Western European Journal of Linguistics and Education
Volume 4, Issue 2, February 2026
https://westerneuropeanstudies.com/index.php/2
ISSN (E): 2942-190X Open Access| Peer Reviewed

EXEMPLAR BASED WORD MEANING IN THE AGE OF CONTEXTUAL EMBEDDINGS

Zulfiya Tukhtakhodjaeva
PhD, Associate professor
Uzbekistan state world languages university

Abstract. This article discusses a closely related alternative to prototype-based accounts of lexical meaning: exemplar-based and usage-centered approaches to words and senses. Instead of treating a word as a container of a small set of discrete dictionary meanings, exemplar-oriented models assume that language users retain memory traces of encountered uses and interpret new instances by similarity to stored experiences. The article traces the cognitive foundations of exemplar models in categorization research and shows how their logic aligns with corpus-driven semantics and the long-standing critique that "word senses" are task-dependent abstractions from usage. It then connects these ideas to contemporary computational semantics, where contextualized language models (e.g., ELMo, BERT) produce token-level representations that vary with context, and where word sense disambiguation is increasingly approached through nearest-neighbor decisions in embedding space or by aligning contextual vectors to knowledge-base senses. Finally, the paper outlines unresolved questions on sense granularity, interpretability, and the relation between psychologically plausible exemplars and engineering-oriented embeddings, arguing that the most productive view treats "senses" as emergent clusters within a continuum shaped by frequency, discourse goals, and domain constraints.

Keywords: Exemplar-based semantics, usage-based lexical meaning, contextual embeddings, polysemy, word sense disambiguation, word sense induction, sense inventories

Introduction

Prototype-based categorization has been influential in lexical semantics because it offers a principled way to describe central and peripheral meanings within a word's semantic network. Yet many empirical and methodological debates about polysemy suggest that not every word behaves as if it had a stable "center" from which all extensions radiate. In some cases, related senses appear only weakly connected, and speakers may treat them as separate enough that a single core meaning becomes hard to defend. Experimental work on polysemous words has provided evidence that at least some sense distinctions are represented separately and that any shared "core" can be minimal, which makes purely prototype-centered analyses less straightforward for these items.

A natural adjacent perspective is to shift attention from prototypes to exemplars. Exemplar-based theories propose that categories are represented through remembered instances and that classification depends on similarity to these stored cases. In this view, "category structure" emerges from the distribution of exemplars rather than from a fixed summary representation. Classic formulations such as Medin and Schaffer's context theory and Nosofsky's exemplar

ISSN (E): 2942-190X Open Access| Peer Reviewed

modeling tradition provide formal mechanisms where similarity and selective attention determine how new stimuli are categorized.

When transported to lexical meaning, exemplar logic supports a usage-centered ontology: the basic objects are not pre-given senses but actual occurrences of words in context, and "senses" are clusters we infer for particular descriptive or applied purposes. This resonates with a major critique from lexicography and corpus linguistics that word senses are often "slippery" and become stable only relative to a task and a set of usage examples.

Discussion

A key advantage of exemplar-based thinking is that it treats variability as fundamental rather than exceptional. For many words, especially high-frequency items, meaning is not well captured by a small list of discrete senses. Instead, interpretation depends on contextual constraints (syntactic frame, collocates, discourse topic, speaker intentions) and on how similar the current usage is to previous experiences with the word. In an exemplar view, a word's "meaning" is not a single object but a structured memory of uses, with dense regions (frequent, conventional patterns) and sparse regions (rare, creative, domain-specific patterns). This picture fits well with the idea that senses are abstractions from clustered citations, not entities that exist independently of usage.

Exemplar models in cognitive science also clarify why sharp boundaries between polysemy and homonymy are difficult to maintain. If a learner stores instances and organizes them by similarity, then the "distance" between clusters matters more than a binary label. Under this logic, strongly separated clusters behave like homonymy; while overlapping or bridged clusters behave like polysemy, and many real cases fall between. This gradient perspective is supported by research showing that similarity-based mechanisms and attention play central roles in categorization, and that the apparent discreteness of categories can be an emergent effect of how exemplars distribute in psychological space.

However, exemplar approaches are not simply the mirror image of prototype theory; they make different commitments about what is stored and what is computed. Medin and Schaffer's context theory emphasizes how classification may be determined by similarity to stored instances under contextual weighting, while Nosofsky's framework connects similarity, attention, and decision rules into a quantitative account that can be fit to behavioral data. These ideas also motivated connectionist exemplar models such as ALCOVE, which incorporates learned attention to dimensions and can model category learning dynamics over time.

For lexical semantics, the central question becomes: what is an "exemplar" of a word sense? One natural answer is that the exemplar is a token usage – an attested instance in a particular context. This moves analysis toward corpora and toward methods that cluster usages by distributional similarity. Yet it also exposes the sense granularity problem: how many clusters should we posit, and at what level of abstraction? WordNet and related resources provide explicit sense inventories, which are extremely useful for standardizing annotation and evaluation, but their fine granularity can diverge from how speakers naturally group meanings in everyday processing.

This tension is visible in the computational literature on word sense disambiguation (WSD). WSD is typically defined as selecting the correct sense label for a word in context from a predefined inventory, and survey work has long emphasized how the problem is shaped by the chosen sense inventory and evaluation setup. A major step toward comparability was the development of unified evaluation frameworks, which make explicit that performance differences can reflect not only modeling choices but also dataset and annotation decisions.

Western European Journal of Linguistics and Education
Volume 4, Issue 2, February 2026
https://westerneuropeanstudies.com/index.php/2
ISSN (E): 2942-190X Open Access| Peer Reviewed

In parallel, computational semantics has increasingly adopted representations that look strikingly exemplar-like. Contextualized language models such as ELMo and BERT produce token-level vectors whose geometry varies with context; in practical terms, a word type no longer has one representation but many – one per occurrence. The original ELMo work explicitly frames its goal as capturing how word use varies across contexts, including polysemy, by learning representations from deep bidirectional language modeling. BERT similarly learns contextual representations through masked language modeling and bidirectionality.

Once meanings are represented at the token level, a direct exemplar-style decision rule becomes plausible: interpret a new token by comparing it to stored (or labeled) token vectors. This intuition underlies approaches that perform WSD via nearest-neighbor classification in contextual embedding space, treating previous labeled occurrences as exemplars for senses. Work explicitly asking whether "BERT makes any sense" reports that contextualized embeddings can separate polysemous words into sense regions and that simple similarity-based classifiers can reach strong results on standard benchmarks.

At the same time, not all contextual representations behave equally "sense-like." Analyses of contextualized embeddings show that context sensitivity changes by layer and that embedding spaces exhibit anisotropy, complicating naïve similarity judgments. The question of whether contextual models yield infinitely many idiosyncratic token meanings or instead map tokens onto a smaller set of sense-like attractors has become a research topic in its own right.

From the perspective of sense representations, this development is often described as a response to the "meaning conflation deficiency" of single-vector word embeddings. A major survey on sense embeddings argues that representing a word with all its meanings in one vector blurs distinctions that matter for downstream tasks, motivating approaches that shift from word-level to sense-level representations, either via clustering (unsupervised) or via lexical knowledge bases (knowledge-based).

A practical compromise has emerged: rather than choosing between purely usage-based clustering and fixed lexical inventories, some methods align contextual token vectors with knowledge-base senses. For instance, approaches to contextualized sense embeddings aim to place sense representations (e.g., from a lexical resource) into the same space as contextual word vectors, enabling models to select a sense that is both interpretable and context-compatible.

This is where exemplar semantics also connects to a different tradition in lexical theory: generative and underspecification approaches. Pustejovsky's Generative Lexicon argues that lexical meaning supports systematic sense generation via typed structures and compositional mechanisms, emphasizing the "multiplicity of word meaning" and how context can coerce or refine interpretation. While this tradition is not exemplar-based in a memory-first way, it converges on a similar conclusion: meanings in context are not always best treated as a lookup of a fixed sense list.

A balanced synthesis is to treat lexical meaning as having both memory-based and rule-/structure-based components. Exemplar mechanisms plausibly account for fast, experience-driven expectations (what sense is most likely here, given similar past contexts), while generative mechanisms account for systematic productive reinterpretations that go beyond memorized instances. In real language use, both are likely active: frequent, conventional uses form dense exemplar clouds, and compositional constraints shape which regions of that cloud are activated or how new regions are constructed.

Methodologically, exemplar-oriented lexical semantics encourages three research habits. First, it treats corpus evidence as primary: senses should be justified by attested usage distributions rather than by introspection alone. Second, it requires clarity about task goals: the "right" sense partition for lexicography, translation, language teaching, and WSD may legitimately differ. Third, it recommends evaluation strategies that go beyond sense-label accuracy and include analyses of cluster stability across domains, interpretability, and the ability to generalize to unseen contexts – issues strongly emphasized in modern work on sense and contextual representations.

Conclusion

A prototype-based model explains many facts about lexical categories, especially centrality and graded membership, but an exemplar-based and usage-centered perspective offers a complementary lens that is often closer to the data reality of polysemy. Cognitive categorization research provides formal tools for understanding how categories can be built from remembered instances, while corpus linguistics reminds us that "senses" are frequently abstractions constructed for particular purposes rather than fixed natural kinds. Contemporary NLP, through contextualized embeddings and embedding-space WSD, has effectively operationalized a form of exemplar semantics: words are represented as token-specific vectors, and meaning selection can be approximated by similarity to prior instances or to aligned sense representations. The next step is not to declare a winner among prototypes, exemplars, inventories, or generative structures, but to specify which combination best fits a given explanatory target: psychological plausibility, lexicographic usability, or computational performance.

References

1. Camacho-Collados, J., & Pilehvar, M. T. (2018). From word to sense embeddings: A survey on vector representations of meaning. *Journal of Artificial Intelligence Research, 63*, 743–788.
2. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT 2019* (pp. 4171–4186). Association for Computational Linguistics.
3. Ethayarajh, K. (2019). How contextual are contextualized word representations? Dissecting the geometry of BERT. In *Proceedings of EMNLP-IJCNLP 2019* (pp. 2769–2783). Association for Computational Linguistics.
4. Hadiwinoto, C., Ng, H. T., & Gan, W. C. (2019). Improved word sense disambiguation using pre-trained contextualized word representations. In *Proceedings of EMNLP-IJCNLP 2019* (pp. 5297–5306). Association for Computational Linguistics.
5. Kilgarriff, A. (1997). I don't believe in word senses. *Computers and the Humanities, 31*, 91–113.
6. Klein, D. E., & Murphy, G. L. (2001). The representation of polysemous words. *Journal of Memory and Language, 45*(2), 259–282.
7. Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99*(1), 22–44.
8. Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85*(3), 207–238.

9. Miller, G. A. (1995). WordNet: A lexical database for English. *Communications of the ACM, 38*(11), 39–41.
10. Navigli, R. (2009). Word sense disambiguation: A survey. *ACM Computing Surveys, 41*(2), Article 10.
11. Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115*(1), 39–61.
12. Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. In *Proceedings of NAACL-HLT 2018* (pp. 2227–2237). Association for Computational Linguistics.
13. Pustejovsky, J. (1995). *The generative lexicon*. MIT Press.
14. Pustejovsky, J. (1998). The semantics of lexical underspecification. *Folia Linguistica, 32*, 323–347.
15. Raganato, A., Camacho-Collados, J., & Navigli, R. (2017). Word sense disambiguation: A unified evaluation framework and empirical comparison. In *Proceedings of EACL 2017* (pp. 99–110). Association for Computational Linguistics.
16. Scarlini, B., Pasini, T., & Navigli, R. (2020). Contextualized sense embeddings for all-round word sense disambiguation. In *Proceedings of EMNLP 2020* (pp. 10399–10413). Association for Computational Linguistics.
17. Wiedemann, G., Remus, S., Chawla, A., & Biemann, C. (2019). Does BERT make any sense? Interpretable word sense disambiguation with contextualized embeddings. *arXiv preprint*.
18. Loureiro, D., Rezaee, K., & Camacho-Collados, J. (2021). Analysis and evaluation of language models for word sense disambiguation. *Computational Linguistics*.