

# Evaluating Distributed Representations for Multi-Level Lexical Semantics: A Research Proposal

**Zhu Liu**

Tsinghua University

Computational Linguistics, School of Humanities

liuzhu22@mails.tsinghua.edu

## Abstract

Modern neural networks (NNs), trained on extensive raw sentence data, construct distributed representations by compressing individual words into dense, continuous, high-dimensional vectors. These representations are expected to capture multi-level lexical meaning. In this thesis, our objective is to examine the efficacy of distributed representations from NNs in encoding lexical meaning. Initially, we identify and formalize three levels of lexical semantics: *local*, *global*, and *mixed* levels. Then, for each level, we evaluate language models by collecting or constructing multilingual datasets, leveraging various language models, and employing linguistic analysis theories. This thesis builds a bridge between computational models and lexical semantics, aiming to complement each other.

## 1 Introduction

A key issue in lexical semantics is the many-to-many relation between form and meaning is multifaceted. One form can encompass multiple senses, such as in cases of homonymy and polysemy. Conversely, one meaning can correspond to multiple forms, e.g., synonymy. Together, these phenomena organize a lexicon into a semantic field (Jackson and Amvela, 2000) with varying levels (Liétard et al., 2024). From a local perspective, each word possesses potential meanings influenced by different contexts. These meanings may be unrelated (homonymy) or related, with varying degrees of relatedness. At a global level, words are interconnected through relationships such as word analogy (*man* – *woman* = *king* – *queen*) and compositionality (*Italy* + *noodles* = *spaghetti*). At a mixed level, we can conceive an overarching conceptual space that captures cross-lingual universality, integrating both words and their local meanings.

Modern computational language models based on transformer architectures (Vaswani et al., 2017)

represent a word <sup>1</sup> as a real-valued, contextual, high-dimensional vector (Petersen and Potts, 2023) (or representation). This approach is grounded in the distributional hypothesis (Harris, 1954), which posits that words appearing in similar contexts have similar meanings. Vectors from pre-trained language models (PLMs) and large language models (LLMs) have been utilized as initial inputs for downstream tasks related to lexical semantics, achieving outstanding performance in areas such as word sense disambiguation (Bevilacqua et al., 2021), lexicon induction (Li et al., 2023), and reverse dictionary tasks (Tian et al., 2024). Natural and important questions arise: how well do these distributed representations from PLMs and LLMs convey contextual meaning? Do they truly possess the same lexical knowledge as humans, given their impressive linguistic performance? If not, to what extent do they fall short? Answering these questions is significant for both the computational and linguistic communities. For computational researchers, these insights could enhance the transparency and reliability of black-box models. For linguists, they could aid in constructing meaning systems and discovering novel meanings.

This thesis explores these research questions by evaluating current pre-trained language models (PLMs) and large language models (LLMs) on a variety of lexical meaning tasks with differing levels of granularity. To incorporate a cross-lingual perspective, we design benchmarks spanning multiple languages. Furthermore, we test our hypotheses by integrating theories and models from both computational linguistics and traditional linguistic frameworks.

**Research Proposal** In this thesis, we first formalize our research problems, including key concepts and the three levels of analysis (see Section 2).

---

<sup>1</sup>The actual representation for a model corresponds to a token, i.e., subword. We can represent a word by aggregating the contextual embeddings from subwords in practice.

Next, we collect datasets and design experiments to evaluate each aspect (see Section 3). Finally, we address potential methodological challenges and present the conclusions drawn from our findings (see Section 4).

## 2 Formalization

### 2.1 Basic Notions

We consider a lexicon  $\mathcal{W}$  consisting of a finite and countable set of items. These items include all lemmatized words  $w_i \in \mathcal{W}$ , each associated with a prototypical meaning  $m(w_i)$ . This prototypical meaning represents the most frequent and typical sense recognized by speakers of a given language community (Rosch, 1975), and is typically listed first in a dictionary. For instance, the prototypical meaning of  $m_{\text{bank}}$  refers to a financial institution, rather than the side of a river or a shape of fog. These prototypical meanings form a meaning space  $\mathcal{M}$ .

Words are then sampled and combined into sequences, forming a corpus  $\mathcal{R}$ , consisting of sentences  $s_i$ . For a word  $w_i$ , its occurrence  $o(w'_i)$  is realized as a sentence  $s_i$  within the corpus  $\mathcal{R}$ , where  $w'_i$  represents a conjugated form of  $w_i$ , such as the addition of inflectional morphemes<sup>2</sup>. The word  $w'_i$  in its context  $s_i$  carries a conventional meaning, referred to as its sense  $e(w'_i)$ , which is typically listed in other dictionary entries, distinct from the prototypical meaning  $m(w_i)$ . These senses, across lexical items, constitute a conceptual field  $\mathcal{C}$ , where each concept represents a distinct type of meaning. Thus, a concept is the minimal unit through which we understand the world. There are several examples of the combination of  $\mathcal{W}$ ,  $\mathcal{R}$ , and  $\mathcal{C}$ .

**Example 1** A common dictionary for learners, such as the Oxford Dictionary<sup>3</sup>, includes words  $w$  with their associated meanings  $m$ , and sentences derived from a corpus  $\mathcal{R}$ . While the entire conceptual space is implicit, it is suggested through sets of synonyms.

**Example 2** WordNet (Miller et al., 1990) contains words, their meanings, and sentences, similar to a conventional dictionary. Additionally, synsets in WordNet, as nodes, explicitly represent a conceptual space by grouping synonymous words.

<sup>2</sup>This process occurs in the surface structure derived from the deep structure, according to generative grammar (Chomsky, 1965).

<sup>3</sup><https://en.oxforddictionaries.com/>

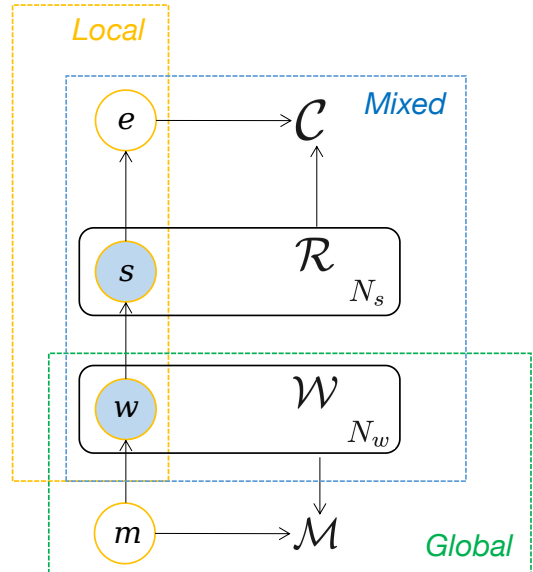


Figure 1: Graph models for four spaces: lexicon  $\mathcal{W}$ , corpus  $\mathcal{R}$ , concept  $\mathcal{C}$ , and prototypical meaning  $\mathcal{M}$ . The three levels and their respective scopes are identified. The shaded circle represents an observable variable, while the unshaded one indicates an unobservable variable.

**Example 3** The mental lexicon (Klepousniotou, 2002) is inherently present in our minds. Words in the vocabulary are learned incrementally from fragments of utterances in  $\mathcal{R}$ . The conceptual space helps us recognize equivalent concepts and facilitate categorization.

### 2.2 Three Levels

We propose a graphical model to illustrate the relationships among four spaces:  $\mathcal{W}$ ,  $\mathcal{M}$ ,  $\mathcal{R}$ , and  $\mathcal{C}$ , as shown in Figure 1. Further examples are provided in Figure 2. Based on these, we identify three hierarchical levels.

**Local Level** This level focuses on the sense space within a word, addressing issues such as the number of senses, how senses are divided and evolve, the similarity and relatedness of potential senses, as well as conventional or temporal meanings. Tasks like word sense disambiguation (Navigli, 2009), word sense induction (Van de Cruys and Apidianaki, 2011), and lexical semantic changes (Schlechtweg, 2023) reflect the local space of a word. We formalize this level using the following likelihood probability:

$$p(e) = p(e|w, s), \quad (1)$$

where  $w \in \mathcal{W}$ ,  $s \in \mathcal{R}$ , and  $e \in \mathcal{C}$ .

**Global Level** From another perspective, we consider all the words in a lexicon and examine their semantic relationships. Since no specific context is provided at this level, we focus solely on the prototypical meaning of words. Relations among words include word analogy (Mikolov et al., 2013), word similarity (Huang et al., 2012), compositionality (Yamagiwa et al., 2023), and others. Structures such as vectors and graphs are commonly used to represent these relationships. We formalize this level using the following likelihood probability:

$$p(\mathcal{M}) = p([e_i]_N | \mathcal{W}), \quad (2)$$

where  $N$  is the number of words in the lexicon, and  $[e_i]_N = \{e_1, e_2, \dots, e_i, \dots, e_N\}$ .

**Mixed Level** The mixed level considers the senses across all words in the lexicon. A semantic field can form a set of concepts shared within a language community. Building on the local level, problems such as concept induction and the similarity of concepts are classical challenges at this level. A conceptual space (Haspelmath, 2003) that discovers relations among concepts, often in a cross-lingual context, represents the structure of concepts. The process at the mixed level can be realized in two steps: first, local-level sense induction, followed by global concept induction. We formalize this level using the following likelihood probability:

$$p(\mathcal{C}) = p([c_i]_M | \mathcal{W}, \mathcal{R}) = p(e|w, s) \cdot p([c_i]_M | e, \mathcal{W}), \quad (3)$$

where  $M$  is the number of concepts, and  $[c_i]_M = \{c_1, c_2, \dots, c_i, \dots, c_M\}$ .

In this paper, we evaluate to what extent distributed representations convey lexical semantics at these three levels. At the local level (Subsection 3.2), we focus on the phenomenon of continuous relatedness of senses and investigate whether language models can reflect the continuous variation in different senses. At the global level (Subsection 3.3), we construct a network of word embeddings learned by large language models and analyze their geometric relations among different words. For the mixed level (Section 3.4), we build conceptual spaces from a subset of the lexicon based on typological theories in language.

### 3 Evaluation

We begin by utilizing various types of language models to extract representations, followed by an evaluation of these representations across the three levels.

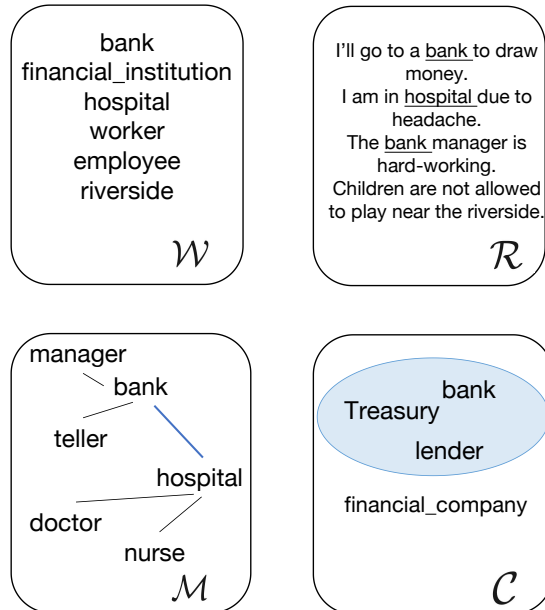


Figure 2: Examples of the four spaces: lexicon  $\mathcal{W}$ , corpus  $\mathcal{R}$ , concept  $\mathcal{C}$ , and prototypical meaning  $\mathcal{M}$ .

### 3.1 Models and Representations

Drawing on the Distributional Hypothesis (Harris, 1954), language models employ neural networks to derive continuous vectors from large-scale corpora. Models such as Word2Vec (Mikolov et al., 2013) and GloVe (Pennington et al., 2014) generate static word representations that do not account for the different senses a word may take in varying contexts. In contrast, transformer-based models (Kenton and Toutanova, 2019) acquire layer-wise contextual representations. Both static and contextual representations leverage the distribution of neighboring words but do not explicitly model meaning in a nuanced manner. Some studies, however, aim to build sense representations in an unsupervised (Liu et al., 2015) or knowledge-based (Chen et al., 2014) fashion.

To evaluate the representations from these models, we consider different configurations, including BERT-like bidirectional models (primarily used in pre-trained language models, PLMs) and GPT-like generative models (commonly employed in large language models, LLMs). We propose to investigate how LLMs encode lexical semantics (BehnamGhader et al., 2024; Liu et al., 2024) in comparison to the more established research on BERT (BERTology (Rogers et al., 2021)). Specifically, we aim to identify where lexical semantics are encoded within the model, which will guide how we extract vectors for further evaluation.

Additionally, we propose transforming the original semantic space into a more discrete form to decouple interdependent features. Techniques such as independent component analysis (Yamagiwa et al., 2023) can be employed for this purpose. By analyzing the values of the transformed feature axes, we seek to identify meaningful dimensions of semantics that better capture the underlying structure of lexical meaning.

### 3.2 Local Level: Sense Continuity

At the local level, we focus on the distribution of possible senses for words. A well-known phenomenon in lexical semantics is the continuity of sense distinctions. That is, the relatedness of senses for different words varies continuously. Linguists distinguish between homonymy and polysemy, for example. However, even within polysemy, the distance between senses for some words in different contexts can vary. Furthermore, variations in semantic roles for words within the same context can also differ. These challenges are reflected in the high annotator disagreement found in lexical semantics-related tasks, such as word sense disambiguation (WSD) (Navigli, 2009).

**Uncertainty in WSD** Existing supervised methods (Bevilacqua et al., 2021) treat WSD as a classification task and have achieved remarkable performance. However, they often overlook uncertainty estimation (UE) in real-world settings, where sense labeling involves substantial disagreement and varying degrees of uncertainty for different words (Liu and Liu, 2023). We propose addressing this issue by formalizing it as a probabilistic inference problem (Gal et al., 2016), which outputs a well-calibrated probability distribution over the candidate sense space.

Incorporating uncertainty estimation also helps differentiate between homonymy and polysemy. Homonymy tends to exhibit higher certainty in its sense labeling, whereas polysemy is often characterized by greater uncertainty. This perspective aligns with the inherent subjectivity of language, which includes phenomena such as underspecification, vagueness, and context sensitivity (Sennet, 2023).

**Semantic Roles** In a basic SVO sentence structure based on predicate-argument relationships, the degree of semantic roles assigned to different grammatical constituents (S, V, and O) can vary. This

phenomenon is linguistically universal. For instance, the verb *shot at* in the sentence *The hunter shot at the bear* exhibits weaker transitivity than *shot* in *The hunter shot the bear*. In Chinese, word order plays a significant role in determining semantic roles; typically, the first argument is the agent, and the second is the patient. Changing the subject and object order can alter the agency and objectivity, although in some cases, the semantic roles remain unchanged. For example, 张三打李四 (Tom hit Alice) is fundamentally different from 李四打张三 (Alice hit Tom). However, 十个人吃一顿饭 (Ten people eat a meal) essentially carries the same meaning as 一顿饭吃十个人 (A meal provides ten people to eat). In the former case, the verb *hit* has the same transitivity in both contexts, while *eat* does not. We propose collecting minimal pairs (where only the subject and object change) that represent different shifts in semantic roles and investigating whether language models can capture such nuances. The forms of minimal pairs may vary across different languages because distinct languages may employ different strategies, such as morphological or syntactic variations, to reflect changes in semantic roles. We emphasize that the form of the minimal pair is specific to each language, as languages may have distinct ways of modifying the degree of semantic roles.

This work is related to the classical NLP task of semantic role labeling (Jurafsky and Martin, 2020), but with several key differences. First, we treat the degree of semantic roles as a continuous variable, rather than a binary choice. Thus, we utilize representations to calculate the similarity between corresponding items. Second, we collect minimal pairs where only one factor changes, such as the transitivity of verbs, to facilitate causal analysis. Finally, we consider cross-lingual universality and compare the behavior of semantic role assignment across different languages.

### 3.3 Global Level: Word Network

At the global level, we propose constructing a word network by leveraging the embeddings learned in large language models (LLMs). In this network, nodes represent words, and edges reflect the similarity between corresponding words. This structure illustrates how the internal representations of models capture the relationships among words. We also aim to compare networks built using models of varying scales. Several factors need to be



addressed. First, how should we extract the representation of a word from raw **token** embeddings? Does the conventional mean-pooling approach remain effective? Second, what similarity measure should be employed? While cosine similarity is commonly used, it may not suffice in higher-dimensional spaces. Third, how should we prune a fully connected graph, particularly when comparing different models? What constitutes a fair strategy for pruning across different model types? Finally, how can we design effective indicators to zoom in and out of the network? When zooming in (Li et al., 2024), we can observe detailed examples such as meaning analogies (e.g., manking-woman-queen). Conversely, zooming out allows us to characterize the entire network through graph statistics, such as the number of connected components.

### 3.4 Mixed Level: Conceptual Spaces

At the mixed level, we consider the senses across words and construct a conceptual space that reflects the similarity of concepts. Given the cross-lingual universality of concepts, we adopt the theory of language typology, specifically semantic map modeling (Haspelmath, 2003) (SMM), where a conceptual space is built based on a connectivity hypothesis.

Function words, affixes, and certain adverbs play a crucial role in SMM due to their multifunctionality. They exhibit a broader range of nuanced semantics or functions compared to content words and are often not exhaustively listed in dictionaries. For example, repetitive grams (such as "and" or "again") in various languages can demonstrate over 20 distinct functions (Zhang, 2017). Linguists use Semantic Map Models (SMM) (Haspelmath, 2003) to visually represent these functions within a conceptual or semantic space, interconnected by lines to form a network. Functions with greater similarity are positioned closer together on the map. SMM is grounded in cross-linguistic comparison, following the "semantic connectivity hypothesis," which suggests that functions expressed by a language-specific category should occupy contiguous areas on the semantic map.

In our approach, we utilize representations from language models to measure the similarity between different occurrences of a target word. Subsequently, we design a graph algorithm to construct the semantic map, adhering to the connectivity

principle. We aim to assess the quality of the automatically generated graph against a human-annotated one using designated metrics.

## 4 Conclusion and Challenges

Distributed representations encode rich lexical semantics, capturing not only word meanings but also contextual associations, allowing for a nuanced understanding of language. This thesis evaluates the extent to which vectorized representations reflect word meaning across three levels: the local level (word-relatedness), the global level (multilingual lexicon relations), and the mixed level (cross-lingual conceptual space). These levels offer both micro and macro perspectives of semantic fields, providing a comprehensive framework for evaluating lexical semantics. By considering these multiple levels, we aim to better understand how models capture the complexity of language. We propose evaluating model representations using common benchmarks and custom datasets, with performance serving as an indicator of semantic quality and the ability to reflect the depth of meaning encoded in these representations.

However, several challenges arise in this probing approach. The first is the "extraction dilemma": can poor task performance be attributed to a model's failure to capture semantics, or does it result from suboptimal representation extraction strategies? Without isolating other factors, we cannot conclusively assess the semantic capacity of the representations. The second issue is probe selectivity (Hewitt and Liang, 2019), where it's unclear if the probe extracts representations or simply learns the task. The third challenge is dataset bias, as contextual meanings are often more subjective than static meanings. Disagreements among human annotators on tasks like polysemy disambiguation can introduce uncertainty, affecting metric design and reliability.

Lastly, while scaling laws (Kaplan et al., 2020) show that larger models improve performance, they also increase opacity, raising questions about model interpretability and trustworthiness. How do these models arrive at their conclusions? Can we explain their predictions in a way that is meaningful to humans? Our research aims to enhance the transparency of modern language models and bridge the gap between computer science and linguistics, fostering a better understanding of how these models represent linguistic information.

## References

- Parishad BehnamGhader, Vaibhav Adlakha, Marius Mosbach, Dzmitry Bahdanau, Nicolas Chapados, and Siva Reddy. 2024. [LLM2Vec: Large language models are secretly powerful text encoders](#). In *First Conference on Language Modeling*.
- Michele Bevilacqua, Tommaso Pasini, Alessandro Raganato, and Roberto Navigli. 2021. Recent trends in word sense disambiguation: A survey. In *International Joint Conference on Artificial Intelligence*, pages 4330–4338. International Joint Conference on Artificial Intelligence, Inc.
- Xinxiong Chen, Zhiyuan Liu, and Maosong Sun. 2014. A unified model for word sense representation and disambiguation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1025–1035.
- Noam Chomsky. 1965. *Aspects of the Theory of Syntax*. The MIT Press, Cambridge, MA.
- Yarin Gal et al. 2016. Uncertainty in deep learning.
- Zellig S Harris. 1954. Distributional structure. *Word*, 10(2-3):146–162.
- Martin Haspelmath. 2003. The geometry of grammatical meaning: Semantic maps and cross-linguistic comparison. In *The new psychology of language*, pages 211–242. Psychology Press.
- John Hewitt and Percy Liang. 2019. Designing and interpreting probes with control tasks. 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 2733–2743, Hong Kong, China. Association for Computational Linguistics.
- Eric H. Huang, Richard Socher, Christopher D. Manning, and Andrew Y. Ng. 2012. Improving word representations via global context and multiple word prototypes. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 873–882, Jeju Island, Korea.
- Howard Jackson and Etienne Zé Amvela. 2000. *Words, Meaning, and Vocabulary*. Continuum.
- Daniel Jurafsky and James H Martin. 2020. *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Ekaterini Klepousniotou. 2002. The processing of lexical ambiguity: Homonymy and polysemy in the mental lexicon. *Brain and Language*, 81(1-3):205–223.
- Yaoyiran Li, Anna Korhonen, and Ivan Vulić. 2023. On bilingual lexicon induction with large language models. *arXiv preprint arXiv:2310.13995*.
- Yuxiao Li, Eric J Michaud, David D Baek, Joshua Engels, Xiaoqing Sun, and Max Tegmark. 2024. The geometry of concepts: Sparse autoencoder feature structure. *arXiv preprint arXiv:2410.19750*.
- Bastien Liétard, Pascal Denis, and Mikaela Keller. 2024. [To word senses and beyond: Inducing concepts with contextualized language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2684–2696, Miami, Florida, USA. Association for Computational Linguistics.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2015. Learning context-sensitive word embeddings with neural tensor skip-gram model. In *IJCAI*, pages 1284–1290.
- Zhu Liu, Cunliang Kong, Ying Liu, and Maosong Sun. 2024. [Fantastic semantics and where to find them: Investigating which layers of generative LLMs reflect lexical semantics](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 14551–14558, Bangkok, Thailand. Association for Computational Linguistics.
- Zhu Liu and Ying Liu. 2023. [Ambiguity meets uncertainty: Investigating uncertainty estimation for word sense disambiguation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3963–3977, Toronto, Canada. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26.
- George A Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, and Katherine J Miller. 1990. Introduction to wordnet: An on-line lexical database. *International journal of lexicography*, 3(4):235–244.
- Roberto Navigli. 2009. Word sense disambiguation: A survey. *ACM computing surveys (CSUR)*, 41(2):1–69.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Erika Petersen and Christopher Potts. 2023. [Lexical semantics with large language models: A case study of English “break”](#). In *Findings of the Association*

- for *Computational Linguistics: EACL 2023*, pages 490–511, Dubrovnik, Croatia. Association for Computational Linguistics.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2021. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- E. Rosch. 1975. Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104(3):192–233.
- Dominik Schlechtweg. 2023. *Human and Computational Measurement of Lexical Semantic Change*. Ph.D. thesis, University of Stuttgart, Stuttgart, Germany.
- Adam Sennet. 2023. Ambiguity. In Edward N. Zalta and Uri Nodelman, editors, *The Stanford Encyclopedia of Philosophy*, Summer 2023 edition. Metaphysics Research Lab, Stanford University.
- 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.
- Tim Van de Cruys and Marianna Apidianaki. 2011. [Latent semantic word sense induction and disambiguation](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1476–1485, Portland, Oregon, USA. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Hiroaki Yamagiwa, Momose Oyama, and Hidetoshi Shimodaira. 2023. [Discovering universal geometry in embeddings with ICA](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4647–4675, Singapore. Association for Computational Linguistics.
- Ying Zhang. 2017. Semantic map approach to universals of conceptual correlations: a study on multifunctional repetitive grams. *Lingua Sinica*, 3(1):7.