

Augmenting BERT-style Models with Predictive Coding to Improve Discourse-level Representations

Vladimir Araujo^{1,2}, Andrés Villa¹, Marcelo Mendoza³, Marie-Francine Moens², Alvaro Soto¹

¹Pontificia Universidad Católica de Chile, ²KU Leuven,

³Universidad Técnica Federico Santa María

{vgaraujo, afvilla}@uc.cl, mmendoza@inf.ut fsm.cl

sien.moens@kuleuven.be, asoto@ing.puc.cl

Abstract

Current language models are usually trained using a self-supervised scheme, where the main focus is learning representations at the word or sentence level. However, there has been limited progress in generating useful discourse-level representations. In this work, we propose to use ideas from predictive coding theory to augment BERT-style language models with a mechanism that allows them to learn suitable discourse-level representations. As a result, our proposed approach is able to predict future sentences using explicit top-down connections that operate at the intermediate layers of the network. By experimenting with benchmarks designed to evaluate discourse-related knowledge using pre-trained sentence representations, we demonstrate that our approach improves performance in 6 out of 11 tasks by excelling in discourse relationship detection.

1 Introduction

Pre-trained language models are among the leading methods to learn useful representations for textual data. Several pre-training objectives have been proposed in recent years, such as causal language modeling (Radford et al., 2018, 2019), masked language modeling (Devlin et al., 2019), and permutation language modeling (Yang et al., 2019). However, these approaches do not produce suitable representations at the discourse level (Huber et al., 2020).

Simultaneously, neuroscience studies have suggested that predictive coding (PC) plays an essential role in language development in humans (Ylinen et al., 2016; Zettersten, 2019). PC postulates that the brain is continually making predictions of incoming sensory stimuli (Rao and Ballard, 1999; Friston, 2005; Clark, 2013; Hohwy, 2013), with word prediction being the main mechanism (Berkum et al., 2005; Kuperberg and Jaeger, 2015). However, recent studies speculate that the predictive process could occur within and across utter-

ances, fostering discourse comprehension (Kandylaki et al., 2016; Pickering and Gambi, 2018).

In this work, we propose to extend BERT-type models with recursive bottom-up and top-down computation based on PC theory. Specifically, we incorporate top-down connections that, according to PC, convey predictions from upper to lower layers, which are contrasted with bottom-up representations to generate an error signal that is used to guide the optimization of the model. Using this approach, we attempt to build feature representations that capture discourse-level relationships by continually predicting future sentences in a latent space. We evaluate our approach on DiscoEval (Chen et al., 2019) and SciDTB for discourse evaluation (Huber et al., 2020) to assess whether the embeddings produced by our model capture discourse properties of sentences without finetuning. Our model achieves competitive performance compared to baselines, especially in tasks that require to discover discourse-level relations.

2 Related Work

2.1 BERT for Sentence Representation

Pre-trained self-supervised language models have become popular in recent years. BERT (Devlin et al., 2019) adopts a transformer encoder using a masked language modeling (MLM) objective for word representation. It also proposes an additional loss called next-sentence prediction (NSP) to train a model that understands sentence relationships. On the other hand, ALBERT (Lan et al., 2020) proposes a loss based primarily on coherence called sentence-order prediction (SOP).

SBERT (Reimers and Gurevych, 2019) uses a siamese structure to obtain semantically meaningful sentence embeddings, focusing on textual similarity tasks. ConveRT (Henderson et al., 2020) uses a dual-encoder to improve sentence embeddings for response selection tasks. These models focus on obtaining better representations for specialized sen-

tence pair tasks, so they are not comparable with our which intended to be general-purpose.

More recently, SLM (Lee et al., 2020) proposes a sentence unshuffling approach for a fine understanding of the relations among the sentences at the discourse level. CONPONO (Iter et al., 2020) considers a discourse-level objective to predict the surrounding sentences given an anchor text. This work is related to our approach; the key difference is that our model predicts future sentences sequentially using a top-down pathway. We consider CONPONO as our main baseline.

2.2 Predictive Coding and Deep Learning

Recent work in computer vision takes inspiration from PC theory to build models for accurate (Han et al., 2018) and robust (Huang et al., 2020) image classification. PredNet (Lotter et al., 2017) proposes a network capable of predicting future frames in a video sequence by making local predictions at each level using top-down connections. CPC (Oord et al., 2018) is an unsupervised learning approach to extract useful representations by predicting text in a latent space. Our method takes inspiration from these models, considering top-down connections and predictive processing in a latent space.

3 Proposed Method

3.1 Model Details

Our model consists of a BERT-style model as a sentence encoder (ALBERT and BERT are used in this work) and a GRU model (Cho et al., 2014) that predicts next sentences (see Figure 1). Our intuition is that by giving the model the ability to predict future sentences using a top-down pathway, it will learn better relationships between sentences, thus improving sentence-level representations of each layer for downstream tasks.

The input is a sequence s_1, s_2, \dots, s_n of sentences extracted from a paragraph. We encode sentence s_t with encoder g_{enc} that generates output z_t^l at time step t and layer l (l is from 1 to L). Note that vector z_t^l is obtained from the special token [CLS], which is commonly used as sentence representation. Next, an autoregressive model g_{ar} produces a context vector c_t^l as a function of z_t^l and the context vectors of the upper layer c_t^{l+1} and the previous step c_{t-1}^l .

$$z_t^l = g_{enc}(s_t), c_t^l = g_{ar}(z_t^l, c_t^{l+1}, c_{t-1}^l) \quad (1)$$

Then we introduce a predictive function $f(\cdot)$ to predict a future sentence. In other words, $f(\cdot)$ takes as input the context representation c_t^l from time step

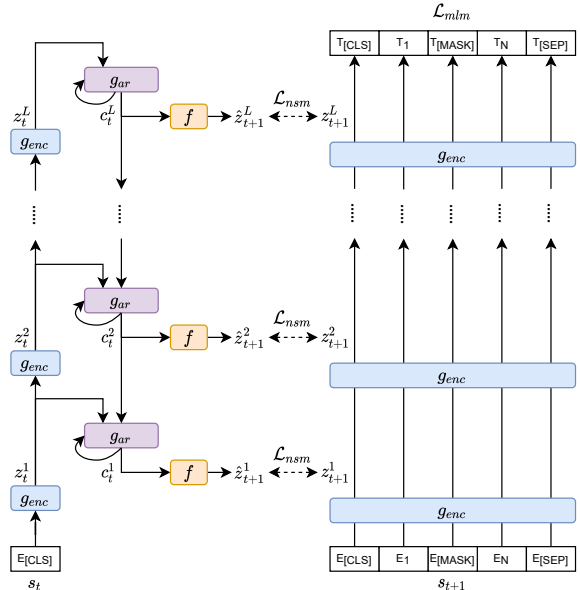


Figure 1: Example of predicting one future sentence. Given an input sentence s_t at time step t , the corresponding representation z_t^l is calculated at layer l . Then a context vector c_t^l is computed via a top-down pathway (left). Afterwards, a future sentence \hat{z}_{t+1}^l is predicted to be compared to the actual representation z_{t+1}^l (right).

t and layer l , and predicts the latent representation \hat{z}_{t+1}^l at time step $t + 1$, *i.e.*:

$$\hat{z}_{t+1}^l = f(c_t^l) \quad (2)$$

In the spirit of Seq2Seq (Sutskever et al., 2014), representations are predicted sequentially, which differs from the CONPONO model that predicts k future sentences with a unique context vector.

3.2 Loss Function

We rely on the InfoNCE loss proposed for the CPC model (Oord et al., 2018). This constructs a binary task in which the goal is to classify one real sample among many noise samples. InfoNCE encourages the predicted representation \hat{z} to be close to the ground truth z .

In the forward pass, the ground truth representation z and the predicted representation \hat{z} are computed at each layer of the model. So we denote the corresponding feature vectors as z_i^j and \hat{z}_i^j where i denotes the temporal index and j is the layer index. A dot product computes the similarity between the predicted and ground truth pair. Then, we optimize a cross-entropy loss that distinguishes the positive pair out of all other negative pairs:

$$\mathcal{L}_{nsm} = - \sum_{i,j} \left[\log \frac{\exp(\hat{z}_i^j \cdot z_i^j)}{\sum_m \exp(\hat{z}_i^j \cdot z_m^j)} \right] \quad (3)$$

There is only one positive pair (z_i^j, z_i^j) for a predicted sentence z_i^j , which are the features at the same time step and the same layer. The rest of pairs (z_i^j, z_m^j) are negative pairs, where $(i, j) \neq (m, j)$. In practice, we draw negative samples from the batch. This is a simple method and a more complex generation of negative samples could improve results. Our loss function, which we refer to as next-sentence modeling (\mathcal{L}_{nsm}), is used in conjunction with the BERT masked language model loss (\mathcal{L}_{mlm}). Accordingly, we train to minimize:

$$\mathcal{L} = \mathcal{L}_{nsm} + \mathcal{L}_{mlm} \quad (4)$$

3.3 Pre-training and Implementation Details

We extend ALBERT and BERT models, obtaining PredALBERT and PredBERT as a result. As mentioned above, our models are fed with a set of contiguous sentences s_n that are processed one-at-a-time. Note that the length of the conventional BERT input is 512 tokens. However, it is unlikely that a sentence will have that many tokens.

We join 3 contiguous sentences to create a long sequence. Longer sequences are truncated, and shorter sequences are padded. We use an overlapping sentence between contiguous sentence groups. For instance, given a paragraph s_1, s_2, \dots, s_9 , the 1st sequence is s_1, s_2, s_3 , the 2nd sequence is s_3, s_4, s_5 , and so on. Our early experiments show that this setting improves the model’s predictive ability in the validation set. We hypothesize that the model can predict up to 3 sentences by using information from the overlapping sentences.

We pre-train our models with the predictive mechanism set to predict the next 2 future sentences ($k = 2$). At time 1, our model represents sequence 1, then this vector feeds the top-down flow (GRU) generating a context representation in each layer that is used to predict sequence 2. Then, the model represents sequence 2 to contrast it with the predicted one. This is repeated one more time to reach $k = 2$ predicted future sequences. For a fair comparison, we train using the BookCorpus (Zhu et al., 2015) and Wikipedia datasets, as well as the BERT, ALBERT, and CONPONO models.

Note that top-down connections are only available during pre-training. At evaluation time, we discard the top-down connections keeping only the encoder, thus obtaining a model equivalent to BERT or ALBERT in terms of the parameters. Table 1 shows the number of parameters in our models.

Model	Parameters		Layers
	Training	Inference	
PredALBERT-B	17.7M	11.7M	12
PredALBERT-L	28.3M	17.7M	24
PredBERT-B	114.2M	109.5M	12

Table 1: Number of parameters of our models.

We used the Huggingface library (Wolf et al., 2020) to implement our models. We initialize the encoder model with BERT or ALBERT weights depending on the version. The autoregressive model was initialized with random weights. For model efficiency in both versions, we use parameter-sharing across layers in the autoregressive model. We trained the models for 1M steps with batch size 8. We use Adam optimizer with weight decay and learning rate of 5e-5. For the masked language modeling, we consider dynamic masking, where the masking pattern is generated every time we feed a sequence to the model. Unlike BERT, we mask 10% of all tokens in each sequence at random.

4 Experiments

4.1 Datasets

Our focus is to evaluate if the discourse properties of sentences are captured by our model without finetuning. DiscoEval (Chen et al., 2019) and SciDTB-DE (Huber et al., 2020) datasets include probing tasks designed for discourse evaluation, thus letting us know what discourse-related knowledge our model is capturing effectively.

DiscoEval: Suite of tasks to evaluate discourse-related knowledge in sentence representations. It includes 7 tasks: Sentence position (SP), Binary sentence ordering (BSO), Discourse coherence (DC), Sentence section prediction (SSP), Penn discourse tree bank (PDTB-E/I), and Rhetorical structure theory (RST). SP, BSO, DC, and SSP assess discourse coherence with binary classification, while PDTB and RST assess discourse relations between sentences through multi-class classification.

SciDTB-DE: Set of tasks designed to determine whether an encoder captures discourse properties from scientific texts. It considers 4 tasks: Swapped units detection (Swapped), Scrambled sentence detection (Scrambled), Relation detection (BinRel), and Relation semantics detection (SemRel). Both Swapped and Scrambled tasks were designed for clause coherence verification, while BinRel and SemRel for discourse relationship detection.

Model	DiscoEval						SciDTB-DE			
	SP	BSO	DC	SSP	PDTB-E/I	RST	BinRel	SemRel	Swap	Scram
ALBERT-B	52.11	70.10	53.89	80.10	37.10 / 39.28	54.03	71.48	56.32	92.04	84.57
ALBERT-L	52.61	70.91	53.85	80.45	37.44 / 38.15	54.71	72.03	54.86	92.67	87.79
BERT-B	53.80	72.33	59.27	80.37	42.52 / 41.97	59.63	76.17	64.93	92.07	86.91
CONPONO	54.52	72.81	58.62	80.87	41.37 / 41.27	59.74	77.86	68.75	91.76	83.62
PredALBERT-B	50.59	69.64	61.69	79.60	40.71 / 41.40	58.54	77.83	65.49	92.51	83.38
PredALBERT-L	51.70	70.35	61.87	79.85	42.66 / 42.92	59.93	78.11	70.15	92.67	84.79
PredBERT-B	50.80	69.94	62.25	79.83	40.10 / 42.20	59.58	76.21	72.92	91.06	80.18

Table 2: Accuracy results in the DiscoEval and SciDTB-DE datasets. We carry out the evaluation 10 times with different seeds and report the average across the trials. B and L indicate the base and large versions, respectively.

4.2 Experimental Setup

Baselines: Following Chen et al. (2019); Huber et al. (2020), we include the results of BERT Base (Devlin et al., 2019). We also evaluate CONPONO (Iter et al., 2020), which is the most related model to our approach. Because these models have more parameters than PredBERT, we also include ALBERT (Lan et al., 2020) Base and Large, which are directly comparable to our model. For a fair and consistent comparison, we rerun all baseline evaluations. We use the pre-trained Huggingface models (Wolf et al., 2020) for BERT and ALBERT. In the case of CONPONO, we use a version pre-trained to predict 2 next surrounding sentences¹.

Evaluation: In the case of DiscoEval, we use the original code provided by Chen et al. (2019). We observe that this configuration leads to CONPONO model results that differ from the reported on the original paper. On the other hand, following Huber et al. (2020), we use SentEval (Conneau and Kiela, 2018) toolkit for SciDTB-DE evaluation. In both cases, the process involves loading a pre-trained model with frozen weights and training a logistic regression on top of the sentence embeddings. To train, we use the average of sentence representations ([CLS]) from all the layers.

5 Results

Table 2 shows the results of our models. We observe improvements in discourse relation detection (PDTB, RST, BinRel, SemRel) and discourse coherence (DC) tasks compared to the best baseline (CONPONO). Across these tasks, PredALBERT-L outperforms by ~ 1.34 points on average, while PredBERT-B by ~ 0.94 . PredALBERT-B achieves competitive performance but does not outperform CONPONO. However, if we compare our models with their direct baselines (ALBERT-B/L, BERT-B), the increase is greater. PredALBERT-

B by ~ 3.92 , PredALBERT-L by ~ 7.43 , and PredBERT-B by ~ 1.46 points on average. The Stuart–Maxwell tests demonstrated a significant difference between our best model PredALBERT-L and ALBERT-L ($p = 0.009$) or CONPONO ($p = 0.05$). We also highlight that PredALBERT-B/L achieves competitive performance with fewer parameters than BERT and CONPONO.

Decreased performance of our models in the SP, BSO, SSP, Swap, and Scram tasks is due to the fact they are closely related to the baselines optimization objectives, which consist of sentence order prediction (ALBERT), topic prediction (BERT), or a combination of them (CONPONO). In contrast, our approach uses a next sentence prediction task in a generative way that encourages the capture of discourse relationships, improving its performance on PDTB, RST, BinRel, and SemRel tasks.

6 Ablation Study

In order to verify the influence of PC mechanism on the pre-training result, we carry out ablation experiments. We use our PredALBERT-B as the Default model, which includes top-down connections and recurrence from layer 12 to layer 1. Ablations involve removing top-down connections and the recurrence of certain layers. Table 3 shows performance across all tasks for each benchmark.

The first experiment uses the PC mechanism on Half the layers, *i.e.*, the GRU and predictive layer are present from layer 12 to layer 6. This variation exceeds the Default model by ~ 0.03 in DiscoEval and ~ 0.14 in SciDTB-DE. The second experiment uses the PC mechanism only on the Last layer of the transformer. It means that the combination of the GRU and prediction layer is only present in layer 12. This reduces the performance by ~ 0.91 in DiscoEval and ~ 2.41 in SciDTB-DE.

Also, we conducted an additional experiment where we removed the top-down connections (w/o TDC) to the Default model. This is equiv-

¹<https://bit.ly/3mn3LQI>

Model	DiscoEval							SciDTB-DE				
	SP	BSO	DC	SSP	PDTB-E/I	RST	Avg	BinRel	SemRel	Swap	Scram	Avg
Default	50.59	69.64	61.69	79.60	40.71 / 41.40	58.54	<u>57.45</u>	77.83	65.49	92.51	83.38	<u>79.80</u>
Half	51.43	70.03	62.71	79.88	40.26 / 40.67	57.37	<u>57.48</u>	79.73	63.54	91.54	84.95	<u>79.94</u>
Last	49.73	68.36	62.88	79.33	38.80 / 39.62	57.06	<u>56.54</u>	74.00	59.51	92.64	83.42	<u>77.39</u>
w/o TDC	45.67	65.76	57.94	78.38	37.49 / 39.32	57.37	<u>54.56</u>	72.14	60.55	88.81	79.99	<u>75.37</u>

Table 3: Results of ablation experiments with PredALBERT-B as the Default model.

alent to modifying equation 1 by $g_{ar}(z_t^l, c_{t-1}^l)$. We found that this ablation severely affects the performance of the Default model performance by ~ 2.89 in DiscoEval and ~ 4.43 in SciDTB-DE.

Our findings indicate that top-down pathway is beneficial for improving discourse representations of BERT-type models. However, it is not clear in which layers it is crucial to have the PC mechanism. We hypothesize that this is related to the fact that the BERT-style models encode syntactic and semantic features in different layers (Jawahar et al., 2019; Aspillaga et al., 2021), so a specialized PC mechanism for syntax or semantics would be desirable. We left this study for future work.

7 What Does The Model Learn?

Because our model excels at detecting discourse relations, in this section, we explore whether the resulting vectors actually represent the role of a sentence in its discursive context. To illustrate what PredBERT learns, we follow the methodology proposed by Lee et al. (2020). We use labeled sentences with discourse relations as queries to retrieve the top 3 most similar sentences from an unlabeled corpus using cosine similarity. We obtained the queries from the MIT Discourse Relations Annotation Guide² and the unlabeled sentences from the Gutenberg dataset (Lahiri, 2014). We compute the representations as mentioned in Section 4.2. This process allowed us to verify that similar vectors share the same or equivalent discourse relations.

Temporal relation: Query = He **knows** a **tasty meal** **when** he eats one.

1. The last five words took Tuppence’s fancy mightily, especially **after** a **meagre breakfast** and a supper of buns the night before.
2. I **know** a disinterested man **when** I see him.
3. He had about ten pounds **when** I found him.

Sentence 1 has a *succession* relation due to the use of the word **after**. Sentence 3 shows a *synchrony* relation because it uses **when** as the query. Sentence 2 does not have a *temporal* relation.

²<https://bit.ly/3z45IG2>

Comparison relation: Query = IBM’s **stock price** rose, **but** the overall **market** fell.

1. The **stock markets** of the world gambled upon its chances, and its bonds at one time were high in favor.
2. Tommy’s heart beat faster, **but** his casual pleasantness did not waver.
3. I guess I was just a mite hasty, **but** I’ve been feeling bad about this **money** question.

Sentence 1 matched **stock** and **market** words but does not contain a *comparison* relation. Sentences 2 and 3 include a *counter-expectation* relation similar to the query sentence that uses the word **but**.

Contingency relation: Query = I refused to **pay** the cobbler the full \$95 **because** he did **poor work**.

1. I did the labor of writing one address this year, **and** got thunder for my **reward**.
2. I don’t believe in a law to prevent a man from getting **rich**; it would do more harm than good.
3. **When** I fixed a plan for an election in Arkansas I did it in ignorance that your convention was doing the same **work**.

All sentences contain semantically related words like **pay/reward** and **poor/rich**. Sentences 1 and 2 include a *cause* relation explicit and implicit. This is related to the query that has *pragmatic cause* relation. Sentence 3 shows a *temporal* relation.

8 Conclusions

We introduce an approach based on PC theory, which extends BERT-style models with recursive bottom-up and top-down computation along with a discourse representation objective. Our models achieve competitive results in discourse analysis tasks, excelling in relations detection.

Acknowledgements

This work was supported in part by the Millennium Institute for Foundational Research on Data (IMFD), the European Research Council Advanced Grant 788506 and the TPU Research Cloud (TRC) program. The first author is grateful to Rodrigo Tufiño for helpful technical advice.

References

- Carlos Aspillaga, Marcelo Mendoza, and Alvaro Soto. 2021. [Inspecting the concept knowledge graph encoded by modern language models](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2984–3000, Online. Association for Computational Linguistics.
- Jos J. A. Van Berkum, Colin M. Brown, Pienie Zwitserlood, Valesca Kooijman, and Peter Hagoort. 2005. [Anticipating upcoming words in discourse: Evidence from ERPs and reading times](#). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(3):443–467.
- Mingda Chen, Zewei Chu, and Kevin Gimpel. 2019. [Evaluation benchmarks and learning criteria for discourse-aware sentence representations](#). 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 649–662, Hong Kong, China. Association for Computational Linguistics.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Andy Clark. 2013. [Whatever next? predictive brains, situated agents, and the future of cognitive science](#). *Behavioral and Brain Sciences*, 36(3):181–204.
- Alexis Conneau and Douwe Kiela. 2018. [SentEval: An evaluation toolkit for universal sentence representations](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- 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.
- Karl Friston. 2005. [A theory of cortical responses](#). *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456):815–836.
- Kuan Han, Haiguang Wen, Yizhen Zhang, Di Fu, Eugenio Culurciello, and Zhongming Liu. 2018. [Deep predictive coding network with local recurrent processing for object recognition](#). In *Advances in Neural Information Processing Systems*, volume 31, pages 9201–9213. Curran Associates, Inc.
- Matthew Henderson, Iñigo Casanueva, Nikola Mrkšić, Pei-Hao Su, Tsung-Hsien Wen, and Ivan Vulić. 2020. [ConveRT: Efficient and accurate conversational representations from transformers](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2161–2174, Online. Association for Computational Linguistics.
- Jakob Hohwy. 2013. *The Predictive Mind*. Oxford University Press.
- Yujia Huang, James Gornet, Sihui Dai, Zhiding Yu, Tan Nguyen, Doris Tsao, and Anima Anandkumar. 2020. [Neural networks with recurrent generative feedback](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 535–545. Curran Associates, Inc.
- Laurine Huber, Chaker Memmadi, Mathilde Dargnat, and Yannick Toussaint. 2020. [Do sentence embeddings capture discourse properties of sentences from scientific abstracts ?](#) In *Proceedings of the First Workshop on Computational Approaches to Discourse*, pages 86–95, Online. Association for Computational Linguistics.
- Dan Iter, Kelvin Guu, Larry Lansing, and Dan Jurafsky. 2020. [Pretraining with contrastive sentence objectives improves discourse performance of language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4859–4870, Online. 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.
- K. D. Kandylaki, A. Nagels, S. Tune, T. Kircher, R. Wiese, M. Schlesewsky, and I. Bornkessel-Schlesewsky. 2016. [Predicting "when" in discourse engages the human dorsal auditory stream: An fMRI study using naturalistic stories](#). *Journal of Neuroscience*, 36(48):12180–12191.
- Gina R. Kuperberg and T. Florian Jaeger. 2015. [What do we mean by prediction in language comprehension?](#) *Language, Cognition and Neuroscience*, 31(1):32–59.
- Shibamouli Lahiri. 2014. [Complexity of word collocation networks: A preliminary structural analysis](#). In *Proceedings of the Student Research Workshop at the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 96–105, Gothenburg, Sweden. Association for Computational Linguistics.

- 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 *International Conference on Learning Representations*.
- Haejun Lee, Drew A. Hudson, Kangwook Lee, and Christopher D. Manning. 2020. [SLM: Learning a discourse language representation with sentence unshuffling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1551–1562, Online. Association for Computational Linguistics.
- William Lotter, Gabriel Kreiman, and David D. Cox. 2017. [Deep predictive coding networks for video prediction and unsupervised learning](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Martin J. Pickering and Chiara Gambi. 2018. [Predicting while comprehending language: A theory and review](#). *Psychological Bulletin*, 144(10):1002–1044.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Rajesh P. N. Rao and Dana H. Ballard. 1999. [Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects](#). *Nature Neuroscience*, 2(1):79–87.
- 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.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- 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, pages 5753–5763. Curran Associates, Inc.
- Sari Ylinen, Alexis Bosseler, Katja Juntila, and Minna Huotilainen. 2016. [Predictive coding accelerates word recognition and learning in the early stages of language development](#). *Developmental Science*, 20(6):e12472.
- Martin Zettersten. 2019. [Learning by predicting: How predictive processing informs language development](#). In Beatrix Busse and Ruth Moehlig-Falke, editors, *Patterns in Language and Linguistics*, pages 255–288. De Gruyter.
- Y. Zhu, R. Kiros, R. Zemel, R. Salakhutdinov, R. Urtasun, A. Torralba, and S. Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 19–27.