

DomBERT: Domain-oriented Language Model for Aspect-based Sentiment Analysis

Hu Xu¹, Bing Liu¹, Lei Shu¹ and Philip S. Yu^{1,2}

¹Department of Computer Science, University of Illinois at Chicago, Chicago, IL, USA

²Institute for Data Science, Tsinghua University, Beijing, China

{hXu48, liub, lshu3, psyu}@uic.edu

Abstract

This paper focuses on learning domain-oriented language models driven by end tasks, which aims to combine the worlds of both general-purpose language models (such as ELMo and BERT) and domain-specific language understanding. We propose DomBERT, an extension of BERT to learn from both in-domain corpus and relevant domain corpora. This helps in learning domain language models with low-resources. Experiments are conducted on an assortment of tasks in aspect-based sentiment analysis (ABSA), demonstrating promising results.¹

1 Introduction

Pre-trained language models (LMs) (Peters et al., 2018; Radford et al., 2018, 2019; Devlin et al., 2019) aim to learn general (or mixed-domain) knowledge for end tasks. Recent studies (Xu et al., 2019; Gururangan et al., 2020) show that learning domain-specific LMs are equally important. This is because the training corpus of general LMs is out-of-domain for end tasks in a particular domain and, more importantly, because general LMs may not capture the long-tailed and underrepresented domain details (Xu et al., 2018). An intuitive example related to corpus of aspect-based sentiment analysis (ABSA) can be found in Table 1, where all masked words *sky*, *water*, *idea*, *screen* and *picture* can appear in a mixed-domain corpus. A general-purpose LM may favor frequent examples and ignore long-tailed choices in certain domains.

In contrast, although domain-specific LMs can capture fine-grained domain details, they may suffer from insufficient training corpus (Gururangan et al., 2020) to strengthen general knowledge within a domain. To this end, we propose a domain-

Example	Domain
The [MASK] is clear .	
The <i>sky</i> is clear .	Astronomy [Irrelevant Domain]
The <i>water</i> is clear .	Liquids [Irrelevant Domain]
The <i>idea</i> is clear .	Concepts [Irrelevant Domain]
The <i>screen</i> is clear .	Desktop [Relevant Domain]
The <i>picture</i> is clear .	Laptop [Target Domain]

Table 1: Multiple choices to recover a masked token (an aspect in ABSA) for different domains: a target domain needs more examples from a relevant domain.

oriented learning task that aims to combine the benefits of both general and domain-specific worlds:

Domain-oriented Learning: Given a target domain t and a set of diverse source domains $S = \{s_1, s_2, \dots\}$, perform (language model) learning that focusing on t and all its relevant domains in S .

This learning task resolves the issues in both general and domain-specific worlds. On one hand, the training of LM does not need to focus on unrelated domains anymore (e.g., *Books* is one big domain but not very related to *laptop*); on the other hand, although an in-domain corpus may be limited, other relevant domains can share a great amount of knowledge (e.g., *Desktop* in Table 1) to make in-domain corpus more diverse and general.

This paper proposes an extremely simple extension of BERT (Devlin et al., 2019) called DomBERT to learn domain-oriented language models. DomBERT simultaneously learns masked language modeling and discovers relevant domains (with a built-in retrieval model (Lewis et al., 2020)) to draw training examples, where the later are computed from domain embeddings learned from an auxiliary task of domain classification. We apply DomBERT to end tasks in aspect-based sentiment analysis (ABSA) in low-resource settings, demonstrating promising results.

Related Work Pre-trained language models gain significant improvements over a wide spectrum of NLP tasks (Minaee et al., 2020), including ELMo (Peters et al., 2018), GPT/GPT2 (Radford

¹The code will be released on <https://github.com/howardhsu/BERT-for-RRC-ABSA>.

et al., 2018, 2019), BERT(Devlin et al., 2019), XL-Net(Yang et al., 2019), RoBERTa(Liu et al., 2019), ALBERT(Lan et al., 2019), ELECTRA(Clark et al., 2019). This paper extends BERT’s masked language model (MLM) with domain knowledge learning. Following RoBERTa, the proposed DomBERT leverages dynamic masking, removes the next sentence prediction (NSP) task (which is proved to have negative effects on pre-trained parameters), and allows for max-length MLM to fully utilize the computational power. This paper also borrows ALBERT’s removal of dropout since pre-trained LM, in general, is an underfitting task that requires more parameters instead of avoiding overfitting.

The proposed domain-oriented learning task can be viewed as one type of transfer learning(Pan and Yang, 2009), which learns a transfer strategy implicitly that transfer training examples from relevant (source) domains to the target domain. This transfer process is conducted throughout the training process of DomBERT.

The experiment of this paper focuses on aspect-based sentiment analysis (ABSA), which typically requires a lot of domain-specific knowledge. Reviews serve as a rich resource for sentiment analysis (Pang et al., 2002; Hu and Liu, 2004; Liu, 2012, 2015). ABSA aims to turn unstructured reviews into structured fine-grained aspects (such as the “battery” or aspect category of a laptop) and their associated opinions (e.g., “good battery” is *positive* about the aspect battery). This paper focuses on three (3) popular tasks in ABSA: aspect extraction (AE) (Hu and Liu, 2004; Li and Lam, 2017), aspect sentiment classification (ASC) (Hu and Liu, 2004; Dong et al., 2014; Nguyen and Shirai, 2015; Li et al., 2018; Tang et al., 2016; Wang et al., 2016a,b; Ma et al., 2017; Chen et al., 2017; Ma et al., 2017; Tay et al., 2018; He et al., 2018; Liu et al., 2018) and end-to-end ABSA (E2E-ABSA) (Li et al., 2019a,b). AE aims to extract aspects (e.g., “battery”), ASC identifies the polarity for a given aspect (e.g., *positive* for *battery*) and E2E-ABSA is a combination of AE and ASC that detects the aspects and their associated polarities simultaneously. This paper focuses on self-supervised methods² to improve ABSA.

²We assume domain tags are largely available online without extra supervised annotation.

2 DomBERT

This section presents DomBERT, which is an extension of BERT for domain knowledge learning. The goal of DomBERT is to discover relevant domains from the pool of source domains and uses the training examples from relevant source domains for masked language model learning. As a result, DomBERT has a sampling process over a categorical distribution on all domains (including the target domain) to retrieve relevant domains’ examples. Learning such a distribution needs to detect the domain similarities between all source domains and the target domain. DomBERT learns an embedding for each domain and computes such similarities. The domain embeddings are learned from an auxiliary task called domain classification.

2.1 Domain Classification

Given a pool of source and target domains, one can easily form a classification task on domain tags. As such, each text document has its domain label l . Following RoBERTa(Liu et al., 2019)’s max-length training examples, we pack different texts from the same domain up to the maximum length into a single training example.

Let the number of source domains be $|S| = n$. Then the number of domains (including the target domain) is $n + 1$. Let $h_{[CLS]}$ denote the hidden state of the [CLS] token of BERT, which indicates the document-level representations of one example. We first pass this hidden states into a dense layer to reduce the size of hidden states. Then we pass this reduced hidden states to a dense layer $D \in \mathbb{R}^{(n+1)*m}$ to compute the logits over all domains \hat{l} :

$$\hat{l} = D \cdot (W \cdot h_{[CLS]} + b), \quad (1)$$

where m is the size of the dense layer, D , W and b are trainable weights. Besides a dense layer, D is essentially a concatenation of domain embeddings: $D = d_t \circ d_1 \circ \dots \circ d_n$. Then we apply cross-entropy loss to the logits and label to obtain the loss of domain classification:

$$\mathcal{L}_{CLS} = \text{CrossEntropyLoss}(\hat{l}, l). \quad (2)$$

To encourage the diversity of domain embeddings, we further compute a regularizer among domain embeddings as following:

$$\Delta = \frac{1}{|D|^2} \|\cos(D, D^T) - I\|_2^2. \quad (3)$$

Minimizing this regularizer encourages the learned embeddings to be more orthogonal (thus diverse) to each other. Finally, we add the loss of domain classification, BERT’s masked language model and regularizer together:

$$\mathcal{L} = \lambda \mathcal{L}_{\text{MLM}} + (1 - \lambda) \mathcal{L}_{\text{CLS}} + \Delta, \quad (4)$$

where λ controls the ratio of losses between masked language model and domain classification.

2.2 Domain Sampler

As a side product of domain classification, DomBERT has a built-in data sampling process to draw examples from both the target domain and relevant domains for future learning. This process follows a unified categorical distribution over all domains, which ensures a good amount of examples from both the target domains and relevant domains are sampled. As such, it is important to always have the target domain t with the highest probability for sampling.

To this end, we use cosine similarity as the similarity function, which has the property to always let $\cos(d_t, d_t) = 1$. For an arbitrary domain i , the probability P_i of domain i being sampled is computed from a softmax function over domain similarities as following:

$$P_i = \frac{\exp(\cos(d_t, d_i)/\tau)}{\sum_{j=0}^{n+1} \exp(\cos(d_t, d_j)/\tau)}, \quad (5)$$

where τ is the temperature (Hinton et al., 2015) to control the importance of highly-ranked domains vs long-tailed domains.

To form a mini-batch for the next training step, we sample domains following the categorical distribution of $s \sim P$ and retrieve the next available example from each sampled domain. As such, we maintain a randomly shuffled queue of examples for each domain. When the examples of one domain are exhausted, a new randomly shuffled queue will be generated for that domain.

3 Experiments

3.1 Datasets

We apply DomBERT to end tasks in aspect-based sentiment analysis from the SemEval dataset, which focusing on *Laptop*, *Restaurant*. We choose 3 end tasks: aspect extraction (AE), aspect sentiment classification (ASC), and end2end ABSA.

For AE, we choose SemEval 2014 Task 4 for laptop and SemEval-2016 Task 5 for restaurant to be consistent with (Xu et al., 2018) and other previous works. For ASC, we use SemEval 2014 Task 4 for both laptop and restaurant as existing research frequently uses this version. We use 150 examples from the training set of all these datasets for validation. For E2E-ABSA, we adopt the formulation of (Li et al., 2019a) where the laptop data is from SemEval-2014 task 4 and the restaurant domain is a combination of SemEval 2014-2016.

Based on the domains of end tasks from SemEval dataset, we explore the capability of the large-scale unlabeled corpus from Amazon review datasets(He and McAuley, 2016) and Yelp dataset³. Following (Xu et al., 2019), we select all laptop reviews from the electronics department. This ends with about 100 MB corpus. Similarly, we simulate a low-resource setting for restaurants and randomly select about 100 MB reviews tagged as *Restaurants* as their first category from Yelp reviews. For source domains S , we choose all reviews from the 5-core version of Amazon review datasets and all Yelp reviews excluding *Laptop* and *Restaurants*. Note that Yelp is not solely about restaurants but has other location-based domains such as car service, bank, theatre, etc. This ends with a total of $|D| = 4680$ domains, and $n = 4679$ are source domains. The total size of the corpus is about 20 GB.

The number of examples for each domain is plotted in Figure 1, where the distribution of domains is heavily long-tailed.

3.2 Hyper-parameters

We adopt BERT_{BASE} (uncased) as the basis for all experiments due to the limits of computational power in our academic setting. We choose the hidden size of domain embeddings $m = 64$ to ensure the regularizer term in the loss doesn’t consume too much GPU memory. We choose $\tau = 0.1$ for

³<https://www.yelp.com/dataset/challenge>, 2019 version.

	Laptop	Restaurant
Training		
Sentence	3045	2000
Aspect	2358	1743
Testing		
Sentence	800	676
Aspect	654	622

Table 2: Summary of datasets on aspect extraction.

	Laptop	Restaurant
Training		
Positive	987	2164
Negative	866	805
Neutral	460	633
Testing		
Positive	341	728
Negative	128	196
Neutral	169	196

Table 3: Summary of datasets on aspect sentiment classification.

	Laptop	Restaurant
Training		
Positive	987	2407
Negative	860	1035
Neutral	450	664
Testing		
Positive	339	1524
Negative	130	500
Neutral	165	263

Table 4: Summary of datasets on end-to-end aspect-based sentiment analysis.

laptop and $\tau = 0.13$ for restaurant domain and $\lambda = 0.9$. We assume the number of training examples per epoch is the number of examples in the target domains. Then, we train DomBERT for 400 epochs to get enough training examples from relevant domains. The full batch size is set to 288. The maximum length of DomBERT is consistent with BERT as 512. We use Adamax(Kingma and Ba, 2014) as the optimizer. Lastly, the learning rate is to be $5e-5$.

3.3 Compared Methods

BERT this is the vanilla **BERT**_{BASE} pre-trained model from (Devlin et al., 2019), which is used to show the performance of BERT without any domain adaption.

BERT-Review post-train BERT on all (mixed-domain) Amazon review datasets and Yelp datasets in a similar way of training BERT. Following (Liu et al., 2019), we train the whole corpus for 4 epochs, which took about 10 days of training (much longer than DomBERT).

BERT-DK is a baseline borrowed from (Xu et al., 2019) that trains an LM per domain. Note that the restaurant domain is trained from 1G of reviews that aligns well with the types of restaurants in

Domain	Laptop	Restaurant
Methods	F1	F1
BERT(Devlin et al., 2019)	79.28	74.1
BERT-Review	83.64	76.20
BERT-DK(Xu et al., 2019)	83.55	77.02
DomBERT	83.89	77.21

Table 5: AE in F1.




Figure 1: Rank of domains by number of examples.

Domain	Laptop		Rest.	
Methods	Acc.	MF1	Acc.	MF1
BERT(Devlin et al., 2019)	75.29	71.91	81.54	71.94
BERT-Review	78.62	75.5	83.35	74.9
BERT-DK(Xu et al., 2019)	77.01	73.72	83.96	75.45
DomBERT	76.72	73.46	83.14	75.00

Table 6: ASC in Accuracy and Macro-F1(MF1).

SemEval, which is not a low-resource case. We use this baseline to show that DomBERT can reach competitive performance.

DomBERT is the model proposed in this paper⁴.

3.4 Evaluation Metrics

For AE, we use F1 score. For ASC, we compute both accuracy and Macro-F1 over 3 classes of polarities, where Macro-F1 is the major metric as the imbalanced classes introduce biases on accuracy. Examples belonging to the *conflict* polarity are dropped as in (Tang et al., 2016). For E2E-ABSA, we adopt the evaluation script from (Li et al., 2019a), which reports precision, recall, and F1 score. Results are as averages of 10 runs.

3.5 Result Analysis and Discussion

AE: In Table 5, we notice that AE is a very domain-specific task. DomBERT further improves the performance of BERT-DK that only uses domain-specific corpus. Note that BERT-DK for restaurant uses 1G of restaurant corpus. But DomBERT’s target domain corpus is just 100 MB. So DomBERT further learns domain-specific knowledge from relevant domains. Although Yelp data contain a great portion of restaurant reviews, a mixed-domain training as BERT-Review does not yield enough domain-specific knowledge.

ASC: ASC is a more domain agnostic task because most of the sentiment words are sharable across all domains (e.g., “good” and “bad”). As such, in Table 6, we notice ASC for restaurant is more domain-specific than laptop. DomBERT is worse than BERT-Review in laptop because a 20+ G can

⁴We do not compare DomBERT with LMs that require extra (directly or indirectly) annotated data.

	Laptop			Restaurant		
	P	R	F1	P	R	F1
Existing Models						
(Li et al., 2019a)	61.27	54.89	57.90	68.64	71.01	69.80
(Luo et al., 2019)	-	-	60.35	-	-	72.78
(He et al.)	-	-	58.37	-	-	-
LSTM-CRF						
(Lample et al., 2016)	58.61	50.47	54.24	66.10	66.30	66.20
(Ma and Hovy, 2016)	58.66	51.26	54.71	61.56	67.26	64.29
(Liu et al., 2018)	53.31	59.40	56.19	68.46	64.43	66.38
BERT+Linear(Li et al., 2019b)	62.16	58.90	60.43	71.42	75.25	73.22
BERT(Devlin et al., 2019)	61.97	58.52	60.11	68.86	73.00	70.78
BERT-Review	65.80	63.12	64.37	69.92	75.36	72.47
BERT-DK(Xu et al., 2019)	63.95	61.18	62.45	71.88	74.07	72.88
DomBERT	66.96	65.58	66.21	72.17	74.96	73.45

Table 7: Results of E2E ABSA: baselines are borrowed from (Li et al., 2019b).

Laptop	Restaurant
Tablets	Food
Boot Shop (Men)	Coffee & Tea
Laptop & Netbook Computer Accessories	Bakeries
Computers & Accessories	Bars
Microsoft Windows	Nightlife
Electronics Warranties	Arts & Entertainment
Desktops	Grocery
Antivirus & Security	Venues & Event Spaces
Aviation Electronics	Lounges
Watch Repair	Beer
Orthopedists	Casinos
Compact Stereos	Hotels
Unlocked Cell Phones	Dance Clubs
Power Strips	Tea Rooms
Mobile Broadband	Pubs
Cleaners	Cinema
No-Contract Cell Phones	Event Planning & Services
Video Games/PC/Accessories	Sports Bars
Antivirus	Specialty Food
MP3 Players & Accessories	Desserts

Table 8: Top-20 relevant domains

learn general-purpose sentiment better. BERT-DK is better than DomBERT because a much larger in-domain corpus is more important for performance. **E2E ABSA:** By combining AE and ASC together, E2E ABSA exhibit more domain-specificity, as shown in Table 7. In this case, we can see the full performance of DomBERT because it can learn both general and domain-specific knowledge well. BERT-Review is poor probably because it focuses on irrelevant domains such as *Books*.

We further investigate relevant domains discovered by DomBERT in Table 8. The results are closer to our intuition because most domains are very close to laptop and restaurant, respectively.

4 Conclusions

We propose DomBERT, which automatically exploits the power of training corpus from relevant domains for a target domain. Experiments demonstrate that DomBERT is promising for ABSA.

Acknowledgments

This work was partially supported by four grants from National Science Foundation: IIS-1910424, IIS-1838770, III-1763325, III-1909323, and SaTC-1930941.

References

- Peng Chen, Zhongqian Sun, Lidong Bing, and Wei Yang. 2017. Recurrent attention network on memory for aspect sentiment analysis. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, pages 452–461.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2019. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*.
- 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.
- Li Dong, Furu Wei, Chuanqi Tan, Duyu Tang, Ming Zhou, and Ke Xu. 2014. Adaptive recursive neural network for target-dependent twitter sentiment classification. In *Proceedings of the 52nd annual meeting of the association for computational linguistics (volume 2: Short papers)*, volume 2, pages 49–54.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don’t stop pretraining: Adapt language models to domains and tasks. In *Proceedings of ACL*.
- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. An interactive multi-task learning network for end-to-end aspect-based sentiment analysis. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2018. Effective attention modeling for aspect-level sentiment classification. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1121–1131.
- Ruining He and Julian McAuley. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *World Wide Web*.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177. ACM.

- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. **Neural architectures for named entity recognition**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*.
- Mike Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. 2020. Pre-training via paraphrasing. *arXiv preprint arXiv:2006.15020*.
- Xin Li, Lidong Bing, Wai Lam, and Bei Shi. 2018. Transformation networks for target-oriented sentiment classification. *arXiv preprint arXiv:1805.01086*.
- Xin Li, Lidong Bing, Piji Li, and Wai Lam. 2019a. A unified model for opinion target extraction and target sentiment prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6714–6721.
- Xin Li, Lidong Bing, Wenxuan Zhang, and Wai Lam. 2019b. Exploiting bert for end-to-end aspect-based sentiment analysis. *arXiv preprint arXiv:1910.00883*.
- Xin Li and Wai Lam. 2017. Deep multi-task learning for aspect term extraction with memory interaction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2886–2892.
- Bing Liu. 2012. Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies*, 5(1):1–167.
- Bing Liu. 2015. *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge University Press.
- L. Liu, J. Shang, F. Xu, X. Ren, H. Gui, J. Peng, and J. Han. 2018. Empower Sequence Labeling with Task-Aware Neural Language Model. In *AAAI*.
- Qiao Liu, Haibin Zhang, Yifu Zeng, Ziqi Huang, and Zufeng Wu. 2018. Content attention model for aspect based sentiment analysis. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*, pages 1023–1032. International World Wide Web Conferences Steering Committee.
- 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 preprint arXiv:1907.11692*.
- Huaishao Luo, Tianrui Li, Bing Liu, and Junbo Zhang. 2019. **DOER: Dual cross-shared RNN for aspect term-polarity co-extraction**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 591–601, Florence, Italy. Association for Computational Linguistics.
- Dehong Ma, Sujian Li, Xiaodong Zhang, and Houfeng Wang. 2017. Interactive attention networks for aspect-level sentiment classification. *arXiv preprint arXiv:1709.00893*.
- Xuezhe Ma and Eduard Hovy. 2016. **End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1064–1074, Berlin, Germany. Association for Computational Linguistics.
- Shervin Minaee, Nal Kalchbrenner, Erik Cambria, Narjes Nikzad, Meysam Chenaghlu, and Jianfeng Gao. 2020. Deep learning based text classification: A comprehensive review. *arXiv preprint arXiv:2004.03705*.
- Thien Hai Nguyen and Kiyooki Shirai. 2015. **PhraseRNN: Phrase recursive neural network for aspect-based sentiment analysis**. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2509–2514, Lisbon, Portugal. Association for Computational Linguistics.
- Sinno Jialin Pan and Qiang Yang. 2009. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359.
- Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up?: sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10*, pages 79–86. Association for Computational Linguistics.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of NAACL-HLT*, pages 2227–2237.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. URL https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/languageunsupervised/language_understanding_paper.pdf.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.

- Duyu Tang, Bing Qin, and Ting Liu. 2016. Aspect level sentiment classification with deep memory network. *arXiv preprint arXiv:1605.08900*.
- Yi Tay, Luu Anh Tuan, and Siu Cheung Hui. 2018. Learning to attend via word-aspect associative fusion for aspect-based sentiment analysis. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Wenya Wang, Sinno Jialin Pan, Daniel Dahlmeier, and Xiaokui Xiao. 2016a. Recursive neural conditional random fields for aspect-based sentiment analysis. *arXiv preprint arXiv:1603.06679*.
- Yequan Wang, Minlie Huang, Li Zhao, et al. 2016b. Attention-based lstm for aspect-level sentiment classification. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 606–615.
- Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu. 2018. Double embeddings and cnn-based sequence labeling for aspect extraction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu. 2019. Bert post-training for review reading comprehension and aspect-based sentiment analysis. In *Proceedings of the 2019 Conference of the North American Chapter of the 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*, pages 5754–5764.