

# Word-Embeddings Distinguish Denominal and Root-Derived Verbs in Semitic

Ido Benbaji\*  
MIT  
Cambridge, MA, USA  
ibenbaji@mit.edu

Omri Doron\*  
MIT  
Cambridge, MA, USA  
omrid@mit.edu

Adèle Hénot-Mortier\*  
MIT  
Cambridge, MA, USA  
mortier@mit.edu

Proponents of the Distributed Morphology framework have posited the existence of two levels of morphological word formation: a lower one, leading to loose input-output semantic relationships; and an upper one, leading to tight input-output semantic relationships. In this work, we propose to test the validity of this assumption in the context of Hebrew word embeddings. If the two-level hypothesis is borne out, we expect state-of-the-art Hebrew word embeddings to encode (1) a noun, (2) a denominal derived from it (*via* an upper-level operation), and (3) a verb related to the noun (*via* a lower-level operation on the noun’s root), in such a way that the denominal (2) should be closer in the embedding space to the noun (1) than the related verb (3) is to the same noun (1). We report that this hypothesis is verified by four embedding models of Hebrew: fastText, GloVe, Word2Vec and AlephBERT. This suggests that word embedding models are able to capture complex and fine-grained semantic properties that are morphologically motivated.

## 1 Introduction: A few basic principles of word formation

### 1.1 Morphological processes sometimes appear “irregular”

A common assumption in generative morphology is that word formation differs essentially from the formation of larger phrasal structures [33, 9, 10]. Phrase formation on the one hand, is generally productive in the sense that it is seldom subject to arbitrary constraints; it also proves to be semantically compositional. Word formation on the other hand, is compositional at times and non-compositional at others, and is known to exhibit arbitrary paradigmatic gaps [3].

The lack of compositionality in word formation has been observed in certain English “berry-words” [23, 3]. The berry-words in (1a) seem semantically compositional as they involve concatenation of two independent morphemes, each of which contributes its meaning to the word as a whole. In contrast, the berry-words in (1b) cannot be said to be compositional, as one of their morphemes fails to convey any meaning when uttered independently from the other (i.e., the morphemes *cran*, *boysen*, and *huckle* are not meaningful units in English).

- (1) a. *blackberry*, *blueberry* “compositional” berries  
b. *cranberry*, *boysenberry*, *huckleberry* “non-compositional” berries

For an example of the arbitrariness in the application of morphological processes, consider the English nominals in (2). (2a) illustrates a regular morphological process that merges the morpheme *-ity* to an adjective ending with the phonemes *l-(i)ous/*, resulting in a noun. As is illustrated in (2b), some arbitrary constraint prevents this process from applying to the adjective *atrocious* [3]. Instead, a different process seems to apply in the derivation of *atrocidity* from *atrocious* – perhaps one which involves truncation of part of the adjective, followed by the merger of the nominalizing morpheme.

---

\*All authors contributed equally.

- (2) a. *curiosity, monstrosity, pomposity* regular *ity*-nominalization  
 b. *atrocitiy. \*atrociousity* irregular *-ity*-nominalization

## 1.2 The two-level model

To account for this ambivalent nature of morphological processes, linguists working in the tradition of the *Distributed Morphology* (DM) framework ([21, 13, 6, 22]) have posited the existence of two “levels” of morphological derivation; a “lower” level, where word formation may be irregular, arbitrary, and non-productive, and an “upper” level, where word formation is mostly regular and productive (cf. [26]).

The DM framework assumes that while morphological processes apply as part of the syntactic derivation, the “upper” and “lower” levels of morphological derivation are distinguished by the merger of a so-called “functional head” (*n, v, a, etc.*). A functional head sets the semantic, syntactic and phonological features of the word it creates. It may be merged directly with a *root*, which is an atomic element devoid of functional material [2]; or it may be merged with some other non-atomic constituent already dominating a functional head. More specifically, it has been assumed that any operation that applies directly to a root, (i.e., an operation that applies at the “lower”, non-word level), should remain impenetrable to upper-level operations. Fig. 2 below is an illustration of how this model accounts for the difference between the compositional berry-names (a.) and the non-compositional ones (b.).






Figure 1: Basic decomposition of the word formation process



(a) Compositional case: the 2 sub-words are merged with their respective heads (*a, n*) before being compounded.



(b) Non-compositional case: 1 of the sub-words is not merged with a functional head before compounding, and thus remains semantically opaque.

Figure 2: Compositional vs non-compositional berry-words

In the compositional case, the roots of the first morpheme combine with an adjectivizing head (*a*) to form an adjective, and the root of the second morpheme combines with a nominalizing head (*n*) to form a noun. The two are then conjoined to form a word whose meaning is compositional in that it simply involves intersecting the meaning of the adjective with that of the noun. In the non-compositional case, the root that does not convey any meaning in the context of the whole compound ( $\sqrt{\text{cran}}$ ,  $\sqrt{\text{boysen}}$ ,  $\sqrt{\text{huckle}}$ ), is not merged with any functional head before it combines with the noun (*berry*).

The derivation is in that case a “low” level process. The divide between pre-word “lower” level and post-word “upper” level processes leads to two related predictions:

**Prediction (a): Elements derived from the same *root* via a lower-level morphological operation may arbitrarily differ semantically;**

**Prediction (b): Elements that derive from the same *word* via upper-level morphological operations should be closely related semantically.**

### 1.3 The relevance of Semitic morphology

Semitic languages provide a useful testing ground for the predictions of the two-level model. This is because in many of these languages, words can be decomposed into consonantal roots and fixed morphophonemic patterns that introduce functional information, like part of speech (*n, a, v* etc.) and valence in the case of verbs ([36]). As is illustrated in Tab. 1, patterns can be seen as recipes to “fill the gaps” between the consonants of the root. These roots and patterns, which only form independent words when merged together, have nevertheless been argued to be mentally represented as accessible independent morphological units (cf. [34]).

Verbal pattern \ Root		$\sqrt{\text{x}f\text{v}}$		$\sqrt{\text{ktv}}$	
(i)	CaCaC	xafav	‘thought’	katav	‘wrote’
(ii)	niCCaC	nexfav	‘was well considered’	nixtav	‘was written’
(iii)	CiCCeC	xifev	‘calculated’	kitev	‘CC-ed’
(iv)	CuCCaC	xufav	‘was calculated’	kutav	‘was CC-ed’
(v)	hiCCiC	hixfiv	‘considered’	hixtiv	‘dictated’
(vi)	huCCaC	huxfav	‘was considered’	huxtav	‘was dictated’
(vii)	hitCaCCeC	hitxafev	‘was considerate’	hitkatev	‘corresponded’

Table 1: Varieties of meaning for the same root

While Semitic roots do seem to signal some broad semantic field, the meaning that results from combining a certain root with a certain pattern is highly unpredictable and non-systematic. This is also illustrated in Tab. 1, where two tri-consonantal roots ( $\sqrt{\text{x}f\text{v}}$  and  $\sqrt{\text{ktv}}$ ), are combined with each of the language’s seven verbal patterns. Crucially however, there is no transparent semantic operation that the patterns seem to denote. While niCCaC (ii) seems to constitute the passive version of CaCaC (i) when the two patterns are combined with the root  $\sqrt{\text{ktv}}$ , this is not the case when they are combined with the root  $\sqrt{\text{x}f\text{v}}$ . Some patterns seem to behave more systematically than others (for instance, huCCaC (vi) is usually just the passive of hiCCiC (v)), but this is not generally the case.

It is also worth noting that Hebrew has three verbal patterns containing four consonant slots – CiCCeC (iii), CuCCaC (iv), and hitCaCCeC (vii). As pointed out in [2], no gemination exists in these patterns in Modern Hebrew. However, we nevertheless have evidence that these patterns contain an extra consonant slot, as they *productively* combine with roots of four consonant. For instance, the root  $\sqrt{\text{h}v\text{b}}$  can combine with these patterns to form  $\text{h}v\text{b}\text{v}$  (‘liberated’),  $\text{h}v\text{b}\text{av}$  (‘was liberated’), and  $\text{h}v\text{b}\text{iv}$  (‘liberated oneself’). Other patterns (such as hiCCiC (v)) can combine with four non-templatic consonants, but this process is less productive and is generally restricted to loan-words (for instance, the Hebrew verb  $\text{h}v\text{p}\text{r}\text{v}$  (‘splashed’) is derived from the German verb with the same meaning *spritzen*). This will be relevant to our discussion of the behavior of denominal verbs in Hebrew in the next section.

## 2 The case of Hebrew denominal verbs

Given our two-level model of morphological processes, root-pattern combination is a lower-level operation by definition. Therefore, it is not surprising that the meanings obtained from combining the same root with different patterns differ in arbitrary ways. Furthermore, a given root can yield forms with potentially divergent syntactic categories, as the root can combine with functional heads of various labels (*n*, *v*, *a*, etc.). This is illustrated in Tab. 2, which combines the same roots used in Tab. 1 with nominal and adjectival patterns, rather than verbal ones.

	Root	$\sqrt{xjv}$	$\sqrt{ktv}$
	Pattern		
(viii)	<b>miCCaCa</b> (n)	<b>maxj</b> ava ‘thought’	<b>mixt</b> ava ‘desk’
(ix)	<b>maCCeC</b> (n)	<b>maxj</b> ev ‘computer’	NA
(x)	<b>miCCuC</b> (n)	<b>mixj</b> uv ‘computing’	NA
(xi)	<b>CCiCut</b> (n)	<b>xaj</b> ivut ‘importance’	NA
(xii)	<b>taCCiC</b> (n)	<b>taxj</b> iv ‘calculation’	<b>taxt</b> iv ‘decree’
(xiii)	<b>CeCCon</b> (n)	<b>xej</b> bon ‘bill’	NA
(xiv)	<b>CCiCa</b> (n)	<b>xaj</b> iva ‘thinking’	<b>kt</b> iva ‘writing’
(xv)	<b>CCaC</b> (n)	NA	<b>kt</b> av ‘hand-writing’
(xvi)	<b>CaCaC</b> (n)	<b>xaj</b> av ‘accountant’	<b>kat</b> av ‘correspondent’
(xvii)	<b>miCCaC</b> (n)	NA	<b>mixt</b> av ‘letter’
(xviii)	<b>CaCuC</b> (a)	<b>xaj</b> uv ‘important’	<b>kat</b> uv ‘written’
(xix)	<b>CaCCan</b> (a)	NA	<b>kat</b> van ‘pulp writer’

Table 2: Nominal and adjectival patterns

Recall that the two-level model made two predictions: (a) Elements derived from the same *root* via a lower-level morphological operation may arbitrarily differ semantically; and (b) elements that derive from the same *word* via upper-level morphological operations should be closely related semantically. The data in Tab. 1 and 2 seem to confirm the former prediction, as the meanings that result from combining the different roots with nominal and adjectival patterns seem quite unpredictable. Arad [2] claims that Hebrew provides us with the possibility to test the second prediction. More specifically, Hebrew comes with a morphophonemic diagnostic that can help identify elements derived *via* upper-level operations, such as denominal and de-adjectival verbs.<sup>1</sup> Denominal verbs are derived by merging a verbal pattern with a *noun*, rather than with a *root*. The base noun is itself derived from a root:

$$\sqrt{\xrightarrow{\text{lower-level}} N \xrightarrow{\text{upper-level}} V_{\text{denom}}}$$

In Hebrew, certain nominal patterns include templatic consonants (marked in blue in the Figures), in addition to the templatic vowels. For instance, templates (viii-x) and (xvii) in Tab. 2 have a templatic /m/, templates (xi-xii) have a templatic /t/, and templates (xiii) and (xix) have a templatic /n/. Arad notes that certain *verbs* (such as *mixjev* and *hitxafben* in Tab. 3 below) seem to behave as if a consonant has been adjoined to their root. She argues that this consonant is a templatic consonant originating from the root-derived noun that served as a base for the derivation of the corresponding verb (*denominal*). In other words, it is argued that *mixjev* and *hitxafben* do not derive directly from the root  $\sqrt{xjv}$ , but rather, from the root-derived nouns *maxjev* and *xejbon*, respectively. This is illustrated in Fig. 3a. Other verbs, such as *xisev* and *hitxafev*, which are derived using the same pattern as *mixjev* and *hitxafben* respectively

<sup>1</sup>In the prose, we will refer going forward only to denominal verbs, but the observations apply to de-adjectival ones as well.

(cf. Tab. 3), lack a templatic consonant and are therefore argued to derive *directly* from their respective roots. This is illustrated in Fig. 3b. Also note in Tab. 3 that verbal forms involving a templatic consonant (mixfev, hitxafben) have a meaning which seems closely related to the meaning of the nouns derived from the same root (maxfev, xefbon). If the presence of a templatic consonant is indeed the marker of an upper-level operation mapping a noun to a verb, as Arad argues, this semantic property is in line with **Prediction (b)** of two-level model (cf. p. 3). In contrast, verbal forms derived using the same patterns, but devoid of any templatic consonant (xifev, hitxafev) seem not as close to the corresponding root-derived nouns in terms of meaning, this time in line with **Prediction (a)** of the two-level model.

N pattern	$\sqrt{\text{ }}$ -derived Noun	V pattern	Possible Verbs
maCCeC (ix)	maxfev ‘computer’	CiCCeC	mixfev ‘computerized’ xifev ‘calculated’
CeCCon (xiii)	xefbon ‘bill’	hitCaCCeC	hitxafben ‘settled up financially’ hitxafev ‘was considerate’

Table 3: The root  $\sqrt{\text{x}f\text{v}}$  and the templatic consonants

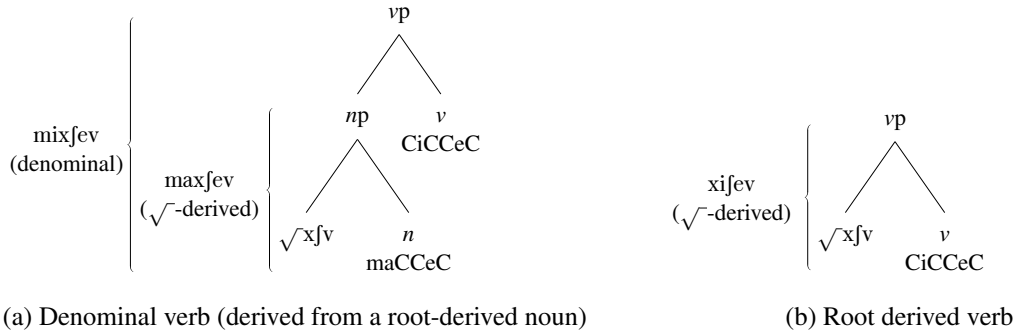


Figure 3: The structure of denominal vs root-derived verbs

In brief, the two-level model, together with Arad’s claim that verbs with an extra nominal-templatic consonant are denominal rather than root-derived, predicts that these denominal verbs exhibit a closer semantic similarity to the noun from which they were derived, than root-derived verbs do.

### 2.1 Previous empirical investigations into denominals in the two-level model

As far as we know, our paper is the first to investigate the semantic encoding of Hebrew denominal verbs within word embedding models. However, an experimental study on the processing of Hebrew denominal verbs has already been conducted by Brice on human participants [8]. A priming experiment was used to provide further evidence for the two-level model, based on the background assumption that morphological inclusion corresponds to a strong priming connection. This claim is supported by a series of papers ([16], [15], [17], [39], [40]), showing that roots have a strong priming effect for verbs that are derived from them. More generally, it can be assumed that if some language component *A* is contained in the morphological structure of a component *B*, then the string that orthographically represents *A* should prime the string that represents *B*. In [8], Brice tests this property on Hebrew nouns and denominals, by comparing the priming effect of a noun on the denominal verb derived from it, to that of a verb derived from the same root. An example for such triplets is given in Tab. 4. Crucially, the two stimuli were orthographically equally similar to the target, so any difference in priming could only be attributed to deeper connections between the words.

Given these assumptions, the two-level model’s prediction is that the noun stimulus should have a stronger priming effect on the denominal than the root-derived verb stimulus. That is because the noun is contained in the denominal verb’s morphological structure, while the root-derived verb is not. Brice indeed found that the noun stimuli yield significantly shorter reaction times among the participants than the verb stimuli, i.e., demonstrate a stronger priming effect. We view our study as complementing this result. However, Brice argued that a given noun’s priming its corresponding denominal verb could not be explained by a semantic connection between them, since priming has been shown to be unaffected by the meaning of the stimulus (see [32], [29], a.o.). Therefore, Brice took his results as evidence for the morphosyntactic aspect of the two-level model. Our study, on the other hand, aims to provide evidence in favor of the existence of semantic consequences of the the two-level model; i.e., the idea that upper-level derivational operations entail a high degree of semantic similarity between their input and output.

Target	letaxkeɁ	‘to debrief’
Noun stimulus	taxkiɁ	‘debriefing’
Verb stimulus	laxkoɁ	‘to inquire’

Table 4: Example stimulus from [8]

### 3 Testing the predictions using Hebrew word-embedding models

#### 3.1 Words embeddings capture meaningful semantic generalizations

Static word embedding models such as Word2Vec [27], GloVe [31] and fastText [7] map each word of a given lexicon to a dense, high-dimensional vector. This mapping is obtained by training a neural network on specific language-related tasks, such as predicting the context of any given word (Skip-gram), or predicting a word given a context (CBOW) [27]. By contrast, contextualized word embeddings like BERT [12], and its variant AlephBERT [35], pretrained on Hebrew data, can take whole sentences as input, and may map the same word to different representations, depending on the context. BERT in particular, adopts a deep encoder-decoder architecture (“Transformer”, [38]). A stack of encoders (12 for the basic model) use attention mechanisms to forward a more complete picture of the whole sequence to the decoder.

Word embeddings in general have been argued to encode lexical meaning, in the sense that word-vectors that are close to each other in the embedding space are expected to be close in meaning [24]. The relevant metric is usually taken to be the cosine similarity (measure of the angle) between two vectors. In particular, word embeddings have been shown to encode specific morphological/semantic relationships (such as comparative/superlative formation, masculine/feminine nominalizations, some part-whole relationships) as stable linear transformations, a somewhat surprising result that has led to various explanations in the recent years [4, 18, 14, 1]. Contextualized embeddings like BERT were shown to be especially good in capturing polysemy and homonymy in natural language [28]. Given that word embeddings provide a quantitative measure of semantic similarity, and that our morphological model makes specific semantic predictions, those language models appear as an interesting testing ground for the predictions of the two-level model.

#### 3.2 Word embeddings and the two-level model

We propose here that the abstract linguistic notion of *root* be modeled as a subspace of the word embedding. This region should contain (at least) all the vectors corresponding to words derived from the

root through a merger with a functional head (more specifically in our case, a pattern).<sup>2</sup> For instance, the root  $\sqrt{x}f\bar{v}$  designates a region that contains, among others, the vectors for *xafuv* (‘important’), *maxjava* (‘thought’), and *hitxafev* (‘was considerate’).

In the two-level framework, the generation of root-derived elements is semantically opaque (cf. **Prediction (a)**). In other words, elements derived from the same root *via* a lower-level process (i.e. the merger of a head directly with the root), are expected to exhibit arbitrary semantic differences. Assuming that roots denote regions in the embedding space, semantic opacity corresponds to an expectation that for any root  $\sqrt{x}$ , and any set of templates  $\{t_1, t_2, \dots, t_i\}$ , the vectors corresponding to the words derived from applying the templates to the root (i.e.,  $\{\overrightarrow{t_1(\sqrt{x})}, \overrightarrow{t_2(\sqrt{x})}, \dots, \overrightarrow{t_i(\sqrt{x})}\}$ ) will be arbitrarily distributed over the region designated by the root.

The generation of word-derived elements, on the other hand, is argued to be semantically restricted and transparent (cf. **Prediction (b)**). More specifically, if  $Y$  is an element derived by merging a functional head with an element that already contains a functional head  $X$ , then the meaning of  $Y$  is expected to be close to the meaning of  $X$  in a systematic way [2]. In other words, if the merger of the first functional head can lead to a vector  $\vec{X}$  located anywhere within the region denoted by the root (Prediction (a)), the merger of a second head on  $X$  (yielding  $Y$ ) should lead to a representation  $\vec{Y}$  that is in the close vicinity of  $\vec{X}$  (Prediction (b)). An illustration of this interpretation of the two-level model predictions is provided in Fig. 4.




Figure 4: Expected distribution of root-derived *vs* word-derived elements within a simplified embedding space.

### 3.3 Testing the two-level model in the context of Hebrew word embeddings

When it comes to Hebrew denominal verbs, the prediction of our interpretation of the two-level model in the context of word embeddings is the following. Given a root ( $\sqrt{r}$ ); a noun derived from it *via* a lower-level operation ( $N_{\sqrt{r}}$ ); a denominal verb derived from that noun *via* an upper-level operation ( $V_{N_{\sqrt{r}}}$ ); and finally, a verb derived directly from the root *via* a lower-level operation ( $V_{\sqrt{r}}$ ), we expect  $\overrightarrow{V_{N_{\sqrt{r}}}}$  to be generally<sup>3</sup> closer to  $\overrightarrow{N_{\sqrt{r}}}$  within the embedding space, than  $\overrightarrow{V_{\sqrt{r}}}$  is.

The exact nature of the root-derived verb  $V_{\sqrt{r}}$  remains to be fleshed out however. Indeed, if a given root  $\sqrt{r}$  normally yields a single root-derived noun ( $N_{\sqrt{r}}$ ), it can in principle give rise to many different root-derived verbs  $\{V_{\sqrt{r}}^{(1)}, \dots, V_{\sqrt{r}}^{(k)}\}$ , some of them being closer to  $N_{\sqrt{r}}$  than others. Assuming that  $\{\overrightarrow{V_{\sqrt{r}}^{(1)}}, \dots, \overrightarrow{V_{\sqrt{r}}^{(k)}}\}$  are somewhat uniformly distributed across the region defined by  $\sqrt{r}$ , we expect the mean similarity between  $\overrightarrow{N_{\sqrt{r}}}$  and each of the  $\overrightarrow{V_{\sqrt{r}}^{(i)}}$  to be lower than the similarity between  $\overrightarrow{N_{\sqrt{r}}}$  and the denominal derived from it,  $\overrightarrow{V_{N_{\sqrt{r}}}}$ . This is formalized in Eq. 1 below (where  $\mathcal{S}$  stands for cosine similarity):

<sup>2</sup>A possible way to define this region in mathematical terms would be to use the notion of convex hull, although our analysis does not depend on any specific implementation thereof.

<sup>3</sup>We think that it is too restrictive to state that this property should hold *for any given root*, due to the arbitrary character of lower-level morphological operations. Since lower-level operations are assumed to map root-derived elements somewhat randomly within the subspace defined by the root, a root-derived verb might accidentally end up very close to a root-derived noun (meaning, closer to this noun than the noun-derived denominal is). Our modeling however, predicts that this configuration should be rare enough for our inequalities to have some statistical significance.

$$\forall \sqrt{\phantom{x}} : \frac{1}{k} \sum_{i=1}^k \mathcal{S} \left( \overrightarrow{N_{\sqrt{\phantom{x}}}}, \overrightarrow{V_{\sqrt{\phantom{x}}}^{(k)}}} \right) < \mathcal{S} \left( \overrightarrow{N_{\sqrt{\phantom{x}}}}, \overrightarrow{V_{N_{\sqrt{\phantom{x}}}}} \right) \quad (1)$$

**Hypothesis 1**

Assuming that vectors of root-derived verbs are arbitrarily distributed across the region defined by the root, some root-derived verbs might be accidentally closer to the root-derived noun than the denominal verb derived from that noun. By using the mean similarity between the root-derived noun and the various root-derived verbs, the prediction is rendered compatible with this possibility. However, it might also be worthwhile to test a stronger prediction, according to which the similarity between a noun and a denominal verb derived from it should be greater than the similarity between the noun and the root-derived verb that is *maximally similar* to the noun. This is formalized in Eq. 2 below.

$$\forall \sqrt{\phantom{x}} : \max_{i \in [1, k]} \mathcal{S} \left( \overrightarrow{N_{\sqrt{\phantom{x}}}}, \overrightarrow{V_{\sqrt{\phantom{x}}}^{(k)}}} \right) < \mathcal{S} \left( \overrightarrow{N_{\sqrt{\phantom{x}}}}, \overrightarrow{V_{N_{\sqrt{\phantom{x}}}}} \right) \quad (2)$$

**Hypothesis 2**

## 4 Implementation and results

### 4.1 Dataset creation

We generated and tested Hebrew data in order to validate our two hypotheses. Each data point in our dataset contained (1) a noun with a templatic consonant, (2) a denominal verb containing the templatic consonant from the noun, in addition to the three root consonants (cf. Tab. 3), and (3) a list of verbs derived directly from the same root as the noun (and thus devoid of templatic consonants).<sup>4</sup> Each data point contained between one and five root-derived verbs, depending on how productively the given root combined with the verbal templates.

A list of nominal patterns containing templatic consonants was constructed using introspection and previous linguistic papers on the subject ([5, 2, 8]). Each nominal template was mapped to the verbal template which incorporates the nominal consonant into the verbal form. This is illustrated in Tab. 5 (partial list). Given that Modern Hebrew lacks vowel marking in the orthography, and therefore involves a high rate of ambiguity (cf. [37]), we used the infinitival form of the verbal templates, which are not ambiguous with nominal elements in the language even in the absence of vowel markings.

Nominal pattern	Denominal pattern
<b>tiCCoCet</b> <b>tiCCoCa</b> <b>taCCiC</b>	<b>letaCCeC</b>
<b>CeCCon</b>	<b>leCaCCen, lehitCaCCen</b>
<b>maCCeC</b> <b>miCCeCet</b> <b>miCCaC</b>	<b>lemaCCeC, lehitmaCCeC</b>
<b>šaCCeCet</b>	<b>lešaCCeC, lehištaCCeC</b>
<b>CaCaCat</b>	<b>leCaCCet, lehitCaCCet</b>

Table 5: Templates used for data generation

A list of nouns instantiating the relevant nominal templates was generated by matching nouns from a PoS-tagged Hebrew corpus, *The Knesset Meetings Corpus*<sup>5</sup>, against the various nominal templates.

<sup>4</sup>It is worth mentioning that a single noun could in certain cases give rise to two data points, when the noun’s root happened to be compatible with two denominal templates.

<sup>5</sup>This corpus gathers protocols of sessions in the Israeli parliament between January 2004 and November 2005. The particular archive we used was `kneset16`.



Candidate denominal verbs corresponding to each noun could be subsequently created, using the template mapping established in Tab. 5. Root-derived verbs were generated using the five infinitival verbal templates of Hebrew, corresponding to the seven inflected templates in Tab. 1. This process generated a list of 1435 potential data points. Given that not all nouns can productively give rise to denominal verbs, most of these were not actual verbs of Hebrew, or did not have a root that could productively combine with other verbal templates. We first eliminated the data points that were obviously not part of the grammar, by checking if the denominal (or any inflected form thereof) could be found in the list of verbs extracted from the Knesset dataset. This first filtering step ruled out 1322 data points, and left us with 113 data points to inspect further. We manually discarded the remaining defective items, ending up with a list of 66 denominal verbs.

## 4.2 Word embeddings

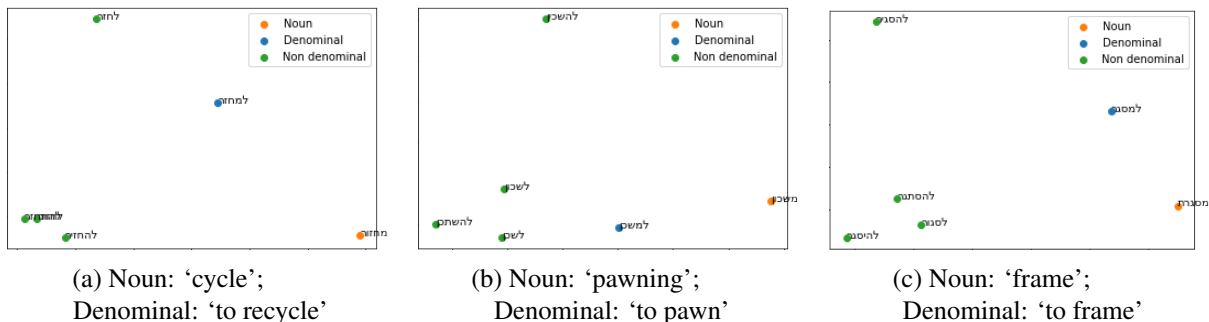
To convert the words in our dataset into vectors, we used pretrained models from fastText [19], and BERT (AlephBERT) [35], and trained GloVe and Word2Vec<sup>8</sup> models on Hebrew Wikipedia dumps. The characteristics of those various embeddings are given in Tab. 6.

Model	Word2Vec	GloVe	fastText	BERT
# vectors	584 160	584 162	2 billion	NA <sup>6</sup>
Initial dimension	100	50/100 <sup>7</sup>	300	768
PCA-reduced dimension	27	28/46	50	107

Table 6: Characteristics of the models

The AlephBERT embedding was obtained by summing the last 4 layers obtained after feeding the model with individual tokenized words, one at a time. Tokenization was performed using a dedicated function from the BERT library. If a given word was represented using several tokens, then, the representation of the word was obtained by averaging the representations of the individual tokens.

Since the dimensions of the embeddings were close to our dataset’s total size, we reduced the space using Principal Component Analysis (PCA, [30]) prior to computing the similarities and testing the hypotheses. We relied on the Guttman-Kaiser Criterion [20] to determine the target dimension for each space. Before computing the similarities, we chose to plot a few data points in a 2D space. For the visualizations to be as meaningful and readable as possible, we used PCA with a cosine kernel on each separate data point. A few such plots are represented in Fig 5.



<sup>7</sup>We trained GloVe using the default dimension of 50, then switched to 100 to allow for a better comparison with the other models, which have similar dimensions.

<sup>7</sup>As BERT models do not assign words to a fixed vector.

<sup>8</sup>For training the Word2Vec model, we used the Skip-Gram architecture.

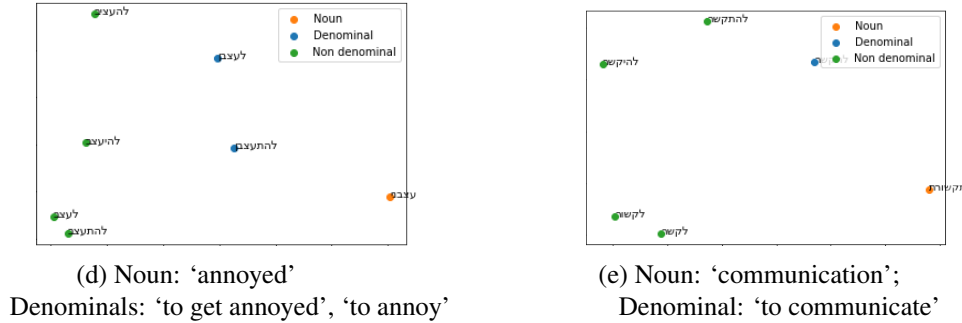


Figure 5: 2D projection of a few data points (PCA, cosine kernel, fastText model)

Those 2D plots bring preliminary evidence in support of the denominal verbs (blue dots) being closer to their respective nouns (orange dots) in a cosine space, than the other root-derived verbs (green dots).

### 4.3 Computation of the similarities

For each data point, our main hypothesis predicts that the cosine similarity between the noun and the denominal verb should be higher than the mean cosine similarity between the noun and all the other verbs sharing the same root. The stronger hypothesis predicts that the same kind of inequality holds when the “mean” operator is replaced by a “max” operator. In both cases, for each data point, the difference between two measures of similarities (noun/denominal; noun/other root derived element) is expected to be positive. The distributions of those two measures of similarity are represented in Fig. 6, for the main hypothesis, and the various models we tested. Those plots suggest that the main hypothesis is verified, since all distributions seem to have a mean and median above 0. In the next section, we test the significance of those empirical observations.

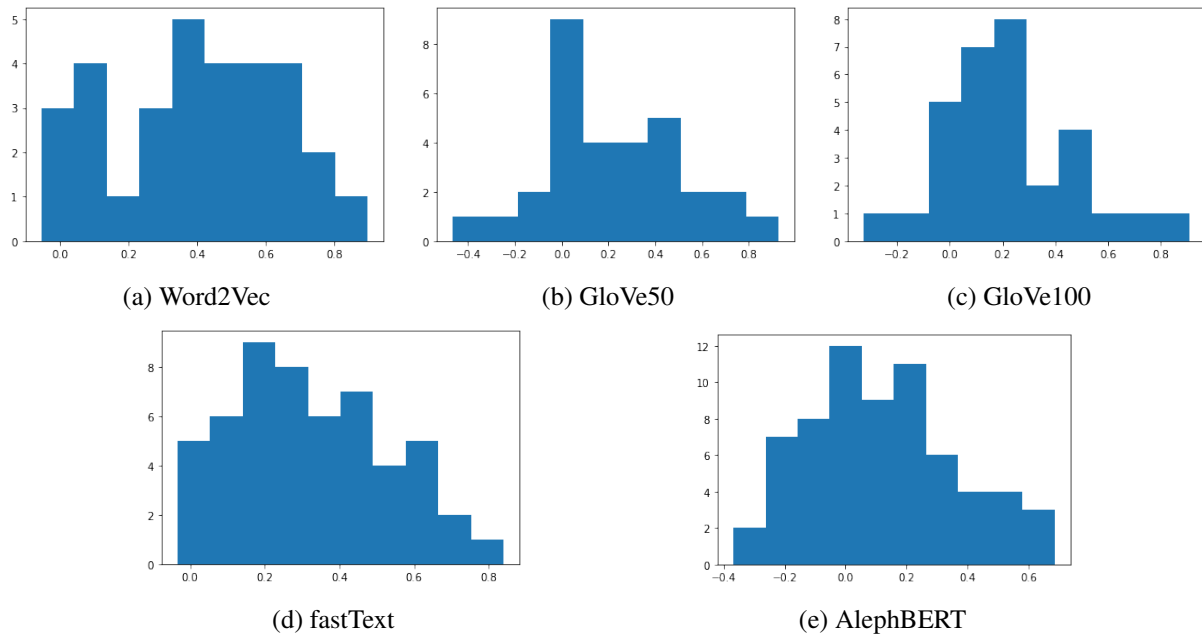


Figure 6: Paired similarity differences for various models under H1 (main hypothesis)

#### 4.4 Tests and results

To test whether denominals were significantly more similar to their corresponding nouns than other root-derived verbs, we performed (non-parametric) one-tailed Wilcoxon tests for matched-pairs<sup>9</sup> on the data described and plotted in the previous section. We did not feel the need perform any correction on the  $p$ -values, even though 2 hypotheses were tested for each model, because H1 was entailed by H2 by construction (so, a significant  $p$ -value for H2 implies a significant  $p$ -value for H1 as well). The  $p$ -values and effect sizes are compiled in Tab. 7. The effect sizes correspond to Cliff’s  $\Delta$  [11], which is a robust measure for non-parametric samples.

	Word2Vec	GloVe50	GloVe100	fastText	AlephBERT
# data points	31	31	31	53	66
H1	$1.06 \times 10^{-6}$ 0.86 (Large)	$2.43 \times 10^{-4}$ 0.52 (Large)	$6.64 \times 10^{-5}$ 0.66 (Large)	$1.42 \times 10^{-10}$ 0.79 (Large)	$4.84 \times 10^{-4}$ 0.30 (Small)
H2	$3.77 \times 10^{-5}$ 0.66 (Large)	$1.68 \times 10^{-1}$ 0.06 (Negligible)	$2.87 \times 10^{-2}$ 0.20 (Small)	$1.39 \times 10^{-8}$ 0.62 (Large)	$3.59 \times 10^{-1}$ 0.02 (Negligible)

Table 7:  $p$ -values and effect sizes (Cliff’s  $\Delta$ ) for H1 and H2 and 4 embedding models

As shown above, our main hypothesis (H1) was verified in all the models tested, and the stronger hypothesis (H2) appeared to be verified in all but two models (GloVe50 and AlephBERT). This indicates that our main prediction is quite robust across various language models. Overall large effect sizes also suggest that the effect, whenever present, is quite strong.

The fact that GloVe100, but not GloVe50, verified both hypotheses is somewhat interesting, as this might mean that a certain richness is needed in the original embedding space (despite the fact that this space is subsequently reduced *via* PCA), in order to capture the relevant generalization. This claim however, seems to be disproved by the results of AlephBERT, since this model failed to verify H2 just like GloVe50, while having the highest original dimension (768). Moreover, the failure of AlephBERT on H2 does not seem to be caused by the dimension reduction process (PCA) being too permissive in keeping too many irrelevant dimensions. Indeed, reducing the space further to the arbitrary dimension of 50 (which retained 71% of the explained variance) did not change the overall outcome – H2 was still rejected. Rather, the inefficiency of AlephBERT may be caused by the relative misuse of that model in the context of our experiment; sequential models such as BERT perform well on contextualized data, but here, the model was fed with isolated words, and could not benefit from the presence of surrounding words to enrich its representations.

## 5 Discussion, conclusion, and further questions

Our results seem to corroborate the claim of the two-level model that word-derived elements are systematically closer semantically to the word from which they derive, than elements derived from the same root are similar to that word. This prediction was verified in all the word embeddings we tested, provided that the original dimension was high enough. This result complements Brice’s contribution (cf. [8]), by establishing that morphological inclusion has some real semantic import, in addition to having consequences in terms of priming. It also gives support to the claim that word embedding models might

<sup>9</sup>We preferred a non-parametric test as opposed to a standard t-test because the similarity plots suggested that the distributions did not satisfy the t-test assumptions. This was confirmed by Levene tests conducted prior to performing the main tests, for all models but AlephBERT.

not only capture superficial linguistic regularities; and that, at least in some very specific corners of the language, those models show behaviors similar to humans’.

However, there is one potential confound in the way we tested our prediction. Static word embedding models obtain a single global representation for each word [25]. Therefore, those models fail to capture the different meanings of ambiguous words. As mentioned earlier, the absence of vowel marking in Modern Hebrew orthography (coupled with the “double life” of certain letters that have multiple pronunciations) renders many written words in Hebrew highly ambiguous. When it comes to denominal verbs matters get even worse, as certain nominal templates that give rise to such verbs are systematically ambiguous with certain inflections of a corresponding root derived verb or with an inflection of the denominal verb itself. The systematically ambiguous templates are listed in Tab. 8.

Templates <sup>10</sup>	Possible PoS & Morphological features	Example
<b>t(a)CCiC</b>	Noun with templatic consonant 2.M.SG IMPERATIVE <b>root-derived</b> verb in hiCCiC	taklit ‘record’ taklit ‘record!’
<b>m(a)CC(e)C</b> , <b>m(i)CC(a)C</b>	Noun with templatic consonant 3.M.SG PRESENT <b>root-derived</b> verb in CiCCeC 3.M.SG PAST <b>denominal</b> verb in CiCCeC	maxšev ‘computer’ mexašev ‘calculates’ mixšev ‘computerized’
<b>C(a)CC(a)n</b> , <b>C(a)C(a)C(a)t</b>	Noun with templatic consonant 3.M.SG PAST <b>denominal</b> verb in CiCCeC	tsalaxat ‘plate’ tsilxet ‘plated’

Table 8: Systematically ambiguous patterns

This makes testing predictions regarding Hebrew in a static word embedding model problematic to begin with, and even more so when it comes to predictions regarding Hebrew denominal verbs. One way to minimize ambiguity in our data was to use forms of each word that are least ambiguous. For instance, to resolve the ambiguity in (i), we used the plural inflections of nouns in the taCCiC template, which are no longer ambiguous with any root-derived verbs. This fix, however, does not resolve the ambiguity in (ii), as the plural form of the noun in that case is still ambiguous – this time with a plural root-derived verb. We also tried to minimize ambiguity by using verbs in their infinitival form, as infinitival verbs are not ambiguous with any nouns. However, this does not fully resolve the confound, as the infinitival form of verbs in Hebrew involves concatenation of the prefix /le-/ to some combination of the root consonants with other templatic information, and the prefix /le-/ itself is ambiguous between an infinitival marker and the preposition ‘to’. Therefore, the orthographic representation of some of our denominal verbs is ambiguous between a verbal interpretation and a the prepositional interpretation ‘to N’, where N is the noun from which the denominal is derived. It is unclear to us how the possible ambiguities discussed here influenced our results.

Another way to avoid the problems imposed on us by the high ambiguity rates in Modern Hebrew could be to test our hypotheses in a proper contextual word embedding, obtained by feeding AlephBERT with words put in disambiguating contexts. However, choosing this solution raises the issue of the design of a relevant context for each target word. First, choosing the right context for a given word is a subjective task that might make our experimental set up biased, or, at least, less controlled. Second, verbs derived from the same root but associated to different templates have different valences and senses, and require complements of different kinds. Those fundamental structural differences would add extra noise that would influence the similarity comparisons, and could not be reasonably counterbalanced. We have yet to investigate how to use a proper contextual embedding model while overcoming this concern.

<sup>10</sup>Vowels that are not orthographically represented are put in parentheses.

## References

- [1] Carl Allen & Timothy M. Hospedales (2019): *Analogies Explained: Towards Understanding Word Embeddings*. CoRR abs/1901.09813, doi:10.48550/arXiv.1901.09813. arXiv:1901.09813.
- [2] Maya Arad (2003): *Locality Constraints on the Interpretation of Roots: The Case of Hebrew Denominal Verbs*. *Natural Language and Linguistic Theory* 21(4), pp. 737–778, doi:10.1023/a:1025533719905.
- [3] Mark Aronoff (1976): *Word Formation in Generative Grammar*. Linguistic Inquiry monographs, MIT press.
- [4] Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma & Andrej Risteski (2016): *A Latent Variable Model Approach to PMI-based Word Embeddings*. *Transactions of the Association for Computational Linguistics* 4, pp. 385–399, doi:10.1162/tacl\_a\_00106. Available at <https://aclanthology.org/Q16-1028>.
- [5] Outi Bat-El (1994): *Stem Modification and Cluster Transfer in Modern Hebrew*. *Natural Language and Linguistic Theory* 12(4), pp. 571–596, doi:10.1007/BF00992928. Available at <http://www.jstor.org/stable/4047868>.
- [6] J.D. Bobaljik (2012): *Universals in Comparative Morphology: Suppletion, Superlatives, and the Structure of Words*. Current Studies in Linguistics, MIT Press, doi:10.7551/mitpress/9069.001.0001.
- [7] Piotr Bojanowski, Edouard Grave, Armand Joulin & Tomas Mikolov (2016): *Enriching Word Vectors with Subword Information*. arXiv preprint arXiv:1607.04606, doi:10.48550/arXiv.1607.04606. Available at <https://arxiv.org/abs/1607.04606>.
- [8] Henry Brice (2016): *The root and word distinction: an experimental study of Hebrew denominal verbs*. *Morphology* 27(2), pp. 159–177, doi:10.1007/s11525-016-9297-0.
- [9] Noam Chomsky (1970): *Remarks on Nominalization*. In R. Jacobs & P. S. Rosenbaum, editors: *Reading in English Transformational Grammar*, Ginn, Waltham, pp. 184–221.
- [10] Noam Chomsky (1973): *Conditions on Transformations*. In S. R. Anderson & R. P. V. Kiparsky, editors: *A Festschrift for Morris Halle*, Holt, Rinehart & Winston, New York, pp. 232–286.
- [11] Norman Cliff (1993): *Dominance statistics: Ordinal analyses to answer ordinal questions*. *Psychological Bulletin* 114(3), pp. 494–509, doi:10.1037/0033-2909.114.3.494.
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee & Kristina Toutanova (2018): *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*. CoRR abs/1810.04805, doi:10.48550/arXiv.1810.04805. arXiv:1810.04805.
- [13] David Embick (2010): *Localism versus globalism in morphology and phonology*. MIT Press, Cambridge, doi:10.7551/mitpress/9780262014229.001.0001.
- [14] Kawin Ethayarajh, David Duvenaud & Graeme Hirst (2018): *Towards Understanding Linear Word Analogies*, doi:10.48550/ARXIV.1810.04882. Available at <https://arxiv.org/abs/1810.04882>.
- [15] Ram Frost, Avital Deutsch & Kenneth I. Forster (2000): *Decomposing morphologically complex words in a nonlinear morphology*. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26(3), pp. 751–765, doi:10.1037/0278-7393.26.3.751.
- [16] Ram Frost, Kenneth I. Forster & Avital Deutsch (1997): *What can we learn from the morphology of Hebrew? A masked-priming investigation of morphological representation*. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23(4), pp. 829–856, doi:10.1037/0278-7393.23.4.829.
- [17] Ram Frost, Tamar Kugler, Avital Deutsch & Kenneth I. Forster (2005): *Orthographic Structure Versus Morphological Structure: Principles of Lexical Organization in a Given Language*. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31(6), pp. 1293–1326, doi:10.1037/0278-7393.31.6.1293.
- [18] Alex Gittens, Dimitris Achlioptas & Michael W. Mahoney (2017): *Skip-Gram - Zipf + Uniform = Vector Additivity*. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Vancouver, Canada, pp. 69–76, doi:10.18653/v1/P17-1007. Available at <https://aclanthology.org/P17-1007>.

- [19] Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin & Tomas Mikolov (2018): *Learning Word Vectors for 157 Languages*. CoRR abs/1802.06893, doi:10.48550/arXiv.1802.06893. arXiv:1802.06893.
- [20] Louis Guttman (1954): *Some necessary conditions for common-factor analysis*. *Psychometrika* 19(2), pp. 149–161, doi:10.1007/bf02289162.
- [21] Morris Halle & Alec Marantz (1993): *Distributed Morphology and the Pieces of Inflection*. In: *The View from Building 20*, MIT Press, Cambridge, MA, pp. 111–176.
- [22] Heidi Harley (2014): *On the identity of roots*. *Theoretical Linguistics* 40(3-4), pp. 225–276, doi:10.1515/tl-2014-0010.
- [23] Sandor G. J. Hervey & Jan W. F. Mulder (1973): *Pseudo-Composites and Pseudo-Words: Sufficient and Necessary Criteria for Morphological Analysis*. *La Linguistique* 9(1), pp. 41–70. Available at <http://www.jstor.org/stable/30248840>.
- [24] D. Jurafsky & J.H. Martin (2000): *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall series in artificial intelligence, Pearson Prentice Hall.
- [25] Qi Liu, Matt J. Kusner & Phil Blunsom (2020): *A Survey on Contextual Embeddings*. CoRR abs/2003.07278, doi:10.48550/arXiv.2003.07278. arXiv:2003.07278.
- [26] Alec Marantz (2000): *Roots: the universality of root and pattern morphology*. In: *conference on Afro-Asiatic languages, University of Paris VII*, 3, p. 14.
- [27] Tomas Mikolov, Kai Chen, Greg Corrado & Jeffrey Dean (2013): *Efficient Estimation of Word Representations in Vector Space*. In Yoshua Bengio & Yann LeCun, editors: *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*, doi:10.48550/arXiv.1301.3781. Available at <http://arxiv.org/abs/1301.3781>.
- [28] Sathvik Nair, Mahesh Srinivasan & Stephan C. Meylan (2020): *Contextualized Word Embeddings Encode Aspects of Human-Like Word Sense Knowledge*. CoRR abs/2010.13057, doi:10.48550/arXiv.2010.13057. arXiv:2010.13057.
- [29] James H Neely (2012): *Semantic priming effects in visual word recognition: A selective review of current findings and theories*. *Basic processes in reading*, pp. 272–344.
- [30] Karl Pearson (1901): *LIII. On lines and planes of closest fit to systems of points in space*. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2(11), pp. 559–572, doi:10.1080/14786440109462720.
- [31] Jeffrey Pennington, Richard Socher & Christopher Manning (2014): *GloVe: Global Vectors for Word Representation*. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, Doha, Qatar, pp. 1532–1543, doi:10.3115/v1/D14-1162. Available at <https://aclanthology.org/D14-1162>.
- [32] Manuel Perea & Eva Rosa (2002): *The effects of associative and semantic priming in the lexical decision task*. *Psychological Research* 66(3), pp. 180–194, doi:10.1007/s00426-002-0086-5.
- [33] Paul Martin Postal (1969): *Anaphoric Islands*. In: *Chicago Linguistic Society* 5, pp. 205–239.
- [34] Jean-Franois Prunet, Renee Beland & Ali Idrissi (2000): *The Mental Representation of Semitic Words*. *Linguistic Inquiry* 31(4), pp. 609–648, doi:10.1162/002438900554497. Available at <http://www.jstor.org/stable/4179126>.
- [35] Amit Seker, Elron Bandel, Dan Bareket, Idan Brusilovsky, Refael Shaked Greenfeld & Reut Tsarfaty (2021): *AlephBERT: A Hebrew Large Pre-Trained Language Model to Start-off your Hebrew NLP Application With*. CoRR abs/2104.04052, doi:10.48550/arXiv.2104.04052. arXiv:2104.04052.
- [36] Yishai Tobin (2004): *Hebrew (Semitic)*. In Geert Booij, Christian Lehmann, Joachim Mugdan, & Stavros Skopeteas, editors: *Morphologie/Morphology*, de Gruyter, Berlin and New York, pp. 1343–58, doi:10.1515/9783110172782.2.16.1343.

- [37] Eran Tomer (2012): *Automatic Hebrew Text Vocalization*. Master's thesis, Ben Gurion University of the Negev.
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser & Illia Polosukhin (2017): *Attention Is All You Need*. *CoRR* abs/1706.03762, doi:10.48550/arXiv.1706.03762. arXiv:1706.03762.
- [39] Hadas Velan & Ram Frost (2007): *Cambridge University versus Hebrew University: The impact of letter transposition on reading English and Hebrew*. *Psychonomic Bulletin & Review* 14(5), pp. 913–918, doi:10.3758/bf03194121.
- [40] Hadas Velan & Ram Frost (2011): *Words with and without internal structure: What determines the nature of orthographic and morphological processing?* *Cognition* 118(2), pp. 141–156, doi:10.1016/j.cognition.2010.11.013.