

Article

Not peer-reviewed version

Domain Adaptation with Sentiment Domain Adapter

Aiden Carter , [Wyne Nasir](#) , Ethan Parker *

Posted Date: 18 April 2024

doi: 10.20944/preprints202404.1031.v1

Keywords: Domain Adaptation; Sentiment Analysis; Interpretability



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Domain Adaptation with Sentiment Domain Adapter

Aiden Carter, Wyne Nasir and Ethan Parker *

Tufts University

* Correspondence: eparker@tufts.edu

Abstract: The field of domain adaptation, particularly in cross-domain sentiment classification, leverages labeled data from a source domain alongside unlabeled or sparsely labeled data from a target domain. The objective is to enhance performance in the target domain by mitigating the discrepancies in data distributions. Traditional methods in this area have focused on identifying and differentiating between *pivotal sentiment words* (shared across domains) and *domain-specific sentiment words*. In our work, we introduce a novel framework called Sentiment Domain Adapter (SDA), which incorporates a Category Attention Network (CAN) alongside a Convolutional Neural Network (CNN). This approach treats *pivotal* and *domain-specific* words as part of a collective category of attributes, which SDA learns to discern automatically, thereby enhancing domain adaptation. Additionally, SDA seeks to provide interpretative insights by learning these category attributes. Our model's optimization targets three main goals: (1) supervised classification accuracy, (2) minimizing the disparity in category feature distribution, and (3) maintaining domain invariance. Evaluations on three benchmark datasets for sentiment analysis affirm that SDA surpasses several established baselines.

Keywords: domain adaptation; sentiment analysis; interpretability

1. Introduction

Sentiment analysis [1,2], also known as opinion mining, is a subfield of natural language processing (NLP) that involves identifying and categorizing opinions expressed in a piece of text to determine the writer's attitude towards a particular topic, product, or service. This process typically classifies sentiments into categories such as positive, negative, or neutral, but can also extend to more nuanced emotions like joy, anger, or disappointment [3]. The utility of sentiment analysis is vast, ranging from businesses assessing customer reviews and feedback to gauge public opinion, to social media platforms monitoring user content to understand prevailing attitudes and trends. [4,5] The challenge lies in the subtleties of human language, including sarcasm, irony, and context-dependent meanings, which can skew straightforward computational interpretations.

Technologically, sentiment analysis involves various computational techniques, from simple rule-based algorithms that scan for positive or negative keywords to sophisticated machine learning models that leverage large datasets to understand context and linguistic nuances. With the advent of deep learning, models such as Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) have become particularly effective, as they can process textual data in sequence and capture temporal dependencies and contextual clues that are essential for accurate sentiment **interpretation** [3,11]. These models are often trained on vast corpora of labeled text data, where they learn to associate specific linguistic patterns with sentiment labels. As NLP continues to evolve, the integration of contextual embeddings and transformer models like BERT (Bidirectional Encoder Representations from Transformers) has further revolutionized sentiment analysis, offering even greater precision in capturing the complexities of human emotion expressed through text.

Domain adaptation is a critical technique in machine learning, aimed at addressing the problem of applying an algorithm trained in one setting (the source domain) to a different but related setting (the target domain). This is especially prevalent in fields like natural language processing, computer vision, and sentiment analysis, where data distribution can vary significantly across domains due to differences in language usage, image backgrounds, or contextual nuances [16,17]. The fundamental challenge in domain adaptation is to minimize the domain shift—whereby the model trained on

the source domain underperforms in the target domain due to differences in feature distribution. Techniques like transfer learning, fine-tuning, and domain invariant feature extraction are commonly employed to mitigate this shift, improving the model's ability to generalize across domains without the need for extensive labeling in each new domain.

Domain adaptation strategies benefit immensely from the abundance of labeled data in a source domain, enabling effective adaptation to similar, albeit unlabeled, data distributions in a target domain. In sentiment analysis, the articulation of emotions can vary significantly across domains [18]. For instance, *delicious* might convey positivity in the *Food* domain, while *heartwarming* might serve a similar purpose in the *Movies* domain. This variance often renders classifiers trained in one domain less effective in another.

In more advanced settings, domain adaptation strategies involve both theoretical and algorithmic innovations to create models that can automatically adjust to new environments. For instance, researchers have developed methods that align the statistical properties of data distributions between domains using techniques such as Maximum Mean Discrepancy (MMD) or adversarial training approaches [19,20]. These methods not only adjust the underlying distributions but also enhance the model's interpretability by identifying which features are most relevant for both domains. Recent approaches have also explored the use of deep learning architectures, which can learn complex representations of data that are more adaptable to different domains. Such models often incorporate elements of feature disentanglement and attention mechanisms to focus on the most transferable aspects of the data, further refining the adaptation process for better accuracy and robustness in diverse applications.

Traditional domain adaptation approaches have emphasized the need to identify shared *pivotal* sentiment words and domain-specific *non-pivotal* words. Early methods like Structural Correspondence Learning (SCL) by Blitzer et al. [25,26] and Spectral Feature Alignment (SFA) by Pan et al. [27] have attempted to bridge these domains by aligning *non-pivotal* words with *pivotal* words. However, these methods treated the categories separately and lacked a unified approach.

With the advent of deep learning, new strategies have emerged to decrease the distributional shift in domain adaptation. Techniques like the Maximum Mean Discrepancy (MMD) and various adversarial training methods have become popular for their effectiveness in aligning domain characteristics [20,28–35]. Yet, these approaches often fall short in terms of interpretability, particularly in understanding the role of *pivotal* and *non-pivotal* sentiment words within the adaptation process.

To bridge this gap, we introduce the Sentiment Domain Adapter (SDA), a novel integration of CAN and CNN. This model views *pivotal* and *non-pivotal* words as a unified set of category attributes. SDA includes a Category Memory Module (CMM), a Dynamic Matching (DM) process, and a Category Attention (CA) layer within its architecture. The CMM stores a predefined set of category attributes, which are dynamically matched to each sample through the DM process. The CA layer then focuses on these attributes within the sample, enhancing the model's attention to relevant features for domain adaptation [28,34]. SDA is applied to both the source domain, where CMM is specifically tailored, and the target domain, where CMM starts with a random initialization, offering insights into the transferability of domain knowledge.

Through extensive optimization of the model's objectives, SDA not only focuses on relevant category features but also disregards non-pertinent ones, thus boosting performance in the target domain. The domain-aware CMM and the CA within SDA provide valuable interpretative insights into the process of domain adaptation. Our comprehensive evaluations across multiple real-world datasets reveal that SDA outperforms other models, with further analyses confirming that the domain-aware CMM effectively facilitates knowledge transfer from the source to the target domain, offering a robust interpretation of adaptable domain features.

2. Related Work

The exploration of domain adaptation techniques often revolves around the challenge of distinguishing between pivotal and non-pivotal elements within datasets. One seminal approach, the Structural Correspondence Learning (SCL) method, introduced by Blitzer et al. [25,26], employs a strategy to induce a shared low-dimensional feature space across domains, capitalizing on the co-occurrence of pivotal and non-pivotal elements. This method relies on multiple pivot prediction tasks to uncover the underlying connections between these elements, setting a foundational framework for further studies in this field.

Another notable technique, the Spectral Feature Alignment (SFA) method [27], seeks to cluster non-pivotal features from varied domains into unified groups using pivots as a cohesive link. This alignment facilitates a more seamless domain adaptation by bridging gaps between distinct data characteristics. Additionally, the Adversarial Memory Network (AMN) [42] introduces a dynamic mechanism to identify and leverage non-pivotal features, enhancing the model's adaptability across domains with minimal pivot overlap.

Further advancements have been made with the Hierarchical Attention Transfer Network (HATN) [18], which innovatively captures both pivotal and non-pivotal features without segregating them into separate networks. This model utilizes a dual-network system, P-net and NP-net, each dedicated to handling specific types of features, yet it strives for a unified approach in feature processing, marking a significant step towards integrated domain adaptation models.

In the realm of deep learning, substantial progress has been made towards automating the extraction of robust feature representations for cross-domain applications, particularly in sentiment classification. Early work by Glorot et al. employed a Stacked Denoising Autoencoder (SDA) to derive meaningful, unsupervised feature representations, subsequently applying these features to train a specialized classifier. This approach paved the way for the integration of advanced techniques such as the Maximum Mean Discrepancy (MMD) measure, which has been widely adopted as a regularization strategy to minimize distribution mismatches across domains [28–30,50].

Long et al. enhanced this methodology by incorporating a multiple kernel variant of MMD (MK-MMD), originally proposed by Gretton et al. , into Convolutional Neural Networks (CNNs). This integration significantly improved dataset bias reduction and boosted the transfer capabilities within task-specific layers of CNNs. Furthermore, Dong and de Melo introduced a novel approach for inducing sentiment embeddings through supervision on out-of-domain data, integrating these embeddings into the model via a dedicated memory-based component to further refine the adaptation process.

The recent surge in popularity of adversarial training methods and Generative Adversarial Networks (GANs) has opened new avenues for domain adaptation. Key methodologies in this area include Domain Adaptation with Adversarial Training (DAAT) [34], Domain-Adversarial Neural Network (DANN) [31], Domain Separation Networks (DSN) [32], Selective Adversarial Networks (SAN) [33], among others [20,35,58]. These approaches, while effective, often grapple with issues of interpretability, as the direct learning of transferred category attribute words remains elusive. The lack of clear interpretability can pose significant challenges in user comprehension and trust, particularly when deploying these models in real-world applications, where understanding the basis of model decisions is crucial.

3. Methodology

In this section, we present the conceptual framework and the computational details of our proposed model, the Sentiment Domain Adapter (SDA).

3.1. Problem Formulation

Consider the problem of domain adaptation between a source domain \mathcal{D}_s and a target domain \mathcal{D}_t , both concerned with a binary classification task. We have a set $\mathbf{X}_s = \{(x_1^s, y_1^s), \dots, (x_{N_s}^s, y_{N_s}^s)\}$ of N_s

labeled samples from \mathcal{D}_s , where each sample $\mathbf{x}_i^s \in \mathbb{R}^L$ corresponds to a feature vector of length L and $y_i^s \in \{1, 2, \dots, C\}$ represents the associated category label. The target domain \mathcal{D}_t provides a dataset $\mathbf{X}_t = \{\mathbf{x}_1^t, \dots, \mathbf{x}_{N_t}^t\}$ of N_t unlabeled samples, each also of length L . The objective is to leverage \mathbf{X}_s^{train} , the labeled data from \mathcal{D}_s , along with \mathbf{X}_t^{train} , the unlabeled data from \mathcal{D}_t , to train a model that performs effectively on the test data \mathbf{X}_t^{test} from the target domain.

3.2. Category Attention Network (CAN)

The Category Attention Network (CAN) is designed to emphasize significant categorical features, such as specific sentiment-bearing words in sentiment analysis. Words like *excellent* and *terrible* frequently indicate positive and negative sentiments, respectively, and are crucial for classification accuracy. The CAN comprises several components: a Category Memory Module (CMM), a Dynamic Matching (DM) process, and a Category Attention (CA) layer, each contributing uniquely to the model's performance.

3.2.1. Category Memory Module (CMM)

The CMM is a repository of category-specific attribute words, extracted from labeled data in the source domain. These words are distinctly frequent in their respective categories but rare in others. For instance, we extract attribute words for the positive category by identifying words with a significantly higher frequency in positive contexts than in negative ones, as per the formula:

$$a_1^{pos} = \arg \max_i \left(\log T_i^{pos} - \log T_i^{neg} \right), \quad (1)$$

where T_i^{pos} and T_i^{neg} are the occurrences of the i -th word in positive and negative samples, respectively. This process is mirrored to identify the most significant negative attribute words. The CMM thus comprises the top- M attribute words for each category.

3.2.2. Dynamic Matching (DM)

The DM process dynamically matches category attribute words from the CMM to each sample during training, using cosine similarity to identify the most relevant words for each instance. This is formulated as:

$$s_m^c = \max_{l \in \{1, \dots, L\}} \cos(\mathbf{w}_l, \mathbf{a}_m^c), \quad (2)$$

$$q_1^c = \arg \max_{m \in \{1, \dots, M\}} s_m^c, \quad (3)$$

where \mathbf{w}_l is the embedding vector of the l -th word in a sample, and \mathbf{a}_m^c is the embedding of the m -th attribute word for category c . The top- K attribute words are selected for each category, providing a focused subset for further analysis.

3.2.3. Category Attention (CA)

The CA mechanism applies attention to the dynamically matched words, enhancing the model's focus on relevant category-specific features. It calculates attention weights for each word in a sentence relative to the matched attribute words, significantly highlighting the most indicative features for classification. This is described by:

$$s_{k,l}^c = \frac{\exp(\mathbf{q}_k^{c\top} \tanh(\mathbf{W}_a \mathbf{w}_l + \mathbf{b}))}{\sum_{l=1}^L \exp(\mathbf{q}_k^{c\top} \tanh(\mathbf{W}_a \mathbf{w}_l + \mathbf{b}))}, \quad (4)$$

where \mathbf{W}_a and \mathbf{b} are the weight and bias parameters of the attention mechanism, respectively. This focused attention helps in isolating the most discriminative features within the input data.

3.3. Integration of CAN and CNN for Domain Adaptation

The integration of the Category Attention Network (CAN) with a Convolutional Neural Network (CNN) forms the core of our domain adaptation approach. The CNN, following the architecture of TextCNN [59], extracts broad contextual features from the text, while the CAN focuses on category-specific attributes. The combined features from both networks are then used to predict the final category labels.

The optimization of our model involves a composite loss function that addresses classification accuracy, domain adaptation, and category-specific feature alignment. This includes a supervised classification loss $\mathcal{L}_C(\theta)$ for labeled data in the source domain, a domain adaptation loss $\mathcal{L}_I(\theta)$ using Maximum Mean Discrepancy (MMD) to minimize the difference between source and target feature distributions, and a distribution loss $\mathcal{L}_D(\theta)$ that aligns the attention weights of category attribute words across domains. Each component is essential for ensuring that the model not only performs well on the source domain data but also adapts effectively to the target domain. It is worth noting that our approach is a general framework and the optimization objective acts on the CAN and the features extractor CNN. So the CNN can be replaced by any other efficient feature extractors (e.g., LSTM [60], Transformer [61]).

Table 1. Top five category-specific attribute words identified by CMM in three sentiment analysis domains: Consumer Reviews (CR), Amazon Fine Foods (AFF), and Movie Reviews (MR).

CR		AFF		MR	
Positive	Negative	Positive	Negative	Positive	Negative
excellent	poor	tasty	awful	captivating	dull
satisfied	problematic	yummy	bad	compelling	pointless
best	disappointing	flavorful	disappointing	fascinating	lackluster
superb	worst	scrumptious	unappealing	masterpiece	bland
perfect	terrible	mouthwatering	horrible	inspiring	dreary

Table 2. Comparative accuracy performance of the proposed SDA model against various baseline models, utilizing 10-fold cross-validation across different domain adaptation scenarios.

	Model	MR→CR	AFF→CR	CR→AFF	MR→AFF	CR→MR	AFF→MR
Direct Transfer	fastText-random	0.6290	0.6720	0.6790	0.6900	0.5750	0.5850
	fastText-finetuned	0.6680	0.7470	0.7240	0.7480	0.6550	0.6890
	CNN-char	0.5600	0.6670	0.7140	0.6620	0.5610	0.5930
	CNN-random	0.6070	0.6990	0.7130	0.6750	0.5900	0.6010
	CNN-finetuned	0.6900	0.7580	0.7520	0.7630	0.6680	0.6920
Domain Adaptation	SDA	0.6080	0.6650	0.6750	0.6930	0.6250	0.6350
	mSDA	0.5960	0.6430	0.6810	0.7060	0.6210	0.6390
	SDA-fine-tuned	0.6230	0.6940	0.6900	0.7150	0.6310	0.6430
	DAAT	0.6990	0.7310	0.7220	0.7440	0.6240	0.6530
	SDA (shared CMM)	0.7150	0.7500	0.7660	0.7810	0.6550	0.6970
	SDA	0.7320	0.7650	0.7890	0.7930	0.6800	0.7100

4. Experiments

We assess the efficacy of the SDA model against several benchmarks within the scope of three distinct sentiment analysis datasets. The experimental setup includes diverse datasets and multiple domain adaptation scenarios, reflecting real-world challenges in sentiment analysis.

4.1. Datasets

The datasets employed in our study are: 1) **CR**: Customer Review dataset from Amazon, covering various products. 2) **AFF**: Amazon Fine Foods Review dataset, a subset of which was randomly selected. 3) **MR**: Movie Review dataset from Cornell University, containing diverse film critiques.

These datasets lack a predefined train/test split; therefore, we implement a 10-fold cross-validation approach. This method ensures robustness and generalizability of our findings. The chosen datasets, rich in domain-specific nuances, are ideal for evaluating the adaptability of the SDA model.

4.2. Implementation Details

We construct six domain adaptation tasks from the aforementioned datasets, forming combinations such as MR→CR, AFF→CR, and so on. To ensure a fair comparison across all experiments, we standardize model parameters such as filter widths, feature maps, and embedding dimensions based on the established TextCNN configurations. Hyperparameters such as batch size and learning rate are also unified across all setups.

The CAN component utilizes a configuration of 50 attribute words per category, dynamically selecting 5 for matching in real-time processing. The hyperparameters α and β , determining the strength of the loss function components, are optimized through grid search on the validation set of the CR dataset.

4.3. Baseline Models

Our benchmarks include direct transfer models such as fastText and CNN, each with variations in word vector initialization (random vs. fine-tuned). We also consider domain adaptation models like SDA, mSDA, and adversarial training approaches like DAAT, comparing their performance without needing to explicitly handle pivots and non-pivots separately.

4.4. Performance Comparison

We present our findings in Table 2, where the SDA model generally outperforms all baselines across different adaptation scenarios, highlighting its robustness and effectiveness in leveraging domain-specific knowledge. The results illustrate that pre-trained embeddings and fine-tuning significantly contribute to performance improvements in domain adaptation tasks. Further, the shared CMM approach underlines the importance of domain-aware adaptation, confirming that direct application of source domain knowledge to target domains without adjustments tends to reduce performance.

Additionally, to dissect the impact of individual components of the SDA model, we conduct an ablation study, presented in Table 3. This study confirms that each component of the loss function contributes meaningfully to the overall effectiveness of the model, with the full configuration achieving the best results.

Table 3. Ablation study results showing target domain accuracy of the SDA model with various component configurations.

Model Configuration	MR→CR	CR→AFF	AFF→MR
SDA (without α, β)	0.7164	0.7661	0.7008
SDA (without α)	0.7281	0.7700	0.7044
SDA (without β)	0.7148	0.7867	0.6989
Full SDA Model	0.7302	0.7882	0.7098

4.5. Interpretability Analysis

As mentioned before, the Category Memory Module (CMM) in our Sentiment Domain Adapter (SDA) model utilizes labeled data to distill category-defining attribute words, which poses a limitation when dealing with unlabeled target domain data. Notably, as seen in Table 1, distinct domains manifest unique sets of category attribute words for identical categories (e.g., the words "delicious" and "fast" for positive sentiments in AFF differ from "heartwarming" and "vividly" in MR). Therefore, the transfer of a CMM configured for one domain directly to another, as attempted in the CAN-CNN-shared model, results in suboptimal performance, as evidenced by the comparative outcomes in Table 2.

In the SDA model, the target domain's CMM is initialized randomly and undergoes adaptive refinement during training, aligning with domain-specific semantic contexts. Table 4 showcases the

evolution of vocabulary words closely aligned with category attribute words in the target domain's CMM from their initial random state to their post-training, contextually enriched state. Initially, the resemblance to random words is apparent, but post-training, a clear thematic alignment emerges, reflecting effective domain adaptation. This dynamic adaptation not only illustrates the transfer of contextual knowledge between domains but also underscores the model's ability to internally recalibrate its interpretative focus, aligning it with emergent domain-specific semantics.

Table 4. Evolution of vocabulary words closely associated with category-defining attributes in the target domain, before and after model training.

	MR→CR		CR→AFF		AFF→MR	
	Before	After	Before	After	Before	After
pos.	random1	great	random2	superb	random3	exceptional
	random4	excellent	random5	delicious	random6	captivating
	random7	stunning	random8	perfect	random9	thrilling
	random10	impressive	random11	amazing	random12	enthraling
neg.	random13	poor	random14	dreadful	random15	miserable
	random16	terrible	random17	bad	random18	disappointing
	random19	worst	random20	awful	random21	unwatchable
	random22	pathetic	random23	horrendous	random24	lackluster

4.6. Case Study on Target Domain Sentiment Analysis

We further elucidate the interpretability of the SDA model through a detailed case study, focusing on the visualization of category attention weights in the target domain of Movie Reviews (MR). The weights are averaged from the top- K matched category attribute words, denoted as $s^c = \frac{1}{K} \sum_k s_k^c$, and depicted in Table 5. Here, the intensity of the color overlay on each word corresponds to its computed category attention weight, offering a visual representation of the model's focus.

The analysis reveals nuanced insights into the model's operational dynamics. For instance, in the positive example (1), the words "funny" and "heartbreaking" are highlighted as significant, aligning well with the sample's positive sentiment. Conversely, in the negative example (1), the word "boring" receives substantial emphasis, accurately reflecting its negative sentiment impact. This visual analysis not only confirms the model's effectiveness in identifying sentiment-critical words but also showcases its ability to dynamically adjust focus within different contextual frames, ensuring robust domain adaptation.

Table 5. Visualization of category attention weights for selected sentences from the target domain MR, highlighting the model's focus on sentiment-defining words.

Category	Sentences with Highlighted Words	
Positive (1)	a	wonderfully engaging, narrative
Negative (1)	a	dismal, tale of woe

5. Conclusion and Future Work

In our study, we introduced the Sentiment Domain Adapter (SDA), an innovative model designed to enhance the efficiency of domain adaptation tasks while also providing a mechanism for interpretability. This was achieved through the integration of a Category Attention Network (CAN) with a conventional Convolutional Neural Network (CNN). Our approach simplifies the complex dynamics of learning in domain adaptation by treating both pivots and non-pivots as unified category attributes. This obviates the need for separate network designs for different types of words, streamlining the learning process.

The core innovation of our model lies in its ability to dynamically learn and adjust the category attribute words within the target domain. This functionality not only aids in the adaptive learning

process but also enhances the model's capacity to interpretatively determine the most salient features to transfer from the source to the target domain. The empirical validation of our model on three distinct sentiment analysis datasets—spanning across different domains—showcases significant improvements in performance when compared against a range of existing baseline models.

The findings underscore the SDA model's effectiveness in handling domain discrepancies and its adeptness at learning domain-specific nuances without extensive manual feature engineering. The integration of category attention mechanisms has particularly proven beneficial in refining the feature representations to be more domain-adaptive.

Looking ahead, several avenues remain open for further enhancing the SDA model. Future work could explore the incorporation of more granular attention mechanisms that could fine-tune the interpretation capabilities of the model. Additionally, extending the model's architecture to support multi-lingual datasets could vastly increase its applicability in global sentiment analysis tasks. Moreover, experimenting with different forms of neural network architectures, such as Transformers, may provide deeper insights and potentially yield improvements in both adaptability and accuracy.

Another promising direction would be to enhance the model's ability to handle larger and more diverse datasets, possibly incorporating unsupervised or semi-supervised learning elements to reduce the dependency on labeled data. Finally, a deeper exploration into the interpretability aspect could involve developing visualization tools that allow users to see and understand how the model makes its predictions, thereby increasing trust and transparency in automated decision-making processes.

In conclusion, the SDA model represents a significant step forward in the domain adaptation field, particularly within the realm of sentiment analysis. Its ability to seamlessly adapt and interpret across domains holds great promise for real-world applications, where understanding and reacting to user sentiments across various platforms and demographics is crucial.

References

1. Chi Sun, Luyao Huang, and Xipeng Qiu. Utilizing bert for aspect-based sentiment analysis via constructing auxiliary sentence. *arXiv preprint arXiv:1903.09588*, 2019.
2. Bing Liu. *Sentiment analysis: mining opinions, sentiments, and emotions*. Cambridge University Press, 2015.
3. Aniruddha Tammewar, Alessandra Cervone, and Giuseppe Riccardi. Emotion carrier recognition from personal narratives. *Accepted for publication at INTERSPEECH*, 2021. URL <https://arxiv.org/abs/2008.07481>.
4. Priyank Sonkiya, Vikas Bajpai, and Anukriti Bansal. Stock price prediction using bert and gan, 2021.
5. Hao Fei, Meishan Zhang, and Donghong Ji. Cross-lingual semantic role labeling with high-quality translated training corpus. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7014–7026, 2020.
6. Shengqiong Wu, Hao Fei, Fei Li, Meishan Zhang, Yijiang Liu, Chong Teng, and Donghong Ji. Mastering the explicit opinion-role interaction: Syntax-aided neural transition system for unified opinion role labeling. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence*, pages 11513–11521, 2022.
7. Wenxuan Shi, Fei Li, Jingye Li, Hao Fei, and Donghong Ji. Effective token graph modeling using a novel labeling strategy for structured sentiment analysis. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4232–4241, 2022.
8. Hao Fei, Yue Zhang, Yafeng Ren, and Donghong Ji. Latent emotion memory for multi-label emotion classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7692–7699, 2020.
9. Fengqi Wang, Fei Li, Hao Fei, Jingye Li, Shengqiong Wu, Fangfang Su, Wenxuan Shi, Donghong Ji, and Bo Cai. Entity-centered cross-document relation extraction. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9871–9881, 2022.
10. Ling Zhuang, Hao Fei, and Po Hu. Knowledge-enhanced event relation extraction via event ontology prompt. *Inf. Fusion*, 100:101919, 2023.
11. Hao Fei, Yafeng Ren, and Donghong Ji. Retrofitting structure-aware transformer language model for end tasks. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 2151–2161, 2020.

12. Hao Fei, Shengqiong Wu, Jingye Li, Bobo Li, Fei Li, Libo Qin, Meishan Zhang, Min Zhang, and Tat-Seng Chua. Lasuie: Unifying information extraction with latent adaptive structure-aware generative language model. In *Proceedings of the Advances in Neural Information Processing Systems, NeurIPS 2022*, pages 15460–15475, 2022.
13. Guang Qiu, Bing Liu, Jiajun Bu, and Chun Chen. Opinion word expansion and target extraction through double propagation. *Computational linguistics*, 37(1):9–27, 2011.
14. Hao Fei, Yafeng Ren, Yue Zhang, Donghong Ji, and Xiaohui Liang. Enriching contextualized language model from knowledge graph for biomedical information extraction. *Briefings in Bioinformatics*, 22(3), 2021.
15. Shengqiong Wu, Hao Fei, Wei Ji, and Tat-Seng Chua. Cross2StrA: Unpaired cross-lingual image captioning with cross-lingual cross-modal structure-pivoted alignment. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2593–2608, 2023.
16. Aili Shen, Xudong Han, Trevor Cohn, Timothy Baldwin, and Lea Frermann. Contrastive learning for fair representations, 2021.
17. Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. Next-gpt: Any-to-any multimodal llm. *CoRR*, abs/2309.05519, 2023.
18. Zheng Li, Ying Wei, Yu Zhang, and Qiang Yang. Hierarchical attention transfer network for cross-domain sentiment classification. In *AAAI*, 2018.
19. Jing Han, Zixing Zhang, and Bjorn Schuller. Adversarial training in affective computing and sentiment analysis: Recent advances and perspectives. *IEEE Computational Intelligence Magazine*, 14(2):68–81, 2019.
20. Young-Bum Kim, Karl Stratos, and Dongchan Kim. Adversarial adaptation of synthetic or stale data. In *ACL*, 2017.
21. Hao Fei, Fei Li, Chenliang Li, Shengqiong Wu, Jingye Li, and Donghong Ji. Inheriting the wisdom of predecessors: A multiplex cascade framework for unified aspect-based sentiment analysis. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI*, pages 4096–4103, 2022.
22. Shengqiong Wu, Hao Fei, Yafeng Ren, Donghong Ji, and Jingye Li. Learn from syntax: Improving pair-wise aspect and opinion terms extraction with rich syntactic knowledge. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 3957–3963, 2021.
23. Bobo Li, Hao Fei, Lizi Liao, Yu Zhao, Chong Teng, Tat-Seng Chua, Donghong Ji, and Fei Li. Revisiting disentanglement and fusion on modality and context in conversational multimodal emotion recognition. In *Proceedings of the 31st ACM International Conference on Multimedia, MM*, pages 5923–5934, 2023.
24. Hao Fei, Qian Liu, Meishan Zhang, Min Zhang, and Tat-Seng Chua. Scene graph as pivoting: Inference-time image-free unsupervised multimodal machine translation with visual scene hallucination. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5980–5994, 2023.
25. John Blitzer, Ryan McDonald, and Fernando Pereira. Domain adaptation with structural correspondence learning. In *Proceedings of the 2006 conference on empirical methods in natural language processing*, pages 120–128. *ACL*, 2006.
26. John Blitzer, Mark Dredze, and Fernando Pereira. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *ACL*, pages 440–447, 2007.
27. Sinno Jialin Pan, Xiaochuan Ni, Jian-Tao Sun, Qiang Yang, and Zheng Chen. Cross-domain sentiment classification via spectral feature alignment. In *WWW*, pages 751–760. *ACM*, 2010.
28. Muhammad Ghifary, W Bastiaan Kleijn, and Mengjie Zhang. Domain adaptive neural networks for object recognition. In *Pacific Rim international conference on artificial intelligence*, 2014.
29. Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv:1412.3474*, 2014.
30. Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. In *NeurIPS*, 2016.
31. Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 17(1): 2096–2030, 2015.
32. Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. In *NeurIPS*, 2016.
33. Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Michael I Jordan. Partial transfer learning with selective adversarial networks. In *CVPR*, 2018.

34. Muhammad Imran Firoj Alam, Shafiq Joty. Domain adaptation with adversarial training and graph embeddings. In *ACL*, 2018.
35. Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. Wasserstein distance guided representation learning for domain adaptation. In *AAAI*, 2018.
36. Jingye Li, Kang Xu, Fei Li, Hao Fei, Yafeng Ren, and Donghong Ji. MRN: A locally and globally mention-based reasoning network for document-level relation extraction. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1359–1370, 2021.
37. Hao Fei, Shengqiong Wu, Yafeng Ren, and Meishan Zhang. Matching structure for dual learning. In *Proceedings of the International Conference on Machine Learning, ICML*, pages 6373–6391, 2022.
38. Hu Cao, Jingye Li, Fangfang Su, Fei Li, Hao Fei, Shengqiong Wu, Bobo Li, Liang Zhao, and Donghong Ji. OneEE: A one-stage framework for fast overlapping and nested event extraction. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1953–1964, 2022.
39. Hao Fei, Fei Li, Bobo Li, and Donghong Ji. Encoder-decoder based unified semantic role labeling with label-aware syntax. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 12794–12802, 2021.
40. Bobo Li, Hao Fei, Fei Li, Yuhan Wu, Jinsong Zhang, Shengqiong Wu, Jingye Li, Yijiang Liu, Lizi Liao, Tat-Seng Chua, and Donghong Ji. DiaASQ: A benchmark of conversational aspect-based sentiment quadruple analysis. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13449–13467, 2023.
41. Shengqiong Wu, Hao Fei, Yixin Cao, Lidong Bing, and Tat-Seng Chua. Information screening whilst exploiting! multimodal relation extraction with feature denoising and multimodal topic modeling. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14734–14751, 2023.
42. Zheng Li, Yun Zhang, Ying Wei, Yuxiang Wu, and Qiang Yang. End-to-end adversarial memory network for cross-domain sentiment classification. In *IJCAI*, pages 2237–2243, 2017.
43. Hao Fei, Shengqiong Wu, Yafeng Ren, Fei Li, and Donghong Ji. Better combine them together! integrating syntactic constituency and dependency representations for semantic role labeling. In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021*, pages 549–559, 2021.
44. Shengqiong Wu, Hao Fei, Hanwang Zhang, and Tat-Seng Chua. Imagine that! abstract-to-intricate text-to-image synthesis with scene graph hallucination diffusion. *Advances in Neural Information Processing Systems*, 36, 2024.
45. Hao Fei, Shengqiong Wu, Wei Ji, Hanwang Zhang, and Tat-Seng Chua. Empowering dynamics-aware text-to-video diffusion with large language models. *arXiv preprint arXiv:2308.13812*, 2023.
46. Leigang Qu, Shengqiong Wu, Hao Fei, Liqiang Nie, and Tat-Seng Chua. Layoutllm-t2i: Eliciting layout guidance from llm for text-to-image generation. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 643–654, 2023.
47. Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *ICML*, 2011.
48. Hao Fei, Yafeng Ren, and Donghong Ji. Boundaries and edges rethinking: An end-to-end neural model for overlapping entity relation extraction. *Information Processing & Management*, 57(6):102311, 2020.
49. Jingye Li, Hao Fei, Jiang Liu, Shengqiong Wu, Meishan Zhang, Chong Teng, Donghong Ji, and Fei Li. Unified named entity recognition as word-word relation classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10965–10973, 2022.
50. Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I Jordan. Learning transferable features with deep adaptation networks. *arXiv:1502.02791*, 2015.
51. Hao Fei, Tat-Seng Chua, Chenliang Li, Donghong Ji, Meishan Zhang, and Yafeng Ren. On the robustness of aspect-based sentiment analysis: Rethinking model, data, and training. *ACM Transactions on Information Systems*, 41(2):50:1–50:32, 2023.
52. Yu Zhao, Hao Fei, Yixin Cao, Bobo Li, Meishan Zhang, Jianguo Wei, Min Zhang, and Tat-Seng Chua. Constructing holistic spatio-temporal scene graph for video semantic role labeling. In *Proceedings of the 31st ACM International Conference on Multimedia, MM*, pages 5281–5291, 2023.
53. Hao Fei, Yafeng Ren, Yue Zhang, and Donghong Ji. Nonautoregressive encoder-decoder neural framework for end-to-end aspect-based sentiment triplet extraction. *IEEE Transactions on Neural Networks and Learning Systems*, 34(9):5544–5556, 2023.

54. Yu Zhao, Hao Fei, Wei Ji, Jianguo Wei, Meishan Zhang, Min Zhang, and Tat-Seng Chua. Generating visual spatial description via holistic 3D scene understanding. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7960–7977, 2023.
55. Arthur Gretton, Dino Sejdinovic, Heiko Strathmann, Sivaraman Balakrishnan, Massimiliano Pontil, Kenji Fukumizu, and Bharath K Sriperumbudur. Optimal kernel choice for large-scale two-sample tests. In *NeurIPS*, 2012.
56. Hao Fei, Bobo Li, Qian Liu, Lidong Bing, Fei Li, and Tat-Seng Chua. Reasoning implicit sentiment with chain-of-thought prompting. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1171–1182, 2023.
57. Xin Dong and Gerard de Melo. A helping hand: Transfer learning for deep sentiment analysis. In *ACL*, pages 2524–2534, 2018.
58. Ming-Yu Liu and Oncel Tuzel. Coupled generative adversarial networks. In *NeurIPS*, 2016.
59. Yoon Kim. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, 2014.
60. Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
61. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. pages 5998–6008, 2017.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.