

Will it Unblend?

Yuval Pinter

School of Interactive Computing
Georgia Institute of Technology
Atlanta, GA, USA
uvp@gatech.edu

Cassandra L. Jacobs

Department of Psychology
University of Wisconsin
Madison, WI, USA
cjacobs2@wisc.edu

Jacob Eisenstein

Google Research
Seattle, WA, USA
jeisenstein@google.com

Abstract

Natural language processing systems often struggle with out-of-vocabulary (OOV) terms, which do not appear in training data. Blends, such as *innovator*, are one particularly challenging class of OOV, as they are formed by fusing together two or more bases that relate to the intended meaning in unpredictable manners and degrees. In this work, we run experiments on a novel dataset of English OOV blends to quantify the difficulty of interpreting the meanings of blends by large-scale contextual language models such as BERT. We first show that BERT’s processing of these blends does not fully access the component meanings, leaving their contextual representations semantically impoverished. We find this is mostly due to the loss of characters resulting from blend formation. Then, we assess how easily different models can recognize the structure and recover the origin of blends, and find that context-aware embedding systems outperform character-level and context-free embeddings, although their results are still far from satisfactory.

1 Introduction

For the token-based architectures that dominate contemporary natural language processing, a particularly difficult form of linguistic generalization arises from unseen phenomena at the word level, where novel sequences of characters, morphemes, or phonemes are known as out-of-vocabulary (OOV) terms (Brants, 2000; Plank, 2016; Heigold et al., 2017). Pretrained transformers like BERT (Devlin et al., 2019) handle OOV terms by **subtokenization**: segmenting all whitespace-delimited tokens into smaller units, from which any OOV term can be constructed (Sennrich et al., 2016).¹ But while this ap-

¹Another approach is to operate directly at the character level (e.g., Ling et al., 2015), but this has not been widely

proach is well suited for phenomena like concatenative English morphology, many linguistic processes generate OOV terms that cannot be cleanly decomposed into meaningful subtoken segments.

In this paper we address a particularly interesting and challenging source of OOV terms: novel **blends** (Algeo, 1977), also known as portmanteaux (Deri and Knight, 2015). Blends are constructed from the combination of multiple **bases** into a new form, in which some characters is shared across both bases: for example, *shop + optics = shoptics*. In this way, blends differ from other lexical compounds (e.g., *watermelon = water + melon*), which are formed by simple concatenation. Examples of OOV blends and their bases from our novel English blends dataset, collected from a natural source linked to the blends’ originating contexts (§2), are presented in Table 1. OOV blends are especially challenging to process, due to their combination of function-level semantic novelty with the form-level pathology of an unexpected character sequence.

Our main contribution is to offer what is to our knowledge the first analysis of how transformer-based contextual embedding models process novel blends and the representations they are able to produce for these challenging forms. First, we examine the impact of blends’ wordforms by comparing the ability of contextualized models to represent blends against the minimally-different case of novel lexical compounds. In §3, we show that the limited ability of contextual language models to represent novel blends’ components faithfully is primarily attributable to their form properties, whereas semantic differences between compounds and blends play a much smaller role. We

adopted within the mainstream framework of transformer-based models (Vaswani et al., 2017), due in part to the difficulty of scaling transformer-based models to sufficiently large contexts when operating at the character level.

Blend	PAXOBS	Bases	Semantic relation	Definition
<u>hatriotism</u>	AXXBBBBSSS	hate patriotism	ATTRIBUTE	Hate disguised as patriotism.
<u>shoptics</u>	AAXXBBBS	shop optics	LOC-PART-WHOLE	The social image projected when shopping.
<u>innoventor</u>	XXAAXBBXXX	innovator inventor	CAUSAL	A person who innovates by inventing.
<u>thrupple</u>	AAABOBBB	three couple	CONTAINMENT	A group of three people acting as a couple.

Table 1: A sample of the blends from the dataset, with definitions and our full annotation as described in §2. Linear blends are underlined.

then investigate how well several methods are able to recover the morphological boundaries within blends, which could mitigate the impact by splitting blends into segments contributed by each constituent base (§4.1). Finally, we attempt to recover the constituent bases given a segmentation (§4.2). Even under favorable conditions, we find that systems proposed previously for similar tasks struggle on blends, showing limitations of form-based and distributional similarity approaches. We propose a novel unsupervised base recovery method using contextualized masked language models, BERT RANKER. While this system performs well relative to others, we find substantial room for improvement. In our view, these results demonstrate the need for future work on our novel dataset and associated tasks.²

2 Complex Words Dataset

Our proposed investigation of the behavior of NLP systems on novel complex words requires a high-quality, reliable resource of truly novel blends and compounds in their original contexts, annotated for character sequence composition and semantic properties. The NYTWIT dataset (Pinter et al., 2020) contains English words new to the New York Times extracted by a bot³ between the dates of November 2017 and March 2019 with associated news article contexts. Words were annotated for their type of novelty.

We extract and further annotate three types from this dataset (version 1.1): blends (142 items), transparent compounds (121), and opaque compounds (49).⁴ The difference between the compound classes is semantic and somewhat subjective: transparent compounds have meanings which are comprehensible with little context (e.g. *quiz-maker*, a person who makes quizzes), while

²We release our code and data at <http://github.com/yuvalpinter/unblend>.

³www.twitter.com/NYT_First_Said

⁴Originally annotated as “compositional” and “new” compounds, respectively.

opaque compounds exhibit metaphoric or allusive semantics (e.g. *deathbox*, a dangerous car).

The first two authors annotated each word for its constituent bases, the character locations in which each base is represented, and the semantic relation between the bases. A sample of annotated blends is presented in Table 1. All disagreements resulting from the first round of blend base annotation (7%) were resolved by discussion, with the help of the words’ originating context. These contexts vary considerably in their length and informativity,⁵ but typically contain direct or indirect disambiguating information, and sometimes the component bases themselves: for blends, 40.3% of documents contained at least one of the bases within sentences where the blends appear (e.g. only *shop* or *optics*, for *shoptics*), while 10.2% contained both.

Semantic relations. An author annotated all blends and compounds in the dataset according to the well-studied semantic taxonomy of Tratz and Hovy (2010), which was designed for multiword nominal compounds (e.g. *cooking pot*). These relations were not intended to be applied to other types of phrases, blends, or lexical compounds. However, by referring to the official taxonomy, the expanded definitions in Dima and Hinrichs (2015), and the coarse- and fine-grained relation training data, the annotator was able to assign one of twelve coarse-grained relation classes to each word.⁶

As a preliminary check, we trained a relation classifier following the approach of Dima and Hinrichs (2015), a single-hidden-layer classifier over GloVe embeddings (Pennington et al., 2014), on the RANDOM partition of the Tratz and Hovy

⁵Compare, for example, the following context sentences: “*Blaspy?*”; “The *procrastibaker* must believe that it is possible to be simultaneously working on a document, buttering pans and separating eggs.”

⁶For example, ATTRIBUTE was applied to the adjective-noun blend *fitfluencers*. 69% of the words contain a non-noun base.

(2010) data.⁷ This model achieved .203 accuracy and .173 macro-F1 on our dataset for all 311 items, substantially higher than baselines such as majority class (.087 acc. / .013 F1) and random prediction calibrated to the marginal label distributions (.106 acc. / .078 F1 for the best of ten runs), indicating credible annotation.⁸ This performance is still poor relative to multiword compounds, possibly due to the fundamentally different linguistic processes governing lexical compounding and blending processes as opposed to multiword compounding.

Character-level labels. We introduce a character-level labeling schema to help classify blend types and evaluate and train blend segmentation models, called **PAXOBS**.⁹ Each character is labeled as P or S if it is in a prefix or suffix, respectively; as X or O if it is contributed by more than one or none of the bases, respectively; and by successive letters of the alphabet for characters from only each base, typically just A or B.¹⁰ This schema covers the full range of processes undergone in blending, except for annotation of characters removed altogether from the bases (e.g. the *e* from *hate* in *hatriotism*), and may be trivially applied to lexical compounds as well.

Blends may be classified into further subcategories based on the correspondence between their form and the bases. For example, **linear** blends are similar to compounds, as each base’s portion appears uninterrupted in the blend (underlined in Table 1). Formally, a blend is linear if its label sequence contains: no O; no A preceded by a B or X; and no B followed by an A or X. In our dataset 59% of blends are linear, though prior work has reported up to 95% linear blends among blends extracted from a curated lexicon (Cook and Stevenson, 2010). One possible explanation for this discrepancy is that words that make it into common use may be simpler in their surface quality.

⁷The model trained on this split, set to $d_e = d_h = 50$, slightly outperformed an identical one trained on the LEXICAL split, and its test set accuracy on the original dataset is .721, close to replication. The LEXICAL split was created to correct an over-representation of some compound bases in RANDOM, one which biases statistical models toward lexical memorization (see details in §4.1 of Shwartz and Waterson (2018)), but has no bearing on our dataset.

⁸The numbers for blends only are .148 / .079 vs. .063 / .020 for majority class.

⁹May be pronounced like “pack sobs”.

¹⁰This framework is loosely similar to *edit scripts* (Chrupala et al., 2008), but rather than transducing one string into another, the task is to combine two strings into a third.

3 Blends in Context

Novel blends are a unique linguistic phenomenon, posing challenges for automated systems on many different levels. However, the sparsity of their appearances in real-world text, as well as the expertise required for creating a natural language understanding task which uses specific documents from a large variety of domains as supporting information, make the evaluation of the effect of novel blends on this type of downstream task an impractical goal for the scope of this work. Instead, we assess the treatment of novel blends at the representational step of contemporary contextualized language models, by performing an analysis of their processing by BERT (Devlin et al., 2019). To gauge how well BERT represents blends, we conduct a comparison with its treatment of a minimally-different control class of novel words, namely **lexical compounds**. These are forms where at least two bases are concatenated in full (e.g. *quizmaker*), without the character loss incurred in blends.

Our analysis begins with the assumption that in any given context, the meaning representation of a complex word (blend or lexical compound) must be **composed** from its bases, which we can estimate using the representational similarity between a complex word and its bases in the same context. This criterion can be viewed as a form of linguistic generalization, and if satisfied, enables downstream models to produce consistent results across related words. To test this criterion, we compute the vector similarities between the contextualized representations of complex words and their components, a method that coheres with human judgments of contextual semantic similarity (Giulianelli et al., 2020). We probe BERT¹¹ with synthetic inputs constructed by replacing each complex word with its space-delimited bases. Formally, given a sentence $S = (w_1, \dots, w_{i-1}, x, w_{i+1}, \dots, w_n)$ where x denotes a blend or compound with contributing bases b_1, b_2 , we record the average vector across x ’s wordpiece tokens for each layer output in BERT’s transformer stack, $e^{(l)}(x)$, $l \in 0, \dots, 12$, and compute its cosine similarity with the averaged vectors $\frac{1}{2}(e^{(l)}(b_1) + e^{(l)}(b_2))$ in the sequence $S' = (w_1, \dots, w_{i-1}, b_1, b_2, w_{i+1}, \dots, w_n)$.¹²

¹¹We use the `base-uncased` flavor and the Huggingface implementation (Wolf et al., 2019) throughout the paper.

¹²Four forms in our dataset have three bases: “fanimatic”

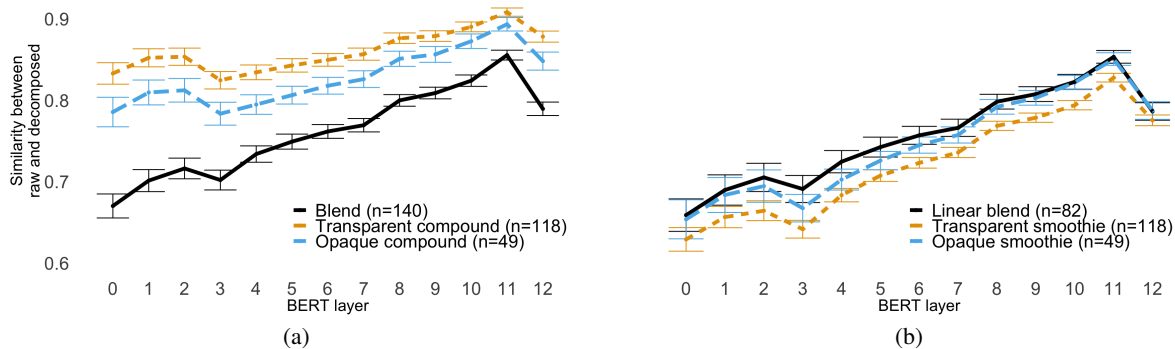


Figure 1: Pretrained BERT’s layer-wise similarity between representations of (a) complex OOVs and their base components; and (b) linear blends and “smoothies” (§3.2), lexical compounds forced to lose characters while remaining linear. All representations are computed using the original context in which the words appear. Error bars represent standard error of the mean over the class.

Figure 1(a) compares the per-layer similarities for blends with the two types of compounds described in §2: We find a clear distinction between blends and both compound classes. For compounds, BERT induces representations that are very similar to those of the components at all layers of the model. For blends, these representations diverge greatly, especially in the lower layers of the model, which capture surface-form characteristics of the input (Jawahar et al., 2019).

Since the difference between classes exists across all layers, we first wish to perform a more thorough analysis of possible reasons for it.

3.1 Semantics

One possible explanation for the difference in BERT’s treatment of blends and lexical compounds is that blends arise in lexical situations that are qualitatively different from those in which compounds are formed. This would lead to a different distribution of semantic relationships between bases of blends and compounds. In our annotated dataset we were able to witness such differences; for example, the ATTRIBUTE relation accounts for 23% of compounds but 38% of blends.

If BERT’s divergent treatment of blends and compounds is explained by the distribution over

= fan + animation + cinematic, “shaggydoodle” = *shaggy + labrador + poodle*, “frenemesis” = *friend + enemy + nemesis*, “orchestraits” = *orca + orchestrates + straits*. In these cases, we include the vectors for all three bases. One blend, “pregret”, has only one base, against which it is compared.

Five words are missing from the analysis as they no longer appeared in their original contexts at scraping time due to editorial actions on the NYT website: the blends “humailiation” and “crapberg”; and the compounds “cybersensation”, “garagerock”, and “storytale”.

semantic categories for each complex word type, then we would expect the similarity scores *within* categories to be identical. Repeating the contextual similarity analysis within each semantic category, we find that there are substantial divergences between blends and compounds in several of the semantic categories. Figure 2 presents the similarity scores for the six relations containing at least 15 observations; blend representations are less similar to their decomposed versions compared to compounds regardless of the relation. A linear model trained to predict similarity confirms that blends are less similar to their components than compounds ($\rho = -.128, p < .001$).

3.2 Form

Another potential explanation is that differences in BERT’s treatment of blends and lexical compounds are driven by the form of each compound, rather than the meaning. On this view, the choice of whether to create a compound or a blend is a stylistic one (Renner, 2015), and so controlling for the character loss incurred in blends would produce the same processing difficulty for compounds.

Smoothies. If differences in surface form are what drives differences in contextualized representations, then transforming the compounds into mock-blends, which we term “smoothies”, should eliminate the differences between the two complex word types: we would expect the function of the similarity of a contextual encoding of a blend to its bases given a context it naturally occurs in to be approximately the same function of simi-

larity of a contextual encoding of a smoothie to its bases given the context the original compound occurs in. We create our smoothies using COPYCAT (Kulkarni and Wang, 2018), a model which generates blends from two base forms via a sequence of character `copy` and `delete` actions learned over features extracted from an language model, an LSTM, and length-based heuristics. We train an ensemble of 50 COPYCAT models on the blends from Deri and Knight (2015) and apply them to our novel compounds.¹³ Since COPYCAT can produce only linear blends, we compare the BERT correspondence for smoothies against linear blends only (whose aggregate similarities are notably similar to those of blends as a whole). In creating the smoothies,¹⁴ we made sure that the overall rate of lost characters (`delete` operations) is comparable to that of the true linear blends. We show in Figure 1(b) that smoothies pose similar generalization challenges as blends: the gap between linear blends and smoothies is small, while generalization for smoothies is far below that of the original compounds.

Tokenization. Having established that surface form is a main driver of the representational differences between blends and compounds, we now assess the specific impact of BERT’s tokenization model, WordPiece (WP). WordPiece is a trained model, consisting of a subword vocabulary constructed by identifying units (pieces) that appear repeatedly in a corpus. It distinguishes between word-initial pieces, which may be whole words, and word-noninitial pieces which are marked by a special “##” prefix. A word is then assigned a sequence of pieces whose characters matches it when concatenated. For example, WP(“segmenting”)=[‘segment’, ‘##ing’]. Such a model might be poorly suited to novel blends, which by definition reuse characters across bases, and which cannot be analyzed by traditional patterns of morphology.

To test the effect of segmentation, we provide WP with base-congruent segmentation points informed by their PAXOBS tags: for example, *shop-tics* is fed to BERT as $\rightarrow sh+##op+##tics$. We find that this change does little to bridge the gap between blends and compounds: a redrawn ver-

¹³We run the model ten times, and average the BERT differences over each base pair’s resulting smoothies before aggregating for categories.

¹⁴Examples include *bow + person = boerson* and *junk + time = junime*.




Figure 2: Pretrained BERT’s similarity measures for each semantic relation with $n > 15$ instances.

sion of Figure 1(a) using this tokenization is almost identical to the original. Upon further examination, we find that while pre-tokenizing with PAXOBS results in a larger number of wordpiece tokens (an average of 4.55 vs. 3.30), a similar leap occurs in compounds (3.41 vs. 2.48), suggesting that WP does not produce morphologically accurate segments for compounds either (Bostrom and Durrett, 2020). The crux of the issue must therefore lie within BERT.

In conclusion, we have shown that the root cause of blend mistreatment in large contextual transformer models is their form, although knowing only their sequence structure is not sufficient. Therefore, in the following section we suggest models which attempt to identify blend segmentation points, but also ones which attempt recovery of their original bases, in order to place them in an appropriate topical context.

4 Will it Unblend?

We next test to what extent existing models can help systems understand the meaning of novel blends, an aspect of human language understanding that has been little explored in NLP evaluation tasks. As demonstrated in §3, successfully

representing blends requires the capability to both properly decompose their form and identify the original constituents. We therefore cast blend understanding as two tasks: segmenting blends into character sequences (§4.1), and recovering blends’ bases post-segmentation (§4.2). We leave the task of recognizing blends to future work.¹⁵

Compounds. Compounds were used as a comparative class in §3, but for the purpose of form understanding we focus on blends. For compounds with a known segmentation, base recovery is trivial, as each side of the segmentation point is always a base. As for segmentation, we have shown in §3.2 that BERT’s transformer layers are capable of recovering from poor WordPiece performance, and so the utility of segmenting compounds is limited compared to blends.¹⁶ Knowing that words are kept in their original form can define much simpler and more effective systems of discovery than the ones described below, such as a dictionary lookup of both sides for each possible single segmentation point.

4.1 Blend Segmentation

We approximate the problem of inferring blend structure by defining a segmentation task over the character sequence which is the blend form, on the rationale that supplying a downstream system with character segments, each coherently representing a single known word or morpheme, would improve its ability to represent the input sequence. For example, a character-aware system familiar with non-complex words might understand that the initial *hat* from *hatriotism* is related to *hate* if given in isolation; but with *hatr* it would be at a loss.

Metrics. We draw on our PAXOBS schema (§2) to define segment-level precision and recall scores for a given blend (e.g. *shoptics*: AAXXBBS). A system’s prediction is a set of character indices where segmentation should occur. We count any index which separates characters of the same label as a false positive, towards precision (e.g. the seg-

¹⁵Note that the best OOV classification baseline in Pinter et al. (2020), a ridge classifier using character ngram features, reaches .305 F1 on the blends class (its macro-F1 on all 18 classes is .323).

¹⁶We nevertheless evaluated the segmenters on compounds (compare with Table 2). WP performs about as well as on blends in F1 (.558), and better in exact match (34%), and Domain Unigram LM outperforms it on both (.636, 39% respectively). In compounds, lenient and strict metrics converge.

Model	#segs	Prec.	L F1	S F1	L EM
All-chars	10.15	.272	.427	.427	0%
Seq. tagger	4.70	.291	.400	.376	5%
WordPiece	3.30	.450	.562	.484	22%
Domain BPE	3.18	.408	.517	.441	24%
Domain ULM	4.08	.428	.556	.492	22%

Table 2: Results for segmentation ($N = 142$, micro-aggregation). “L” – lenient, “S” – strict, “EM” – exact match. The desired number of segments varies between 2.05 and 3.40, depending on affixes and recall policy.

mentation in [*shopti;cs*]). False negatives may be defined **strictly** or **leniently**: in the strict evaluation, a false negative is any segment that contains characters belonging to more than one base, or to a base and shared material (X), while a lenient evaluation permits the inclusion of shared material: [*shop;tics*] is leniently sound, but [*sh;op;tics*] is strictly sound as well. We report micro-level precision, as well as F1 computed with both lenient and strict recall, and lenient exact match. We ignore prefixes and suffixes, and allow models to freely separate or include them in the adjacent base.

Systems. We compare the following systems (see Appendix A.2 for implementation details):

- **All-chars.** A baseline which marks every character as its own segment (perfect recall).
- **Sequence Tagger.** We annotated the 1,579 blends in Gangal et al. (2017)’s dataset¹⁷ for PAXOBS tags and used them for training a supervised neural character-level tagger, whose results are converted into segmentations.
- **WordPiece.** We run WP “out of the box”.
- **In-domain Subwords.** We train BPE (Sennrich et al., 2016) and Unigram LM (Kudo, 2018) subword tokenizers on news data from the Corpus of Contemporary American English (1990–2015; Davies, 2008) using the sentencepiece package (Kudo and Richardson, 2018), set to the same vocabulary size as the WP model.

Results. The results in Table 2 show that all models struggle to find correct segmentation, even compared to the all-chars baseline. The low performance of the supervised tagger suggests that little can be inferred from relative character placement, demonstrating the highly variable nature of novel blends. Corpus-based segmentation models manage to segment over 20% of the

¹⁷After filtering duplicates from its 1,624 lines.

blends successfully. This number is slightly higher when looking at the subset of linear blends: Domain BPE matches 29% of them exactly. Further analysis of the WP segmentations reveals a weakness in cases where the first post-A characters suggest a plausible continuation to base A , common enough to appear in WP’s vocabulary, e.g. [males;tream] (true bases *male*, *mainstream*; labels XXAABBBBBB), or [chip;ster] (true bases *chicano*, *hipster*; labels AXXBBBBB).

We next consider the challenge of reconstructing the base components for segmented blends.

4.2 Blend Component Recovery

We tasked different models with identifying the contributing bases (A , B)¹⁸ out of all possible words given a gold-segmented blend and an input vocabulary. We create sets of candidate words for each blend which could, in principle, create the same blend as the true bases. For example, the blend *thrupple* = *three* + *couple* will induce candidates such as *thrash* for A and *example* for B . We report the following metrics:

- **MRR**-(A , B , ω) is the mean reciprocal rank of true base A/B across all possible candidates for that side (*single-side prediction*), or (ω) of the true base pair out of all possible base pair candidates (*pair prediction*);
- **Precision @1** is the proportion of blends for which the top candidate is the true base pair.¹⁹

In order to maintain a fair comparison between the models (see below), we extracted the candidate lists for all model evaluations from the GloVe (Pennington et al., 2014) model’s vocabulary, as it is the only one restricted for invocabulary testing. In total, 33 of the candidate lists (12%) are singletons, including two blends where neither base has negative samples. Six blends (4%) lacked the correct base for one of the sides, and these cases were treated as ranked last among candidates; three of these lists were empty, translating to in a #1 rank for all systems. The lower bounds on the metrics resulting from these

¹⁸In three-based blends, we denote the last base as B for this task, surmising it is more important than the second base given the right-headedness tendency of English.

¹⁹This contrasts with precision as reported by Cook and Stevenson (2010), who count a pair as correct if at least one base is correct.

candidate list limitations are presented at the top of Table 3.²⁰

BERT RANKER. We propose a contextual representational approach for ranking two-sided base candidates using iterative piece prediction: we replace each appearance of a blend b in its context sentence $(w_1, \dots, w_{i-1}, b, w_{i+1}, \dots, w_n)$ with two successive [MASK] tokens: $(w_1, \dots, w_{i-1}, m_1, m_2, w_{i+1}, \dots, w_n)$. Then, we use a pretrained BERT masked language model to compute wordpiece prediction distributions for these masked tokens. We sort all possible candidate base pairs $\langle l, r \rangle$ according to the sum of probabilities for their bases’ first pieces, $P(m_1 = l_0) + P(m_2 = r_0)$,²¹ and record the rank of the true base pair. When candidate pairs have the same first-piece pair, we break ties by iteratively predicting the next pieces after inserting the shared pieces into the input. We also implement an ablation (–CONTEXT) where no context is added to the masks, in order to evaluate the contribution of the sentence contexts in isolation.

For single-side metrics, we report the rank of the true base in the prediction distribution of a single [MASK] token (instead of two); in another variant (+OTHER-BASE) we add the true base from the other side to the context, in order to level the playing field with the baselines, which we describe next (see Appendix A.3 for implementation details):

- **Character RNN.** We separately train a forward and a backward character-level RNN on over 100,000 documents from the Westbury corpus (Shaoul, 2010). We feed the blend’s left (right) context to the forward (backward) RNN, then record the probability of each A (B) candidate as a continuation of the context, computed as the average of character log-likelihoods.
- **Edit distance (ED).** Following Cook and Stevenson (2010), we compare the string similarity (Levenshtein distance) between base candidate pairs’ orthographic forms.²² Single-side

²⁰The full lists are available on the repository.

²¹This is crucial, since predicting a ‘##’-initial suffix token effectively attaches it to the preceding token

²²A variant using phonological forms, extracted from a phonological lexicon (Lee et al., 2020), was limited by only having pronunciations for a fraction of bases and candidates. In cases where both base pronunciations were found the ranking was good, hinting at a promising avenue for future work by implementing automatic text-to-phone modeling.

Model	MRR-			P@1
	A	B	ω	
Lower bound	.115	.257	.036	.014
Character RNN	.162	.368	.060	.021
Edit distance	.176*	.432*	.066	.014
fastText	.357*	.610*	.167	.127
GloVe	.449*	.734*	.188	.127
BERT RANKER	.392	.711	.288	.264
+OTHER-BASE	.403*	.703*		
-CONTEXT	.379	.675	.147	.127
+OTHER-BASE	.379*	.668*		

Table 3: Results for component recovery. *Results dependent on knowledge of the correct base on the other side.

prediction fixes one base and ranks candidates from the other side based on similarity.

- **Static embeddings.** We calculate cosine similarity between candidate base pairs’ embeddings in fastText (Mikolov et al., 2018) and GloVe (Pennington et al., 2014). fastText includes character n-grams, allowing an assessment of the utility of subword information.

To summarize, both **ED** and **Static** methods are contextless pair-matchers which operate in the +OTHER-BASE knowledge setup when evaluated for MRR-A and MRR-B; **Character RNN** is a single-base ranker which uses context from one side only and cannot be helped by knowledge of the other base.

Results. Results are presented in Table 3. We note the higher performance on B bases achieved by all models, a fact which advantages WordPiece which leaves word-initial pieces unmarked (see §3.2), as opposed to models such as XLM (Conneau et al., 2019), which mark word-final pieces. If the beginning of the blend bears more resemblance to the base it originated from, there’s a better chance of properly representing that base in the overall blend. These findings suggest an iterative setup, where first the B base is predicted and only then A is matched, might prove more successful. We leave this variant to future work.

Our BERT RANKER model outperforms all baselines in the more realistic full-word setting (MRR- ω , P@1). When ranking single bases, it does not benefit much from awareness of the true other base (the oscillations recorded in the table are too small to be meaningful), suggesting that

most of its power lies in processing context and not in word form representation. This conclusion is further supported by the superior performance of the static type-level GloVe embeddings, whose lead over fastText and BERT-CONTEXT in all MRR measures suggests that word form is less helpful even in uncontextualized settings. The particularly poor performance of the character RNN and edit distance model shows that it is difficult to learn the task without any semantic signal.

Error Analysis. A qualitative assessment of the contexts which help BERT RANKER to predict bases perfectly relative to the -CONTEXT variant shows that they typically contain one or more of the bases in their entirety (e.g., *eggcessories* appears near multiple occurrences of the word *eggs*). By contrast, in some longer contexts containing diverse topics, the inclusion of context wipes out the accessibility of the component bases, typically the first one (e.g. *chesticle*, in which the context does not mention body parts, or *cancerchondria* which mentions the word *condition* but neither of the bases).

5 Related Work

Prior work on blends has largely focused on **generation** (e.g., Das and Ghosh, 2017; Simon, 2018; Deri and Knight, 2015; Kulkarni and Wang, 2018; Smith et al., 2014). While Gangal et al. (2017) provide a unified dataset of 1,579 blends, annotated for bases, they do not provide contexts for real-world appearances of the blends, nor a breakdown of the semantic relationship between their constituents. Moreover, some are synthetically generated by a seq2seq model. In addition, these works all restrict their models to linear two-word blends. Our PAXOBS scheme handles nonlinear and multi-base blends.

Cook and Stevenson (2010) presented a non-contextual method for blend base detection using a dictionary-based lexicon, evaluated over an unreleased dataset, and Ek (2018) used features from static embeddings to unblend words in Swedish. We adopt the candidate-ranking approach of these works to evaluate component recovery, but incorporate context with context-sensitive language models, and add the task of blend segmentation.

Extracting the semantics of constituents from larger phrases is not a problem unique to single-token blends. Shwartz and Waterson (2018) worked on multi-word compounds; Maddela et al.

(2019) segment hashtags, roughly half of which are akin to our notion of compounds, by training a neural scoring system over features extracted from word form, dictionary lookup and language model probabilities. Another connection is to the learning of morphological rules, e.g., for processes such as derivation (Kondratyuk, 2019) and lemmatization (Chrupala, 2006; Ullman et al., 1976; Hirschberg, 1977). Cotterell and Schütze (2018) present a supervised model of derivational morphology that jointly accounts for segmentation as well as composition of static word embeddings from the embeddings of morphemes, thereby touching on two of the main tasks undertaken in our paper. However, the application of such a model to blends is complicated by the relative lack of labeled training data, as well as the irregularity of the underlying phenomenon.

Novel blends are an example of linguistic creativity, which frequently operates at the sub-word level. Related phenomena include **eggcorns**, which are alternative spellings that yield an apparently more transparent relationship between form and function (Reddy, 2009); **puns**, which substitute words in new contexts based on phonological similarity (Jaech et al., 2016); respellings that attempt to reintroduce prosodic expression into spelling (Brody and Diakopoulos, 2011); intentional obfuscation (Zalmout et al., 2019); and typographical errors (Heigold et al., 2017). We therefore view blends as an instance of a broad set of creative phenomena that poses challenges for the token-based approaches that currently dominate natural language processing.

6 Conclusion

This work focuses on the challenge of interpreting novel blends, which requires integrating sub-word structure and contextual features. We present a new dataset annotated using a novel character-level schema as well as for semantic tags, and offer preliminary evaluations showing that (a) blends are handled differently than compounds by BERT, due mostly to the phenomenon of character loss; (b) existing tokenizers generally do not respect blend boundaries; and (c) recovering the components of a blend is a difficult task which challenges word-form, distributional, and contextual approaches. Our results further highlight that annotation schemata such as those of [Tratz and Hovy \(2010\)](#), which were designed for noun com-

pounds, are generalizable to other relational word types. In future work, we plan to integrate these signals into a better blend processor, and to further address the effect of blends on downstream tasks from semantic and syntactic viewpoints. In addition, we aim to further examine the methods from our experiments on other classes of novel words and in other languages. We also plan to add phonetic resources for improving treatment of nonlinear blends.

Acknowledgments

We thank Max Bittker for providing us with details and logs from the NYT bot. We thank Lelia Glass, Kyle Gorman, Arya McCarthy, Ian Stewart, Kristina Toutanova, Sarah Wiegrefe, Yuwei Wu, Wei Xu, Diyi Yang, and the anonymous reviewers for their valuable notes. YP is a Bloomberg Data Science PhD Fellow. CJ is a postdoctoral research associate supported by NSF grant 1849236 awarded to Maryellen MacDonald.

References

- John Algeo. 1977. Blends, a structural and systemic view. *American speech*, 52(1/2):47–64.
- Kaj Bostrom and Greg Durrett. 2020. Byte pair encoding is suboptimal for language model pretraining. *arXiv preprint arXiv:2004.03720*.
- Thorsten Brants. 2000. Tnt: a statistical part-of-speech tagger. In *Proceedings of the sixth conference on Applied natural language processing*, pages 224–231. Association for Computational Linguistics.
- Samuel Brody and Nicholas Diakopoulos. 2011. Coo: using word lengthening to detect sentiment in microblogs. In *Proceedings of the conference on empirical methods in natural language processing*, pages 562–570. Association for Computational Linguistics.
- Grzegorz Chrupala. 2006. Simple data-driven context-sensitive lemmatization. *Procesamiento del lenguaje natural*, (37):121–130.
- Grzegorz Chrupala, Georgiana Dinu, and Josef van Genabith. 2008. [Learning morphology with Morfette](#). In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC’08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.

- Paul Cook and Suzanne Stevenson. 2010. [Automatically identifying the source words of lexical blends in English](#). *Computational Linguistics*, 36(1):129–149.
- Ryan Cotterell and Hinrich Schütze. 2018. [Joint semantic synthesis and morphological analysis of the derived word](#). *Transactions of the Association for Computational Linguistics*, 6:33–48.
- Kollol Das and Shaona Ghosh. 2017. [Neuramanteau: A neural network ensemble model for lexical blends](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 576–583.
- M. Davies. 2008. [The corpus of contemporary american english: 450 million words, 1990-present](#).
- Aliya Deri and Kevin Knight. 2015. [How to make a frenemy: Multitape FSTs for portmanteau generation](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 206–210, Denver, Colorado. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Corina Dima and Erhard Hinrichs. 2015. [Automatic noun compound interpretation using deep neural networks and word embeddings](#). In *Proceedings of the 11th International Conference on Computational Semantics*, pages 173–183, London, UK. Association for Computational Linguistics.
- Adam Ek. 2018. [Identifying source words of lexical blends in swedish](#). *Thanks to our sponsors: Gold: Stora Skuggans Vårdshus Silver: TT Nyhetsbyrå and Lingsoft Bronze: Convertus, Digital Grammars, IQVIA and Voice Provider*, page 82.
- Varun Gangal, Harsh Jhamtani, Graham Neubig, Eduard Hovy, and Eric Nyberg. 2017. [Charmanteau: Character embedding models for portmanteau creation](#). *arXiv preprint arXiv:1707.01176*.
- Mario Giulianelli, Marco Del Tredici, and Raquel Fernández. 2020. [Analysing lexical semantic change with contextualised word representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Seattle, USA. Association for Computational Linguistics.
- Georg Heigold, Günter Neumann, and Josef van Genabith. 2017. [How robust are character-based word embeddings in tagging and mt against word scrambling or random noise?](#) *arXiv preprint arXiv:1704.04441*.
- Daniel S. Hirschberg. 1977. [Algorithms for the longest common subsequence problem](#). *J. ACM*, 24(4):664–675.
- Aaron Jaech, Rik Koncel-Kedziorski, and Mari Ostendorf. 2016. [Phonological pun-derstanding](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 654–663.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *arXiv preprint arXiv:1412.6980*.
- Dan Kondratyuk. 2019. [Cross-lingual lemmatization and morphology tagging with two-stage multilingual BERT fine-tuning](#). In *Proceedings of the 16th Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 12–18, Florence, Italy. Association for Computational Linguistics.
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. [Sentence-Piece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Vivek Kulkarni and William Yang Wang. 2018. [Simple models for word formation in slang](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1424–1434, New Orleans, Louisiana. Association for Computational Linguistics.
- Jackson L. Lee, Lucas F.E. Ashby, M. Elizabeth Garza, Yeonju Lee-Sikka, Sean Miller, Alan Wong, Arya D. McCarthy, and Kyle Gorman. 2020. [Massively multilingual pronunciation modeling with WikiPron](#). In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 4223–4228, Marseille, France. European Language Resources Association.

- Wang Ling, Chris Dyer, Alan W Black, Isabel Trancoso, Ramón Fernández, Silvio Amir, Luís Marujo, and Tiago Luís. 2015. [Finding function in form: Compositional character models for open vocabulary word representation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1520–1530, Lisbon, Portugal. Association for Computational Linguistics.
- Mounica Maddela, Wei Xu, and Daniel Preotjiuc-Pietro. 2019. [Multi-task pairwise neural ranking for hashtag segmentation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2538–2549, Florence, Italy. Association for Computational Linguistics.
- Tomas Mikolov, Edouard Grave, Piotr Bojanowski, Christian Puhersch, and Armand Joulin. 2018. Advances in pre-training distributed word representations. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [PyTorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Yuval Pinter, Cassandra L Jacobs, and Max Bittker. 2020. NYTWIT: A dataset of novel words in the New York Times. In *Proceedings of the 28th International Conference on Computational Linguistics*, Online. Association for Computational Linguistics.
- Barbara Plank. 2016. What to do about non-standard (or non-canonical) language in NLP. *arXiv preprint arXiv:1608.07836*.
- Sravana Reddy. 2009. [Understanding eggcorns](#). In *Proceedings of the Workshop on Computational Approaches to Linguistic Creativity*, pages 17–23, Boulder, Colorado. Association for Computational Linguistics.
- Vincent Renner. 2015. [Lexical Blending as Wordplay](#), pages 119 – 134. De Gruyter, Berlin, Boston.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Cyrus Shaoul. 2010. The Westbury lab Wikipedia corpus. *Edmonton, AB: University of Alberta*, page 131.
- Vered Shwartz and Chris Waterson. 2018. [Olive oil is made of olives, baby oil is made for babies: Interpreting noun compounds using paraphrases in a neural model](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 218–224, New Orleans, Louisiana. Association for Computational Linguistics.
- Jonathan A Simon. 2018. Entendpreneur: Generating humorous portmanteaus using word-embeddings. *NeurIPS*.
- Michael R Smith, Ryan S Hintze, and Dan Ventura. 2014. Nehovah: A neologism creator nomen ipsum. In *ICCC*, pages 173–181.
- Stephen Tratz and Eduard Hovy. 2010. [A taxonomy, dataset, and classifier for automatic noun compound interpretation](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 678–687, Uppsala, Sweden. Association for Computational Linguistics.
- J. D. Ullman, A. V. Aho, and D. S. Hirschberg. 1976. [Bounds on the complexity of the longest common subsequence problem](#). *J. ACM*, 23(1):1–12.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rmi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Transformers: State-of-the-art natural language processing](#).
- Nasser Zalmout, Kapil Thadani, and Aasish Pappu. 2019. [Unsupervised neologism normalization using embedding space mapping](#). In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 425–430, Hong Kong, China. Association for Computational Linguistics.