

PAPER

Unified Likelihood Ratio Estimation for High- to Zero-Frequency N -Grams

Masato KIKUCHI^{†a)}, Member, Kento KAWAKAMI^{††}, Nonmember, Kazuho WATANABE^{†††}, Member, Mitsuo YOSHIDA^{†††}, Nonmember, and Kyoji UMEMURA^{†††}, Member

SUMMARY Likelihood ratios (LRs), which are commonly used for probabilistic data processing, are often estimated based on the frequency counts of individual elements obtained from samples. In natural language processing, an element can be a continuous sequence of N items, called an N -gram, in which each item is a word, letter, etc. In this paper, we attempt to estimate LRs based on N -gram frequency information. A naive estimation approach that uses only N -gram frequencies is sensitive to low-frequency (rare) N -grams and not applicable to zero-frequency (unobserved) N -grams; these are known as the low- and zero-frequency problems, respectively. To address these problems, we propose a method for decomposing N -grams into item units and then applying their frequencies along with the original N -gram frequencies. Our method can obtain the estimates of unobserved N -grams by using the unit frequencies. Although using only unit frequencies ignores dependencies between items, our method takes advantage of the fact that certain items often co-occur in practice and therefore maintains their dependencies by using the relevant N -gram frequencies. We also introduce a regularization to achieve robust estimation for rare N -grams. Our experimental results demonstrate that our method is effective at solving both problems and can effectively control dependencies.

key words: likelihood ratio, the low-frequency problem, the zero-frequency problem, uLSIF

1. Introduction

Likelihood ratios (LRs) — statistical measures defined as ratios of likelihood functions — are used in numerous tasks, including relation extraction [1], [2] and data analysis [3], [4]. In practice, LRs are estimated using samples drawn from populations. In natural language processing (NLP), observed frequencies of discrete elements, e.g., word frequencies observed from language resources, are often used for estimation [5], [6]. In general, however, there will be a limited number of observed elements, and many of them will be infrequent. In this paper, we focus on the problems that the well-used estimation framework encounters in the presence of low- or zero-frequency elements.

To clarify our motivation, we suppose that the prediction task for the left context of named entities (NEs), which

Table 1 Examples of occurrence frequencies.

N -gram	Frequency				$r_{\text{MLE}}(w)$
	w	n_{de}	$c_{\text{de}}(w)$	n_{nu}	
w_A	10^7	5,000	10^4	100	20
w_B	10^7	50	10^4	1	20
w_C	10^7	50	10^4	2	40
w_D	10^7	14	10^4	0	0

are proper nouns (e.g., person names and location names), and calculate the following likelihood ratio to predict the context:

$$r(w) = \frac{p_{\text{nu}}(w)}{p_{\text{de}}(w)},$$

where w is an element sampled from a probability distribution with density $p_*(w)$, $* \in \{\text{de}, \text{nu}\}$. Let $w = a_1 a_2 \cdots a_N$ be a continuous sequence of N items (called an N -gram) in which a_k is an item such as a word or a letter at the k -th position. $p_{\text{de}}(w)$ is the probability density where w occurs at any position in the corpus, and $p_{\text{nu}}(w)$ is the probability density where w occurs on the left side of the NEs. That is, if w is the left context of the NEs, $r(w)$ is high. Otherwise, $r(w)$ is low. A naive estimation approach would be to first obtain the maximum likelihood estimate (MLE) $\hat{p}_*(w)$ of each distribution and then take the ratio of MLEs*:

$$\hat{p}_*(w) = \frac{c_*(w)}{n_*}, \quad (1)$$

$$r_{\text{MLE}}(w) = \frac{\hat{p}_{\text{nu}}(w)}{\hat{p}_{\text{de}}(w)},$$

where $c_*(w)$ is the frequency count of w sampled from the distribution with density $p_*(w)$ and n_* is the sum of $c_*(w)$ for all w . In this task, $c_{\text{de}}(w)$ is the frequency of w across the entire corpus, and $c_{\text{nu}}(w)$ is the frequency of w on the left side of the NEs. We call this type of approach “indirect estimation” because it involves probability estimation. As an example, we examine the frequencies listed in Table 1. The estimates $r_{\text{MLE}}(w)$ for each w are listed in the rightmost column of the table. Although this approach is simple, it faces the following two problems:

1) Zero-frequency problem.

$r_{\text{MLE}}(w)$ for zero-frequency

Strictly speaking, $\hat{p}_(w)$ indicates the MLE for the occurrence probability of w when the occurrence is modeled using a multinomial distribution.

Manuscript received July 16, 2020.

Manuscript revised November 16, 2020.

Manuscript publicized February 8, 2021.

[†]The author is with Department of Computer Science, Graduate School of Engineering, Nagoya Institute of Technology, Nagoya-shi, 466-8555 Japan.

^{††}The author is with LINE Corporation, Tokyo, 160-0004 Japan.

^{†††}The authors are with Department of Computer Science and Engineering, Toyohashi University of Technology, Toyohashi-shi, 441-8580 Japan.

a) E-mail: kikuchi@nitech.ac.jp

DOI: 10.1587/transfun.2020EAP1088

(unobserved) elements cannot be estimated. For example, $r_{\text{MLE}}(w_{\text{D}})$ is zero regardless of the other frequencies because $c_{\text{nu}}(w_{\text{D}})$ is zero. Furthermore, if $c_{\text{de}}(w)$ is zero, $r_{\text{MLE}}(w)$ will be infinity owing to division by zero. Owing to the nature of the language, the more items, $\{a_k\}_{k=1}^N$, that compose w (i.e., the larger N), the less likely the same w occurs in the language resource. However, in pre-learning systems, such as those of speech recognition and machine translation, long N -grams are often given as input to the trained models in the form of text or speech. Then, most are not contained in the training set. Thus, it is necessary to treat them as unobserved elements. As described, if we set the estimates of the statistical measures for all unobserved elements to zeros or infinities, the systems often do not work well. Therefore, it is important to properly estimate the measures, even for the zero-frequency elements. Additionally, a real-time processing system must handle such elements without expanding the training set.

- 2) **Low-frequency problem.** $r_{\text{MLE}}(w)$ estimated for low-frequency (rare) elements can have very large estimation errors. For example, the estimation uncertainty for w_{B} is larger than that for w_{A} because the frequencies of w_{B} are lower than those of w_{A} ; that is, $c_{\text{de}}(w_{\text{A}}) > c_{\text{de}}(w_{\text{B}})$ and $c_{\text{nu}}(w_{\text{A}}) > c_{\text{nu}}(w_{\text{B}})$. However, $r_{\text{MLE}}(w_{\text{B}})$ is high and equal to $r_{\text{MLE}}(w_{\text{A}})$ (both 20). In particular, $c_{\text{nu}}(w_{\text{B}}) = 1$ may be a coincidence. Therefore, $r_{\text{MLE}}(w)$ should be adjusted depending on the uncertainty. Moreover, $r_{\text{MLE}}(w)$ is sensitive to the number of observations; $c_{\text{nu}}(w_{\text{C}})$ has only one more observation than $c_{\text{nu}}(w_{\text{B}})$, but $r_{\text{MLE}}(w_{\text{C}})$ is twice $r_{\text{MLE}}(w_{\text{B}})$ (40 and 20, respectively). In such cases, $r_{\text{MLE}}(w)$ could be more robust. Most N -grams are infrequent, and, as N increases, the infrequent N -grams tend to increase further. Therefore, to handle them, it is essential to address the low-frequency problem.

A method for mitigating the low-frequency problem was proposed by [7]. This method applies *unconstrained Least-Squares Importance Fitting* (uLSIF) [8], which is used to estimate $r(w)$ without estimates of $p_*(w)$ by solving a least-squares minimization and is therefore called “direct estimation.” The resulting model in [7] underestimates $r(w)$ properly by using a regularization parameter introduced in the minimization process, thereby making the estimator robust (see Sect. 3.3). To the best of our knowledge, no other method has been found to mitigate the low-frequency problem in LR estimation. Unfortunately, this method does not address the zero-frequency problem. In fact, in cases in which many rare elements are observed, it is expected that many unobserved elements will also be present. Thus, the low- and zero-frequency problems should both be addressed at the same time if possible.

This paper presents a method for mitigating both problems based on the following approaches:

Approach for Problem 1. One possible approach for estimating $r(w)$ is to “itemize” or decompose w into item units

of $\{a_k\}_{k=1}^N$ and then estimate $r_{\text{item}}(w)$ using their frequencies. Itemizing w is commonly performed in probability estimation and is here applied to LR estimation. If each a_k is observed separately, $r_{\text{item}}(w)$ can be estimated even for unobserved w . For example, we suppose that the entity type is person name and that $w = \text{“to Mr.”}$ ($a_1 = \text{“to”}$ and $a_2 = \text{“Mr.”}$). If the two words a_1 and a_2 occur separately, we can estimate $r_{\text{item}}(w)$, even if w does not occur. However, the standard estimation of $r_{\text{item}}(w)$ assumes that all items occur independently, and the dependencies between items are ignored. Here, we give another example. We suppose that the entity type is location name, and that $w = \text{“Hospital in.”}$ In this example, the dependency of a_1 and a_2 is important. This is because the preposition “in” occurs in various contexts, and the noun “Hospital” at the left of “in” is the key to predicting whether w is the left context of location names. To utilize dependency, we introduce the term $t_{\text{d}}(w)$, which is estimated using the frequencies of w . In this manner, our method can incorporate dependency and successfully treat observed and unobserved w .

Approach for Problem 2. An important challenge is determining the term $t_{\text{d}}(w)$. For w with $c_{\text{de}}(w) = c_{\text{nu}}(w) = 0$, we estimate $r_{\text{item}}(w)$ and then set $t_{\text{d}}(w)$ to zero. For rare w , we can use an approach similar to that used in [7]. Both $r_{\text{item}}(w)$ and $t_{\text{d}}(w)$ are estimated using regularization parameters, which can mitigate the low-frequency problem. More importantly, these parameters are shown to also be capable of controlling the dependency between items. We explain this advantage by reusing the examples used in the approach for problem 1. If $w = \text{“to Mr.”}$ we can imagine that a person name will be on the right side of $a_2 = \text{“Mr.”}$ regardless of a_1 . Therefore, the required dependency between a_1 and a_2 is inferred to be weak. In contrast, in the case of $w = \text{“Hospital in.”}$ the sequence of two words is important for context prediction, and the required dependency is inferred to be strong. In such examples, where the required dependencies are different, the proposed method adjusts the strength of the dependency and incorporates the appropriate dependency into the estimate. Consequently, our method can obtain an informative estimate, even for rare w .

Experimental results demonstrate that our method is effective at addressing both problems and further enables the control of dependency between items, which leads to performance improvement.

2. Related Work

Naive LR estimation approaches face both the low- and the zero-frequency problems. As a first step toward finding a mitigation approach, in this section we review several methods that have been developed to address these problems for non-LR statistical measures.

Studies addressing the use of estimation uncertainty to mitigate the low-frequency problem include Johnson [9], who used confidence interval bounds for pointwise mutual information (PMI) estimation. Such bounds can reflect the estimation variance because the interval varies with it. Pantel

and Ravichandran [10] proposed a penalized estimator based on frequencies. Because PMI includes LR in its definition, these two measures are closely related. As a conditional probability estimation method, Rudin et al. [11] proposed an approach involving obtaining the MLE and then adding a positive constant to the denominator. As conditional probability can be defined as a probability ratio, it can also be considered a type of LR. Under this definition, the method in [11] is formally equivalent to the existing method [7] described in Sect. 1. Although the approaches described above all mitigate the low-frequency problem, none address the zero-frequency problem.

Many probability estimation approaches have attempted to address low- and zero-frequency problems through the use of smoothing techniques, which discount probability estimates for observed elements and distribute them for unobserved elements. Here, we focus on smoothing for a probability $p(w)$ of an N -gram w [12]–[18]. This type of smoothing uses lower order N -gram (partial elements of w) frequencies in addition to original N -gram frequencies. Even $p(w)$ for unobserved w can be estimated by using them alternatively. Although w -decomposition is related to our approach, these techniques are not directly applicable to LR estimation. Because the LR is given as a ratio of probability distributions, representing it as a ratio of smoothed estimates appears to be a natural approach; however, we would expect that this approach would still face the low-frequency problem as a result of indirect estimation and, therefore, we must verify it experimentally.

Other approaches for mitigating the two problems have been used to estimate odds ratios, which are represented as ratios of probability distributions, for correcting individual probability estimators. In [19] and [20], unbiased and Bayes estimators were used for cell probability estimation. These estimators reduce the estimation bias for rare elements and can avoid the zero-frequency problem in cases in which all observed elements are infrequent. However, we treat elements ranging from high- to zero-frequency in LR estimation. Moreover, approaches involving ratios of corrected estimates have been found to face the low-frequency problem when used for LR estimation [7], which leads us to suspect that the smoothing approach described above also faces the low-frequency problem.

Some classifiers decompose item sequences. For instance, naive Bayes classifiers are based on a strong (naive) independence assumption and Bayes' theorem. The classifiers approximate the probability $p(D|C)$ that a document $D = a_1 a_2 \cdots a_N$ belongs to a class C as follows:

$$p(D|C) \approx \prod_i^N p(a_i|C),$$

where a_i is an item, e.g., a word or letter, in D and $p(a_i|C)$ is estimated using smoothing techniques. In this approach, the low- and zero-frequency problems are mitigated. The approximation is based on a between-item conditional independence assumption; i.e., $p(a_i|C, a_j) = p(a_i|C)$, where

$i \neq j$. This assumption, however, often does not hold in real-world data, and extended methods to alleviate it have been proposed [21]–[26]. Our approach for itemizing N -grams and incorporating dependency between items was inspired by this framework.

3. Direct Estimation of Likelihood Ratio: Existing Methods

Several direct LR estimation methods, including the use of kernel mean matching [27], logistic regression [28], Kullback-Leibler divergence minimization [29], and least-squares [8] approaches, have been proposed. These methods are used to estimate LRs represented by continuous distributions. A method for estimating them represented by discrete distributions, was proposed by [7]. This method is a variant of the least-squares approach called *unconstrained Least-Squares Importance Fitting* (uLSIF). Our proposed method, described in Sect. 4, follows the uLSIF framework, although we use the variant to treat discrete elements. Thus, in this section we explain the problem definition and then describe the original uLSIF and its variant.

3.1 Problem Definition

We define the overall data domain as D . In Sects. 3.2 and 3.3, the domains are defined as $D \subset \mathbb{R}^d$ and $D \subset \mathbb{U}$, respectively. \mathbb{R}^d is a real d -dimensional space. \mathbb{U} is a set with v discrete elements and is also called a finite alphabet in information theory. We examine a case in which we obtain an independent and identically distributed (i.i.d.) sample $\{x_i^{\text{de}}\}_{i=1}^{n_{\text{de}}}$ from a probability distribution with density $p_{\text{de}}(x)$ and an i.i.d. sample $\{x_j^{\text{nu}}\}_{j=1}^{n_{\text{nu}}}$ from a probability distribution with density $p_{\text{nu}}(x)$:

$$\{x_i^{\text{de}}\}_{i=1}^{n_{\text{de}}} \stackrel{\text{i.i.d.}}{\sim} p_{\text{de}}(x), \quad \{x_j^{\text{nu}}\}_{j=1}^{n_{\text{nu}}} \stackrel{\text{i.i.d.}}{\sim} p_{\text{nu}}(x).$$

Following previous studies, we assume that $p_{\text{de}}(x)$ satisfies the condition

$$p_{\text{de}}(x) > 0 \quad \text{for all } x \in D,$$

which makes it possible to define an LR for all x . We then attempt to estimate the following LR $r(x)$ directly from two samples, $\{x_i^{\text{de}}\}_{i=1}^{n_{\text{de}}}$ and $\{x_j^{\text{nu}}\}_{j=1}^{n_{\text{nu}}}$:

$$r(x) = \frac{p_{\text{nu}}(x)}{p_{\text{de}}(x)}.$$

3.2 Unconstrained Least-Squares Importance Fitting (uLSIF)

In uLSIF [8], $r(x)$ is modeled by the following linear model:

$$\hat{r}(x) = \sum_{l=1}^b \beta_l \varphi_l(x), \quad (2)$$

where $\boldsymbol{\beta} = (\beta_1, \beta_2, \dots, \beta_b)^\top$ are the parameters learned

from data samples and $\{\varphi_l\}_{l=1}^b$ are the basis functions, which are always non-negative. Note that b and $\{\varphi_l\}_{l=1}^b$ are independent of the samples $\{x_i^{\text{de}}\}_{i=1}^{n_{\text{de}}}$ and $\{x_i^{\text{nu}}\}_{i=1}^{n_{\text{nu}}}$. The goal of uLSIF is to solve the optimization problem[†]:

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^b} \left[\frac{1}{2} \boldsymbol{\beta}^T \widehat{\mathbf{H}} \boldsymbol{\beta} - \widehat{\mathbf{h}}^T \boldsymbol{\beta} + \frac{\lambda}{2} \boldsymbol{\beta}^T \boldsymbol{\beta} \right], \quad (3)$$

where $\widehat{\mathbf{H}}$ is a $b \times b$ matrix in which the (l, l') -th element is given by

$$\widehat{H}_{l,l'} = \frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_l(x_i^{\text{de}}) \varphi_{l'}(x_i^{\text{de}}), \quad (4)$$

and $\widehat{\mathbf{h}}$ is a b -dimensional vector in which the l -th element is given by

$$\widehat{h}_l = \frac{1}{n_{\text{nu}}} \sum_{j=1}^{n_{\text{nu}}} \varphi_l(x_j^{\text{nu}}). \quad (5)$$

In Eq. (3), a penalty term $\frac{\lambda}{2} \boldsymbol{\beta}^T \boldsymbol{\beta}$ is introduced for regularization to $\boldsymbol{\beta}$; λ (≥ 0) is a regularization parameter and $\boldsymbol{\beta}^T \boldsymbol{\beta} / 2$ is an ℓ_2 -regularizer. As the objective function in Eq. (3) is a convex quadratic function, its solutions can be written as

$$\widetilde{\boldsymbol{\beta}}(\lambda) = (\widehat{\mathbf{H}} + \lambda \mathbf{1}_b)^{-1} \widehat{\mathbf{h}},$$

where $\mathbf{1}_b$ is a b -dimensional vector containing all ones. As the optimization problem does not have non-negativity constraints on the parameters, some of them can be negative. Therefore, $\widetilde{\boldsymbol{\beta}}(\lambda)$ are modified as follows by taking into account the non-negativity of $r(x)$:

$$\widehat{\boldsymbol{\beta}}(\lambda) = \max(\mathbf{0}_b, \widetilde{\boldsymbol{\beta}}(\lambda)),$$

where the ‘‘max’’ operation is applied in the element-wise manner and $\mathbf{0}_b$ is a b -dimensional vector containing all zeros. $\widehat{\boldsymbol{\beta}}(\lambda)$ are the final solutions of Eq. (3).

uLSIF uses the structure of a sample space to model $r(x)$. For this reason, the selection of basis functions is important. Moreover, the ℓ_2 -regularizer gives prior information that $\boldsymbol{\beta}$ are uniform under the condition that no samples are observed from the space.

3.3 A Variant of uLSIF for Discrete Distributions

The primary alteration to the original uLSIF involves a change of basis functions. In natural language processing (NLP), observed samples often comprise discrete elements such as words and letters rather than continuous elements. uLSIF uses Gaussian kernels in its basis functions, but these cannot take into account the discreteness of the sample space. To apply uLSIF in discrete cases, Kikuchi et al. [7] proposed

the use of the basis functions $\{\varphi_l\}_{l=1}^v$:

$$\varphi_l(x) = \begin{cases} 1 & x = x_{(l)}, \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

In the above, x is a discrete element, v is the number of possible element types, and l is an index specifying the specific type from among v types; that is, $x_{(l)}$ indicates the l -th element among v different elements. The functions can also be regarded as Kronecker delta functions. In addition, $r(x_{(m)})$, $1 \leq m \leq v$, is modeled as Eq. (2):

$$\widehat{r}_{\mathbf{K}}(x_{(m)}) = \sum_{l=1}^v \beta_l \varphi_l(x_{(m)}) = \beta_m.$$

As the basis function $\varphi_l(x)$ becomes zero unless x is identical to $x_{(l)}$, the $\widehat{H}_{l,l'}$ and \widehat{h}_l corresponding to Eqs. (4) and (5), respectively, are given by

$$\begin{aligned} \widehat{H}_{l,l'} &= \begin{cases} \frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_l(x_i^{\text{de}}) \varphi_{l'}(x_i^{\text{de}}) & (l = l') \\ 0 & (l \neq l') \end{cases} \\ &= \frac{1}{n_{\text{de}}} c_{\text{de}}(x_{(l)}) \quad (l = l'), \\ \widehat{h}_l &= \frac{1}{n_{\text{nu}}} \sum_{j=1}^{n_{\text{nu}}} \varphi_l(x_j^{\text{nu}}) = \frac{1}{n_{\text{nu}}} c_{\text{nu}}(x_{(l)}), \end{aligned}$$

where $c_*(x_{(l)})$, $* \in \{\text{de}, \text{nu}\}$ is the frequency of $x_{(l)}$ sampled from a probability distribution with density $p_*(x_{(l)})$ and $\widehat{\mathbf{H}}$ is a $v \times v$ diagonal matrix. From the above, $\widehat{r}_{\mathbf{K}}(x_{(m)})$ can be calculated as

$$\begin{aligned} \widehat{r}_{\mathbf{K}}(x_{(m)}) &= \widetilde{\boldsymbol{\beta}}_m(\lambda) \\ &= (\widehat{H}_{m,m} + \lambda)^{-1} \widehat{h}_m \\ &= \left(\frac{1}{n_{\text{de}}} c_{\text{de}}(x_{(m)}) + \lambda \right)^{-1} \frac{1}{n_{\text{nu}}} c_{\text{nu}}(x_{(m)}), \quad (7) \end{aligned}$$

which is a non-negative value and therefore becomes the solution^{††}. The regularization parameter λ results in an underestimation of $r(x_{(m)})$. If λ is zero, $\widehat{r}_{\mathbf{K}}(x_{(m)})$ is $r_{\text{MLE}}(w) = \frac{\widehat{p}_{\text{nu}}(x_{(m)})}{\widehat{p}_{\text{de}}(x_{(m)})}$, which is equivalent to Eq. (1) with $\widehat{p}_*(x_{(m)})$ representing the MLE[†] of a probability distribution with density $p_*(x_{(m)})$. Eq. (7) adjusts only the denominator $\frac{1}{n_{\text{de}}} c_{\text{de}}(x_{(m)})$ of $r_{\text{MLE}}(w)$ using λ .

3.4 Regularization Effect on $\widehat{r}_{\mathbf{K}}(x_{(m)})$

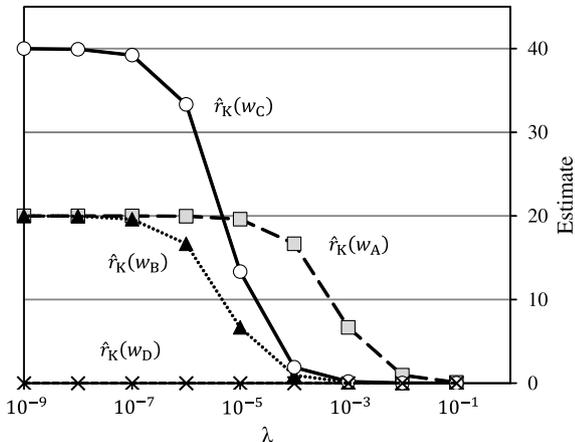
As we use the estimator $\widehat{r}_{\mathbf{K}}(x_{(m)})$ in our proposed method, understanding its function will be important. Accordingly, in this section we explain $\widehat{r}_{\mathbf{K}}(x_{(m)})$ from the perspective of regularization. To do so, we assume that each discrete element x is an N -gram w . As in Sect. 1, we calculate the likelihood ratio, $r(w) = \frac{p_{\text{nu}}(w)}{p_{\text{de}}(w)}$, to predict the left context of

[†]The objective function in Eq. (3) is derived from the square loss of $\widehat{r}(x)$ and $r(x)$. See Appendix A for a derivation of the first two terms in this function.

^{††}As the original paper [7] was written in Japanese, in Sect. 3.4 we present an estimation example for further clarification.

Table 2 Frequency examples (identical to Table 1).

N -gram w	Frequency				$r_{\text{MLE}}(w)$
	n_{de}	$c_{\text{de}}(w)$	n_{nu}	$c_{\text{nu}}(w)$	
w_A	10^7	5,000	10^4	100	20
w_B	10^7	50	10^4	1	20
w_C	10^7	50	10^4	2	40
w_D	10^7	14	10^4	0	0

**Fig. 1** Regularization effect on $\hat{r}_K(w)$ given in Eq. (7). The horizontal and vertical axes of the graph represent the regularization parameter λ and estimate of $\hat{r}_K(w)$, respectively.

NEs. $p_{\text{de}}(w)$ is the probability density where w occurs at any position in the corpus, and $p_{\text{nu}}(w)$ is the probability density where w occurs on the left side of the NEs. That is, if w is the left context of the NEs, $r(w)$ is high. Otherwise, $r(w)$ is low. We also assume that we observe the frequencies shown in Table 2 (identical to Table 1) and use them to estimate $r(w)$ for each w . Here, $c_{\text{de}}(w)$ is the frequency of w in the entire corpus, and $c_{\text{nu}}(w)$ is the frequency of w on the left side of NEs. $\hat{r}_K(w)$ with $\lambda = 0$ is equal to $r_{\text{MLE}}(w)$ as defined by Eq. (1). Figure 1 shows the estimated $\hat{r}_K(w)$ as the regularization parameter λ is varied over the range $10^{-9}, 10^{-8}, \dots, 10^{-1}$.

At $\lambda = 10^{-9}$, each $\hat{r}_K(w)$ is nearly equal to $r_{\text{MLE}}(w)$ in Table 2 as a result of the weak regularization. At nearly $\lambda = 10^{-7}$, $\hat{r}_K(w_B)$ and $\hat{r}_K(w_C)$ both begin to decrease, while $\hat{r}_K(w_A)$ remains close to 20. The frequencies of w_B and w_C are lower than those of w_A , indicating that the regularization effect acts more strongly on $\hat{r}_K(w_B)$ and $\hat{r}_K(w_C)$ even when $r_{\text{MLE}}(w_A) = r_{\text{MLE}}(w_B)$. As λ increases further, $\hat{r}_K(w_A)$ begins to decrease, and all estimators go to zero, suggesting that, for all w , $\hat{r}_K(w)$ becomes robust for frequencies with large λ . Therefore, if we use $\hat{r}_K(w)$ in the prediction task of the left N -grams of NEs, we can preferentially predict high-frequency left N -grams, and after that we can predict low-frequency left N -grams. Note that $\hat{r}_K(w_D)$ is always zero. This method addresses the low-frequency problem but not the zero-frequency problem. When dealing with high-order (long) N -grams, there are many w s having $c_{\text{de}}(w) = 0$ or $c_{\text{nu}}(w) = 0$. Therefore, in the next section, we propose a

new estimator that can address the zero-frequency problem by also using the composed items of w .

4. Proposed Likelihood Ratio Estimator

Our proposed method follows the uLSIF framework. Our framework models LR as a linear sum of basis functions and their weighting parameters, which are optimized using least-squares minimization. Note that we select the basis functions applicable to discrete elements, as well as the method described in Sect. 3.3. We briefly explain our estimation procedure as follows. For an “itemizable” element w , $r(w)$ is first roughly approximated as $r_{\text{item}}(w)$ using the component frequencies of w under the assumption of independence between components. Then, the term $t_d(w)$ based on the frequency of w is added to $\hat{r}_{\text{item}}(w)$ for independence alleviation. Finally, an estimate of $t_d(w)$ for which the square loss of $\hat{r}_{\text{item}}(w) + \hat{t}_d(w)$ and $r(w)$ is minimized is made. Below, we describe our method in detail.

4.1 Overall Framework

We examine a case in which we observe N -grams as itemizable elements. Each N -gram w is a continuous sequence of N items $\{a_k\}_{k=1}^N$:

$$w = a_1 a_2 \cdots a_N,$$

where a_k is a discrete item at the k -th position in w . The items can be words, letters, etc. In particular, for $N = 1, 2, 3$, w is called a unigram, bigram, and trigram, respectively. Before going further, we define the two operators Ψ_k and $\Psi_{k(l)}$ for simplicity. These generate item patterns from w .

Definition 1 (Operator Ψ_k): An operator Ψ_k replaces all items except a_k with a wildcard symbol \bullet . This symbol allows any item type match at the k -th position. If w is $a_1 a_2 a_3$, $\Psi_k w$ ($k = 1, 2, 3$) is $\Psi_1 w = a_1 \bullet \bullet$, $\Psi_2 w = \bullet a_2 \bullet$, and $\Psi_3 w = \bullet \bullet a_3$.

Definition 2 (Operator $\Psi_{k(l)}$): An operator $\Psi_{k(l)}$ replaces all items except a_k with \bullet and a_k with $a_{(l)}$, $1 \leq l \leq v$, where l is an index specifying an item type and v is the number of item types that can exist; that is, $a_{(l)}$ indicates the l -th item among all different v items. If w is $a_1 a_2 a_3$, $\Psi_{k(l)} w$ ($k = 1, 2, 3$) is $\Psi_{1(l)} w = a_{(l)} \bullet \bullet$, $\Psi_{2(l)} w = \bullet a_{(l)} \bullet$, and $\Psi_{3(l)} w = \bullet \bullet a_{(l)}$.

Our initial approach is to assume the statistical independence of items $\{a_k\}_{k=1}^N$ in w and to itemize w into item units. Based on this, $r(w)$ can be represented by $r_{\text{item}}(w)$, which is the product of the LRs of units $\{\Psi_k w\}_{k=1}^N$:

$$\begin{aligned} r_{\text{item}}(w) &= \prod_{k=1}^N r(\Psi_k w) \\ &= r(a_1 \bullet \bullet \cdots \bullet) r(\bullet a_2 \bullet \cdots \bullet) \\ &\quad \cdots r(\bullet \bullet \cdots \bullet a_N). \end{aligned} \quad (8)$$

All of $\Psi_k w$ match many N -grams including w . Therefore, if $r(w)$ can be defined, $r(\Psi_k w)$ can also be defined. By itemizing each w —even those sequences not observed in samples—we can estimate $r_{\text{item}}(w)$ based on individual unit frequencies. However, the independence assumption is inaccurate and might cause a significant estimation error. Therefore, we add a dependency term $t_d(w)$ to Eq. (8) to incorporate the dependency between items. Our overall model is then:

$$r_{\text{ours}}(w) = r_{\text{item}}(w) + t_d(w), \quad (9)$$

which is estimated in two steps. We first estimate $r_{\text{item}}(w)$ using each unit frequency $c_*(\Psi_k w)$, $* \in \{\text{de}, \text{nu}\}$, and we then estimate $t_d(w)$ using each frequency $c_*(w)$. $t_d(w)$ is defined for each N -gram w and estimated so that the squared loss between our model $r_{\text{ours}}(w)$ and the ideal $r(w)$ is minimized. When at least $c_{\text{de}}(w) > 0$ or $c_{\text{nu}}(w) > 0$ and the estimate of $r_{\text{item}}(w) + t_d(w)$ is non-negative, we can obtain the exact solution that minimizes the squared error.

4.2 Estimation of Itemized LR $r_{\text{item}}(w)$

Given an N -gram w , we estimate $r(\Psi_k w)$ in Eq. (8) by applying the method described in Sect. 3.3. We define each of the basis functions $\{\varphi_l(\Psi_k w)\}_{l=1}^v$ by

$$\varphi_l(\Psi_k w) = \begin{cases} 1 & \Psi_k w = \Psi_{k(l)} w, \\ 0 & \text{otherwise,} \end{cases}$$

where l is the index specifying an item type that is not replaced by the wildcard symbol \bullet and v is the number of item types that can exist. If $\Psi_k w$ is identical to $\Psi_{k(m)} w$, $1 \leq m \leq v$, that is, if a_k is identical to $a_{(m)}$, $\widehat{r}_K(\Psi_{k(m)} w)$ can be calculated as

$$\begin{aligned} \widehat{r}_K(\Psi_{k(m)} w) &= \left(\frac{1}{n_{\text{de}}} c_{\text{de}}(\Psi_{k(m)} w) + \lambda_{\text{item}} \right)^{-1} \\ &\quad \times \frac{1}{n_{\text{nu}}} c_{\text{nu}}(\Psi_{k(m)} w), \end{aligned}$$

where $\lambda_{\text{item}} (\geq 0)$ is a regularization parameter and n_{de} and n_{nu} are the total frequencies of all N -grams observed from samples $\{w_i^{\text{de}}\}_{i=1}^{n_{\text{de}}}$ and $\{w_j^{\text{nu}}\}_{j=1}^{n_{\text{nu}}}$, respectively. From the above, we estimate $r_{\text{item}}(w)$ as the product of $\{\widehat{r}_K(\Psi_k w)\}_{k=1}^N$:

$$\widehat{r}_{\text{item}}(w) = \prod_{k=1}^N \widehat{r}_K(\Psi_k w). \quad (10)$$

In our experiments, n_{de} and n_{nu} are the total frequencies of N -grams in the training set and on the left side of the NEs, respectively. From these definitions, $n_{\text{de}} \gg n_{\text{nu}}$ holds. Additionally, many N -grams comprise low-frequency items. Considering these, $\widehat{p}_{\text{de}}(\Psi_k w) \ll \widehat{p}_{\text{nu}}(\Psi_k w)$ holds for most $\Psi_k w$. The value of $\widehat{p}_{\text{de}}(\Psi_k w)$ is strongly reflected in the estimate of $r(\Psi_k w)$. Therefore, we correct the denominator of $r(\Psi_k w)$. $\widehat{r}_{\text{item}}(w)$ has N regularization parameters. However, searching for the optimal combination of

all λ_{item} is computationally expensive. On the other hand, $\widehat{p}_{\text{de}}(\Psi_k w)$ is the occurrence probability of $\Psi_k w$ in the training set. Here, if the words a_k and $a_{k'}$ are identical, then $\widehat{p}_{\text{de}}(\Psi_k w) \approx \widehat{p}_{\text{de}}(\Psi_{k'} w)$, $k \neq k'$. That is, $\widehat{p}_{\text{de}}(\Psi_k w)$ is almost determined by the word type of a_k , not k . For these reasons, we consider them as common parameters and set the parameters all to the same value.

4.3 Estimation of Dependency Term $t_d(w)$.

Given an N -gram w and $\widehat{r}_{\text{item}}(w)$, we estimate $t_d(w)$ in Eq. (9). We model $t_d(w)$ as the following linear model:

$$\widehat{t}_d(w) = \sum_{l=1}^v \beta_l \varphi_l(w),$$

where $\boldsymbol{\beta} = (\beta_1, \beta_2, \dots, \beta_v)^T$ are the parameters learned from data samples and $\{\varphi_l(w)\}_{l=1}^v$ are the basis functions, which are always non-negative. Similar to Eq. (6), we define each of the basis functions for w as

$$\varphi_l(w) = \begin{cases} 1 & w = w_{(l)}, \\ 0 & \text{otherwise,} \end{cases}$$

where l is the index specifying a particular N -gram type from among all N -gram types; that is, $w_{(l)}$ indicates the l -th N -gram from among all N -grams. Note that v is not defined as in Sect. 4.2 but here represents the number of possible N -gram types. For identical w with $w_{(m)}$, $1 \leq m \leq v$, our model $r_{\text{ours}}(w_{(m)})$ is given as follows:

$$\begin{aligned} r_{\text{ours}}(w_{(m)}) &= \widehat{r}_{\text{item}}(w_{(m)}) + \widehat{t}_d(w_{(m)}) \\ &= \widehat{r}_{\text{item}}(w_{(m)}) + \beta_m. \end{aligned} \quad (11)$$

$r_{\text{item}}(w_{(m)})$ is replaced by $\widehat{r}_{\text{item}}(w_{(m)})$, as shown in Eq. (10), and treated as a constant. To determine $\boldsymbol{\beta}$ from the data samples, we solve the following problem:

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^v} \left[\widehat{J}(\boldsymbol{\beta}) + \frac{\lambda_d}{2} \boldsymbol{\beta}^T \boldsymbol{\beta} \right], \quad (12)$$

where $\lambda_d (\geq 0)$ is a regularization parameter and $\widehat{J}(\boldsymbol{\beta})$ are the squared-loss-based terms[†]:

$$\begin{aligned} \widehat{J}(\boldsymbol{\beta}) &= \frac{1}{2} \sum_{l=1}^v \beta_l^2 \left(\frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_l(w_i^{\text{de}}) \right) \\ &\quad + \widehat{r}_{\text{item}}(w_{(m)}) \sum_{l=1}^v \beta_l \left(\frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_l(w_i^{\text{de}}) \right) \\ &\quad - \sum_{l=1}^v \beta_l \left(\frac{1}{n_{\text{nu}}} \sum_{j=1}^{n_{\text{nu}}} \varphi_l(w_j^{\text{nu}}) \right). \end{aligned} \quad (13)$$

As the objective function in Eq. (12) is a convex quadratic function, we can differentiate it with respect to β_m and set it to zero:

[†] See Appendix B for the derivation of Eq. (13).

Table 3 Examples of original and itemized frequencies.

Bigram	Original Freq.				$r_{\text{MLE}}(w)$	Itemized Freq.				$\hat{r}_{\text{item}}(w)$ with $\lambda_{\text{item}} = 0$
	$w = a_1 a_2$	n_{de}	$c_{\text{de}}(w)$	n_{nu}		$c_{\text{nu}}(w)$	$c_{\text{de}}(a_1 \bullet)$	$c_{\text{nu}}(a_1 \bullet)$	$c_{\text{de}}(\bullet a_2)$	
w_A	10^7	5,000	10^4	100	20	8,000	1,000	5,000	500	12,500
w_B	10^7	50	10^4	1	20	240	30	150	15	12,500
w_C	10^7	50	10^4	2	40	80	10	50	5	12,500
w_D	10^7	14	10^4	0	0	80	10	50	5	12,500

$$\begin{aligned}
& \frac{\partial}{\partial \beta_m} \left(\hat{J}(\beta) + \frac{\lambda_d}{2} \beta^T \beta \right) \\
&= \beta_m \left(\frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_m(w_i^{\text{de}}) + \lambda_d \right) + \\
& \hat{r}_{\text{item}}(w(m)) \left(\frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_m(w_i^{\text{de}}) \right) - \frac{1}{n_{\text{nu}}} \sum_{j=1}^{n_{\text{nu}}} \varphi_m(w_j^{\text{nu}}) \\
&= 0.
\end{aligned}$$

Solving the above for β_m , we obtain

$$\begin{aligned}
& \tilde{\beta}_m(\lambda_d) \\
&= \left(\frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_m(w_i^{\text{de}}) + \lambda_d \right)^{-1} \times \\
& \left\{ \frac{1}{n_{\text{nu}}} \sum_{j=1}^{n_{\text{nu}}} \varphi_m(w_j^{\text{nu}}) - \hat{r}_{\text{item}}(w(m)) \left(\frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_m(w_i^{\text{de}}) \right) \right\} \\
&= \left(\frac{c_{\text{de}}(w(m))}{n_{\text{de}}} + \lambda_d \right)^{-1} \times \\
& \left\{ \frac{c_{\text{nu}}(w(m))}{n_{\text{nu}}} - \hat{r}_{\text{item}}(w(m)) \left(\frac{c_{\text{de}}(w(m))}{n_{\text{de}}} \right) \right\}. \quad (14)
\end{aligned}$$

This is the solution of the optimization problem shown in Eq. (12)[†]. Substituting $\tilde{\beta}_m(\lambda_d)$ into Eq. (11), our model can be calculated as

$$\tilde{r}_{\text{ours}}(w(m)) = \hat{r}_{\text{item}}(w(m)) + \tilde{\beta}_m(\lambda_d).$$

As this can be negative because $\tilde{\beta}_m(\lambda_d)$ includes a subtraction term, we modify $\tilde{r}_{\text{ours}}(w(m))$ as

$$\hat{r}_{\text{ours}}(w(m)) = \max(0, \tilde{r}_{\text{ours}}(w(m))) \quad (15)$$

to obtain our estimator. The estimator $\hat{r}_{\text{ours}}(w(m))$ has two parameters, λ_{item} and λ_d , which balance the impacts of $\hat{r}_{\text{item}}(w)$ and $t_d(w)$ in Eq. (9) while mitigating the low-frequency problem. As a result, the estimator allows us to control the dependency strength between items.

4.4 Regularization Effect on $\hat{r}_{\text{ours}}(w)$

In this section, we explain our estimator $\hat{r}_{\text{ours}}(w)$ in Eq. (15)

[†] As shown in Eq. (14), $\hat{t}_d(w(m))$ is calculated from word-based frequencies, and we set the largest N -gram order N to 4 in our experiments. Therefore, it is feasible to estimate t_d within a realistic time.

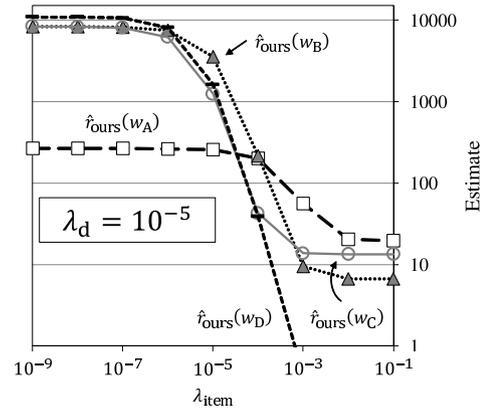


Fig. 2 Regularization effect on our estimator $\hat{r}_{\text{ours}}(w)$ shown in Eq. (15). The horizontal and vertical axes of the graph represent the regularization parameter λ_{item} and estimate of $\hat{r}_{\text{ours}}(w)$, respectively. To confirm the effect of λ_{item} , we have fixed λ_d at 10^{-5} .

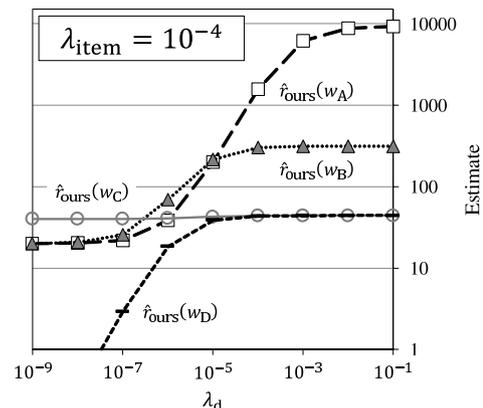


Fig. 3 Regularization effect on our estimator $\hat{r}_{\text{ours}}(w)$ shown in Eq. (15). The horizontal and vertical axes of the graph represent the regularization parameter λ_d and estimate of $\hat{r}_{\text{ours}}(w)$, respectively. To confirm the effect of λ_d , we have fixed λ_{item} at 10^{-4} .

from the perspective of the regularization process. As in Sect. 1, we calculate the likelihood ratio, $r(w) = \frac{p_{\text{nu}}(w)}{p_{\text{de}}(w)}$, to predict the left context of the NEs. $p_{\text{de}}(w)$ is the probability density where w occurs at any position in the corpus, and $p_{\text{nu}}(w)$ is the probability density where w occurs on the left side of the NEs. That is, if w is the left context of NEs, $r(w)$ is high. Otherwise, $r(w)$ is low. We assume that we observe the frequencies shown in Table 3 and use them to estimate $r(w)$ for each w . $\hat{r}_{\text{ours}}(w)$ has two regularization parameters: λ_{item} and λ_d . To clarify the effect of each parameter, we fix one and change the other over the range $10^{-9}, 10^{-8}, \dots, 10^{-1}$.

Figures 2 and 3 show the regularization effects on $\hat{r}_{\text{ours}}(w)$ by λ_{item} and λ_{d} , respectively.

In Fig. 2, $\hat{r}_{\text{ours}}(w)$ for all w decreases as λ_{item} grows. For $\lambda_{\text{item}} = 10^{-9}$, $\hat{r}_{\text{ours}}(w_{\text{B}})$, $\hat{r}_{\text{ours}}(w_{\text{C}})$, and $\hat{r}_{\text{ours}}(w_{\text{D}})$ are all nearly 10,000 because of the weak regularization. The estimates are close to the $\hat{r}_{\text{item}}(w)$ obtained with $\lambda_{\text{item}} = 0$ in Table 3. We infer that $\hat{r}_{\text{ours}}(w_{\text{A}})$ is lower than the other estimates as a result of the subtraction term in Eq. (14). For $\lambda_{\text{item}} = 10^{-5}$, $\hat{r}_{\text{ours}}(w_{\text{C}})$ is lower than $\hat{r}_{\text{ours}}(w_{\text{B}})$. As the itemized frequencies of w_{C} are lower than those of w_{B} , the regularization acts more strongly on $\hat{r}_{\text{ours}}(w_{\text{C}})$. As λ_{item} grows further, $\hat{r}_{\text{ours}}(w_{\text{A}})$ produces the highest estimate, $\hat{r}_{\text{ours}}(w_{\text{B}})$ and $\hat{r}_{\text{ours}}(w_{\text{C}})$ produce nearly identical estimates, and $\hat{r}_{\text{ours}}(w_{\text{D}})$ goes to zero. Here $\hat{r}_{\text{ours}}(w)$ becomes robust for itemized frequencies and original frequencies $c_*(w)$ are more heavily emphasized because $\lambda_{\text{item}} > \lambda_{\text{d}}$. The setting, $\lambda_{\text{item}} > \lambda_{\text{d}}$, is effective in the case where the required dependency between items is strong. Let us assume that the entity type is location name, and the N -gram w is “Hospital in” ($a_1 = \text{“Hospital,”}$ $a_2 = \text{“in”}$). Because the preposition “in” occurs in various contexts, and “Hospital” at the left of “in” is the key to predicting whether w is the left context of location names. For such an example, the co-occurrence frequency $c_*(w)$ of two words can be used to incorporate the dependency into the estimate.

Unlike in Fig. 2, in Fig. 3 $\hat{r}_{\text{ours}}(w)$ increases for all w as λ_{d} grows. For $\lambda_{\text{d}} = 10^{-9}$, each $\hat{r}_{\text{ours}}(w)$ is close to the corresponding $\hat{r}_{\text{MLE}}(w)$ in Table 3 because of the weak regularization. For $\lambda_{\text{d}} = 10^{-5}$, $\hat{r}_{\text{ours}}(w_{\text{B}})$ is lower than $\hat{r}_{\text{ours}}(w_{\text{A}})$. As the $c_*(w_{\text{B}})$ are lower than the $c_*(w_{\text{A}})$, the regularization acts more strongly on $\hat{r}_{\text{ours}}(w_{\text{B}})$. As λ_{d} grows further, $\hat{r}_{\text{ours}}(w_{\text{A}})$ and $\hat{r}_{\text{ours}}(w_{\text{B}})$ produce the highest and second-highest estimates, respectively, while $\hat{r}_{\text{ours}}(w_{\text{D}})$ produces an estimate nearly equal to that for $\hat{r}_{\text{ours}}(w_{\text{C}})$ despite the fact that $c_{\text{nu}}(w_{\text{D}}) = 0$. Here $\hat{r}_{\text{ours}}(w)$ becomes robust for original frequencies and itemized frequencies are more heavily emphasized because $\lambda_{\text{d}} > \lambda_{\text{item}}$. The setting, $\lambda_{\text{d}} > \lambda_{\text{item}}$, is effective in the case where the required dependency between items is weak. Let us assume that the entities type is person name, and $w = \text{“to Mr.”}$ Here, we can imagine that a person name will be on the right side of $a_2 = \text{“Mr.”}$ regardless of a_1 . For such an example, the frequency information for each word a_k can be strongly reflected in the estimate.

Overall, these results indicate a trade-off between the original and itemized frequencies; as such, we can control the dependency between items by tuning λ_{item} and λ_{d} .

5. Objective Evaluation

In this section, we analyze the behavior of our method and verify its effectiveness in mitigating the low- and zero-frequency problems. To do so, we evaluate a task involving the use of LR estimators to predict the left context of named entities (NEs) in datasets. We use two types of NEs—names of people and locations—which are tagged PER and LOC, respectively, and adopt word N -grams as their contexts. Figure 4 shows examples of the two types of NEs

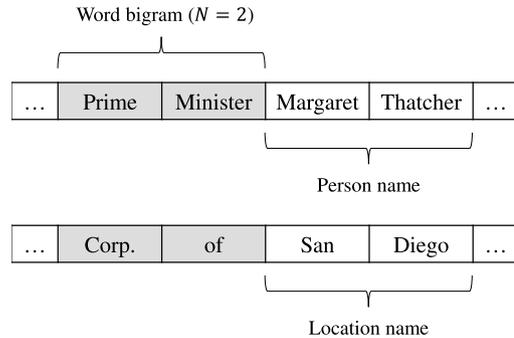


Fig. 4 Examples of two types of named entities (NEs) and their left bigrams. The PER context (left bigram) contains the noun representing a person “Minister.” The LOC context contains the preposition “of,” which tends to appear in various contexts.

and their left bigram ($N = 2$). The left N -grams of PER tend to contain specific words, e.g., honorific titles or nouns representing a person. As such, the words are crucial to predicting the PER context. By contrast, LOC N -grams tend to contain prepositions, which are contained in various contexts. Therefore, unlike in the PER context, contextual sequences (here indicating N -grams) are crucial. These NEs with different context features are helpful for behavioral analysis. Furthermore, the N -grams to be predicted are uniquely determined, which facilitates objective evaluation. Finally, the N -gram frequency distributions are long-tail distributions, which forces us to handle numerous rare or unobserved items. Therefore, we must be aware that methods used to handle them will significantly affect the prediction performance, allowing the differences between our method and the comparison methods described in Sect. 5.3 to be observed clearly.

5.1 Dataset Descriptions and Experimental Conditions

We used NE-tagged datasets created by applying the following steps. First, we classified and tagged entities in the 1987 Wall Street Journal corpus[†] using the Stanford Named Entity Recognizer [30]. We then randomly sampled articles from the tagged corpus and allocated them to the training, validation, and test sets. The validation and test set sizes were fixed at 1,000 articles in all cases, while the training set size S_{tr} was changed according to the following conditions, all of which had two choices^{††}.

- N -gram order N : 2, 4
- training set size S_{tr} : 10,000 articles, 2,500 articles
- NE type: location names (LOC), person names (PER)

We then conducted experiments at eight condition settings. The first two conditions were for controlling training set

[†]<https://catalog.ldc.upenn.edu/LDC2000T43>

^{††}We conducted experiments with N varied from two to four. We obtained similar results for the $N = 2$ and $N = 3$ cases. Owing to the large degree of sparseness, we expected that the results for $N > 4$ would not differ significantly from those of the $N = 4$ cases; therefore, we set $N = 2$ and 4 as the N -gram order choices.

Table 4 # of bigrams contained in each dataset.

Size of Train S_{tr}	Dataset	O_{ALL}		O_{LOC}		O_{PER}	
		types	total	types	total	types	total
10,000	Train	1,468,292	4,002,930	31,294	64,072	41,123	67,780
	Valid	230,528	401,445	4,318	6,116	5,795	7,463
	Test	231,931	403,145	4,164	5,876	6,035	7,635
2,500	Train	493,368	1,000,627	9,983	15,965	12,458	17,292
	Valid	237,805	415,369	4,752	6,736	5,522	6,973
	Test	233,143	402,320	4,448	6,372	5,318	6,622

Table 5 # of 4-grams contained in each dataset.

Size of Train S_{tr}	Dataset	O_{ALL}		O_{LOC}		O_{PER}	
		types	total	types	total	types	total
10,000	Train	3,618,822	3,982,930	59,167	63,665	64,729	67,437
	Valid	385,604	399,445	5,931	6,064	7,355	7,430
	Test	388,401	401,145	5,713	5,832	7,497	7,593
2,500	Train	947,710	995,627	15,305	15,851	16,919	17,219
	Valid	400,336	413,369	6,574	6,699	6,861	6,940
	Test	387,489	400,320	6,212	6,327	6,537	6,589

sparseness; the last was for predicting different types of context. The dataset descriptions are shown in Tables 4 and 5. Note that articles in the validation and test sets were re-sampled for different training set sizes. O_{ALL} indicates that N -gram occurrences are at any position in a dataset. O_{LOC} and O_{PER} indicate occurrences at the left side of LOC and PER, respectively, in a dataset.

5.2 Experimental Procedure

For any N -gram w in the training set, we counted the frequencies on both the left side of the NEs and at any position. Similarly, we counted the frequencies of $\{\Psi_k w\}_{k=1}^N$ to estimate $r_{item}(w)$. We then estimated $r(w)$ using the frequencies for each w in the test set as follows:

$$r(w) = \frac{p_{NE}(w)}{p_{ALL}(w)},$$

where $p_{NE}(w)$ and $p_{ALL}(w)$ are the occurrence probabilities of w on the left side of the NEs and at any position in the training set, respectively.

We sorted and ranked w in descending order of their estimates and evaluated the top 8,000. If w occurred on the left side of the NEs more than once in the test set, it was labeled “True”; otherwise, it was labeled “False.” Finally, we constructed a rank–recall curve on a graph in which the horizontal and vertical axes represent rank and recall, respectively. From the curve, we can see the approximate precision. The slope of a straight line connecting the origin to a point on the curve is proportional to the precision at that point. The recall and precision are defined as follows:

$$\text{recall} = \frac{|\{w|w \in R\}|}{|R|}, \quad \text{precision} = \frac{|\{w|w \in R\}|}{|\{w\}|},$$

where w is the N -gram to be evaluated and R is the set of N -grams occurring on the left side of the NEs in the test set, i.e., the correction set.

5.3 Comparison Methods

In this section, we describe the comparison methods and the parameter tuning process. We used the following comparison methods, which have different behaviors for low-frequency (rare) or zero-frequency (unobserved) elements or for both types of elements.

A baseline (B for short, Eq. (1)) first obtains an MLE of each probability distribution and then takes the ratio of the MLEs. This simple approach is often called “indirect estimation” because it involves probability estimation; however, it faces both the low- and zero-frequency problems.

Kikuchi:19 (K for short, Eq. (7)) is a variant of uLSIF that is applied to discrete distributions [7]. K introduces a regularization parameter λ to mitigate the low-frequency problem but does not use frequencies of item units[†]. Thus, it cannot estimate LRs from unobserved elements.

In probability estimation, Kneser–Ney smoothing [18] is widely used to solve the zero-frequency problem. This smoothing uses lower order N -gram frequencies in addition to original N -gram frequencies. A possible way to estimate $r(w)$ even from unobserved elements might be to smooth each probability estimate and then take the ratio of the estimates. Therefore, we add the following method, which we call the Kneser–Ney-based method (KN for short), for comparison in $N = 4$ settings:

$$\hat{r}_{KN}(w) = \frac{\hat{p}_{nu}(w)}{\hat{p}_{de}(w)},$$

[†]For details of K and ITEM, see Sects. 3.3 and 4.2, respectively.

Table 6 Regularization parameters chosen for $N = 2$.

Size of Train S_{tr}	NE	K	ITEM	Ours	
		λ	λ'	λ_{item}	λ_d
10,000	LOC	10^{-2}	10^{-5}	10^{-1}	10^{-4}
	PER	10^{-3}	10^{-5}	10^{-4}	10^{-5}
2,500	LOC	10^{-1}	10^{-4}	10^{-1}	10^{-3}
	PER	10^{-1}	10^{-5}	10^{-4}	10^{-5}

Table 7 Regularization parameters chosen for $N = 4$.

Size of Train S_{tr}	NE	K	ITEM	Ours	
		λ	λ'	λ_{item}	λ_d
10,000	LOC	10^{-6}	10^{-6}	10^{-5}	10^{-8}
	PER	10^{-6}	10^{-5}	10^{-5}	10^{-6}
2,500	LOC	10^{-1}	10^{-5}	10^{-5}	10^{-6}
	PER	10^{-1}	10^{-5}	10^{-5}	10^{-5}

$$\begin{aligned} \hat{p}_*(w) = & p_*^S(a_1)p_*^S(a_2 | a_1) \\ & p_*^S(a_3 | a_1a_2)p_*^S(a_4 | a_1a_2a_3), \end{aligned}$$

where $\hat{p}_{de}(w)$ and $\hat{p}_{nu}(w)$ are modeled with smoothed conditional probabilities for $\{a_k\}_{k=1}^4$, and p_*^S is the (interpolated) Kneser–Ney estimator and has a discount parameter d . That is to say, KN has eight parameters in total. Note that $*$ \in {de, nu} is an index that distinguishes the denominator and numerator of a likelihood ratio.

Itemized LR (ITEM for short, Eq. (10)) mitigates the zero-frequency problem by itemizing N -grams to item units[†]. The low-frequency problem can be mitigated by using K for the LR estimation of each unit. As this approach does not use original N -gram frequencies, the itemization sacrifices dependencies between items. This is equivalent to our model without the dependency term $t_d(w)$, and we therefore include it in our comparison to confirm the effectiveness of $t_d(w)$. To avoid duplication in the model’s notation, the regularization parameter is replaced by λ' .

With the exception of B, all of the evaluated methods have tuning parameters. We therefore regarded the validation sets as test sets and constructed rank–recall curves with the regularization parameters changed over the range $10^{-9}, 10^{-8}, \dots, 10^{-1}$. Finally, we chose the optimal value that maximized the area under the rank–recall curve. As our method uses two parameters— λ_{item} and λ_d —the optimal pair was found by changing them simultaneously. The chosen regularization parameters are listed in Tables 6 and 7. As described above, KN has eight tuning parameters. For each Kneser–Ney estimator, we varied the parameter d over the range 0.1, 0.2, \dots , 0.9 and chose the optimal value that maximized its likelihood. This is a natural way to set d s when using the smoothing.

5.4 Results

We first show the respective rank–recall curves, which are summarized for different N . Then, for behavioral analysis,

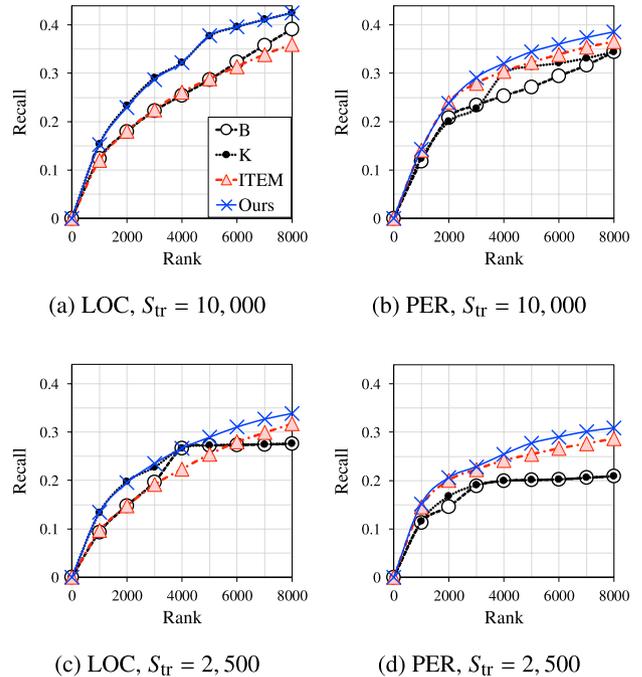


Fig. 5 Rank–recall curves for $N = 2$. PER and LOC indicate that prediction targets are the left PER and LOC contexts, respectively. S_{tr} is the training set size. The horizontal and vertical axes of curves represent rank and recall, respectively. The slope of a straight line connecting the origin to a point on the curve is proportional to the precision at that point. For a given rank, the method with the largest value on the vertical axis, i.e., the highest recall, has the best performance.

we show some ranked N -grams. Owing to space limitations, only the N -grams for $N = 2$, $S_{tr} = 10,000$ and $N = 4$, $S_{tr} = 2,500$ are shown.

Rank-Recall Curves for $N = 2$

Figure 5 shows the rank–recall curves for bigrams ($N = 2$). From this figure, we can see the following:

- The baseline (B) performed worse than our method as a result of the low-frequency problem. In addition, Kikuchi:19 (K) performed generally better than B for $S_{tr} = 10,000$ because of its regularization. In particular, K was competitive with our method (Fig. 5(a)). However, both methods suffered from unobserved bigrams for $S_{tr} = 2,500$ and could not improve recall in the low-rank range (Figs. 5(c) and 5(d)).
- Itemized LR (ITEM) demonstrated improving recall up to rank 8,000. This suggests that itemizing N -grams is effective in mitigating the zero-frequency problem. Moreover, this method produced the second-best performance results in PER context prediction, as shown in Figs. 5(b) and 5(d). However, in terms of LOC context prediction, ITEM performed worse than our method and by far the worst for $S_{tr} = 10,000$ (Fig. 5(a)) despite having the richest training context.
- Our method, ITEM with $t_d(w)$, was able to mitigate the zero-frequency problem while maintaining dependency between items. Moreover, its regularization mitigated

Table 8 Ranked bigrams $w = a_1 a_2$ for $S_{tr} = 10,000$.

w	$a_1 a_2$	# w		Rank				T/F	$\widehat{r}_{ours}(w)$	
		O_{ALL}	O_{LOC}	B	K	ITEM	Ours		$\widehat{r}_{item}(w)$	$\widehat{t}_d(w)$
w_A	Hospital in	22	22	1	190	86	115	T	4.88×10^{-3}	3.25
w_B	court in	186	156	602	16	180	3	T	3.40×10^{-2}	16.61
w_C	chair of	1	1	1	4,153	–	4,119	F	9.94×10^{-5}	0.16
w_D	Europe in	12	0	–	–	147	–	F	9.54×10^{-3}	-2.78×10^{-4}

w	$a_1 a_2$	# w		Rank				T/F	$\widehat{r}_{ours}(w)$	
		O_{ALL}	O_{PER}	B	K	ITEM	Ours		$\widehat{r}_{item}(w)$	$\widehat{t}_d(w)$
w_E	Prime Minister	228	228	1	11	1	19	T	614.42	-472.42
w_F	says Mr.	632	537	1,210	2	89	159	T	287.75	-223.42
w_G	\$53.9 million,	1	1	1	3,045	–	6,069	F	2.18×10^{-2}	1.44
w_H	So President	0	0	–	–	260	82	T	81.09	0

Table 9 Itemized frequencies of unobserved bigrams for $S_{tr} = 10,000$.

w	$a_1 a_2$	#($a_1 \bullet$)		#($\bullet a_2$)	
		O_{ALL}	O_{LOC}	O_{ALL}	O_{LOC}
w_D	Europe in	325	45	73,816	10,312

w	$a_1 a_2$	#($a_1 \bullet$)		#($\bullet a_2$)	
		O_{ALL}	O_{PER}	O_{ALL}	O_{PER}
w_H	So President	559	32	1,214	1,172

the low-frequency problem. Overall, our method outperformed the other methods on all graphs excluding Fig. 5(a), in which it was competitive with K.

Ranked Bigram Examples for $S_{tr} = 10,000$

To clarify the factors underlying differences in performance among the methods, we analyzed the behavior of each method qualitatively. Table 8 shows word bigrams ranked by method for $S_{tr} = 10,000$. In the table, “–” indicates that the bigram is ranked below the top 8,000; “#” indicates the frequencies in the training set. Table 9 shows itemized frequencies of the unobserved bigrams in Table 8. The rank–recall curves corresponding to these tables are shown in Figs. 5(a) and 5(b). From the tables, we can see the following:

- B ranked w_A , w_C , w_E and w_G at the top. These always occurred on the left side of NEs in the training set. While B allowed placing of the true bigrams w_A and w_E within high ranks, the rare and false bigrams w_C and w_G , respectively, were also placed within high ranks because it overestimated the LR of rare bigrams. The frequent bigrams w_B and w_F , which should have been in the higher ranks, were ranked lower than they were by the other methods. Therefore, B does not have good performance for any S_{tr} or NE type.
- K placed frequent and true bigrams w_A , w_B , w_E and w_F within the higher ranks while placing the rare and false bigrams w_C and w_G in the lower ranks. This suggests that the regularization worked well and resulted in better performance than B.

- ITEM behaved in a manner similar to K except for the unobserved bigrams w_D and w_H . This suggests that itemizing w was not harmful for observed bigrams. However, this method could not take into account the dependency between a_1 and a_2 and therefore computed item pair estimates without considering their co-occurrence. Consequently, the false bigram w_D , which is not a pair of co-occurring items but rather a pair of separate, frequent items, as shown in Table 9, was placed in a high rank. We can infer that this degraded the prediction performance of LOC context, which required a strong dependency.
- Our method can adjust the dependency strength according to usage and application. Because LOC context prediction requires a strong dependency, the parameters were chosen so that $\lambda_{item} > \lambda_d$, as shown in Table 6. Thus, $\widehat{t}_d(w)$ was much larger than $\widehat{r}_{item}(w)$ except for w_D . Meanwhile, $\widehat{t}_d(w_D)$ was smaller than $\widehat{r}_{item}(w_D)$ owing to the low dependency; i.e., $\#w_D$ of O_{LOC} was zero. As a result, w_D had the small estimate and was placed in a lower rank. In the PER context prediction process, by contrast, the parameters were chosen so that the difference between them was smaller than in the LOC case. Unlike LOC cases, $\widehat{r}_{item}(w)$ was much larger than $\widehat{t}_d(w)$ except for w_G . In particular, $\widehat{r}_{ours}(w_H)$ was large even though the $\#w_H$ of O_{ALL} and O_{PER} were both zero. (\bullet President) was the most frequent and crucial unit, as shown in Table 9. Here, our method paid attention to occurring individual items and not their dependencies. As a result, w_H had the large estimate and was placed in a high rank. Overall, we can hypothesize that this behavior resulted in the best performance for predictions in different contexts.

Rank-Recall Curves for $N = 4$

Figure 6 shows the rank–recall curves for four-grams ($N = 4$). In this case, the context was much sparser than for $N = 2$, and therefore all methods suffered from a lack of frequency. From the figure, we can see the following:

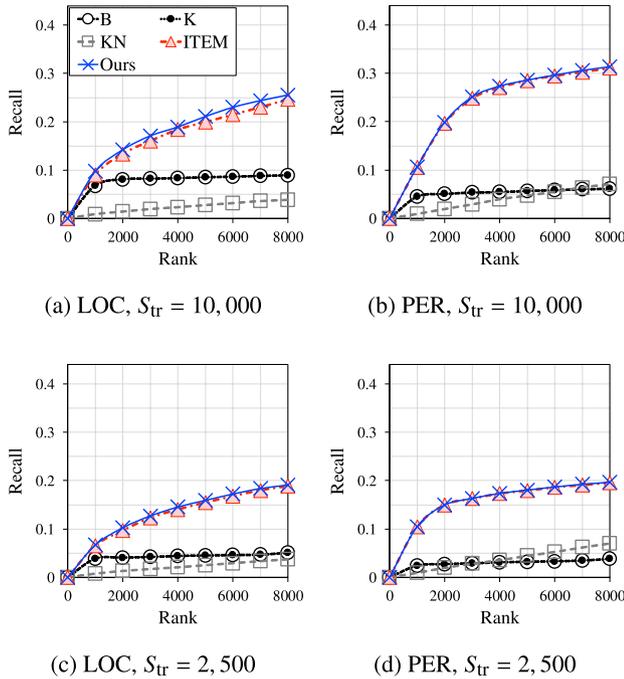


Fig. 6 Rank–recall curves for $N = 4$. PER and LOC indicate that prediction targets are left PER and LOC contexts, respectively. S_{tr} is the training set size. The horizontal and vertical axes of the curves represent rank and recall, respectively. The slope of a straight line connecting the origin to a point on the curve is proportional to the precision at that point. For a given rank, the method with the largest value on the vertical axis, i.e., the highest recall, has the best performance.

- B and K performed in a similarly poor manner in all rank–recall curves. Their precisions improved slightly in the high-rank ranges, but did not improve thereafter even at $S_{tr} = 10,000$ cases. This suggests that the methods discovered very few contexts in the sparse data settings.
- The Kneser–Ney-based method (KN) also performed poorly. In probability estimation, Kneser–Ney smoothing is known to be an effective method for solving the zero-frequency problem. However, we found that taking the ratio of smoothed estimates significantly degraded performance.
- In the high-rank ranges of all curves, ITEM and our method performed with the highest precisions. Furthermore, they continued to improve their recall up to rank 8,000, a result that confirms the effectiveness of itemizing N -grams. Unlike the bigram cases, these two methods performed similarly; we can infer from this that the effect of $t_d(w)$ was weakened. However, our method performed slightly better than ITEM in LOC context prediction, as shown in Figs. 6(a) and 6(c). This result confirms that incorporating the dependency is also effective for large- N settings.

Ranked Four-gram Examples for $S_{tr} = 2,500$

We analyzed the behavior of each method for $N = 4$.

Table 10 shows the word four-grams ranked by methods for $S_{tr} = 2,500$. Table 11 shows the itemized frequencies of the unobserved four-grams in Table 10. $\Psi_k w$ for $w = a_1 a_2 \cdots a_N$ represents $\bullet \cdots \bullet a_k \cdots \bullet$ as defined in Sect. 4.1. The rank–recall curves corresponding to these tables are shown in Figs. 6(c) and 6(d). From the tables, we can see the following:

- For B and K, all unobserved four-grams w_C , w_D , w_G , and w_H were ranked at 611 and 415. This indicates that there were only 610 and 414 observed four-grams for LOC and PER, respectively, that could be LRs estimated by the two methods. Thus, B and K performed in a similarly poor manner.
- The behavior of KN was very different from that of the other approaches. The observed four-grams w_A and w_E , which always occurred on the left side of the NEs in the training set, were not ranked in the top 8,000; instead, the higher ranks included the unobserved and false four-grams w_D and w_H . Both of these four-grams are composed of unobserved units, as shown in Table 11, and are unreliable. This implies that taking the ratio of smoothed estimates resulted in poor performance.
- In focusing on unobserved four-grams, both ITEM and our method could ignore the false four-grams w_D and w_H and ranked the true four-grams w_C and w_G in high ranks. Nearly all units of w_C and w_G were frequent, as seen from Table 11. These results further confirm the effectiveness of itemizing N -grams. We could not confirm any difference in the results with and without $\widehat{t}_d(w)$, as the chosen λ_{item} and λ_d were close values even in the LOC cases, as seen from Table 7. In addition, we confirmed many cases where $\widehat{t}_d(w)$ was zero even in high ranks, as seen from w_C and w_G . From the above, we can infer that $\widehat{t}_d(w)$ does not work well in sparse data settings such as the 4-gram cases, and we leave the development of an effective approach to incorporating dependency to future work.

5.5 Discussion

In this section, we summarize the experimental results and discuss the benefits of our method and future works.

B was found to face both the low- and zero-frequency problems, and therefore its effectiveness is narrow and limited. We confirmed that K handled low frequencies well but suffered from the zero-frequency problem. We found that KN, which involves taking the ratio of smoothed estimates, was useless. KN has eight parameters, which can conceivably be tuned simultaneously. However, varying some probability estimators simultaneously no longer satisfies the smoothing scheme, which corrects a single estimator.

ITEM, involving the itemization of N -grams, was effective for rare and unobserved context prediction but sacrificed the dependency between items. Thus, our method introduced the term $t_d(w)$ to ITEM to incorporate the dependency. This resulted in the best performance for two

Table 10 Ranked four-grams $w = a_1 a_2 a_3 a_4$ for $S_{tr} = 2, 500$.

w	$a_1 a_2 a_3 a_4$	# w		Rank					T/F	$\widehat{r}_{ours}(w)$	
		O_{ALL}	O_{LOC}	B	K	KN	ITEM	Ours		$\widehat{r}_{item}(w)$	$\widehat{r}_d(w)$
w_A	University of California at	7	7	1	8	–	307	678	T	204.63	–124.16
w_B	a large number of	5	1	487	518	–	–	4,850	F	5.07	6.25
w_C	in West Germany, France,	0	0	611	611	–	57	41	T	1655.59	0
w_D	Christ Church in Grosse	0	0	611	611	169	–	–	F	0	0

w	$a_1 a_2 a_3 a_4$	# w		Rank					T/F	$\widehat{r}_{ours}(w)$	
		O_{ALL}	O_{PER}	B	K	KN	ITEM	Ours		$\widehat{r}_{item}(w)$	$\widehat{r}_d(w)$
w_E	In an interview, Mr.	7	7	1	4	–	38	67	T	1560.35	–620.29
w_F	Procter & Gamble Co.	4	1	290	314	–	3,122	2,799	F	7.60	1.96
w_G	promised by Prime Minister	0	0	415	415	–	14	13	T	3153.59	0
w_H	Mr. Taniguchi said Denmark	0	0	415	415	18	–	–	F	0	0

Table 11 Itemized frequencies of unobserved four-grams for $S_{tr} = 2, 500$.

w	$a_1 a_2 a_3 a_4$	# $\Psi_1 w$		# $\Psi_2 w$		# $\Psi_3 w$		# $\Psi_4 w$	
		O_{ALL}	O_{LOC}	O_{ALL}	O_{LOC}	O_{ALL}	O_{LOC}	O_{ALL}	O_{LOC}
w_C	in West Germany, France,	18,456	395	365	23	41	12	22	11
w_D	Christ Church in Grosse	0	0	11	0	18,618	1,034	0	0

w	$a_1 a_2 a_3 a_4$	# $\Psi_1 w$		# $\Psi_2 w$		# $\Psi_3 w$		# $\Psi_4 w$	
		O_{ALL}	O_{PER}	O_{ALL}	O_{PER}	O_{ALL}	O_{PER}	O_{ALL}	O_{PER}
w_G	promised by Prime Minister	46	1	5,445	128	69	59	100	99
w_H	Mr. Taniguchi said Denmark	5,704	141	0	0	6,170	161	3	0

types of context prediction, which required different dependency strengths, suggesting that our method can control the dependency strength. In experiments with the different N -gram order N s, the best comparison method was different for N , whereas our method consistently showed good performance. In NLP, the proper N sometimes varies depending on practical tasks. Considering this fact, our method has the advantage that it can be used for various N s. Furthermore, our method was found to be able to treat the entire high- to zero-frequency range within a unified framework that does not depend on a particular application. This leads us to believe that our method will be helpful in various analyses and applications, including feature extraction, relation detection, and classification.

By contrast, $t_d(w)$ measures dependency using only original N -gram frequencies and, for this reason, does not work well at sparse frequencies. As an approach to incorporating dependency more effectively in future work, we could assume a multiple Markov independence in item occurrence. We could also refer to previous research on alleviating the conditional independence assumption in naive Bayes classifiers [21]–[26] and consider whether these approaches are applicable to our method. This might allow us to provide effective estimates from sparse data or sequences of long- N (≥ 4) items. As described in Sect. 1, likelihood ratios were used in numerous tasks. Among these, our method can be used in cases where observation targets are continuous sequences (e.g., N -grams). Another future work is to find further practical applications that use the sequences and

verify the practicality of our method.

6. Conclusion

This paper presented a method for mitigating the low- and zero-frequency problems of LR estimation. One possible approach for mitigating the zero-frequency problem is itemizing w and roughly approximating $r(w)$ as $r_{item}(w)$. Although this is a straightforward approach, it ignores dependencies between items. Therefore, our proposed method adds the term $t_d(w)$, which is estimated from the N -gram frequencies, to $r_{item}(w)$ to incorporate dependency and allow for the successful treatment of observed and unobserved N -grams. To mitigate the low-frequency problem, we combined this framework with an existing method [7]. The two regularization parameters, λ_{item} and λ_d , introduced by this combination, allowed us to control the dependency. Our experimental results demonstrated that our method could mitigate both problems. On the other hand, $t_d(w)$ measured dependency using only the original N -gram frequencies. Thus, we found that it did not work well when the frequency was sparse. Hence, our future work includes the development of an approach to incorporate dependency more effectively.

References

- [1] Y. Jo, C. Lagoze, and C.L. Giles, “Detecting research topics via the correlation between graphs and texts,” *SIGKDD*, pp.370–379, 2007.
- [2] M. Huang, B. Ye, Y. Wang, H. Chen, J. Cheng, and X. Zhu, “New word detection for sentiment analysis,” *ACL*, pp.531–541, 2014.

- [3] M. Ravichandran, G. Kulanthaivel, and T. Chellatamilan, “Intelligent topical sentiment analysis for the classification of e-learners and their topics of interest,” *The Scientific World Journal*, pp.1–8, 2015.
- [4] T.K. Tran and T.T. Phan, “Deep learning application to ensemble learning—The simple, but effective, approach to sentiment classifying,” *Applied Sciences*, vol.9, no.13, pp.1–18, 2019.
- [5] T. Dunning, “Accurate methods for the statistics of surprise and coincidence,” *Computational Linguistics*, vol.19, no.1, pp.61–74, 1993.
- [6] C.D. Manning and H. Schütze, *Foundations of Statistical Natural Language Processing*, MIT press, 1999.
- [7] M. Kikuchi, K. Kawakami, M. Yoshida, and K. Umemura, “Conservative direct estimation for likelihood ratios based on observed frequencies,” *IEICE Trans. Inf. & Syst. (Japanese Edition)*, vol.J102-D, no.4, pp.289–301, 2019.
- [8] T. Kanamori, S. Hido, and M. Sugiyama, “A least-squares approach to direct importance estimation,” *J. Machine Learning Research*, vol.10, pp.1391–1445, July 2009.
- [9] M. Johnson, “Confidence intervals on likelihood estimates for estimating association strengths,” unpublished technical report, 1999.
- [10] P. Pantel and D. Ravichandran, “Automatically labeling semantic classes,” *NAACL-HLT*, pp.321–328, 2004.
- [11] C. Rudin, B. Letham, and D. Madigan, “Learning theory analysis for association rules and sequential event prediction,” *J. Machine Learning Research*, vol.14, no.1, pp.3441–3492, 2013.
- [12] F. Jelinek and R.L. Mercer, “Interpolated estimation of Markov source parameters from sparse data,” *Workshop on Pattern Recognition in Practice*, pp.21–23, 1980.
- [13] S.M. Katz, “Estimation of probabilities from sparse data for the language model component of a speech recognizer,” *IEEE Trans. Acoust., Speech, Signal Process.*, vol.35, no.5, pp.400–401, 1987.
- [14] L.H. Witten and T.C. Bell, “The zero-frequency problem: estimating the probabilities of novel events in adaptive text compression,” *IEEE Trans. Inf. Theory*, vol.37, no.4, pp.1085–1094, 1991.
- [15] H. Ney and U. Essen, “On smoothing techniques for bigram-based natural language modelling,” *ICASSP*, pp.825–828, 1991.
- [16] H. Ney, U. Essen, and R. Kneser, “On structuring probabilistic dependences in stochastic language modeling,” *Computer Speech and Language*, vol.8, no.1, pp.1–38, 1994.
- [17] D.J.C. MacKay and L.C.B. Peto, “A hierarchical Dirichlet language model,” *Nat. Lang. Eng.*, vol.1, no.3, pp.289–308, 1995.
- [18] R. Kneser and H. Ney, “Improved backing-off for M-gram language modeling,” *ICASSP*, pp.181–184, 1995.
- [19] M. Parzen, S. Lipsitz, J. Ibrahim, and N. Klar, “An estimate of the odds ratio that always exists,” *J. Comput. Graph. Stat.*, vol.11, no.2, pp.420–436, 2002.
- [20] K. Raweesawat, Y. Areepong, K. Jampachaisri, and S. Sukparungsee, “Odds ratios estimation of rare event in binomial distribution,” *J. Probability and Statistics*, vol.2016, pp.1–8, 2016.
- [21] E. Frank, M. Hall, and B. Pfahringer, “Locally weighted naive bayes,” *UAI*, pp.249–256, 2002.
- [22] L. Jiang, D. Wang, and Z. Cai, “Discriminatively weighted naive bayes and its application in text classification,” *Int. J. Artif. Intell. Tools*, vol.21, no.1, 2012.
- [23] J. Wu, S. Pan, X. Zhu, P. Zhang, and C. Zhang, “SODE: Self-adaptive one-dependence estimators for classification,” *Pattern Recogn.*, vol.51, pp.358–377, 2016.
- [24] B. Tang, S. Kay, and H. He, “Toward optimal feature selection in naive bayes for text categorization,” *IEEE Trans. Knowl. Data Eng.*, vol.28, no.9, pp.2508–2521, 2016.
- [25] D.M. Diab and K.M.E. Hindi, “Using differential evolution for fine tuning naive bayesian classifiers and its application for text classification,” *Applied Soft Computing*, vol.54, pp.183–199, 2017.
- [26] L. Jiang, L. Zhang, C. Li, and J. Wu, “A correlation-based feature weighting filter for naive bayes,” *IEEE Trans. Knowl. Data Eng.*, vol.31, no.2, pp.201–213, 2018.
- [27] J. Huang, A.J. Smola, A. Gretton, K.M. Borgwardt, and B. Schölkopf,

“Correcting sample selection bias by unlabeled data,” *Advances in Neural Information Processing Systems 19*, pp.601–608, 2007.

- [28] S. Bickel, M. Brückner, and T. Scheffer, “Discriminative learning for differing training and test distributions,” *ICML*, pp.81–88, 2007.
- [29] M. Sugiyama, S. Nakajima, H. Kashima, P. von Büna, and M. Kawanabe, “Direct importance estimation with model selection and its application to covariate shift adaptation,” *Advances in Neural Information Processing Systems 20*, pp.1433–1440, 2008.
- [30] J.R. Finkel, T. Grenager, and C. Manning, “Incorporating non-local information into information extraction systems by Gibbs sampling,” *ACL*, pp.363–370, 2005.

Appendix A: Derivation of the First Two Terms of the Objective Function in uLSIF

In uLSIF, β are determined so that the squared loss $J_0(\beta)$ is minimized:

$$\begin{aligned} J_0(\beta) &= \frac{1}{2} \int (\widehat{r}(x) - r(x))^2 p_{de}(x) dx \\ &= \frac{1}{2} \int \widehat{r}(x)^2 p_{de}(x) dx - \int \widehat{r}(x) r(x) p_{de}(x) dx \\ &\quad + \frac{1}{2} \int r(x)^2 p_{de}(x) dx \\ &= \frac{1}{2} \int \widehat{r}(x)^2 p_{de}(x) dx - \int \widehat{r}(x) p_{nu}(x) dx \\ &\quad + \frac{1}{2} \int r(x)^2 p_{de}(x) dx, \end{aligned}$$

where $p_{de}(x)$ is canceled in the second term. In addition, the last term of $J_0(\beta)$ is a constant and therefore can be safely ignored. Let the first two terms be $J(\beta)$:

$$\begin{aligned} J(\beta) &= \frac{1}{2} \int \widehat{r}(x)^2 p_{de}(x) dx - \int \widehat{r}(x) p_{nu}(x) dx \\ &= \frac{1}{2} \sum_{l,l'=1}^b \beta_l \beta_{l'} \left(\int \varphi_l(x) \varphi_{l'}(x) p_{de}(x) dx \right) \\ &\quad - \sum_{l=1}^b \beta_l \left(\int \varphi_l(x) p_{nu}(x) dx \right). \end{aligned}$$

Approximating the expectations in $J(\beta)$ by empirical averages, we obtain

$$\begin{aligned} \widehat{J}(\beta) &= \frac{1}{2n_{de}} \sum_{i=1}^{n_{de}} \widehat{r}(x_i^{de})^2 - \frac{1}{n_{nu}} \sum_{j=1}^{n_{nu}} \widehat{r}(x_j^{nu}) \\ &= \frac{1}{2} \sum_{l,l'=1}^b \beta_l \beta_{l'} \left(\frac{1}{n_{de}} \sum_{i=1}^{n_{de}} \varphi_l(x_i^{de}) \varphi_{l'}(x_i^{de}) \right) \\ &\quad - \sum_{l=1}^b \beta_l \left(\frac{1}{n_{nu}} \sum_{j=1}^{n_{nu}} \varphi_l(x_j^{nu}) \right) \\ &= \frac{1}{2} \beta^T \widehat{H} \beta - \widehat{h}^T \beta, \end{aligned}$$

which are the first two terms of the objective function in Eq. (3).

Appendix B: Derivation of $\widehat{J}(\boldsymbol{\beta})$ in Our Method

Let the squared loss be $J_0(\boldsymbol{\beta})$:

$$\begin{aligned} J_0(\boldsymbol{\beta}) &= \frac{1}{2} \int (r_{\text{ours}}(w) - r(w))^2 p_{\text{de}}(w) dw \\ &= \frac{1}{2} \int r_{\text{ours}}(w)^2 p_{\text{de}}(w) dw \\ &\quad - \int r_{\text{ours}}(w) r(w) p_{\text{de}}(w) dw \\ &\quad + \frac{1}{2} \int r(w)^2 p_{\text{de}}(w) dw \\ &= \frac{1}{2} \int r_{\text{ours}}(w)^2 p_{\text{de}}(w) dx \\ &\quad - \int r_{\text{ours}}(w) p_{\text{nu}}(w) dw \\ &\quad + \frac{1}{2} \int r(w)^2 p_{\text{de}}(w) dw, \end{aligned}$$

where $p_{\text{de}}(w)$ is canceled in the second term and the last term can be safely ignored. Then, let the first two terms be $J(\boldsymbol{\beta})$:

$$\begin{aligned} J(\boldsymbol{\beta}) &= \frac{1}{2} \int r_{\text{ours}}(w)^2 p_{\text{de}}(w) dw \\ &\quad - \int r_{\text{ours}}(w) p_{\text{nu}}(w) dw \\ &= \frac{1}{2} \sum_{l,l'=1}^v \beta_l \beta_{l'} \left(\int \varphi_l(w) \varphi_{l'}(w) p_{\text{de}}(w) dw \right) \\ &\quad + \widehat{r}_{\text{item}}(w) \sum_{l=1}^v \beta_l \left(\int \varphi_l(w) p_{\text{de}}(w) dw \right) \\ &\quad - \sum_{l=1}^v \beta_l \left(\int \varphi_l(w) p_{\text{nu}}(w) dw \right) \\ &\quad + \frac{1}{2} \widehat{r}_{\text{item}}(w)^2 \int p_{\text{de}}(w) dw \\ &\quad - \widehat{r}_{\text{item}}(w) \int p_{\text{nu}}(w) dw, \end{aligned}$$

where $r_{\text{item}}(w)$ is replaced by $\widehat{r}_{\text{item}}(w)$, as shown in Eq. (10), and treated as a constant. Thus, the last two terms are also constants and can be safely ignored. Approximating the expectations of the first three terms in $J(\boldsymbol{\beta})$ by empirical averages, we obtain

$$\begin{aligned} \widehat{J}(\boldsymbol{\beta}) &= \frac{1}{2} \sum_{l,l'=1}^v \beta_l \beta_{l'} \left(\frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_l(w_i^{\text{de}}) \varphi_{l'}(w_i^{\text{de}}) \right) \\ &\quad + \widehat{r}_{\text{item}}(w) \sum_{l=1}^v \beta_l \left(\frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_l(w_i^{\text{de}}) \right) \\ &\quad - \sum_{l=1}^v \beta_l \left(\frac{1}{n_{\text{nu}}} \sum_{j=1}^{n_{\text{nu}}} \varphi_l(w_j^{\text{nu}}) \right), \end{aligned}$$

where the first term becomes zero unless $l = l'$ and therefore can be rewritten as

$$\begin{aligned} \widehat{J}(\boldsymbol{\beta}) &= \frac{1}{2} \sum_{l=1}^v \beta_l^2 \left(\frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_l(w_i^{\text{de}}) \right) \\ &\quad + \widehat{r}_{\text{item}}(w) \sum_{l=1}^v \beta_l \left(\frac{1}{n_{\text{de}}} \sum_{i=1}^{n_{\text{de}}} \varphi_l(w_i^{\text{de}}) \right) \\ &\quad - \sum_{l=1}^v \beta_l \left(\frac{1}{n_{\text{nu}}} \sum_{j=1}^{n_{\text{nu}}} \varphi_l(w_j^{\text{nu}}) \right). \end{aligned}$$

From the above, we obtain $\widehat{J}(\boldsymbol{\beta})$ as in Eq. (13).

Appendix C: Precision–Recall Curves

Figures A·1 and A·2 show precision–recall curves representing different forms of the rank–recall curves shown in Figs. 5 and 6, respectively. In Sect. 5, the parameters for each method were chosen based on the top 8,000 N -grams ranked in descending order of estimates. Therefore, the curves shown here are also depicted using only the top 8,000.

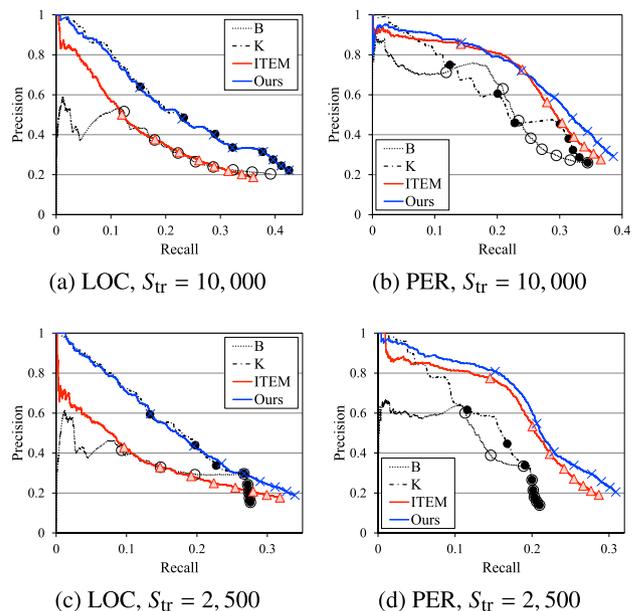


Fig. A·1 Precision–recall curves for $N = 2$. PER and LOC indicate that prediction targets are the left PER and LOC contexts, respectively. S_{tr} is the training set size. The horizontal and vertical axes of curves represent recall and precision, respectively. The markers on each curve indicate points of ranks 1,000, 2,000, \dots , 8,000. For a given rank, the method with the largest values on the horizontal and vertical axes, i.e., the highest recall and precision, has the best performance.

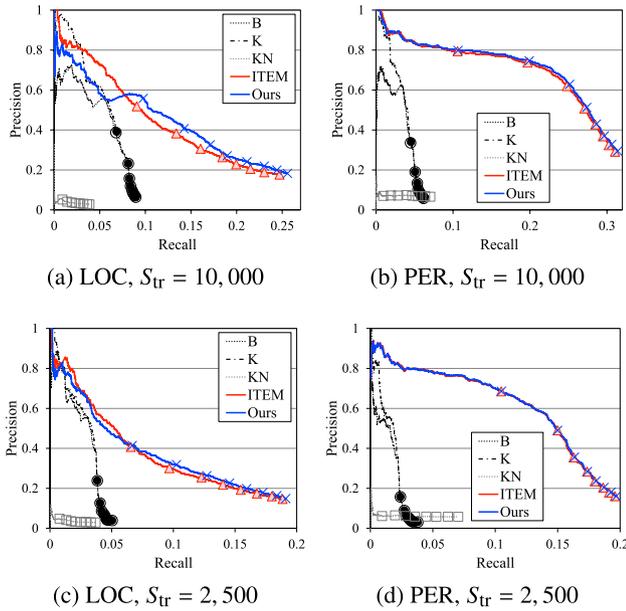
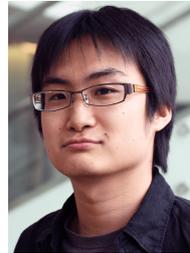


Fig. A·2 Precision–recall curves for $N = 4$. PER and LOC indicate that prediction targets are the left PER and LOC contexts, respectively. S_{tr} is the training set size. The horizontal and vertical axes of curves represent recall and precision, respectively. The markers on each curve indicate points of ranks 1,000, 2,000, . . . , 8,000. For a given rank, the method with the largest values on the horizontal and vertical axes, i.e., the highest recall and precision, has the best performance.



Kazuho Watanabe received the B.E., M.E. and Ph.D. degrees from Tokyo Institute of Technology, Japan, in 2002, 2004 and 2006, respectively. He is currently an associate professor of the Department of Computer Science and Engineering, Toyohashi University of Technology. His research interests include statistical machine learning theory and algorithms.



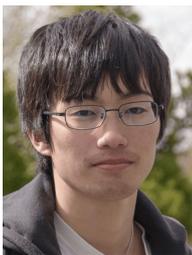
Mitsuo Yoshida received the Ph.D. degree from University of Tsukuba, Japan, in 2014. He is currently an assistant professor of the Department of Computer Science and Engineering, Toyohashi University of Technology. He is also the founder of TechTech Inc., which provides a news search engine and others. His research interests include the science of science, computational social science, and natural language processing.



Kyoji Umemura received the B.E., M.E., and Ph.D. degrees from the University of Tokyo, Japan, in 1981, 1983 and 1991, respectively. He is currently a professor of the Department of Computer Science and Engineering, Toyohashi University of Technology. His research interests include information retrieval, Lisp and symbolic computation, compiler, operating system, and natural language processing.



Masato Kikuchi received the B.E., M.E. and Ph.D. degrees from Toyohashi University of Technology, Japan, in 2016, 2018 and 2021, respectively. He is currently an assistant professor of the Department of Computer Science, Graduate School of Engineering, Nagoya Institute of Technology. His research interests include data mining, natural language processing, and web intelligence.



Kento Kawakami received the B.E. and M.E. degrees from Toyohashi University of Technology, Japan, in 2017 and 2019, respectively. His research interests include data mining, natural language processing, and network security.