

# Language in a (Search) Box: Grounding Language Learning in Real-World Human-Machine Interaction

**Federico Bianchi\***  
Bocconi University  
Milano, Italy  
f.bianchi@unibocconi.it

**Ciro Greco**  
Coveo Labs  
New York, USA  
cgreco@coveo.com

**Jacopo Tagliabue**  
Coveo Labs  
New York, USA  
jtagliabue@coveo.com

## Abstract

We investigate grounded language learning through real-world data, by modelling a teacher-learner dynamics through the natural interactions occurring between users and search engines; in particular, we explore the emergence of semantic generalization from unsupervised dense representations outside of synthetic environments. A grounding domain, a denotation function and a composition function are learned from user data only. We show how the resulting semantics for noun phrases exhibits compositional properties while being fully learnable without any explicit labelling. We benchmark our grounded semantics on compositionality and zero-shot inference tasks, and we show that it provides better results and better generalizations than SOTA non-grounded models, such as word2vec and BERT.

## 1 Introduction

Most SOTA models in NLP are only *intra-textual*. Models based on distributional semantics – such as standard and contextual word embeddings (Mikolov et al., 2013; Peters et al., 2018; Devlin et al., 2019) – learn representations of word meaning from patterns of co-occurrence in big corpora, with no reference to extra-linguistic entities.

While successful in a range of cases, this approach does not take into consideration two fundamental facts about language. The first is that language is a referential device used to refer to extra-linguistic objects. Scholarly work in psycholinguistics (Xu and Tenenbaum, 2000), formal semantics (Chierchia and McConnell-Ginet, 2000) and philosophy of language (Quine, 1960) show that (at least some aspects of) linguistic meaning can be represented as a sort of mapping between linguistic and extra-linguistic entities. The second is

that language may be learned based on its usage and that learners draw part of their generalizations from the observation of teachers’ behaviour (Tomasello, 2003). These ideas have been recently explored by work in grounded language learning, showing that allowing artificial agents to access human actions providing information on language meaning has several practical and scientific advantages (Yu et al., 2018; Chevalier-Boisvert et al., 2019).

While most of the work in this area uses toy worlds and synthetic linguistic data, we explore grounded language learning offering an example in which unsupervised learning is combined with a language-independent grounding domain in a real-world scenario. In particular, we propose to use the interaction of users with a search engine as a setting for grounded language learning. In our setting, users produce search queries to find products on the web: queries and clicks on search results are used as a model for the teacher-learner dynamics.

We summarize the contributions of our work as follows:

1. we provide a grounding domain composed of dense representations of extra-linguistic entities constructed in an unsupervised fashion from user data collected in the real world. In particular, we learn neural representations for our domain of objects leveraging prod2vec (Grbovic et al., 2015): crucially, building the grounding domain does not require any linguistic input and it is independently justified in the target domain (Tagliabue et al., 2020a). In this setting, lexical denotation can also be learned without explicit labelling, as we use the natural interactions between the users and the search engine to learn a noisy denotation for the lexicon (Bianchi et al., 2021). More specifically, we use *DeepSets* (Cotter et al., 2018) constructed from user behavioural signals as the extra-linguistic reference of words. For in-

\*Corresponding author. Authors contributed equally and are listed alphabetically.

stance, the denotation of the word “shoes” is constructed from the clicks produced by real users on products that are in fact *shoes* after having performed the query “shoes” in the search bar. Albeit domain specific, the resulting language is significantly richer than languages from agent-based models of language acquisition (Słowik et al., 2020; Fitzgerald and Tagliabue, 2020), as it is based on 26k entities from the inventory of a real website.

2. We show that a dense domain built through unsupervised representations can support *compositionality*. By replacing a discrete formal semantics of noun phrases (Heim and Kratzer, 1998) with functions learned over DeepSets, we test the generalization capability of the model on zero-shot inference: once we have learned the meaning of “Nike shoes”, we can reliably predict the meaning of “Adidas shorts”. In this respect, this work represents a major departure from previous work on the topic, where compositional behavior is achieved through either discrete structures built manually (Lu et al., 2018; Krishna et al., 2016), or embeddings of such structures (Hamilton et al., 2018).
3. To the best of our knowledge, no dataset of this kind (product embeddings from shopping sessions and query-level data) is publicly available. As part of this project, we release our code and a curated dataset, to broaden the scope of what researchers can do on the topic<sup>1</sup>.

Methodologically, our work draws inspiration from research at the intersection between Artificial Intelligence and Cognitive Sciences: as pointed out in recent papers (Bisk et al., 2020; Bender and Koller, 2020), extra-textual elements are crucial in advancing our comprehension of language acquisition and the notion of “meaning”. While synthetic environments are popular ways to replicate child-like abilities (Kosoy et al., 2020; Hill et al., 2020), our work calls the attention on real-world Information Retrieval systems as experimental settings: cooperative systems such as search engines offer new ways to study language grounding, in between the oversimplification of toy models and the

---

<sup>1</sup>Please refer to the project repository for additional information: <https://github.com/coveooss/naacl-2021-grounded-semantics>.

daunting task of providing a general account of the semantics of a natural language. The chosen IR domain is rich enough to provide a wealth of data and possibly to see practical applications, whereas at the same time it is sufficiently self-contained to be realistically mastered without human supervision.

## 2 Methods

Following our informal exposition in Section 1, we distinguish three components, which are learned separately in a sequence: learning a language-independent grounding domain, learning noisy denotation from search logs and finally learning functional composition. While only the first model (prod2vec) is completely unsupervised, it is important to remember that the other learning procedures are only weakly supervised, as the labelling is obtained by exploiting an existing user-machine dynamics to provide noisy labels (i.e. no human labeling was necessary at any stage of the training process).

**Learning a representation space.** We train product representation to provide a “dense ontology” for the (small) world we want our language to describe. Those representations are known in product search as product embeddings (Grbovic et al., 2015): *prod2vec* models are word2vec models in which words in a sentence are replaced by products in a shopping session. For this study, we pick CBOW (Mu et al., 2018) as our training algorithm, and select  $d = 24$  as vector size, optimizing hyperparameters as recommended by Bianchi et al. (2020); similar to what happens with word2vec, related products (e.g. two pairs of sneakers) end up closer in the embedding space. In the overall picture, the product space just constitutes a grounding domain, and re-using tried and tested (Tagliabue et al., 2020b) neural representations is an advantage of the proposed semantics.

**Learning lexical denotation.** We interpret clicks on products in the search result page, after a query is issued, as a noisy “pointing” signal (Tagliabue and Cohn-Gordon, 2019), i.e., a map between text (“shoes”) and the target domain (a portion of the product space). In other words, our approach can be seen as a neural generalization of model-theoretic semantics, where the extension of “shoes” is not a discrete set of objects, but a region in the grounding space. Given a list of products clicked by shoppers after queries, we represent meaning through an order-invariant op-

eration over product embeddings (average pooling weighted by empirical frequencies, similar to Yu et al. (2020)); following Cotter et al. (2018), we refer to this representation as a *DeepSet*. Since words are now grounded in a dense domain, set-theoretic functions for NPs (Chierchia and McConnell-Ginet, 2000) need to be replaced with matrix composition, as we explain in the ensuing section.

**Learning functional composition.** Our functional composition will come from the composition of DeepSet representations, where we want to learn a function  $f : \text{DeepSet} \times \text{DeepSet} \rightarrow \text{DeepSet}$ . We address functional composition by means of two models from the relevant literature (Hartung et al., 2017): one, *Additive Compositional Model (ADM)*, sums vectors together to build the final DeepSet representation. The second model is instead a *Matrix Compositional Model (MDM)*: given in input two DeepSets (for example, one for “Nike” and one for “shoes”) the function we learn as the form  $Mv + Nu$ , where the interaction between the two vectors is mediated through the learning of two matrices,  $M$  and  $N$ . Since the output of these processes is always a DeepSet, both models can be recursively composed, given the form of the function  $f$ .

### 3 Experiments

**Data.** We obtained catalog data, search logs and detailed behavioral data (anonymized product interactions) from a partnering online shop, *Shop X*. *Shop X* is a mid-size Italian website in the sport apparel vertical<sup>2</sup>. Browsing and search data are sampled from one season (to keep the underlying catalog consistent), resulting in a total of 26,057 distinct product embeddings, trained on more than 700,000 anonymous shopping sessions. To prepare the final dataset, we start from comparable literature (Baroni and Zamparelli, 2010) and the analysis of linguistic and browsing behavior in *Shop X*, and finally distill a set of NP queries for our compositional setting.

In particular, we build a rich, but tractable set by excluding queries that are too rare (<5 counts), queries with less than three different products clicked, and queries for which no existing product embedding is present. Afterwards, we zoom into NP-like constructions, by inspecting which features are frequently used in the query log (e.g. shoppers

<sup>2</sup>For convenience of exposition, all queries and examples cited in the paper are translated into English.

search for sport, not colors), and matching logs and NPs to produce the final set. Based on our experience with dozens of successful deployments in the space, NPs constitute the vast majority of queries in product search: thus, even if our intent is mainly theoretical, we highlight that the chosen types overlap significantly with real-world frequencies in the relevant domain. Due to the power-law distribution of queries, one-word queries are the majority of the dataset (60%); to compensate for sparsity we perform data augmentation for rare compositional queries (e.g. “Nike running shoes”): after we send a query to the existing search engine to get a result set, we simulate  $n = 500$  clicks by drawing products from the set with probability proportional to their overall popularity (Bianchi et al., 2021)<sup>3</sup>.

The final dataset consists of 104 “activity + sortal”<sup>4</sup> queries – “running shoes” –; 818 “brand + sortal” queries – “Nike shoes” –, and 47 “gender + sortal” queries – “women shoes”; our testing data consists of 521 “brand + activity + sortal” (BAS) triples, 157 “gender + activity + sortal” (GAS) triples, 406 “brand + gender + activity + sortal” (BGAS) quadruples.<sup>5</sup>

**Tasks and Metrics.** Our evaluation metrics are meant to compare the *real* semantic representation of composed queries (“Nike shoes”) with the one predicted by the tested models: in the case of the proposed semantics, that means evaluating how it predicts the DeepSet representation of “Nike shoes”, given the representation of “shoes” and “Nike”. Comparing target vs predicted representations is achieved by looking at the nearest neighbors of the predicted DeepSet, as intuitively complex queries behave as expected only if the two representations share many neighbors. For this reason, quantitative evaluation is performed using two well-known ranking metrics: *nDCG* and *Jac-*

<sup>3</sup>Since the only objects users can click on are those returned by the search box, query representation may in theory be biased by the idiosyncrasies of the engine. In practice, we confirmed that the embedding quality is stable even when a sophisticated engine is replaced by simple Boolean queries over TF-IDF vectors, suggesting that any bias of this sort is likely to be very small and not important for the quality of the compositional semantics.

<sup>4</sup>“Sortal” refers to a *type* of object: *shoes* and *polo* are sortals, while *black* and *Nike* are not; “activity” is the sport activity for a product, e.g. *tennis* for a racket.

<sup>5</sup>Dataset size for our compositional tests is in line with intra-textual studies on compositionality (Baroni and Zamparelli, 2010; Rubinstein et al., 2015); moreover, the lexical atoms in our study reflect a real-world distribution that is independently generated, and not frequency on general English corpora.

<i>LOBO</i>	$ADM_p$	$MDM_p$	$ADM_v$	$MDM_v$	<b>UM</b>	<b>W2V</b>
nDCG	0.1821	<b>0.2993</b>	0.1635	0.0240	0.0024	0.0098
Jaccard	<u>0.0713</u>	<b>0.1175</b>	0.0450	0.0085	0.0009	0.0052

Table 1: Results on **LOBO** (**bold** are best, underline second best).

<i>ZT</i>	$ADM_p$	$MDM_p$	$ADM_v$	$MDM_v$	<b>UM</b>	<b>W2V</b>
<b>BAS</b> ( <i>brand + activity + sortal</i> )						
nDCG	<u>0.0810</u>	<b>0.0988</b>	0.0600	0.0603	0.0312	0.0064
Jaccard	<u>0.0348</u>	<b>0.0383</b>	0.0203	0.0214	0.0113	0.0023
<b>GAS</b> ( <i>gender + activity + sortal</i> )						
nDCG	<b>0.0221</b>	0.0078	0.0097	0.0160	<u>0.0190</u>	0.0005
Jaccard	<b>0.0083</b>	0.0022	0.0029	<u>0.0056</u>	<u>0.0052</u>	0.0001
<b>BGAS</b> ( <i>brand + gender + activity + sortal</i> )						
nDCG	<u>0.0332</u>	<b>0.0375</b>	0.0118	0.0177	0.0124	0.0059
Jaccard	<u>0.0162</u>	<b>0.0163</b>	0.0042	0.0061	0.0044	0.0019

Table 2: Results on **ZT** (**bold** are best, underline second best).

*card* (Vasile et al., 2016; Jaccard, 1912). We focus on two tasks: leave-one-brand-out (**LOBO**) and zero-shot (**ZT**). In **LOBO**, we train models over the “brand + sortal” queries but we exclude from training a specific brand (e.g., “Nike”); in the test phase, we ask the models to predict the DeepSet for a seen sortal and an unseen brand. For **ZT** we train models over queries with two terms (“brand + sortal”, “activity + sortal” and “gender + sortal”) and see how well our semantics generalizes to compositions like “brand + activity + sortal”; the complex queries that we used at test time are new and unseen.

**Models.** We benchmark our semantics (tagged as  $p$  in the results table) based on **ADM** and **MDM** against three baselines: one is another grounded model, where prod2vec embeddings are replaced by image embeddings (tagged as  $v$  in the results table), to test the representational capabilities of the chosen domain against a well-understood modality – image vectors are extracted with ResNet-18, taking the average pooling of the last layer to obtain 512-dimensional vectors; two are intra-textual models, where word embeddings are obtained from state-of-the-art distributional models, BERT (**UM**) (the Umberto model<sup>6</sup>) and Word2Vec (**W2V**), trained on textual metadata from *Shop X* catalog. For **UM**, we extract the 768 dimensional representation from the [CLS] embedding of the 12th layer of the query and learn a linear projection

<sup>6</sup><https://huggingface.co/Musixmatch/umberto-commoncrawl-cased-v1>

to the product-space (essentially, training to predict the DeepSet representation from text). The generalization to different and longer queries for **UM** comes from the embeddings of the queries themselves. Instead, for **W2V**, we learn a compositional function that concatenates the two input DeepSets, projects them to 24 dimensions, pass them through a Rectified Linear Unit, and finally project them to the product space.<sup>7</sup> We run every model 15 times and report average results; RMSProp is the chosen optimizer, with a batch size of 200, 20% of the training set as validation set and early stopping with *patience* = 10.

**Results.** Table 1 shows the results on **LOBO**, with grounded models outperforming intra-textual ones, and prod2vec semantics (tagged as  $p$ ) beating all baselines. Table 2 reports performance for different complex query types in the zero-shot inference task: grounded models are superior, with the proposed model outperforming baselines across all types of queries.

**MDM** typically outperforms **ADM** as a composition method, except for *GAS*, where all models suffer from *gender* sparsity; in that case, the best model is **ADM**, i.e. the one without an implicit bias from the training. In general, grounded models outperform intra-textual models, often by a wide margin, and prod2vec-based semantics outperforms image-based semantics, proving that the

<sup>7</sup>First results with the same structure as **ADM** and **MDM** showed very low performances, thus we made the architecture more complex and non-linear.



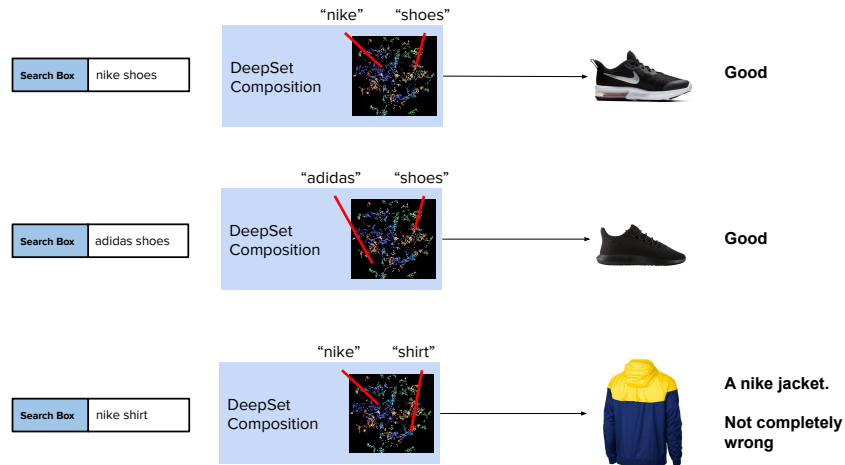


Figure 1: Examples of qualitative predictions made by MDM on the LOBO task.

chosen latent grounding domain supports rich representational capabilities. The quantitative evaluations were confirmed by manually inspecting nearest neighbors for predicted DeepSets in the **LOBO** setting – as an example, **MDM** predicts for “Nike shoes” a DeepSet that has (correctly) all *shoes* as neighbors in the space, while, for the same query, **UM** suggests *shorts* as the answer. Figure 1 shows some examples of compositions obtained by the MDM model on the LOBO task; the last example shows that the model, given in input the query “Nike shirt”, does not reply with a shirt, but with a Nike jacket: even if the correct meaning of “shirt” was not exactly captured in this contest, the model ability to identify a similar item is remarkable.

#### 4 Conclusions and Future Work

In the spirit of [Bisk et al. \(2020\)](#), we argued for grounding linguistic meaning in artificial systems through experience. In our implementation, all the important pieces – domain, denotation, composition – are learned from behavioral data. By grounding meaning in (a representation of) objects and their properties, the proposed noun phrase semantics can be learned “bottom-up” like distributional models, but can generalize to unseen examples, like traditional symbolic models: the implicit, dense structure of the domain (e.g. the relative position in the space of *Nike* products and *shoes*) underpins the explicit, discrete structure of queries picking objects in that domain (e.g. “Nike shoes”) – in other words, compositionality is an emergent phenomenon. While encouraging, our results are still preliminary: first, we plan on extending our semantics, starting with Boolean operators (e.g. “shoes

*NOT Nike*”); second, we plan to improve our representational capabilities, either through symbolic knowledge or more discerning embedding strategies; third, we wish to explore transformer-based architectures ([Lee et al., 2019](#)) as an alternative way to produce set-like representations.

We conceived our work as a testable application of a broader methodological stance, loosely following the agenda of the child-as-hacker ([Rule et al., 2020](#)) and child-as-scientist ([Gopnik, 2012](#)) programs. Our “search-engine-as-a-child” metaphor may encourage the use of abundant real-world search logs to test computational hypotheses about language learning inspired by cognitive sciences ([Carey and Bartlett, 1978](#)).

#### Acknowledgments

We wish to thank Christine Yu, Patrick John Chia and the anonymous reviewers for useful comments on a previous draft. Federico Bianchi is a member of the Bocconi Institute for Data Science and Analytics (BIDSA) and the Data and Marketing Insights (DMI) unit.

#### Ethical Considerations

User data has been collected in the process of providing business services to the clients of *Coveo*: user data is collected and processed in an anonymized fashion, in full compliance with existing legislation (GDPR). In particular, the target dataset uses only anonymous uuids to label sessions and, as such, it does not contain any information that can be linked to individuals.

## References

- Marco Baroni and Roberto Zamparelli. 2010. [Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space](#). In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1183–1193, Cambridge, MA. Association for Computational Linguistics.
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Federico Bianchi, Jacopo Tagliabue, and Bingqing Yu. 2021. [Query2Prod2Vec: Grounded Word Embeddings for eCommerce](#). In *NAACL-HLT*. Association for Computational Linguistics.
- Federico Bianchi, Jacopo Tagliabue, Bingqing Yu, Luca Bigon, and Ciro Greco. 2020. [Fantastic embeddings and how to align them: Zero-shot inference in a multi-shop scenario](#). In *Proceedings of the SIGIR 2020 eCom workshop, July 2020, Virtual Event, published at <http://ceur-ws.org> (to appear)*.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. [Experience grounds language](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8718–8735, Online. Association for Computational Linguistics.
- S. Carey and E. Bartlett. 1978. Acquiring a single new word.
- Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, L. Willems, Chitwan Saharia, T. Nguyen, and Yoshua Bengio. 2019. [BabyAI: A platform to study the sample efficiency of grounded language learning](#). In *ICLR*.
- Gennaro Chierchia and Sally McConnell-Ginet. 2000. *Meaning and Grammar (2nd Ed.): An Introduction to Semantics*. MIT Press, Cambridge, MA, USA.
- Andrew Cotter, Maya R. Gupta, Heinrich Jiang, James Muller, Taman Narayan, Serena Wang, and Tao Zhu. 2018. [Interpretable set functions](#). *ArXiv*, abs/1806.00050.
- 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.
- Nicole J M Fitzgerald and Jacopo Tagliabue. 2020. [On the plurality of graphs](#). In *NETREASON Workshop at 24th European Conference on Artificial Intelligence*.
- Alison Gopnik. 2012. [Scientific thinking in young children: Theoretical advances, empirical research, and policy implications](#). *Science (New York, N.Y.)*, 337:1623–7.
- Mihajlo Grbovic, Vladan Radosavljevic, Nemanja Djuric, Narayan Bhamidipati, Jaikit Savla, Varun Bhagwan, and Doug Sharp. 2015. [E-commerce in your inbox: Product recommendations at scale](#). In *Proceedings of KDD '15*.
- William L. Hamilton, Payal Bajaj, Marinka Zitnik, Dan Jurafsky, and Jure Leskovec. 2018. [Embedding logical queries on knowledge graphs](#). In *NeurIPS*.
- Matthias Hartung, Fabian Kaupmann, Soufian Jebbara, and Philipp Cimiano. 2017. [Learning compositionality functions on word embeddings for modelling attribute meaning in adjective-noun phrases](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 54–64.
- Irene Heim and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Wiley-Blackwell.
- Felix Hill, O. Tieleman, Tamara von Glehn, N. Wong, Hamza Merzic, and Stephen Clark. 2020. [Grounded language learning fast and slow](#). *ArXiv*, abs/2009.01719.
- Paul Jaccard. 1912. [The distribution of the flora in the alpine zone.1](#). *New Phytologist*, 11(2):37–50.
- Eliza Kosoy, Jasmine Collins, David M. Chan, Jessica B. Hamrick, Sandy H. Huang, A. Gopnik, and J. Canny. 2020. [Exploring exploration: Comparing children with RL agents in unified environments](#). In *ICLR Workshop on "Bridging AI and Cognitive Science"*.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2016. [Visual genome: Connecting language and vision using crowdsourced dense image annotations](#). *International Journal of Computer Vision*, 123:32–73.
- Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiosek, Seungjin Choi, and Yee Whye Teh. 2019. [Set transformer: A framework for attention-based permutation-invariant neural networks](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 3744–3753. PMLR.
- Di Lu, Spencer Whitehead, Lifu Huang, Heng Ji, and Shih-Fu Chang. 2018. [Entity-aware image caption generation](#). In *EMNLP*.

- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS'13*, page 3111–3119, Red Hook, NY, USA. Curran Associates Inc.
- Cun Mu, Guang Yang, and Zheng Yan. 2018. [Revisiting skip-gram negative sampling model with regularization](#). In *Proceedings of the 2019 Computing Conference*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). 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 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- W. V. O. Quine. 1960. *Word & Object*. MIT Press.
- Dana Rubinstein, Effi Levi, Roy Schwartz, and Ari Rappoport. 2015. [How well do distributional models capture different types of semantic knowledge?](#) 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 726–730, Beijing, China. Association for Computational Linguistics.
- Joshua S Rule, Joshua B Tenenbaum, and Steven T Piantadosi. 2020. [The child as hacker](#). *Trends in cognitive sciences*, 24(11):900–915.
- Agnieszka Słowik, Abhinav Gupta, William L. Hamilton, Mateja Jamnik, Sean B. Holden, and Christopher Joseph Pal. 2020. [Exploring structural inductive biases in emergent communication](#). *ArXiv*, abs/2002.01335.
- Jacopo Tagliabue and Reuben Cohn-Gordon. 2019. [Lexical learning as an online optimal experiment: Building efficient search engines through human-machine collaboration](#). *ArXiv*, abs/1910.14164.
- Jacopo Tagliabue, Bingqing Yu, and Marie Beaulieu. 2020a. [How to grow a \(product\) tree: Personalized category suggestions for eCommerce type-ahead](#). In *Proceedings of The 3rd Workshop on e-Commerce and NLP*, pages 7–18, Seattle, WA, USA. Association for Computational Linguistics.
- Jacopo Tagliabue, Bingqing Yu, and Federico Bianchi. 2020b. [The embeddings that came in from the cold: Improving vectors for new and rare products with content-based inference](#). In *Fourteenth ACM Conference on Recommender Systems, RecSys '20*, page 577–578, New York, NY, USA. Association for Computing Machinery.
- Michael Tomasello. 2003. *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Harvard University Press, Cambridge, MA.
- Flavian Vasile, Elena Smirnova, and Alexis Conneau. 2016. [Meta-prod2vec: Product embeddings using side-information for recommendation](#). In *Proceedings of the 10th ACM Conference on Recommender Systems*.
- Fei Xu and Joshua B. Tenenbaum. 2000. [Word learning as bayesian inference](#). In *In Proceedings of the 22nd Annual Conference of the Cognitive Science Society*, pages 517–522. Erlbaum.
- Bingqing Yu, Jacopo Tagliabue, Ciro Greco, and Federico Bianchi. 2020. [“An Image is Worth a Thousand Features”: Scalable Product Representations for In-Session Type-Ahead Personalization](#). In *Companion Proceedings of the Web Conference 2020, WWW '20*, page 461–470, New York, NY, USA. Association for Computing Machinery.
- Haonan Yu, H. Zhang, and W. Xu. 2018. [Interactive grounded language acquisition and generalization in a 2d world](#). In *ICLR*.