

# KNOWLEDGE ADAPTATION: TEACHING TO ADAPT

**Sebastian Ruder**

Insight Centre for Data Analytics  
National University of Ireland, Galway  
sebastian.ruder@insight-centre.org

**Parsa Ghaffari**

Aylien Ltd.  
Dublin, Ireland  
parsa@aylien.com

**John G. Breslin**

Insight Centre for Data Analytics  
National University of Ireland, Galway  
john.breslin@insight-centre.org

## ABSTRACT

Domain adaptation is crucial in many real-world applications where the distribution of the training data differs from the distribution of the test data. Previous Deep Learning-based approaches to domain adaptation need to be trained jointly on source and target domain data and are therefore unappealing in scenarios where models need to be adapted to a large number of domains or where a domain is evolving, e.g. spam detection where attackers continuously change their tactics.

To fill this gap, we propose *Knowledge Adaptation*, an extension of *Knowledge Distillation* (Bucilua et al., 2006; Hinton et al., 2015) to the domain adaptation scenario. We show how a student model achieves state-of-the-art results on unsupervised domain adaptation from multiple sources on a standard sentiment analysis benchmark by taking into account the domain-specific expertise of multiple teachers and the similarities between their domains.

When learning from a single teacher, using domain similarity to gauge trustworthiness is inadequate. To this end, we propose a simple metric that correlates well with the teacher’s accuracy in the target domain. We demonstrate that incorporating high-confidence examples selected by this metric enables the student model to achieve state-of-the-art performance in the single-source scenario.

## 1 INTRODUCTION

In many real-world applications such as sentiment classification (Pang & Lee, 2008), a model trained on one domain may not work well when directly applied to another domain due to the difference in the data distribution between the domains. At the same time, labeled data in new domains is scarce or non-existent and manual labeling of large amounts of target domain data is expensive. Domain adaptation allows models to reduce the domain discrepancy and adapt to new domains. While fine-tuning is a commonly used method for supervised domain adaptation, there is no cheap equivalent in the unsupervised case as existing Deep Learning-based approaches need to be trained jointly on source and target domain data. This is prohibitive in scenarios with a large number of domains, such as sentiment classification on the plethora of real-world review categories, blog types, or communities (Hamilton et al., 2016). Additionally, re-training a model on source data is unfeasible for evolving domains, such as spam detection where attackers continuously adapt their strategy, scene classification where the scene changes over time (Hoffman et al., 2014), or a conversational agent for a user with a rapidly evolving style, such as a child or second language learner.

Rather than re-training, we would like to be able to leverage our trained model in the source domain to inform the predictions of a new model trained on the target domain. This objective aligns organically with the idea of *Knowledge Distillation* (Bucilua et al., 2006; Hinton et al., 2015), which we extend as *Knowledge Adaptation* to the domain adaptation scenario. While Knowledge Distillation concentrates on training a *student* model on the predictions of a (possibly larger) *teacher* model, Knowledge Adaptation focuses on determining *what part* of the teacher’s expertise can be trusted and applied to the target domain.

In this context, determining when to trust the teacher is key. This circumstance is paralleled in real-world teacher-student and adviser-advisee relationships: Children learn early on to trust familiar advisers but to moderate that trust depending on the adviser’s recent history of accuracy or inaccuracy (Corriveau & Harris, 2009), while adults may surround themselves with advisers, e.g. to make a financial investment and gradually learn whose expertise to trust (Johnson & Grayson, 2005).

We demonstrate how domain similarity metrics can be used as a measure of relative trust in a teacher for unsupervised domain adaptation with multiple source domains and show state-of-the-art results for a student model that learns from multiple domain-specific teachers.

When learning from a single teacher in the single-source scenario, using a general measure of domain similarity is inadequate as the student has no other, more relevant teacher to turn to for advice in case its teacher is untrustworthy. To this end, we propose a simple measure, which correlates well with the teacher’s accuracy in the target domain and allows the student to gauge the teacher’s confidence in its predictions. We demonstrate that by incorporating high-confidence examples selected by this metric in the training process, the student model is able to outperform the state-of-the-art in single-source unsupervised domain adaptation.

Crucially, our models are the first Deep Learning-based models for domain adaptation that perform adaptation without expensive re-training on the source domain data. They are thus able to make use of readily available trained source domain models and are particularly apt for scenarios where domains change or occur in large numbers.

## 2 RELATED WORK

**Distilling knowledge.** Bucilua et al. (2006) first proposed a method to compress the knowledge of a source model, which was later improved by Hinton et al. (2015). Romero et al. (2015) showed how this method can be adapted to train deep and thin models, while Kim & Rush (2016) apply the technique to sequence-level models. In addition, Hu et al. (2016) use it to constrain a student model with logic rules. Our goal differs from the previous methods due to the difference in data distributions between source and target data, which necessitates to learn from the teacher’s knowledge *only insofar* as it is useful for the target domain. Similar in spirit to Knowledge Distillation is the KL-divergence based objective by Yu et al. (2013) and Li et al. (2014) for adapting an acoustic model and the Adaptive Mixture of Experts model (Nowlan & Hinton, 1990), which also learns which expert to trust for a given example. Both, however, require labeled examples, which are scarce for domain adaptation, while our model is entirely unsupervised.

**Domain adaptation.** Domain adaptation has a long history of research: Blitzer et al. (2006) proposed a structural correspondence learning algorithm. Daumé III (2007) introduced a kernel function that maps source and target domain data to a space that encourages in-domain similarity, while Pan et al. (2010) proposed a spectral feature alignment algorithm to align domain-specific words into meaningful clusters, while Long & Wang (2015) use multi-task learning to avoid negative transfer.

**Deep learning-based domain adaptation.** Deep learning-based approaches to domain adaptation are more recent and have focused mainly on learning domain-invariant representations: Glorot et al. (2011) first employed stacked Denoising Auto-encoders (SDA) to extract meaningful representations. Chen et al. (2012) in turn extended SDA to marginalized SDA by addressing SDA’s high computational cost and lack of scalability to high-dimensional features, while Zhuang et al. (2015) proposed to use deep auto-encoders for transfer learning. Ajakan et al. (2016) added a Gradient Reversal Layer that hinders the model’s ability to discriminate between domains. Finally, Zhou et al. (2016) transferred the source examples to the target domain and vice versa using Bi-Transferring Deep Neural Networks, while Bousmalis et al. (2016) propose Domain Separation Networks that employ domain-specific and general-domain encoders. All of these approaches, however, require to jointly train the model on source and target data for every new target domain.

**Domain adaptation from multiple sources.** For domain adaptation from multiple sources, Mansour (2009) proposed a distribution weighted hypothesis with theoretical guarantees. Duan et al. (2009) proposed a method to learn a least-squares SVM classifier by leveraging source classifiers, while Chattopadhyay et al. (2012) assign pseudo-labels to the target data. Finally, Wu & Huang (2016) exploit general sentiment knowledge and word-level sentiment polarity relations for multi-source domain adaptation.

### 3 KNOWLEDGE ADAPTATION

#### 3.1 PROBLEM DEFINITION

In the following, we describe domain adaptation within the knowledge adaptation framework: We are provided with one or multiple source domains  $\mathcal{D}_{S_i}$  and a target domain  $\mathcal{D}_T$ . For each of the source domains, we are provided with a teacher model  $T_i$  that was trained on examples  $X_{S_i} = \{x_1^{S_i}, \dots, x_n^{S_i}\}$  and their labels  $\{y_1^{S_i}, \dots, y_n^{S_i}\}$  from  $\mathcal{D}_{S_i}$ . In the target domain  $\mathcal{D}_T$ , we only have access to the examples  $\{x_1^T, \dots, x_n^T\}$  without knowledge of their labels. Note that we omit source and target domain indexes in the following for simplicity in cases where examples are unambiguous. Our task is now to train a student model  $S$  that performs well on unseen examples from the target domain  $\mathcal{D}_T$ .

#### 3.2 SINGLE TEACHER-STUDENT MODEL

Our teacher and student models are simple multilayer perceptrons (MLP). The basic MLP consists of an input layer, one or multiple intermediate layers, and an output layer. Each intermediate layer  $\ell$  learns to embed the output of the previous layer  $x$  into a latent representation  $h_\ell = f_\ell(W_\ell x + b_\ell)$  where  $W_\ell$  and  $b_\ell$  are the weights and bias of the  $\ell^{th}$  layer, while  $f_\ell$  is the activation, typically ReLU  $f_\ell(x) = \max\{0, x\}$  for hidden layers and softmax units  $f_\ell(x) = \text{softmax}(x) = e^x / \sum_{i=1}^{|x|} e^{x_i}$  for the output layer.

In the single source setting, the teacher  $T$  has an output softmax  $P_T = \text{softmax}(z_T)$  where  $z_T$  are the logits of the teacher’s output layer.  $T$  is trained to minimize the loss  $\mathcal{L}_T = \mathcal{H}(y_i, P_T)$  where  $\mathcal{H}$  refers to the cross-entropy and  $y_i$  is the label of the  $i^{th}$  training example in the source domain  $\mathcal{D}_S$ .

The student  $S$  similarly models an output probability  $P_S = \text{softmax}(z_S)$  where  $z_S$  are the logits of the student’s output layer. In the context of knowledge distillation (Hinton et al., 2015), the student  $S$  is trained so that its output  $P_S$  is similar to the teacher’s output  $P_T$  and to the true labels. In practice, the output probability of the teacher is smoothed with a temperature  $\tau$  to soften the signal and provide more information during training. The same temperature  $\tau$  is applied to the output of the student network for the comparison:

$$P_T^\tau = \text{softmax}\left(\frac{z_T}{\tau}\right), \quad P_S^\tau = \text{softmax}\left(\frac{z_S}{\tau}\right). \quad (1)$$

For unsupervised domain adaptation, true labels in the target domain  $\mathcal{D}_T$  are not available. Thus the student  $S$  is trained solely to mimic the teacher’s softened output with the following loss, which is similar to treating source input modalities as privileged information (Lopez-Paz et al., 2016):

$$\mathcal{L}_S = \mathcal{H}(P_T^\tau, P_S^\tau). \quad (2)$$

#### 3.3 MULTIPLE TEACHER-STUDENT MODEL

The teacher-student paradigm lends itself naturally to the scenario with multiple source domains. Intuitively, the trust that a student should place in a teacher should be proportional to the degree of similarity between the teacher’s domain and the student’s domain.

To this end, we consider three measures of domain similarity, which have been successfully used in domain adaptation research: Jensen-Shannon divergence (Remus, 2012) and Renyi divergence (Van Asch & Daelemans, 2010), which are both based on Kullback-Leibler divergence and are computed with regard to the domains’ term distributions; and Maximum Mean Discrepancy (Tzeng et al., 2014), which we compute with respect to the teacher’s latent representation. These measures are computed between the target domain  $\mathcal{D}_T$  and every source domain  $\mathcal{D}_S$  (additional information with regard to our choice and use of domain similarity measures can be found in the appendix A.1).

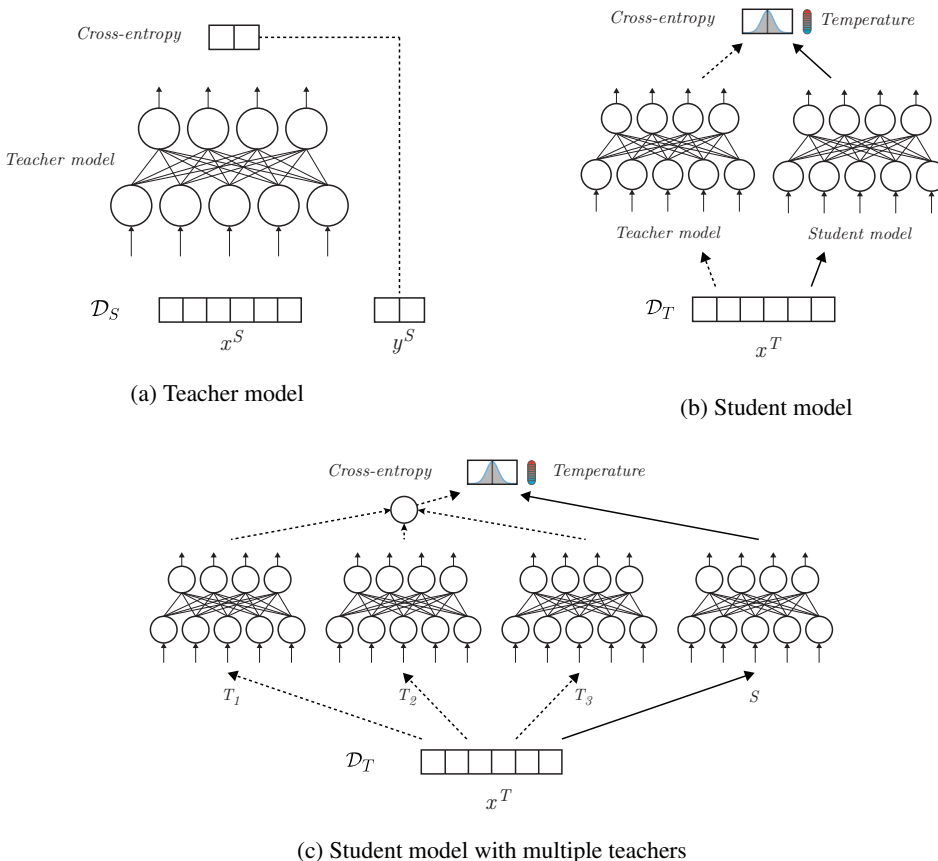


Figure 1: Training procedures for a) the teacher model, b) the student model, and c) the student model with multiple teachers. The teacher is trained on examples  $x^S$  and their true labels  $y^S$  in the source domain  $\mathcal{D}_S$ , while the student is trained on the softened predictions of one or multiple teachers of examples  $x^T$  in the target domain  $\mathcal{D}_T$ .

The student model with multiple teachers is then trained to imitate the sum of the teacher’s individual predictions weighted with the normalized similarity  $sim(\mathcal{D}_S, \mathcal{D}_T)$  of their respective source domain  $\mathcal{D}_S$  to the target domain  $\mathcal{D}_T$ :

$$\mathcal{L}_{MUL} = \mathcal{H}\left(\sum_{i=1} sim(\mathcal{D}_{S_i}, \mathcal{D}_T) \cdot P_{T_i}^T, P_S^T\right). \tag{3}$$

### 3.4 LEVERAGING A SINGLE TEACHER’S KNOWLEDGE

General measures of domain similarity are useful in the multi-source setting, where we can rely on multiple teachers and choose to trust one more than the others. In the scenario with a single teacher, it is not helpful to know whether we can trust the teacher *in general*. We rather want a measure that allows us to determine if we can trust the teacher for a *specific example*.

To arrive at such a measure, we revisit the representations the teacher learns from the input data: In order to make accurate predictions, the teacher model learns to separate the representation of different output classes in its hidden representation (we use a one-layer MLP in our experiments as detailed in §4.2; in deeper networks, this would be an intermediate layer). Even though the teacher model is trained on the source domain, this separation still holds – albeit with decreased accuracy – in the target domain. This can be seen in Figure 2, where examples in the target domain that were predicted as positive and negative by the teacher form distinct clusters (refer to §4.1 for details with regard to the data and task). Importantly, many of these predictions are incorrect.

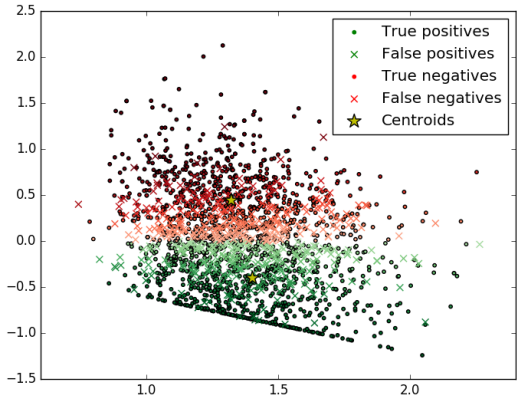


Figure 2: PCA visualization of a teacher’s latent representations of target domain examples for the K->D domain pair (see §4.1 for details). A darker color reflects a higher MCD value. Best viewed in close-up.




Figure 3: Accuracy of the teacher’s predictions on the top  $n$  target domain examples with the highest MCD value for the K->D domain pair.

As evidenced in Figure 2, incorrect predictions are frequent along the decision boundary and infrequent along the cluster edges, where examples are less ambiguous. More precisely, the accuracy of the teacher’s predictions on the target domain is proportional to the absolute difference in similarity of the teacher’s representation  $h$  with the cluster centroids, which we refer to as Maximum Cluster Difference (MCD) and define as follows:

$$MCD_h = |\cos(c_p, h) - \cos(c_n, h)| \tag{4}$$

where  $c_p$  and  $c_n$  are the centroids of the positive and negative cluster respectively as predicted by the teacher, i.e. the mean representation of all examples assigned to the cluster by the teacher. Note that while we are focusing on binary classification involving two clusters, the measure is equally applicable to the multi-class setting, as demonstrated in Appendix A.2.

Evidence of the efficacy of this measure for obtaining the trustworthiness of a teacher for an example can be found in the PCA visualization<sup>1</sup> in Figure 2, where incorrect predictions are far less common for (more darkly colored) examples with higher MCD values. Additionally, the MCD score of a target domain example and the accuracy of the teacher’s prediction correlate with an average Pearson’s  $r$  of 0.33 and  $p < 0.05$  across all domain pairs of the data described in §4.1. We furthermore plot the teacher’s accuracy for the top  $n$  target domain examples with the highest MCD values in Figure 3. While the measure becomes less accurate as  $n$  increases, it is very accurate for low  $n$ .

For this reason, rather than weighing all examples with MCD, we propose to add  $n$  unlabeled training examples with the highest MCD with their teacher-assigned label as pseudo-supervised examples on which we train the student with the following objective:

$$\mathcal{L}_S = \mathcal{H}((1 - \lambda) \cdot y_{\text{teacher}} + \lambda P_T^T, P_S^T) \tag{5}$$

where  $y_{\text{teacher}}$  is the indicator array containing 1 at the index  $\text{argmax}(P_T)$  and 0 at all other indexes, while  $\lambda$  determines the contribution of the soft targets. This can be seen as a representation-based variant of instance adaptation (Jiang & Zhai, 2007), which uses MCD as a measure of confidence as it correlates better with teacher accuracy than teacher prediction probability. In practice, we alternate unsupervised training with the objective in equation 2 and pseudo-supervised training with the objective in equation 5, although other curricula are imaginable.

<sup>1</sup>A visualization using t-SNE revealed the same cluster. However, PCA showed a clearer decision boundary.

## 4 EXPERIMENTS

### 4.1 DATA SET

We use the Amazon product reviews sentiment analysis dataset of Blitzer et al. (2006), a common benchmark for domain adaptation. The dataset consists of 4 different domains: Book (B), DVDs (D), Electronics (E) and Kitchen (K). We follow the conventions of past work and evaluate on the binary classification task where reviews with more than 3 stars are considered positive and reviews with 3 stars or fewer are considered negative. Each domains contains 1,000 positive, 1,000 negative, and approximately 4,000 unlabeled reviews. For fairness of comparison, we use the raw bag-of-words unigram/bigram features pre-processed with tf-idf as input (Blitzer et al., 2006).

For single-source adaptation, we replicate the set-up of previous methods and train our teacher models on all 2,000 labeled examples, of which we reserve 200 as dev set. For domain adaptation from multiple sources, we follow the conventions of Bollegala et al. (2011) and limit the total number of training examples for all teachers to 1,600, i.e. given three source domains, each teacher is only trained on about 533 labeled samples. We also train a general teacher on the same 1,600 examples of the three domains. In both scenarios, the student is evaluated on all 2,000 labeled samples of the target domain. As we have not found a universally applicable way to optimize hyperparameters or perform early stopping for unsupervised domain adaptation, we choose to use a small number of unlabeled examples as a labeled validation set similar to (Bousmalis et al., 2016).

### 4.2 HYPERPARAMETERS

Both student and teacher models are one-layer MLPs with 1,000 hidden dimensions. We use a vocabulary size of 10,000, a temperature of 5, a batch size of 10, and Adam (Kingma & Ba, 2015) as optimizer with a learning rate of 0.001. For every experiment, we report the average of 10 runs.

### 4.3 DOMAIN ADAPTATION FROM MULTIPLE SOURCES

As it is easier for the student to assign trust when learning from multiple teachers, we first conduct experiments on the sentiment analysis benchmark for domain adaptation from multiple sources. For each experiment, one of the four domains is used as the target domain, while the remaining ones are treated as source domains.

**Domain similarity.** We first evaluate the performance of our student depending on different measures of domain similarity, with which we interpolate the predictions of the teachers. As evidenced in Table 2 provided in the appendix, Jensen-Shannon divergence generally performs best. We thus use this measure for the remainder of the experiments.

**Our models.** For multi-source domain adaptation, we first consider a teacher-only baseline (Teacher-only), where teacher sentiment probabilities are combined, weighted with Jensen-Shannon divergence, and the most likely sentiment is chosen. We further train our student on a) the source domain-specific teachers as detailed in §3.3, b) the general teacher trained on all source domains as described in §4.1, and on c) the combination of source domain and general teachers.

**Comparison models.** We compare our models against the following methods: domain adaptation with structural correspondence learning (SCL) (Blitzer et al., 2006); domain adaptation based on spectral feature alignment (SFA) (Pan et al., 2010); adaptations of SCL and SFA via majority voting to the multi-source scenario (SCL-com and SFA-com); cross-domain sentiment classification by constructing a sentiment-sensitive thesaurus (SST) (Bollegala et al., 2011); multiple-domain sentiment analysis by identifying domain dependent/independent word polarity (IDDIWP) (Yoshida et al., 2011); three general-purpose multiple source domain adaptation methods (DWHC, Mansour (2009)), (DAM, Duan et al. (2009)), (CP-MDA, Chattopadhyay et al. (2012)); cross-domain sentiment classification by transferring sentiment along a sentiment graph with hinge loss and logistic loss respectively (SDAMS-SVM and SDAMS-Log) (Wu & Huang, 2016). Numbers are used as reported by Wu & Huang (2016).

**Results.** All results are depicted in Table 1. Evaluating the combination of the source teacher models directly on the target domain (Teacher-only) produces the worst results, which underscores the need for methods that allow adaptation to the target domain. Training the student model on the soft targets

	<b>Book</b>	<b>DVD</b>	<b>Electronics</b>	<b>Kitchen</b>
SCL (Blitzer et al., 2006)	0.7457	0.7630	0.7893	0.8207
SFA (Pan et al., 2010)	0.7598	0.7848	0.7808	0.8210
SCL-com	0.7523	0.7675	0.7918	0.8247
SFA-com	0.7629	0.7869	0.7864	0.8258
SST (Bollegala et al., 2011)	0.7632	0.7877	0.8363	0.8518
IDDIWP (Yoshida et al., 2011)	0.7524	0.7732	0.8167	0.8383
DWHC (Mansour, 2009)	0.7611	0.7821	0.8312	0.8478
DAM (Duan et al., 2009)	0.7563	0.7756	0.8284	0.8419
CP-MDA (Chattopadhyay et al., 2012)	0.7597	0.7792	0.8331	0.8465
SDAMS-SVM (Wu & Huang, 2016)	0.7786	0.7902	<b>0.8418</b>	0.8578
SDAMS-Log (Wu & Huang, 2016)	0.7829	0.7913	0.8406	0.8629
Teacher-only	0.7565	0.7765	0.7960	0.8210
Student (source teachers)	0.7918	0.7968	0.8203	0.8523
Student (general teacher)	<b>0.8014</b>	0.8062	0.8365	<b>0.8675</b>
Student (source teachers + general)	0.8010	<b>0.8088</b>	0.8311	0.8647

Table 1: Average results for domain adaptation from multiple sources for the comparison models and ours on the sentiment analysis benchmark. For the results in each column, the domain in the column header is used as target domain and the remaining three domains are used as source domains.

of the teachers allows us to improve upon the teacher-only baseline significantly, thereby demonstrating the appropriateness of the teacher-student paradigm to the domain adaptation scenario. The student model outperforms comparison methods that rely on source model predictions by combining (Mansour, 2009) or predicting (Duan et al., 2009) them. This showcases the usefulness of learning from soft targets in the domain adaptation scenario. Training on a general teacher model as well as on a combination of the general teacher and the source domain teachers allows us to improve results even further. Both models improve over existing approaches to domain adaptation from multiple sources and outperform approaches that rely on sentiment analysis-specific information (Wu & Huang, 2016) in all but the electronics domain.

#### 4.4 SINGLE-SOURCE DOMAIN ADAPTATION

We additionally evaluate the ability of the student to only learn from a single teacher. This scenario is more challenging as the student cannot consider other teachers that might provide more relevant predictions. For each target domain, each of the three other domains is used as source domain, yielding 12 domain pairs.

**Our models.** On these domain pairs, we firstly evaluate our student-teacher (TS) model. For training a model that incorporates high-confidence predictions of the teacher (TS-MCD), we cross-validate the interpolation parameter  $\lambda$  in equation 5 and the number of examples with the highest MCD scores  $n$ . We find that a low  $\lambda$  (around 0.2) generally yields the best results in the domain adaptation setting, as the high-confidence predictions are helpful to guide the student’s learning during training. Additionally, using the top 500 unlabeled target domain examples with the highest MCD scores for pseudo-supervised training of the student produces the best results.

**Comparison models.** For the single-source case, we similarly compare against SCL (Blitzer et al., 2006) and SFA (Pan et al., 2010), as well as against multi-label consensus training (MCT), which combines base classifiers trained with SCL (Li & Zong, 2008) and against an approach that links heterogeneous input features with points via non-negative matrix factorization (PNMF) (Zhou et al., 2015). We additionally compare against the following deep learning-based approaches: stacked denoising auto-encoders (SDA) (Glorot et al., 2011); marginalized SDA (mSDA) (Chen et al., 2012); transfer learning with deep auto-encoders (TLDA) (Zhuang et al., 2015); and bi-transferring deep neural networks (BTDNN) (Zhou et al., 2016). Numbers are used as reported by Zhou et al. (2016).

**Results.** The results can be seen in Figure 4. The student trained on the source domain teacher (TS) achieves convincing results and outperforms the state-of-the-art on three domain pairs – twice with the Book domain as source domain, showing that knowledge acquired from the Book domain

might perhaps be more easily transferable to a student model. For many domain pairs, the student still falls significantly short compared to the performance of the state-of-the-art, which highlights that solely relying on a single teacher’s predictions is insufficient to bridge the discrepancy between the domains. Instead, additional methods are necessary to provide evidence for the student when to trust the teacher’s predictions. Leveraging the teacher’s knowledge by incorporating high-confidence examples selected by MCD into the training (TS-MCD) improves the performance of the student in almost all cases significantly. This allows the student to outperform the state-of-the-art on 8 out of 12 domain pairs without expensive joint training on source and target data and with the sole dependence of a single model trained on the source domain, which is typically readily available.

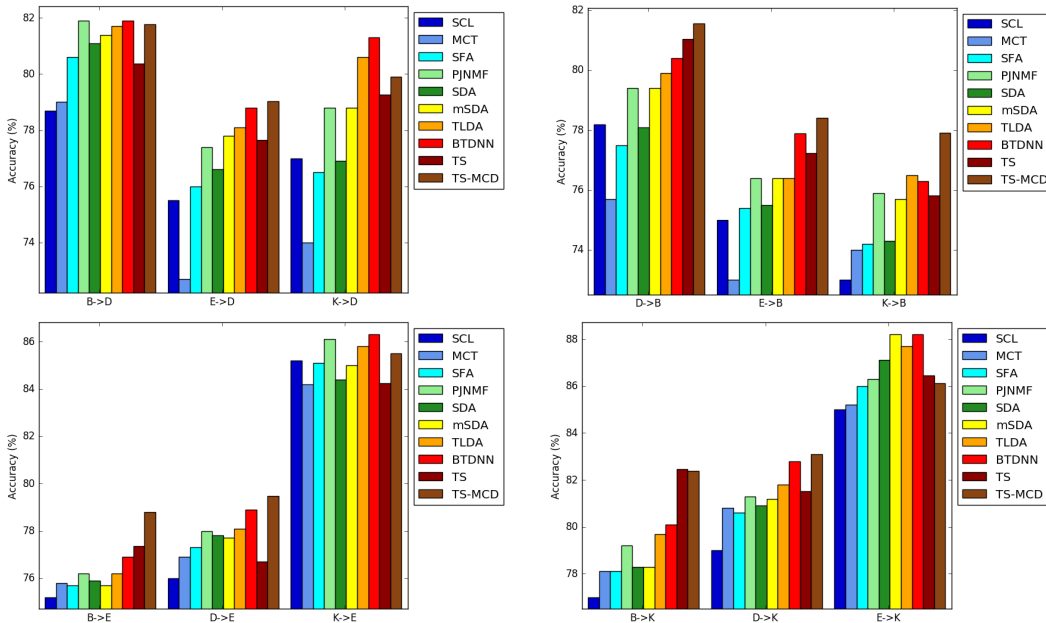


Figure 4: Average results for single-source domain adaptation for the comparison models and our models on the sentiment analysis benchmark. B: Book. D: DVD. E: Electronics. K: Kitchen.

## 5 CONCLUSION

In this work, we have proposed Knowledge Adaptation, an extension of the Knowledge Distillation idea to the domain adaptation scenario. This method – in contrast to prevalent domain adaptation methods – is able to perform adaptation without re-training. We firstly demonstrated the benefit of this paradigm by showing that a student model that takes into account the predictions of multiple teachers and their domain similarities is able to outperform the state-of-the-art for multi-source unsupervised domain adaptation on a standard sentiment analysis benchmark. We additionally introduced a simple measure to gauge the trustworthiness of a single teacher and showed how this measure can be used to achieve state-of-the-art results on 8 out of 12 domain pairs for single-source unsupervised domain adaptation.

## ACKNOWLEDGMENTS

We thank John Glover and Chris Hokamp for fruitful discussions. This publication has emanated from research conducted with the financial support of the Irish Research Council (IRC) under Grant Number EBPPG/2014/30 and with Aylien Ltd. as Enterprise Partner as well as from research supported by a research grant from Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289.



## REFERENCES

- Hana Ajakan, Hugo Larochelle, Mario Marchand, and Victor Lempitsky. Domain-Adversarial Training of Neural Networks. *Journal of Machine Learning Research*, 17:1–35, 2016. doi: 10.1088/1475-7516/2015/08/013.
- Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. *Advances in Neural Information Processing Systems*, 19:137–144, 2007. ISSN 10495258.
- John Blitzer, Ryan McDonald, and Fernando Pereira. Domain Adaptation with Structural Correspondence Learning. *EMNLP '06 Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, (July):120–128, 2006. doi: 10.3115/1610075.1610094.
- John Blitzer, Mark Dredze, and Fernando Pereira. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. *Annual Meeting-Association for Computational Linguistics*, 45(1):440, 2007. ISSN 0736587X. doi: 10.1109/IRPS.2011.5784441.
- Danushka Bollegala, David Weir, and John Carroll. Using Multiple Sources to Construct a Sentiment Sensitive Thesaurus for Cross-Domain Sentiment Classification. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pp. 132–141, 2011.
- Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain Separation Networks. *NIPS*, 2016.
- Cristian Bucilua, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '06*, pp. 535, 2006. doi: 10.1145/1150402.1150464. URL <http://portal.acm.org/citation.cfm?doid=1150402.1150464>.
- Rita Chattopadhyay, Qian Sun, Jieping Ye, Sethuraman Panchanathan, W E I Fan, and I A N Davidson. Multi-Source Domain Adaptation and Its Application to Early Detection of Fatigue. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(4), 2012.
- Minmin Chen, Zhixiang Xu, Kilian Q. Weinberger, and Fei Sha. Marginalized Denoising Autoencoders for Domain Adaptation. *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pp. 767–774, 2012. ISSN 0960-3174. doi: 10.1007/s11222-007-9033-z.
- Kathleen Corriveau and Paul L Harris. Choosing your informant: weighing familiarity and recent accuracy. *Developmental science*, 12(3):426–437, 2009.
- Hal Daumé III. Frustratingly Easy Domain Adaptation. *Association for Computational Linguistic (ACL)s*, (June):256–263, 2007. ISSN 0736587X. doi: 10.1.1.110.2062. URL <https://arxiv.org/pdf/0907.1815.pdf>.
- Lixin Duan, Ivor W. Tsang, Dong Xu, and Tat-Seng Chua. Domain Adaptation from Multiple Sources via Auxiliary Classifiers. In *Proceedings of the 26th Annual International Conference on Machine Learning*, 2009.
- Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Domain Adaptation for Large-Scale Sentiment Classification: A Deep Learning Approach. *Proceedings of the 28th International Conference on Machine Learning*, (1):513–520, 2011. URL [http://www.icml-2011.org/papers/342{}\\_icmlpaper.pdf](http://www.icml-2011.org/papers/342{}_icmlpaper.pdf).
- William L. Hamilton, Kevin Clark, Jure Leskovec, and Dan Jurafsky. Inducing Domain-Specific Sentiment Lexicons from Unlabeled Corpora. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 2016. URL <http://arxiv.org/abs/1606.02820>.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the Knowledge in a Neural Network. *arXiv preprint arXiv:1503.02531*, pp. 1–9, 2015. ISSN 0022-2488. doi: 10.1063/1.4931082. URL <http://arxiv.org/abs/1503.02531>.

- Judy Hoffman, Trevor Darrell, and Kate Saenko. Continuous manifold based adaptation for evolving visual domains. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 867–874, 2014. ISSN 10636919. doi: 10.1109/CVPR.2014.116.
- Zhiting Hu, Xuezhe Ma, Zhengzhong Liu, Eduard Hovy, and Eric Xing. Harnessing Deep Neural Networks with Logic Rules. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pp. 1–18, 2016. URL <http://arxiv.org/abs/1603.06318>.
- Jing Jiang and ChengXiang Zhai. Instance Weighting for Domain Adaptation in NLP. *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, (October):264–271, 2007. ISSN 0736-587X. doi: 10.1145/1273496.1273558. URL <http://aclanthology.info/papers/instance-weighting-for-domain-adaptation-in-nlp>.
- Devon Johnson and Kent Grayson. Cognitive and Affective Trust in Service Relationships. *Journal of Business research*, 58(4):500–507, 2005. doi: 10.1016/S0148-2963(03)00140-1.
- Yoon Kim and Alexander M Rush. Sequence-Level Knowledge Distillation. *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP-16)*, 2016.
- Diederik P. Kingma and Jimmy Lei Ba. Adam: a Method for Stochastic Optimization. *International Conference on Learning Representations*, pp. 1–13, 2015.
- Jinyu Li, Rui Zhao, Jui Ting Huang, and Yifan Gong. Learning small-size DNN with output-distribution-based criteria. *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, (September):1910–1914, 2014. ISSN 19909772.
- Shoushan Li and Chengqing Zong. Multi-domain Adaptation for Sentiment Classification: Using Multiple Classifier Combining Methods. In *International Conference on Natural Language Processing and Knowledge Engineering (NLP-KE'08)*. IEEE, 2008. ISBN 9781424427802.
- Mingsheng Long and Jianmin Wang. Learning Multiple Tasks with Deep Relationship Networks. *Arxiv*, pp. 1–9, 2015. URL <http://arxiv.org/abs/1506.02117>.
- David Lopez-Paz, Léon Bottou, Bernhard Schölkopf, and Vladimir Vapnik. Unifying distillation and privileged information. *ICLR*, 2016. URL <http://arxiv.org/abs/1511.03643>.
- Yishay Mansour. Domain Adaptation with Multiple Sources. *NIPS*, 2009.
- Steven J. Nowlan and Geoffrey E. Hinton. Evaluation of Adaptive Mixture of Competing Experts. In *NIPS*, 1990.
- Sinno Jialin Pan, Xiaochuan Ni, Jian-tao Sun, Qiang Yang, and Zheng Chen. Cross-Domain Sentiment Classification via Spectral Feature Alignment. In *Proceedings of the 19th International Conference on World Wide Web*, pp. 751–760, 2010. ISBN 9781605587998.
- Bo Pang and Lillian Lee. Opinion Mining and Sentiment Analysis. *Foundations and trends in information retrieval*, 2(1-2):1–135, 2008. ISSN 1554-0669. doi: 10.1561/1500000001.
- Robert Remus. Domain adaptation using Domain Similarity- and Domain Complexity-based Instance Selection for Cross-Domain Sentiment Analysis. In *IEEE ICDM SENTIRE-2012*, 2012. URL [http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=6406510](http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=6406510).
- Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for Thin Deep Nets. *ICLR*, pp. 1–13, 2015. URL <http://arxiv.org/pdf/1412.6550.pdf>.
- Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep Domain Confusion: Maximizing for Domain Invariance. *CoRR*, 2014. URL <https://arxiv.org/pdf/1412.3474.pdf>.
- Vincent Van Asch and Walter Daelemans. Using Domain Similarity for Performance Estimation. *Computational Linguistics*, (July):31–36, 2010. ISSN 9781932432800. URL <http://eprints.pascal-network.org/archive/00007014/>.

- Fangzhao Wu and Yongfeng Huang. Sentiment Domain Adaptation with Multiple Sources. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL 2016)*, pp. 301–310, 2016. URL <https://pdfs.semanticscholar.org/09f0/885d1727a0b82300e94856e0be2f2f72561c.pdf>.
- Yasuhisa Yoshida, Tsutomu Hirao, Tomoharu Iwata, Masaaki Nagata, and Yuji Matsumoto. Transfer Learning for Multiple-Domain Sentiment Analysis - Identifying Domain Dependent/Independent Word Polarity. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence Transfer*, pp. 1286–1291, 2011.
- Dong Yu, Kaisheng Yao, Hang Su, Gang Li, and Frank Seide. KL-divergence regularized deep neural network adaptation for improved large vocabulary speech recognition. *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pp. 7893–7897, 2013. ISSN 15206149. doi: 10.1109/ICASSP.2013.6639201.
- Guangyou Zhou, Tingting He, Wensheng Wu, and Xiaohua Tony Hu. Linking Heterogeneous Input Features with Pivots for Domain Adaptation. *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pp. 1419–1425, 2015.
- Guangyou Zhou, Zhiwen Xie, Jimmy Xiangji Huang, and Tingting He. Bi-Transferring Deep Neural Networks for Domain Adaptation. *ACL*, pp. 322–332, 2016. URL <https://www.aclweb.org/anthology/P/P16/P16-1031.pdf>.
- Fuzhen Zhuang, Xiaohu Cheng, Ping Luo, Sinno Jialin Pan, and Qing He. Supervised Representation Learning: Transfer Learning with Deep Autoencoders. *IJCAI International Joint Conference on Artificial Intelligence*, pp. 4119–4125, 2015. ISSN 10450823.

## A APPENDIX

### A.1 DOMAIN SIMILARITY MEASURES

We use three measures of domain similarity in our experiments: Jensen-Shannon divergence, Renyi divergence, and Maximum Mean Discrepancy (MMD).

Jensen-Shannon divergence is a smoothed, symmetric variant of KL divergence. The Jensen-Shannon divergence between two different probability distributions  $P$  and  $Q$  can be written as:

$$D_{JS}(P||Q) = \frac{1}{2}[D_{KL}(P||M) + D_{KL}(Q||M)] \quad (6)$$

where  $M = \frac{1}{2}(P + Q)$ , i.e. the average distribution of  $P$  and  $Q$ , and  $D_{KL}$  is the KL divergence:

$$D_{KL}(P||Q) = \sum_{i=1}^n p_i \frac{p_i}{q_i}. \quad (7)$$

Renyi divergence similarly generalizes KL divergence by assigning different weights to the probability distributions of the source and target domain and is defined as follows:

$$D_R(P||Q) = \frac{1}{\alpha - 1} \log\left(\sum_{i=1}^n \frac{p_i^\alpha}{q_i^{\alpha-1}}\right). \quad (8)$$

If  $\alpha = 1$ , Renyi divergence reduces to KL divergence. In our experiments, we set  $\alpha = 0.99$  following Van Asch & Daelemans (2010).

These domain similarity measures are typically based on the term distributions of the source and target domains, i.e. the probability distribution  $P$  of a domain is the term distribution  $t \in \mathcal{R}^{|V| \times 1}$  where  $t_i$  is the relative probability of word  $w_i$  appearing in the domain and  $|V|$  is the size of the vocabulary of the domain. The intuition behind using term distributions is that similar domains

	<b>Book</b>	<b>DVD</b>	<b>Electronics</b>	<b>Kitchen</b>
None	0.7821	0.7913	0.8181	<b>0.8529</b>
Renyi divergence	0.7722	0.7727	0.8133	0.8420
Maximum Mean Discrepancy	0.7811	0.7839	0.7890	0.8273
Jensen-Shannon divergence	<b>0.7918</b>	<b>0.7968</b>	<b>0.8203</b>	0.8523

Table 2: Comparison of the impact of different domain similarity measures on the student’s performance when used for interpolating the predictions of the source domain teacher models. For the results in each column, the domain in the column header is used as target domain and the remaining three domains are used as source domains.

usually have more terms in common than dissimilar domains. While term distributions are efficient to compute and have proven effective in previous work (Van Asch & Daelemans, 2010; Wu & Huang, 2016), they only capture shallow occurrence statistics.

Another form of similarity metrics such as MMD are based on representations. MMD measures the distance between a source and target distribution with respect to a particular representation  $\phi$ . The MMD between the source data  $X_S$  and the target data  $X_T$  is defined as follows:

$$MMD(X_S, X_T) = \left\| \frac{1}{|X_S|} \sum_{x_S \in X_S} \phi(x_S) - \frac{1}{|X_T|} \sum_{x_T \in X_T} \phi(x_T) \right\|. \quad (9)$$

The representation  $\phi$  is usually obtained by embedding the source data and target data in a Reproducing Kernel Hilbert Space via a specifically chosen kernel, e.g. Bousmalis et al. (2016) use a linear combination of RBF kernels. Similarly to Tzeng et al. (2014), we use the hidden representation of a neural network as basis for  $\phi$ , as we are interested in how well the teacher’s representation captures difference in domain.

In our experiments, MMD does not outperform the more traditional term distribution-based similarity measure, which we attribute to two reasons: 1) Due to the limited amount of data, our teacher model is not deep enough to capture the difference in domain in its single hidden layer; Tzeng et al. (2014) in contrast identify the fully-connected layer *fc7* in the AlexNet architecture as the layer minimizing MMD. 2) The teacher is only trained on the source domain data. Its representation is thus not sensitive to detect the domain shift to the target domain. Training a separate model to minimize MMD alleviates this, but incurs additional computational costs and requires retraining on the source data during adaptation, which we set out to avoid to enable efficient adaptation.

Another commonly used measure of domain similarity is  $\mathcal{A}$ -distance. Ben-David et al. (2007) show that computing the  $\mathcal{A}$ -distance between two domains reduces to minimizing the empirical risk of a classifier that tries to discriminate between the examples in those domains. Previous work (Blitzer et al., 2007) uses the Huber loss and a linear classifier for computing the  $\mathcal{A}$ -distance. In our experiments,  $\mathcal{A}$ -distance did not outperform Jensen-Shannon divergence, while its reliance on training a classifier is a downside in our scenario with multiple or changing target domains, where we would prefer more efficient measures of domain similarity.

## A.2 MULTI-CLASS MCD

Maximum cluster difference can be easily extended to the multi-class setting. For  $n$  classes, we compute  $n$  cluster centroids for the clusters whose members have been assigned the same class by the model. We then create a set  $C$  containing all  $n(n - 1)/2$  unique pairs of cluster centroids. Finally, we compute the sum of pair-wise differences of the model’s representation  $h$  with regard to the cluster centroid pairs:

$$MCD_{\text{multi}} = \sum_{c_1, c_2 \in C} |\cos(c_1, h) - \cos(c_2, h)|. \quad (10)$$