

Adaptation of Hierarchical Structured Models for Speech Act Recognition in Asynchronous Conversation

Tasnim Mohiuddin^{1*} and Thanh-Tung Nguyen^{1*} and Shafiq Joty^{1,2*}

^{*}Nanyang Technological University, Singapore

[†]Salesforce Research Asia, Singapore

{mohi0004@e., NG0155NG@e., srjoty@}ntu.edu.sg

Abstract

We address the problem of speech act recognition (SAR) in asynchronous conversations (forums, emails). Unlike synchronous conversations (e.g., meetings, phone), asynchronous domains lack large labeled datasets to train an effective SAR model. In this paper, we propose methods to effectively leverage abundant unlabeled conversational data and the available labeled data from synchronous domains. We carry out our research in three main steps. First, we introduce a neural architecture based on hierarchical LSTMs and conditional random fields (CRF) for SAR, and show that our method outperforms existing methods when trained on in-domain data only. Second, we improve our initial SAR models by semi-supervised learning in the form of pretrained word embeddings learned from a large unlabeled conversational corpus. Finally, we employ adversarial training to improve the results further by leveraging the labeled data from synchronous domains and by explicitly modeling the distributional shift in two domains.

1 Introduction

With the ever-increasing popularity of Internet and mobile technologies, communication media like emails and forums have become an integral part of people’s daily life where they discuss events, issues and experiences. Participants interact with each other *asynchronously* in these media by writing at different times, generating a type of conversational discourse that is different from synchronous conversations such as meeting and phone conversations (Louis and Cohen, 2015). In the course of the interactions, the participants perform certain communicative acts like asking questions, requesting information, or suggesting something, which are known as **speech acts** (Austin,

- C₁:hoping to do the XinJiang Tibet Highway. [Statement]
Has anyone done it? [Question]
Am hoping to hire a 4-wheel drive. [Statement]
I know the roads are bad and I would need an experienced driver and guide - any recommendations? [Question]
- C₂:I never done this routine,however i been to Xinjiang twice,in my opinion the local people not friendly, not safe to do this. [Response]
I still have relative stay in Xinjiang, however don't know what they can offer for help... [Response]
- C₃:I'm not sure if travelling overland from Xinjiang to Tibet is officially legal yet. [Response]
You might want to post your question on the North-East Asia branch of Lonely Planet's ThornTree forum for more (useful) answers. [Suggestion]
- C₄:a friend and i are trying this route as well, we will likely be in urumuqi and northern part of xinjiang from 8th apr to end apr; looking at doing the xin jiang tibet highway from end apr. (truncated) [Statement]
contact me at [email] if you want to hook up for possible transport sharing [Suggestion]
cheers. [Polite]

Figure 1: Example of speech acts in a forum thread.

1962). For example, consider the forum conversation in Figure 1. The participant who posted the initial comment C₁, describes his situation and asks a couple of *questions*. Other participants *respond* to the initial post with more information and provide *suggestions*. In this process, the participants get into a conversation by taking turns, each of which consists of one or more speech acts.

Speech act recognition (SAR) is an important step towards deep conversational analysis, and can benefit many downstream applications. Availability of large labeled datasets such as the Switchboard-DAMSL (SWBD) (Jurafsky et al., 1997) and the Meeting Recorder Dialog Act (MRDA) (Dhillon et al., 2004) corpora has fostered research in data-driven SAR methods in synchronous domains. However, such large corpora

[‡]All authors contributed equally.

are not available in the asynchronous domains, and many of the existing (small-sized) ones use task-specific tagsets as opposed to a standard one. The unavailability of large annotated datasets with standard tagsets is one of the main reasons for SAR not getting much attention in asynchronous domains, and it is often quite expensive to annotate such datasets for each domain of interest.

SAR methods proposed before the neural ‘tsunami’, *e.g.*, (Qadir and Riloff, 2011; Jeong et al., 2009; Tavafi et al., 2013), used mostly bag-of-ngram representation (*e.g.*, unigram, bigram) of a sentence, and most of these methods disregard conversational dependencies (discourse structure) between sentences. Recently, Joty and Hoque (2016) proposed a neural-CRF framework for SAR in forum conversations. In their approach, a bi-LSTM (trained on the SAR task) first encodes the sentences separately into task-specific embeddings, which are then used in a separate CRF model to capture the conversational dependencies between sentences. They also use labeled data from the MRDA meeting corpus, without which their LSTMs perform worse than simple feed-forward networks. Although their method attempts to model sentence structure (using LSTM) and conversational dependencies (using CRF), the approach has several limitations.

First, the LSTM-CRF framework was disjoint, and thus cannot be trained end-to-end. Second, when using the MRDA meeting data, their method simply concatenates it with the target domain data assuming they have the same distribution. However, asynchronous domains (forum, email) differ from synchronous (MRDA) in their underlying conversational structure (Louis and Cohen, 2015), in style (spoken vs. written), and in vocabulary usage (meetings on some focused agenda vs. conversations on any topic of interests in a public forum). Therefore, we hypothesize that to make the best use of labeled data from synchronous domains, one needs to model the shift in domains.

In this work, we advance the state-of-the-art of SAR in asynchronous conversations in three main steps. First, we introduce an end-to-end neural architecture based on a hierarchical LSTM encoder with a Softmax or CRF output layer. Second, we improve our initial SAR model by semi-supervised learning in the form of word embeddings learned from a large unlabeled conversational corpus. Most importantly, we adapt our hi-

erarchical LSTM encoder using domain adversarial training (Ganin et al., 2016) to leverage the labeled data from synchronous domains by explicitly modeling the shift in the two domains.

We evaluate our models on three different asynchronous datasets containing forum and email conversations, and on the MRDA meeting corpus. Our main findings are: (i) the hierarchical LSTMs outperform existing methods when trained on in-domain data for both synchronous and asynchronous domains, setting a new state-of-the-art; (ii) conversational word embeddings yield significant improvements over off-the-shelf ones; and (iii) domain adversarial training improves the results by inducing domain-invariant features. The source code, the conversational word embeddings, and the datasets are available at <https://ntunlp.sg.github.io/demo/project/speech-act/>.

2 Related Work

Previous studies on **SAR in asynchronous conversation** have used supervised, semi-supervised and unsupervised methods. Cohen et al. (2004) classify emails into acts like ‘deliver’ and ‘meeting’. Their approach however does not take email context into account. Carvalho and Cohen (2005) use an iterative algorithm containing two different classifiers: the content classifier that only looks at the content of the message, and the context classifier that takes into account both the content and contextual speech acts in the email thread structure. Other supervised approaches use classifiers and sequence taggers with hand-crafted features (Qadir and Riloff, 2011; Tavafi et al., 2013).

Jeong et al. (2009) use semi-supervised boosting to induce informative patterns from labeled spoken domains (MRDA, SWBD). Given a sentence represented as a set of trees (dependency, POS tags, n-grams), the boosting algorithm iteratively learns the sub-tree features. This approach does not consider the dependencies between the act types, something we successfully exploit in our work. Also, we leverage labeled data from synchronous conversations while adapting our model to account for the domain shift. Joty and Hoque (2016) use a bi-LSTM to encode a sentence, then use a separate CRF to model conversational dependencies. To learn an effective bi-LSTM model, they use the MRDA meeting data; however, without modeling the domain differences.

The unsupervised methods use variations of Hidden Markov Models (HMM) including HMM-Topic (Ritter et al., 2010), HMM-Mix (Joty et al., 2011), and Mixed Membership (Paul, 2012).

Several neural methods have been proposed in recent years for **SAR in synchronous conversations**. Kalchbrenner and Blunsom (2013) use a simple recurrent neural network (RNN) to model sequential dependencies between act types in phone conversations. They use a convolutional network to compose sentence representations from word vectors. Lee and Deroncourt (2016) use a similar model, but also experiment with RNNs to compose sentence representations. Khanpour et al. (2016) use a stacked LSTM to compose word vectors into a sentence vector. Kumar et al. (2018) also use a hierarchical LSTM-CRF. However, none of these methods were applied to asynchronous conversations, where not much labeled data is available. Also to the best of our knowledge, no prior work attempted to do domain adaptation from the synchronous conversation, which is our main contribution in this paper.

3 The Base Model

We use a bidirectional long short-term memory or bi-LSTM (Hochreiter and Schmidhuber, 1997) to encode each sentence into a vector representation. Given an input sentence $\mathbf{x}_i = (w_1, \dots, w_m)$ of length m , we first map each word w_t to its corresponding vector representation \mathbf{v}_t by looking up the word embedding matrix. The LSTM recurrent layer then computes a compositional representation \mathbf{z}_t at every time step t by performing nonlinear transformations of the current input \mathbf{v}_t and the output of the previous time step \mathbf{z}_{t-1} . The output of the last time step \mathbf{z}_m is considered as the representation of the sentence. A bi-LSTM composes a sentence in two directions: left-to-right and right-to-left, yielding a representation $\mathbf{h}_i = [\overrightarrow{\mathbf{z}}_m; \overleftarrow{\mathbf{z}}_m]$, where ‘;’ denotes concatenation. Similar to (Joty and Hoque, 2016), we could use \mathbf{h}_i to classify sentence \mathbf{x}_i into one of the speech act types using a Softmax output layer. However, in that case, we would disregard the discourse-level dependencies between sentences in a conversation. To take conversational dependencies into account, we explore two methods as we describe below.

3.1 Hierarchical LSTM

We consider a conversation as a sequence of utterances (sentences). Given an input sequence of n

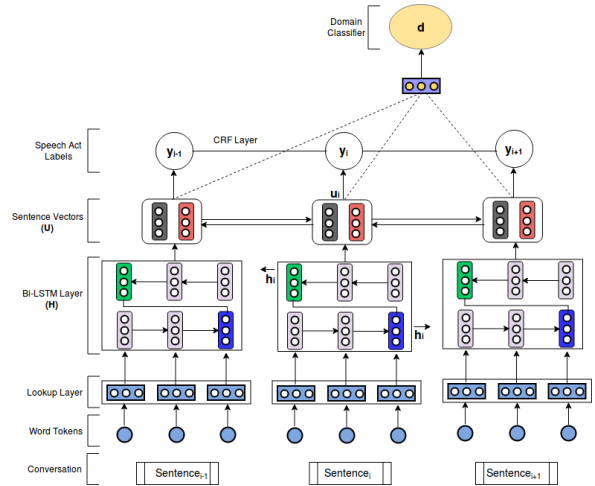


Figure 2: Hierarchical bi-LSTM-CRF model with domain adversarial training for speech act recognition.

sentences $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, the sentence-level bi-LSTM generates a sequence of n vectors $\mathbf{H} = (\mathbf{h}_1, \dots, \mathbf{h}_n)$. To consider interdependencies between sentences, we place another bi-LSTM layer on top of \mathbf{H} to connect the sentence vectors sequentially in both directions, and encode each sentence within its left and right contexts. As shown in Figure 2, the upper bi-LSTM combines the current input \mathbf{h}_i with its previous hidden state $\overrightarrow{\mathbf{u}}_{i-1}$ (resp., $\overleftarrow{\mathbf{u}}_{i+1}$) to generate a representation for the current sentence $\overrightarrow{\mathbf{u}}_i$ (resp., $\overleftarrow{\mathbf{u}}_i$). The hierarchically encoded sentence vectors $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_n)$ (where $\mathbf{u}_i = [\overrightarrow{\mathbf{u}}_i; \overleftarrow{\mathbf{u}}_i]$) are fed into a Softmax classifier for speech act classification.

$$p(y_i = k | \mathbf{X}, W, \theta) = \frac{\exp(\mathbf{w}_k^T \mathbf{u}_i)}{\sum_{k=1}^K \exp(\mathbf{w}_k^T \mathbf{u}_i)}$$

where W are the classifier weights, and θ are the parameters of the hierarchical LSTM encoder. We train the model by minimizing the cross entropy:

$$\mathcal{L}_c(W, \theta) = - \sum_{i=1}^n \sum_{k=1}^K y_{i,k} \log p(y_i = k | \mathbf{X}, W, \theta)$$

with $y_{i,k}$ being the one-hot encoding of the label.

3.2 Hierarchical LSTM with CRF

The hierarchical LSTM (H-LSTM) captures contextual information by propagating information through hidden layers, and has been shown to be effective in similar tasks such as context encoding in dialog systems (Serban et al., 2016). Despite this, its modeling strength is limited compared

to structured models that use global inference to model consistency in the output, especially when there are strong dependencies between output labels (Collobert et al., 2011). Therefore, instead of classifying sentences independently with a Softmax layer, our second method is to model them jointly with a CRF layer (Lafferty et al., 2001). For an input-output sequence pair (\mathbf{X}, \mathbf{y}) , we define the joint probability distribution:

$$p(\mathbf{y}|\mathbf{X}) = \frac{1}{Z(\mathbf{U}, A, V, \theta)} \prod_{i=1}^n \underbrace{\psi_n(y_i|\mathbf{u}_i, V)}_{\text{node factor}} \prod_{i=0}^n \underbrace{\psi_e(y_{i,i+1}|A)}_{\text{edge factor}}$$

where $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_n)$ is the hierarchically encoded sentence vectors as before, and $\psi_n(y_i = k|\mathbf{u}_i, V) = \exp(V_k^T \mathbf{u}_i)$ is the node-level score with V being the weight matrix, ψ_e is the transition matrix parameterized by A , and $Z(\cdot)$ is the global normalization constant that ensures a valid probability distribution. The cross entropy loss for the (\mathbf{X}, \mathbf{y}) sequence pair can be written as:

$$\mathcal{L}_c(V, A, \theta) = -\sum_{i=1}^n \log \psi_n(y_i|\mathbf{u}_i, V) - \sum_{i=0}^n \log A_{i,i+1} + \log Z$$

We use Viterbi decoding to infer the most probable tag sequence for an input sequence of sentences, $\mathbf{y}^* = \operatorname{argmax}_{\mathbf{y}} p(\mathbf{y}|\mathbf{X}, V, A, \theta)$. We will demonstrate later in our experiments that a CRF layer helps the H-LSTM to adapt quickly (*i.e.*, with less labeled data) to a target domain by exploiting the tag dependencies in the source domain.

4 Adaptation Methods

The hierarchical models have many parameters. Given enough training data, they should be able to encode a sentence, capturing its syntactic and semantic properties, and discourse-level dependencies. However, when it comes to SAR in asynchronous domains, not many large annotated corpora are available. Because of the large number of parameters, the models usually overfit when trained on small datasets of asynchronous conversations (shown in Sec. 6). We propose two solutions to address this problem. Our first (simple but effective) solution is to leverage large unlabeled conversational corpus to learn better task-agnostic word embeddings, and use it to initialize our models for better generalization. In the interests of coherence, we present this method in Section 5.

Our second solution is to leverage data from synchronous domains for which large annotated

corpus is available (*e.g.*, MRDA corpus). However, as we will see, simple concatenation of the datasets is not quite effective in our case, because the conversations in synchronous and asynchronous domains differ in their conversational structures, modality (spoken vs. written), and vocabulary usage. To get the best out of the available synchronous domain data, we need to adapt our models by explicitly modeling the domain shift. More precisely, our goal is to adapt the hierarchical encoder so that it learns to encode sentence representations \mathbf{U} (*i.e.*, features used for classification) that is not only discriminative for the act classification, but also invariant across the domains. We propose to use the **domain adversarial training** proposed by Ganin et al. (2016).

Let $\mathcal{D}_S = \{\mathbf{X}_p, \mathbf{y}_p\}_{p=1}^P$ denote the set of P labeled training conversations in the **source** domain (MRDA). We consider two adaptation scenarios.

(i) Unsupervised adaptation: In this scenario, we have only *unlabeled* examples in the **target** domain (*e.g.*, forum). Let $\mathcal{D}_T^u = \{\mathbf{X}_p\}_{p=P+1}^Q$ be the set of $(Q - P - 1)$ unlabeled training instances in the target domain with Q being the total number of training instances in the two domains.

(ii) Semi-supervised/supervised adaptation: In addition to the unlabeled instances \mathcal{D}_T^u , here we have access to some *labeled* training instances in the target domain, $\mathcal{D}_T^l = \{\mathbf{X}_p, \mathbf{y}_p\}_{p=Q+1}^R$, with R being the total number of training examples in the two domains. Depending on the amount of labeled data in the target domain, this setting is referred to as semi-supervised or supervised adaptation.

4.1 Unsupervised Adaptation

The dashed lines in Figure 2 show the extension of our base model for adaptation. The input conversation \mathbf{X} is sampled either from a synchronous domain (*e.g.*, meeting) or from an asynchronous domain (*e.g.*, forum). Our goal is to adapt the H-LSTM encoder (parameterized by θ) to generate \mathbf{U} such that it is not only informative for the SAR task but also invariant across domains. Upon achieving this, we can use the adapted encoder to encode a target sentence, and use the source classifier (Softmax or CRF) to classify the sentences.

We achieve this by adding a domain discriminator (dashed lines in Figure 2), another neural network that takes \mathbf{U} as input, and tries to discriminate the domains of the input conversation \mathbf{X} (*e.g.*, meeting vs. forum). The output of the discrimina-

tor is defined by a sigmoid function:

$$\hat{d}_\omega = p(d = 1 | \mathbf{u}_i, \omega, \theta) = \text{sigm}(\mathbf{w}_d^T \mathbf{h}_d) \quad (1)$$

where $d \in \{0, 1\}$ denotes the domain (1 for meeting, 0 for forum), \mathbf{w}_d are the final layer weights of the discriminator, and $\mathbf{h}_d = g(U_d \mathbf{u}_i)$ defines the hidden layer of the discriminator with U_d being the layer weights, and $g(\cdot)$ being the activations. We use cross entropy as the discrimination loss:

$$\mathcal{L}_d(\omega, \theta) = -d \log \hat{d}_\omega - (1 - d) \log (1 - \hat{d}_\omega) \quad (2)$$

The composite network has three players: the hierarchical **LSTM encoder**, the **classifier** (Softmax or CRF), and the domain **discriminator**. During training, the encoder and the classifier play a co-operative game, while the encoder and the discriminator play an adversarial game. The training objective $\mathcal{L}(W, \theta, \omega)$ of the composite model is:

$$\underbrace{\sum_{p=1}^P \mathcal{L}_c^p(W, \theta)}_{\text{act classif (src)}} - \lambda \left[\underbrace{\sum_{p=1}^P \mathcal{L}_d^p(\omega, \theta)}_{\text{domain disc (src)}} + \underbrace{\sum_{p=P+1}^Q \mathcal{L}_d^p(\omega, \theta)}_{\text{domain disc (tar)}} \right] \quad (3)$$

where θ are the parameters of the encoder, W are the classifier weights, and $\omega = \{U_d, \mathbf{w}_d\}$ are the parameters of the discriminator.¹ The hyperparameter λ controls the relative strength of the act classifier and the discriminator. We learn θ that optimizes the following **min-max** criterion:

$$\theta^* = \underset{W, \theta}{\text{argmin}} \max_{U_d, \mathbf{w}_d} \mathcal{L}(W, \theta, \omega) \quad (4)$$

Note that the updates of the shared encoder for the two networks (classifier and discriminator) work *adversarially* with respect to each other. Algorithm 1 provides pseudocode of our training method. The main challenge in adversarial training is to balance the networks (Arjovsky et al., 2017). In our experiments, we found the discriminator to be weaker initially. To balance the two components, we would need the error signals from the discriminator to be fairly weak initially, with full power unleashed only as the classification errors start to dominate. We follow the weighting schedule proposed in (Ganin et al., 2016, p. 21), which initializes λ to 0, and then changes it gradually to 1 as training progresses.

¹For simplicity, we describe adaptation of the encoder with Softmax output, but this generalizes naturally to CRF.

Algorithm 1: Adversarial training with SGD.

Input : Data $\mathcal{D}_S, \mathcal{D}_T^u$, and batch size b

Output: Adapted model parameters θ, W

1. Initialize model parameters;

2. **repeat**

(a) Randomly sample $\frac{b}{2}$ labeled examples from \mathcal{D}_S

(b) Randomly sample $\frac{b}{2}$ unlabeled examples from \mathcal{D}_T^u

(c) Compute $\mathcal{L}_c(W, \theta)$ and $\mathcal{L}_d(\omega, \theta)$

(d) Set $\lambda = \frac{2}{1 + \exp(-10 * p)} - 1$; p is the training progress linearly changing from 0 to 1.

// Classifier & Encoder

(e) Take a gradient step for $\frac{2}{b} \nabla_{W, \theta} \mathcal{L}_c(W, \theta)$

// Discriminator

(f) Take a gradient step for $\frac{2\lambda}{b} \nabla_{U_d, \mathbf{w}_d} \mathcal{L}_d(\omega, \theta)$

// Gradient reversal

(g) Take a gradient step for $-\frac{2\lambda}{b} \nabla_{\theta} \mathcal{L}_d(\omega, \theta)$

until convergence;

4.2 Semi-supervised/supervised Adaptation

It is straight-forward to extend our adaptation method to a semi-supervised/supervised setting. Similar to the instances in the source domain, the labeled instances in the target domain \mathcal{D}_T^l are used for act classification and domain discrimination. The total training loss in this case is

$$\mathcal{L}(W, \theta, \omega) = \underbrace{\sum_{p=1}^P \mathcal{L}_c^p(W, \theta)}_{\text{act classif. (source)}} + \underbrace{\sum_{p=Q+1}^R \mathcal{L}_c^p(W, \theta)}_{\text{act classif. (target)}} - \lambda \left[\underbrace{\sum_{p=1}^P \mathcal{L}_d^p(\omega, \theta)}_{\text{dom classif. (source)}} + \underbrace{\sum_{n=P+1}^R \mathcal{L}_d^p(\omega, \theta)}_{\text{dom classif. (target)}} \right]$$

where the second term is the classification loss on the target dataset \mathcal{D}_T^l , and the last term is the discrimination loss on both labeled and unlabeled data in the target domain.

5 Corpora

We now describe the datasets and the act tagset that we use, and the conversational word embeddings that we learn from a large unlabeled corpus.

5.1 Labeled Datasets

As mentioned, asynchronous domains lack large corpora that are annotated with a standard speech act tagset. Jeong et al. (2009) annotated sentences in TripAdvisor (**TA**) **forum** threads with the standard 12 act types defined in MRDA. They also remapped the **BC3 email** corpus (Ulrich et al., 2008) according to these tags. Subsequent studies (Tavafi et al., 2013; Oya and Carenini, 2014; Joty

	Asynchronous			Synchronous
	TA	BC3	QC3	MRDA
Total # of conversations	200	39	47	73
Avg. # of comments/conv	4.02	6.54	13.32	N.A
Avg. # of sentences/conv	18.56	34.15	33.28	955.10
Avg. # of words/sen	14.90	12.61	19.78	10.11

Table 1: Basic statistics about our corpora.

Tag	Description	Asynchronous			Synchronous
		TA	BC3	QC3	MRDA
SU	Suggestion	7.71	5.48	17.38	5.97
R	Response	2.4	3.75	5.24	15.63
Q	Questions	14.71	8.41	12.59	8.62
P	Polite	9.57	8.63	6.13	3.77
ST	Statement	65.62	73.72	58.66	66.00

Table 2: Distribution of speech acts in our corpora.

and Hoque, 2016) used these datasets but grouped the 12 acts into 5 coarser classes. Joty and Hoque (2016) also created a new dataset of QatarLiving² forum threads called QC3.³ We use these three asynchronous datasets in our experiments. For our experiments on synchronous domains, we use the MRDA meeting corpus that was also used in related studies (Jeong et al., 2009; Joty and Hoque, 2016). Tables 1 and 2 show some basic statistics of the datasets and the tag distributions. Note that the tagset used by us and other related studies in asynchronous (written) conversation is different from the one used in synchronous spoken conversations (Lee and Deroncourt, 2016; Khanpour et al., 2016; Kumar et al., 2018). The later tagset contains acts like *backchannel*, *filter* and *disruption* that are more specific to speech.

The train-dev-test splits of the asynchronous datasets are done uniformly at random at the **conversation level**. Since the asynchronous datasets are quite small in size, to have a reliable test set, we create the train:test splits with an equal number of conversations (Table 3). Joty and Hoque (2016) also created conversation level datasets to train and test their CRF models. Their test sets however contain only 20% of the conversations, providing only 5 conversations for QC3 and BC3, and 20 for TA. Our experiments on these small test sets showed unstable results for all the models. Therefore, we use a larger test set (50%), and we report more general results on the whole corpus based on **2-fold** cross-validation, where the second fold was

²<http://www.qatarliving.com/>

³<https://ntunlp.github.io/project/speech-act/>

Source	MRDA	Train	Dev.	Test
		59 (50865)	6 (8366)	8 (10492)
Target	TA	90 (1667)	20 (427)	90 (1617)
	QC3	20 (675)	7 (230)	20 (660)
	BC3	16 (557)	7 (233)	16 (542)
Total		126 (2899)	34 (890)	126 (2819)

Table 3: Train, dev. and test sets for the datasets. Numbers in parentheses indicate the number of sentences.

	# of Threads	# of Tokens	# of Words
W3C	23,940	21,465,830	546,921
TripAdvisor	25,000	2,037,239	127,333
Qatar Living	219,690	103,255,922	1,157,757
MRDA	-	675,110	18,514
SWBD	-	1,131,516	57,075

Table 4: Datasets and their statistics used for training the conversational word embeddings.

created by interchanging the train and test splits in Table 3. The same development set was used to tune the hyperparameters of the models for experiments on each fold. For experiments on MRDA, we use the same train:test:dev split as in (Jeong et al., 2009; Joty and Hoque, 2016).

5.2 Conversational Word Embeddings

One simple and effective approach to semi-supervised learning is to use word embeddings pretrained from a large unlabeled corpus. In our work, we use generic *off-the-shelf* pretrained embeddings to boost the performance of our models. In addition, we have also trained word embeddings from a large conversational corpus to get more relevant **conversational** word embeddings.

We use Glove (Pennington et al., 2014) to train our word embeddings from a corpus that contains 24K email threads from W3C (w3c.org), 25K threads from TripAdvisor, 220K threads from QatarLiving, and all conversations from SWBD and MRDA (a total of 120M tokens). Table 4 shows some statistics of the datasets used for training the conversational word embeddings. We also trained skip-gram word2vec (Mikolov et al., 2013), but its performance was worse than Glove.

6 Experiments

We followed similar preprocessing steps as Joty and Hoque (2016); specifically: normalize all characters to lower case, spell out digits and URLs, and tokenize the texts using TweetNLP (Gimpel et al., 2011). For performance compar-

	QC3	TA	BC3	MRDA
SVM _{c-gl}	16.96±0.00	20.17±0.00	17.20±0.00	31.47±0.00
FFN _{c-gl}	48.29±0.25	61.36±0.21	39.58±0.26	71.12±0.13
FFN _{skip-th}	50.80±1.21	61.44±0.92	47.67±0.74	71.73±0.48
B-LSTM _{rand}	50.25±0.57	62.11±0.64	45.08±1.03	70.72±0.02
B-LSTM _{gl}	53.21±0.77	63.23±0.80	49.04±0.90	72.23±0.18
B-GRU _{c-gl}	60.50±0.36	67.23±0.76	55.45±1.05	72.04±0.35
B-LSTM _{c-gl}	61.01±0.60	67.23±0.70	55.32±0.68	72.42±0.14
S-LSTM _{c-gl}	56.70±0.58	62.28±1.23	52.31±0.86	71.32±0.28
H-LSTM _{c-gl}	60.76±0.99	68.38±0.65	57.17±0.87	72.91±0.14
H-LSTM-CRF _{c-gl}	59.83±1.27	68.10±0.68	56.37±0.61	72.77±0.17

Table 5: Macro- F_1 scores for in-domain training.

ison, we use **accuracy** and **macro- F_1** . Like other related studies, we consider macro- F_1 as the main metric (more appropriate when class distributions are imbalanced), and select our model based on the best F_1 on the development set. Due to space limitations, we report only macro- F_1 here. Please refer to the Appendix for the accuracy numbers.

6.1 Experiments on In-domain Training

We first evaluate our base models on in-domain datasets by comparing with state-of-the-art models. In the next subsection, we evaluate our adaptation method in the three adaptation scenarios.

Settings. To validate the efficacy of our model, we compare it with two baselines: a Support Vector Machine (SVM) and a feed-forward network (FFN). In one setting, we use the concatenated word vectors as the input sentence representation, while in another, we use the pretrained skip-thought vectors (Kiros et al., 2015). We also compare our models with the bi-LSTM (B-LSTM) model of Joty and Hoque (2016) and the stacked LSTM (S-LSTM) of Khanpour et al. (2016).

We use the Adam optimizer (Kingma and Ba, 2014) with a learning rate of 0.001, and use dropout to avoid over-fitting. We use the Xavier initializer (Glorot and Bengio, 2010) to initialize the weights, and uniform $\mathcal{U}(-0.05, 0.05)$ to initialize the word vectors randomly. For pretrained word embeddings, we experiment with off-the-shelf embeddings that come with Glove as well as with our conversational word embeddings. For both random and pretrained initialization, we **fine-tune** our word embeddings on the SAR task.

We construct sequences from the **chronological order** of the sentences in a conversation. Since MRDA conversations are much longer compared to those in asynchronous domains (955 vs. 18-34

sentences in Table 1), we split the MRDA conversations into smaller parts containing a maximum of 100 sentences.⁴ The number of epochs and batch size were fixed to 30 and 5 (conversations), respectively. We ran each experiment 5 times, each time with a different random seed, and report the average of the (2-fold×5=10) runs along with the standard deviation. Recently, Crane (2018) show that the main source of variability in results for neural models come from the random seed, and the author has recommended to report the distribution of results from a range of seeds.

Results. We present the results in Table 5. From the **first block** of results, we notice that both SVM and FFN baselines perform poorly compared to other models that tune the word embeddings and learn the sentence representation on the SAR task.

The **second block** contains five LSTM variants: (i) **B-LSTM_{rand}**, referring to bi-LSTM with *random* initialization; (ii) **B-LSTM_{gl}**, referring to bi-LSTM initialized with *off-the-shelf* Glove embeddings; (iii) **B-GRU_{c-gl}**, referring to bidirectional Gated Recurrent Unit (Cho et al., 2014) initialized with our *conversational* Glove; (iv) **B-LSTM_{c-gl}**, referring to bi-LSTM initialized with *conversational* Glove, and (v) **S-LSTM_{c-gl}**, referring to a 2-layer stacked LSTM with *conversational* Glove.⁵ From the results, we can make the following conclusions. First, B-LSTM_{rand} overfits extremely on

⁴In a different setting, we created sequences by connecting each non-initial comment with the initial comment generating many 2-comment sequences. This is considering the fact that in many QA forums, users mostly answer to the questions asked in the initial post. In our experiments on in-domain training, we found this competitive with our ‘one long-chain’ structure. However, the adaptation in this setting was much worse because of the mismatch in discourse structures of synchronous and asynchronous conversations.

⁵Increasing the number of layers in S-LSTM_{c-gl} did not give any gain (see Table 2 in the Appendix).

the asynchronous datasets, giving the worst results among the LSTMs. Second, pretrained vectors help to achieve better results, however, compared to the off-the-shelf vectors, our conversational word vectors yield much higher F_1 , especially, in the asynchronous datasets that are smaller in size (5 - 11% absolute gains). This demonstrates that pretrained word embeddings provide an effective method to perform semi-supervised learning, when they are learned from relevant datasets.

The **last block** shows the results of our models. It is evident that both H-LSTM and H-LSTM-CRF outperform other models in all the datasets except QC3 where the difference is very small. They also give the best F_1 reported so far on MRDA, outperforming the B-LSTM models of Joty and Hoque (2016) and S-LSTM model of Khanpour et al. (2016). When we compare the two models, we notice that H-LSTM outperforms H-LSTM-CRF in all the datasets. A reason for this could be that the contextual dependency is already captured by the upper LSTM layer and the data may be too small for the CRF to offer anything more.

6.2 Experiments on Domain Adaptation

Settings. We compare our adversarial adaptation method with three baseline methods: Transfer, Merge and Fine-tune. **Transfer** models are trained on the source (MRDA) and tested on the target (QC3, TA, BC3). Our **adversarial unsupervised adaptation** method is comparable to the transfer method as they use labeled data only from the source domain. In **Merge**, models are trained on the concatenated training set of source and target datasets. **Fine-tune** is a widely used adaptation method for neural models (Chu and Wang, 2018). In this method, we first train a model on the source domain until convergence, then we fine-tune it on the target by training it further on the target dataset. Both merge and fine-tune are comparable to our **semi-supervised/supervised adaptation** as these methods use labeled data from the target domain. For semi-supervised experiments, we take smaller subsets (*e.g.*, 25%, 50%, and 75% of the labeled data) from the target domain.

We also compare our method with **Neural SCL** (Ziser and Reichart, 2017), which is another domain adaptation method in the neural framework. We used the implementation made available by the authors.⁶ For training our adaptation models,

⁶<https://github.com/yftah89/structural-correspondence-learning-SCL>

Method	Model	QC3	TA	BC3
Transfer	SVM	17.78 \pm 0.00	20.44 \pm 0.00	17.85 \pm 0.00
	FFN	46.91 \pm 0.00	56.30 \pm 0.00	46.74 \pm 0.00
	S-LSTM	49.89 \pm 1.29	62.52 \pm 1.49	36.36 \pm 1.28
	B-LSTM	50.50 \pm 0.91	65.47\pm0.62	35.92 \pm 0.62
	H-LSTM	50.22 \pm 0.64	64.43 \pm 0.52	35.11 \pm 1.64
	H-LSTM-CRF	50.83\pm0.70	63.80 \pm 0.81	34.45 \pm 1.42
Unsup. adapt	Neural SCL	37.73 \pm 0.92	53.98 \pm 0.33	46.90\pm0.89
	Adv-S-LSTM	43.36 \pm 1.44	48.51 \pm 0.51	42.05 \pm 0.21
	Adv-B-LSTM	47.39 \pm 0.74	58.49 \pm 1.29	32.86 \pm 1.35
	Adv-H-LSTM	46.53 \pm 1.48	52.90 \pm 1.20	31.36 \pm 1.91
	Adv-H-LSTM-CRF	47.06 \pm 1.24	61.58 \pm 0.78	29.54 \pm 1.06
Merge (50%)	S-LSTM	55.39 \pm 0.29	67.79 \pm 0.15	50.82 \pm 0.98
	B-LSTM	55.08 \pm 0.67	68.99 \pm 0.35	51.05 \pm 0.60
	H-LSTM	51.74 \pm 0.44	69.09 \pm 0.64	47.82 \pm 1.47
	H-LSTM-CRF	50.92 \pm 0.48	68.66 \pm 0.12	48.58 \pm 0.59
Fine-tune (50%)	S-LSTM	53.94 \pm 1.04	66.07 \pm 0.67	51.73 \pm 1.53
	B-LSTM	54.81 \pm 0.48	68.43 \pm 0.51	52.26 \pm 0.82
	H-LSTM	54.34 \pm 0.71	69.16 \pm 0.66	50.81 \pm 0.97
	H-LSTM-CRF	54.97 \pm 0.87	69.91 \pm 0.53	51.42 \pm 1.24
Semisup. adapt (50%)	Neural SCL	41.46 \pm 0.75	58.85 \pm 0.27	48.32 \pm 0.19
	Adv-S-LSTM	60.20 \pm 0.32	68.71 \pm 0.75	57.97 \pm 0.39
	Adv-B-LSTM	58.57 \pm 0.80	66.51 \pm 0.81	54.39 \pm 1.28
	Adv-H-LSTM	60.19 \pm 1.10	69.43 \pm 0.64	58.39 \pm 1.12
	Adv-H-LSTM-CRF	61.81\pm0.63	70.34\pm0.62	59.43\pm1.41
Merge (100%)	S-LSTM	59.18 \pm 1.40	66.93 \pm 1.70	54.87 \pm 1.47
	B-LSTM	58.33 \pm 0.84	70.12 \pm 0.39	55.89 \pm 0.89
	H-LSTM	59.89 \pm 0.57	70.40 \pm 0.41	57.19 \pm 0.87
	H-LSTM-CRF	59.53 \pm 0.66	69.88 \pm 0.68	56.04 \pm 1.15
Fine-tune (100%)	S-LSTM	56.67 \pm 0.85	67.41 \pm 0.34	56.40 \pm 0.44
	B-LSTM	59.74 \pm 0.53	69.87 \pm 0.82	57.09 \pm 1.14
	H-LSTM	60.12 \pm 0.44	70.96 \pm 0.61	58.09 \pm 1.03
	H-LSTM-CRF	59.99 \pm 0.59	70.44 \pm 0.76	57.17 \pm 0.97
Sup. adapt	Neural SCL	43.35 \pm 0.30	60.40 \pm 0.23	48.88 \pm 0.70
	Adv-S-LSTM	61.19 \pm 0.48	70.33 \pm 0.66	59.19 \pm 0.67
	Adv-B-LSTM	60.60 \pm 0.68	69.30 \pm 0.50	60.12 \pm 1.26
	Adv-H-LSTM	63.10\pm0.83	72.82 \pm 0.59	60.38\pm1.07
	Adv-H-LSTM-CRF	62.24 \pm 0.74	73.04\pm0.38	59.84 \pm 0.75

Table 6: Domain adaptation results on our datasets. All models use conversational word embeddings. Results are averaged over (2 folds \times 5) 10 runs.

we use SGD (Algorithm 1 in the Appendix) with a momentum term of 0.9 and a dynamic learning rate as suggested by Ganin et al. (2016).

Results. The adaptation results are shown in Table 6. We observe that without any labeled data from the target (**Unsup. adapt**), our adversarial adapted models (Adv-H-LSTM, Adv-H-LSTM-CRF) perform worse than the transfer baseline in all three datasets. In this case, since the out-of-domain labeled dataset (MRDA) is much larger, it overwhelms the model inducing features that are not relevant for the task in the target domain. However, when we provide the models with some labeled in-domain examples in the **semi-supervised (50%)** setting, we observe about 11% absolute gains in QC3 and BC3 over the corresponding Merge baselines, and 7 - 8% gains over the corresponding Fine-tune baselines. As we add more target labels (**100%**), performance of our adapted models (**Sup. adapt**) improve further, yielding sizable improvements (\sim 3% absolute) over the corresponding baselines in all datasets.




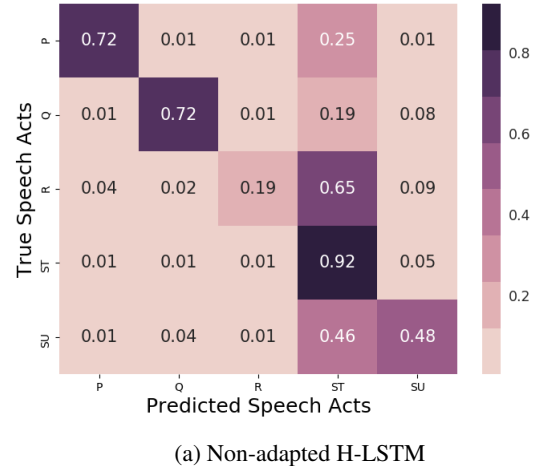
Figure 3: F_1 with varying amount of target labels.

Also notice that our adversarial adaptation outperforms Merge and Fine-tune methods for all models over all datasets, showing its effectiveness.

Figure 3 presents the F_1 scores of our adapted models with varying amount of labeled data in the target domain. We notice that the largest improvements for all three datasets come from the first 25% of the target labels. The gains from the second quartile are also relatively higher than the last two quartiles for TA and BC3. Another interesting observation is that H-LSTM-CRF performs better in unsupervised and semi-supervised settings (*i.e.*, with less target labels). In other words, H-LSTM-CRF adapts better than H-LSTM with small target datasets by exploiting the tag dependencies in the source. As we include more labeled data from the target, H-LSTM catches up with H-LSTM-CRF. Surprisingly, Neural SCL performs the worst. We suspect this is due to the mismatches between pivot features of the source and target domains.

If we compare our **adaptation results** with the **in-domain results** in Table 5, we notice that using the same amount of labeled data in the target, our supervised adaptation gives 3-4% gains across the datasets. Our semi-supervised adaptation using half of the target labels (50%) also outperforms the in-domain models that use all the target labels.

To further analyze the cases where our adapted models make a difference, Figure 4 shows the confusion matrices for the adapted H-LSTM and the non-adapted H-LSTM on the concatenated test-sets of QC3, TA, and BC3. In general, our classifiers get confused between *Response* and *Statement*, and between *Suggestion* and *Statement* the most. We noticed similar phenomena in the human annotations, where annotators had difficulties



(a) Non-adapted H-LSTM



(b) Adapted H-LSTM

Figure 4: Confusion matrices on the combined test sets.

with these three acts. It is however noticeable that the adapted H-LSTM is less affected by class imbalance, and it can detect the *Suggestion* and *Polite* acts more correctly than the non-adapted one.

7 Conclusion

We proposed an adaptation framework for speech act recognition in asynchronous conversation. Our base model is a hierarchical LSTM encoder with a Softmax or CRF output layer, which achieves state-of-the-art results for in-domain training. Crucial to its performance is the conversational word embeddings. We adapted our base model with adversarial training to effectively leverage out-of-domain meeting data, and to improve the results further. A comparison with existing methods and baselines in different training scenarios demonstrates the effectiveness of our approach.

Acknowledgments

Shafiq Joty would like to thank the funding support from MOE Tier-1 (Grant M4011897.020).

References

- Martín Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein GAN. *CoRR*, abs/1701.07875.
- John Langshaw Austin. 1962. How to do things with words. *Harvard University Press*.
- Vitor R. Carvalho and William W. Cohen. 2005. On the collective classification of email "speech acts". In *SIGIR '05*, pages 345–352, New York, NY, USA.
- Kyunghyun Cho, Bart van Merriënboer, Çalar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *EMNLP'14*, Doha, Qatar.
- Chenhui Chu and Rui Wang. 2018. A survey of domain adaptation for neural machine translation. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1304–1319. Association for Computational Linguistics.
- William W. Cohen, Vitor R. Carvalho, and Tom M. Mitchell. 2004. Learning to classify email into "speech acts". In *EMNLP'04*, pages 309–316.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *JMLR*, 12:2493–2537.
- Matt Crane. 2018. Questionable answers in question answering research: Reproducibility and variability of published results. *TACL'18*, 6:241–252.
- Rajdip Dhillon, Sonali Bhagat, Hannah Carvey, and Elizabeth Shriberg. 2004. Meeting Recorder Project: Dialog Act Labeling Guide. Technical report, ICSI Tech. Report.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *JMLR*, 17(1):2096–2030.
- Kevin Gimpel, Nathan Schneider, Brendan O'Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A Smith. 2011. Part-of-speech tagging for twitter: Annotation, features, and experiments. In *HLT'11: short papers-Volume 2*, pages 42–47.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *JMLR*, Sardinia, Italy.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Minwoo Jeong, Chin-Yew Lin, and Gary Geunbae Lee. 2009. Semi-supervised speech act recognition in emails and forums. In *EMNLP'09*, pages 1250–1259, Singapore.
- Shafiq Joty, Giuseppe Carenini, and Chin-Yew Lin. 2011. Unsupervised Modeling of Dialog Acts in Asynchronous Conversations. In *IJCAI'11*, pages 1–130, Barcelona.
- Shafiq Joty and Enamul Hoque. 2016. Speech act modeling of written asynchronous conversations with task-specific embeddings and conditional structured models. In *ACL '16*, pages 1746–1756, Berlin, Germany.
- Dan Jurafsky, Liz Shriberg, and Debra Biasca. 1997. Switchboard SWBD-DAMSL Shallow-Discourse-Function Annotation Coders Manual. Technical report, UC Boulder & SRI International.
- Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent convolutional neural networks for discourse compositionality. In *CVSC'13*, pages 119–126.
- Hamed Khanpour, Nishitha Guntakandla, and Rodney Nielsen. 2016. Dialogue act classification in domain-independent conversations using a deep recurrent neural network. In *COLING'16*, pages 2012–2021.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2015. Skip-thought vectors. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'15, pages 3294–3302, Cambridge, MA, USA.
- Harshit Kumar, Arvind Agarwal, Riddhiman Dasgupta, Sachindra Joshi, and Arun Kumar. 2018. Dialogue act sequence labeling using hierarchical encoder with CRF. In *AAAI'18*.
- John Lafferty, Andrew K. McCallum, and Fernando Pereira. 2001. Conditional Random Fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*.
- Ji Young Lee and Franck Dernoncourt. 2016. Sequential short-text classification with recurrent and convolutional neural networks. In *HLT'16*, pages 515–520.
- Annie P Louis and Shay B Cohen. 2015. Conversation trees: A grammar model for topic structure in forums. In *ACL '15*.
- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013. Linguistic Regularities in Continuous Space Word Representations. In *NAACL-HLT '13*, pages 746–751, Atlanta, GA, USA.
- Tatsuro Oya and Giuseppe Carenini. 2014. Extractive summarization and dialogue act modeling on email threads: An integrated probabilistic approach. In *SIGDIAL'14*, page 133–140, Philadelphia, PA, U.S.A.

- Michael J. Paul. 2012. Mixed Membership Markov Models for Unsupervised Conversation Modeling. In *EMNLP-CoNLL '12*, pages 94–104, Stroudsburg, PA, USA.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. *Glove: Global vectors for word representation*. In *EMNLP'14*, pages 1532–1543, Doha, Qatar.
- Ashequl Qadir and Ellen Riloff. 2011. *Classifying sentences as speech acts in message board posts*. In *EMNLP'11*, pages 748–758, Edinburgh, Scotland, UK.
- A. Ritter, C. Cherry, and B. Dolan. 2010. Unsupervised modeling of twitter conversations. In *HLT: NAACL'10*, LA, California.
- Iulian V. Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. *Building end-to-end dialogue systems using generative hierarchical neural network models*. In *AAAI'16*.
- Maryam Tavafi, Yashar Mehdad, Shafiq Joty, Giuseppe Carenini, and Raymond Ng. 2013. *Dialogue act recognition in synchronous and asynchronous conversations*. In *SIGDIAL'13*, Metz, France.
- Jan Ulrich, Gabriel Murray, and Giuseppe Carenini. 2008. A publicly available annotated corpus for supervised email summarization. In *AAAI'08*, Chicago, USA.
- Yftah Ziser and Roi Reichart. 2017. Neural structural correspondence learning for domain adaptation. In *CoNLL 2017*, pages 400–410.