

Multi-Channel Reverse Dictionary Model

Lei Zhang,^{2*†} Fanchao Qi,^{1,2*} Zhiyuan Liu,^{1,2‡} Yasheng Wang,³ Qun Liu,³ Maosong Sun^{1,2}

¹Department of Computer Science and Technology, Tsinghua University
²Institute for Artificial Intelligence, Tsinghua University
 Beijing National Research Center for Information Science and Technology
³Huawei Noah’s Ark Lab
 zhanglei9003@gmail.com, qfc17@mails.tsinghua.edu.cn
 {liuzy, sms}@tsinghua.edu.cn, {wangyasheng, qun.liu}@huawei.com

Abstract

A reverse dictionary takes the description of a target word as input and outputs the target word together with other words that match the description. Existing reverse dictionary methods cannot deal with highly variable input queries and low-frequency target words successfully. Inspired by the description-to-word inference process of humans, we propose the multi-channel reverse dictionary model, which can mitigate the two problems simultaneously. Our model comprises a sentence encoder and multiple predictors. The predictors are expected to identify different characteristics of the target word from the input query. We evaluate our model on English and Chinese datasets including both dictionary definitions and human-written descriptions. Experimental results show that our model achieves the state-of-the-art performance, and even outperforms the most popular commercial reverse dictionary system on the human-written description dataset. We also conduct quantitative analyses and a case study to demonstrate the effectiveness and robustness of our model. All the code and data of this work can be obtained on <https://github.com/thunlp/MultiRD>.

Introduction

A regular (forward) dictionary maps words to definitions while a *reverse dictionary* (Sierra 2000) does the opposite and maps descriptions to corresponding words. In Figure 1, for example, a regular dictionary tells you that “expressway” is “a wide road that allows traffic to travel fast”, and when you input “a road where cars go very quickly without stopping” to a reverse dictionary, it might return “expressway” together with other semantically similar words like “freeway”.

Reverse dictionaries have great practical value. First and foremost, they can effectively address the tip-of-the-tongue problem (Brown and McNeill 1966), which severely afflicts many people, especially those who write a lot such as researchers, writers and students. Additionally, reverse dictionaries can render assistance to new language learners who

*Indicates equal contribution

†Work done during internship at Tsinghua University

‡Corresponding author

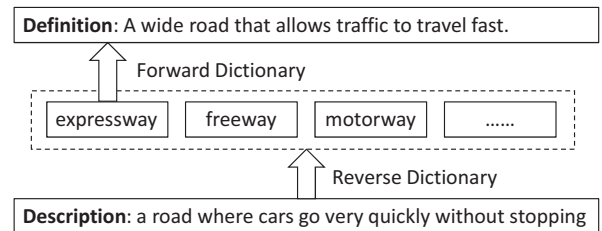


Figure 1: An example illustrating what a forward and a reverse dictionary are.

know a limited number of words. Moreover, reverse dictionaries are believed to be helpful to word selection (or word dictionary) anomia patients, people who can recognize and describe an object but fail to name the object due to neurological disorder (Benson 1979). In terms of natural language processing (NLP), reverse dictionaries can be used to evaluate the quality of sentence representations (Hill et al. 2016). They are also beneficial to the tasks involving text-to-entity mapping including question answering and information retrieval (Kartsaklis, Pilehvar, and Collier 2018)

There have been some successful commercial reverse dictionary systems such as OneLook¹, the most popular one, but their architecture is usually undisclosed proprietary knowledge. Some scientific researches into building reverse dictionaries have also been conducted. Early work adopts sentence matching based methods, which utilize hand-engineered features to find the words whose stored definitions are most similar to the input query (Bilac et al. 2004; Zock and Bilac 2004; Méndez, Calvo, and Moreno-Armendáriz 2013; Shaw et al. 2013). But these methods cannot successfully cope with the main difficulty of reverse dictionaries that human-written input queries might differ widely from target words’ definitions.

Hill et al. (2016) propose a new method based on neural language model (NLM). They employ a NLM as the sentence encoder to learn the representation of the input query, and return those words whose embeddings are closest to the

¹<https://onelook.com/thesaurus/>

input query’s representation. The NLM based reverse dictionary model alleviates the above-mentioned problem of variable input queries, but its performance is heavily dependent on the quality of word embeddings. According to Zipf’s law (Zipf 1949), however, quite a few words are low-frequency and usually have poor embeddings, which will undermine the overall performance of ordinary NLM based models.

To tackle the issue, we propose the **multi-channel reverse dictionary model**, which is inspired by the description-to-word inference process of humans. Taking “expressway” as an example, when we forget what word means “a road where cars go very quickly”, it may occur to us that the *part-of-speech tag* of the target word should be “noun” and it belongs to the *category* of “entity”. We might also guess that the target word probably contains the *morpheme* “way”. When having knowledge of these characteristics, it is much easier for us to search the target word out. Correspondingly, in our multi-channel reverse dictionary model, we employ multiple predictors to identify different characteristics of target words from input queries. By doing this, the target words with poor embeddings can still be picked out by their characteristics and, moreover, the words which have close embeddings to the correct target word but contradictory characteristics to the given description will be filtered out.

We view each characteristic predictor as an information channel of searching the target word. Two types of channels involving internal and external channels are taken into consideration. The internal channels correspond to the characteristics of words themselves including the part-of-speech (POS) tag and morpheme. The external channels reflect characteristics of target words related to external knowledge bases. We take account of two external characteristics including the word category and *sememe*. The word category information can be obtained from word taxonomy systems and it usually corresponds to the genus words of definitions. A sememe is defined as the minimum semantic unit of human languages (Bloomfield 1926), which is similar to the concept of *semantic primitive* (Wierzbicka 1996). Sememes of a word depict the meaning of the word atomically, which can be also predicted from the description of the word.

More specifically, we adopt the well-established bi-directional LSTM (BiLSTM) (Hochreiter and Schmidhuber 1997) with attention (Bahdanau, Cho, and Bengio 2015) as the basic framework and add four feature-specific characteristic predictors to it. In experiments, we evaluate our model on English and Chinese datasets including both dictionary definitions and human-written descriptions, finding that our model achieves the state-of-the-art performance. It is especially worth mentioning that for the first time OneLook is outperformed when input queries are human-written descriptions. In addition, to test our model under other real application scenarios like crossword game, we provide our model with prior knowledge about the target word such as the initial letter, and find it yields substantial performance enhancement. We also conduct detailed quantitative analyses and a case study to demonstrate the effectiveness of our model as well as its robustness in handling polysemous and low-frequency words.

Related Work

Reverse Dictionary Models

Most of existing reverse dictionary models are based on sentence-sentence matching methods, i.e., comparing the input query with stored word definitions and return the word whose definition is most similar to the input query (Zock and Bilac 2004; Bilac et al. 2004). They usually use some hand-engineered features, e.g., *tf-idf*, to measure sentence similarity, and leverage well-established information retrieval techniques to search the target word (Shaw et al. 2013). Some of them utilize external knowledge bases like WordNet (Miller 1995) to enhance sentence similarity measurement by finding synonyms or other pairs of related words between the input query and stored definitions (Méndez, Calvo, and Moreno-Armendáriz 2013; Lam and Kalita 2013; Shaw et al. 2013).

Recent years have witnessed a growing number of reverse dictionary models which conduct sentence-word matching. Thorat and Choudhari (2016) present a node-graph architecture which can directly measure the similarity between the input query and any word in a word graph. However, it works on a small lexicon (3,000 words) only. Hill et al. (2016) propose a NLM based reverse dictionary model, which uses a bag-of-words (BOW) model or an LSTM to embed the input query into the semantic space of word embeddings, and returns the words whose embeddings are closest to the representation of the input query.

Following the NLM model, Morinaga and Yamaguchi (2018) incorporate category inference to eliminate irrelevant results and achieve better performance; Kartsaklis, Pilehvar, and Collier (2018) employ a graph of WordNet synsets and words in definitions to learn target word representations together with a multi-sense LSTM to encode input queries, and they claim to deliver state-of-the-art results; Hedderich et al. (2019) use multi-sense embeddings when encoding the queries, aiming to improve sentence representations of input queries; Pilehvar (2019) adopt sense embeddings to disambiguate senses of polysemous target words.

Our multi-channel model also uses a NLM to embed input queries. Compared with previous work, our model employs multiple predictors to identify characteristics of target words, which is consistent with the inference process of humans, and achieves significantly better performance.

Applications of Dictionary Definitions

Dictionary definitions are handy resources for NLP research. Many studies utilize dictionary definitions to improve word embeddings (Noraset et al. 2017; Tissier, Gravier, and Habrard 2017; Bahdanau et al. 2017; Bosc and Vincent 2018; Scheepers, Kanoulas, and Gavves 2018). In addition, dictionary definitions are utilized in various applications including word sense disambiguation (Luo et al. 2018), knowledge representation learning (Xie et al. 2016), reading comprehension (Long et al. 2017) and knowledge graph generation (Silva, Freitas, and Handschuh 2018; Prokhorov, Pilehvar, and Collier 2019).

Methodology

In this section, we first introduce some notations. Then we describe our basic framework, i.e., BiLSTM with attention. Next we detail our multi-channel model and its two internal and two external predictors. The architecture of our model is illustrated in Figure 2.

Notations

We define \mathbb{W} as the vocabulary set, \mathbb{M} as the whole morpheme set and \mathbb{P} as the whole POS tag set. For a given word $w \in \mathbb{W}$, its morpheme set is $\mathbb{M}_w = \{m_1, \dots, m_{|\mathbb{M}_w|}\}$, where each of its morpheme $m_i \in \mathbb{M}$ and $|\cdot|$ denotes the cardinality of a set. A word may have multiple senses and each sense corresponds to a POS tag. Supposing w has n_w senses, all the POS tags of its senses form its POS tag set $\mathbb{P}_w = \{p_1, \dots, p_{n_w}\}$, where each POS tag $p_i \in \mathbb{P}$. In subsequent sections, we use lowercase boldface symbols to stand for vectors and uppercase boldface symbols for matrices. For instance, \mathbf{w} is the word vector of w and \mathbf{W} is a weight matrix.

Basic Framework

The basic framework of our model is essentially similar to a sentence classification model, composed of a sentence encoder and a classifier. We select Bidirectional LSTM (BiLSTM) (Schuster and Paliwal 1997) as the sentence encoder, which encodes an input query into a vector. Different words in a sentence have different importance to the representation of the sentence, e.g., the genus words are more important than the modifiers in a definition. Therefore, we integrate attention mechanism (Bahdanau, Cho, and Bengio 2015) into BiLSTM to learn better sentence representations.

Formally, for an input query $Q = \{q_1, \dots, q_{|Q|}\}$, we first pass the pre-trained word embeddings of its words $\mathbf{q}_1, \dots, \mathbf{q}_{|Q|} \in \mathbb{R}^d$ to the BiLSTM, where d is the dimension of word embeddings, and obtain two sequences of directional hidden states:

$$\begin{aligned} & \{\vec{\mathbf{h}}_1, \dots, \vec{\mathbf{h}}_{|Q|}\}, \{\overleftarrow{\mathbf{h}}_1, \dots, \overleftarrow{\mathbf{h}}_{|Q|}\} \\ & = \text{BiLSTM}(\mathbf{q}_1, \dots, \mathbf{q}_{|Q|}), \end{aligned} \quad (1)$$

where $\vec{\mathbf{h}}_i, \overleftarrow{\mathbf{h}}_i \in \mathbb{R}^l$ and l is the dimension of directional hidden states. Then we concatenate bi-directional hidden states to obtain non-directional hidden states:

$$\mathbf{h}_i = \text{Concatenate}(\vec{\mathbf{h}}_i, \overleftarrow{\mathbf{h}}_i). \quad (2)$$

The final sentence representation is the weighted sum of non-directional hidden states:

$$\mathbf{v} = \sum_{i=1}^{|Q|} \alpha_i \mathbf{h}_i, \quad (3)$$

where α_i is the attention item serving as the weight:

$$\begin{aligned} \alpha_i &= \mathbf{h}_i \cdot \mathbf{h}_i, \\ \mathbf{h}_t &= \text{Concatenate}(\vec{\mathbf{h}}_{|Q|}, \overleftarrow{\mathbf{h}}_1). \end{aligned} \quad (4)$$

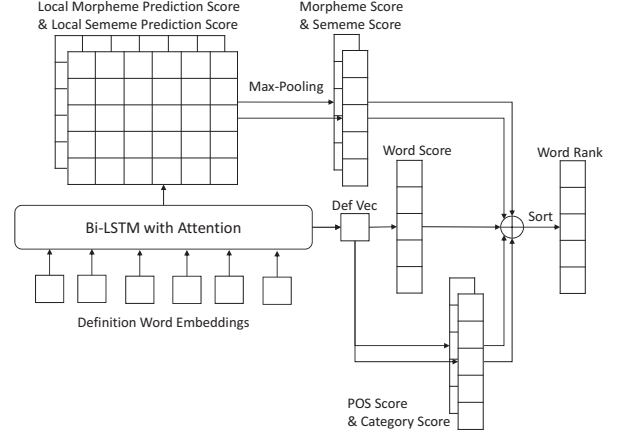


Figure 2: Multi-channel reverse dictionary model.

Next we map \mathbf{v} , the sentence vector of the input query, into the space of word embeddings, and calculate the confidence score of each word using dot product:

$$\begin{aligned} \mathbf{v}_{word} &= \mathbf{W}_{word} \mathbf{v} + \mathbf{b}_{word}, \\ sc_{w,word} &= \mathbf{v}_{word} \cdot \mathbf{w}, \end{aligned} \quad (5)$$

where $sc_{w,word}$ indicates the confidence score of w , $\mathbf{W}_{word} \in \mathbb{R}^{d \times 2l}$ is a weight matrix, $\mathbf{b}_{word} \in \mathbb{R}^d$ is a bias vector.

Internal Channel: POS Tag Predictor

A dictionary definition or human-written description of a word is usually able to reflect the POS tag of the corresponding sense of the word. We believe that predicting the POS tag of the target word can alleviate the problem of returning words with POS tags contradictory to the input query in existing reverse dictionary models.

We simply pass the sentence vector of the input query \mathbf{v} to a single-layer perceptron:

$$\mathbf{sc}_{pos} = \mathbf{W}_{pos} \mathbf{v} + \mathbf{b}_{pos}, \quad (6)$$

where $\mathbf{sc}_{pos} \in \mathbb{R}^{|\mathbb{P}|}$ records the prediction score of each POS tag, $\mathbf{W}_{pos} \in \mathbb{R}^{|\mathbb{P}| \times 2l}$ is a weight matrix, and $\mathbf{b}_{pos} \in \mathbb{R}^{|\mathbb{P}|}$ is a bias vector.

The confidence score of w from the POS tag channel is the sum of the prediction scores of w 's POS tags:

$$sc_{w,pos} = \sum_{p \in \mathbb{P}_w} [\mathbf{sc}_{pos}]_{\text{index}^{pos}(p)}, \quad (7)$$

where $[\mathbf{x}]_i$ denotes the i -th element of \mathbf{x} , and $\text{index}^{pos}(p)$ returns the POS tag index of p .

Internal Channel: Morpheme Predictor

Most words are complex words consisting of more than one morphemes. We find there exists a kind of local semantic correspondence between the morphemes of a word and its definition or description. For instance, the word “expressway” has two morphemes “express” and “way” and its dictionary definition is “a wide **road** in a city on which cars

can travel very **quickly**". We can observe that the two words "road" and "quickly" semantically correspond to the two morphemes "way" and "express" respectively. By predicting morphemes of the target word from the input query, a reverse dictionary can capture compositional information of the target word, which is complementary to contextual information of word embeddings.

We design a special morpheme predictor. Different from the POS tag predictor, we allow each hidden state to be involved in morpheme prediction directly, and do max-pooling to obtain final morpheme prediction scores. Specifically, we feed each non-directional hidden state to a single-layer perceptron and obtain local morpheme prediction scores:

$$\mathbf{sc}_{mor}^i = \mathbf{W}_{mor} \mathbf{h}_i + \mathbf{b}_{mor}, \quad (8)$$

where $\mathbf{sc}_{mor}^i \in \mathbb{R}^{|\mathbb{M}|}$ measures the semantic correspondence between i -th word in the input query and each morpheme, $\mathbf{W}_{mor} \in \mathbb{R}^{|\mathbb{M}| \times 2l}$ is a weight matrix, and $\mathbf{b}_{mor} \in \mathbb{R}^{|\mathbb{M}|}$ is a bias vector. Then we do max-pooling over all the local morpheme prediction scores to obtain global morpheme prediction scores:

$$[\mathbf{sc}_{mor}]_j = \max_{1 \leq i \leq |Q|} [\mathbf{sc}_{mor}^i]_j. \quad (9)$$

And the confidence score of w from the morpheme channel is:

$$sc_{w,mor} = \sum_{m \in \mathbb{M}_w} [\mathbf{sc}_{mor}]_{\text{index}^{mor}(m)}, \quad (10)$$

where $\text{index}^{mor}(m)$ returns the morpheme index of m .

External Channel: Word Category Predictor

Semantically related words often belong to different categories, although they have close word embeddings, e.g., "car" and "road". Word category information is helpful in eliminating semantically related but not similar words from the results of reverse dictionaries (Morinaga and Yamaguchi 2018). There are many available word taxonomy systems which can provide hierarchical word category information, e.g., WordNet (Miller 1995). Some of them provides POS tag information as well, in which case POS tag predictor can be removed.

We design a hierarchical predictor to calculate prediction scores of word categories. Specifically, each word belongs to a certain category in each layer of word hierarchy. We first compute the word category prediction score of each layer:

$$\mathbf{sc}_{cat,k} = \mathbf{W}_{cat,k} \mathbf{v} + \mathbf{b}_{cat,k}, \quad (11)$$

where $\mathbf{sc}_{cat,k} \in \mathbb{R}^{c_k}$ is the word category prediction score distribution of k -th layer, $\mathbf{W}_{cat,k} \in \mathbb{R}^{c_k \times 2l}$ is a weight matrix, $\mathbf{b}_{cat,k} \in \mathbb{R}^{c_k}$ is a bias vector, and c_k is the category number of k -th layer. Then the final confidence score of w from the word category channel is the weighted sum of its category prediction scores of all the layers:

$$sc_{w,cat} = \sum_{k=1}^K \beta_k [\mathbf{sc}_{cat,k}]_{\text{index}_k^{cat}(w)}, \quad (12)$$

where K is the total layer number of the word hierarchy, β_k is a hyper-parameter controlling the relative weights, and $\text{index}_k^{cat}(w)$ returns the category index of w in the k -th layer.

External Channel: Sememe Predictor

In linguistics, a sememe is the minimum semantic unit of natural languages (Bloomfield 1926). Sememes of a word can accurately depict the meaning of the word. HowNet (Dong and Dong 2003) is the most famous sememe knowledge base. It defines about 2,000 sememes and uses them to annotate more than 100,000 Chinese and English words by hand. HowNet and its sememe knowledge has been widely applied to various NLP tasks including sentiment analysis (Fu et al. 2013), word representation learning (Niu et al. 2017), semantic composition (Qi et al. 2019a), sequence modeling (Qin et al. 2019) and textual adversarial attack (Zang et al. 2019).

Sememe annotation of a word in HowNet includes hierarchical sememe structures as well as relations between sememes. For simplicity, we extract a set of unstructured sememes for each word, in which case sememes of a word can be regarded as multiple semantic labels of the word. We find there also exists local semantic correspondence between the sememes of a word and its description. Still taking "expressway" as an example, its annotated sememes in HowNet are *route* and *fast*, which semantically correspond to the words in its definition "road" and "quickly" respectively.

Therefore, we design a sememe predictor similar to the morpheme predictor. Formally, we use \mathbb{S} to represent the set of all sememes. The sememe set of a word w is $\mathbb{S}_w = \{s_1, \dots, s_{|\mathbb{S}_w|}\}$. We pass each hidden state to a single-layer perceptron to calculate local sememe prediction scores:

$$\mathbf{sc}_{sem}^i = \mathbf{W}_{sem} \mathbf{h}_i + \mathbf{b}_{sem}, \quad (13)$$

where $\mathbf{sc}_{sem}^i \in \mathbb{R}^{|\mathbb{S}|}$ indicates how corresponding between i -th word in the input query and each sememe, $\mathbf{W}_{sem} \in \mathbb{R}^{|\mathbb{S}| \times 2l}$ is a weight matrix, and \mathbf{b}_{sem} is a bias vector. Final sememe prediction scores are computed by doing max-pooling:

$$[\mathbf{sc}_{sem}]_j = \max_{1 \leq i \leq |Q|} [\mathbf{sc}_{sem}^i]_j. \quad (14)$$

The confidence score of w from the sememe channel is:

$$sc_{w,sem} = \sum_{s \in \mathbb{S}_w} [\mathbf{sc}_{sem}]_{\text{index}^{sem}(s)}, \quad (15)$$

where $\text{index}^{sem}(s)$ returns the sememe index of s .

Multi-channel Reverse Dictionary Model

By combining the confidence scores of direct word prediction and indirect characteristic prediction, we obtain the final confidence score of a given word w in our multi-channel reverse dictionary model:

$$sc_w = \lambda_{word} sc_{w,word} + \sum_{c \in \mathbb{C}} \lambda_c sc_{w,c}, \quad (16)$$

where $\mathbb{C} = \{pos, mor, cat, sem\}$ is the channel set, and λ_{word} and λ_c are the hyper-parameters controlling relative weights of corresponding terms.

As for training loss, we simply adopt the one-versus-all cross-entropy loss inspired by the sentence classification models.

Model	Seen Definition			Unseen Definition			Description		
OneLook	0	.66/.94/.95	200	-	-	-	5.5	.33 /.54/.76	332
BOW	172	.03/.16/.43	414	248	.03/.13/.39	424	22	.13/.41/.69	308
RNN	134	.03/.16/.44	375	171	.03/.15/.42	404	17	.14/.40/.73	274
RDWECI	121	.06/.20/.44	420	170	.05/.19/.43	420	16	.14/.41/.74	306
SuperSense	378	.03/.15/.36	462	465	.02/.11/.31	454	115	.03/.15/.47	396
MS-LSTM	0	.92 /.98/.99	65	276	.03/.14/.37	426	1000	.01/.04/.18	404
BiLSTM	25	.18/.39/.63	363	101	.07/.24/.49	401	5	.25/.60/.83	214
+Mor	24	.19/.41/.63	345	80	.08/.26/.52	399	4	.26/.62/.85	198
+Cat	19	.19/.42/.68	309	68	.08/.28/.54	362	4	.30/.62/.85	206
+Sem	19	.19/.43/.66	349	80	.08/.26/.53	393	4	.30/. 64 /.87	218
Multi-channel	16	.20/.44/.71	310	54	.09 /.29/.58	358	2	.32/. 64 /.88	203
				<i>median rank</i>		<i>accuracy@1/10/100</i>		<i>rank variance</i>	

Table 1: Overall reverse dictionary performance of all the models.

Experiments

In this section, we evaluate the performance of our multi-channel reverse dictionary model. We also conduct detailed quantitative analyses as well as a case study to explore the influencing factors in the reverse dictionary task and demonstrate the strength and weakness of our model. We carry out experiments on both English and Chinese datasets. But due to limited space, we present our experiments on the Chinese dataset in the appendix.

Dataset

We use the English dictionary definition dataset created by Hill et al. (2016)² as the training set. It contains about 100,000 words and 900,000 word-definition pairs. We have three test sets including: (1) **seen** definition set, which contains 500 pairs of words and WordNet definitions existing in the training set and is used to assess the ability to recall previously encoded information; (2) **unseen** definition set, which also contains 500 pairs of words and WordNet definitions but the words together with all their definitions have been excluded from the training set; and (3) **description set**, which consists of 200 pairs of words and human-written descriptions and is a benchmark dataset created by Hill et al. (2016) too.

To obtain the morpheme information our model needs, we use Morfessor (Virpioja et al. 2013) to segment all the words into morphemes. As for the word category information, we use the *lexical names* from WordNet (Miller 1995). There are 45 lexical names and the total layer number of the word category hierarchy is 1. Since the lexical names have included POS tags, e.g., *noun.animal*, we remove the POS tag predictor from our model. We use HowNet as the source of sememes. It contains 43,321 English words manually annotated with 2,148 different sememes in total. We employ OpenHowNet (Qi et al. 2019b), the open data accessing API of HowNet, to obtain sememes of words.

²The definitions are extracted from five electronic resources: *WordNet*, *The American Heritage Dictionary*, *The Collaborative International Dictionary of English*, *Wiktionary* and *Webster’s*.

Experimental Settings

Baseline Methods We choose the following models as the baseline methods: (1) OneLook, the most popular commercial reverse dictionary system, whose 2.0 version is used; (2) BOW and RNN with rank loss (Hill et al. 2016), both of which are NLM based and the former uses a bag-of-words model while the latter uses an LSTM; (3) RDWECI (Morinaga and Yamaguchi 2018), which incorporates category inference and is an improved version of BOW; (4) SuperSense (Pilehvar 2019), an improved version of BOW which uses pretrained sense embeddings to substitute target word embeddings; (5) MS-LSTM (Kartsaklis, Pilehvar, and Collier 2018), an improved version of RNN which uses graph-based WordNet synset embeddings together with a multi-sense LSTM to predict *synsets* from descriptions and claims to produce state-of-the-art performance; and (6) BiLSTM, the basic framework of our multi-channel model.

Hyper-parameters and Training For our model, the dimension of non-directional hidden states is 300×2 , the weights of different channels are equally set to 1, and the dropout rate is 0.5. For all the models except MS-LSTM, we use the 300-dimensional word embeddings pretrained on GoogleNews with *word2vec*³, and the word embeddings are fixed during training. For all the other baseline methods, we use their recommended hyper-parameters. For training, we adopt Adam as the optimizer with initial learning rate 0.001, and the batch size is 128.

Evaluation Metrics Following previous work, we use three evaluation metrics: the median rank of target words (lower better), the accuracy that target words appear in top 1/10/100 ($\text{acc}@1/10/100$, higher better) and the standard deviation of target words’ ranks (rank variance, lower better). Notice that MS-LSTM can only predict WordNet synsets. Thus, we map the target words to corresponding WordNet synsets (target synsets) and calculate the accuracy and rank variance of the target synsets.

³<https://code.google.com/archive/p/word2vec/>

Prior Knowledge	Seen Definition			Unseen Definition			Description		
None	16	.20/.44/.71	310	54	.09/.29/.58	358	2.5	.32/.64/.88	203
POS Tag	13	.21/.45/.72	290	45	.10/.31/.60	348	3	.35/.65/.91	174
Initial Letter	1	.39/.73/.90	270	4	.26/.63/.85	348	0	.62/.90/.97	160
Word Length	1	.40/.71/.90	269	6	.25/.56/.84	346	0	.55/.85/.95	163
				<i>median rank</i>		<i>accuracy@1/10/100</i>		<i>rank variance</i>	

Table 2: Reverse dictionary performance with prior knowledge.

Overall Experimental Results

Table 1 exhibits reverse dictionary performance of all the models on the three test sets, where “Mor”, “Cat” and “Sem” represent the morpheme, word category and sememe predictors respectively. Notice that the performance of OneLook on the unseen test set is meaningless because we cannot exclude any definitions from its definition bank, hence we do not list corresponding results. From the table, we can see:

(1) Compared with all the baseline methods other than OneLook, our multi-channel model achieves substantially better performance on the unseen definition set and the description set, which verifies the absolute superiority of our model in generalizing to the novel and unseen input queries.

(2) OneLook significantly outperforms our model when the input queries are dictionary definitions. This result is expected because the input dictionary definitions are already stored in the database of OneLook and even simple text matching can easily handle this situation. However, the input queries of a reverse dictionary cannot be exact dictionary definitions in reality. On the description test set, our multi-channel model achieves better overall performance than OneLook. Although OneLook yields slightly higher acc@1, it has limited value in terms of practical application, because people always need to pick the proper word from several candidates, not to mention the fact that the acc@1 of OneLook is only 0.33.

(3) MS-LSTM performs very well on the seen definition set but badly on the description set, which manifests its limited generalization ability and practical value. Notice that when testing MS-LSTM, the searching space is the whole synset list rather than the synset list of the test set, which causes the difference in performance on the unseen definition set measured by us and recorded in the original work (Kartsaklis, Pilehvar, and Collier 2018).

(4) All the BiLSTM variants enhanced with different information channels (+Mor, +Cat and +Sem) perform better than vanilla BiLSTM. These results prove the effectiveness of predicting characteristics of target words in the reverse dictionary task. Moreover, our multi-channel model achieves further performance enhancement as compared with the single-channel models, which demonstrates the potency of characteristic fusion and also verifies the efficacy of our multi-channel model.

(5) BOW performs better than RNN, which is consistent with the findings from Hill et al. (2016). However, BiLSTM far surpasses BOW as well as RNN. This verifies the necessity for bi-directional encoding in RNN models, and also shows the potential of RNNs.

Performance with Prior Knowledge

In practical application of reverse dictionaries, extra information about target words in addition to descriptions may be known. For example, we may remember the initial letter of the word we forget, or the length of the target word is known in crossword game. In this subsection, we evaluate the performance of our model with the prior knowledge of target words, including POS tag, initial letter and word length. More specifically, we extract the words satisfying given prior knowledge from the top 1,000 results of our model, and then reevaluate the performance. The results are shown in Table 2.

We can find that any prior knowledge improves the performance of our model to a greater or lesser extent, which is an expected result. However, the performance boost brought by the initial letter and word length information is much bigger than that brought by the POS tag information. The possible reasons are as follows. For the POS tag, it has been already predicted in our multi-channel model, hence the improvement it brings is limited, which also demonstrates that our model can do well in POS tag prediction. For the initial letter and word length, they are hard to predict according to a definition or description and not considered in our model. Therefore, they can filter many candidates out and markedly increase performance.

Analyses of Influencing Factors

In this subsection, we conduct quantitative analyses of the influencing factors in reverse dictionary performance. To make results more accurate, we use a larger test set consisting of 38,113 words and 80,658 **seen** pairs of words and WordNet definitions. Since we are interested in the features of target words, we exclude MS-LSTM that predicts WordNet synsets.

Sense Number Figure 3 exhibits the acc@10 of all the models on the words with different numbers of senses. It is obvious that performance of all the models declines with the increase in the sense number, which indicates that polysemy is a difficulty in the task of reserve dictionary. But our model displays outstanding robustness and its performance hardly deteriorates even on the words with the most senses.

Word Frequency Figure 4 displays all the models’ performance on the words within different ranges of word frequency ranking. We can find that the most frequent and infrequent words are harder to predict for all the reverse dictionary models. The most infrequent words usually have




Figure 3: Acc@10 on words with different sense numbers. The numbers of words are 21, 582, 8, 266, 3, 538, 1, 691, 953 and 2, 083 respectively.

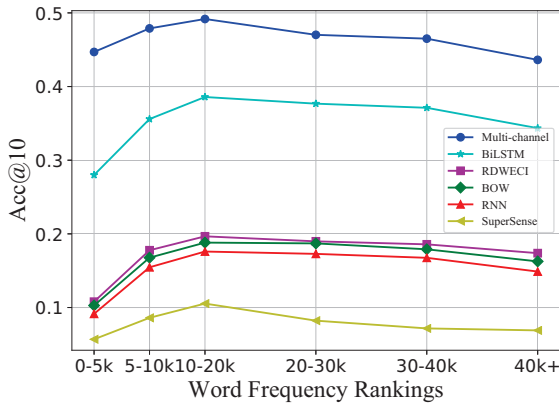


Figure 4: Acc@10 on ranges of different word frequency rankings. The number of words in each range is 3, 299, 3, 243, 3, 515, 3, 300, 5, 565 and 19, 191 respectively.

poor embeddings, which may damage the performance of NLM based models. For the most frequent words, on the other hand, although their embeddings are better, they usually have more senses. We count the average sense numbers of all the ranges, which are 5.6, 3.2, 2.6, 2.1, 1.7 and 1.4 respectively. The first range has a much larger average sense number, which explains its bad performance. Moreover, our model also demonstrates remarkable robustness.

Query Length The effect of query length on reverse dictionary performance is illustrated in Figure 5. When the input query has only one word, the system performance is strikingly poor, especially our multi-channel model. This is easy to explain because the information extracted from the input query is too limited. In this case, outputting the synonyms of the query word is likely to be a better choice.




Figure 5: Acc@10 on ranges of different query length. The number of queries in each range is 672, 3, 609, 21, 113, 31, 013, 14, 684, 5, 803, 2, 316 and 1, 448 respectively.

Word	Input Query
postnuptial	relating to events after a marriage.
takeaway	concession made by a labor union to a company.

Table 3: Two reverse dictionary cases.

Case study

In this subsection, we give two cases in Table 3 to display the strength and weakness of our reverse dictionary model. For the first word “postnuptial”, our model correctly predicts its morpheme “post” and sememe “GetMarried” from the words “after” and “marriage” in the input query. Therefore, our model easily finds the correct answer. For the second case, the input query describes a rare sense of the word “takeaway”. HowNet has no sememe annotation for this sense, and morphemes of the word are not semantically related to any words in the query either. Our model cannot solve this kind of cases, which is in fact hard to handle for all the NLM based models. In this situation, the text matching methods, which return the words whose stored definitions are most similar to the input query, may help.

Conclusion and Future Work

In this paper, we propose a multi-channel reverse dictionary model, which incorporates multiple predictors to predict characteristics of target words from given input queries. Experimental results and analyses show that our model achieves the state-of-the-art performance and also possesses outstanding robustness.

In the future, we will try to combine our model with text matching methods to better tackle extreme cases, e.g., single-word input query. In addition, we are considering extending our model to the cross-lingual reverse dictionary task. Moreover, we will explore the feasibility of transferring our model to related tasks such as question answering.

Acknowledgements

This work is funded by the Natural Science Foundation of China (NSFC) and the German Research Foundation (DFG) in Project Crossmodal Learning, NSFC 61621136008 / DFG TRR-169. Furthermore, we thank the anonymous reviewers for their valuable comments and suggestions.

References

- Bahdanau, D.; Bosc, T.; Jastrzebski, S.; Grefenstette, E.; Vincent, P.; and Bengio, Y. 2017. Learning to compute word embeddings on the fly. *arXiv preprint arXiv:1706.00286*.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of ICLR*.
- Benson, D. F. 1979. Neurologic correlates of anomia. In *Studies in neurolinguistics*. Elsevier. 293–328.
- Bilac, S.; Watanabe, W.; Hashimoto, T.; Tokunaga, T.; and Tanaka, H. 2004. Dictionary search based on the target word description. In *Proceedings of NLP*.
- Bloomfield, L. 1926. A set of postulates for the science of language. *Language* 2(3):153–164.
- Bosc, T., and Vincent, P. 2018. Auto-encoding dictionary definitions into consistent word embeddings. In *Proceedings of EMNLP*.
- Brown, R., and McNeill, D. 1966. The “tip of the tongue” phenomenon. *Journal of verbal learning and verbal behavior* 5(4):325–337.
- Dong, Z., and Dong, Q. 2003. Hownet-a hybrid language and knowledge resource. In *Proceedings of NLP-KE*.
- Fu, X.; Liu, G.; Guo, Y.; and Wang, Z. 2013. Multi-aspect sentiment analysis for chinese online social reviews based on topic modeling and hownet lexicon. *Knowledge-Based Systems* 37:186–195.
- Hedderich, M. A.; Yates, A.; Klakow, D.; and de Melo, G. 2019. Using multi-sense vector embeddings for reverse dictionaries. In *Proceedings of IWCS*.
- Hill, F.; Cho, K.; Korhonen, A.; and Bengio, Y. 2016. Learning to understand phrases by embedding the dictionary. *TKDE* 4:17–30.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Kartsaklis, D.; Pilehvar, M. T.; and Collier, N. 2018. Mapping text to knowledge graph entities using multi-sense lstms. In *Proceedings of EMNLP*.
- Lam, K. N., and Kalita, J. K. 2013. Creating reverse bilingual dictionaries. In *Proceedings of HLT-NAACL*.
- Long, T.; Bengio, E.; Lowe, R.; Cheung, J. C. K.; and Precup, D. 2017. World knowledge for reading comprehension: Rare entity prediction with hierarchical lstms using external descriptions. In *Proceedings of EMNLP*.
- Luo, F.; Liu, T.; Xia, Q.; Chang, B.; and Sui, Z. 2018. Incorporating glosses into neural word sense disambiguation. In *Proceedings of ACL*.
- Méndez, O.; Calvo, H.; and Moreno-Armendáriz, M. A. 2013. A reverse dictionary based on semantic analysis using wordnet. In *Proceedings of MICAI 2013*.
- Miller, G. A. 1995. Wordnet: a lexical database for english. *Communications of the Acm* 38(11):39–41.
- Morinaga, Y., and Yamaguchi, K. 2018. Improvement of reverse dictionary by tuning word vectors and category inference. In *Proceedings of ICIST*.
- Niu, Y.; Xie, R.; Liu, Z.; and Sun, M. 2017. Improved word representation learning with sememes. In *Proceedings of ACL*.
- Noraset, T.; Liang, C.; Birnbaum, L.; and Downey, D. 2017. Definition modeling: Learning to define word embeddings in natural language. In *Proceedings of AAAI*.
- Pilehvar, M. T. 2019. On the importance of distinguishing word meaning representations: A case study on reverse dictionary mapping. In *Proceedings of NAACL-HLT*.
- Prokhorov, V.; Pilehvar, M. T.; and Collier, N. 2019. Generating knowledge graph paths from textual definitions using sequence-to-sequence models. In *Proceedings NAACL-HLT*.
- Qi, F.; Huang, J.; Yang, C.; Liu, Z.; Chen, X.; Liu, Q.; and Sun, M. 2019a. Modeling semantic compositionality with sememe knowledge. In *Proceedings of ACL*.
- Qi, F.; Yang, C.; Liu, Z.; Dong, Q.; Sun, M.; and Dong, Z. 2019b. Openhownet: An open sememe-based lexical knowledge base. *arXiv preprint arXiv:1901.09957*.
- Qin, Y.; Qi, F.; Ouyang, S.; Liu, Z.; Yang, C.; Wang, Y.; Liu, Q.; and Sun, M. 2019. Enhancing recurrent neural networks with sememes. *arXiv preprint arXiv:1910.08910*.
- Scheepers, T.; Kanoulas, E.; and Gavves, E. 2018. Improving word embedding compositionality using lexicographic definitions. In *Proceedings of WWW*.
- Schuster, M., and Paliwal, K. K. 1997. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing* 45(11):2673–2681.
- Shaw, R.; Datta, A.; VanderMeer, D. E.; and Dutta, K. 2013. Building a scalable database-driven reverse dictionary. *TKDE* 25:528–540.
- Sierra, G. 2000. The onomasiological dictionary: a gap in lexicography. In *Proceedings of the ninth Euralex international congress*.
- Silva, V.; Freitas, A.; and Handschuh, S. 2018. Building a knowledge graph from natural language definitions for interpretable text entailment recognition. In *Proceedings LREC*.
- Thorat, S., and Choudhari, V. 2016. Implementing a reverse dictionary, based on word definitions, using a node-graph architecture. In *Proceedings of COLING*.
- Tissier, J.; Gravier, C.; and Habrard, A. 2017. Dict2vec : Learning word embeddings using lexical dictionaries. In *Proceedings of EMNLP*.
- Virpioja, S.; Smit, P.; Grönroos, S. A.; and Kurimo, M. 2013. Morfessor 2.0: Python implementation and extensions for morfeessor baseline. *Aalto University Publication*.
- Wierzbicka, A. 1996. *Semantics: Primes and universals: Primes and universals*. Oxford University Press, UK.
- Xie, R.; Liu, Z.; Jia, J.; Luan, H.; and Sun, M. 2016. Representation learning of knowledge graphs with entity descriptions. In *Proceedings of AAAI*.
- Zang, Y.; Yang, C.; Qi, F.; Liu, Z.; Zhang, M.; Liu, Q.; and Sun, M. 2019. Textual adversarial attack as combinatorial optimization. *arXiv preprint arXiv:1910.12196*.
- Zipf, G. K. 1949. Human behavior and the principle of least effort. *SERBIULA (sistema Librum 2.0)*.
- Zock, M., and Bilac, S. 2004. Word lookup on the basis of associations: from an idea to a roadmap. In *Proceedings of the Workshop on Enhancing and Using Electronic Dictionaries*.