

Unsupervised Adversarial Domain Adaptation for Implicit Discourse Relation Classification

Hsin-Ping Huang

Department of Computer Science
The University of Texas at Austin
hsinping@cs.utexas.edu

Junyi Jessy Li

Department of Linguistics
The University of Texas at Austin
jessy@austin.utexas.edu

Abstract

Implicit discourse relations are not only more challenging to classify, but also to annotate, than their explicit counterparts. We tackle situations where training data for implicit relations are lacking, and exploit domain adaptation from explicit relations (Ji et al., 2015). We present an unsupervised adversarial domain adaptive network equipped with a reconstruction component. Our system outperforms prior works and other adversarial benchmarks for unsupervised domain adaptation. Additionally, we extend our system to take advantage of labeled data if some are available.

1 Introduction

Discourse relations capture the relationship between units of text—e.g., sentences and clauses—and are an important aspect of text coherence. While some relations are expressed explicitly with a discourse connective (e.g., “for example”, “however”), relations are equally often expressed implicitly without an explicit connective (Prasad et al., 2008); in these cases, the relation needs to be inferred.

Resources for implicit discourse relations are scarce compared to the explicit ones, since they are harder to annotate (Miltsakaki et al., 2004). For example, among corpora annotated with discourse relations such as Arabic (Al-Saif and Markert, 2010), Czech (Poláková et al., 2013), Chinese (Zhou and Xue, 2015), English (Prasad et al., 2008), Hindi (Oza et al., 2009), and Turkish (Zeyrek et al., 2013), only the Chinese, English and Hindi corpora include implicit discourse relations (Prasad et al., 2014). In this low-resource scenario, Ji et al. (2015) proposed training with explicit relations via unsupervised domain adaptation, viewing explicit relations as a source domain with labeled training data, and implicit relations

as a target domain with no labeled data. The domain gap between explicit and implicit relations is acknowledged by prior observations that the two types of discourse relations are linguistically dissimilar (Sporleder and Lascarides, 2008; Rutherford and Xue, 2015).

We present a new system for the unsupervised domain adaptation setup on the Penn Discourse Treebank (Prasad et al., 2008). Our system is based on Adversarial Discriminative Domain Adaptation (Tzeng et al., 2017), which decouples source domain training and representation mapping between source and target. We improve this framework by proposing a reconstruction component to preserve the discriminability of target features, and incorporating techniques for stabler training on textual data.

Experimental results show that even with a simple architecture for representation learning, our unsupervised domain adaptation system outperforms prior work by 1.4-2.3 macro F1, with substantial improvements on Temporal and Contingency relations. It is also superior to DANN (Ganin et al., 2016), an adversarial framework widely used in NLP (Chen et al., 2018; Gui et al., 2017; Zhang et al., 2017; Fu et al., 2017; Joty et al., 2017; Xu and Yang, 2017), by 5.7 macro F1.

Finally, we extend the system to incorporate in-domain supervision as it is sometimes feasible resource-wise to build a seed corpus that may not be large enough to train a fully supervised system. We simulate this scenario by enabling the system to jointly optimize over a varying number of labeled examples of implicit relations. Our system consistently outperforms two strong baselines.

2 Related Work

Sporleder and Lascarides (2008) and Rutherford

and Xue (2015) observed that explicit and implicit relations are linguistically dissimilar, warranting an unsupervised domain adaptation approach in Ji et al. (2015). They used a marginalized denoising autoencoder to obtain generalized feature representations across the source and target domains with a linear SVM as the classification model. Our system improves upon this work using an adversarial network; we further generalize our network to semi-supervised settings.

To supplement the training data of implicit discourse relations, prior works have used weak supervision from sentences with discourse connectives (Marcu and Echihiabi, 2002; Sporleder and Lascarides, 2008; Braud and Denis, 2014; Ji et al., 2015), by analyzing connectives (Zhou et al., 2010a,b; Biran and McKeown, 2013; Rutherford and Xue, 2015; Braud and Denis, 2016; Wu et al., 2017), using a multi-task framework with other corpora (Lan et al., 2013; Liu et al., 2016; Lan et al., 2017), or utilizing cross-lingual data (Wu et al., 2016; Shi et al., 2017). The important distinction between this work and the research above is that these are supervised systems that used all of the annotated *implicit* annotation from PDTB during training, while exploring non-PDTB corpora for additional, noisy discourse cues; on the contrary, our main goal is to assume no labeled training data for implicit discourse relations.

Unsupervised domain adaptation with adversarial networks has become popular in recent years; this type of approach generates a representation for the target domain with the goal that the discriminator unable to distinguish between the source and target domains. Prior works proposed both generative approaches (Liu and Tuzel, 2016; Bousmalis et al., 2017; Sankaranarayanan et al., 2018; Russo et al., 2018) and discriminative approaches (Ganin et al., 2016; Tzeng et al., 2015, 2017). The discriminative DANN algorithm from Ganin et al. (2016) is frequently used in NLP tasks (Chen et al., 2018; Gui et al., 2017; Zhang et al., 2017; Fu et al., 2017; Joty et al., 2017; Xu and Yang, 2017). Our method builds upon Adversarial Discriminative Domain Adaptation (Tzeng et al., 2017), shown to outperform DANN in visual domain adaptation but has not been used in NLP tasks. The key differences between the two are discussed in Section 3.

Qin et al. (2017) adopted adversarial strategies to supervised implicit discourse classification.

They train an adversarial model using implicit discourse relations with and without expert-inserted connectives. Note again that theirs is a fully supervised system using signals in addition to the implicit relation annotations themselves, while our main focus is unsupervised domain adaptation that does not train on implicit relations.

3 Model Architecture

To classify discourse relations, our system takes a pair of sentence arguments x as input, and outputs the discourse relation y between these two arguments. With unsupervised domain adaptation, we have examples (X_s, Y_s) from the source domain, i.e., explicit discourse relations, and unlabeled examples (X_t) from a target domain, i.e., implicit discourse relations.

We use ADDA (Tzeng et al., 2017) as our underlying framework for domain adaptation. ADDA first learns a discriminative representation for the classification task in the source domain, then learns a representation for the target domain that mimics the distribution of the source domain. The key insight here is asymmetric mapping, where the target representation is “updated” until it matches with the source, a process more similar to the original Generative Adversarial Networks (Goodfellow et al., 2014) than joint training as in DANN (Ganin et al., 2016). Intuitively, since ADDA learns distinct feature encoders for the source and target domains instead of using a shared encoder, the same network doesn’t have to handle instances from different domains.

Summarized in Figure 1, we first pre-train a source encoder M_s and source classifier C (Section 3.1), then train the target encoder M_t (initialized with M_s) and discriminator D in an adversarial way, to minimize the domain discrepancy distance between the target representation distribution $M_t(X_t)$ and that of the source $M_s(X_s)$ (Section 3.2). Eventually, the target feature space is trained to match the source, and the source classifier C can be directly used on the target domain.

3.1 Base encoder and classifier

The source and target **encoders** M_s and M_t follow the same architecture; M_t is initialized to be M_s during adaptation. The encoders encode relation arguments into latent representations, and then feeds the representations into a **classifier** C to predict the discourse relation.

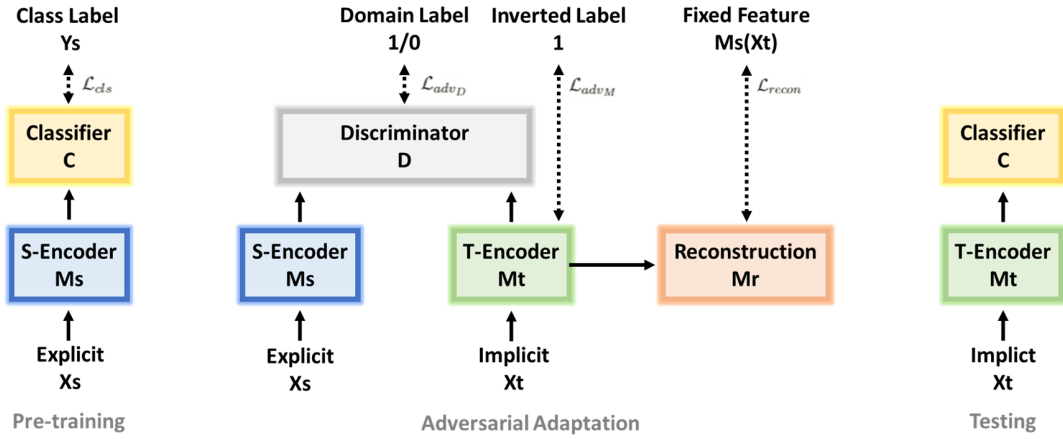


Figure 1: The framework of our proposed adversarial domain adaptation model, containing the pre-training stage, the adversarial adaptation stage, and the testing stage. The dashed box shows the supervised component.

Encoder The encoder generates a representation for each argument with an inner-attention Bidirectional LSTM (Yang et al., 2016) shared between the two arguments. Then, the representations of the two arguments are concatenated to form the final representation, shown in Figure 2.

Specifically, we encode each word in an argument into its word embeddings, which are fed into a BiLSTM, to get the hidden representations z_i using a fully-connected layer W_c on top of the concatenated hidden states $h_i = [\vec{h}_i, \overleftarrow{h}_i]$. We then apply an attention mechanism to induce a distribution of weights over all tokens in the argument; the final argument representation Arg is a weighted sum of z_i based on the attention weights α_i :

$$\begin{aligned}
 z_i &= W_c h_i + b_c \\
 u_i &= \tanh(W_w h_i + b_w) \\
 \alpha_i &= \frac{\exp(u_i^T u_w)}{\sum_i \exp(u_i^T u_w)} \\
 Arg &= \sum_i \alpha_i z_i
 \end{aligned} \quad (1)$$

Where W_c, b_c, W_w, b_w, u_w are model parameters.

Classifier The classifier consists of a single fully-connected layer on top of the encoder, finished with a softmax classification layer.

The source encoder M_s and the classifier C are trained using a supervised loss:

$$\begin{aligned}
 \min_{M_s, C} \mathcal{L}_{cls}(X_s, Y_s) = \\
 \mathbb{E}_{(x_s, y_s)} - \sum_k \mathbb{1}[k = y_s] \log C(M_s(x_s))
 \end{aligned} \quad (2)$$




Figure 2: Neural structure of inner-attention BiLSTM to encode relation argument pairs.

3.2 Unsupervised adversarial domain adaptation

We then learn a target encoder M_t to generate features for the target data which can be classified with classifier C , without assuming labels Y_t in the target domain. This is achieved by training a domain discriminator D , which classifies whether a feature is from the source or the target domain, and the target encoder M_t , that produces features similar to the source domain features and tries to fool the discriminator to predict the incorrect domain label.

The discriminator D is optimized according to a standard supervised loss:

$$\begin{aligned}
 \min_D \mathcal{L}_{advD}(X_s, X_t, M_s, M_t) = \\
 - \mathbb{E}_{x_s} [\log D(M_s(x_s))] \\
 - \mathbb{E}_{x_t} [\log (1 - D(M_t(x_t)))]
 \end{aligned} \quad (3)$$

D consists of two fully-connected layers on top of the encoder, finished with a softmax classification layer.

The target encoder M_t is optimized according to a standard GAN loss with inverted labels:

$$\min_{M_t} \mathcal{L}_{adv_M}(X_s, X_t, D) = -\mathbb{E}_{x_t}[\log D(M_t(x_t))] \quad (4)$$

Spectral normalization To stabilize the training of the discriminator, we employ spectral normalization, a weight normalization technique (Miyato et al., 2018), which controls the Lipschitz constant of the discriminator function by constraining the spectral norm of each layer. Spectral normalization is easy to implement without tuning any hyper-parameters and has a small additional computational cost.

Label smoothing We utilize label smoothing (Szegedy et al., 2016) to regularize the classifier during pre-training, which prevents the largest logit from becoming much larger than all others, and therefore prevents overfitting and makes the classifier, trained in the source domain, more adaptable.

For a source domain training example x_s with ground-truth label y_s and ground-truth distribution $q(k|x_s)$, the classifier computes the classification probability over relation classes as $p(k|x_s)$ for $k \in \{1 \dots K\}$. With label smoothing, we replace the ground-truth label distribution $q(k|x_s)$ in the standard cross-entropy loss as a linear combination of $q(k|x_s)$ and a uniform distribution over classes $u(k) = 1/K$.

$$\min_{M_s, C} \mathcal{L}_{cls}(X_s, Y_s) = -\sum_k q'(k) \log(p(k))$$

$$q'(k|x_s) = (1 - \epsilon)\mathbb{1}[k = y_s] + \epsilon/K \quad (5)$$

Reconstruction loss In order to classify the target representations using the source classifier, the target encoder is trained to produce representations that mimic the source domain representations in the adversarial training stage. Since there is no supervised loss applied in this stage, the target encoder may lose its ability to produce discriminative features that are helpful during classification. We propose a reconstruction loss to preserve the discriminability of the target encoder when adversarially adapting its features.

Since we initialize the target encoder with the source encoder, the initial representation (before domain adaptation) of a target instance x_t is the




Figure 3: The reconstruction loss component augmenting our unsupervised adversarial domain adaptor.

representation of *target* instances produced by the *source* encoder $M_s(x_t)$ (which is then fixed). After training, $M_t(x_t)$ adapts to the source domain and becomes dissimilar to $M_s(x_t)$. The reconstruction loss encourages the target encoder to produce features that can be reconstructed back to $M_s(x_t)$ (Figure 3).

For a target example x_t , we learn a reconstruction mapping M_r that maps the target representation $M_t(x_t)$ to $M_s(x_t)$:

$$x_t \rightarrow M_t(x_t) \rightarrow M_r(M_t(x_t)) \approx M_s(x_t) \quad (6)$$

The target encoder M_t and the reconstruction mapping M_r are optimized jointly with a reconstruction loss:

$$\min_{M_t, M_r} \mathcal{L}_{recon}(X_t, M_s) = \mathbb{E}_{x_t}[\|M_r(M_t(x_t)) - M_s(x_t)\|_2^2] \quad (7)$$

M_r consists of three fully-connected layers on top of the encoder.

Unsupervised objective For unsupervised domain adaptation, our full objective is:

$$\mathcal{L}^{unsup}(X_s, Y_s, X_t, M_s, M_t, D) = \min_{M_s, C} \mathcal{L}_{cls}(X_s, Y_s) + \min_D \mathcal{L}_{adv_D}(X_s, X_t, M_s, M_t) + \min_{M_t} \mathcal{L}_{adv_M}(X_s, X_t, D) + \min_{M_t, M_r} \mathcal{L}_{recon}(X_t, M_s) \quad (8)$$

3.3 Training

Summarized in Algorithm 1, the training procedure consists of three stages: pre-training, adversarial adaptation, and testing. During pre-training, we train the source encoder M_s and the classifier C according to Eq.(2). In the adversarial adaptation stage, we alternately train the discriminator D , target encoder M_t , and reconstruction mapping

Algorithm 1: Adversarial Adaptation

Input: explicit sentences with labels $\{x_s, y_s\}$
implicit sentences without labels $\{x_t\}$

Notations: source encoder M_s , classifier C ,
target encoder M_t ,
reconstruction mapping M_r ,
domain discriminator D

- 1 Train M_s, C through Eq.(2) with $\{x_s, y_s\}$;
 - 2 Initialize M_t as M_s ;
 - 3 **Repeat**
 - 4 Train D through Eq.(3) with $\{x_s, x_t\}$ and train M_t
through Eq.(4) with $\{x_t\}$;
 - 5 Train M_t, M_r through Eq.(7) with $\{x_t\}$;
- Output:** M_t and C for relation prediction
-

M_r according to Eq.(3), Eq.(4), Eq.(7). Finally, we test the model using the target encoder M_t and classifier C . The steps (lines 4 and 5) in the Repeat loop execute once in one iteration, and we optimized the model in two-step units.

4 Unsupervised Domain Adaptation Experiments

We first evaluate our model for the default task: unsupervised domain adaptation from explicit discourse relations to implicit discourse relations.

4.1 Settings

Data We train and test our model on the PDTB, following the experimental setup of Ji et al. (2015). The test examples are implicit relation instances from PDTB sections 21-22. The explicit training set consists of explicit examples from sections 02-20 and 23-24, and the explicit development set consists of the explicit examples from sections 00-01. The implicit training set and the implicit development set consist of examples from the same sections as the explicit sets. Evaluation is done for the first-level relations—Temporal, Comparison, Contingency, and Expansion. Table 1 summarizes the statistics of the four top-level implicit and explicit discourse relations in the PDTB.

Training Details We early-stopped training for both stages before total convergence if the macro F1 on the development set does not improve. During pre-training (Line 1 in Algorithm 1), we train and validate the model on the explicit training and development set. Early stopping happened after around 20 epochs. During adversarial adaptation (lines 4 and 5 in Algorithm 1), we train the model on the explicit and implicit training sets *without relation labels* Y_t , and validate on the implicit de-

Relation	Explicit		Implicit		
	Train	Dev	Train	Dev	Test
Temporal	2904	288	704	68	54
Contingency	2792	181	3622	276	287
Comparison	4674	366	2104	146	191
Expansion	5342	450	7394	556	651

Table 1: The number of examples of the four top level discourse relations in PDTB 2.0.

velopment set¹. Early stopping happened after around 5 epochs (with lines 4 and 5 executed once in each epoch).

Model configuration The hyperparameters, as well as the number of fully connected layers for the classifier C , discriminator D and the reconstruction mapping M_r , are all set according to the performance on the development sets. We first set the hyper-parameters of the encoders M_s, M_t and classifier C based on development performance during the pre-training stage. Then, we set the hyper-parameters of D and M_r based on development performance of the adaptation stage.

We use GloVe (Pennington et al., 2014) for word embeddings with dimension 300. The maximum argument length is set to 80. The encoder contains an inner-attention BiLSTM with dimension 50, producing a representation with dimension 200 for each example. The discriminator D consists of 2 hidden layers with 200 and 200 neurons on each layer. The reconstruction mapping M_r contains 3 hidden layers with 120, 15 and 120 neurons on each layer. The label smoothing parameter ϵ is 0.1. We use Adam (Kingma and Ba, 2015) with learning rate $1e-4$ for the base encoder and classifier, and $1e-6$ for the adversarial domain adapter. We use SGD optimizer with learning rate $1e-2$ for the reconstruction component. All the models were implemented using PyTorch (Paszke et al., 2017) and adapted from Conneau et al. (2017).

4.2 Systems

We experiment with three settings:

Implicit \rightarrow **Implicit** A supervised implicit discourse relation classifier using the base encoder and classifier, optimizing the standard cross-entropy loss, using the full implicit training and

¹Using a development set in the target domain is common in unsupervised domain adaptation (Ganin et al., 2016; Liu and Tuzel, 2016; Tzeng et al., 2017; Bousmalis et al., 2017; Russo et al., 2018)

	Temporal	Contingency	Comparison	Expansion	Macro F_1
Implicit \rightarrow Implicit	25.53	41.02	30.35	65.38	40.57
Explicit \rightarrow Implicit	22.22	22.35	23.06	57.86	31.37
+Domain Adaptation	30.62	42.71	25.00	52.23	37.64
+Spectral Normalization	30.20	45.42	21.90	58.72	39.06
+Label Smoothing	31.58	46.40	24.64	58.03	40.16
+Reconstruction	31.25	48.04	25.15	59.15	40.90
(Ji et al., 2015)	19.26	41.39	25.74	68.08	38.62
+weak supervision	20.35	42.25	26.32	68.92	39.46
DANN	26.19	34.20	25.74	54.70	35.21

Table 2: Per-class and macro average F1 (%) of unsupervised domain adaptation from explicit to implicit relations.

development sets. This model does not use the explicit relations.

Explicit \rightarrow Implicit A discourse relation classifier using the explicit training set and implicit development set, optimizing the standard cross-entropy loss. This serves as a baseline without any domain adaptation.

Domain adaptation Our full adaptation model trained on the explicit training set, and adapted to the implicit training set without relation labels. We perform ablation study with different extensions describe in Section 3: spectral normalization, label smoothing, and reconstruction loss.

For *benchmarking*, we train an unsupervised domain adaptation system using **DANN** (Ganin et al., 2016), which jointly learns domain-invariant representations and the classifier and is often used in NLP (c.f. Section 2). We use the same encoder, classifier and discriminator structures, with parameters tuned on the implicit development data. The system is optimized using Adam with learning rate $2e-4$ and the adaptation parameter 0.25, chosen between 0.01 and 1 on a logarithmic scale.

4.3 Results

To evaluate our model, we train four-way classifiers and report **per-class and macro F1 scores**. Table 2 tabulates the experimental results for unsupervised domain adaptation.

We also show reported results from Ji et al. (2015). Even though they trained four binary classifiers (instead of doing multi-class classification), it is the only prior work exploring unsupervised domain adaptation for implicit discourse relation classification. We include two settings: their best system with labeled data from PDTB explicit relations only (and an implicit development set), and their system with additional weak supervision

from non-PDTB sources.²

Our full system achieves the best average F1 measure, a 9.53% absolute increase from Explicit \rightarrow Implicit. It also performs 2.28% better than Ji et al. (2015)’s model trained without weak supervision, and 1.44% better than their model trained with weak supervision. The full system achieved an average F1 comparable to the supervised Implicit \rightarrow Implicit, while Ji et al. (2015)’s models did not. Comparing with DANN, our system achieved superior performance for 3 of the 4 relations, showing that training target representations and the classifier in two stages outperforms doing both jointly.

The largest improvements from the Explicit \rightarrow Implicit baseline are from Temporal (from 22.22 to 31.25) and Contingency (from 22.35 to 48.04) relations. Our system performs $\sim 11\%$ better for Temporal, and $\sim 6\%$ better for Contingency, than Ji et al. (2015)’s binary classifiers. The Comparison and Expansion relations improved by about 2% from the baseline, a smaller improvement compared to the other two relations. Our Comparison performance is comparable with Ji et al. (2015)’s model without weak supervision.

Notably, the performance for Expansion dropped after domain adaptation (without extensions) by about 5%. We suspect that this is because the distributions of Expansion among other relations are very different (33% for explicit and 53% for implicit, c.f. Table 1). By applying Spectral Normalization, the performance improved and surpassed Explicit \rightarrow Implicit.

Component-wise, Spectral Normalization helps two of the four relations (Contingency and Expansion), but hurts the performance of Comparison.

²The weakly labeled data includes sentences extracted from 1000 CNN articles, with explicit discourse connectives but without annotated discourse relations.




Figure 4: Normalized confusion matrices before and after unsupervised domain adaptation.

Label smoothing improves performance for all relations except Expansion; applying the reconstruction loss improves performance for all relations except Temporal. Overall, the best result on this task is to incorporate all components.

Error analysis Figure 4 shows the normalized confusion matrices before and after unsupervised domain adaptation. Before adaptation, Temporal and Contingency relations are often misclassified as Expansion, which is substantially improved after adaptation. The improvement in F score for Comparison is milder due to lower precision and higher recall, which is also reflected in the matrices. Finally, the drop of performance in Expansion adaptation can be traced through increased confusion between Expansion and Contingency.

5 How about a little supervision?

We have so far presented an unsupervised domain adaptation system that is not trained on any labels Y_t in the target domain. However, it is sometimes feasible to have some seed annotation that can be used to improve prediction. Hence we extend the model with an optional supervised component. We evaluate this extension by gradually

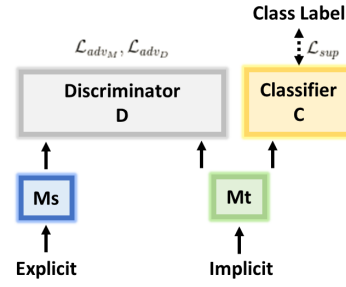


Figure 5: An extension component to incorporate supervision with our unsupervised adversarial domain adaptor.

adding labeled examples of implicit discourse relations, simulating situations when different numbers of labeled examples are available.

5.1 Incorporating supervision

We extend the model with a supervised component, where a subset $X_t^L \subseteq X_t$ has labels Y_t^L . Illustrated in Figure 5, we jointly optimize the target encoder M_t and the classifier C according to an additional supervised loss:

$$\min_{M_t, C} \mathcal{L}_{sup}(X_t^L, Y_t^L) = \mathbb{E}_{(x_t^L, y_t^L)} - \sum_k \mathbb{1}[k = y_t^L] \log C(M_t(x_t^L)) \quad (9)$$

Effectively, we encourage the target encoder to jointly extract more discriminative features for all target examples (X_t), and learning target domain representations close to the source.

The full objective incorporates supervision of the in-domain labels by adding $\mathcal{L}_{sup}(X_t^L, Y_t^L)$ to the unsupervised objective:

$$\mathcal{L}^{sup}(X_s, Y_s, X_t, M_s, M_t, D) = \mathcal{L}^{unsup}(X_s, Y_s, X_t, M_s, M_t, D) + \min_{M_t, C} \mathcal{L}_{sup}(X_t^L, Y_t^L) \quad (10)$$

5.2 Data and settings

We synthesize the labeled target subset (X_t^L, Y_t^L) ($X_t^L \subseteq X_t$) by randomly extracting subsets from the implicit training set and get their labels. The sizes of this subset range from 1382 to 13824 with a stepsize of 1382. Note that we use the entire implicit training set (X_t without relation labels Y_t) in the adversarial adaptation process as unlabeled data in the target domain, and the sampled labeled data is used in the supervised component only.

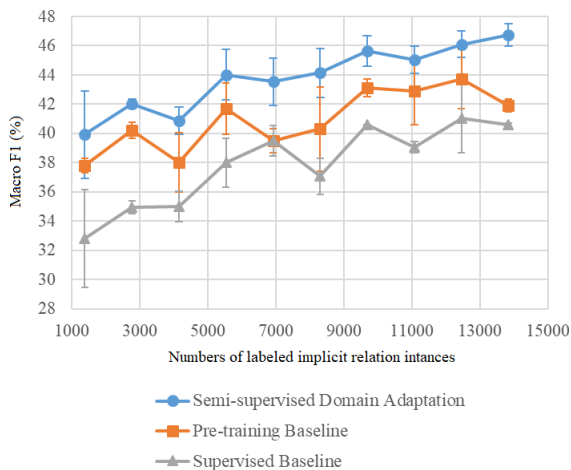


Figure 6: The average F1 (%) with varying numbers of labeled implicit relation training data.

We use the same hyper-parameters as the unsupervised experiment, except that we tune the learning rate on the implicit development set.

5.3 Systems

We compare three settings:

Supervised baseline The encoder and classifier trained on the sampled implicit instances (X_t^L, Y_t^L), optimizing Eq.(9).

Pre-training baseline Our model with the supervised component, but without the domain adaptation component. This setting is equivalent to pre-training on the explicit instances then fine-tuning on the sampled implicit instances. This is trained on the explicit training set (X_s, Y_s), plus the sampled implicit instances (X_t^L, Y_t^L), optimizing Eq.(5) and Eq.(9).

Semi-supervised domain adaptation Our full model with both the supervised and adaptation components, optimizing Eq.(10). The supervised component uses the sampled implicit instances (X_t^L, Y_t^L) for training.

5.4 Results

Since the added training data is randomly sampled, we average the performance across 3 different runs. Figure 6 shows the average F1 measure (y-axis) of the above three supervised systems, with varying numbers of labeled implicit relation training data (x-axis). Standard errors are also shown in the graph.

Our full system outperforms both the supervised baseline and the pre-training baseline, re-

gardless of the amount of labeled target data. This evaluation also reveals that the pre-training baseline also improves upon the supervised baseline across the board, which means that the performance of implicit relation classification can be improved with pre-training on explicit relations.

Finally, the macro F1 of our system using full supervision is 47.50. Since we focus on domain adaptation and used very simple encoders, we do not attempt to achieve state-of-the-art (e.g., Dai and Huang (2018), Bai and Zhao (2018)). However this performance is on-par with many recent work using multi-task or GANs, including Lan et al. (2017) (47.80), Qin et al. (2017) (44.38, Reproduced results on four-way classification), and Liu et al. (2016) (44.98). These results confirm that our framework generalizes well with respect to the amount of supervision in the target domain.

6 Conclusion

Our work tackles implicit discourse relation classification in a low resource setting that is flexible to the amount of supervision. We present a new system based on the adversarial discriminative domain adaptation framework (Tzeng et al., 2017) for unsupervised domain adaptation from explicit discourse relation to implicit discourse relation. We propose a reconstruction loss to preserve the discriminability of features during adaptation, and we generalize the framework to make use of possibly available seed data by jointly optimizing it with a supervised loss. Our system outperforms prior work and strong adversarial baselines on unsupervised domain adaptation, and works effectively with varying amount of supervision.

Acknowledgments

We thank Ray Mooney, Eric Holgate, and the anonymous reviewers for their helpful feedback. This work was partially supported by the NSF Grant IIS-1850153.

References

- Amal Al-Saif and Katja Markert. 2010. The leeds arabic discourse treebank: Annotating discourse connectives for arabic. In *LREC*.
- Hongxiao Bai and Hai Zhao. 2018. Deep enhanced representation for implicit discourse relation recognition. In *COLING*.

- Or Biran and Kathleen McKeown. 2013. Aggregated word pair features for implicit discourse relation disambiguation. In *ACL*.
- Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. 2017. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *CVPR*.
- Chloé Braud and Pascal Denis. 2014. Combining natural and artificial examples to improve implicit discourse relation identification. In *COLING*.
- Chloé Braud and Pascal Denis. 2016. Learning connective-based word representations for implicit discourse relation identification. In *EMNLP*.
- Xilun Chen, Yu Sun, Ben Athiwaratkun, Claire Cardie, and Kilian Weinberger. 2018. Adversarial deep averaging networks for cross-lingual sentiment classification. *Transactions of the Association for Computational Linguistics*, 6:557–570.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *EMNLP*.
- Zeyu Dai and Ruihong Huang. 2018. Improving implicit discourse relation classification by modeling inter-dependencies of discourse units in a paragraph. In *NAACL*.
- Lisheng Fu, Thien Huu Nguyen, Bonan Min, and Ralph Grishman. 2017. Domain adaptation for relation extraction with domain adversarial neural network. In *IJCNLP*.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(1):2096–2030.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *NIPS*.
- Tao Gui, Qi Zhang, Haoran Huang, Minlong Peng, and Xuanjing Huang. 2017. Part-of-speech tagging for twitter with adversarial neural networks. In *EMNLP*.
- Yangfeng Ji, Gongbo Zhang, and Jacob Eisenstein. 2015. Closing the gap: Domain adaptation from explicit to implicit discourse relations. In *EMNLP*.
- Shafiq Joty, Preslav Nakov, Lluís Màrquez, and Israa Jaradat. 2017. Cross-language learning with adversarial neural networks. In *CoNLL*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Man Lan, Jianxiang Wang, Yuanbin Wu, Zheng-Yu Niu, and Haifeng Wang. 2017. Multi-task attention-based neural networks for implicit discourse relationship representation and identification. In *EMNLP*.
- Man Lan, Yu Xu, and Zhengyu Niu. 2013. Leveraging synthetic discourse data via multi-task learning for implicit discourse relation recognition. In *ACL*.
- Ming-Yu Liu and Oncl Tuzel. 2016. Coupled generative adversarial networks. In *NIPS*.
- Yang Liu, Sujian Li, Xiaodong Zhang, and Zhifang Sui. 2016. Implicit discourse relation classification via multi-task neural networks. In *AAAI*.
- Daniel Marcu and Abdessamad Echihabi. 2002. An unsupervised approach to recognizing discourse relations. In *ACL*.
- Eleni Miltsakaki, Rashmi Prasad, Aravind Joshi, and Bonnie Webber. 2004. The Penn Discourse Treebank. In *LREC*.
- Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. 2018. Spectral normalization for generative adversarial networks. In *ICLR*.
- Umangi Oza, Rashmi Prasad, Sudheer Kolachina, Dipti Misra Sharma, and Aravind Joshi. 2009. The Hindi Discourse Relation Bank. In *The Third Linguistic Annotation Workshop*.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. In *NIPS-W*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- Lucie Poláková, Jiří Mírovský, Anna Nedoluzhko, Pavlína Jínová, Šárka Zikánová, and Eva Hajičová. 2013. Introducing the prague discourse treebank 1.0. In *IJCNLP*.
- Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Miltsakaki, Livio Robaldo, Aravind Joshi, and Bonnie Webber. 2008. The Penn Discourse TreeBank 2.0. In *LREC*.
- Rashmi Prasad, Bonnie Webber, and Aravind Joshi. 2014. Reflections on the Penn Discourse Treebank, comparable corpora, and complementary annotation. *Computational Linguistics*, 40(4):921–950.
- Lianhui Qin, Zhisong Zhang, Hai Zhao, Zhiting Hu, and Eric Xing. 2017. Adversarial connective-exploiting networks for implicit discourse relation classification. In *ACL*.
- Paolo Russo, Fabio Maria Carlucci, Tatiana Tommasi, and Barbara Caputo. 2018. From source to target and back: symmetric bi-directional adaptive gan. In *CVPR*.

- Attapol Rutherford and Nianwen Xue. 2015. Improving the inference of implicit discourse relations via classifying explicit discourse connectives. In *NAACL*.
- Swami Sankaranarayanan, Yogesh Balaji, Carlos D. Castillo, and Rama Chellappa. 2018. Generate to adapt: Aligning domains using generative adversarial networks. In *CVPR*.
- Wei Shi, Frances Yung, Raphael Rubino, and Vera Demberg. 2017. Using explicit discourse connectives in translation for implicit discourse relation classification. In *IJCNLP*.
- Caroline Sporleder and Alex Lascarides. 2008. Using automatically labelled examples to classify rhetorical relations: An assessment. *Natural Language Engineering*, 14(3):369–416.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *CVPR*.
- Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. 2015. Simultaneous deep transfer across domains and tasks. In *ICCV*.
- Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. 2017. Adversarial discriminative domain adaptation. In *CVPR*.
- Changxing Wu, Xiaodong Shi, Yidong Chen, Yanzhou Huang, and Jinsong Su. 2016. Bilingually-constrained synthetic data for implicit discourse relation recognition. In *EMNLP*.
- Changxing Wu, Xiaodong Shi, Yidong Chen, Jinsong Su, and Boli Wang. 2017. Improving implicit discourse relation recognition with discourse-specific word embeddings. In *ACL*.
- Ruo Chen Xu and Yiming Yang. 2017. Cross-lingual distillation for text classification. In *ACL*.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alexander J. Smola, and Eduard H. Hovy. 2016. Hierarchical attention networks for document classification. In *NAACL*.
- Deniz Zeyrek, Işın Demirşahin, AB Sevdik-Çallı, and Ruket Çakıcı. 2013. Turkish discourse bank: Porting a discourse annotation style to a morphologically rich language. *Dialogue and Discourse*, 4(2):174–184.
- Yuan Zhang, Regina Barzilay, and Tommi Jaakkola. 2017. Aspect-augmented adversarial networks for domain adaptation. *Transactions of the Association for Computational Linguistics*, 5:515–528.
- Yuping Zhou and Nianwen Xue. 2015. The Chinese Discourse TreeBank: a Chinese corpus annotated with discourse relations. *Language Resources and Evaluation*, 49(2):397–431.
- Zhi Min Zhou, Man Lan, Zheng Yu Niu, Yu Xu, and Jian Su. 2010a. The effects of discourse connectives prediction on implicit discourse relation recognition. In *SIGDIAL*.
- Zhi-Min Zhou, Yu Xu, Zheng-Yu Niu, Man Lan, Jian Su, and Chew Lim Tan. 2010b. Predicting discourse connectives for implicit discourse relation recognition. In *COLING*.