

KinGDOM: Knowledge-Guided DOMain Adaptation for Sentiment Analysis

Deepanway Ghosal[†], Devamanyu Hazarika^Φ, Abhinaba Roy[⊙],
Navonil Majumder[†], Rada Mihalcea[□], Soujanya Poria[†]

[†]Singapore University of Technology and Design, Singapore

^ΦNational University of Singapore, Singapore

[⊙]Nanyang Technological University, Singapore

[□]University of Michigan, USA

{deepanway_ghosal@mymail., sporia@, nmder@}sutd.edu.sg,
{hazarika@comp.nus, abhinaba.roy@ntu}.edu.sg, mihalcea@umich.edu

Abstract

Cross-domain sentiment analysis has received significant attention in recent years, prompted by the need to combat the domain gap between different applications that make use of sentiment analysis. In this paper, we take a novel perspective on this task by exploring the role of external commonsense knowledge. We introduce a new framework, *KinGDOM*, which utilizes the ConceptNet knowledge graph to enrich the semantics of a document by providing both domain-specific and domain-general background concepts. These concepts are learned by training a graph convolutional autoencoder that leverages inter-domain concepts in a domain-invariant manner. Conditioning a popular domain-adversarial baseline method with these learned concepts helps improve its performance over state-of-the-art approaches, demonstrating the efficacy of our proposed framework.

1 Introduction

Sentiment Analysis (SA) is a popular NLP task used in many applications (Zhang et al., 2018). Current models trained for this task, however, cannot be reliably deployed due to the distributional mismatch between the training and evaluation domains (Daumé III and Marcu, 2006). *Domain adaptation*, a case of transductive transfer learning, is a widely studied field of research that can be effectively used to tackle this problem (Wilson and Cook, 2018).

Research in the field of cross-domain SA has proposed diverse approaches, which include learning domain-specific sentiment words/lexicons (Sarma et al., 2018; Hamilton et al., 2016b), co-occurrence based learning (Blitzer et al., 2007a), domain-adversarial learning (Ganin et al., 2016), among

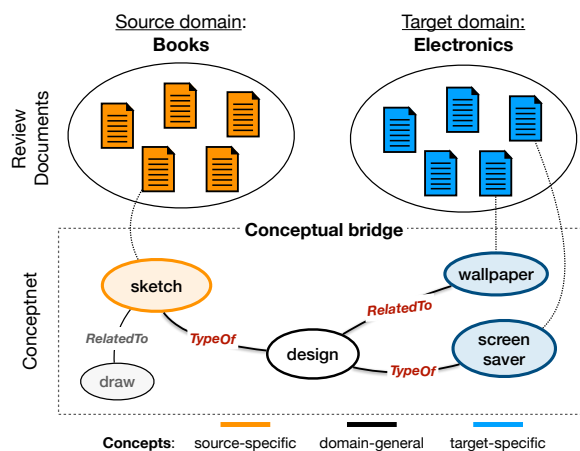


Figure 1: *ConceptNet* provides networks with background concepts that enhance their semantic understanding. For example, for a target sentence from electronics domain, *The software came with decent screen savers*, comprising domain-specific terms like *screen saver* or *wallpaper*, *ConceptNet* helps connecting them to general concepts like *design*, thus allowing a network better understand their meaning. Furthermore, inter-domain conceptual bridge can also be established to connect source and target domains (*wallpaper*–*sketch* have similar conceptual notions under the link *design*).

others. In this work, we adopt the domain-adversarial framework and attempt to improve it further by infusing commonsense knowledge using *ConceptNet* – a large-scale knowledge graph (Speer et al., 2017).

Augmenting neural models with external knowledge bases (KB) has shown benefits across a range of NLP applications (Peters et al., 2019; Li et al., 2019; IV et al., 2019; liu et al., 2019; Bi et al., 2019). Despite their popularity, efforts to incorporate KBs into the domain-adaptation framework has been sporadic (Wang et al., 2008; Xiang et al., 2010). To this end, we identify multiple advantages of using commonsense KBs for domain adaptation.

First, KBs help in grounding text to real enti-

ties, factual knowledge, and commonsense concepts. Commonsense KBs, in particular, provide a rich source of background concepts—related by commonsense links—which can enhance the semantics of a piece of text by providing both domain-specific and domain-general concepts (Yang et al., 2019; Zhong et al., 2019; Agarwal et al., 2015; Zhong et al., 2019) (see Fig. 1). For cross-domain SA, word polarities might vary among different domains. For example, *heavy* can be a positive feature for a truck, but a negative feature for a smartphone. It is, however, difficult to assign contextual-polarities solely from data, especially when there is no supervision (Boia et al., 2014). In this domain-specific scenario, commonsense knowledge provides a dynamic way to enhance the context and help models understand sentiment-bearing terms and opinion targets through its structural relations (Cambria et al., 2018). They also often aid in unearthing implicitly expressed sentiment (Balahur et al., 2011).

Second, domains often share relations through latent semantic concepts (Kim et al., 2017a). For example, notions of *wallpaper* (from electronics) and *sketch* (from books) can be associated via related concepts such as *design* (see Fig. 1). Multi-relational KBs provide a natural way to leverage such inter-domain relationships. These connections can help models understand target-specific terms by associating to known domain-general or even source-specific concepts.

Following these intuitions, we propose a two-step modular framework, *KinGDOM* (**K**nowledge-**G**uided **D**omain adaptation), which utilizes commonsense KB for domain adaptation. *KinGDOM* first trains a shared graph autoencoder using a graph convolution network (GCN) on ConceptNet, so as to learn: 1) inter-domain conceptual links through multiple inference steps across neighboring concepts; and 2) domain-invariant concept representations due to shared autoencoding. It then extracts document-specific sub-graph embeddings and feeds them to a popular domain-adversarial model DANN (Ganin et al., 2016). Additionally, we also train a shared autoencoder on these extracted graph embeddings to promote further domain-invariance (Glorot et al., 2011).

Our main contributions in this work are:

1. We propose *KinGDOM*, a domain-adversarial framework that uses an external KB (ConceptNet) for unsupervised domain adaptation. *KinG-*

DOM learns domain-invariant features of KB concepts using a graph autoencoding strategy.

2. We demonstrate, through experiments, that *KinGDOM* surpasses state-of-the-art methods on the Amazon-reviews dataset (Blitzer et al., 2007b), thus validating our claim that external knowledge can aid the task of cross-domain SA.

In the remaining paper, §2 explains related works and compares *KinGDOM* to them; §3 presents task definition and preliminaries; §4 introduces our proposed framework, *KinGDOM*; §5 discusses experimental setup followed by results and extensive analyses in §6; finally, §7 concludes this paper.

2 Related Work

Domain adaptation methods can be broadly categorized into three approaches: a) *instance-selection* (Jiang and Zhai, 2007; Chen et al., 2011; Cao et al., 2018), b) *self-labeling* (He and Zhou, 2011) and c) *representation learning* (Glorot et al., 2011; Chen et al., 2012; Tzeng et al., 2014). Our focus is on the third category which has emerged as a popular approach in this deep representation learning era (Ruder, 2019; Poria et al., 2020).

Domain-adversarial Training. Our work deals with domain-adversarial approaches (Kouw and Loog, 2019), where we extend DANN Ganin et al. (2016). Despite its popularity, DANN cannot model domain-specific information (e.g. indicators of *tasty*, *delicious* for kitchen domain) (Peng et al., 2018b). Rectifications include shared-private encoders that model both domain-invariant and -specific features (Li et al., 2012; Bousmalis et al., 2016a; Kim et al., 2017b; Chang et al., 2019), using adversarial and orthogonality losses (Liu et al., 2017; Li et al., 2018). Although we do not use private encoders, we posit that our model is capable of capturing domain-specificity via the sentence-specific concept graph. Also, our approach is flexible enough to be adapted to the setup of shared-private encoders.

External Knowledge. Use of external knowledge has been explored in both inductive and transductive settings (Banerjee, 2007; Deng et al., 2018). Few works have explored external knowledge in domain adaptation based on Wikipedia as auxiliary information, using co-clustering (Wang et al., 2008) and semi-supervised learning (SSL) (Xiang et al., 2010). SSL has also been explored by Alam

et al. (2018) in the Twitter domain. Although we share a similar motivation, there exist crucial differences. Primarily, we learn graph embeddings at the concept level, not *across* complete instances. Also, we do not classify each concept node in the graph, which renders SSL inapplicable to our setup.

Domain Adaptation on Graphs. With the advent of graph neural networks, graph-based methods have become a new trend (Ghosal et al., 2019) in diverse NLP tasks such as emotion recognition in conversations (Poria et al., 2019). Graph-based domain adaptation is categorized based on the availability of cross-domain connections. For domain-exclusive graphs, approaches include SSL with GCNs (Shen and Chung, 2019) and domain-adversarial learning (Dai et al., 2019). For cross-domain connected graphs, co-regularized training (Ni et al., 2018) and joint-embedding (Xu et al., 2017) have been explored. We also utilize GCNs to learn node representations in our cross-domain ConceptNet graph. However, rather than using explicit divergence measures or domain-adversarial losses for domain invariance, we uniquely adopt a shared-autoencoder strategy on GCNs. Such ideas have been explored in vector-based approaches (Glorot et al., 2011; Chen et al., 2012).

Sentiment Analysis. One line of work models domain-dependent word embeddings (Sarma et al., 2018; Shi et al., 2018; K Sarma et al., 2019) or domain-specific sentiment lexicons (Hamilton et al., 2016a), while others attempt to learn representations based on co-occurrences of domain-specific with domain-independent terms (Blitzer et al., 2007a; Pan et al., 2010; Sharma et al., 2018). Our work is related to approaches that address domain-specificity in the target domain (Peng et al., 2018b; Bhatt et al., 2015). Works like Liu et al. (2018) attempts to model target-specificity by mapping domain-general information to domain-specific representations by using domain descriptor vectors. In contrast, we address relating domain-specific terms by modeling their relations with the other terms in knowledge bases like ConceptNet.

3 Background

3.1 Task Definition

Domain adaptation deals with the training of models that can perform inference reliably in multiple domains. Across domains, it is assumed that the feature and label spaces are the same but with dis-

crepancies in their feature distributions. In our setup, we consider two domains: source \mathcal{D}_s and target domain \mathcal{D}_t with different marginal data distributions, i.e., $P_{\mathcal{D}_s}(x) \neq P_{\mathcal{D}_t}(x)$. This scenario, also known as the *covariate shift* (Elsahar and Gallé, 2019), is predominant in SA applications and arises primarily with shifts in topics – causing a difference in vocabulary usage and their corresponding semantic and sentiment associations.

We account for *unsupervised* domain adaptation, where we are provided with labeled instances from the source domain $D_s^l = \{(x_i, y_i)\}_{i=1}^{N_s}$ and unlabeled instances from the target domain $D_t^u = \{(x_i)\}_{i=1}^{N_t}$.¹ This is a realistic setting as curating annotations for the target domain is often expensive as well as time consuming. Given this setup, our goal is to train a classifier that can achieve good classification performance on the target domain.

3.2 Domain-Adversarial Neural Network

We base our framework on the domain-adversarial neural network (DANN) proposed by Ganin et al. (2016). DANN learns a shared mapping of both source and target domain instances $M(\mathbf{x}_{s/t})$ such that a classifier C trained for the source domain can be directly applied for the target domain. Training of C is performed using the cross-entropy loss:

$$\mathcal{L}_{\text{cls}} = \mathbb{E}_{(x_s, y_s)} \left(- \sum_{k=1}^K \mathbb{1}_{[k=y_s]} \log C(M(x_s)) \right),$$

where K is the number of labels. Both the mapping function M and the classifier C are realized using neural layers with parameters θ_M and θ_C .

Adversarial Loss. The core idea of DANN is to reduce domain gap by learning common representations that are indistinguishable to a domain discriminator. To learn a domain-invariant mapping, DANN uses an adversarial discriminator D_{adv} with parameters θ_D , whose job is to distinguish between source and target instances, $M(x_s)$ vs. $M(x_t)$. It is trained using the cross-entropy loss:

$$\mathcal{L}_{\text{adv}_D} = -\mathbb{E}_{\mathbf{x}_s} (\log D_{adv}(M(\mathbf{x}_s))) - \mathbb{E}_{\mathbf{x}_t} (\log (1 - D_{adv}(M(\mathbf{x}_t)))) .$$

The mapping function then learns domain invariance by pitting against the discriminator in a minimax optimization with loss $\mathcal{L}_{\text{adv}_M} = -\mathcal{L}_{\text{adv}_D}$ (Tzeng et al., 2017). This setup forces the features to become discriminative to the main

¹For our case, each instance is a review document




Figure 2: Illustration of *KinGDOM*: Step 1 uses GCN to learn concept representations. Step 2 feeds concept features to DANN.

All domains	DVD	Books	Kitchen	Electronics
RelatedTo (580k)	RelatedTo	RelatedTo	RelatedTo	RelatedTo
HasContext (80k)	HasContext	HasContext	IsA	IsA
IsA (60k)	IsA	IsA	Synonym	Synonym
DerivedFrom (42k)	Synonym	Synonym	DerivedFrom	DerivedFrom
Synonym (40k)	DerivedFrom	DerivedFrom	HasContext	HasContext
AtLocation (14k)	AtLocation	CapableOf	AtLocation	AtLocation
UsedFor (12k)	CapableOf	AtLocation	UsedFor	UsedFor
CapableOf (11k)	UsedFor	SimilarTo	SimilarTo	SimilarTo
SimilarTo (10k)	SimilarTo	UsedFor	CapableOf	CapableOf
Etymologically (5k)	Antonym	Antonym	Antonym	Antonym

Table 1: Top-10 relations of \mathcal{G}' based on frequency. Top relations for each domain are also mentioned.

learning task and indistinguishable across domains. The point estimates of the parameters are decided at a saddle point using the minimax objective:

$$\theta^* = \arg \min_{\theta_{M,C}} \max_{\theta_D} (\mathcal{L}_{cls} + \lambda \mathcal{L}_{adv_D}),$$

where λ is a hyper-parameter. The minimax objective is realized by reversing the gradients of \mathcal{L}_{adv_D} when back-propagating through M .

4 Our Proposed Method

KinGDOM aims to improve the DANN approach by leveraging an external knowledge source i.e., ConceptNet. Such a knowledge base is particularly useful for domain adaptation as it contains both domain specific and domain general knowledge. Unlike traditional word embeddings and semantic knowledge graphs (e.g. WordNet), ConceptNet is unique as it contains commonsense related information. We posit that both these properties of

ConceptNet will be highly useful for domain adaptation. *KinGDOM* follows a two-step approach described below:

Step 1: This step deals with training a domain-aggregated sub-graph of ConceptNet. In particular, it involves: a) Creating a sub-graph of ConceptNet based on all domains (§4.1). b) Training a graph-convolutional autoencoder to learn concept embeddings (Schlichtkrull et al., 2018) (§4.2).

Step 2: After the graph autoencoder is trained, a) we extract and pool document-relevant features from the trained graph for each instance in the dataset (§4.3). b) The corresponding graph feature vector is then fed into the DANN architecture for adversarial training (Ganin et al., 2016). To further enforce domain invariance, we also introduce a shared autoencoder to reconstruct the graph features (§4.4).

4.1 Step 1a) Domain-Aggregated Commonsense Graph Construction

We construct our domain-aggregated graph from ConceptNet (Speer et al., 2017). First, we introduce the following notation: the ConceptNet graph is represented as a directed labeled graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{R})$, with concepts/nodes ² $v_i \in \mathcal{V}$ and labeled edges $(v_i, r_{ij}, v_j) \in \mathcal{E}$, where $r_{ij} \in \mathcal{R}$ is the relation type of the edge between v_i and v_j . The concepts in ConceptNet are unigram words or n-gram phrases. For instance one such triplet from ConceptNet is *[baking-oven, AtLocation, kitchen]*.

ConceptNet has approximately 34 million edges, from which we first extract a subset of edges. From the training documents of all domains in our dataset, we first extract the set of all the unique nouns, adjectives, and adverbs.³ These extracted words are treated as the seeds that we use to filter ConceptNet into a sub-graph. In particular, we extract all the triplets from \mathcal{G} which are within a distance of 1 to any of those seed concepts, resulting in a sub-graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}', \mathcal{R}')$, with approximately 356k nodes and 900k edges. This sub-graph would thus contain concepts across all domains along with inter-concept links. Looking at the sub-graph \mathcal{G}' from the lens of each domain, we can observe the top-10 relations within the domain in Table 1.

4.2 Step 1b) Knowledge Graph Pre-training

To utilize \mathcal{G}' in our task, we first need to compute a representation of its nodes. We do this by training a graph autoencoder model to perform link prediction. The model takes as input an incomplete set of edges $\hat{\mathcal{E}}'$ from \mathcal{E}' in \mathcal{G}' and then assign scores to possible edges (c_1, r, c_2) , determining how likely are these edges to be in \mathcal{E}' . Following Schlichtkrull et al. (2018), our graph autoencoder model consists of: a R-GCN entity encoder and a DistMult scoring decoder.

Encoder Module. We employ the Relational Graph Convolutional Network (R-GCN) encoder from Schlichtkrull et al. (2018) as our graph encoder network. The power of this model comes from its ability to accumulate relational evidence in multiple inference steps from the local neighborhood around a given concept. The neighborhood-based convolutional feature transformation process always ensures that distinct domains are connected

²We use *node*, *concept*, and *entity* interchangeably

³We use the Spacy POS Tagger: <https://spacy.io/usage/linguistic-features#pos-tagging>

via underlying concepts and influence each other to create enriched domain-aggregated feature vectors.

Precisely, our encoder module consists of two R-GCN encoders stacked upon one another. The initial concept feature vector \mathbf{g}_i is initialized randomly and thereafter transformed into the domain-aggregated feature vector $\mathbf{h}_i \in \mathbb{R}^d$ using the two-step graph convolution process. The transformation process is detailed below:

$$f(\mathbf{x}_i, l) = \sigma\left(\sum_{r \in \mathcal{R}} \sum_{j \in N_i^r} \frac{1}{c_{i,r}} W_r^{(l)} \mathbf{x}_j + W_0^{(l)} \mathbf{x}_i\right),$$

$$\mathbf{h}_i = \mathbf{h}_i^{(2)} = f(\mathbf{h}_i^{(1)}, 2) \quad ; \quad \mathbf{h}_i^{(1)} = f(\mathbf{g}_i, 1),$$

where N_i^r denotes the neighbouring concepts of concept i under relation $r \in \mathcal{R}$; $c_{i,r}$ is a normalization constant which either can be set in advance, such that, $c_{i,r} = |N_i^r|$, or can be learned in a gradient-based learning setup. σ is an activation function such as ReLU, and $W_r^{(1/2)}$, $W_0^{(1/2)}$ are learnable parameters of the transformation.

This stack of transformations effectively accumulates the normalized sum of the local neighborhood i.e. the neighborhood information for each concept in the graph. The self-connection ensures self-dependent feature transformation.

Decoder Module. DistMult factorization (Yang et al., 2014) is used as the scoring function. For a triplet (c_i, r, c_j) , the score s is obtained as follows:

$$s(c_i, r, c_j) = \sigma(\mathbf{h}_{c_i}^T R_r \mathbf{h}_{c_j}),$$

where σ is the logistic function; \mathbf{h}_{c_i} , $\mathbf{h}_{c_j} \in \mathbb{R}^d$ are the R-GCN encoded feature vectors for concepts c_i , c_j . Each relation $r \in \mathcal{R}$ is also associated with a diagonal matrix $R_r \in \mathbb{R}^{d \times d}$.

Training. We train our graph autoencoder model using negative sampling (Schlichtkrull et al., 2018). For triplets in $\hat{\mathcal{E}}'$ (positive samples), we create an equal number of negative samples by randomly corrupting the positive triplets. The corruption is performed by randomly modifying either one of the constituting concepts or the relation, creating the overall set of samples denoted by \mathcal{T} .

The task is set as a binary classification between the positive/negative triplets, where the model is trained with the standard cross-entropy loss:

$$\mathcal{L}_{\mathcal{G}'} = -\frac{1}{2|\hat{\mathcal{E}}'|} \sum_{(c_i, r, c_j, y) \in \mathcal{T}} (y \log s(c_i, r, c_j) + (1 - y) \log(1 - s(c_i, r, c_j))).$$

Once we train the autoencoder graph model, it will ensure that target domain-specific concepts (crucial for KG) can possibly be explained via domain-general concepts and further via inter-domain knowledge. In other words, the encoded node representations \mathbf{h}_i will capture commonsense graph information in the form of domain-specific and domain-general features and thus will be effective for the downstream task when there is a distributional shift during evaluation.

4.3 Step 2a) Commonsense Graph Feature Extraction

The trained graph autoencoder model as explained in the previous section §4.2, can be used for feature extraction. We now describe the methodology to extract the document-specific commonsense graph features for a particular document x :

- 1) The first step is to extract the set of all unique nouns, adjectives, and adverbs present in the document. We call this set \mathcal{W} .
- 2) Next, we extract a subgraph from \mathcal{G}' , where we take all triplets for which both the constituting nodes are either in \mathcal{W} or are within the vicinity of radius 1 of any of the words in \mathcal{W} . We call this graph $\mathcal{G}'_{\mathcal{W}}$.
- 3) We then make a forward pass of $\mathcal{G}'_{\mathcal{W}}$ through the encoder of the pre-trained graph autoencoder model. This results in feature vectors \mathbf{h}_j for all unique nodes j in $\mathcal{G}'_{\mathcal{W}}$.
- 4) Finally, we average over the feature vectors \mathbf{h}_j for all unique nodes in $\mathcal{G}'_{\mathcal{W}}$, to obtain the commonsense graph features \mathbf{x}_{cg} for document x .

We surmise that since most documents will have both domain-specific and domain-general words in \mathcal{W} , \mathbf{x}_{cg} will inherently capture the commonsense information likely to be helpful during domain adaptation.

4.4 Step 2b) Domain-adversarial Training

We feed the commonsense graph feature \mathbf{x}_{cg} pooled from $\mathcal{G}'_{\mathcal{W}}$ for document x (§4.3) into the DANN architecture (see §3.2). We proceed by learning an encoder function for the graph vector $\mathbf{z}_{grp} = M'_{\theta_G}(\mathbf{x}_{cg})$ and combine its representation with the DANN encoder $\mathbf{z}_{dann} = M_{\theta_M}(\mathbf{x})$ to get the final feature representation $[\mathbf{z}_{dann}; \mathbf{z}_{grp}]$, of the document x . Here, $[a; b]$ represents concatenation.

The task classifier C and domain-discriminator D_{adv} now takes this modified representation, $[\mathbf{z}_{dann}; \mathbf{z}_{grp}]$, as its input instead of only \mathbf{z}_{dann} . To further enforce domain-invariance into the encoded graph representation \mathbf{z}_{grp} , we consider it as a hidden code in a traditional autoencoder and consequently add a shared decoder D_{recon} (with parameters θ_R) with a reconstruction loss (mean-squared error):

$$\begin{aligned} \mathcal{L}_{recon}(X_s, X_t) &= \mathcal{L}_{recon}(X_s) + \mathcal{L}_{recon}(X_t), \\ \text{s.t. } \mathcal{L}_{recon} &= -\mathbb{E}_{\mathbf{x}_{cg}} (\|D_{recon}(\mathbf{z}_{grp}) - \mathbf{x}_{cg}\|_2^2). \end{aligned}$$

We hypothesize that if θ_R can reconstruct graph features for both domains, then it would ensure stronger domain invariance constraints in \mathbf{z}_{grp} . The final optimization of this domain-adversarial setup is based on the minimax objective:

$$\theta^* = \arg \min_{\theta_G, M, C, R} \max_{\theta_D} (\mathcal{L}_{cls} + \lambda \mathcal{L}_{adv_D} + \gamma \mathcal{L}_{recon}),$$

where λ and γ are hyper-parameters.

5 Experimental Setup

5.1 Dataset

We consider the Amazon-reviews benchmark dataset for domain adaptation in SA (Blitzer et al., 2007b). This corpus consists of Amazon product reviews and ranges across four domains: Books, DVDs, Electronics, and Kitchen appliances. Each review is associated with a rating denoting its sentiment polarity. Reviews with rating up to 3 stars are considered to contain negative sentiment and 4 or 5 stars as positive sentiment. The dataset follows a balanced distribution between both labels yielding $2k$ unlabelled training instances for each domain. Testing contains $3k - 6k$ samples for evaluation. We follow similar pre-processing as done by Ganin et al. (2016); Ruder and Plank (2018) where each review is encoded into a 5000-dimensional tfidf weighted bag-of-words (BOW) feature vector of unigrams and bigrams.

5.2 Training Details

We follow Ganin et al. (2016) in training our network. Our neural layers i.e., DANN encoder (M), graph feature encoder (M'), graph feature reconstructor (D_{recon}), task classifier (C) and domain discriminator (D_{adv}) are implemented with 100 dimensional fully connected layers. We use a cyclic λ as per (Ganin et al., 2016) and $\gamma = 1$ after validating with $\gamma \in \{0.5, 1, 2\}$. 25% dropout is used in

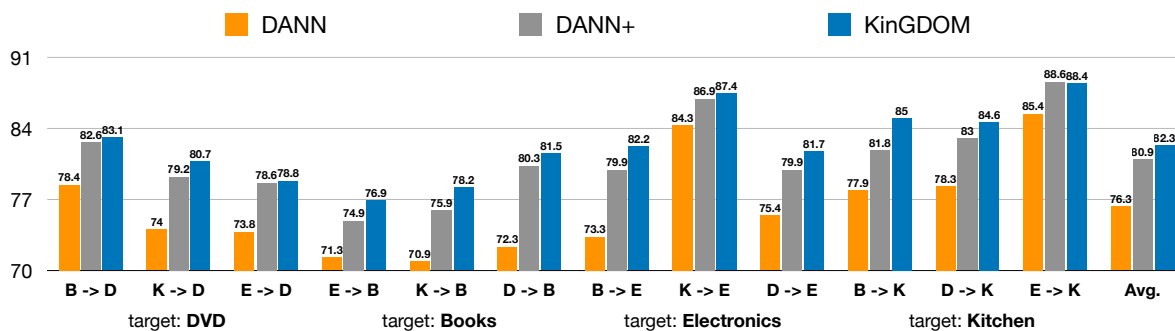


Figure 3: Results of DANN vs DANN+ vs *KinGDOM* across different target domains. Best viewed in colour.

the fully connected layers and the model is trained with Adam (Kingma and Ba, 2015) optimizer.

5.3 Baseline Methods

In this paper, to inspect the role of external commonsense knowledge and analyze the improvement in performance it brings, we intentionally use BOW features and compare them against other baseline models that also use BOW features. This issue has also been addressed by Poria et al. (2020). The flexibility of *KinGDOM* allows other approaches, such as mSDA, CNN, etc. to be easily incorporated in it, which we plan to analyze in the future.

We compare *KinGDOM* with the following unsupervised domain adaptation baseline methods: **DANN** (Ganin et al., 2016) is a domain-adversarial method, based on which we develop *KinGDOM* (§3.2); **DANN+** The DANN model where we use an Adam optimizer instead of the original SGD optimizer. The network architecture and the rest of the hyperparameters are kept same; Variational Fair Autoencoder (**VFAE**) (Louizos et al., 2015) learns latent representations independent from sensitive domain knowledge, while retaining enough task information by using a MMD-based loss; Central Moment Discrepancy (**CMD**) (Zellinger et al., 2017) is a regularization method which minimizes the difference between feature representations by utilizing equivalent representation of probability distributions by moment sequences; **Asym** (Saito et al., 2017) is the asymmetric tri-training framework that uses three neural networks asymmetrically for domain adaptation; **MT-Tri** (Ruder and Plank, 2018) is similar to **Asym**, but uses multi-task learning; Domain Separation Networks (**DSN**) (Bousmalis et al., 2016b) learns to extract shared and private components of each domain. As per Peng et al. (2018a), it stands as the present state-of-the-art method for unsupervised domain adaptation; Task

Refinement Learning (**TRL**) (Ziser and Reichart, 2019) Task Refinement Learning is an unsupervised domain adaptation framework which iteratively trains a Pivot Based Language Model to gradually increase the information exposed about each pivot; **TAT** (Liu et al., 2019) is the transferable adversarial training setup to generate examples which helps in modelling the domain shift. TAT adversarially trains classifiers to make consistent predictions over these transferable examples; **CoCMD** (Peng et al., 2018a) is a co-training method based on the CMD regularizer which trains a classifier on simultaneously extracted domain specific and invariant features. CoCOMD, however, is SSL-based as it uses labeled data from the target domain. Although it falls outside the regime of unsupervised domain adaptation, we report its results to provide a full picture to the reader.

6 Results and Analysis

As mentioned in §5.3, we reimplemented the baseline DANN model using Adam optimizer and observed that its results has been notably underreported in many of the unsupervised domain adaptation literature for sentiment analysis (see Table 2). In the original DANN implementation (Ganin et al., 2016), Stochastic Gradient Descent (SGD) was used as the optimizer. However, in DANN+, using Adam optimizer leads to substantial performance jump that outperforms many of the recent advanced domain adaptation methods – CMD (Zellinger et al., 2017), VFAE (Louizos et al., 2015), ASym (Saito et al., 2017), and MT-Tri (Ruder and Plank, 2018).

We compare the performance of *KinGDOM* with its base models – DANN and DANN+. As observed in Fig. 3, *KinGDOM* surpasses DANN+ by 1.4% which asserts the improvement in domain-invariance due to the incorporation of external com-

Source ↓ Target	DANN (5k)	DANN+ (5k)	VFAE (5k)	CMD (5k)	Asym (5k)	MT-Tri (5k)	TRL* (5k)	DSN (5k)	CoCMD* (5k)	KinGDOM (5k)	DANN+ (30k)	TAT (30k)	KinGDOM (30k)
B → D	78.4	82.6	79.9	80.5	80.7	81.2	82.2	82.8	83.1	83.1	84.7	84.5	85.0
B → E	73.3	79.9	79.2	78.7	79.8	78.0	-	81.9	83.0	82.2	83.0	80.1	83.9
B → K	77.9	81.8	81.6	81.3	82.5	78.8	82.7	84.4	85.3	85.0	84.0	83.6	86.6
D → B	72.3	80.3	75.5	79.5	73.2	77.1	-	80.1	81.8	81.4	82.7	81.9	82.7
D → E	75.4	79.9	78.6	79.7	77.0	81.0	-	81.4	83.4	81.7	83.4	81.9	83.9
D → K	78.3	83.0	82.2	83.0	82.5	79.5	-	83.3	85.5	84.6	85.3	84.0	87.1
E → B	71.3	74.9	72.7	74.4	73.2	73.5	-	75.1	76.9	76.9	77.1	83.2	78.4
E → D	73.8	78.6	76.5	76.3	72.9	75.4	75.8	77.1	78.3	78.8	79.6	77.9	80.3
E → K	85.4	88.6	85.0	86.0	86.9	87.2	-	87.2	87.3	88.4	89.0	90.0	89.4
K → B	70.9	75.9	72.0	75.6	72.5	73.8	72.1	76.4	77.2	78.2	77.1	75.8	80.0
K → D	74.0	79.2	73.3	77.5	74.9	77.8	-	78.0	79.6	80.7	81.3	77.7	82.3
K → E	84.3	86.9	83.8	85.4	84.6	86.0	-	86.7	87.2	87.4	88.0	88.2	88.6
Avg.	76.3	80.9	78.4	79.8	78.4	79.1	-	81.2	82.4	82.3	82.9	82.4	84.0

Table 2: Comparison with different baseline and state-of-the-art models (§5.3). TRL* reported results on four combinations. CoCMD* is a semi-supervised domain adaptation method. DSN is the current state-of-the-art for unsupervised domain adaptation on the Amazon reviews dataset. Scores for MT-Tri are extrapolated from the graphs illustrated in Ruder and Plank (2018). Note: **B**: Books, **D**: DVD, **E**:Electronics, and **K**: Kitchen domains. 5k, 30k signify 5000 and 30,000 dimensional BOW features.

monsense knowledge.

Next, we look at Table 2 where comparisons are made with other baselines, including the state-of-the-art DSN approach. As observed, *KinGDOM* outperforms DSN in all the task scenarios, indicating the efficacy of our approach. Blitzer et al. (2007b), in their original work, noted that domain transfer across the two groups of *DVD*, *Books* and *Electronics*, *Kitchen* is particularly challenging. Interestingly, in our results, we observe the highest gains when the source and target domains are from these separate groups (e.g., Kitchen → DVD, Kitchen → Books, Electronics → Books).

In Table 2, we also compare *KinGDOM* against CoCMD and TAT. Although CoCMD is a semi-supervised method, *KinGDOM* surpasses its performance in several of the twelve domain-pair combinations and matches its overall result without using any labelled samples from the target domain. TAT is the state-of-the-art method for unsupervised domain adaptation in the Amazon reviews dataset when used with 30,000 Bag-Of-Words (BOW) features. Interestingly, *KinGDOM* used with 5000 BOW features can match TAT with 30,000 BOW features and outperforms TAT by around 1.6% overall when used with the same 30,000 BOW features. The reimplementation of DANN – DANN+ with 30,000 BOW also surpasses the result of TAT by 0.5%. The results indicate that external knowledge,

when added to a simple architecture such as DANN, can surpass sophisticated state-of-the-art models, such as DSN and TAT. Our primary intention to utilize DANN as the base model is to highlight the role of knowledge base infusion in domain adaptation, devoid of sophisticated models, and complex neural maneuvering. Nevertheless, the flexibility of *KinGDOM* allows it to be associated with advanced models too (e.g., DSN, TAT), which we believe could perform even better. We intend to analyze this in the future.

6.1 Ablation Studies

We further analyze our framework and challenge our design choices. Specifically, we consider three variants of our architecture based on alternative ways to condition DANN with the graph features. Each of these variants reveals important clues regarding the invariance properties and task appropriateness of \mathbf{z}_{grp} . **Variation 1** denotes separate decoders D_{recon} for source and target domains. In **Variation 2**, domain classifier D_{adv} takes only \mathbf{z}_{dann} as input whereas the sentiment classifier C takes the concatenated feature $[\mathbf{z}_{dann}; \mathbf{z}_{grp}]$. Finally, in **Variation 3**, D_{adv} takes input $[\mathbf{z}_{dann}; \mathbf{z}_{grp}]$ whereas C only takes \mathbf{z}_{dann} . As seen in Fig. 4, all the variants perform worse than *KinGDOM*. For Variation 1, the performance drop indicates that having a shared decoder D_{recon} in *KinGDOM* facilitates




Figure 4: Average accuracy (%) on target domains across different variants defined in §6.1. Best viewed in colour.

learning invariant representations and helps target domain classification. For Variant 2, removal of \mathbf{z}_{grp} from domain classifier diminishes the domain-invariance capabilities, thus making the domain classifier stronger and leading to a drop in sentiment classification performance. For Variant 3, removal of \mathbf{z}_{grp} from sentiment classifier C degrades the performance. This indicates that in *KinGDOM*, \mathbf{z}_{grp} contain task appropriate features retrieved from external knowledge (see §1).

Besides ablations, we also look at alternatives to the knowledge graph and bag-of-words representation used for the documents. For the former, we consider replacing ConceptNet with WordNet (Fellbaum, 2010), which is a lexical knowledge graph with conceptual-semantic and lexical connections. We find the performance of *KinGDOM* with WordNet to be 1% worse than ConceptNet in terms of average accuracy score. This indicates the compatibility of ConceptNet with our framework. However, the competitive performance with WordNet also suggests the usability of our framework with any structural resource comprising inter-domain connections. For the latter, we use Glove-averaged embeddings with DANN. Glove is a popular word embedding method which captures semantics using co-occurrence statistics (Pennington et al., 2014). Results in Fig. 4 show that using only Glove does not provide the amount of conceptual semantics available in ConceptNet.

6.2 Case Studies

We delve further into our results and qualitatively analyze *KinGDOM*. We look at a particular test document from DVD domain, for which *KinGDOM* predicts the correct sentiment, both when the

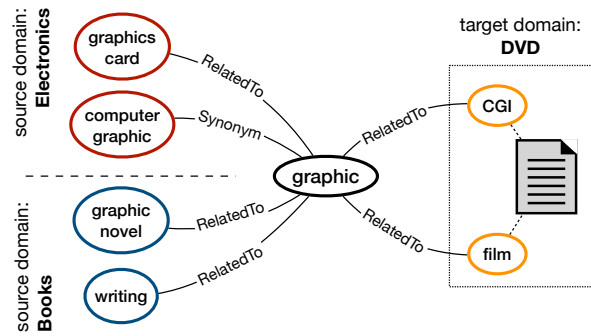


Figure 5: Domain-general term *graphic* bridges the common-sense knowledge between domain-specific terms in Electronics, Books and DVD.

source domain is Electronics and also Books. In similar settings, DANN mispredicts the same document. Looking at the corresponding document-specific sub-graph for this document, we observe conceptual links to both domain-general concepts and domain-specific concepts from the source domain. In Fig. 5, we can see the domain-specific terms CGI and film to be related to the general concept *graphic* which is further linked to domain-specific concepts like *graphics card*, *writing*, etc. from Electronics, Books, respectively. This example shows how *KinGDOM* might use these additional concepts to enhance the semantics as required for sentiment prediction.

7 Conclusion

In this paper, we explored the role of external commonsense knowledge for domain adaptation. We introduced a domain-adversarial framework called *KinGDOM*, which relies on an external commonsense KB (ConceptNet) to perform unsupervised domain adaptation. We showed that we can learn domain-invariant features for the concepts in the KB by using a graph convolutional autoencoder. Using the standard Amazon benchmark for domain adaption in sentiment analysis, we showed that our framework exceeds the performance of previously proposed methods for the same task. Our experiments demonstrate the usefulness of external knowledge for the task of cross-domain sentiment analysis. Our code is publicly available at <https://github.com/declare-lab/kingdom>.

Acknowledgments

This research is supported by A*STAR under its RIE 2020 Advanced Manufacturing and Engineering (AME) programmatic grant, Award No. - A19E2b0098.

References

- Basant Agarwal, Namita Mittal, Pooja Bansal, and Sonal Garg. 2015. [Sentiment analysis using common-sense and context information](#). *Comp. Int. and Neurosc.*, 2015:715730:1–715730:9.
- Firoj Alam, Shafiq R. Joty, and Muhammad Imran. 2018. [Domain adaptation with adversarial training and graph embeddings](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1077–1087.
- Alexandra Balahur, Jesús M. Hermida, and Andrés Montoyo. 2011. [Detecting implicit expressions of sentiment in text based on commonsense knowledge](#). In *Proceedings of the 2nd Workshop on Computational Approaches to Subjectivity and Sentiment Analysis, WASSA@ACL 2011, Portland, OR, USA, June 24, 2011*, pages 53–60. Association for Computational Linguistics.
- Somnath Banerjee. 2007. [Boosting inductive transfer for text classification using wikipedia](#). In *The Sixth International Conference on Machine Learning and Applications, ICMLA 2007, Cincinnati, Ohio, USA, 13-15 December 2007*, pages 148–153. IEEE Computer Society.
- Himanshu Sharad Bhatt, Deepali Semwal, and Shourya Roy. 2015. [An iterative similarity based adaptation technique for cross-domain text classification](#). In *Proceedings of the 19th Conference on Computational Natural Language Learning, CoNLL 2015, Beijing, China, July 30-31, 2015*, pages 52–61.
- Bin Bi, Chen Wu, Ming Yan, Wei Wang, Jiangnan Xia, and Chenliang Li. 2019. [Incorporating external knowledge into machine reading for generative question answering](#). 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 2521–2530, Hong Kong, China. Association for Computational Linguistics.
- John Blitzer, Mark Dredze, and Fernando Pereira. 2007a. [Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification](#). In *ACL 2007, Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, June 23-30, 2007, Prague, Czech Republic*.
- John Blitzer, Mark Dredze, and Fernando Pereira. 2007b. [Biographies, Bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 440–447, Prague, Czech Republic. Association for Computational Linguistics.
- Marina Boia, Claudiu Cristian Musat, and Boi Faltings. 2014. [Acquiring commonsense knowledge for sentiment analysis through human computation](#). In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada*, pages 901–907. AAAI Press.
- Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. 2016a. [Domain separation networks](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 343–351.
- Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. 2016b. [Domain separation networks](#). In *Advances in neural information processing systems*, pages 343–351.
- Erik Cambria, Soujanya Poria, Devamanyu Hazarika, and Kenneth Kwok. 2018. [Senticnet 5: Discovering conceptual primitives for sentiment analysis by means of context embeddings](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 1795–1802. AAAI Press.
- Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Michael I. Jordan. 2018. [Partial transfer learning with selective adversarial networks](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 2724–2732.
- Wei-Lun Chang, Hui-Po Wang, Wen-Hsiao Peng, and Wei-Chen Chiu. 2019. [All about structure: Adapting structural information across domains for boosting semantic segmentation](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 1900–1909.
- Minmin Chen, Kilian Q. Weinberger, and John Blitzer. 2011. [Co-training for domain adaptation](#). In *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain*, pages 2456–2464.
- Minmin Chen, Zhixiang Eddie Xu, Kilian Q. Weinberger, and Fei Sha. 2012. [Marginalized denoising autoencoders for domain adaptation](#). In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*.
- Quanyu Dai, Xiao Shen, Xiao-Ming Wu, and Dan Wang. 2019. [Network transfer learning via adversarial domain adaptation with graph convolution](#). *CoRR*, abs/1909.01541.

- Hal Daumé III and Daniel Marcu. 2006. [Domain adaptation for statistical classifiers](#). *J. Artif. Intell. Res.*, 26:101–126.
- Yang Deng, Ying Shen, Min Yang, Yaliang Li, Nan Du, Wei Fan, and Kai Lei. 2018. [Knowledge as A bridge: Improving cross-domain answer selection with external knowledge](#). In *Proceedings of the 27th International Conference on Computational Linguistics, COLING 2018, Santa Fe, New Mexico, USA, August 20-26, 2018*, pages 3295–3305.
- Hady Elsahar and Matthias Gallé. 2019. [To annotate or not? predicting performance drop under domain shift](#). 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 2163–2173, Hong Kong, China. Association for Computational Linguistics.
- Christiane Fellbaum. 2010. Wordnet. In *Theory and applications of ontology: computer applications*, pages 231–243. Springer.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. 2016. [Domain-adversarial training of neural networks](#). *Journal of Machine Learning Research*, 17(59):1–35.
- Deepanway Ghosal, Navonil Majumder, Soujanya Poria, Niyati Chhaya, and Alexander Gelbukh. 2019. [Dialoguecn: A graph convolutional neural network for emotion recognition in conversation](#). 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 154–164.
- Xavier Glorot, Antoine Bordes, and Yoshua Bengio. 2011. [Domain adaptation for large-scale sentiment classification: A deep learning approach](#). In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pages 513–520.
- William L. Hamilton, Kevin Clark, Jure Leskovec, and Dan Jurafsky. 2016a. [Inducing domain-specific sentiment lexicons from unlabeled corpora](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 595–605. The Association for Computational Linguistics.
- William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2016b. [Diachronic word embeddings reveal statistical laws of semantic change](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1501, Berlin, Germany. Association for Computational Linguistics.
- Yulan He and Deyu Zhou. 2011. [Self-training from labeled features for sentiment analysis](#). *Inf. Process. Manage.*, 47(4):606–616.
- Robert L. Logan IV, Nelson F. Liu, Matthew E. Peters, Matt Gardner, and Sameer Singh. 2019. [Barack’s wife hillary: Using knowledge graphs for fact-aware language modeling](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5962–5971. Association for Computational Linguistics.
- Jing Jiang and ChengXiang Zhai. 2007. [Instance weighting for domain adaptation in NLP](#). In *ACL 2007, Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, June 23-30, 2007, Prague, Czech Republic*.
- Prathusha K Sarma, Yingyu Liang, and William Sethares. 2019. [Shallow domain adaptive embeddings for sentiment analysis](#). 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 5548–5557, Hong Kong, China. Association for Computational Linguistics.
- Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. 2017a. [Learning to discover cross-domain relations with generative adversarial networks](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 1857–1865.
- Young-Bum Kim, Karl Stratos, and Dongchan Kim. 2017b. [Adversarial adaptation of synthetic or stale data](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1297–1307.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Wouter Marco Kouw and Marco Loog. 2019. A review of domain adaptation without target labels. *IEEE transactions on pattern analysis and machine intelligence*.
- Lianghao Li, Xiaoming Jin, and Mingsheng Long. 2012. [Topic correlation analysis for cross-domain text classification](#). In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada*. AAAI Press.
- Pengfei Li, Kezhi Mao, Xuefeng Yang, and Qi Li. 2019. [Improving relation extraction with knowledge-attention](#). 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 229–239, Hong Kong, China. Association for Computational Linguistics.

- Yitong Li, Timothy Baldwin, and Trevor Cohn. 2018. What’s in a domain? learning domain-robust text representations using adversarial training. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pages 474–479.
- Hong Liu, Mingsheng Long, Jianmin Wang, and Michael Jordan. 2019. Transferable adversarial training: A general approach to adapting deep classifiers. In *International Conference on Machine Learning*, pages 4013–4022.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2017. Adversarial multi-task learning for text classification. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1–10.
- Qi Liu, Yue Zhang, and Jiangming Liu. 2018. Learning domain representation for multi-domain sentiment classification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 541–550.
- zhibin liu, Zheng-Yu Niu, Hua Wu, and Haifeng Wang. 2019. Knowledge aware conversation generation with explainable reasoning over augmented graphs. 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 1782–1792, Hong Kong, China. Association for Computational Linguistics.
- Christos Louizos, Kevin Swersky, Yujia Li, Max Welling, and Richard Zemel. 2015. The variational fair autoencoder. *arXiv preprint arXiv:1511.00830*.
- Jingchao Ni, Shiyu Chang, Xiao Liu, Wei Cheng, Haifeng Chen, Dongkuan Xu, and Xiang Zhang. 2018. Co-regularized deep multi-network embedding. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France, April 23-27, 2018*, pages 469–478.
- Sinno Jialin Pan, Xiaochuan Ni, Jian-Tao Sun, Qiang Yang, and Zheng Chen. 2010. Cross-domain sentiment classification via spectral feature alignment. In *Proceedings of the 19th International Conference on World Wide Web, WWW 2010, Raleigh, North Carolina, USA, April 26-30, 2010*, pages 751–760.
- Minlong Peng, Qi Zhang, Yu-gang Jiang, and Xuanjing Huang. 2018a. Cross-domain sentiment classification with target domain specific information. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2505–2513.
- Minlong Peng, Qi Zhang, Yu-Gang Jiang, and Xuanjing Huang. 2018b. Cross-domain sentiment classification with target domain specific information. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2505–2513.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar; A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543. ACL.
- Matthew E. Peters, Mark Neumann, Robert Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A. Smith. 2019. Knowledge enhanced contextual word 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 43–54, Hong Kong, China. Association for Computational Linguistics.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, and Rada Mihalcea. 2020. Beneath the tip of the iceberg: Current challenges and new directions in sentiment analysis research. *arXiv preprint arXiv:2005*.
- Soujanya Poria, Navonil Majumder, Rada Mihalcea, and Eduard Hovy. 2019. Emotion recognition in conversation: Research challenges, datasets, and recent advances. *IEEE Access*, 7:100943–100953.
- Sebastian Ruder. 2019. *Neural Transfer Learning for Natural Language Processing*. Ph.D. thesis, NATIONAL UNIVERSITY OF IRELAND, GALWAY.
- Sebastian Ruder and Barbara Plank. 2018. Strong baselines for neural semi-supervised learning under domain shift. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1044–1054. Association for Computational Linguistics.
- Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. 2017. Asymmetric tri-training for unsupervised domain adaptation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2988–2997. JMLR. org.
- Prathusha K. Sarma, Yingyu Liang, and Bill Sethares. 2018. Domain adapted word embeddings for improved sentiment classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 2: Short Papers*, pages 37–42.
- Michael Sejr Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max

- Welling. 2018. [Modeling relational data with graph convolutional networks](#). In *The Semantic Web - 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3-7, 2018, Proceedings*, volume 10843 of *Lecture Notes in Computer Science*, pages 593–607. Springer.
- Raksha Sharma, Pushpak Bhattacharyya, Sandipan Dandapat, and Himanshu Sharad Bhatt. 2018. [Identifying transferable information across domains for cross-domain sentiment classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 968–978.
- Xiao Shen and Fu-Lai Chung. 2019. [Network embedding for cross-network node classification](#). *CoRR*, abs/1901.07264.
- Bei Shi, Zihao Fu, Lidong Bing, and Wai Lam. 2018. [Learning domain-sensitive and sentiment-aware word embeddings](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2494–2504. Association for Computational Linguistics.
- Robert Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. 2017. [Adversarial discriminative domain adaptation](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 2962–2971. IEEE Computer Society.
- Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. 2014. [Deep domain confusion: Maximizing for domain invariance](#). *CoRR*, abs/1412.3474.
- Pu Wang, Carlotta Domeniconi, and Jian Hu. 2008. [Using wikipedia for co-clustering based cross-domain text classification](#). In *Proceedings of the 8th IEEE International Conference on Data Mining (ICDM 2008), December 15-19, 2008, Pisa, Italy*, pages 1085–1090. IEEE Computer Society.
- Garrett Wilson and Diane J. Cook. 2018. [Adversarial transfer learning](#). *CoRR*, abs/1812.02849.
- Evan Wei Xiang, Bin Cao, Derek Hao Hu, and Qiang Yang. 2010. [Bridging domains using world wide knowledge for transfer learning](#). *IEEE Trans. Knowl. Data Eng.*, 22(6):770–783.
- Linchuan Xu, Xiaokai Wei, Jiannong Cao, and Philip S. Yu. 2017. [Embedding of embedding \(EOE\): joint embedding for coupled heterogeneous networks](#). In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM 2017, Cambridge, United Kingdom, February 6-10, 2017*, pages 741–749. ACM.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2014. [Embedding entities and relations for learning and inference in knowledge bases](#). *arXiv preprint arXiv:1412.6575*.
- Pengcheng Yang, Lei Li, Fuli Luo, Tianyu Liu, and Xu Sun. 2019. [Enhancing topic-to-essay generation with external commonsense knowledge](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2002–2012, Florence, Italy. Association for Computational Linguistics.
- Werner Zellinger, Thomas Grubinger, Edwin Lughofer, Thomas Natschläger, and Susanne Saminger-Platz. 2017. [Central moment discrepancy \(cmd\) for domain-invariant representation learning](#). *arXiv preprint arXiv:1702.08811*.
- Lei Zhang, Shuai Wang, and Bing Liu. 2018. [Deep learning for sentiment analysis: A survey](#). *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.*, 8(4).
- Peixiang Zhong, Di Wang, and Chunyan Miao. 2019. [Knowledge-enriched transformer for emotion detection in textual conversations](#). 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 165–176, Hong Kong, China. Association for Computational Linguistics.
- Yftah Ziser and Roi Reichart. 2019. [Task refinement learning for improved accuracy and stability of unsupervised domain adaptation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5895–5906.