

# Do Pretrained Contextual Language Models Distinguish between Hebrew Homograph Analyses?

Avi Shmidman<sup>1,2,†</sup> Cheyn Shmuel Shmidman<sup>2,‡</sup>  
Dan Bareket<sup>1,‡</sup> Moshe Koppel<sup>1,2,‡</sup> Reut Tsarfaty<sup>1,†</sup>

<sup>1</sup>Bar Ilan University      <sup>2</sup>ICTA / Jerusalem, Israel

<sup>†</sup>{avi.shmidman, reut.tsarfaty}@biu.ac.il

<sup>‡</sup>{cheynshmuel, dbareket, moishk}@gmail.com

## Abstract

Semitic morphologically-rich languages (MRLs) are characterized by extreme word ambiguity. Because most vowels are omitted in standard texts, many of the words are homographs with multiple possible analyses, each with a different pronunciation and different morphosyntactic properties. This ambiguity goes *beyond* word-sense disambiguation (WSD), and may include token segmentation into multiple word units. Previous research on MRLs claimed that standardly trained pre-trained language models (PLMs) based on word-pieces may not sufficiently capture the internal structure of such tokens in order to distinguish between these analyses. Taking Hebrew as a case study, we investigate the extent to which Hebrew homographs can be disambiguated and analyzed using PLMs. We evaluate all existing models for contextualized Hebrew embeddings on a novel Hebrew homograph challenge sets that we deliver. Our empirical results demonstrate that contemporary Hebrew contextualized embeddings outperform non-contextualized embeddings; and that they are most effective for disambiguating segmentation and morphosyntactic features, less so regarding pure word-sense disambiguation. We show that these embeddings are more effective when the number of word-piece splits is limited, and they are more effective for 2-way and 3-way ambiguities than for 4-way ambiguity. We show that the embeddings are equally effective for homographs of both balanced and skewed distributions, whether calculated as masked or unmasked tokens. Finally, we show that these embeddings are as effective for homograph disambiguation with extensive supervised training as with a few-shot setup.

## 1 Introduction

Semitic morphologically-rich languages (MRLs) such as Arabic, Hebrew, and Aramaic are characterized by extreme ambiguity at the word level

(Wintner, 2014; Tsarfaty et al., 2020). In a standard text, many (and often most) of the words will be homographs with multiple possible analyses. The high ambiguity derives from several factors. First, prepositions, conjunctions, accusative pronouns, and possessive pronouns are often seamlessly affixed to words. Next, vowels are generally omitted in written texts. Finally, proper nouns are not differentiated from common nouns (no capital letters).

The task of distinguishing between Hebrew homograph analyses is related to the general task of Word Sense Disambiguation (WSD) (Agirre and Edmonds, 2006; Navigli, 2009), yet it is more challenging. In the standard case of WSD, a single orthographic form is associated with a single word that can be analyzed in terms of two or more senses; also, the analyses are generally pronounced identically, and often have the same morphosyntactic properties (e.g bank of a river vs. savings bank). In contrast, in Semitic languages, the need for disambiguation often goes beyond a determination of sense. Hebrew word ambiguities can be divided into three primary categories (Table 1): 1. Segmentation ambiguities, in which a given orthographic form may (or may not) be segmented into multiple word units each bearing its own role (POS tag) in the sentence. 2. Morphosyntactic ambiguities, in which the segmentation of the form is not ambiguous, but multiple analyses of the word reflect different morphosyntactic properties of each word unit(s). 3. Sense ambiguities (the aforementioned standard case of WSD), in which the analyses of the unit(s) do not differ in their morphosyntactic properties, but rather in their *sense*. One orthographic form may exhibit multiple types of ambiguity simultaneously.

Pretrained contextualized language models with standard word-piece tokenization mechanisms have been shown to excel at WSD in English and other Indo-European languages (Yaghoobzadeh et al., 2019). However, for Hebrew and other semitic lan-

Type	Form	Word (translation)	Morphology
Segmentation	הקפה	הַקֶּפֶה (the+coffee)	DET + Noun [M,S,abs]
		הַקֶּפֶד (credit)	Noun [F,S,abs]
	שארף	שָׂאֵף (for+even)	Sconj + Cconj
Morph	אלימות	אֶלִּימוּת (violent)	Adj [F,P,abs]
		אַלְיָמוּת (violence)	Noun [F,S,abs/cons]
	הרים	הִרִים (he lifted)	Verb [M,S,3,PAST]
		הַרִים (mountains)	Noun [M,P,abs]
		הַרִים (mountains)	Noun [M,P,abs]
Semantic	הזמר	הַזֶּמֶר (the+song)	DET + Noun [M,S,abs]
		הַזֹּמֵר (the+singer)	DET + Noun [M,S,abs]
	הסופר	הַסּוֹפֵר (the+author)	DET + Noun [M,S,abs]
		הַסּוֹפֵר (the+market)	DET + Noun [M,S,abs]

Table 1: Examples of Hebrew ambiguity types

guages it has been argued that such models would not sufficiently capture the structure of MRLs in order to distinguish between internally-complex homograph analyses (Klein and Tsarfaty, 2020; Tsarfaty et al., 2020). In this work, we take Modern Hebrew, a Semitic language with rich and highly ambiguous morphology, as a case study, and we investigate the extent to which homographs can be disambiguated by contextualized embeddings, regarding all three levels of ambiguity. Regarding Arabic — a sister language to Hebrew — a wide survey of WSD methods is presented by [Abderahim and Abderrahim \(2022\)](#). They raise the possibility of utilizing pretrained contextualized embeddings, yet leave its evaluation to future work.<sup>1</sup>

Hebrew is a particularly challenging language on which to perform a homograph disambiguation due to the limited available corpora. First of all, currently existing Hebrew treebanks are severely limited in size, such that most of the words in the language are not amply represented. Furthermore, even regarding common Hebrew words, this corpus is problematic, because the nature of language is such that many homographs are skewed in their distribution; thus, even if the primary analysis is sufficiently represented within a tagged corpus, the secondary analysis will often be hopelessly under-represented. For instance, one common Hebrew homograph is **מהם** (*mhm*), which can be analyzed as a preposition with pronominal suffix, or as an interrogative. The ratio of these two analyses in naturally-occurring Hebrew text is over 50:1; thus, occurrences of the secondary analysis within existing tagged corpora are insufficient.

In theory, these homograph ambiguities could be addressed using POS tagging systems. For instance, [Habash and Rambow \(2005\)](#) consider the

<sup>1</sup>Additional studies in Arabic WSD include [Merhben et al. \(2010\)](#), [Merhbene et al. \(2013\)](#), and [Shah et al. \(2010\)](#).

use of a morphological tagging system to solve WSD in Arabic. A number of Hebrew POS tagging systems have been published as well ([Yona and Wintner, 2005](#); [Adler and Elhadad, 2006](#); [Shacham and Wintner, 2007](#)). The current SOTA for Hebrew POS tagging is the YAP morpho-syntactic parser ([Tsarfaty et al., 2019](#)). However, as we have shown in a previous study ([Shmidman et al., 2020](#), p. 3318, table 2), although YAP produces high accuracy overall on normal Hebrew text, its scores drop drastically regarding homographs of skewed distribution.

For analogous cases of skewed distribution in other languages, researchers have proposed the creation of dedicated challenge sets, containing hard-to-classify sentences not easily found in naturally-occurring text ([Gardner et al., 2020](#); [Elkahky et al., 2018](#)). In the aforementioned previous study, we produced 22 such challenge sets for Hebrew homographs, and demonstrated that a Bi-LSTM of non-contextualized embeddings can obtain high accuracy on this task, establishing the current SOTA for Hebrew homograph disambiguation ([Shmidman et al., 2020](#)). In this paper, we extend the investigation by considering whether contextualized embeddings from pretrained language models (PLMs) can provide a more optimal solution. We consider all existing contextualized Hebrew PLMs: multilingual BERT ("mBERT") ([Devlin et al., 2019](#)); HeBERT ([Chriqui and Yahav, 2021](#)); and AlephBERT ([Seker et al., 2021](#)) (Table 2). Moreover, we evaluate and verify these on a new dataset, substantially larger than all previous datasets for Hebrew homograph disambiguation.

Our experiments demonstrate that contextualized PLMs pre-trained on sufficiently large unlabeled data and vocabulary size are excellent at disambiguating the word-internal structures of homographs, yet face some challenge with pure sense disambiguation. We show the efficacy of these models in cases of homographs with skewed distribution, and in a few-shot setup. Finally, we establish new state-of-the-art results on the challenging task of homograph disambiguation for a morphologically-rich language printed without vowels, along with a novel benchmark for assessing the morphological reach of future PLMs in Hebrew.

## 2 The Data

The challenge sets for Hebrew homograph disambiguation from our previous study were limited in

Model	Vocab (Heb. tokens)	Corpus Size (Heb. sentences)
mBERT	2.5K	6.3M
HeBERT	30K	27.2M
AlephBERT	52K	98.7M

Table 2: Comparison of available Hebrew BERT models

number (only 22 sets) and insufficiently representative regarding types of ambiguities; only one of the sets was a prefix-segmentation ambiguity. Further, they were limited to binary cases, where only two analyses exist.

In contrast, for this study we employed field experts to choose the most critical homographs in the language. The experts chose 75 homographs from a list of the 3600 most frequent words in the language, balancing frequency of word occurrence with practical need for its disambiguation. All of the homographs occur with a minimum frequency of 27 words per million in naturally occurring Hebrew text. Our challenge sets include homographs with 2-5 possible analyses. Our sets contain a wide representation of segmentation ambiguities (15 in number), as well as 5 cases of purely semantic ambiguities. For each of the 75 homographs, we collect hundreds of naturally-occurring sentences attesting to each analysis. In almost all cases, we succeed in collecting 1000 sentences for the primary analysis, at least 500 sentences for the secondary analysis, and at least 250 for each additional analysis. The sentences were culled from newspapers, Wikipedia, literature, and social media. We employed a team of annotators who chose the relevant homograph analysis for each case.<sup>2</sup> All in all, our 75 challenge sets contain 150K tagged sentences. The full list of homographs and analyses is provided in Appendix A.<sup>3</sup>

### 3 Experimental Setup

To evaluate the ability of pre-trained language models (PLMs) to disambiguate the in-context analyses of morphologically rich and highly ambiguous homographs in Hebrew, we adopt a "word expert" approach, producing dedicated classifiers for each individual homograph (Zhao et al., 2020).

We use two types of PLMs, contextualized and non-contextualized. For the non-contextualized case, we replicate our previous method detailed

<sup>2</sup>The annotation process is detailed in Appendix B.

<sup>3</sup>The dataset is downloadable at: [https://github.com/Dicta-Israel-Center-for-Text-Analysis/EACL\\_2023](https://github.com/Dicta-Israel-Center-for-Text-Analysis/EACL_2023)

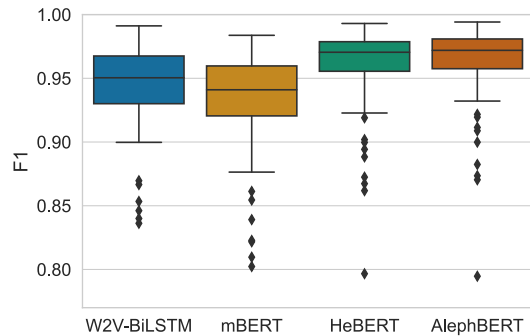


Figure 1: Comparison of previous SOTA (w2v-based Bi-LSTM method) versus BERT-based approaches.

in Shmidman et al. (2020). For each training example, we use a BiLSTM on top of the word2vec embeddings of all of the words in the sentence (other than the homograph itself) to produce an encoding for disambiguation.<sup>4</sup> An MLP is trained to predict the correct homograph analysis based on this encoding.<sup>5</sup> For the contextualized case, we run the sentence through a pretrained contextualized language model and retrieve the 768-dimension embedding representing the homograph in question. An MLP is trained to predict the correct analysis based on the homographs embeddings alone. In the standard "unmasked" scenario, the sentence is fed into the model as is, including the homograph in question. In the "masked" scenario, the homograph is replaced with a [MASK] token.

We evaluate the performance of each given method on each given challenge set using 10-fold cross-validation. We calculate an F1 score for each homograph analysis, based upon the precision and recall scores micro-averaged across all folds. We then calculate the macro-average of the F1 scores for all possible analyses for a given homograph, and this is the score reported in the charts herein.

### 4 Results and Analysis

**Standard (Unmasked) Scenario** Figure 1 presents the cumulative F1 score obtained by the models for all challenge sets. Our results show that HeBERT and AlephBERT far outperform mBERT, with AlephBERT achieving the higher score. The poor performance of mBERT is likely due to its smaller pre-training data size and exceedingly lean Hebrew vocabulary (cf. Table 2). Furthermore, the

<sup>4</sup>We also tested fastText, but results were inferior.

<sup>5</sup>For implementation details, see Appendix D.

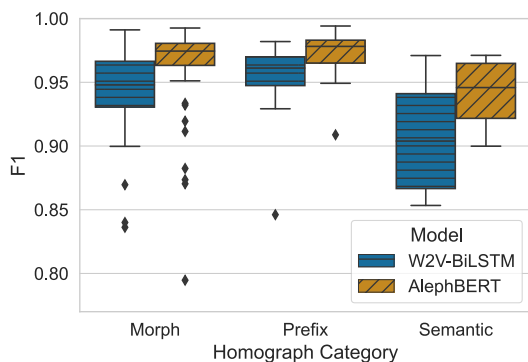


Figure 2: Disambiguation accuracy across various categories of homograph ambiguity.

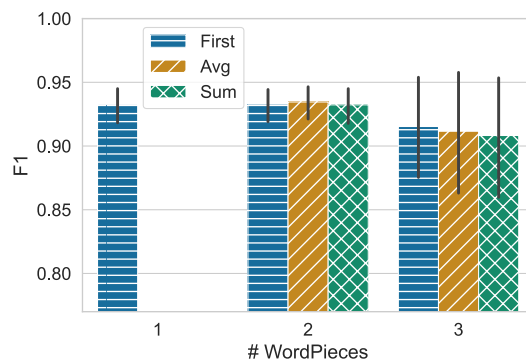


Figure 4: Disambiguation accuracy across varying degrees of word piece splits within the target homograph, using mBERT.




Figure 3: AlephBERT's disambiguation accuracy across homographs with differing counts of possible analyses.

HeBERT and AlephBERT models both substantially outperform the previous word2vec-based SOTA. It is thus apparent that *contextualized* language models do effectively capture Hebrew homograph distinctions, even those based on word-pieces, even for an MRL, and they do so more effectively than non-contextualized models.

Figure 2 demonstrates AlephBERT's performance on different ambiguity types. AlephBERT performs equally well on cases of segmentation ambiguity and morphosyntactic ambiguity. In contrast, when it comes to ambiguities that are purely semantic, the scores are noticeably lower. This is in line with the findings of Ettinger (2020), who shows that BERT is stronger with syntax than semantics; Goldberg (2019) also notes BERT's strong syntactic abilities. Interestingly, the same gap exists with the W2V-based method. Thus, both contextualized and non-contextualized embeddings struggle to differentiate between senses which are morphologically equivalent. Although such cases are only of minimal import when it comes to sentence parsing,

they are critical for downstream tasks such as coreference resolution and relation extraction. It thus remains a desideratum to improve disambiguation of purely semantic Hebrew homographs.

The results in Figure 3 demonstrate that AlephBERT performs equally well on cases of binary homographs as on cases of three-way homograph classification. However, when faced with cases of 4-way or 5-way classification, accuracy declines.

**The Effect of Word-Pieces** Previous studies have hypothesized that word-pieces are not adequate for capturing complex morphosyntactic structures due to arbitrary (non-linguistic) word-splits. To probe into this we investigate the question, do such splits affect performance. Our 75 homographs are all treated as single tokens in HeBERT and AlephBERT. However, many of the homographs are broken up into word pieces in mBERT, due to its meager Hebrew vocabulary. We thus compare mBERT's results on words treated as single tokens versus those that are broken up into two or three pieces, which are aggregated using first, sum, or average of the vectors. With regard to cases of split words, we train models using three separate methods: providing the MLP with only the embedding of the first word piece; with an average of the word piece embeddings; or with the sum of the embeddings. As shown in Figure 4, the splitting of a homograph into three word-pieces appears to have a negative impact on the ability of the resulting embedding to differentiate between homograph analyses, for all aggregation methods.

**Masked Scenario** We consider whether AlephBERT embeddings are more effective if we replace the homograph word with [MASK] when running




Figure 5: Comparison of unmasked vs. masked scenarios, across varying degrees of skewed homographs. The ratio axis indicates the relative distribution between the least and most frequent analyses of each plotted homograph within naturally-occurring Hebrew text.

the challenge set sentences through AlephBERT. The motivation behind this experiment is that, as explained above, many of the homographs are skewed in their natural proportion. In such cases, we worry whether AlephBERT might be disproportionately influenced by the skewed distribution; replacing the word with [MASK] would prevent the model from being influenced as such. As shown in figure 5, AlephBERT achieves high scores both with balanced homographs as well as with homographs of highly skewed distribution. Using a [MASK] token instead of the actual word does not generally improve performance, whether or not the homographs are of skewed proportion.

**Few-Shot Scenarios** In our experiments thus far, the 10-fold cross-validation allows the MLP to leverage 90% of the data in each fold (hundreds of sentences for each analysis) in order to learn the difference between the analyses. We now consider whether the AlephBERT embeddings can suffice on a few-shot basis, where the training stage has access to only 100, 50, 25, 10 or even 5 examples of each analysis. In these cases, we train an MLP based only on these few samples, and we use the rest of the sentences for evaluation. Astoundingly, as demonstrated in Figure 6, the AlephBERT embeddings provide a highly accurate solution even on this few-shot basis. Even when training with only 5 examples of each homograph analysis, AlephBERT reaches an accuracy that is not far below the accuracy achieved when performing full 10-fold CV across hundreds of sentences of each analysis.

**Probing Scenarios** Finally, we probe the pre-trained AlephBERT embeddings (Yaghoobzadeh et al., 2019; Tenney et al., 2019; Klafka and Ettinger, 2020; Belinkov, 2021) to see whether in




Figure 6: Use of AlephBERT embeddings to differentiate between homographs on a few-shot basis, contrasted with scores from the full 10-fold CV ("All").

and of themselves they reflect clusters which correspond to different homograph analyses. We skip the MLP, and instead use the raw embeddings directly, classifying sentences based on their proximity to the centroid of the training samples for each homograph analysis. As shown in the orange bars in Figure 6, this method generally does not perform as well as the MLP-based method; however, the degradation is limited to only a few percentage points, indicating that the raw embeddings are generally clustered in groups which indeed reflect the distinctions between the analyses.

## 5 Conclusion

In this study we have utilized a wide-ranging collection of Hebrew homograph challenge sets in order to evaluate the extent to which raw contextualized embeddings can be leveraged to disambiguate Hebrew homographs. We found that contextualized embeddings can effectively disambiguate analyses of homographs, much more so than non-contextualized ones, regarding multiple types of ambiguity: segmentation, morphosyntactic and sense. Yet, efficacy on pure sense ambiguity is lower than on the other two types. Additionally, an increasing number of splits, or an increasing number of different possible analyses of a token, each lower efficacy. Finally, we found that contextualized embeddings can function effectively for this purpose on a few-shot basis, with as little as 5 examples of each analysis. This indicates that with relatively modest effort, highly ambiguous homographs may be effectively treated.

## Limitations

One of the major strengths of this paper is its new and comprehensive dataset for the training and benchmarking Hebrew homograph disambiguation. The dataset is incomparable in size, quality and balance to all previous Hebrew homograph datasets; we have made every effort to be as inclusive as possible in the creation of the dataset, making sure to include data from a widely diverse set of genres. Nevertheless, a perennial challenge in corpus-based studies is that the lion's share of the available data tends to be authored by male writers. In order to offset this bias, we bolstered our corpus with a large corpus of texts specifically taken from blog sites devoted entirely to female bloggers. Even so, we cannot escape the fact that female writing and feminine conjugations are underrepresented in our dataset.

A further limitation derives from our filter for sentences with offensive language. We perceived early on that our human annotators were not comfortable tagging sentences with offensive language, and we therefore took steps to remove such sentences from our corpus. Nevertheless, this means that our resulting dataset is limited in that it does not properly reflect the use of offensive language in naturally-occurring Hebrew sentences. Furthermore, our resulting tests and reported scores may not accurately reflect the performance of our models when applied to sentences with offensive language.

## Ethics Statement

**Creation of the Dataset** As noted, our dataset contains over 150K sentences in all. Every sentence was reviewed and tagged by our team of human annotators, who chose the relevant homograph analysis for each instance of each of our 75 homographs. Our annotator team included members of diverse genders and sexual orientations. They were paid hourly wages with legal pay stubs. Their hourly wage was well above the minimum wage required by law. The entirety of the dataset will be made available for research purposes with the acceptance of this article, together with the tagging information.

**Risks of the Research** Ultimately, this data will enable end-users to automatically diacritize and parse large corpora of Hebrew text. For the most part, this will provide a beneficial contribution to

the world: for the visually impaired, this technology will enable the development of more precise text-to-speech products; teachers will be able to provide children and second-language learners with accessible diacritized texts; and humanities and linguistics researchers can bolster their research with big-data analysis. However, there also is a risk of nefarious use, for instance, if an end user were to leverage these capabilities in order to produce anonymous texts or recordings containing threats to life, liberty, or happiness.

## Acknowledgments

The work of the third and last authors has been funded by a grant from the Israeli Ministry of Science and Technology (MOST) grant no. 3-17992.

## References

- Mohammed Alaeddine Abderrahim and Mohammed El-Amine Abderrahim. 2022. [Arabic word sense disambiguation for information retrieval](#). *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 21(4).
- Meni Adler and Michael Elhadad. 2006. [An unsupervised morpheme-based hmm for hebrew morphological disambiguation](#). In *COLING/ACL 2006 - 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference*, COLING/ACL 2006 - 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, pages 665–672, United States. Association for Computational Linguistics (ACL). 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, COLING/ACL 2006 ; Conference date: 17-07-2006 Through 21-07-2006.
- Eneko Agirre and Philip Edmonds, editors. 2006. [Word Sense Disambiguation: Algorithms and Applications](#), volume 33 of *Text, Speech and Language Technology*. Springer.
- Yonatan Belinkov. 2021. [Probing Classifiers: Promises, Shortcomings, and Advances](#). *Computational Linguistics*, pages 1–13.
- Avihay Chriqui and Inbal Yahav. 2021. [Hebert hebemo: a hebrew bert model and a tool for polarity analysis and emotion recognition](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).

- Ali Elkahky, Kellie Webster, Daniel Andor, and Emily Pitler. 2018. [A challenge set and methods for noun-verb ambiguity](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2562–2572, Brussels, Belgium. Association for Computational Linguistics.
- Allyson Ettinger. 2020. [What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Matt Gardner, Yoav Artzi, Victoria Basmova, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hanna Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. [Evaluating models’ local decision boundaries via contrast sets](#).
- Yoav Goldberg. 2019. [Assessing bert’s syntactic abilities](#). *CoRR*, abs/1901.05287.
- Nizar Habash and Owen Rambow. 2005. [Arabic tokenization, part-of-speech tagging and morphological disambiguation in one fell swoop](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pages 573–580, Ann Arbor, Michigan. Association for Computational Linguistics.
- Josef Klafka and Allyson Ettinger. 2020. [Spying on your neighbors: Fine-grained probing of contextual embeddings for information about surrounding words](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4801–4811, Online. Association for Computational Linguistics.
- Stav Klein and Reut Tsarfaty. 2020. [Getting the ##life out of living: How adequate are word-pieces for modelling complex morphology?](#) In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 204–209, Online. Association for Computational Linguistics.
- Laroussi Merhben, Zouaghi Anis, and Mounir Zrigui. 2010. [Ambiguous arabic words disambiguation](#). pages 157–164.
- Laroussi Merhbene, Anis Zouaghi, and Mounir Zrigui. 2013. [A semi-supervised method for Arabic word sense disambiguation using a weighted directed graph](#). In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 1027–1031, Nagoya, Japan. Asian Federation of Natural Language Processing.
- Roberto Navigli. 2009. [Word sense disambiguation: A survey](#). *ACM Comput. Surv.*, 41.
- Amit Seker, Elron Bandel, Dan Bareket, Idan Brusilovsky, Refael Shaked Greenfeld, and Reut Tsarfaty. 2021. [Alephbert: a hebrew large pre-trained language model to start-off your hebrew nlp application with](#).
- Danny Shacham and Shuly Wintner. 2007. [Morphological disambiguation of Hebrew: A case study in classifier combination](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 439–447, Prague, Czech Republic. Association for Computational Linguistics.
- Rushin Shah, Paramveer S. Dhillon, Mark Liberman, Dean Foster, Mohamed Maamouri, and Lyle Ungar. 2010. [A new approach to lexical disambiguation of Arabic text](#). In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 725–735, Cambridge, MA. Association for Computational Linguistics.
- Avi Shmidman, Joshua Guedalia, Shaltiel Shmidman, Moshe Koppel, and Reut Tsarfaty. 2020. [A novel challenge set for Hebrew morphological disambiguation and diacritics restoration](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3316–3326, Online. Association for Computational Linguistics.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R. Bowman, Dipanjan Das, and Ellie Pavlick. 2019. [What do you learn from context? probing for sentence structure in contextualized word representations](#).
- Reut Tsarfaty, Dan Bareket, Stav Klein, and Amit Seker. 2020. [From SPMRL to NMRL: What did we learn \(and unlearn\) in a decade of parsing morphologically-rich languages \(MRLs\)?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7396–7408, Online. Association for Computational Linguistics.
- Reut Tsarfaty, Amit Seker, Shoval Sadde, and Stav Klein. 2019. [What’s wrong with hebrew nlp? and how to make it right](#).
- Shuly Wintner. 2014. [Morphological processing of semitic languages](#). In *NLP of Semitic Languages*.
- Yadollah Yaghoobzadeh, Katharina Kann, T. J. Hazen, Eneko Agirre, and Hinrich Schütze. 2019. [Probing for semantic classes: Diagnosing the meaning content of word embeddings](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5740–5753, Florence, Italy. Association for Computational Linguistics.
- Shlomo Yona and Shuly Wintner. 2005. [A finite-state morphological grammar of Hebrew](#). In *Proceedings of the ACL Workshop on Computational Approaches to Semitic Languages*, pages 9–16, Ann Arbor, Michigan. Association for Computational Linguistics.

Mengjie Zhao, Philipp Dufter, Yadollah Yaghoobzadeh, and Hinrich Schütze. 2020. [Quantifying the contextualization of word representations with semantic class probing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1219–1234, Online. Association for Computational Linguistics.



## A Appendix A: Table of Homographs

We present three tables of homographs, corresponding to the three categories of homographs discussed within the paper (segmentation ambiguity, morphosyntactic ambiguity, and semantic ambiguity).<sup>6</sup> In each table, the first column ("form") indicates the homograph as it is found in naturally-occurring non-diacritized Hebrew text.<sup>7</sup> The second column ("word") indicates the possible diacritizations of each form. In cases where the diacritized word includes an attached prefix, a plus sign indicates the segmentation point between the prefix letters and the primary word. Regarding all homographs considered in this paper, different segmentation options are diacritized differently. Thus, for each sentence in the dataset, our human annotators were asked simply to choose the correct diacritization for the target homograph (that is, they were asked to choose among the options listed in the "word" column). There was no need for the annotators to separately tag the segmentation, because in all cases the choice of diacritization itself indicates the segmentation.<sup>8</sup> The third column indicates the morphology of each of the possible diacritizations.<sup>9</sup> The fourth column lists the translation.<sup>10</sup> Within each table the homographs are listed alphabetically.<sup>11</sup>

### A.1 Homographs with Segmentation Ambiguity

Form	Word	Morphology <sup>12</sup>	Translation	Sentences
האם	הָאֵם	Interrogative	does	1,000
	הַאֵם	Det + Noun [F,S,abs]	the + mother	1,000
המראה	הַמְרָאָה	SConj + Ptcp [M,S] / Det + Noun [M,S,abs]	that + indicates / the + sight	1,000
	הַמְרָאָה	SConj + Ptcp [F,S] / Det + Noun [F,S,abs]	that + indicates / the + mirror	882
	הַמְרָאָה	Noun [F,S,abs]	takeoff	735

<sup>6</sup>It should be noted that the higher level of ambiguity are supersets of the lower levels: segmentation ambiguities generally entail differences on the morphosyntactic and semantic levels as well, and morphosyntactic ambiguities generally also entail semantic ambiguities. Furthermore, because many of the homographs admit to more than two analyses, it is often the case that a subset of the analyses may form a lower level of ambiguity (e.g., just a semantic ambiguity), while other analyses form a higher level of ambiguity (e.g. segmentation ambiguity). For the purposes of this paper, we categorize each homograph according to the highest level of ambiguity involved: First, if a segmentation ambiguity is indicated anywhere across the possible analyses, then we include the homograph in the "segmentation ambiguity" category. Next, if there is no segmentation ambiguity, but a morphosyntactic ambiguity is indicated anywhere across the possible analyses, then we include the homograph in the "morphosyntactic ambiguity" category. Finally, if the analyses all differ only on the semantic level, then we include the homograph in the "semantic ambiguity" category.

<sup>7</sup>The ranking of analyses is based on a frequency analysis of our in-house annotated corpus. It is worth emphasizing that the paper as a whole relates to each of the 75 homographs specifically as they are spelled in this list, and does not relate to cases where further prefixes are attached to the homographs. As a result, the frequency analysis may sometimes seem counterintuitive. For instance, regarding the form רֵאשִׁי, a native Hebrew speaker might intuit that the adjectival form is primary (רֵאשִׁי). However, in practice, that sense is common only when prefixed with a definitive marker (הַרֵאשִׁי). In contrast, the homograph considered here involves the form רֵאשִׁי as is, without any prefixes; in this case, the other analyses are far more common.

<sup>8</sup>It should be noted that in certain sentences, an exceedingly rare diacritization was warranted, which was not among the options listed in the "word" column. The annotators were instructed to tag such cases as "none of the above", and all such sentences were removed from the corpus. Similarly, some sentences do not provide enough context to determine the correct diacritization; the annotators were asked to tag such sentences as "unclear", and these sentences too were removed from the corpus.

<sup>9</sup>In most cases, a diacritized form has one specific morphological analysis. However, in other cases, the diacritized form can admit to multiple morphologies. In such cases, we list all of the practically relevant morphological analyses in the third column, separated by a slash (thus for instance in the case of זָר). Rare analyses which hardly ever occur in practice are not listed.

<sup>10</sup>Naturally, a given Hebrew term often captures a substantial range of potential English translations, and it would not be practical to list them all in this column; therefore, we generally present only a single representative translation.

<sup>11</sup>Ideally, we might have grouped the homographs based on part-of-speech instead. However, as can be seen from the following tables, the 75 homographs vary so widely in terms of the parts of speech that they can represent, such that alphabetical listing was deemed most useful.

<sup>12</sup>Key to morphology abbreviations: F=Feminine; M=Masculine; S=Singular; P=Plural; abs=absolute; cons=construct; 1,2,3=First Person, Second Person, Third Person; suf=suffix; Det = determiner; SConj=Subordinating Conjunction; Adj=Adjective; Conj=Conjunction; Prep=Preposition; Propn=Proper Noun; Adv=Adverb; Ptcp=Participle; ACC=Accusative marker.

הקפה	ה+קפה	Det + Noun [M,S,abs]	the + coffee	1,000
	הקפה	Noun [F,S,abs]	credit	500
הקשר	ה+קשר	Det + Noun [M,S,abs]	the + connection	1,000
	ה+קשר	Det + Noun [M,S,abs]	the + signaler	1,000
	הקשר	Noun [M,S,abs/cons]	context	588
השלמה	השלמה	Noun [F,S,abs]	completion	1,000
	ה+שלמה	Det + Adj [F,S]	the + complete	906
ועד	ו+עד	Conj + Prep	and + until	1,000
	ועד	Noun [M,S,abs/cons]	committee	519
לשם	לשם	Prep	for the purpose of	1,000
	ל+שם	Prep + Adverb	to + there	1,000
מבחינה	מ+בחינה	Prep + Noun [F,S,abs]	from + point of view	1,000
	מבחינה	Participle [F,S,abs]	she notices	607
מסיבות	מסיבות	Noun [F,P,abs]	parties	1,000
	מ+סיבות	Prep + Noun [F,P,abs/cons]	due to + reasons	1,000
מפתח	מפתח	Noun [M,S,abs]	key	1,000
	מפתח	Participle [M,S,abs]	develops / developer	329
	מ+מפתח	Prep + Noun [M,S,abs/cons]	from + opening	206
שאלה	שאלה	Noun [F,S,abs]	question	1,000
	שאלה	Verb [F,S,3,Past]	she asked	1,000
	ש+שאלה	SConj + Pronoun [MF,P,3]	that + these	228
שאף	ש+אף	SConj + CConj	for + even	1,000
	שאף	Verb [M,S,3,Past]	he aspired	665
שבה	ש+בה	Sconj + Prep [suf=F,S,3]	that + in it	1,000
	שבה	Verb [F,S,Present/Past]	she returns / she returned	1,000
שמן	שמן	Noun [M,S,abs/cons]	oil	1,000
	שמן	Adj [M,S,abs]	wide	335
	ש+מן	SConj + Prep	that + from	207
	שמן	Noun [M,S,abs,suf=F,P,3]	their name	149
שמר	שמר	Verb [M,S,3,Past]	he guarded	1,000
	שמר	Propn	Shemer	872
	ש+מר	SConj + Titular [M,S]	that + Mr.	224

## A.2 Homographs with Morphosyntactic Ambiguity

Form	Word	Morphology	Translation	Sentences
אהבה	אהבה	Noun [F,S,abs]	love	1,000
	אהבה	Verb [F,S,3,Past]	she loved	1,000
אוכל	אוכל	Noun [M,S,abs/cons]	food	748
	אוכל	Modal [MF,S,1,Future]	I can	729
	אוכל	Participle [M,S,abs/cons]	eats	640
אחדות	אחדות	Det [F,P,abs]	several	1,000
	אחדות	Noun [F,S,abs]	unity	1,000

אחיו	אָחָיו	Noun [MF,P,abs,suf=M,S,3]	his brothers	1,000
	אָחִיו	Noun [MF,S,abs,suf=M,S,3]	his brother	774
אלימות	אֲלִימוֹת	Adj [F,P]	violent	1,000
	אֲלִימוּת	Noun [F,S,abs/cons]	violence	1,000
אם	אִם	Conj	if	1,000
	אִמָּה	Noun [F,S,abs/cons]	mother	1,000
אמצעי	אֲמָצְעִי	Noun [M,P,cons]	centers of / methods of	1,000
	אֲמָצְעִי	Noun [M,S,abs] / Adj [M,S]	method / central	969
אמרה	אָמְרָה	Verb [F,S,3,Past]	she said	1,000
	אָמְרָה	Noun [F,S,abs]	a saying	343
אפשר	אֶפְשָׁר	Modal / Adv	possible	1,000
	אֶפְשָׁר	Verb [M,S,3,Past]	he allowed	511
אה	אֶת	ACC	accusative	1,000
	אַתָּה	Pronoun [F,S,2]	you	1,000
בהמשך	בְּהַמְשָׁךְ	Prep + Noun [M,S,cons]	in + continuation of	1,000
	בְּהַמְשָׁךְ	Prep [with Det] + Noun [M,S,abs]	in the + future	930
בחי	בְּחַיֵּי	Prep + Noun [M,P,cons]	in + lives of	1,000
	בְּחַיֵּי	Prep + Noun [M,P,abs,suf=MF,S,1]	in + my life	1,000
בעולם	בְּעוֹלָם	Prep [with Det] + Noun [M,S,abs]	in the + world	1,000
	בְּעוֹלָם	Prep + Noun [M,S,cons]	in + a world of	913
	בְּעוֹלָם	Prep + Noun [M,S,abs]	in + a world	485
בקרב	בְּקִרְבִּי	Prep + Noun [M,S,cons]	in + midst of	1,000
	בְּקִרְבִּי	Prep [with Det] + Noun [M,S,abs]	in the + battle	1,000
	בְּקִרְבִּי	Prep + Verb [Bare Infinitive]	in + approaching of	734
	בְּקִרְבִּי	Prep + Noun [M,S,abs]	in + a battle	256
גילו	גִּילּוֹ	Verb [MF,P,3,Past]	they discovered	1,000
	גִּילּוֹ	Noun [M,S,abs,suf=M,S,3]	his age	859
די	דִּי	Adverb	sufficiently	1,000
	דִּי	Det [cons]	enough of	828
	דִּי	Prefix	di-	808
הזקן	הַזָּקֵן	Det + Noun [M,S,abs]	the + beard	739
	הַזָּקֵן	Det + Adj [M,S] / Det + Noun [M,S,abs]	the + old / the + elderly man	533
החל	הִחֵל	Verb [M,S,3,Past]	he began	1,000
	הִחֵל	Verb [Bare Infinitive]	starting (from)	1,000
המשנה	הַמִּשְׁנָה	Det + Noun [M,S,abs]	the + deputy	1,000
	הַמִּשְׁנָה	Det + Propn [F,S,abs]	the + Mishna	1,000
הנחה	הַנְּחָה	Noun [F,S,abs]	placing	1,000
	הַנְּחָה	Verb [M,S,3,Past]	he directed	731
	הַנְּחָה	Noun [F,S,abs]	discount	517
הרים	הִרִים	Verb [M,S,3,Past]	he lifted	1,000
	הִרִים	Noun [M,P,abs]	mountains	1,000
ואת	וְאֶת	Conj + ACC	and + accusative	1,000
	וְאֶת	Conj + Pronoun [F,S,2]	and + you	1,000
זר	זֶר	Noun [M,S,abs/cons]	bouquet	1,000
	זֶר	Adj [M,S] / Noun [M,S,abs]	foreign / stranger	1,000

חברות	חִבְרוֹת	Noun [F,P,abs]	companies	1,000
	חִבְרוֹת	Noun [F,P,cons]	companies of	1,000
	חִבְרוֹת	Noun [F,P,abs/cons]	friends	501
	חִבְרוֹת	Noun [F,S,abs/cons]	friendship	398
	חִבְרוֹת	Noun [F,P,cons]	friends of	261
חדר	חֵדֶר	Noun [M,S,cons]	room of	1,000
	חֵדֶר	Noun [M,S,abs]	room	1,000
	חֵדֶר	Verb [M,S,3,Past]	penetrated	783
טוב	טוֹב	Adj [M,S]	good	1,000
	טוֹב	Noun [M,S,abs/cons]	goodness	357
יהודי	יְהוּדִי	Noun [M,S,abs] / Adj [M,S]	a Jew / Jewish	1,000
	יְהוּדִי	Noun [M,P,cons]	Jews	1,000
כיוון	כִּיּוֹן	Conj	because	487
	כִּיּוֹן	Noun [M,S,abs] / Noun [M,S,cons]	direction	468
	כִּיּוֹן	Verb [M,S,3,Past]	directed	455
לו	לוֹ	Prep [suf=M,S,3]	to him	1,000
	לוֹ	Conj	if only	1,000
לחם	לֶחֶם	Noun [M,S,abs]	bread	1,000
	לֶחֶם	Verb [M,S,3,Past]	he fought	1,000
לפנות	לְפָנוֹת	Prep / Verb [Infinitive]	facing / to turn	1,000
	לְפָנוֹת	Verb [Infinitive]	to clear out	570
מדי	מְדֵי	Det [cons]	every	1,000
	מְדֵי	Adv	too much	802
	מְדֵי	Noun [M,P,cons]	uniforms of	541
מהם	מֵהֶם	Preposition [suf=M,P,3]	from them	1,000
	מֵהֶם	Interrogative	what are	587
מי	מִי	Interrogative / Pronoun [S,3]	who	1,000
	מִי	Noun [M,P,cons]	waters of	1,000
מלך	מֶלֶךְ	Noun [M,S,abs/cons]	king	1,000
	מֶלֶךְ	Verb [M,S,3,Past]	he ruled	522
מעבר	מֵעַבֵּר	Prep	beyond	1,000
	מֵעַבֵּר	Noun [M,S,abs]	passage	1,000
	מֵעַבֵּר	Noun [M,S,cons]	passage of	883
מראה	מֵרְאֶה	Participle [M,S]	he shows	1,000
	מֵרְאֶה	Participle [F,S]	she shows	1,000
מרכז	מֵרְכֵז	Noun [M,S,cons]	center of	1,000
	מֵרְכֵז	Noun [M,S,abs]	center	1,000
	מֵרְכֵז	Participle [M,S,abs/cons]	organizes / organizer	393
משחק	מִשְׁחָק	Noun [M,S,abs]	game	1,000
	מִשְׁחָק	Participle [M,S,abs]	plays / player	479
נעשה	נַעֲשֶׂה	Verb [MF,P,1,Future]	we will do	1,000
	נַעֲשֶׂה	Verb [M,S,3,Past]	was done	1,000
נשים	נָשִׁים	Noun [F,P,abs]	women	1,000
	נָשִׁים	Verb [MF,P,1,Future]	we will put	613
נתן	נָתַן	Verb [M,S,3,Past]	gave	1,000
	נָתַן	Propn	Nathan	681

עבר	עָבַר	Verb [M,S,3,Past]	he passed	1,000
	עֶבֶר	Noun [M,S,abs]	past	1,000
	עֵבֶר	Noun [M,S,abs/cons]	side	456
עד	עַד	Prep	until	1,000
	עֵד	Noun [M,S,abs/cons]	witness	1,000
עובדות	עֹבְדוֹת	Noun [F,P,abs/cons]	facts	1,000
	עֹבְדוֹת	Participle [F,P]	they work / workers	1,000
עם	עִם	Prep	with	1,000
	עַם	Noun [M,S,abs/cons]	nation	1,000
פני	פָּנָיו	Noun [M,P,cons]	faces of	1,000
	פָּנִי	Noun [MF,P,abs,suf=MF,S,1]	my face	338
פרס	פְּרָס	Noun [M,S,abs]	award	1,000
	פֶּרֶס	Propn	Peres	956
	פָּרַס	Verb [M,S,3,Past]	he spread	630
	פְּרָס	Noun [M,S,cons]	award of	290
ציון	צִיּוֹן	Propn	Zion	1,000
	צִיּוֹן	Noun [M,S,abs/cons]	mark	1,000
קודם	קֹדֶם	Adv	before	1,000
	קֹדֶם	Adj [M,S]	previous	1,000
	קָדַם	Verb [M,S,3,Past]	was promoted	284
ראשי	רָאשֵׁי	Noun [M,P,cons]	heads	1,000
	רָאשֵׁי	Noun [M,S,abs,suf=MF,S,1]	my head	881
	רָאשֵׁי	Adj [M,S,abs]	head	399
שירת	שִׁירָתוֹ	Verb [M,S,3,Past]	he served	1,000
	שִׁירָתוֹ	Noun [F,S,cons]	poetry of	896
שכר	שָׂכָר	Noun [M,S,abs]	salary	1,000
	שָׂכָר	Noun [M,S,cons]	salary of	751
	שָׂכָר	Verb [M,S,3,Past]	rented	681
שם	שֵׁם	Noun [M,S,abs/cons]	name	1,000
	שָׁם	Adv	there	1,000
	שָׂם	Verb [M,S,Present/Past]	he placed	1,000
תנאי	תְּנָאִי	Noun [M,P,cons]	conditions of	1,000
	תְּנָאִי	Noun [M,S,abs/cons]	condition	834

### A.3 Homographs with Semantic Ambiguity

Form	Word	Morphology	Translation	Sentences
הזמר	הַ+זָּמֵר	Det + Noun [M,S,abs]	the + music	1,000
	הַ+זָּמֵר	Det + Noun [M,S,abs]	the + musician	1,000

הסופר	ה+סופר	Det + Noun [M,S,abs]	the + market	763
	ה+סופר	Det + Noun [M,S,abs]	the + author	570
זמר	זמר	Noun [M,S,abs]	musician	1,000
	זמר	Noun [M,S,abs]	song	602
חברה	חברה	Noun [F,S,abs]	friend	1,000
	חברה	Noun [F,S,abs]	company	1,000
רשות	רשות	Noun [F,S,abs/cons]	permission	1,000
	רשות	Noun [F,S,abs/cons]	authority	1,000

## B Appendix B: Creation of our Annotated Dataset

We present here further details regarding the creation of our annotated homograph dataset. As noted above, our dataset contains challenge sets for 75 Hebrew homographs. In almost all cases, we collect 1000 sentences for the primary analysis, at least 500 sentences for the secondary analysis, and at least 250 for each additional analysis. Every sentence was tagged by our team of human annotators.

The main challenge that we faced is that it is often extraordinarily difficult to identify a sufficient number of naturally-occurring Hebrew sentences attesting to the non-primary homograph analyses. In many cases, the naturally-occurring ratio of these analyses is 100:1 or worse, meaning that the non-primary analyses only occur once in 100 or more sentences. In such cases, *prima facie*, we would need to tag some 50,000 randomly-selected cases of the homograph in order to gather 500 cases of the alternate analysis. It would not be practical to tag this number of sentences for each of the 75 homographs. Therefore, we devised the following three-step process to allow us to efficiently gather a sufficient number of sentences attesting each possible analysis of each homograph:

1. For each homograph, we first gather 4000 randomly-selected sentences from an untagged Hebrew corpus. For each sentence, our annotation team determines the correct analysis of the homograph. This initial step virtually always provides us with a sufficient number of sentences attesting to the primary analysis, and sometimes for the secondary analysis as well. However, it often leaves us with an insufficient number of sentences regarding the non-primary analyses.

2. In order to find additional sentences for the non-primary homograph analyses, we train classifiers to identify a wide net of potential candidate sentences. For these classifiers, we use the known cases of the primary homograph analysis (identified in the previous step) as the first class, and we use various proxies to represent the opposing class. The proxies used for the opposing class include the following: (A) Hebrew words which unambiguously function morphosyntactically in a manner that is different from the morphosyntactic function of the primary class; (B) Hebrew words whose word2vec embedding is maximally distant from that of the homograph in question (because,

in practice, for homographs of skewed proportions, the word2vec embedding of the word tends to represent the primary usage); (C) Randomly-selected Hebrew words. These classifiers are trained with a BiLSTM encoding of the word2vec embeddings of the four neighbors on either side of the target word (not including the target word itself). Essentially, the point of these classifiers is to identify a generous selection of cases in which the homographs may possibly be functioning in a way that is different than their primary usage.

3. We run these classifiers on all instances of the target homograph in our untagged Hebrew corpus. Cases that are classified as the opposing class are sent to our annotators for human verification. In practice, these proxy-based classifiers reach a precision of 15-30% on the task of identifying sentences in which the target homographs function according to a non-primary analysis. Thus, on average, for any given homograph, tagging an additional 2222 sentences allows us to collect 500 relevant cases of the secondary analysis.

Our procedure ensures diversity while keeping the process cost-effective. All in all it took approximately 1500 hours of annotation time to amass the sentences in the present dataset.

## C Appendix C: Computational Equipment

We performed all computations on a desktop workstation with an i9-10980XE processor and 256GB of memory. This system enabled us to run 36 experiments in parallel (the processor contains 18 hyper-threaded cores), and thus we were able to complete all of the relevant experiments and computations over the course of several weeks of calendar time.

## D Appendix D: Neural Network Implementation Details

All BiLSTMs and MLPs were trained using dynet (<http://dynet.io/>). All MLPs are constructed with a hidden layer of size 100, and with two layers. We train with the Adam optimizer at a learning rate of .001, for 3 epochs.

Our word2vec embeddings for Hebrew are of size 100. We trained them on a 500M-word Hebrew corpus using word2vecf (<https://github.com/BIU-NLP/word2vecf>, adding position info to context words, and excluding words with a frequency of less than seven. All in all, our word2vec vocabulary covers 94.6% of the tokens in our dataset. For

the out-of-vocabulary words, we use a trainable UNK parameter in place of the word2vec embedding, which is trained from scratch for each “word expert” classifier.

As per the "word expert" paradigm, a completely separate MLP is trained for each homograph. In each case, the possible homograph analyses are each treated as a possible class for prediction, and the MLP is trained to choose from among those classes. Thus, for instance, if the homograph has two analyses, we train an MLP to predict one of the two classes; if the homograph has three analyses, then we train an MLP to predict one of the three classes; and so on.

For the Probing Scenarios based on centroid classification, we proceed as follows. For each of the homographs, given the training sample size (100, 50, 25, etc.), we randomly select that amount of training sentences for each of the possible analyses of the homograph. We calculate the centroid for each of the analyses by averaging the embeddings of the target homographs within the corresponding training sentences. The remainder of the available sentences for the homograph forms the test set. We classify them by calculating the dot product of the embedding of the target homograph in each given test sentence with the centroid of each of the homograph analyses. We run this process 200 times, each time selecting a different random set of training sentences. The values plotted in Figure 6 reflect the average of the F1 scores across these 200 rounds. For the corresponding MLP-based experiments presented for comparison in the aforementioned table, we follow an analogous procedure, across 10 rounds.