

GEAR: A Simple GENERATE, EMBED, AVERAGE AND RANK Approach for Unsupervised Reverse Dictionary

Fatemah Almeman^{1,2}, Luis Espinosa-Anke^{1,3}

¹CardiffNLP, School of Computer Science and Informatics, Cardiff University, UK

²College of Computer and Information Sciences, Princess Nourah bint Abdulrahman University, KSA

³AMPLYFI, UK

Correspondence: almemanf@cardiff.ac.uk

Abstract

Reverse Dictionary (RD) is the task of obtaining the most relevant word or set of words given a textual description or dictionary definition. Effective RD methods have applications in accessibility, translation or writing support systems. Moreover, in NLP research we find RD to be used to benchmark text encoders at various granularities, as it often requires word, definition and sentence embeddings. In this paper, we propose a simple approach to RD that leverages LLMs in combination with embedding models. Despite its simplicity, this approach outperforms supervised baselines in well studied RD datasets, while also showing less over-fitting. We also conduct a number of experiments on different dictionaries and analyze how different styles, registers and target audiences impact the quality of RD systems. We conclude that, on average, untuned embeddings alone fare way below an LLM-only baseline (although they are competitive in highly technical dictionaries), but are crucial for boosting performance in combined methods.¹

1 Introduction

Reverse dictionary (RD), conceptual dictionary or concept lookup is the task to returning a word or set of suitable words given a text description or definition (Zock, 2004; Zock and Bilac, 2004). This NLP task broadly equates to the psycholinguistic notion of “lexical access” and “tip of the tongue problem” (Brown and McNeill, 1966), and is crucial for better understanding of the mental lexicon, i.e., how we as humans store meaning and render semantic representations into words (Wanner, 1996; Zock et al., 2010). In practical terms, RD systems are helpful assistants to writers and translators (Yan et al., 2020), and can have significant impact on the linguistic experience of language learners

with limited vocabulary (Zhang et al., 2020) or people affected with anomic aphasia (Benson, 1979). From an NLP perspective, RD serves as a tool for fine-tuning and evaluating text encoders (Hill et al., 2016; Pilehvar, 2019; Zhang et al., 2020; Chen and Zhao, 2022) and, recently, as a probe for understanding Large Language Models’ (LLMs) internal representations (Xu et al., 2024). Despite its usefulness, research in RD is currently limited in two fronts. First, RD benchmarks are mostly sourced from WordNet (Miller, 1995) and the Oxford Dictionary, and little is known about the effectiveness of RD methods on other resources or languages - with a few notable exceptions such as the multilingual experiments in Yan et al. (2020). This is problematic because the generalization ability of models optimized for these two standard and over-utilized resources might not reflect modern, acquired, rare, evolving or technical terminologies. And second, because there is a surprising lack of work exploiting the generative capabilities of LLMs to improve over embedding-only baselines. While Tian et al. (2024) propose to leverage LLMs for RD using prompt engineering and obtain good results, their approach requires fine-tuning a text generation model in the first stage, and the final set of predicted terms may not correspond to the vocabulary of the dictionary in question, making this approach hard to apply on large-scale real-world resources².

In this paper, we make the following contributions. First, we propose GEAR (generate, embed, average and rank), a novel lightweight and unsupervised method for RD that utilizes an LLM for generating a set of candidates given an input definition, and pools their corresponding embeddings

¹The code and data are available at https://github.com/F-Almeman/GEAR_RD

²At the time of writing this manuscript, the English Wiktionary has over 7.5M entries (with over 30M entries across all languages), making embedding search a prerequisite on any realistic RD method. <https://en.wikipedia.org/wiki/Wiktionary>

into a vector used for KNN search. This fast and highly scalable method outperforms heavily tuned supervised baselines (including those leveraging LLMs), setting a new state of the art in two (the two that require generalization) out of three test sets in Hill’s dataset (Hill et al., 2016). We also evaluate GEAR, alongside other unsupervised baselines, on a diverse set of dictionaries, from WordNet to Urban Dictionary³, and perform an in-depth analysis of the strengths and weaknesses of these methods as per the domain, register and target audience of each dictionary.

2 Related Work

This paper is relevant to computational semantics, and more concretely, on the interaction between dictionaries and NLP. Therefore, we cover, first, prominent work in this intersection (Section 2.1), and second, we give an account of RD works (Section 2.2).

2.1 Dictionaries and NLP

Dictionaries and NLP have a healthy relationship. Indeed, dictionaries have proven to be suitable resources for improving text processing pipelines at different stages. For example, for improving Word Sense Disambiguation (WSD) systems based on BERT (Devlin et al., 2018) by fine-tuning them on *context-definition* pairs (Huang et al., 2019), or by combining into one loss function different pre-training strategies, e.g., discriminating between correct and wrong definitions via contrastive learning, or dictionary entry prediction (Yu et al., 2022; Chen et al., 2022). Beyond BERT, Bevilacqua et al. (2020) fine-tuned BART (Lewis et al., 2019) on example-definition pairs, and reported high results in intrinsic benchmarks and, more importantly, used their DM system for downstream NLP, specifically WSD as well as word-in-context (Pilehvar and Camacho-Collados, 2019) classification. This strategy was further adopted to add an interpretability layer to semantic change detection via definition generation (Giulianelli et al., 2023). Other examples of successfully marrying dictionaries and NLP systems include training a unified vector space of words, definitions and “mentions in context” (Gajbhiye et al., 2024), which were successfully used for ontology completion; or fine-tuning an LLM such as LLAMA2 (Touvron et al., 2023) on WordNet’s semantic relations and definitions, which then can

be flexibly used in several lexical semantics tasks (Moskvoretskii et al., 2024).

2.2 Reverse Dictionary

Concerning RD, this is a task with a long tradition in lexical semantics, with early methods exploiting hand-crafted rules (Bila et al., 2004; Shaw et al., 2011) for extracting textual features. It was Hill et al. (2016) who introduced RNNs as suitable architectures that complemented bag of words representations, as well as a dataset specific to RD sourced from different resources, such as WordNet, Webster’s Dictionary and Wiktionary, among others. From here, the usage of neural networks first, and more specific, pre-trained transformer encoders like BERT later, have dominated the RD landscape. Among the former, let us highlight, e.g., Pilehvar (2019), who integrates WordNet senses and supersenses as an additional signal, improving over text-based embeddings alone. Further, Zhang et al. (2020) propose a multi-channel model comprising a sentence encoder based on BiLSTMs and multiple linguistically motivated predictors such as word category (using WordNet’s taxonomy), morpheme or sememe prediction, whereas Chen and Su (2021) directly replaces word embeddings with synset embeddings, optionally leveraging examples of usage. An immediate limitation of the above works is their reliance on a sense inventory such as WordNet, which has proven to work very well for modeling in-domain terminologies, less so for enabling generalization (cf. Table 2).

Other works have exploited multiple but related tasks in the broad “embedding a dictionary” paradigm, e.g., by combining definition generation and RD with reconstruction tasks via autoencoders (Chen and Zhao, 2022), or have fine-tuned T5 (Raffel et al., 2020) with excellent results (Mane et al., 2022). More recently, LLMs have unexpectedly been introduced into RD. For example, in a two-stage approach where a fine-tuned LLM first generates a set of candidates which are then passed in a subsequent prompt to a generator for outputting the final set of predictions (Tian et al., 2024). Finally, from a more “probing” perspective, RD has been used to gain insights into LLMs representations via conceptual inference, showing that they encode information about object categories as well as fine-grained features (Xu et al., 2024).

As we can see from the above survey, dictionaries play an important role in NLP today, and RD is a suitable probe for pre-training and eval-

³<https://www.urbandictionary.com/>

uating text encoders as well as LLMs. However, as we mentioned in Section 1, there has been little work in expanding existing benchmarks beyond Hill’s dataset, especially in terms of different domains and registers. Moreover, no works have explored the seemingly simple generate-then-embed approach, so that many of the practical drawbacks of LLMs (hallucinations and context length limitations, to name a few) could be alleviated, but still using their ability to generate suitable candidate embeddings for KNN search.

3 The GEAR Method

In this section, we give a brief description of GEAR, a novel, *very* simple, lightweight, and, more importantly, unsupervised method for RD. We denote any dictionary as $D = \{(d_i, T_i) | i = 1, \dots, N\}$, where d_i is a definition, $T_i = \{t_{i1}, t_{i2}, \dots, t_{ik_i}\}$ is the set of corresponding terms⁴ (i.e., entries in the dictionary), and $k_i \geq 1$ the number of terms associated with d_i . From here, GEAR consists on four simple steps. First, **generate**, where, given an input definition d_i , an (LLM) generates a set of possible terms $G = \{g_1, g_2, \dots, g_m\}$. All throughout this paper, we use GPT-4O-MINI (Achiam et al., 2023)⁵, which we prompt in three different ways to thoroughly explore any differences in performance that might arise from deeper descriptions of the task. The full details for these prompts can be found in Appendix A), however, at a high level, they can be summarized as follows:

- **Base prompt 1 (bp1)**: includes a short description of the resource: such as *Given the definition {definition}, generate a ranked list of {k} terms, with the first term being the most related to the definition, assuming they are from the {resource} dictionary. {resource} is {resource description}*.
- **Base prompt 2 (bp2)**: the same as bp1 but it also includes a sample of terms and definitions from the specified resource to help the model understand the type of terms it should generate: such as *These are some examples of definitions and terms in this dictionary: {examples}*.

⁴It is important to account for a one-to-many relationship at this point because we will be reporting experiments on combinations of multiple resources.

⁵<https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>.

- **Reasoning prompt (rp)**: this final prompt extends bp2 by requesting the generation of examples alongside the terms to explore whether having the LLM ‘reason’ before answering leads to improved results. Specifically, this part is added: *For each term, provide an example usage in a sentence that matches the style and scope of {dictionary}*.

In the next embed step, a text encoder $f : \mathcal{V} \rightarrow R^n$ maps each term in G to a vector representation in an n -dimensional space. We use SBERT (Reimers and Gurevych, 2019) and the Instructor model (Su et al., 2022) to obtain term embeddings and evaluate their performance for comparison (see Section 6). The resulting matrix $\mathbf{E}_G = [e_1, e_2, \dots, e_m]^T \in R^{m \times n}$, where $e_i = f(g_i)$, is then mean pooled (or averaged) as follows: $\bar{e} = \frac{1}{m} \sum_{i=1}^m e_i$. Finally, in the rank step, given $T = \bigcup_{i=1}^N T_i$, which denotes the set of all unique terms in D , we perform KNN search via cosine similarity over T with \bar{e} . The performance of GEAR, just like any other search-based approach, can be evaluated using Information Retrieval metrics that account for different scenarios, e.g., the rank of the first correct term with Mean Reciprocal Rank (MRR), or the proportion of correct terms at different cutoffs with Precision at k (P@k).

4 Data

In this paper we are concerned not only with exploring the usefulness of GEAR when compared to existing baselines, we are also interested in developing an understanding of what kind of lexical resource poses greater challenges to this method as well as its components alone. For this reason, we perform experiments on two distinct but complementary datasets.

As a first evaluation set, we use the three test sets from the Hill et al. 2016 dataset (Section 5.1) to compare GEAR with other published RD methods: the seen set, which includes 500 word-definition pairs from the training set to evaluate recall; the unseen set, containing 500 pairs where both the words and definitions are excluded from training; and the description set, consisting of 200 words with with human-written descriptions. Both the *unseen* and *human description* datasets are suitable for determining the generalization ability of any tested method. Secondly, we report performance on

Definition	Terms	Sources
Alert and fully informed	[knowing, knowledgeable]	[WN, WN]
River in singapore	[Geylang River, Singapore River]	[WP, WP]
In the middle of the week	[midweek]	[Wik, CHA]
A type of shotgun	[12 gauge, greener]	[Urban, CHA]
An arsonist	[arsonite, torchman, incendiary]	[Wik, Urban, Mul]
Any supply that is running low	[low supply, short supply]	[Hei++, Hei++]
A term used to describe something so awesome the only way it could be better is if it was between two slices of bread	[ass kicking sandwich]	[Urban]
An abnormal accumulation of air in the pleural space (the space between the lungs and the chest cavity) that can result in the partial or complete collapse of a lung	[Primary Spontaneous Pneumothorax]	[Sci]

Table 1: 3D-EX examples. Note that one definition could map to more than one term, which in turn can come from different dictionaries. Also note the range of styles and domains. (WordNet (WN), Wikipedia (WP), Wiktionary (Wik), MultiRD (Mul), Sci-definition (Sci))

the different dictionaries included in 3D-EX (Almeman et al., 2023) (Section 5.2), a comprehensive resource that integrates multiple dictionaries and organizes them into <word, definition> pairs and <word, definition, example> triplets⁶. We convert this dataset into a suitable RD format, namely <definition, list of terms> and perform two types of splits: a *random split* and a *source split*. In the random split, the data is split randomly into 60% for training, 20% for validation, and 20% for testing. In the source split, all definitions from each source in the dataset are extracted into separate datasets, and then each dataset is split randomly into training, validation, and test sets using the same 60%, 20%, and 20% ratio. Despite the unsupervised nature of our work, we still conduct all our experiments on the test splits alone to enable comparison in further iterations with supervised methods. Table 1 shows examples of entries in 3D-EX in its new RD format, and illustrates the diversity in register, domain and style within the resource. For the benefit of the reader, we provide a brief description of each resource integrating 3D-EX below:

- **WordNet**: A lexical database that groups words into synsets (synonym sets) with definitions, lemmas, examples, and word relations (Miller, 1995; Fellbaum, 2013). WordNet is commonly used as a sense inventory in NLP (Agirre and Edmonds, 2007; Zhang et al., 2022).

⁶While dictionary examples are a valuable resource for improving text representations, we leave them out of our experiments in order to limit the number of components to test.

- **CHA**: A dataset containing words, definitions, and examples from the Oxford Dictionary (Chang and Chen, 2019). It has been used for dictionary modeling (Bevilacqua et al., 2020) and to evaluate the quality of WordNet’s examples (Almeman and Espinosa-Anke, 2022).
- **Wikipedia**: A free online encyclopedia collaboratively created by contributors around the world (Yano and Kang, 2016).
- **Wiktionary**: A web-based dictionary that provides detailed information on words, including definitions, examples, and pronunciation (Bajčetić and Declerck, 2022).
- **Urban**: A crowd-sourced platform that focuses on slang and informal language not usually covered in traditional dictionaries (Wilson et al., 2020).
- **CODWOE**: The English dataset from the CODWOE (Comparing Dictionaries and Word Embeddings) SemEval 2022 shared task (Mickus et al., 2022).
- **Sci-definition**: A dataset specifically created to produce definitions for scientific terms with different levels of complexity (August et al., 2022).
- **Webster’s Unabridged**: A version of Webster’s dictionary (Webster, 1900) available through Project Gutenberg (Various, 2009), offering definitions and additional notes for English words.

Model/Method	Seen Definition			Unseen Definition			Description		
	mr	acc@k	rv	mr	acc@k	rv	mr	acc@k	rv
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
Multi-channel	16	.20/.44/.71	310	54	.09/.29/.58	358	2	.32/.64/.88	203
BERT	0	.57/.86/.92	240	18	.20/.46/.64	418	1	.36/.77/.94	94
RoBERTa	0	.57/.84/.92	228	37	.10/.36/.60	405	1	.43/.85/.96	46
GEAR_bp1	0	.66/.84/.96	200.122	0	.70/.88/.97	180.955	0	.89/.99/.99	70.5334
GEAR_bp2	0	.71/.88/.97	170.451	0	.65/.82/.95	225.324	0	.93/.99/1	1.57314
GEAR_rp	0	.70/.87/.96	185.97	0	.66/.86/.96	190.8	0	.91/.99/1	1.7837

Table 2: GEAR results on the Hill’s dataset compared to competitor models (using the Instructor model for embeddings, as it achieves the best performance), according to median rank (mr), accuracy@k (acc @ 1/100/1000), and rank variance (rv). Baselines results are from Zhang et al. (2020) and Yan et al. (2020).

- **MultiRD**: A dataset developed by (Zhang et al., 2019) to evaluate a reverse dictionary model, using the English dictionary data from Hill et al. (2016).
- **Hei++**: A dataset created by Bevilacqua et al. that pairs human-made definitions with adjective-noun phrases, based on the test split of the HeiPLAS dataset (Hartung, 2015).

5 Main Experiments

5.1 GEAR on Hill’s dataset

We introduce two sets of experiments on Hill’s dataset. The first one uses an LLM alone to perform RD by directly generating terms from input definitions. The second experiment integrates the LLM with embedding models to form GEAR, which enhances performance across the board. For evaluation, we compare GEAR results with the following baselines: (1) OneLook, which is the most popular commercial RD system (Zhang et al., 2020); (2) BOW and RNN with rank loss (Hill et al., 2016), which are neural models where BOW uses a bag-of-words approach and RNN employs Long Short-Term Memory (LSTM); (3) RDWECI (Morinaga and Yamaguchi, 2018), which improves BOW by adding category inference; (4) SuperSense (Pilehvar, 2019), which advances BOW by using pre-trained sense embeddings; (5) MS-LSTM (Kartsaklis et al., 2018), which enhances RNN with Word-

Net synset embeddings and a multi-sense LSTM; (6) Multi-channel (Zhang et al., 2020); and (7) BERT and RoBERTa (Yan et al., 2020), which are trained to generate the target word for the RD task. We use three evaluation metrics based on previous work: median rank of target words (lower is better), accuracy of target words in the top K results (higher is better), and rank variance (lower is better). Table 2 shows that GEAR outperforms all baselines on both the unseen definition set and the description set. We can also observe that MS-LSTM performs effectively on the seen definition set but not on the description set, showing its limited ability to generalize (Zhang et al., 2020).

5.2 GEAR on 3D-Ex

Despite the new SoTa established on Hill’s dataset, we are also interested in exploring core components such as LLMs or embeddings on other resources. To this end, and taking 3D-Ex as a test bed, we apply the same set of experiments as described in Section 5.1 and introduce a new experiment that uses only different text embeddings (without the generate step). This will serve as a baseline to what extent GEAR can provide an improvement over these basic and untuned approaches.

In terms of experimental setup, unless otherwise specified, we always evaluate a fixed-length ranked list of k terms for a given definition, which are then compared to the gold terms. As for evaluation

metrics, we use two different metrics:

- *Mean Reciprocal Rank* (MRR), which measures the position of the first correct result in a list of outcomes, is defined as

$$\text{MRR} = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{1}{\text{rank}_i} \quad (1)$$

where Q is a sample of experiment runs and rank_i refers to the rank position of the *first* relevant outcome for the i th run.

- *Precision @ k* (P@k), which calculates the precision of relevant items within the top k positions of a ranked list, is defined as follows:

$$\text{Precision@k} = \frac{1}{k} \sum_{i=1}^k \text{rel}_i \quad (2)$$

where rel_i is 1 if the item at position i is relevant and 0 otherwise. We now introduce the different embedding models we use in our experiments.

Embeddings We evaluate the performance of different embedding models, which we select considering factors such as adoption among the community, performance in open benchmarks such as MTEB (Muennighoff et al., 2023), availability in the HuggingFace⁷ hub, as well as being of manageable size. These models are:

- SBERT (Reimers and Gurevych, 2019) models, namely `all-MiniLM-L6-v2`, `all-distilroberta-v1` and `all-mpnet-base-v2`.
- Jina Embeddings (Günther et al., 2023) which is a language model that has been trained using Jina AI’s Linnaeus-Clean dataset that contains query-document pairs. We use `jina-embedding-b-en-v1` and `jina-embedding-l-en-v1`.
- General Text Embeddings (GTE) model (Li et al., 2023) which is trained on a large-scale corpus of relevance text pairs from different domains. In this work we use `gte-large`.
- Instructor (Su et al., 2022). It generates text embeddings for different tasks (such as classification or retrieval) and domains (such as

science or finance) based on task instructions. We use `instructor-large`. We examined three different variants of instructions for encoding terms and definitions: (1) no instructions provided; (2) using a general description for the target text (i.e., “Represent the sentence:” and “Represent the word:”); and (3) applying dictionary-specific instructions for the target texts (i.e., “Represent the dictionary definition:” and “Represent the sentence: the dictionary entry:”).

- Universal Angle Embedding (Li and Li, 2023), another instruction-based encoder, and which we use with the same configurations as Instructor. We use `UAE-Large-V1`.

6 Results and Analysis

6.1 Hill’s dataset

Table 3 shows how the performance improves with the GEAR method. In the first part of the table, where candidates are evaluated without any embeddings, the likelihood of having the target term in the top 5 candidates is low. Among the models tested with GEAR, the Instructor model, which encodes terms as dictionary entries based on instructions, performs best, which gains of above 10% on the seen and unseen splits, but negligible differences on the human descriptions (presumably because these already get a good sentence embedding from sentence BERT, on one hand, and also because they might not be accurately described as dictionary resources, which is what we used as an instruction for Instructor). For prompt effectiveness, we found that adding the requirement to generate a dictionary example, and despite its usefulness in other settings, does not improve over base prompt 2, which simply provides as input a few exemplars.

6.2 3D-EX Dataset

Table 4 presents the average MRR, P@1, P@3, and P@5 for each prompt across different resources in 3D-EX, comparing two methods: one without embedding models and the other using the GEAR method. As previously demonstrated in Hill’s results (Section 6.1), the GEAR method, particularly with the Instructor model, achieves the best performance, showing improvements in MRR and P@1 ranging from 3 to 6 points. Concerning the type of prompt, interestingly, we found the sophistication of the prompt to matter the most when combined

⁷<https://huggingface.co/>

Model	Split	Prompt	ACC@1	ACC@5
-	S.	bp1	26.8	44.0
		bp2	30.4	50.2
		rp	29.2	46.8
	U.	bp1	30.0	47.2
		bp2	33.4	53.8
		rp	33.4	49.8
	D.	bp1	70.0	77.0
		bp2	72.5	83.0
		rp	72.0	81.5
SBERT	S.	bp1	57.8	68.2
		bp2	62.4	73.0
		rp	60.6	72.2
	U.	bp1	59.0	70.8
		bp2	63.8	77.0
		rp	61.6	75.2
	D.	bp1	90.5	96.5
		bp2	93.5	97.5
		rp	94.0	98.0
Instructor	S.	bp1	66.0	80.6
		bp2	71.4	84.6
		rp	70.4	84.0
	U.	bp1	64.6	79.0
		bp2	70.0	83.4
		rp	66.4	82.4
	D.	bp1	89.5	98.0
		bp2	92.5	98.5
		rp	91.5	99.0

Table 3: Performance comparison of LLMs (no embeddings models, top block) and GEAR methods (middle and bottom block) across various prompts in Hill’s dataset. S.: Seen split, U.: Unseen split, and D.: human description split. Prompt types are bp1 (Base Prompt 1), bp2 (Base Prompt 2), and rp (Reasoning Prompt).

Model	Prompt	MRR	P@1	P@3	P@5
-	bp1	28.27	24.58	10.64	6.91
	bp2	30.21	26.18	11.39	7.42
	rp	30.99	26.98	11.70	75.89
SBERT	bp1	40.61	33.75	17.36	11.96
	bp2	43.01	36.21	18.32	12.59
	rp	44.21	37.09	18.92	12.99
Instructor	bp1	43.47	36.41	18.00	12.24
	bp2	45.58	38.67	18.76	12.73
	rp	46.37	39.31	19.09	12.98

Table 4: Performance comparison of LLMs (no embeddings models) and GEAR methods across different models and prompts in 3D-EX, showing the average score across different dictionaries.

with embedding models (where we see an improvement of around 7% MRR from base to reasoning), but only 2% when prompting alone is considered.

Table 5 shows the average MRR, P@1, P@3, and P@5 for each embedding model mentioned in Section 5.2. Results are much lower compared to

those achieved with the GEAR method, as well as below prompting alone. In Figure 1, we illustrate the performance across these different datasets, and verify that Hei++ and Sci-definition datasets have higher values, while Urban and CHA show lower values. This variation is likely due to the nature of the entries in Hei++ and Sci-definition, designed to capture more specialized and unique terms. We see, interestingly, that while Instructor embeddings alone are consistently outperforming the rest, they particularly shine in WordNet, which suggest that WordNet embeddings may benefit from additional context to the encoder, since it has been shown that WordNet’s definitions and examples are perhaps too short to be informative (Almeman and Anke, 2022; Giulianelli et al., 2023).

Model	MRR	P@1	P@3	P@5
Instructor (dict. - dict.)	24.88	19.80	9.82	6.73
Instructor (gen. - dict.)	24.81	19.40	9.96	6.78
Instructor (gen. - no)	24.73	19.72	9.78	6.68
Instructor (dict. - no)	24.51	19.55	9.68	6.60
Instructor (gen. - gen.)	24.32	19.17	9.68	6.65
Instructor (dict. - gen.)	24.15	19.05	9.55	6.61
Instructor (no - no)	24.03	19.23	9.45	6.47
Instructor (no - dict.)	23.92	18.91	9.51	6.50
Jina (large)	23.78	18.96	9.33	6.40
GTE (large)	23.47	18.18	9.38	6.56
Instructor (no - gen.)	23.24	18.31	9.20	6.37
UAE (gen. - gen.)	22.93	17.63	9.25	6.41
UAE (dict. - gen.)	21.99	16.96	8.77	6.16
UAE (gen. - dict.)	21.79	16.80	8.71	6.12
UAE (dict. - dict.)	20.90	15.87	8.45	5.94
Jina (base)	20.85	16.11	8.38	5.81
all-mpnet-base-v2	20.14	15.96	7.96	5.45
UAE (gen. - no)	18.91	14.48	7.66	5.36
UAE (dict. - no)	17.77	13.55	7.20	5.01
all-MiniLM-L6-v2	17.04	13.07	6.84	4.82
all-distilroberta-v1	16.55	13.05	6.54	4.50
UAE (no - gen.)	8.98	7.25	3.51	2.39
UAE (no - dict.)	8.10	6.50	3.16	2.15
UAE (no - no)	5.34	3.81	2.24	1.58

Table 5: Comparing different embedding models without any support from an LLM-based generation step, showing the average score across 3D-EX dictionaries.

7 Analyzing GEAR Components

Generation with Open Source LLM To ensure that our approach is both effective across different models and accessible for future research, we repeated the GEAR experiment on the Hill’s dataset using Llama (Touvron et al., 2023), specifically Llama 3.1-70B⁸, an open-source, freely available,

⁸<https://huggingface.co/meta-llama/Llama-3.1-70B>

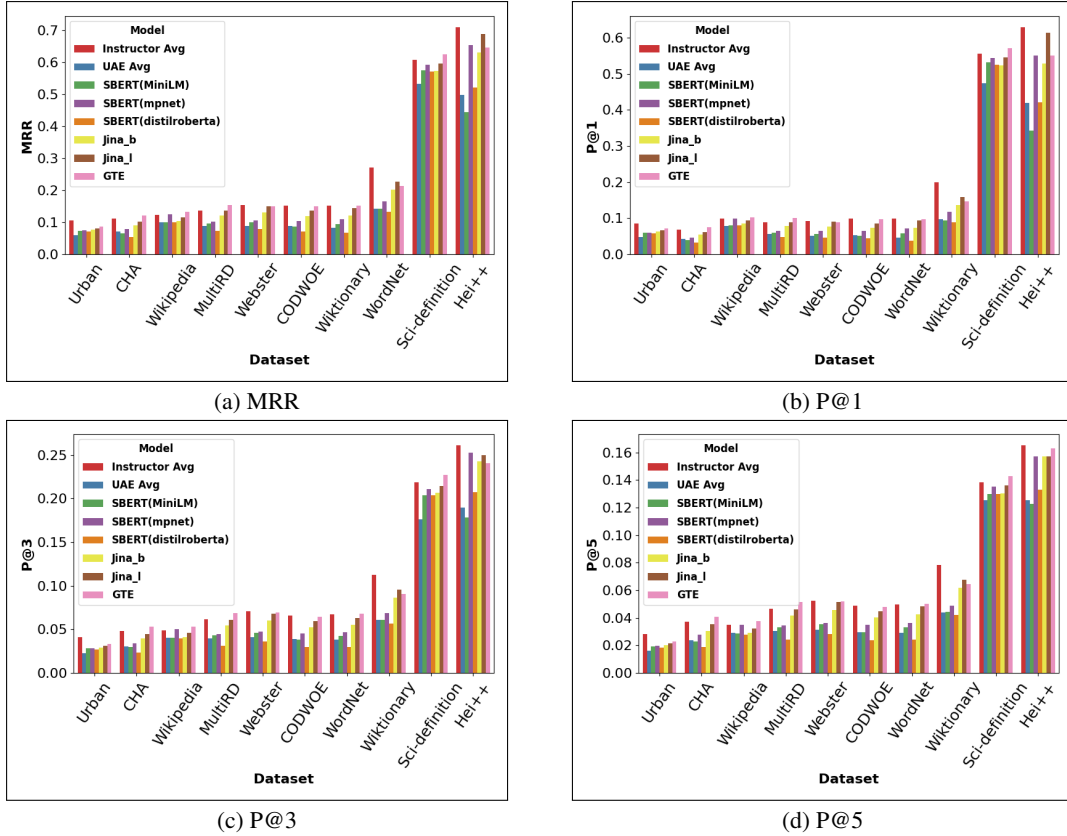


Figure 1: Performance comparison for various embedding models across different metrics in 3D-EX

and highly capable language model. As demonstrated in Table 6, our method continues to outperform the competitor systems in 2 out of 3 datasets.

Split	Prompt	ACC@1	ACC@10	ACC@100
S.	bp1	55.6	67.4	76.8
	bp2	69.8	83.6	94.2
	rp	76.2	89	97
U.	bp1	68.8	84.4	92.4
	bp2	67.2	85	93.2
	rp	69.8	87.8	95.2
D.	bp1	88	97.5	99.5
	bp2	88.5	98.5	99.5
	rp	91.5	100	100

Table 6: GEAR results using Llama for candidates generation and Instructor model for embeddings across different prompts in Hill’s dataset. S.: Seen split, U.: Unseen split, and D.: human description split. Prompt types are bp1 (Base Prompt 1), bp2 (Base Prompt 2), and rp (Reasoning Prompt).

Different Pooling Methods In order to gain a deeper understanding of the effect of the number of candidates produced during the generation step of our GEAR method, we plot the performance metrics across different candidate values for all

three splits and for precision @ k. Figure 2 shows that using just one candidate is not optimal, while averaging over 2 or 3 candidates provides better results, sometimes outperforming all the 5. These results also suggest that while we could have tuned the candidate number on a development set, with the tools we tested (gpt-4o-mini; and Instructor and SBERT), it seems proven that performance plateaus at only a handful of generated terms.

Additionally, We explored the effectiveness of max pooling instead of averaging the generated term embeddings. These experiments did not provide any improvements over the averaging method results shown in Table 2. As we can see in Table 7, for the full GEAR method using **bp1**, the results were around 1-2% worse for all 3 ks in accuracy @ k. Similarly, for the other two prompts, we found a consistent under-performance when compared with averaging, again between 1% and 2% below, with the performance on Hill’s test set going further below, up to 4%.

8 Conclusions and Future Work

We have introduced a *very* simple method for RD, which is based on a simple pipeline where we use

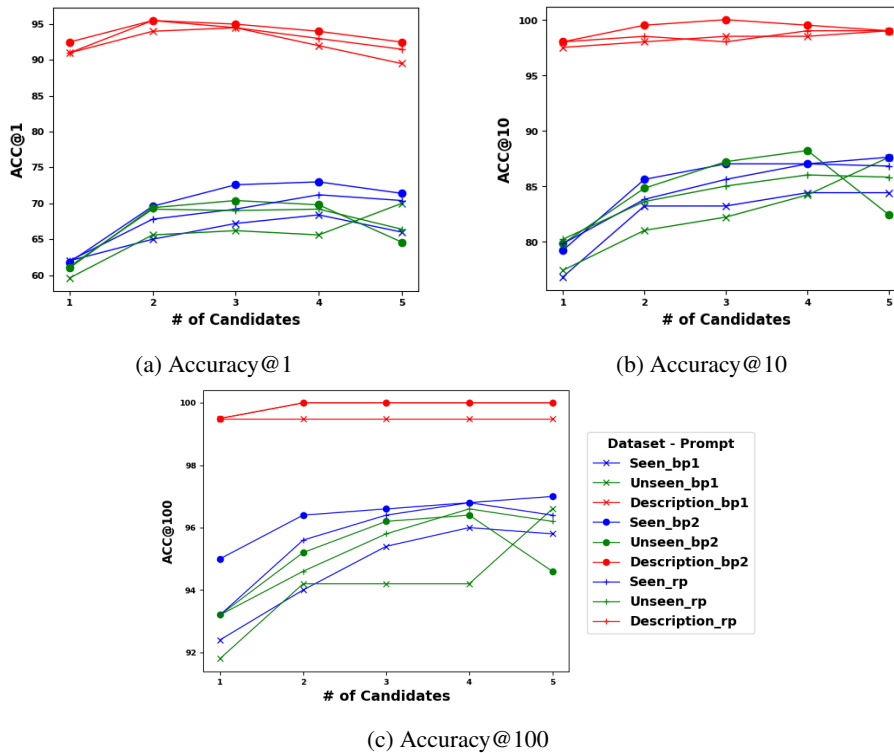


Figure 2: A comparison of the performance on Hill’s splits, evaluating the number of candidates in the generate step, which are then subsequently averaged to produce the input vector for KNN search.

Split	Prompt	ACC@1	ACC@10	ACC@100
S.	bp1	63.6	81.2	94.0
	bp2	69.0	84.6	95.8
	rp	65.6	83.4	4.6
U.	bp1	61.6	80.4	93.8
	bp2	68.2	86.2	95.0
	rp	66.6	85.2	95.2
D.	bp1	90.0	97.5	100
	bp2	90.5	97.0	100
	rp	90.5	98.0	100

Table 7: Max pooling results across different prompts in Hill’s dataset using the Instructor model for embeddings. S.: Seen split, U.: Unseen split, and D.: human description split. Prompt types are bp1 (Base Prompt 1), bp2 (Base Prompt 2), and rp (Reasoning Prompt).

LLMs to generate candidate terms given a definition, and embed them using some pooling technique (the best results are given by simple averaging). This unsupervised method outperforms existing supervised methods on a well known dataset, only falling short on the *seen* split by methods that seem to overfit. In addition to this, we explored different components of this method, and evaluated its performance on various dictionaries and under different evaluation settings. For the future, we would like to explore more LLMs for the genera-

tion step, possibly developing a weighted average approach, which would be task-specific, and that could be learned via simple neural network architectures, similarly to Wang et al. (2021), who used word classification datasets for tuning contextualized word embeddings. We would also like to expand our experiments to a multilingual setting, and test the outputs of an effective RD method as signal for pre-training general purpose embedding models.

Limitations

We have identified some limitations in this work. For instance, assuming that a single embedding model is effective across different registers and audiences might be an oversimplification. Considering various pre-training strategies could provide better insights into which dictionaries perform better or worse in different contexts. Additionally, although we have evaluated the performance of GEAR on the 3D-EX dataset, comparing these results with baseline RD models is needed for a complete evaluation.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Eneko Agirre and Philip Edmonds. 2007. *Word sense disambiguation: Algorithms and applications*, volume 33. Springer Science & Business Media.
- Fatemah Almeman and Luis Espinosa Anke. 2022. Putting wordnet’s dictionary examples in the context of definition modelling: An empirical analysis. In *Proceedings of the Workshop on Cognitive Aspects of the Lexicon*, pages 42–48.
- Fatemah Almeman and Luis Espinosa-Anke. 2022. Putting wordnet’s dictionary examples in the context of definition modelling: An empirical analysis. In *Proceedings of the Workshop on Cognitive Aspects of the Lexicon*, pages 42–48.
- Fatemah Almeman, Hadi Sheikhi, and Luis Espinosa Anke. 2023. [3D-EX: A unified dataset of definitions and dictionary examples](#).
- Tal August, Katharina Reinecke, and Noah A. Smith. 2022. [Generating scientific definitions with controllable complexity](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8298–8317, Dublin, Ireland. Association for Computational Linguistics.
- Lenka Bajčetić and Thierry Declerck. 2022. [Using Wiktionary to create specialized lexical resources and datasets](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3457–3460, Marseille, France. European Language Resources Association.
- D Frank Benson. 1979. Neurologic correlates of anomia. In *Studies in neurolinguistics*, pages 293–328. Elsevier.
- Michele Bevilacqua, Marco Maru, and Roberto Navigli. 2020. [Generatory or “how we went beyond word sense inventories and learned to gloss”](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7207–7221, Online. Association for Computational Linguistics.
- Slaven Bila, Wataru Watanabe, Taiichi Hashimoto, Takenobu Tokunaga, and Hozumi Tanaka. 2004. Dictionary search based on the target word description.
- Roger Brown and David McNeill. 1966. The “tip of the tongue” phenomenon. *Journal of verbal learning and verbal behavior*, 5(4):325–337.
- Ting-Yun Chang and Yun-Nung Chen. 2019. [What does this word mean? explaining contextualized embeddings with natural language definition](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6064–6070, Hong Kong, China. Association for Computational Linguistics.
- Guowei Chen and Jianbo Su. 2021. Towards non-ambiguous reverse dictionary. In *2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 1113–1120. IEEE.
- Pinzhen Chen and Zheng Zhao. 2022. A unified model for reverse dictionary and definition modelling. *arXiv preprint arXiv:2205.04602*.
- Qianglong Chen, Feng-Lin Li, Guohai Xu, Ming Yan, Ji Zhang, and Yin Zhang. 2022. Dictbert: Dictionary description knowledge enhanced language model pre-training via contrastive learning. *arXiv preprint arXiv:2208.00635*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Christiane Fellbaum. 2013. Wordnet. In Carol Chapelle, editor, *The encyclopedia of applied linguistics*, pages 6739–6746. Blackwell Publishing Ltd.
- Amit Gajbhiye, Zied Bouraoui, Luis Espinosa Anke, and Steven Schockaert. 2024. [Amended: Modelling concepts by aligning mentions, definitions and decontextualised embeddings](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 801–811.
- Mario Giulianelli, Iris Luden, Raquel Fernandez, and Andrey Kutuzov. 2023. [Interpretable word sense representations via definition generation: The case of semantic change analysis](#). In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- Michael Günther, Louis Milliken, Jonathan Geuter, Georgios Mastrapas, Bo Wang, and Han Xiao. 2023. [Jina embeddings: A novel set of high-performance sentence embedding models](#). *Preprint, arXiv:2307.11224*.
- Matthias Hartung. 2015. *Distributional Semantic Models of Attribute Meaning in Adjectives and Nouns*. Ph.D. thesis.
- Felix Hill, Kyunghyun Cho, Anna Korhonen, and Yoshua Bengio. 2016. Learning to understand phrases by embedding the dictionary. *Transactions of the Association for Computational Linguistics*, 4:17–30.
- Luyao Huang, Chi Sun, Xipeng Qiu, and Xuan-Jing Huang. 2019. [Glossbert: Bert for word sense disambiguation with gloss knowledge](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3509–3514.

- Dimitri Kartsaklis, Mohammad Taher Pilehvar, and Nigel Collier. 2018. [Mapping text to knowledge graph entities using multi-sense LSTMs](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1959–1970, Brussels, Belgium. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2019. [Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#).
- Xianming Li and Jing Li. 2023. Angle-optimized text embeddings. *arXiv preprint arXiv:2309.12871*.
- Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. 2023. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*.
- Sunil B Mane, Harshal Navneet Patil, Kanhaiya Balaji Madaswar, and Pranav Nitin Sadavarte. 2022. Wordalchemy: a transformer-based reverse dictionary. In *2022 2nd International Conference on Intelligent Technologies (CONIT)*, pages 1–5. IEEE.
- Timothee Mickus, Kees Van Deemter, Mathieu Constant, and Denis Paperno. 2022. [Semeval-2022 task 1: CODWOE – comparing dictionaries and word embeddings](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1–14, Seattle, United States. Association for Computational Linguistics.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Yuya Morinaga and Kazunori Yamaguchi. 2018. [Improvement of Reverse Dictionary by Tuning Word Vectors and Category Inference: 24th International Conference, ICIST 2018, Vilnius, Lithuania, October 4–6, 2018, Proceedings](#), pages 533–545.
- Viktor Moskvoretskii, Ekaterina Neminova, Alina Lobanova, Alexander Panchenko, and Irina Nikishina. 2024. Taxollama: Wordnet-based model for solving multiple lexical semantic tasks. *arXiv preprint arXiv:2403.09207*.
- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2023. Mteb: Massive text embedding benchmark. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2014–2037.
- Mohammad Taher Pilehvar. 2019. On the importance of distinguishing word meaning representations: A case study on reverse dictionary mapping. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2151–2156.
- Mohammad Taher Pilehvar and Jose Camacho-Collados. 2019. [WiC: the word-in-context dataset for evaluating context-sensitive meaning representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1267–1273, Minneapolis, Minnesota. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Ryan Shaw, Anindya Datta, Debra VanderMeer, and Kaushik Dutta. 2011. Building a scalable database-driven reverse dictionary. *IEEE Transactions on Knowledge and Data Engineering*, 25(3):528–540.
- Hongjin Su, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen-tau Yih, Noah A Smith, Luke Zettlemoyer, Tao Yu, et al. 2022. One embedder, any task: Instruction-finetuned text embeddings. *arXiv preprint arXiv:2212.09741*.
- Sicheng Tian, Shaobin Huang, Rongsheng Li, and Chi Wei. 2024. A prompt construction method for the reverse dictionary task of large-scale language models. *Engineering Applications of Artificial Intelligence*, 133:108596.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Various. 2009. *Webster’s Unabridged Dictionary*. Project Gutenberg.
- Yixiao Wang, Zied Bouraoui, Luis Espinosa Anke, and Steven Schockaert. 2021. Deriving word vectors from contextualized language models using topic-aware mention selection. In *Proceedings of the 6th Workshop on Representation Learning for NLP (ReplANLP-2021)*, pages 185–194.
- Leo Wanner. 1996. Lexical choice in text generation and machine translation. *Machine Translation*, 11(1):3–35.
- Noah Webster. 1900. *Webster’s unabridged dictionary of the English language*. Kikwansha.
- Steven R. Wilson, Walid Magdy, Barbara McGillivray, Venkata Rama Kiran Garimella, and Gareth Tyson. 2020. Urban dictionary embeddings for slang nlp applications. In *International Conference on Language Resources and Evaluation*.

Ningyu Xu, Qi Zhang, Menghan Zhang, Peng Qian, and Xuanjing Huang. 2024. On the tip of the tongue: Analyzing conceptual representation in large language models with reverse-dictionary probe. *arXiv preprint arXiv:2402.14404*.

Hang Yan, Xiaonan Li, Xipeng Qiu, and Bocao Deng. 2020. [BERT for monolingual and cross-lingual reverse dictionary](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4329–4338, Online. Association for Computational Linguistics.

Tae Yano and Moonyoung Kang. 2016. Taking advantage of wikipedia in natural language processing.

Wenhao Yu, Chenguang Zhu, Yuwei Fang, Donghan Yu, Shuohang Wang, Yichong Xu, Michael Zeng, and Meng Jiang. 2022. Dict-bert: Enhancing language model pre-training with dictionary. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1907–1918.

Guobiao Zhang, Wenpeng Lu, Xueping Peng, Shoujin Wang, Baoshuo Kan, and Rui Yu. 2022. Word sense disambiguation with knowledge-enhanced and local self-attention-based extractive sense comprehension. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4061–4070.

Lei Zhang, Fanchao Qi, Zhiyuan Liu, Yasheng Wang, Qun Liu, and Maosong Sun. 2019. [Multi-channel reverse dictionary model](#). *Preprint*, arXiv:1912.08441.

Lei Zhang, Fanchao Qi, Zhiyuan Liu, Yasheng Wang, Qun Liu, and Maosong Sun. 2020. Multi-channel reverse dictionary model. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 312–319.

Michael Zock. 2004. Word lookup as an ongoing dialogue between a user and a lexicon. In *Proceedings of the 10th Annual Meeting of the Association for Natural Language Processing*, pages 484–487.

Michael Zock and Slaven Bilac. 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*, pages 29–35.

Michael Zock, Olivier Ferret, and Didier Schwab. 2010. Deliberate word access: an intuition, a roadmap and some preliminary empirical results. *International Journal of Speech Technology*, 13(4):201–218.

A Prompt Types

A.1 Base Prompt 1 (bp1):

Given the definition {definition}, generate a list of {k} terms defined by that definition assuming they are in {dictionary} dictionary. Only give me a list back, do not generate any other text.
{dictionary} is {description}
The returned list should follow the following conditions:

- Terms should be ranked, with the first term being the most related to the definition.
- In a JSON object of the form { "terms": ["term_1", "term_2", ...] }.
- All terms should be in lowercase.

Example:

INPUT: "A piece of furniture for sitting."
OUTPUT: { "terms": ["chair", "stool", "bench", "sofa", "couch"] }

A.2 Base Prompt 2 (bp2):

Given the definition {definition}, generate a list of {k} terms defined by that definition assuming they are in {dictionary} dictionary. Only give me a list back, do not generate any other text.
{dictionary} is {description}
These are some examples of definitions and terms in this dictionary: {examples}
The returned list should follow the following conditions:

- Terms should be ranked, with the first term being the most related to the definition.
- In a JSON object of the form { "terms": ["term_1", "term_2", ...] }.
- All terms should be in lowercase.

Example:

INPUT: "A piece of furniture for sitting."
OUTPUT: { "terms": ["chair", "stool", "bench", "sofa", "couch"] }

A.3 Reasoning Prompt (rp):

Given the definition {definition}, generate a list of {k} terms defined by that definition assuming they are in {dictionary} dictionary. Only give me a list back, do not generate any other text.

{dictionary} is {description}

These are some examples of definitions and terms in this dictionary: {examples}

For each term, provide an example usage in a sentence that matches the style and scope of {dictionary}.

The returned list should follow the following conditions:

- Terms should be ranked, with the first term being the most related to the definition.
- All terms and examples should be in lowercase.
- Return the terms and examples in a JSON object of the form:

```
{ "terms": [ { "term": "term_1", "example": "example_1" }, { "term": "term_2", "example": "example_2" }, ... ] }
```

Example:

INPUT: "A piece of furniture for sitting."

OUTPUT: { "terms": [{ "term": "chair", "example": "he sat on the chair and opened his book." }, { "term": "stool", "example": "she perched on the stool at the bar." }, { "term": "bench", "example": "they rested on the bench after their walk." }, { "term": "sofa", "example": "the family gathered on the sofa to watch TV." }, { "term": "couch", "example": "he stretched out on the couch to take a nap."] } }