

Semantics of Multiword Expressions in Transformer-Based Models: A Survey

Filip Miletic and Sabine Schulte im Walde

Institute for Natural Language Processing, University of Stuttgart, Germany

{filip.miletic, schulte}@ims.uni-stuttgart.de

Abstract

Multiword expressions (MWEs) are composed of multiple words and exhibit variable degrees of compositionality. As such, their meanings are notoriously difficult to model, and it is unclear to what extent this issue affects transformer architectures. Addressing this gap, we provide the first in-depth survey of MWE processing with transformer models. We overall find that they capture MWE semantics inconsistently, as shown by reliance on surface patterns and memorized information. MWE meaning is also strongly localized, predominantly in early layers of the architecture. Representations benefit from specific linguistic properties, such as lower semantic idiosyncrasy and ambiguity of target expressions. Our findings overall question the ability of transformer models to robustly capture fine-grained semantics. Furthermore, we highlight the need for more directly comparable evaluation setups.

1 Introduction

Multiword expressions (MWEs)—such as noun compounds (e.g., *jet lag*), particle verbs (e.g., *take off*), and idioms (e.g., *on the fly*)—are composed of multiple words and exhibit semantic idiosyncrasy, i.e., their overall meaning cannot be directly predicted from the meanings of their constituents. They are ubiquitous, affect various applications, and as such have been extensively addressed in NLP research (Sag et al., 2002; Baldwin and Kim, 2010). Although the widely used transformer-based language models have been analyzed regarding their ability to represent various types of linguistic knowledge, we still lack consolidated insights into their processing of MWE semantics. The present survey provides a critical overview of existing work on this issue.

By definition, the meaning of a MWE is distributed over multiple constituents.¹ In some cases, these can be separated by intervening material (e.g., *turn the volume up*). For any model, it is more challenging to capture the meaning of multiple lexical elements than that of a single word. The overall meaning may further be compositional to various degrees, i.e., similar to the sum of the component parts (e.g., *climate change*) or unrelated to it (e.g., *silver bullet*). Moreover, the same expression may be interpreted more or less compositionally depending on its context (e.g., *kick the bucket*).

Transformer models are assumed to easily overcome some of these challenges since their meaning representations are inherently contextual and successful on a wide array of semantic tasks (Devlin et al., 2019; Brown et al., 2020). But does this imply a robust ability to represent MWE meanings? The answer to this question is not trivial, as such an ability requires the models to systematically (i) capture the semantic contributions of multiple tokens, potentially including figurative and rare meanings; (ii) weigh them based on variable degrees of compositionality; and (iii) further specify the interpretation in context.

In order to establish the current state of knowledge on this issue, we provide the first in-depth survey of the fast-growing body of work on MWE representations in transformer models. We aim to understand the extent to which different pre-trained and optimized models are able to capture MWE semantics, as well as whether they are affected by representational or linguistic factors. Most studies focus on a single expression type (e.g., noun compounds), task (e.g., machine translation), or related wider issue (e.g., phrase representations). We broaden the perspective by

¹This also applies to closed compounds realized as a single orthographic unit (e.g., *flashback*) which comprises clearly identifiable constituents (*flash* and *back*).

systematizing insights from often disjoint strands of research, and identify priorities for future work. Our findings more generally highlight that transformer models represent complex linguistic knowledge inconsistently.

We first provide general background on the surveyed MWE approaches (§2) and then zoom in to transformer models. We explore whether they inherently capture MWE information (§3), where it is localized (§4), and whether it is affected by linguistic properties (§5). We conclude with a summary and directions of future work (§6).

2 Survey Overview

This section motivates our survey by providing background on lexical semantic representations in transformer models, and then summarizing the most frequent implementations and tasks.

2.1 Lexical Semantics in Transformer Models

Language models based on the transformer architecture process information through a series of layers relying on the multi-head attention mechanism, which weighs each token in a sequence based on its similarity to the other tokens (Vaswani et al., 2017). As a sequence progresses through the layers, the representations become gradually more contextualized (Ethayarajh, 2019) and ultimately capture lexical semantic properties such as word senses (Wiedemann et al., 2019). This mechanism should benefit MWE processing by distributing contextually provided semantic information over multiple tokens, but it is unclear where the models encode MWE information and to what extent.

Different types of linguistic information are not represented in the same way in the transformer architecture (Rogers et al., 2020). It has been suggested that surface features are localized in the lower, syntactic features in the middle, and semantic features in the higher layers (Jawahar et al., 2019). But while higher layers do encode word senses (Coenen et al., 2019), type-level lexical information is better accessed in lower layers (Vulić et al., 2020). Moreover, transformer models are affected by spurious effects of sequence position (Mickus et al., 2020), which questions the general robustness of their semantic representations. The models also struggle with logical phenomena such as negation (Ettinger, 2020), and are

not strongly dependent on word order but learn higher-level distributional patterns (Sinha et al., 2021). These observations indicate difficulties in capturing contextually provided meanings of multiple tokens—a key property required for MWE semantics. The mechanisms behind these issues may affect MWE representations, both in terms of overall reliability and features specific to this type of expression.

2.2 Summary of Surveyed Approaches

Our survey aims to bring together insights into how MWE meanings are represented in transformer models. We mainly draw on intrinsic evaluations targeting MWEs, conducted both on pretrained and optimized transformer models. We also consider work on downstream tasks, but only if it has clear implications for model behavior. We purposefully include papers using a variety of models, datasets, and evaluation strategies, whose results may not be directly comparable. However, broad coverage enables us to highlight different perspectives and remaining gaps in the literature.

We analyze the surveyed papers with respect to three lines of questions. (i) In general terms, can transformer representations capture MWE semantics, can they be optimized to more robustly represent these meanings, and can they generalize to unseen expressions? (ii) In terms of localization, which layers, modeled tokens, and contextual elements carry relevant representational information? (iii) Do any linguistic properties of MWEs affect the quality of their representations?

Models. The surveyed papers predominantly use transformer models with an encoder-only architecture. These include BERT and its multilingual version mBERT, trained on the masked language modeling (MLM) and next sentence prediction tasks (Devlin et al., 2019); RoBERTa, which introduced an optimized training procedure (Liu et al., 2019); its multilingual version XLM-R (Conneau et al., 2020); computationally efficient derivatives ALBERT (Lan et al., 2020) and DistilBERT (Sanh et al., 2019); and SBERT, optimized for sentence representations (Reimers and Gurevych, 2019). Further papers use encoder-decoder architectures such as DeBERTa (He et al., 2021), BART (Lewis et al., 2020), and T5 (Raffel et al., 2020); and autoregressive models, in particular XLNet (Yang et al., 2019) and GPT (Radford et al., 2019; Brown et al., 2020).

Phrases	Similarity to other phrase		Attributes
direct link	0.328	access service	IMMEDIACY
formal education	0.508	school board	FORMALITY
common man	0.672	average person	COMMONNESS
Compounds	Compos.	Synonyms	Relations
fairy tale	1.9 ± 1.3	fable	ABOUT
insurance policy	4.4 ± 0.9	insurance plan	ABOUT
birth rate	4.7 ± 0.5	fertility rate	HAVE
Idioms	Literal occurrence		Idiomatic occurrence
in light of	in the light of a bedside lamp		in the light of this success
on the cards	read whatever is written on the card		a decisive victory was on the cards
open one’s eyes	I opened my eyes and looked up		it opened my eyes to the plight

Table 1: Example MWEs with common gold standard information, sampled from datasets on phrase similarity (Asaadi et al., 2019) and attributes (Hartung, 2015); compound compositionality and synonyms (Cordeiro et al., 2019) as well as semantic relations (Ó Séaghdha and Copestake, 2007); and idiomaticity (Haagsma et al., 2020).

Tasks and Datasets. The notion of MWE subsumes varied linguistic phenomena and related tasks. We briefly review the tasks and datasets that are most often used in the surveyed literature (for more extensive categories and definitions, see, e.g., Baldwin and Kim, 2010; Constant et al., 2017). Example expressions and different types of gold standard information are presented in Table 1.

On the most general level, we look into representations of **phrases**. These are groups of words that function as syntactic units and whose overall meanings are therefore derived from multiple lexical elements. Studies adopting this general focus represent the meanings expressed using specific syntactic patterns, e.g., subject–verb–object structures such as *child–read–book*. Evaluation datasets target phrase similarity and paraphrase or attribute detection (Mitchell and Lapata, 2008; Hartung, 2015; Pavlick et al., 2015; Asaadi et al., 2019; Strakatova et al., 2020; Pham et al., 2023).

On a more specific level, we examine **noun compounds** (e.g., *gold mine*). Syntactically, these are also phrases—with at least one modifier and a nominal head—and they are analyzed with a focus on semantic idiosyncrasy. Tasks include predicting the degree of compositionality, i.e., the semantic relatedness of the constituents to the overall meaning; and predicting the meaning of the compound and evaluating it by detecting synonyms, paraphrases, or semantic relations. These tasks rely on a wide range of datasets (Biemann and

Giesbrecht, 2011; Reddy et al., 2011; Hendrickx et al., 2013; Juhasz et al., 2015; Levin et al., 2019; Cordeiro et al., 2019; Pinter et al., 2020a); for a recent analysis, see Schulte im Walde (forthc.).

Idioms are structurally diverse phrases with conventionalized meanings which cannot be deduced from their constituents (e.g., *spill the beans*). Their literal vs. idiomatic interpretation often depends on context, so they are also referred to as potentially idiomatic expressions (PIEs). A standard evaluation is idiomaticity classification on the sentence or token level (Cook et al., 2008; Hashimoto and Kawahara, 2008; Savary et al., 2017; Aharodnik et al., 2018; Moussallem et al., 2018; Haagsma et al., 2020; Saxena and Paul, 2020).

Some papers define their MWE tasks as figurative language or metaphoricity detection. Since figurative language involves non-compositionality and contextual specification, this is largely a matter of perspective. These studies also examine phrases and idioms from the cited or ad-hoc datasets, using tasks such as idiomaticity detection (e.g., is a phrase such as *political storm* used idiomatically?) and plausibility classification (e.g., given a text ending with an idiom, is a candidate continuation plausible?).

3 General Model Properties

In this section, we first assess if MWE meanings are inherently captured by off-the-shelf

transformer models, and if these can be further optimized so as to improve MWE representations. We then look into general mechanisms that may support this ability: recall of memorized information and generalization to unseen data.

3.1 Off-the-shelf Representations

We begin by asking if pretrained transformer models capture MWE meanings without optimization for this type of expression. We examine the extent of this ability based on tasks targeting different compositionality ranges and objectives: predicting an expression’s meaning or its semantic properties, e.g., the degree of compositionality.

A model that encodes MWE semantics should fulfill the necessary (but not sufficient) requirement of **representing compositional phrase meaning**, i.e., the overall meaning that is derived from the meanings of the constituents; this would indicate that the model can capture semantics beyond the level of individual tokens. This issue has been evaluated on the task of predicting phrase similarity. Focusing on English phrases, Gamallo et al. (2021) show that the cosine similarity between phrase-level SBERT embeddings—corresponding to averaging over tokens—is positively correlated with human similarity ratings, reaching $\rho = 0.61$ for noun-verb-noun expressions. In a subsequent study on Galician, their best method uses the contextualized embedding of the verb ($\rho = 0.57$; Gamallo et al., 2022). These studies show that encoding a phrase’s constituent words via transformer architectures can produce meaningful representations of the whole phrase. This in turn indicates that attention-based contextualization effectively distributes (some part of) phrase-level meaning over the constituent tokens.

This tendency is further confirmed by the fact that phrase-level representations can be reconstructed from the representations of their constituents. Even with a straightforward strategy such as vector addition, the mean cosine between the original and reconstructed CLS embeddings reaches 0.92 for BERT and > 0.99 for RoBERTa and DeBERTa (Liu and Neubig, 2022). But successful reconstruction of phrase representations does not entail that they are underpinned by refined compositional processing. If that was the case, they would not be affected by surface factors such as word overlap. Yu and Ettinger (2020)

consider the special case of phrases with inverted constituents (e.g., *law school* and *school law*), corresponding to 12% of phrase pairs they use. The correlation between model-derived phrase similarity and human ratings drops from ≈ 0.6 on the full dataset to ≈ 0.2 on the inverted constituent subset, and indicates a strong effect of word overlap.

MWEs exhibit **variable degrees of compositionality**, which should be reflected by appropriate semantic representations. This has been investigated by predicting the compositionality of noun compounds using features extracted from pretrained models. Nandakumar et al. (2019) report positive Pearson’s correlation with human compositionality ratings, ranging from $r = 0.15$ to 0.60 depending on the dataset and estimation strategy. However, the best BERT results are systematically ≈ 0.2 points behind the strongest methods based on static word embeddings. In a more extensive evaluation, Garcia et al. (2021a) reach $\rho = 0.37$ on English and 0.26 on Portuguese data, similarly lagging behind the SOTA on static word embeddings (0.73 and 0.60, respectively; Cordeiro et al., 2019). These results question BERT’s ability to capture compositionality similarly to humans, but they may be due to a suboptimal use of the information encoded in models which—unlike static word embeddings used on this task—do not learn dedicated MWE representations.

Subsequent work by Miletic and Schulte im Walde (2023) has shown that robust compositionality information *can* be extracted from BERT, but it is not equally accessible across the model architecture. The best results (reaching $\rho = 0.71$) are obtained using embeddings from early layers and comparisons between compounds and their contexts. More generally, Shwartz and Dagan (2019) have examined the potential of transformer representations to capture degrees of compositionality across a variety of tasks. In their supervised classification setup, contextualized representations systematically outperform static word embeddings. However, even implementations that are overall good at distinguishing degrees of compositionality struggle when exposed to more complex semantic mechanisms such as implicit meaning. Accuracy is in the 80 – 90% range on literality-related tasks, but it drops by $\approx 30\%$ points on noun compound relations and adjective–noun attributes.

Another key issue is whether **non-compositional MWE meanings** are represented as such, i.e., more similarly to an independent linguistic unit than a sum of component parts. One line of evidence questioning this ability comes from patterns of similarity between non-compositional expressions. Zeng and Bhat (2022) extract mean-pooled idiom embeddings from BART and find that they cluster together based on surface or syntactic similarity rather than figurative meaning. Garcia et al. (2021b) compare contextualized embeddings of compounds and their synonyms. They assume that their similarity should *not* correlate with compositionality ratings if compound meanings are represented well across compositionality ranges. However, they find moderate-to-strong correlations across models for both English and Portuguese. This indicates that non-compositional compounds are further away from their synonyms in the vector space and to that extent are represented less well than compositional compounds.

A more nuanced picture emerges from attention flows in neural machine translation, as examined by Dankers et al. (2022) on PIEs. In figurative contexts, PIEs exhibit increased self-attention within the expression and reduced interaction with the surrounding context. This suggests they are grouped together more strongly, i.e., processed similarly to a standalone linguistic unit. On the decoder side, this is echoed by lower cross-attention between figurative translations and source PIEs. However, when encoded information is progressively removed through amnesic probing, the model reverts to compositional translations. This brittleness highlights the challenging nature of figurative translations.

Summary. Moderately strong results across tasks of different complexity indicate that pre-trained models capture MWE semantics, but do so inconsistently. This is further shown by their reliance on surface patterns such as word overlap, strong localization of relevant information, and comparatively lower quality of non-compositional meaning representations.

3.2 Optimized Representations

The shortcomings of MWE representations raised in the previous section may be addressed using different approaches. We discuss span representa-

tions, which are optimized to capture the meaning distributed over multiple tokens; task-specific fine-tuning or adapter-tuning, with training strategies that target properties typical of MWE semantics; enhancing models with linguistic knowledge, such as explicit information on potential interpretations of MWEs; and training dedicated neural architectures, which rely on greater model complexity to improve MWE representations.

General-purpose models that are optimized to **represent spans of text** should better capture the meaning of multiple tokens and, by extension, MWEs. SBERT produces sentence-level embeddings following fine-tuning on NLI data with a siamese architecture (Reimers and Gurevych, 2019). SpanBERT is pretrained by masking contiguous spans of text instead of individual tokens, and uses span boundary representations to encode span content (Joshi et al., 2020). PhraseBERT fine-tunes BERT using contrastive learning over positive and negative examples of paraphrases and of contexts (Wang et al., 2021). Among these, SpanBERT obtains the best results on in-context phrase similarity; across evaluations setups, its strongest improvement over BERT is 2.2 accuracy points (Pham et al., 2023). For type-level phrase similarity, better performance can be obtained by aggregating the similarities of multiple pairs of occurrences at inference time. The improvement over BERT stands at 8.6 to 28.7 accuracy points depending on the dataset (Cohen et al., 2022).

Turning to representations more specifically targeting MWEs, different **fine-tuning approaches** have been proposed. Following the findings of Yu and Ettinger (2020) regarding strong effects of surface patterns on phrase representation, Yu and Ettinger (2021) fine-tune models to avoid this effect, for example by predicting if two sentences with high lexical overlap are paraphrases or not. This only leads to minor localized improvements of phrase representations which do not reduce the reliance on word overlap. However, fine-tuning has a stronger effect in other setups. Liu et al. (2022) evaluate figurative language interpretation using a Winograd-style task targeting novel metaphors. Their strongest model is RoBERTa fine-tuned with a contrastive objective, reaching 90.3% accuracy, within 5 points of human performance. This is an improvement of 24.1 points on the zero-shot setup.

A computationally leaner approach consists in **learning an adapter**, as shown by Zeng and Bhat

(2022) on BART idiom embeddings. They evaluate different adapters, with learning objectives that include reconstructing corrupted idiomatic sentences and increasing the similarity between the embeddings of idioms and their dictionary definitions. They obtain clear improvements across evaluations, e.g., accuracy on idiom span detection increases from 50.8 to 76.3. Fine-tuning the full model performs similarly to the directly comparable adapter; it is outperformed by the best adapter variants, trained on additional objectives. An acknowledged limitation is difficulty in generalizing to unseen idioms, stemming from the use of external linguistic knowledge as a supervision signal.

Fine-tuned models can be somewhat further improved using **external linguistic knowledge**. Chakrabarty et al. (2022) evaluate models on a binary classification task targeting plausible continuations of narrative texts whose last sentence contains an idiom or a simile. Their strongest zero-shot approach obtains a relatively high accuracy of 67.7% on unseen idioms. It is strongly outperformed by a model with task-specific fine-tuning (82.0%), but a further improvement of 1.5 accuracy points is obtained by providing knowledge of the literal meaning of idiom constituents.

Targeted improvements have also been obtained using **more complex dedicated architectures**. Zeng and Bhat (2021) approach idiomaticity detection by using the attention flow mechanism (Seo et al., 2017) to fuse BERT-derived representations with static word, character, and POS embeddings. On sentence-level idiomaticity classification across multiple datasets, this method performs similarly or worse than a standard implementation with a linear layer on top of BERT. However, on a stricter accuracy measure—where each token in a sequence is required to be accurately classified for idiomaticity—it outperforms the standard approach by a margin of 20 to 30 points. Focusing on Chinese idiom recommendation in a cloze task, Tan and Jiang (2020) use the MASK embedding to retrieve the correct idiom with a 79.8% accuracy, whereas their “dual embedding” approach—capturing both the immediate context and the broader textual passage—leads to an improvement of 2.6 points. Tan et al. (2021) subsequently propose a dedicated BERT model pretrained on the MLM task by only masking idioms, and then

fine-tuned for multiple-choice recommendation. It reaches 86.3% accuracy, within a point of human performance. These results overall indicate that dedicated architectures can achieve excellent performance on some tasks.

Summary. Different strategies improve MWE representations, with gains over pretrained models varying from marginal to dramatic. The viability of these methods should be carefully weighed against the expected improvements, especially for computationally expensive systems. However, this requirement remains difficult to fulfill: Optimization strategies differ widely in terms of generality (targeting any sequence of tokens vs. a specific type of MWEs), evaluation complexity, generalizability to unseen data, and underlying architectures. More comprehensive evaluations enabling direct comparisons are a priority for future work.

3.3 Memorization and Generalization

Building on our earlier finding that transformer models encode some knowledge of MWEs, we now look into general mechanisms that enable it. We examine the reliance on memorized information and the complementary generalization ability.

Transformer models seem to process MWEs largely based on the **recall of memorized expressions** rather than a sophisticated meaning processing mechanism. When interpreting novel compounds, GPT-3 can provide human-like explanations but it seems to draw on memorized token distributions rather than reason about the underlying conceptual categories (Li et al., 2022). Noun compound paraphrases generated by GPT-3 substantially overlap with web content which likely constitutes its training data; this trend is stronger for existing than for novel compounds. The acceptability of generated paraphrases is lower for novel compounds, which may partly reflect the lack of memorized information (Coil and Shwartz, 2023).

Using the task of predicting the final token of an idiom given its preceding tokens, Haviv et al. (2023) report that GPT-2 has memorized 45%–48%, and BERT 28%–38% of expressions from their set of ≈ 800 items, with higher scores in larger model variants. They further show that memorized information is retrieved in two distinct stages: (i) early layers promote a decrease in the rank of the target token, bringing it closer to the top of the candidate token set; (ii) later layers

promote an increase in its probability. Memorized idioms undergo a slower first stage, i.e., target completions reach the top of the distribution in comparatively later layers, potentially due to the processing of the full input and not only the local context. They also exhibit a more pronounced second stage, with final probabilities around three times higher compared to non-memorized idioms; this is consistent with a smaller set of plausible completions. These findings have important implications for methods that represent MWEs using representations from a specific layer, as the optimal choice may depend on the degree of memorization of the target expression (i.e., memorized expressions may be better represented in comparatively later layers; for further discussion of layers, see §4.1).

Reliance on memorized information is desirable in some settings—for example, when generating highly conventionalized expressions such as idioms—but it may hinder **generalization ability** on other tasks. Falk et al. (2021) evaluate BERT on attribute selection for German adjective–noun phrases, framed as multiclass classification (e.g., *schlauer Junge* ‘smart boy’ has the attribute label *intelligence*). Compared to evaluation on unseen data, performance is stronger when train and validation/test sets have a partial lexical overlap, i.e., the same set of heads or of modifiers. Depending on the dataset variant, this can lead to an improvement of up to 0.26 F1 with modifier overlap. The stronger effect for modifiers (here, adjectives) is consistent with their central role in this task.

These results are echoed by a more general trend that model performance tends to follow: seen data \gg unseen data $>$ cross-lingual data. Fakharian and Cook (2021) evaluate a range of models on PIE idiomaticity classification in English and Russian. The general trend is illustrated by mBERT with task-specific fine-tuning on English. It achieves 83.8% accuracy on seen English data, 74.3% on unseen English data, and 72.4% on Russian data. A similar drop is observed across transformer models, but they remain clearly above baselines, indicating a non-negligible ability to generalize. The same task is investigated for Slovene by Škvorc et al. (2022). They compare mBERT, pretrained on 104 languages, including Slovene; and CroSloEngual-BERT, pretrained only on Croatian, Slovene and English. Looking at sentence-level classification, mBERT is weaker

on seen idioms (0.91 vs. 0.95 F1) but better on unseen idioms (0.90 vs. 0.84). This suggests that it is stronger at generalizing—perhaps due to pretraining on multiple languages related to Slovene—as further confirmed by above-chance cross-lingual performance on Croatian (0.90) and Polish (0.70).

Models can also be optimized for generalization. From a dataset perspective, a BERT-based idiomaticity classifier reaches generalizability faster if the idioms to which it is exposed during training are ordered by decreasing contribution to model performance. The contribution is determined by an idiom’s Shapley value, estimated as the difference between the average performance of multiple models which do vs. do not include a given idiom in training data (Nedumpozhimana et al., 2022). From an architecture perspective, the previously discussed use of attention flow to fuse contextualized and static idiom representations is especially beneficial for generalization to unseen idioms and to other domains (Zeng and Bhat, 2021). This finding—contrasted by competitive performance of standard architectures on seen data—once again indicates that more complex systems are particularly useful in challenging classification scenarios.

Summary. Transformer-based MWE representations strongly rely on memorized information, as observed when generating subparts or paraphrases of target expressions. Novel (non-memorized) expressions yield lower-quality generations and are processed in earlier layers, indicating dominance of the local context. Models generalize to unseen and cross-lingual data with a performance drop in tasks of variable semantic complexity; this can be alleviated by targeted optimization.

4 Impact of Representational Information

Shifting the focus from general insights into MWE semantics captured by different transformer models, we now adopt a finer-grained perspective and examine how MWE representations are impacted by structural factors, i.e., model and input properties that directly affect the representational information extracted from a given transformer architecture. We address three such factors: transformer layers, tokens within the sequence, and the context surrounding the target expression.

4.1 Layers

As previously noted (§2.1), representations from different transformer layers do not capture the same range of linguistic information. We explore the effect that this has on MWE representations by reviewing standard layer choices, their variable effects on performance, interactions with model and linguistic properties, and potential explanations.

When **selecting the layers** to represent MWE meanings, a common choice is the last layer, both as input to a classifier (Nedumpozhimana and Kelleher, 2021; Nedumpozhimana et al., 2022) and as a standalone representation, often constituting a baseline for an optimized model (Wang et al., 2021; Pham et al., 2023). Representations pooled over the last four layers have also been used with variable degrees of success (Gamallo et al., 2021; Garcia et al., 2021a). Other methods learn a scalar mix of layers (Falk et al., 2021); one study has found that a balanced mix of top and bottom layers tends to outperform the individual use of the last layer across tasks (Shwartz and Dagan, 2019).

The impact of layer choice has been assessed on a range of tasks. Most surveyed papers report **better performance in lower layers**; recall that these are the least contextualized representations, assumed to capture surface linguistic features. Brglez (2023) evaluates metaphoricity prediction on 24 Slovene noun phrases, with the best results in the input embedding layer 0. The cosine similarity between constituents is initially higher in literal than metaphorical examples, as expected, but this difference diminishes over the layers. Miletić and Schulte im Walde (2023) predict the degrees of compositionality of 280 English noun compounds, and similarly find the most consistent performance in the low-to-mid range of layers, with the single best result on layer 1. They experiment with pooling contiguous layers, but this penalizes performance. On PIE idiomaticity classification, Tan and Jiang (2021) find that performance stabilizes around layer 4 and generally peaks in the mid-range, indicating that several rounds of contextualization are sufficient for this task. Predicting compositionality or idiomaticity may come down to identifying discrepancies between the target expression and its context, which would support the preference for the less contextualized, lower layers. But similar results have been reported on tasks requiring comparisons of

multiple target expressions. For instance, Burdick et al. (2022) evaluate paraphrase similarity on over 25k phrase pairs and obtain the best individual result with layer 1.

Contrasting this trend, better performance in higher layers has occasionally been reported when predicting PIE idiomaticity (Fakharian and Cook, 2021) and the semantic transparency of closed compounds (Buijtelaar and Pezzelle, 2023). These differences may be explained by the fact that information encoded by different layers is affected by **interactions with other parameters**, such as the choice of the pretrained model architecture and of the target token (e.g., modeled in isolation or in sentence context). Detailed evidence of this trend comes from evaluations of phrase similarity by Yu and Ettinger (2020). They experiment with modeling the target expressions without additional context, and observe the best performance in earlier layers for RoBERTa, XLM-R, and XLNet; middle layers in BERT; and later layers in DistilBERT. By contrast, when sentence context is included, layers in the mid-range are generally strongest for all models. Moreover, the CLS embedding improves in performance as layers progress, pointing to distinct processing of the information it captures.

Layer-level information is also affected by the **linguistic properties** of the modeled expressions. Focusing on the task of paraphrase identification, Tan and Jiang (2021) report an effect of the degree of idiomaticity: When both the target expression and the paraphrase are non-idiomatic, performance is strongest at layer 0 and decreases afterwards; when the target expression is idiomatic (and the paraphrase is either idiomatic or not), performance is relatively stable across the layers. Similarly, Burdick et al. (2022) estimate paraphrase similarity. Within a pair of paraphrases used in the same context, the same words become less similar, and different words more similar, as layers progress; this confirms that later layers capture more contextual information. These findings are echoed by the processing of PIEs in machine translation, examined by Dankers et al. (2022). As layers progress, figuratively used PIEs become less similar to their representations in the preceding layer, compared to their literal counterparts, suggesting a stronger effect of contextualization.

Potential explanations for these trends are provided by the models' **structural features**. Aoyama and Schneider (2022) show that different

types of MWE information do not follow the same distribution over layers. When predicting a MWE token, the model tends to rely on lower layers compared to all tokens, perhaps due to a smaller set of potential candidates. When predicting POS tags, it tends to rely on higher layers compared to all tokens, which may be related to the usefulness of semantic information in resolving POS sequences typical of MWEs. Espinosa Anke et al. (2021) note the effects of anisotropy in BERT, i.e., the tendency for embeddings to concentrate in a narrow cone. On collocate categorization, this leads to overlaps between antonymic collocates—which should ideally be distant in the vector space—with the authors questioning whether the model has inherent knowledge to resolve their task. Klubička et al. (2023) show that, within a given layer, idiomaticity is mostly encoded in vector dimensions rather than the norm, and is somewhat more accessible in the first half of the dimensions. This confirms that linguistic information of interest is not equally distributed over an embedding.

Summary. Later transformer layers are often used to represent MWEs, but lower layers are generally better when predicting both an expression’s meaning and properties such as compositionality. This suggests that MWE semantics are best captured by weakly to moderately contextualized representations, highlighting in turn the relevance of type-level lexical information. This trend is affected by model properties as well as key linguistic features, e.g., figurative and idiomatic expressions benefit from stronger contextualization. But this mirrors the patterns noted for memorized expressions (§3.3); future work should therefore analyze interactions between idiomaticity and memorization. More immediately, the observed patterns indicate that layer choice should be carefully tuned.

4.2 Tokens

After inputting a MWE into a transformer model, embeddings of multiple tokens of interest may be used to represent it. We first address the standard choices and their effects on performance before zooming into attention-based contextualization. We then discuss two implementation issues: the CLS token and subword fragmentation.²

²In what follows, *constituent* denotes an expression’s constituent word without implication for syntactic properties.

When **selecting the modeled token**, frequent choices include the embedding of a constituent, modeled as part of the full expression and thereby contextualized relative to the other constituents (Brglez, 2023); a phrase representation obtained by pooling the constituent embeddings (Pham et al., 2023); the embedding of the MASK token, replacing the target expression in a sequence (Tan and Jiang, 2020); and the embedding of the CLS token, corresponding to the sequence containing the target expression (Fakharian and Cook, 2021). Unsurprisingly, embeddings of different tokens do not capture the same information; they may in fact be complementary. On English and Japanese idiom token classification, Takahashi et al. (2022) gain ≈ 0.025 accuracy points by concatenating contextualized, out-of-context, and MASK representations of constituents; this is opposed to using the contextualized embeddings of constituents.

The **effect of token choice** has been examined on several tasks. On phrase similarity, Yu and Ettinger (2020) show that phrase representations averaged over the constituents obtain better results than alternatives, such as the embedding of the head or of the full sequence. Using a similar approach and task on English, Gamallo et al. (2021) report better results with phrase representations; on Galician, Gamallo et al. (2022) obtain a slight improvement with constituent embeddings (ρ increase of 0.03 compared to phrase embeddings). This may be explained by language-specific patterns or by implementation differences (phrase embeddings obtained using English SBERT vs. simple mean pooling from Galician BERT variants). As for compound compositionality prediction, Miletić and Schulte im Walde (2023) obtain the best results by comparing the embedding of the constituent of interest—the entire compound to predict compound-level compositionality, and the head or modifier to predict their respective contributions to compound meaning—with a pooled embedding of the surrounding sentence context.

Even where post-hoc pooling is required to represent a target item, attention-based **contextualization has beneficial effects**. On the task of adjective attribute classification, Falk et al. (2021) find that phrase embeddings generally outperform constituent embeddings. However, the modifier (i.e., adjective) embeddings are at most slightly behind, indicating that they carry most task-relevant information which is further distributed

through contextualization. In order to predict compound compositionality, Garcia et al. (2021a) compare compound embeddings obtained within a sentence and out-of-context. For the out-of-context setting, they apply pooling over representations of constituents obtained by feeding the model (i) with the entire compound; (ii) with each constituent individually. They obtain better results with (i) ($\rho = 0.37$ vs. 0.16), showing that contextualization via self-attention provides a stronger contribution than a simple composition operation.

In addition to performing well on specific tasks, **pooled representations compared to standalone embeddings** capture some meaningful linguistic information. Nandakumar et al. (2019) predict compound compositionality by comparing constituent and compound BERT embeddings, obtaining a positive correlation with human ratings (up to $r = 0.38$). Garcia et al. (2021b) likewise assess the similarity of a mean-pooled compound representation and that of only one constituent. They obtain generally high cosine scores (≈ 0.8), showing that the representations are closely similar but not identical; and a weak to moderate correlation with compositionality ratings (up to $\rho = 0.45$).

In terms of specific tokens, **the CLS token encodes clearly distinct information** relative to tokens corresponding to MWE constituents. For example, its use on compound compositionality prediction leads to stark drops in performance compared to other tokens of interest (Miletić and Schulte im Walde, 2023). However, this trend may be partly model-specific. On phrase similarity, CLS generally performs poorly except for DistilBERT, where it appears to encode a compositionality signal (Yu and Ettinger, 2020). When reconstructing a phrase representation from its constituents, averaging over all tokens outperforms CLS for BERT, RoBERTa, and DeBERTa—but GPT-2 performs better using the (roughly equivalent) sequence-final token (Liu and Neubig, 2022).

A connected issue is **subword fragmentation**: If a word is not present in a model’s vocabulary, it is tokenized into smaller fragments for which representations exist. Standard solutions include averaging over the subword tokens (e.g., Garcia et al., 2021a) or using only the first (e.g., Gamallo et al., 2021). These solutions are not detrimental when subword fragmentation affects a small subset of target items (Miletić and Schulte im Walde, 2023), but it can be widespread for specific struc-

tures. Focusing on English closed compounds, Pinter et al. (2020b) compare representations obtained by pooling subword-fragmented BERT embeddings vs. those that are first pre-tokenized into gold-standard constituents. Similarity between the two types of pooled representations for a given compound is high overall (cosine reaching ≈ 0.8 – 0.9) but is affected by additional factors: It increases over layers, peaking at layer 11 of 12; and it is stronger for more semantically transparent items. Put differently, subword-fragmented and linguistically motivated representations of constituents recover similar compound-level information, with benefits from attention-based contextualization in both cases. Jenkins et al. (2023) analyze German (closed) compounds and find that pre-tokenization into constituents is beneficial for some evaluations, highlighting the relevance of the target task.

Summary. Multiple lines of evidence converge to indicate that MWEs are best represented by tokens corresponding to linguistic structures of interest, contextualized within the expression and pooled where necessary. This is facilitated by self-attention, which distributes linguistic information over tokens. The CLS token requires cautious implementation due to idiosyncrasies. Subword fragmentation is generally not detrimental.

4.3 Contextual Information

Transformer models can represent sequences of variable length, so we now examine the consequences of modeling MWEs in isolation and in sentence context. We show the benefits of broader context and variants of this information, and then look at how it interacts with model mechanisms.

There is a clear consensus that **contextual information is beneficial** for modeling MWEs. Increasing the amount of linguistic context—including any (rather than none) or including more (rather than some)—improves performance on tasks including phrase similarity estimation (Cohen et al., 2022), idiom translation (Baziotis et al., 2023), metaphoricity prediction on noun-verb phrases (Brglez, 2023), and compositionality prediction on open (Miletić and Schulte im Walde, 2023) and closed compounds (Buijtelaar and Pezzelle, 2023). Contextual information is at the core of some approaches, e.g., idiomaticity detection assuming semantic compatibility between literal MWEs and their context (Zeng and Bhat, 2021). Evidence disputing the usefulness of

context is mostly limited to improvements that are strong overall but absent in a subset of settings, e.g., on phrase similarity (Pham et al., 2023).

Different **variants of contextual information** have been proposed. Representations of phrases—including syntactic structures typical of compounds (noun and adjective phrases) and idioms (verb phrases)—can be obtained through contrastive fine-tuning on paraphrases and further improved by extending the procedure to phrase contexts, i.e., fine-tuning on entire sentences in which phrases appear; accuracy gains reach 8.9 points on longer sequences (Wang et al., 2021). Chinese idiom prediction is similarly improved by including paragraph-level context in addition to the target sentence (Tan and Jiang, 2020). Compound compositionality prediction strongly benefits from modeling paraphrases in addition to compound occurrences, with ρ increasing by 0.5 points compared to only using the targets and their constituents (Nandakumar et al., 2019). More generally, increasing the number of modeled instances per expression leads to an increase in performance; it levels off after ≈ 100 examples on phrase similarity (Cohen et al., 2022). On compound compositionality prediction, the improvement is strongest when shifting from 10 to 100 examples, and minor with a further shift to 1,000 examples (Miletić and Schulte im Walde, 2023).

Linguistic context may **interact with representational properties** such as layers, underlying architectures, and modeled tokens. On phrase similarity, contextual information improves results overall and reduces the impact of individual layers. When only the target expressions are modeled, performance tends to drop as layers progress; when the expressions are in sentence context, it is largely stable (Yu and Ettinger, 2020). Model-specific patterns have been reported on retrieval of idioms vs. compositional phrases. Inclusion of context has no effect on BERT and T5; it reduces surprisal for GPT-2 variants but without altering the patterns relative to other settings (Rambelli et al., 2023). Representations are also affected by word position, shown on pairs of paraphrases in the same context. Same words appearing in different positions are considerably less similar to one another, compared to both same and different words in the same position; this is stronger for larger changes in position (Burdick et al., 2022).

Models draw on **different sources of linguistic information** provided as input. Collocate cat-

egorization improves when MLM predictions—where the target expression is masked in a sentence—are conditioned on the full non-masked sentence by concatenating it (Espinosa Anke et al., 2021). In classification of compound semantic relations and adjective attributes, the strongest results are obtained using the target expression, in sentence context, together with a paraphrase; omitting any of the three elements reduces accuracy by up to 9.4 points (Shwartz and Dagan, 2019). Similarly, probing experiments on idiomaticity classification indicate that BERT relies on information localized mainly in the idiomatic expression itself, but also in the surrounding context (Nedumpozhimana and Kelleher, 2021). This is indirectly echoed by better performance on individual items whose topic distribution is similar to that of the full dataset (Nedumpozhimana et al., 2022).

Summary. A wide array of experimental settings unequivocally show that any increase in contextual information enables better MWE representations. Linguistic context affects the behavior of model structures and is a beneficial source of information on multiple tasks—including those which are not readily reduced to comparisons of target expressions and surrounding context.

5 Impact of Linguistic Properties

MWE representations may be affected by properties of the target expressions themselves. We now provide a breakdown of the reported effects.

Individual expressions vary in terms of their inherent **predictive properties**, as shown in work on the usefulness of individual idioms when training an idiomaticity classifier. Nedumpozhimana et al. (2022) note a positive effect of informativeness, measured by training a classifier on one idiom and evaluating it on the full set of idioms; and ease of prediction, measured by training a classifier on the full set of idioms and evaluating it on one idiom.

Models are affected by the degree of **semantic idiosyncrasy**. Falk et al. (2021) obtain better results for attribute selection on phrases with higher semantic transparency. On idiomaticity detection, Zeng and Bhat (2021) report a slight gain (≈ 0.03 F1) for expressions that are fixed rather than semi-fixed or syntactically flexible. Non-idiomatic expressions are better represented in lower layers (see §4.1; Tan and Jiang, 2021).

As for other semantic properties, lower **polysemy** is associated with better results on attribute selection (Falk et al., 2021) and compositionality prediction (Miletić and Schulte im Walde, 2023), while it does not affect word similarity across paraphrase pairs (Burdick et al., 2022). More **concrete** words are assigned more weight when estimating constituent contributions to compound meaning (Buijtelaar and Pezzelle, 2023). The **type of semantic knowledge** affects metaphor interpretation, with an evaluation on 10k examples showing better results for object and visual commonsense metaphors (referencing common objects and their visual attributes) than for social and cultural commonsense metaphors (referencing human behavior and cultural norms) (Liu et al., 2022).

Varied effects have been noted for **frequency**. They may be partly model-specific, with low-frequency compositional phrases yielding higher surprisal in GPT-2 and T5, but not BERT (Rambelli et al., 2023). Different noun compound analyses have reported that frequency has no effect (Buijtelaar and Pezzelle, 2023), that low frequency is detrimental (Coil and Shwartz, 2023), and that it is beneficial (Miletić and Schulte im Walde, 2023). In the latter case, the trend may be explained by correlation with properties such as **productivity**, with expressions in lower productivity ranges obtaining better representations. More generally, the inconsistencies may stem from the use of different datasets, task formulations, and modeling approaches. This highlights the need for more systematic investigations of frequency effects.

Due to **cross-linguistic variability** in MWE realizations, their optimal computational representations may differ across languages. Yet our ability to assess these trends is limited: Most surveyed work focuses on English, with one to two papers analyzing Chinese, Gallician, German, Japanese, Portuguese, Russian, and Slovenian. We have not identified language-specific patterns for comparable experiments conducted on different languages.

Beyond our survey perspective, direct evidence of cross-linguistic variability is provided by limited studies on two languages in parallel. Some of these report broadly comparable cross-linguistic patterns, e.g., on probing for compound semantics in English and Portuguese (Garcia et al., 2021b) and on idiom token classification in English and Japanese (Takahashi et al., 2022). Varied cross-

linguistic differences have been observed elsewhere. Comparable phrase similarity experiments have obtained the best results using sentence embeddings for English and verb embeddings for Galician (Gamallo et al., 2021, 2022). On compound compositionality prediction, the best model for English is (monolingual) BERT, and for Portuguese (multilingual) SBERT (Garcia et al., 2021a). A PIE identification experiment has found better monolingual generalizability for English than Russian, and better cross-lingual transfer from Russian to English than vice-versa (Fakharian and Cook, 2021). But each of these cases is a single point of reference, making it unclear if the trends are due to model or dataset properties rather than cross-linguistic differences.

Summary. Beyond inherently more informative MWEs, transformer representations are better with lower semantic idiosyncrasy and dispersion (cf. polysemy, productivity). They also appear to be biased towards concrete expressions, while the precise effect of other factors such as frequency remains unclear. Cross-linguistic analyses are limited—both regarding the coverage of different languages and direct comparisons across them—with further work needed to establish reliable trends.

6 Conclusion and Outlook

We have presented a survey of recent work on MWE semantics in transformer-based language models. Starting with a general assessment of pre-trained representations, we have seen that they capture some aspects of MWE meaning, but this ability is neither comprehensive nor consistent. It can in principle be improved with optimization strategies such as fine-tuning and knowledge enhancement, but with highly variable gains. MWE representations rely on memorized information rather than sophisticated meaning processing; this is reflected by suboptimal generalization ability.

Turning to differences in representational information in model architecture and textual input, we find that the most adequate representations are those corresponding to the linguistic structure of interest modeled within broader context; this enables the attention mechanism to efficiently encode expression-level information. There is also broad consensus that lower layers are better at capturing MWE meaning, as observed on tasks such as predicting compound compositionality,

PIE idiomaticity, and paraphrase similarity. However, implementation decisions should always be carefully tuned because of interactions with other factors. This includes a range of linguistic properties, with better representations for expressions exhibiting less semantic idiosyncrasy and dispersion.

The surveyed papers provide varied and valuable insights, but many conclusions are not directly comparable and cannot be extrapolated across MWE types or models. We particularly underscored this issue regarding (i) optimization strategies, which may interact with target expression types, models, and evaluation tasks; (ii) layer-wise processing mechanisms, with similar patterns for memorized and compositional expressions; and (iii) cross-linguistic variability, with insufficient evidence to identify broad trends.

Future studies can address these challenges through several lines of work: (i) Extending the coverage of MWEs beyond the current focus on compounds and idioms, ideally in a comparative setup, and systematically accounting for the effect of their linguistic properties. (ii) Extending the coverage of non-English languages, including in cross-linguistic evaluations. (iii) Broadening the scale of evaluations by multiplying experimental parameters, e.g., by investigating model structures across architectures as well as MWE types. (iv) Formulating tasks that are challenging in terms of both core semantic mechanisms and generalization requirements. We believe that these perspectives will help disentangle interactions between experimental parameters and improve the generalizability of the resulting claims.

Acknowledgments

We thank Michael Roth as well as the action editor and anonymous reviewers for valuable feedback. This research was supported by DFG Research grant SCHU 2580/5-1 (*Computational Models of the Emergence and Diachronic Change of Multi-Word Expression Meanings*).

References

Katsiaryna Aharodnik, Anna Feldman, and Jing Peng. 2018. Designing a Russian idiom-annotated corpus. In *Proceedings of the Eleventh International Conference on*

Language Resources and Evaluation (LREC 2018), Miyazaki, Japan. European Language Resources Association (ELRA).

Tatsuya Aoyama and Nathan Schneider. 2022. Probe-less probing of BERT’s layer-wise linguistic knowledge with masked word prediction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Student Research Workshop*, pages 195–201, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.naacl-srw.25>

Shima Asaadi, Saif Mohammad, and Svetlana Kiritchenko. 2019. Big BiRD: A large, fine-grained, bigram relatedness dataset for examining semantic composition. 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 505–516, Minneapolis, Minnesota. Association for Computational Linguistics. <https://doi.org/10.18653/v1/N19-1050>

Timothy Baldwin and Su Nam Kim. 2010. Multiword expressions. In Nitin Indurkha and Fred J. Damerau, editors, *Handbook of Natural Language Processing*, pages 267–292. CRC Press, Boca Raton, USA.

Christos Baziotis, Prashant Mathur, and Eva Hasler. 2023. Automatic evaluation and analysis of idioms in neural machine translation. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3682–3700, Dubrovnik, Croatia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.eacl-main.267>

Chris Biemann and Eugenie Giesbrecht. 2011. Distributional semantics and compositionality 2011: Shared task description and results. In *Proceedings of the Workshop on Distributional Semantics and Compositionality*, pages 21–28, Portland, Oregon, USA. Association for Computational Linguistics.

Mojca Brglez. 2023. Dispersing the clouds of doubt: Can cosine similarity of word embeddings help identify relation-level metaphors in

- Slovene? In *Proceedings of the 9th Workshop on Slavic Natural Language Processing 2023 (SlavicNLP 2023)*, pages 61–69, Dubrovnik, Croatia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.bsmlp-1.8>
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates.
- Lars Buijtelaar and Sandro Pezzelle. 2023. A psycholinguistic analysis of BERT’s representations of compounds. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2230–2241, Dubrovnik, Croatia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.eacl-main.163>
- Laura Burdick, Jonathan K. Kummerfeld, and Rada Mihalcea. 2022. Using paraphrases to study properties of contextual embeddings. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4558–4568, Seattle, United States. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.naacl-main.338>
- Tuhin Chakrabarty, Yejin Choi, and Vered Shwartz. 2022. It’s not rocket science: Interpreting figurative language in narratives. *Transactions of the Association for Computational Linguistics*, 10:589–606. <https://doi.org/10.1162/tacl.a.00478>
- Andy Coenen, Emily Reif, Ann Yuan, Been Kim, Adam Pearce, Fernanda Viégas, and Martin Wattenberg. 2019. Visualizing and measuring the geometry of BERT. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates.
- Amir Cohen, Hila Gonen, Ori Shapira, Ran Levy, and Yoav Goldberg. 2022. McPhraSy: Multi-context phrase similarity and clustering. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3538–3550, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.findings-emnlp.259>
- Albert Coil and Vered Shwartz. 2023. From chocolate bunny to chocolate crocodile: Do language models understand noun compounds? In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2698–2710, Toronto, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-acl.169>
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.747>
- Mathieu Constant, Gülşen Eryiğit, Johanna Monti, Lonneke van der Plas, Carlos Ramisch, Michael Rosner, and Amalia Todirascu. 2017. Multiword expression processing: A survey. *Computational Linguistics*, 43(4):837–892. <https://doi.org/10.1162/COLIA.00302>
- Paul Cook, Afsaneh Fazly, and Suzanne Stevenson. 2008. The VNC-tokens dataset. In *Proceedings of the LREC Workshop Towards a Shared Task for Multiword Expressions (MWE 2008)*, pages 19–22, Marrakech, Morocco.
- Silvio Cordeiro, Aline Villavicencio, Marco Idiart, and Carlos Ramisch. 2019. Unsupervised compositionality prediction of nominal compounds. *Computational Linguistics*, 45(1):1–57. <https://doi.org/10.1162/colia.00341>

- Verna Dankers, Christopher Lucas, and Ivan Titov. 2022. Can transformer be too compositional? Analysing idiom processing in neural machine translation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3608–3626, Dublin, Ireland. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.acl-long.252>
- 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. <https://doi.org/10.18653/v1/N19-1423>
- Luis Espinosa Anke, Joan Codina-Filba, and Leo Wanner. 2021. Evaluating language models for the retrieval and categorization of lexical collocations. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1406–1417, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.eacl-main.120>
- Kawin Ethayarajh. 2019. How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1006>
- Allyson Ettinger. 2020. What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *Transactions of the Association for Computational Linguistics*, 8:34–48. <https://doi.org/10.1162/tacl.a.00298>
- Samir Fakharian and Paul Cook. 2021. Contextualized embeddings encode monolingual and cross-lingual knowledge of idiomaticity. In *Proceedings of the 17th Workshop on Multilingual Expressions (MWE 2021)*, pages 23–32, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.mwe-1.4>
- Neele Falk, Yana Strakatova, Eva Huber, and Erhard Hinrichs. 2021. Automatic classification of attributes in German adjective-noun phrases. In *Proceedings of the 14th International Conference on Computational Semantics (IWCS)*, pages 239–249, Groningen, The Netherlands (online). Association for Computational Linguistics.
- Pablo Gamallo, Manuel Corral, and Marcos Garcia. 2021. Comparing dependency-based compositional models with contextualized word embeddings. In *Proceedings of the 13th International Conference on Agents and Artificial Intelligence*, pages 1258–1265. SCITEPRESS - Science and Technology Publications. <https://doi.org/10.5220/0010391812581265>
- Pablo Gamallo, Marcos Garcia, and Iria de-Dios-Flores. 2022. Evaluating contextualized vectors from both large language models and compositional strategies. *Procesamiento del Lenguaje Natural*, 153–164. <https://doi.org/10.26342/2022-69-13>
- Marcos Garcia, Tiago Kramer Vieira, Carolina Scarton, Marco Idiart, and Aline Villavicencio. 2021a. Assessing the representations of idiomaticity in vector models with a noun compound dataset labeled at type and token levels. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2730–2741, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.acl-long.212>
- Marcos Garcia, Tiago Kramer Vieira, Carolina Scarton, Marco Idiart, and Aline Villavicencio. 2021b. Probing for idiomaticity in vector space models. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3551–3564, Online. Association for

- Computational Linguistics. <https://doi.org/10.18653/v1/2021.eacl-main.310>
- Hessel Haagsma, Johan Bos, and Malvina Nissim. 2020. MAGPIE: A large corpus of potentially idiomatic expressions. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 279–287, Marseille, France. European Language Resources Association.
- Matthias Hartung. 2015. *Distributional Semantic Models of Attribute Meaning in Adjectives and Nouns*. Ph.D. thesis, Heidelberg University.
- Chikara Hashimoto and Daisuke Kawahara. 2008. Construction of an idiom corpus and its application to idiom identification based on WSD incorporating idiom-specific features. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 992–1001, Honolulu, Hawaii. Association for Computational Linguistics. <https://doi.org/10.3115/1613715.1613844>
- Adi Haviv, Ido Cohen, Jacob Gidron, Roei Schuster, Yoav Goldberg, and Mor Geva. 2023. Understanding transformer memorization recall through idioms. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 248–264, Dubrovnik, Croatia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.eacl-main.19>
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTa: Decoding-enhanced BERT with disentangled attention. In *Proceedings of the Ninth International Conference on Learning Representations (ICLR 2021)*.
- Iris Hendrickx, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Stan Szpakowicz, and Tony Veale. 2013. SemEval-2013 task 4: Free paraphrases of noun compounds. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 138–143, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Ganesh Jawahar, Benoît Sagot, and Djamel 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. <https://doi.org/10.18653/v1/P19-1356>
- Christopher Jenkins, Filip Miletić, and Sabine Schulte im Walde. 2023. To split or not to split: Composing compounds in contextual vector spaces. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 16131–16136, Singapore. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.emnlp-main.1002>
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. SpanBERT: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77. <https://doi.org/10.1162/tacl.a.00300>
- Barbara J. Juhasz, Yun-Hsuan Lai, and Michelle L. Woodcock. 2015. A database of 629 English compound words: Ratings of familiarity, lexeme meaning dominance, semantic transparency, age of acquisition, imageability, and sensory experience. *Behavior Research Methods*, 47(4):1004–1019. <https://doi.org/10.3758/s13428-014-0523-6>, PubMed: 25361864
- Filip Klubička, Vasudevan Nedumpozhimana, and John Kelleher. 2023. Idioms, probing and dangerous things: Towards structural probing for idiomaticity in vector space. In *Proceedings of the 19th Workshop on Multiword Expressions (MWE 2023)*, pages 45–57, Dubrovnik, Croatia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.mwe-1.8>
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A lite BERT for self-supervised learning of language representations. In *Proceedings of the Eighth International Conference on Learning Representations (ICLR 2020)*.
- Beth Levin, Lelia Glass, and Dan Jurafsky. 2019. Systematicity in the semantics of noun

- compounds: The role of artifacts vs. natural kinds. *Linguistics*, 57(3):429–471. <https://doi.org/10.1515/ling-2019-0013>
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.703>
- Siyan Li, Riley Carlson, and Christopher Potts. 2022. Systematicity in GPT-3’s interpretation of novel English noun compounds. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 717–728, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.findings-emnlp.50>
- Emmy Liu, Chenxuan Cui, Kenneth Zheng, and Graham Neubig. 2022. Testing the ability of language models to interpret figurative language. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4437–4452, Seattle, United States. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.naacl-main.330>
- Emmy Liu and Graham Neubig. 2022. Are representations built from the ground up? An empirical examination of local composition in language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9053–9073, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.emnlp-main.617>
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv:1907.11692*. <https://doi.org/10.48550/arXiv.1907.11692>
- Timothee Mickus, Denis Paperno, Mathieu Constant, and Kees van Deemter. 2020. What do you mean, BERT? Assessing BERT as a distributional semantics model. In *Proceedings of the Society for Computation in Linguistics 2020*, pages 279–290, New York, New York. Association for Computational Linguistics.
- Filip Miletić and Sabine Schulte im Walde. 2023. A systematic search for compound semantics in pretrained BERT architectures. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1499–1512, Dubrovnik, Croatia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.eacl-main.110>
- Jeff Mitchell and Mirella Lapata. 2008. Vector-based models of semantic composition. In *Proceedings of ACL-08: HLT*, pages 236–244, Columbus, Ohio. Association for Computational Linguistics.
- Diego Moussallem, Mohamed Ahmed Sherif, Diego Esteves, Marcos Zampieri, and Axel-Cyrille Ngonga Ngomo. 2018. LIdioms: A multilingual linked idioms data set. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Navnita Nandakumar, Timothy Baldwin, and Bahar Salehi. 2019. How well do embedding models capture non-compositionality? A view from multiword expressions. In *Proceedings of the 3rd Workshop on Evaluating Vector Space Representations for NLP*, pages 27–34, Minneapolis, USA. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W19-2004>
- Vasudevan Nedumpozhihana and John Kelleher. 2021. Finding BERT’s idiomatic key. In *Proceedings of the 17th Workshop on Multiword Expressions (MWE 2021)*, pages 57–62, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.mwe-1.7>
- Vasudevan Nedumpozhihana, Filip Klubička, and John D. Kelleher. 2022. Shapley idioms: Analysing BERT sentence embeddings for general idiom token identification. *Frontiers in Artificial Intelligence*, 5:813967. <https://doi.org/10.3389/frai.2022.813967>, PubMed: 35360661

- Diarmuid Ó. Séaghdha and Ann Copestake. 2007. Co-occurrence contexts for noun compound interpretation. In *Proceedings of the Workshop on A Broader Perspective on Multiword Expressions*, pages 57–64, Prague, Czech Republic. Association for Computational Linguistics. <https://doi.org/10.3115/1613704.1613712>
- Ellie Pavlick, Pushpendre Rastogi, Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2015. PPDB 2.0: Better paraphrase ranking, fine-grained entailment relations, word embeddings, and style classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 425–430, Beijing, China. Association for Computational Linguistics. <https://doi.org/10.3115/v1/P15-2070>
- Thang Pham, Seunghyun Yoon, Trung Bui, and Anh Nguyen. 2023. PiC: A phrase-in-context dataset for phrase understanding and semantic search. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1–26, Dubrovnik, Croatia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.eacl-main.1>
- Yuval Pinter, Cassandra L. Jacobs, and Max Bittker. 2020a. NYTWIT: A dataset of novel words in the New York Times. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6509–6515, Barcelona, Spain (Online). International Committee on Computational Linguistics. <https://doi.org/10.18653/v1/2020.coling-main.572>
- Yuval Pinter, Cassandra L. Jacobs, and Jacob Eisenstein. 2020b. Will it unblend? In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1525–1535, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.138>
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(1):5485–5551.
- Giulia Rambelli, Emmanuele Chersoni, Marco S. G. Senaldi, Philippe Blache, and Alessandro Lenci. 2023. Are frequent phrases directly retrieved like idioms? An investigation with self-paced reading and language models. In *Proceedings of the 19th Workshop on Multiword Expressions (MWE 2023)*, pages 87–98, Dubrovnik, Croatia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.mwe-1.13>
- Siva Reddy, Diana McCarthy, and Suresh Manandhar. 2011. An empirical study on compositionality in compound nouns. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 210–218, Chiang Mai, Thailand. Asian Federation of Natural Language Processing.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1410>
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in BERTology: What we know about how BERT works. *Transactions of the Association for Computational Linguistics*, 8:842–866. https://doi.org/10.1162/tacl_a_00349
- Ivan A. Sag, Timothy Baldwin, Francis Bond, Ann Copestake, and Dan Flickinger. 2002. Multiword expressions: A pain in the neck for NLP. In *Computational Linguistics and Intelligent Text Processing*, pages 1–15, Berlin. Springer. https://doi.org/10.1007/3-540-45715-1_1
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: Smaller, faster, cheaper

- and lighter. In *5th Workshop on Energy Efficient Machine Learning and Cognitive Computing – NeurIPS 2019*.
- Agata Savary, Carlos Ramisch, Silvio Cordeiro, Federico Sangati, Veronika Vincze, Behrang QasemiZadeh, Marie Candito, Fabienne Cap, Voula Giouli, Ivelina Stoyanova, and Antoine Doucet. 2017. The PARSEME shared task on automatic identification of verbal multiword expressions. In *Proceedings of the 13th Workshop on Multiword Expressions (MWE 2017)*, pages 31–47, Valencia, Spain. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W17-1704>
- Prateek Saxena and Soma Paul. 2020. EPIE dataset: A corpus for possible idiomatic expressions. In Petr Sojka, Ivan Kopeček, Karel Pala, and Aleš Horák, editors, *Text, Speech, and Dialogue*, volume 12284, pages 87–94, Cham. Springer. https://doi.org/10.1007/978-3-030-58323-1_9
- Sabine Schulte im Walde. forthc. Collecting and investigating features of compositionality ratings. In Voula Giouli and Verginica Barbu Mititelu, editors, *Multiword Expressions in Lexical Resources. Linguistic, Lexicographic and Computational Perspectives*, Phraseology and Multiword Expressions, Language Science Press, Berlin.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *Proceedings of the Fifth International Conference on Learning Representations (ICLR 2017)*.
- Vered Shwartz and Ido Dagan. 2019. Still a pain in the neck: Evaluating text representations on lexical composition. *Transactions of the Association for Computational Linguistics*, 7:403–419. https://doi.org/10.1162/tacl_a_00277
- Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela. 2021. Masked language modeling and the distributional hypothesis: Order word matters pre-training for little. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2888–2913, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.emnlp-main.230>
- Tadej Škvorc, Polona Gantar, and Marko Robnik-Šikonja. 2022. MICE: Mining idioms with contextual embeddings. *Knowledge-Based Systems*, 235:107606. <https://doi.org/10.1016/j.knosys.2021.107606>
- Yana Strakatova, Neele Falk, Isabel Fuhrmann, Erhard Hinrichs, and Daniela Rossmann. 2020. All that glitters is not gold: A gold standard of adjective-noun collocations for German. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4368–4378, Marseille, France. European Language Resources Association.
- Ryosuke Takahashi, Ryohei Sasano, and Koichi Takeda. 2022. Leveraging three types of embeddings from masked language models in idiom token classification. In *Proceedings of the 11th Joint Conference on Lexical and Computational Semantics*, pages 234–239, Seattle, Washington. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.starsem-1.21>
- Minghuan Tan and Jing Jiang. 2020. A BERT-based dual embedding model for Chinese idiom prediction. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1312–1322, Barcelona, Spain (Online). International Committee on Computational Linguistics. <https://doi.org/10.18653/v1/2020.coling-main.113>
- Minghuan Tan and Jing Jiang. 2021. Does BERT understand idioms? A probing-based empirical study of BERT encodings of idioms. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 1397–1407, Held Online. INCOMA Ltd. https://doi.org/10.26615/978-954-452-072-4_156
- Minghuan Tan, Jing Jiang, and Bing Tian Dai. 2021. A BERT-based two-stage model for Chinese chengyu recommendation. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 20(6):1–18. <https://doi.org/10.1145/3453185>
- 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*, volume 30. Curran Associates.
- Ivan Vulić, Edoardo Maria Ponti, Robert Litschko, Goran Glavaš, and Anna Korhonen. 2020. Probing pretrained language models for lexical semantics. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7222–7240, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.586>
- Shufan Wang, Laure Thompson, and Mohit Iyyer. 2021. Phrase-BERT: Improved phrase embeddings from BERT with an application to corpus exploration. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10837–10851, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.emnlp-main.846>
- Gregor Wiedemann, Steffen Remus, Avi Chawla, and Chris Biemann. 2019. Does BERT make any sense? Interpretable word sense disambiguation with contextualized embeddings. In *Proceedings of the 15th Conference on Natural Language Processing (KONVENS 2019): Long Papers*, pages 161–170, Erlangen, Germany. German Society for Computational Linguistics & Language Technology.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R. Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized autoregressive pretraining for language understanding. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates.
- Lang Yu and Allyson Ettinger. 2020. Assessing phrasal representation and composition in Transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4896–4907, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.397>
- Lang Yu and Allyson Ettinger. 2021. On the interplay between fine-tuning and composition in Transformers. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2279–2293, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-acl.201>
- Ziheng Zeng and Suma Bhat. 2021. Idiomatic expression identification using semantic compatibility. *Transactions of the Association for Computational Linguistics*, 9:1546–1562. <https://doi.org/10.1162/tacl.a.00442>
- Ziheng Zeng and Suma Bhat. 2022. Getting BART to ride the idiomatic train: Learning to represent idiomatic expressions. *Transactions of the Association for Computational Linguistics*, 10:1120–1137. <https://doi.org/10.1162/tacl.a.00510>