambiguation (WSD) problem is a very important task in natural language processing (NLP) and various methods have been studied since long ago (Navigli, 2009; Navigli 2012). It is reported that ambiguous words cause performance drops if we do not identify the sense of a word in many tasks (Sanderson, 1994). This problem recognized and shared by many NLP researchers, and they developed many WSD techniques such as supervised learning (Yarowsky, 1995), unsupervised learning (Petersen, 2006; Wawer and Mykowiecka, 2017), knowledge-based (Agirre and Soroa, 2009), and so on.

However, a public tool which automatically determines the sense of a word is not available in Japanese[1]. There are many tools for other tasks such as part-of-speech tagging and named entity recognition. However, no tools are available for WSD. Lopez de Lacalle et al. says that it is difficult to integrate current WSD systems into downstream NLP applications. (de Lacalle and Agirre, 2015).

We thought that it is due to a systematic reason. It is difficult for WSD tool to use it alone, i.e. use it independently to other tasks and language resources. For example, we do not understand well when a word is told to be something like “sense no. 3” in a word dictionary. Although this information is indeed useful, it is not easy to utilize this in the following procedures. In addition, the consistency of word system between WSD and word segmentation tool is required in unsegmented language such as Japanese and Chinese. Consequently, we need to solve these systematic problems to release WSD as a tool.

In this paper, we take a different approach to avoid the above-mentioned problems, and report our WSD tool. We propose a new task of Japanese WSD to use the characteristics of Japanese language: Hiragana-Kanji conversion task. We will describe this task in Section 3.

The purpose of this work is to develop this kind of WSD tool and release to the public, not to propose a new WSD technique. It is important for a tool to be high accuracy when it determine the sense, i.e. not to output wrong results. Hence in this work we investigate the accuracy of conversion word by word, and adopt only high accuracy phenomena into the tool. In addition, we adopt simple mechanism for WSD because of administrative point of view. Each of them is considered less serious in the academy, but we believe they are all important in the engineering point of view.

In this article, we present detail of the Hiragana-Kanji conversion task. We also propose a method to solve semantic ambiguity using only simple pointwise mutual information. We then report an experiment to verify the effectiveness of the method and have some discussions.

We note that our disambiguation tool described in this paper has been released this year on an easy-to-use interface on the Web.

2. Related Works

A number of works have been proposed to disambiguate word sense. WSD methods are mainly classified into two categories: supervised learning and unsupervised learning.

It is known that methods using supervised learning provide higher accuracy than

[1] Besides Japanese we can find some WSD tools such as https://hinoki-project.org/asunaro/index-j.php

unsupervised ones (Raganato et al., 2017). In fact, it is the case in a SemEval workshop (Pradhan, 2007; Mihalcea, 2004). As features we usually use words around the target word that has multiple senses. However, supervised method needs a lot of text where sense of words are annotated. In WSD tasks, learning various patterns by increasing training data is necessary to improve accuracy, because words in the context of the target polysemous word provide clues (Yarowsky, 1995). In order to solve this problem, there is a study that addresses supervised WSD with word embeddings (Sugawara, 2015; Iacobacci, 2016); in particular, word vectors provided by Skip-grams which Mikolov et al. (2013) proposed have superior performance when compared to the conventional count model (Baroni, 2014).

A number of unsupervised method is also proposed (Petersen, 2006). Distributional methods, which is one of the unsupervised WSD method, identify words that appear in similar contexts without regard to any particular underlying sense inventory. Recently, it is reported that the method using word embedding for calculate cosine similarity and separate word sense has better performance than other unsupervised method (Wawer and Mykowiecka, 2017).

Knowledge-based method have been studied since long ago. UKB is knowledge-based WSD system that are using graph information in WordNet (Agirre and Soroa, 2009). UKB system is available in the Web and we can use it in English and Spanish.

Regarding Japanese, there are no WSD systems or tools. A Japanese reading system “ASUNARO” (Abekawa et al., 2002) changes the order of multiple senses automatically by likelihood, but it does not determine word sense. Although Japanese WSD task is available in SemEval 2010 (Okumura et al., 2010), it seems that there is no contribute to other NLP tasks.

3. Hiragana-Kanji Conversion

Japanese is written in four different sets of scripts: Kanji, Hiragana, Katakana, and Romaji (Halpern, 2002). Among them, Kanji is a logographic system consisting of characters borrowed from the Chinese characters. Hiragana is a phonographic character and thus semantically ambiguous, in which they have multiple senses and multiple corresponding Kanji characters. Hence Hiragana-to-Kanji conversion task is indeed a WSD task, but unlike the other WSD tasks, the word sense is *implanted* into the text so that both word sense ID and word dictionary are not necessary. We think this is a big difference to utilize word sense in the following processes.

An example of this is illustrated in Figure 1, where one Hiragana word “かける

(/kakeru/)” corresponds to three (or more, in fact) Kanji words. The meaning of the word “かける” in a sentence “私は野原をかける(/kakeru/) (I run through the field.)” means “駆ける(/kakeru/, to run)”. However, the meaning of the word “かける(/kakeru/)” in a sentence “私はコートをかける(/kakeru/) (I hang a coat.)” means “掛ける(/kakeru/, to hang)”. Hiragana words and Kanji words are always interchangeable in any context. Although the sense looks ambiguous when we use Hiragana word, it is no problem and easily solved for human since they tend to be used only in the unambiguous contexts. However, it is not the case for computer and it is also a kind of WSD problem. This conversion process is also regarded as orthographical normalization, since Hiragana and Kanji are two different orthographical variants that make natural language processing difficult.

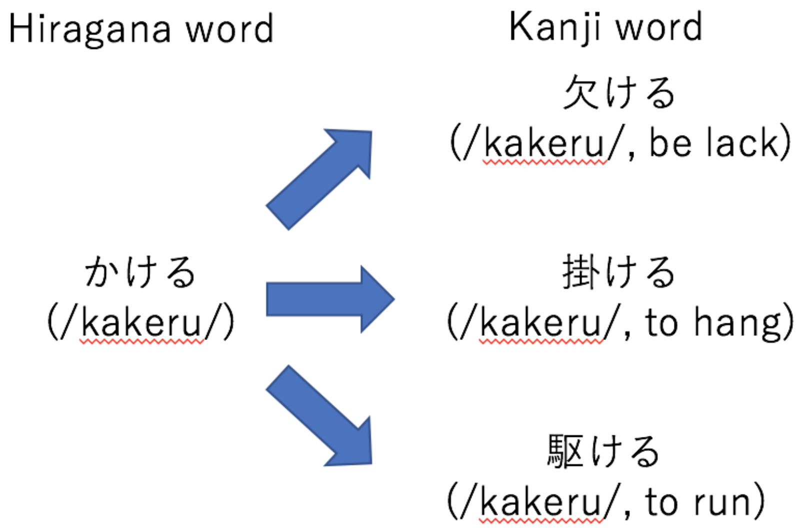

Figure 1: Example of orthographical ambiguity. /…/ in figure represents pronunciation of the word.

4. Disambiguation method

We briefly explain our conversion technique of Hiragana word into Kanji word. Figure 2 is a whole figure of the method. First, we identify a target Hiragana word which is ambiguous and should be disambiguated. We then see the context of the target word to find a clue. We also look for a dictionary to prepare possible Kanji candidates and calculate PMI score between the clue and the candidates. Finally, the highest one is judged to be the disambiguation result.

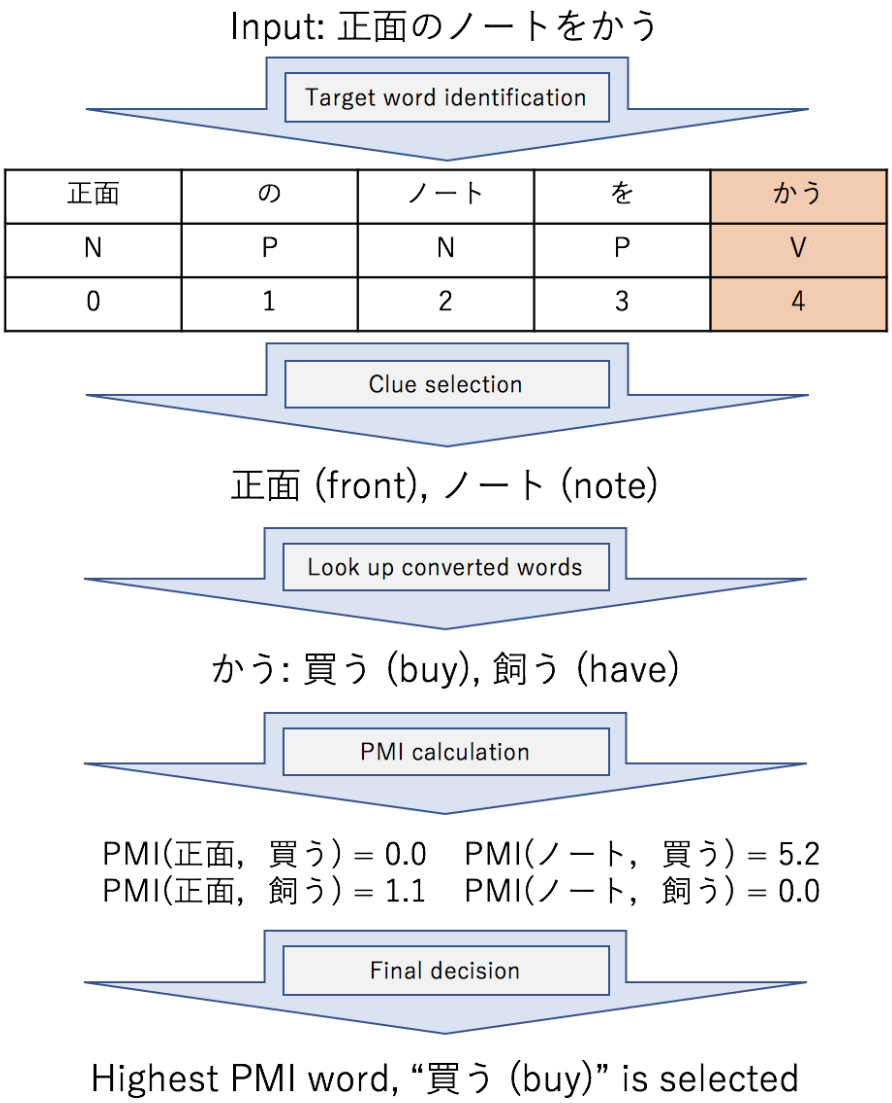

Figure 2: Outline of our Hiragana-Kanji conversion process.

We note that Hiragana words are not always ambiguous, but only some of them. When we see Hiragana words which is not ambiguous, we look for a dictionary to obtain a corresponding Kanji word, and output as it is. We also need this process in terms of orthographical normalization, although it is not WSD.

4.1. SNOWMAN: Japanese Word Analyzer

We develop a WSD tool as a module of SNOWMAN[2] (Yamamoto et al., 2015a), a Japanese word analyzer. SNOWMAN normalizes orthographical variants and each normalized word has a unique word ID. Suppose that a Hiragana word has three possible Kanji candidates. SNOWMAN initially gives different IDs to the Hiragana word and the three Kanji words in the analyzed result. But after the WSD process, SNOWMAN changes ID of the Hiragana word into either of the Kanji words. This means, unlike other WSD tasks, our WSD task can implant word sense information into word ID, thus no outer information such as a word definition dictionary is not necessary to describe word sense.

4.2. Conversion Dictionary Construction

Figure 3 illustrates collection process of word conversion candidates. We first collect all of Hiragana words in the SNOWMAN dictionary and the corresponding Kanji words. The words are limited according to part-of-speech; noun, verb, adjective and adverb is selected. We also use synonym dictionary (Yamamoto et al., 2015b) to reduce the candidates. Words with very high frequency (i.e., 500 or more) in the corpus (BCCWJ[3]) are also out of the conversion target since they are most natural as it is.

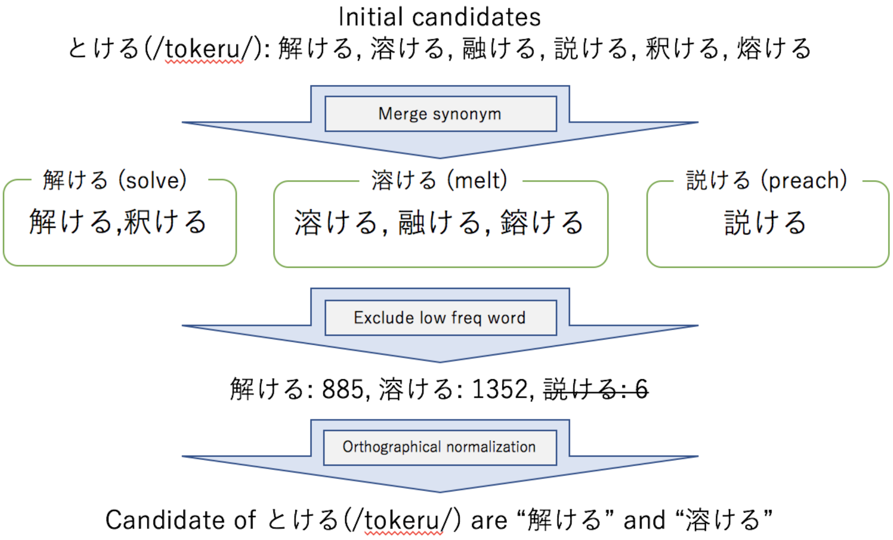

Figure 3: Collection of conversion candidates process.

[2] Japanese analyzer SNOWMAN, http://snowman.jnlp.org/english
[3] The Balanced Corpus of Contemporary Written Japanese(BCCWJ), Ver.1.1, National Institute for Japanese Language and Linguistics.
http://pj.ninjal.ac.jp/corpus center/bccwj/en/

Hiragana	Kanji #1	Kanji #2	Kanji #3
しめる（/shimeru/）	占める（to cover）	締める（to close）	湿る（to dampen）
かう（/kau/）	買う（to buy）	飼う(to feed)	
とける（/tokeru/）	解ける（to solve）	溶ける（to melt）	
なくなる（/nakunaru/）	亡くなる（to die）	無くなる（to lose）	
つかえる（/tsukaeru/）	使える（can use）	仕える（to serve）	

Table 1: Example of Hiragana-Kanji Conversion Table. Note that all Kanji words have the same pronunciation to the Hiragana word, and according to its context, each Hiragana word corresponds to either of Kanji word and its sense.

Kanji words with low frequency are excluded since most of them are regarded as noise. Table 1 shows a part of the conversion table.

Unambiguous Hiragana words are converted to Kanji words by SNOWMAN automatically because these corresponds are stored in SNOWMAN dictionary and processed in the section of orthographical normalization.

4.3. Corpus and Context

In order to construct contextual information, large corpus is required. In this work, we need word frequency information because our disambiguation method needs calculating pointwise mutual information. We use the Japanese Web N-gram Version 1[4] for calculating word frequency. It is the largest word n-gram statistics in Japanese among publicly available resources, and it includes approximately 255 billion words or 570 million word 7-grams. Although it is not “corpus” in a real sense, we use it as a corpus, by using all of word 7-grams as pseudo sentences.

4.4. PMI Calculation

We have calculated PMI score to see degree of co-occurrence of a target word and a clue (that is, a context word). *p(x, y)* is a probability that *x* and *y* appear at the same

[4] Japanese Web N-gram Version 1, GSK2007-C, Language Resources Association (GSK). http://www.gsk.or.jp/en/catalog/gsk2007-c/

time, and *p(x)* is a probability that *x* appears. PMI score between two words *(x, y)* is calculated as follows.

$$PMI(x,y) = \log \frac{p(x,y)}{p(x)p(y)} \tag{1}$$

In our case, *x* corresponds to a target word and y corresponds to a clue word. For each of a clue word we figure out PMI score, and a word with highest PMI score is selected as a clue word of a target word.

When we calculate PMI of a word, the score becomes small in general when it does not have a strong clue, and hence the accuracy tends to be lower. Therefore, we convert it when a word only has a strong clue. We judge them by PMI score and adopt a word when PMI score is higher than a threshold. This is our strong policy that accuracy is more important than coverage, in order to make a tool that is widely used by many users.

In our method, we have to determine two parameters: window size and threshold PMI.

Window size is a parameter that determines how much context to consider. An example of this is illustrated in Figure 4. We define “context” here as surrounding nouns, verbs, adjective and adverb which appears before or after words of the target word. In example of Figure 3, window size is set to 4, so we calculate PMI between candidate Kanji word and content words that appears before or after four words of the target word.

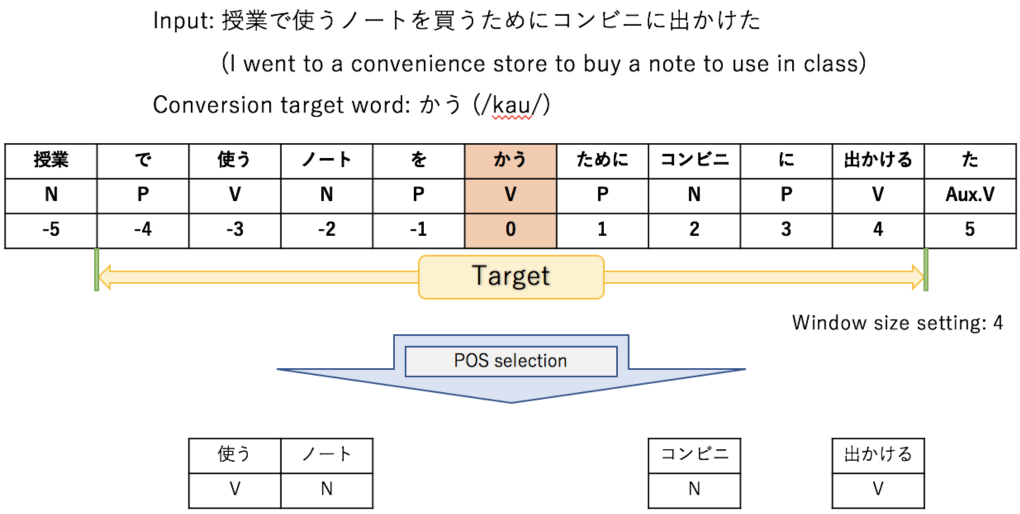

Figure 4: Definition of Context. In this example surrounding 9 words are defined as context of the target word “かう(/kau/)” and only content words are selected.

Threshold of PMI value is also important. As we mentioned earlier, we convert Hiragana to Kanji when a word only has a strong clue for high accuracy. Therefore, we have to determine a threshold PMI value that the precision of conversion is very high and coverage exceed a certain value.

5. Data set creation

Our method does not require training data, but only information of frequency for each word is needed. We created the dictionary which is explained in section 4.3. We also created test dataset to evaluate the performance of our proposed method.

We extracted sentences that contain target Kanji words from BCCWJ and convert target Kanji word to Hiragana word using SNOWMAN dictionary information. Fifty sentences per target Kanji word is extracted for test dataset. The sentences are assumed to be output word of SNOWMAN more than 15 words. There are 311 Hiragana words and 734 Kanji words, which are candidates for conversion. The average number of meanings per Hiragana is 2.36. A part of the sentences extracted as a dataset is shown in Table 2. Underlined parts in the table are target Hiragana words for conversion.

Sentence with the target Hiragana word underlined	Correct Kanji
多発性こうか症・・・私の知人の罹っていた病。 Multiple **sclerosis**, the disease which my acquaintance had suffered.	硬化 (sclerosis)
飲み水の安全性はすべて**だいちょう**菌郡の存否で判定する。 The safety of drinking water is judged by the presence or absence of **a colon** bacillus.	大腸 (a colon)
と、ルミは、少し眠そうな目をしたいしに訊いた。 Rumi asked **the doctor** who had little sleepy eyes.	医師 (the doctor)

Table 2: Example of Hiragana-Kanji conversion dataset.

6. Experiments

In this section we present two experiments. We first evaluate the relation between PMI threshold and the performance. By conducting this experiment, we can find a good threshold for Hiragana-Kanji conversion. In the second experiment, we investigate the effect of window size. In previous research, 2 to 5 words are used as surrounding word features and only a few works consider context of the whole sentence. In this experiment, we clarify the best window size for WSD. We also discuss conversion errors in section 6.3.

6.1 PMI threshold

We use the Japanese analyzer tool SNOWMAN as the morphological analyzer, and we use word frequency dictionary created in section 4.3 for PMI calculation. The window width is fixed at 20, and the change in the precision, the recall, and the F1 value when the PMI threshold is changed from 0 to 20 in increments of 1 is observed. The experimental results are shown in Figure 5.

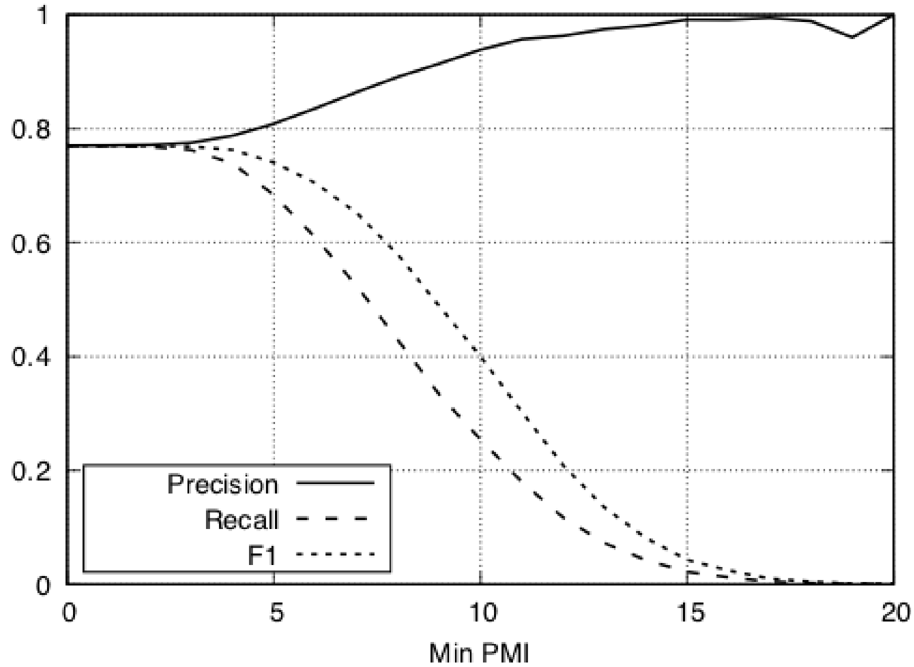

Figure 5: Accuracy in each method when threshold PMI value changes.

Referring to Figure 5, the recall rate and the precision rate are almost the same for PMI thresholds of up to 3. This is because there were no content words which PMI are 3 or less before and after the target Hiragana words. As the PMI threshold is increased from 3, the recall rate decreases. Correspondingly, the precision rate has increased when PMI threshold is increased. When the threshold PMI exceeds 10, the precision exceeds 95%. Generally speaking, by raising the threshold, a Hiragana word that has high PMI value is converted, so the precision rate would be increased. We see that the result of this experiment confirms this hypothesis. On the other hand, the recall rate has decreased with the momentum exceeding the precision rate. As a result, the F1 value is decreasing. This result indicates that setting threshold PMI between 0 and 3 is the best setting from the viewpoint of coverage and accuracy.

The precision of the conversion with the PMI threshold of 0 is approximately 78%, indicating that relatively accurate conversion is possible.

6.2 Window Size

In the window size experiment, we use the same experimental setup as the PMI threshold experiment. We set the PMI threshold to 5 and see the change of the precision, the recall and the F1 value when the window width is changed from 2 to 20 in increments of 1. The experimental results are illustrated in Figure 6.

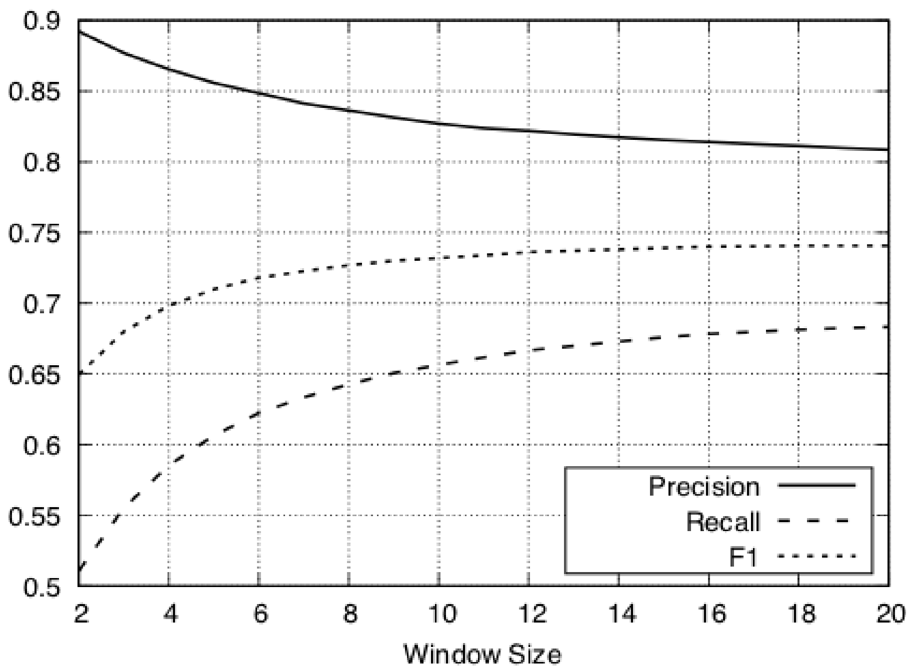

Figure 6: Accuracy in each method when window size changes.

Referring to Figure 6, the precision rate is approximately 90% when window size is 2. By increasing window size, it is possible to increase both the recall rate and the F1 value. This indicates that words that become features also appear in places far from the conversion target word. In other words, we can convert Hiragana to Kanji correctly by considering the whole sentence. However, the precision ratio is decreased by increasing the window size. In order to obtain higher accuracy, we have to consider information such as context and grammatical dependency.

6.3 Error analysis

We analyzed the error when fixing the window width to 5 and fixing the PMI threshold to 5. Example of errors are shown in the Table 3.

Sentence with Hiragana word underlined	Answer	System output
彼が書いたもので、私たちに見せたがっている<u>きじ</u>を紙の束のなかから選び出そうとしているのは、はっきりしていた It was clear that he was trying to select the **<u>articles</u>** he wrote and wanted to show to us from among a bunch of paper	記事 (article)	生地 (texture)
小学校にはいるとき、幼稚園や<u>かてい</u>で学校はどういうところか、よく教えてくれます。 When entering elementary school, he teaches well what kind of school is in kindergarten and **<u>family</u>**.	家庭 (family)	課程 (course)
<u>きたい</u>はただただ無限に薄くなっていくかのようなだけです。 It is just as if the **<u>gas</u>** just gets infinitely thin.	気体 (gas)	期待 (hope)

Table 2: Example of Hiragana-Kanji conversion dataset.

The first sentence is a sentence which was not translated correctly due to insufficient window size. In order to output the correct answer of “記事 (article)”, it is necessary to consider a word “書く (wrote)”. However, when the window size setting is 5, it cannot be regarded as a feature, so “生地 (texture)” which is most relevant to “束 (bundle)” is output as system answer. In order to solve this problem, we have to consider broader context by increasing the window size.

The second sentence is a sentence where PMI cannot correctly solve the target. In this sentence, the word “課程 (course)” co-occurs most with “学校 (school)”, so “課程 (course)” is output instead of “家庭 (family)”. In order to correctly convert such sentences, it is necessary to understand the meaning of sentences and/or using machine learning for considering broader context.

The third sentence is an example where the sentence length for conversion is insufficient. This sentence is ambiguous so that both Kanji candidates are possible for conversion. If we read “きたい (/kitai/)” as “気体 (gas)”, the meaning of this sentence will be like “It is just as if the gas just gets infinitely thin.”. However, if we read “きたい (/kitai/)” as “気体 (gas)”, the meaning of this sentence would be like “Expectations are only becoming impossible.”. Therefore, it is considered very difficult to accurately perform Kanji for this sentence.

7. Tool Implementation for Higher Precision

We exclude conversion words which cannot be obtained high conversion accuracy, since we want to deal with only high accuracy phenomena in the tool. We think low accuracy phenomena should be dealt with other conversion approaches in the future. For each word, we have conducted a preliminary experiment and see the performance. We have excluded words which has word segmentation errors, and which are difficult to judge the performance since they are non-ordinary, obsolete, and so on. As for detail, we have randomly selected each of 20 sentences which has ambiguous Hiragana words, and adopt them if it has 16 or more correct conversions (i.e., if it attains 80% accuracy).

By this process, unambiguous 52 words is selected for Hiragana-Kanji conversion. The accuracy of its conversion attains about 95%, which must be practical level. We have covered 22.4% of ambiguous Hiragana words in the BCCWJ corpus.

We implement Hiragana-Kanji conversation on a Japanese analyzer SNOWMAN. The tool is available in SNOWMAN Web Site[5] (registration are required).

8. Conclusions and Future Work

We propose a new task of Japanese WSD to use the characteristics of Japanese language: Hiragana-Kanji conversion task. We extract datasets from large corpus and evaluate our proposed method. In the experiment, we have confirmed that use of larger window size is important for F1 value because key words to solve WSD correctly may appear in places far from the conversion target. We also investigate PMI threshold; when we increase PMI threshold, the precision rate has increased but the recall rate and F1 value has decreased.

Based on the results, we have implemented a WSD tool of (a part of) Japanese Hiragana words and built into a Japanese analyzer SNOWMAN. This is a first work in Japanese language in which a WSD tool is freely available online. We have illustrated that 71 target words, which has 25% coverage among the ambiguous Hiragana words, have been correctly converted with 94% accuracy. Although the coverage of the conversion becomes partial, we intentionally select the conversion target to which very high accuracy has been expected. For viewpoint of tool development, we assert that specifying high

[5] You can use the system after the registration. http://snowman.jnlp.org/english

accuracy phenomena and releasing them as a tool is important, even though it has low coverage as it stands. As the method is very simple and easy to maintain, it is possible to continuously increasing the coverage by adding words into the target, without any bad effect to the conversion, because the process is conducted word by word and thus independent. We have been continued increasing the target words to cover more.

9. Acknowledgement

This work was supported in part by JSPS KAKENHI Grants-in-Aid for Challenging Research (Exploratory) Grant ID 17K18481.

References

Abekawa, T., Kikuko, N., Okumura, M., Yagi, Y., Totsugi, N., Sugimoto, S., and Fu, L. 2002. Development of Japanese reading system of multi-lingual environment “ASUNARO.” In Proceedings of the 3rd International Conference on Computer Assisted Systems for Teaching & Learning/Japanese, pp. 71–74.

Agirre, E. and Soroa, A. 2009, March. Personalizing PageRank for word sense disambiguation. In Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009), Athens, Greece, pp. 33–41. Association for Computational Linguistics.

Baroni, M., Dinu, G., and Kruszewski, G. 2014, June. Don’t count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Baltimore, Maryland, pp. 238–247. Association for Computational Linguistics.

Halpern, J. 2002. Lexicon-based orthographic disambiguation in CJK intelligent information retrieval. In Proceedings of the 3rd Workshop on Asian Language Resources and International Standardization - Volume 12, COLING ’02, Stroudsburg, PA, USA, pp. 1–7. Association for Computational Linguistics.

Iacobacci, I., Pilehvar, M. T., and Navigli, R. 2016, August. Embeddings for word sense disambiguation: An evaluation study. In Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Berlin, Germany, pp. 897–907. Association for Computational Linguistics.

Lopez de Lacalle, O. and Agirre, E. 2015, June. A methodology for word sense disambiguation at 90% large-scale crowdsourcing. In Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics, Denver, Colorado, pp. 61–70.

Association for Computational Linguistics.

Mihalcea, R., Chklovski, T., and Kilgarriff, A. 2004. The senseval-3 english lexical sample task. In Proceedings of Senseval-3: Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text, Barcelona, Spain, July 2004. Association for Computational Linguistics (ACL).

Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. CoRR abs/1310.4546.

Navigli, R. 2009, February. Word sense disambiguation: A survey. ACM Comput. Surv. 41 (2), 10:1–10:69.

Navigli, R. 2012. A Quick Tour of Word Sense Disambiguation, Induction and Related Approaches, pp. 115–129. Berlin, Heidelberg: Springer Berlin Heidelberg.

Okumura, M., Shirai, K., Komiya, K., and Yokono, H. 2010, July. Semeval-2010 task: Japanese WSD. In Proceedings of the 5th International Workshop on Semantic Evaluation, Uppsala, Sweden, pp. 69–74. Association for Computational Linguistics.

Petersen, T. 2006. Word Sense Disambiguation: Algorithms and Applications, Chapter 6, pp. 133–166. Springer.

Pradhan, S., Loper, E., Dligach, D., and Palmer, M. 2007, June. Semeval-2007 task-17: English lexical sample, SRL and all words. In Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007), Prague, Czech Republic, pp. 87–92. Association for Computational Linguistics.

Raganato, A., Camacho-Collados, J., and Navigli, R. 2017, April. Word sense disambiguation: A unified evaluation framework and empirical comparison. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers, Valencia, Spain, pp. 99–110. Association for Computational Linguistics.

Sanderson, M. 1994. Word Sense Disambiguation and Information Retrieval, pp. 142–151. London: Springer London.

Sugawara, H., Takamura, H., Sasano, R., and Okumura, M. 2015. Context representation with word embeddings for WSD. In International Conference of the Pacific Association for Computational Linguistics, pp. 108–119.

Wawer, A. and Mykowiecka, A. 2017, April. Supervised and unsupervised word sense disambiguation on word embedding vectors of unambigous synonyms. In Proceedings of the 1st Workshop on Sense, Concept and Entity Representations and their Applications, Valencia, Spain, pp. 120–125. Association for Computational Linguistics.

Yamamoto, K., Miyanishi, Y., Takahashi, K., Inomata, Y., Mikami, Y., and Sudo, Y. 2015. What we need is word, not morpheme; constructing word analyzer for Japanese. In

Proceedings of the International Conference on Asian Language Processing (IALP 2015), pp. 49–52.

Yamamoto, K. and Takahashi, K. 2015. Construction of Japanese semantically compatible words resource. In Proceedings of the International Conference on Asian Language Processing (IALP 2015), pp. 61–64.

Yarowsky, D. 1995. Unsupervised word sense disambiguation rivaling supervised methods. In Proceedings of the 33rd Annual Meeting on Association for Computational Linguistics, ACL ’95, Stroudsburg, PA, USA, pp. 189–196. Association for Computational Linguistics.