

Guiding Attention in Sequence-to-Sequence Models for Dialogue Act prediction

Pierre Colombo,^{1,2*} Emile Chapuis,^{1*} Matteo Manica³
Emmanuel Vignon,² Giovanna Varni,¹ Chloe Clavel¹

¹LTCI, Telecom Paris, Institut Polytechnique de Paris, ²IBM GBS France, ³IBM Zurich
¹firstname.lastname@telecom-paris.fr, ³tte@zurich.ibm.com, emmanuel.vignon@fr.ibm.com

Abstract

The task of predicting dialog acts (DA) based on conversational dialog is a key component in the development of conversational agents. Accurately predicting DAs requires a precise modeling of both the conversation and the global tag dependencies. We leverage seq2seq approaches widely adopted in Neural Machine Translation (NMT) to improve the modelling of tag sequentiality. Seq2seq models are known to learn complex global dependencies while currently proposed approaches using linear conditional random fields (CRF) only model local tag dependencies. In this work, we introduce a seq2seq model tailored for DA classification using: a hierarchical encoder, a novel *guided attention* mechanism and beam search applied to both training and inference. Compared to the state of the art our model does not require hand-crafted features and is trained end-to-end. Furthermore, the proposed approach achieves an unmatched accuracy score of 85% on SwDA, and state-of-the-art accuracy score of 91.6% on MRDA.

1 Introduction

In natural language processing research, the dialogue act (DA) concept plays an important role. DAs are semantic labels associated with each utterance in a conversational dialogue that indicate the speaker's intention, e.g., question, backchannel, statement-non-opinion, statement opinion. A key to model dialogue is to detect the intent of the speaker: correctly identifying a question gives an important clue to produce an appropriate response. As can be observed in Table 1, DA classification relies on its conversational aspect, i.e., predicting an utterance's DA requires the knowledge of previous sentences and their associated act labels. For example, if a speaker asks a question, the interlocutor will answer with a response, analogously, a "Greeting" or a "Farwell" will be followed by a similar dialogue act. This means that in a conversation there is a sequential structure in the emitted dialogue acts. This poses the basis for the adoption of a novel perspective on the DA classification problem, i.e., from a multi-classification task to a sequence labeling one.

*Equal contribution

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

| Speaker | Utterance |
|---------|---|
| A | Is there anyone who doesn't know Nancy? |
| A | Do you - Do you know Nancy ? |
| B | Me? |
| B | Mm-hmm |
| B | I know Nancy |

Table 1: Example of conversation from Switchboard Dialogue Act Corpus. A is speaking with B.

Limitations of current models: Current state-of-the-art models rely on the use of linear Conditional Random Field (CRF) combined with a recurrent neural network based encoder (Li et al. 2018a; Chen et al. 2018; Raheja and Tetreault 2019) to model DA sequential dependencies. Unfortunately such approaches only capture local dependencies between two adjacent dialogue acts. For instance, if we consider the example in Table 1 we can see that the last statement "I know Nancy" is a response to the first question "Is there anyone who doesn't know Nancy" and the knowledge of the previous backchannel does not help the prediction of the last dialogue act. Therefore, we must consider dependencies between labels with a scope that is wider than two successive utterances. In Neural Machine Translation (NMT), the problem of global dependencies has been addressed using seq2seq models (Sutskever, Vinyals, and Le 2014) that follow the encoder-decoder framework. The encoder embeds an input sentence into a single hidden vector which contains both global and local dependencies, and the hidden vector is then decoded to produce an output sequence. In this work, we propose a seq2seq architecture tailored towards DA classification paving the way for further innovations inspired by advances in NMT research.

Contributions: In this work (1) we formalise the Dialogue Act Prediction problem in a way that emphasises the relations between DA classification and NMT, (2) we demonstrate that the seq2seq architecture suits better to the DA classification task and (3) we present a seq2seq model leveraging NMT techniques that reaches an accuracy of 85%, outperforming the state of the art by a margin of around 2%, on the Switchboard Dialogue Act Cor-

pus (SwDA) (Stolcke et al. 1998) and a state-of-the-art accuracy score of 91,6% on the Meeting Recorder Dialogue Act (MRDA). This seq2seq model exploits a hierarchical encoder with a novel *guided attention* mechanism that fits with our setting without any handcrafted features. We fine-tune our seq2seq using a sequence level training objective making use of the beam search algorithm. To our knowledge, this is among the first seq2seq model proposed for DA classification.

2 Background

DA classification

Several approaches have been proposed to tackle the DA classification problem. These methods can be divided into two different categories. The first class of methods relies on the independent classification of each utterance using various techniques, such as HMM (Stolcke et al. 2000), SVM (Surendran and Levow 2006) and Bayesian Network (Keizer, op den Akker, and Nijholt 2002). The second class, which achieves better performance, leverages the context, to improve the classifier performance by using deep learning approaches to capture contextual dependencies between input sentences (Bothe et al. 2018; Khanpour, Guntakandla, and Nielsen 2016). Another refinement of input context-based classification is the modelling of inter-tag dependencies. This task is tackled as sequence-based classification where output tags are considered as a DA sequence (Li et al. 2018a; Kumar et al. 2018; Stolcke et al. 2000; Chen et al. 2018; Raheja and Tetreault 2019).

Two classical benchmarks are adopted to evaluate DA classification systems: the Switchboard Dialogue Act Corpus (SwDA)(Stolcke et al. 1998) and the Meeting Recorder Dialogue Act (MRDA) (Janin et al. 2003). State-of-the-art techniques achieve an accuracy of 82.9% (Li et al. 2018a; Raheja and Tetreault 2019). To capture input contextual dependencies they adopt a hierarchical encoder and a CRF to model inter-tag dependencies. The main limitation of the aforementioned architecture is that a linear-CRF model is able to only capture dependencies at a local level and fails to capture non local dependencies. In this paper, we tackle this issue with a sequence-to-sequence using a guided attention mechanism.

Seq2seq models

Seq2seq models have been successfully applied to NMT, where modeling non local dependencies is a crucial challenge. DA classification can be seen as a problem where the goal is to map a sequence of utterances to a sequence of DA. Thus, it can be formulated as sequence to sequence problem very similar to NMT.

The general architecture of our seq2seq models (Sutskever, Vinyals, and Le 2014) follows a classical encoder-decoder approach with attention (Luong, Pham, and Manning 2015). We use GRU cells (Cho et al. 2014), since they are faster to train than LSTM ones (Jozefowicz, Zaremba, and Sutskever 2015). Recent advances have improved both the learning and the inference process, producing sequences that are more coherent by means of sequence level losses Wiseman and Rush

(2016) and various beam search settings (Wu et al. 2016; ?). The closest setting where seq2seq model have been successfully used is dependency parsing (Li et al. 2018b), where output dependencies are crucial to achieve state-of-the-art performance. In our work we adjust NMT techniques to the specifics of DA classification.

3 Problem statement

DA classification as an NMT problem

First, let's define the mathematical notations we will adopt in this work. We have a set D of conversations, i.e $D = (C_1, C_2, \dots, C_{|D|})$ with $Y = (Y_1, Y_2, \dots, Y_{|D|})$ the corresponding set DA labels. A conversation C_i is a sequence of utterances, namely $C_i = (u_1, u_2, \dots, u_{|C_i|})$ with $Y_i = (y_1, y_2, \dots, y_{|C_i|})$ the corresponding sequence of DA labels. Thus, each utterance u_i is associated with a unique DA label $y_i \in \mathcal{Y}$ where \mathcal{Y} is the set of all the possible dialogue acts. Finally, an utterance u_i can be seen as a sequence of words, i.e $u_i = (\omega_1^i, \omega_2^i, \dots, \omega_{|u_i|}^i)$. In NMT, the goal is to associate for any sentence $X^{l_1} = (x_1^{l_1}, \dots, x_{|X^{l_1}|}^{l_1})$ in language l_1 a sentence $X^{l_2} = (x_1^{l_2}, \dots, x_{|X^{l_2}|}^{l_2})$ in language l_2 where $x_i^{l_k}$ is the i word in the sentence in language l_k . Using this formalism, it is straightforward to notice two main similarities $(\mathbb{S}_1, \mathbb{S}_2)$ between DA classification and NMT. (\mathbb{S}_1) In NMT and DA classification, the goal is to maximise the likelihood of the output sequence given the input sequence ($P(X^{l_2}|X^{l_1})$ versus $P(Y_i|C_i)$). (\mathbb{S}_2) For the two tasks, there are strong dependencies between units composing both the input and output sequences. In NMT, those units are words (x_i and y_i), in DA classification those units are utterances and DA labels (u_i and y_i).

Specifics of DA classification

While NMT and DA classification are similar under some point of views, three differences are immediately apparent (\mathbb{D}_i) . (\mathbb{D}_1) In NMT, the input units x_i represent words, in DA classification u_i are input sequences composed with words. Considering the set of all possible sequences as input (context consideration leads to superior performance) implies that the dimension of the input space several order of magnitude larger than compared to a standard NMT. (\mathbb{D}_2) In DA, we have a perfect alignment between input and output sequences (hence $T = T'$). Some languages, e.g., French, English, Italian share a partial alignment, but in DA classification we have a strong mapping between y_i and x_i . (\mathbb{D}_3) In NMT, the input space (number of words in l_1) is approximately the same size of the output space (number of words in l_2). In our case the output space (number of DA tags $|\mathcal{Y}| < 100$ has a limited size, with a dimension that is many order of magnitude smaller than the input space one.

In the following, we propose an end-to-end seq2seq architecture for DA classification that leverages (\mathbb{D}_1) using a hierarchical encoder, (\mathbb{D}_2) through a guided attention mechanism and (\mathbb{D}_3) using beam search during both training and inference, taking advantage of the limited dimension of the output space.

4 Models

In Seq2seq, the encoder takes a sequence of sentences and represents it as a single vector $H_i \in \mathcal{R}^d$ and then pass it to the decoder for tag generations.

Encoders

In this section we introduce the different encoders we consider in our experiments. We exploit the hierarchical structure of the dialogue to reduce the input space size (\mathbb{D}_1) and to preserve word/sentence structure. During both training and inference, the context size is fixed to T . Formally, an encoder takes as input a fixed number of utterances (u_{i-T}, \dots, u_i) and outputs a vector $H_i \in \mathcal{R}^d$ which will serve to initialize the hidden state of the decoder. The first level of the encoder computes \mathcal{E}_{u_t} , an embedding of u_t based on the words composing the utterance, and the next levels compute H_i based on \mathcal{E}_{u_t} .

Vanilla RNN encoder: The vanilla RNN encoder (VGRU_E) introduced by Sutskever, Vinyals, and Le (2014) is considered as a baseline encoder. In the vanilla encoder $\mathcal{E}_{u_i} = \frac{1}{|u_i|} \sum_{k=1}^{|u_i|} \mathcal{E}_{w_k^i}$ where $\mathcal{E}_{w_k^i}$ is an embedding of w_k^i . To better model dependencies between consecutive utterances, we use a bidirectional GRU (Cho et al. 2014):

$$\begin{aligned} \overrightarrow{h_{i-T}^s} &= \overleftarrow{h_{i-T}^s} = \vec{0} \\ \overrightarrow{h_t^s} &= \overrightarrow{GRU}(\mathcal{E}_{u_t}), t \in [i-T, i] \\ \overleftarrow{h_t^s} &= \overleftarrow{GRU}(\mathcal{E}_{u_t}), t \in [i, i-T] \\ H_i &= [\overrightarrow{h_i^s}, \overleftarrow{h_i^s}] \end{aligned} \quad (1)$$

Hierarchical encoders: The vanilla encoder can be improved by computing \mathcal{E}_{u_i} using bi-GRU. This hierarchical encoder (HGRU) is in line with the one introduced by Sorboni et al. (2015). Formally \mathcal{E}_{u_i} is defined as it follows:

$$\begin{aligned} \overrightarrow{h_0^w} &= \overleftarrow{h_0^w} = \vec{0} \\ \overrightarrow{h_t^w} &= \overrightarrow{GRU}(\mathcal{E}_{w_t^i}), t \in [1, |u_i|] \\ \overleftarrow{h_t^w} &= \overleftarrow{GRU}(\mathcal{E}_{w_t^i}), t \in [|u_i|, 1] \\ \mathcal{E}_{u_i} &= [\overrightarrow{h_{|u_i|}^w}, \overleftarrow{h_{|u_i|}^w}] \end{aligned} \quad (2)$$

H_i is then computed using Equation 1. Intuitively, the first GRU layer (Equation 2) models dependencies between words (the hidden state of the word-level GRU is reset at each new utterance), and the second layer models dependencies between utterances.

Persona hierarchical encoders: In SwDA, a speaker turn can be splitted in several utterances. For example, if speaker A is interacting with speaker B we might encounter the sequence (AAABBBAA)¹. We propose a novel Persona Hierarchical encoder (PersoHGRU) to better model speaker-utterance dependencies. We introduce a persona layer be-

¹In SwDA around two third of the sentence have at least a AA or BB

tween the word and the sentence levels, see Figure 1:

$$\begin{aligned} \overrightarrow{h_t^p} &= \begin{cases} \vec{0} & \text{if } t \text{ and } t-1 \text{ have different speakers} \\ \overrightarrow{GRU}(\mathcal{E}_{u_{t-1}}) & \end{cases} \\ \overleftarrow{h_t^p} &= \begin{cases} \vec{0} & \text{if } t \text{ and } t+1 \text{ have different speakers} \\ \overleftarrow{GRU}(\mathcal{E}_{u_{t+1}}) & \end{cases} \\ \mathcal{E}_{u_k}^p &= [\overrightarrow{h_k^p}, \overleftarrow{h_k^p}] \quad \forall k \in [i-T, i] \end{aligned} \quad (3)$$

H_i is then obtained following Equation 1 where \mathcal{E}_{u_i} is replaced by $\mathcal{E}_{u_i}^p$.

Decoders

In this section, we introduce the different decoders we compare in our experiments. We introduce a novel form of attention that we name *guided attention*. *Guided attention* leverages the perfect alignment between input and output sequences (\mathbb{D}_2). The decoder computes the probability of the sequence of output tags based on:

$$p(y_{i-T}, \dots, y_i | u_{i-T}, \dots, u_i) = \prod_{k=i-T}^i p(y_k | H_i, y_{k-1}, \dots, y_{i-T}) \quad (4)$$

see Equation 1.

Vanilla decoder: The vanilla decoder (VGRU_D) is similar to the one introduced by Sutskever, Vinyals, and Le (2014).

Decoders with attention: In NMT, the attention mechanism forces the seq2seq model to learn to focus on specific parts of the sequence each time a new word is generated and let the decoder correctly align the input sequence with output sequence. In our case, we follow the approach described by Bahdanau, Cho, and Bengio (2014) and we define the context vector as:

$$c_k = \sum_{j=i-T}^i \alpha_{j,k} h_j^s \quad (5)$$

where $\alpha_{j,k}$ scores how well the inputs around position k and the output at position j match. Since we have a perfect alignment (\mathbb{D}_2), we know a priori on which sequence the decoder needs to focus more at each time step. Taking into account this aspect of the problem, we propose three different attention mechanisms.

Vanilla attention: This attention represents our baseline attention mechanism and it is the one proposed by Bahdanau, Cho, and Bengio (2014), where:

$$\alpha_{j,k} = \text{softmax}(a(h_{k-1}^{Dec}, h_j^s)) \quad (6)$$

and a is parametrized as a feedforward neural network.

Hard guided attention: The *hard guided attention* forces the decoder to focus only on the u_i while predicting y_i :

$$\alpha_{j,k} = \begin{cases} 0, & \text{if } k \neq j \\ 1, & \text{otherwise} \end{cases} \quad (7)$$

Soft guided attention: The *soft guided attention* guides the decoder to mainly focus on the u_i while predicting y_i ,

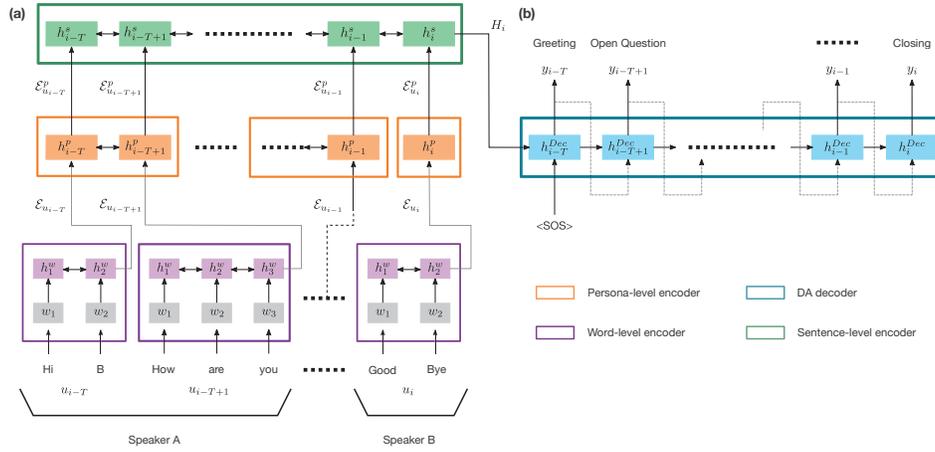


Figure 1: Seq2seq model architecture for DA classification. (a) The encoder is composed with three different levels representing a different hierarchical level in the dialogue. The utterances are encoded at: word level (purple), persona level (orange) and sentence level (green). (b) The decoder (blue) is responsible to generate for each utterance a DA exploiting the last state of the encoder as initial hidden state.

but allows it to have a limited focus on other parts of the input sequence.

$$\tilde{\alpha}_{j,k} = \begin{cases} a(h_{k-1}^{Dec}, h_j^s), & \text{if } k \neq j \\ 1 + a(h_{k-1}^{Dec}, h_j^s), & \text{otherwise} \end{cases} \quad (8)$$

$$\alpha_{j,k} = \text{softmax}(\tilde{\alpha}_{j,k}) \quad (9)$$

where a is parametrised as a feedforward neural network.

Training and inference

In this section, we describe the training and the inference strategies used for our models. A seq2seq model aims to find the best sentence for a given source sentence. This poses a computational challenge when the output vocabulary size is large, since even by using beam search it's expensive to explore multiple paths. Since our output vocabulary size is limited (\mathbb{D}_3), we do not incur in this problem and we can use beam search during both training and inference.

Beam search: In our work we measure the sequence likelihood based on the following formula:

$$s(\tilde{\mathbf{y}}^k, \mathbf{u}_i) = \frac{\log P(\tilde{\mathbf{y}}^k | \mathbf{u}_i)}{lp(\tilde{\mathbf{y}}^k)} \quad (10)$$

where $\mathbf{u}_i = (u_{i-T}, \dots, u_i)$ and $\tilde{\mathbf{y}}^k = (\tilde{y}_{i-T}, \dots, \tilde{y}_{i-T+k})$ is the current target, and $lp(\tilde{\mathbf{y}}) = \frac{(5+|\tilde{\mathbf{y}}|)^\alpha}{(5+1)^\alpha}$ is the length normalisation coefficient (Wu et al. 2016). At each time step the B most likely sequences are kept (B corresponding to the beam size).

Training objective: For training we follow Wiseman and Rush (2016) and train our model until convergence with a token level loss and fine tune it by minimising the expected risk $\mathcal{L}_{\text{RISK}}$ defined as:

$$\mathcal{L}_{\text{RISK}} = \sum_{\tilde{\mathbf{y}} \in U(C_i)} \frac{\text{cost}(\tilde{\mathbf{y}}, \mathbf{y}_i) p(\tilde{\mathbf{y}} | \mathbf{u}_i)}{\sum_{\tilde{\mathbf{y}}' \in U(\mathbf{u}_i)} p(\tilde{\mathbf{y}}' | \mathbf{u}_i)} \quad (11)$$

where $U(\mathbf{u}_i)$ is the set of the sequences generated by the model using a beam search algorithm for the input \mathbf{u}_i , and $\text{cost}(\tilde{\mathbf{y}}, \mathbf{y}_i)$ is defined, for a given a candidate sequence $\tilde{\mathbf{y}}$ and a target \mathbf{y}_i , as:

$$\text{cost}(\tilde{\mathbf{y}}, \mathbf{y}_i) = \begin{cases} 1 & \text{if } \tilde{\mathbf{y}} = \mathbf{y}_i \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

GRU/HGRU CRF baseline

State-of-the-art models use conditional random fields which model dependencies between tags on top of an GRU or a HGRU encoder which computed an embedding of the a variable number of utterances sentences . We have implemented our own CRF (Baseline_{CRF}) following the work of Kumar et al. (2018):

$$p(y_i, \dots, y_{i-T}, u_i, \dots, u_{i-T}; \theta) = \frac{\prod_{t=i}^T \psi(y_{t-1}, y_t, \phi(o_t); \theta)}{\sum_{\mathbf{y}} \prod_{t=i}^T \psi(y_t, y_{t-1}, \phi(o_t); \theta)} \quad (13)$$

Here θ is the set of parameters corresponding to the CRF layer, and ψ is the feature function, providing us with unary and pairwise potentials. Let $\phi: \mathbb{R}^H \rightarrow \mathbb{R}^{|\mathcal{Y}|}$ be the dense representation of each utterance's output provided by the encoder. ϕ can be seen as the unary feature function.

5 Experimental Protocol

In this section we describe the experimental protocols adopted for the evaluation of our approach.

Datasets

We consider two classical datasets for Dialogue Act Classification: The Switchboard Dialogue Act Corpus and the

MRDA. Since our models explicitly generate a sequence of tags we compute the accuracy on the last generated tag. Both datasets are already segmented in utterances and each utterance is segmented in words. For each dataset, we split each conversation C_i in sequence of utterances of length $T = 5^2$.

SwDA: The Switchboard-1 corpus is a telephone speech corpus (Stolcke et al. 1998), consisting of about 2.400 two-sided telephone conversation among 543 speakers with about 70 provided conversation topics. The dataset includes information about the speakers and the topics and has 42 different tags. In this dataset global dependency plays a key role due to the large amount of backchannel (19%), abandoned or turn-exit (5%), uninterpretable acts (1%). In this context, any models that only take into account local dependencies will fail at extracting information to distinguish between ambiguous tags. For the confusion matrix, we follow Li et al. (2018a) and present it for 10 tags only: statement-non-opinion (s_d), backchannel (b), statement-opinion (s_v), conventional-closing (f_c), wh-question (q_w), response acknowledgement (b_k), hedge (h), open-question (q_o), other answers (n_o), thanking (f_t).

MRDA: MRDA: The ICSI Meeting Recorder Dialogue Act corpus (Shriberg et al. 2004) contains 72 hours of naturally occurring multi-party meetings that were first converted into 75 word level conversations, and then hand-annotated with DAs using the Meeting Recorder Dialogue Act Tagset. In this work we use 5 DAs, i.e., statements (s), questions (q), floorgrabber (f), backchannel (b), disruption (d).

Train/Dev/Test Splits: For both SwDA and MRDA we follow the official split introduced by Stolcke et al. (2000). Thus, our model can directly be compared to Li et al.; Chen et al.; Kumar et al.; Raheja and Tetreault (2018a; 2018; 2018; 2019).

Training details

All the hyper-parameters have been optimised on the validation set using accuracy computed on the last tag of the sequence. The embedding layer is initialised with pre-trained fastText word vectors of size 300 (Bojanowski et al. 2017)³, trained with subword information (on Wikipedia 2017, UMBC webbase corpus and statmt.org news dataset), and updated during training. Hyperparameter selection has been done using a random search on a fixed grid. Models have been implemented in PyTorch and trained on a single NVIDIA P100.

Parameters for SwDA: We used Adam optimizer (Kingma and Ba 2014) with a learning rate of 0.01, which is updated using a scheduler with a patience of 20 epochs and a decrease rate of 0.5. The gradient norm is clipped to 5.0, weight decay is set to $1e-5$, and dropout (LeCun, Bengio, and Hinton 2015) is set to 0.2. The maximum sequence length is set to 20. Best performing model is an encoder with size of 128 and a decoder of size 48. For VGRU_E, we use two

² T is an hyperparameter, experiments have shown that 5 leads to the best results.

³In our work we rely on same pretrained embedding word2vect (Mikolov et al. 2013) instead of GloVe (Pennington, Socher, and Manning 2014).

layers for the BiGRU layer. For hierarchical models, we use BiGRU with a single layer.

Parameters for MRDA: We used AdamW optimizer (Loshchilov and Hutter 2017) with a learning rate of 0.001, which is updated using a scheduler with a patience of 15 epochs and a decrease rate of 0.5. The gradient norm is clipped to 5.0, weight decay is set to $5e-5$, and dropout (LeCun, Bengio, and Hinton 2015) is set to 0.3. The maximum sequence length is set to 30. Best performing model is an encoder with size of 40 and a decoder with size 400. For VGRU_E we use two layers for the BiGRU layer, for hierarchical models we use BiGRU with a single layer.

6 Experiments & Results

In this section we propose a set of experiments in order to investigate the performance of our model compared to existing approaches with respect to the difficulties highlighted in the introduction.

Experiment 1: Are Seq2seq better suited to DA prediction than CRF ?

Current state of the art are built on CRF models. In this first section, we aim at comparing a seq2seq with a CRF based model. To provide a fair comparison we perform the same grid search for all models on a fixed grid. At this step, we do not use attention neither use beam search during training or inference. As shown in Table 2, with a vanilla RNN encoder the seq2seq significantly outperforms the CRF on SwDa and MRDA. With an HGRU the seq2seq exhibit significantly higher results on SwDA and reaches comparable performances on MRDA. This behaviour suggests that a model based on a seq2seq architecture tends to be achieve higher score on DA classification than a CRF based model.

| Models | SwDa | MRDA |
|---------------------------------|-------------|-------------|
| Baseline _{CRF} (+GRU) | 77.7 | 88.3 |
| seq2seq (+GRU) | 81.9 | 88.5 |
| Baseline _{CRF} (+HGRU) | 81.6 | 90.0 |
| seq2seq (+HGRU) | 82.4 | 90.0 |

Table 2: Accuracy of a seq2seq on dev test and Baseline_{CRF} on SwDA and MRDA. Bold results exhibit significant differences (p-value < 0.01) according to the Wilcoxon Mann Whitney test performed on 10 runs using different seeds.

Global dependencies analysis: In Table 3 we present two examples where our seq2seq use contextual information to disambiguate the tag and to predict the correct label. In the first example, “It can be a pain” without context can be interpreted both as statement non-opinion (s_d) or statement opinion (s_v). Our seq2seq uses the surrounding context (two sentences before) to disambiguate and assign the s_v label. In the second example, the correct tag assigned to “Oh, okay” is a response acknowledgement (b_k) and not backchannel (b). The key difference between b_k and b is that an utterance labelled with b_k has to be produced within

a question-answer context, whereas b is a *continuer*⁴. In our example, the global context this is a question/reply situation: the first speaker asks a question (“What school is it”), the second replies then, the first speaker answers to the reply. This observation reflects the fact CRF models only handle local dependencies where seq2seq models consider global ones as well.

| Utterances | G. | seq2seq | CRF |
|--|----|---------|-----|
| How long does that take you to get to work? | qw | qw | qw |
| Uh, about forty-five, fifty minutes. | sd | sd | sd |
| How does that work, work out with, uh, storing your bike and showering and all that? | qw | qw | qw |
| Yeah , | b | b | b |
| It can be a pain . | sd | sd | sv |
| It's, it's nice riding to school because it's all along a canal path, uh, Because it's just, | sd | sd | sd |
| Uh, it's along the Erie Canal up here. | sd | sd | sd |
| So, what school is it? | qw | qw | qw |
| Uh, University of Rochester. | sd | sd | sd |
| Oh, okay. | bk | bk | b |

Table 3: Example of predicted sequence of tags taken from SwDA. seq2seq is our best performing model, CRF stands for Baseline_{CRF}, G. is the groundtruth label.

| | | SwDA | | | |
|-------------------|------|-------------------|------|------------|-------------|
| Dec. | Enc. | VGRU _D | att. | soft guid. | hard guid. |
| Beam Size | | 1 | 1 | 1 | 1 |
| VGRU _E | | 81.6 | 82.1 | 82.8 | 82.9 |
| HGRU | | 82.4 | 82.3 | 83.1 | 84.0 |
| PersoHGRU | | 49.8 | 79.4 | 84.0 | 83.5 |
| | | MRDA | | | |
| Dec. | Enc. | VGRU _D | att. | soft guid. | hard guid. |
| Beam Size | | 1 | 1 | 1 | 1 |
| VGRU _E | | 88.5 | 88.5 | 88.5 | 88.5 |
| HGRU | | 90.0 | 89.9 | 90.0 | 90.2 |
| PersoHGRU | | 66.2 | 87.7 | 88.2 | 86.9 |

Table 4: Accuracy on the dev set of the different encoder/decoder combination MRDA and SwDA. For SwDA, Wilcoxon test (10 runs with different seeds) has been performed for an HGRU encoder with a decoder with *hard guided attention* against an HGRU encoder with *soft guided attention*, *soft guided attention*, with attention, without attention pairwise tests exhibit p-value < 0.01.

Experiment 2: What is the best encoder?

In Table 4, we present the results of the three encoders presented in Section 4 on both datasets. For SwDA and MRDA, we observe that a seq2seq equipped with a hierarchical encoder outperforms models with Vanilla RNN encoder, while reducing the number of learned parameters.

⁴This analysis can be supported by 5.1.1 in SwDA coder manual <https://web.stanford.edu/~jurafsky/ws97/manual.august1.html>

The VGRU_D does not play well with the PersoHGRU encoder. When combined with a *guided attention mechanism*, the PersoHGRU exhibits competitive accuracy on SwDA. However on MRDA, adding a persona layer harms the accuracy. This suggests either that the information related to the speaker is irrelevant for our task (no improvement observed while adding persona information)⁵, or that the considered hierarchy is not the optimal structure to leverage this information.

Our final model makes use of the HGRU encoder since in most of the settings it exhibits superior performance.

Experiment 3: Which attention mechanism to use?

The seq2seq encodes a source sentence into a fixed-length vector from which a decoder generates a sequence of tags. Attention forces the decoder to strengthen its focus on the source sentences that are relevant to predicting a label.

In NMT (Luong, Pham, and Manning 2015), complementing a seq2seq with attention contributes to generate better sentences. In Table 4 we see that in most the case, the use of a simple attention mechanism provides a rather small improvement with VGRU and harms a bit the performances with a HGRU encoder. In case of a seq2seq composed with a PersoHGRU and a decoder without attention the learning fails: the decrease of the training loss is relatively small and seq2seq fails to generalise. It appears that in DA classification where sequences are short (5 tags), Vanilla attention does not have as much as impact as in NMT (that have longer sequences with more complex global dependencies).

If we consider an HGRU encoder, we observe that our proposed *guided attention* mechanisms improves dev accuracy which demonstrates the importance of easing the task by using prior knowledge on the alignment between the utterances and the tags. Indeed, while decoding there is a direct correspondence between labels and utterances meaning that y_i is associated with u_i . The soft guided attention will mainly focus on the current utterance with a small additional focus on the context where hard guided attention will only consider the current utterance. Improvement due to *guided attention* demonstrates that the alignment between input/output is a key prior to include in our model.

Attention analysis: Figure 2 shows a representative example of the attention weights of the three different mechanisms. The seq2seq with a normal attention mechanism is characterised by a weight matrix far from the identity (especially the lower right part). While decoding the last tags, this lack of focus leads to a wrongly predicted label for a simple utterance: “Uh-Huh” (backchannel). Both *guided attention* mechanisms focus more on the sentence associated with the tag, at each time step, and predict successfully the last DA.

Since the *hard guided attention* decoder exhibit overall the best results (on both SwDA and MRDA) and does not require any additional parameter we will use it for our final model.

⁵Further investigations with several persona based model inspired from the work of (Li et al. 2016b) shows the same poor improvement (in terms of accuracy)

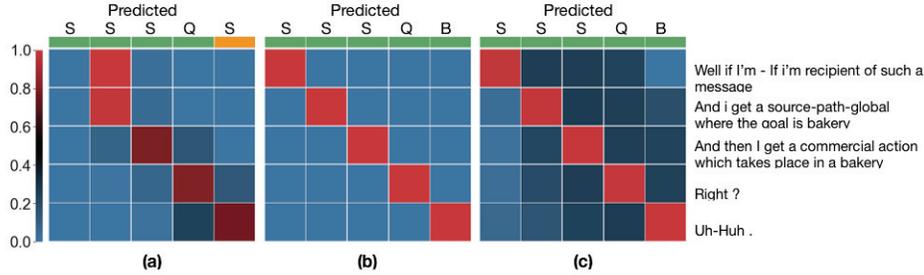


Figure 2: Attention matrix visualisation on MRDA for the fixed context of 5 utterances. Green color for predicted label indicates a correct label, orange color indicates a mistake. (a) stands for the HGRU with attention, (b) stands for the HGRU with hard guided attention, (c) is HGRU with soft guided attention.

Experiment 4: How to leverage beam search to improve the performance?

Beam Search allows the seq2seq model to consider alternative paths in the decoding phase.

Beam Search during inference: Using beam search provides a low improvement (maximum absolute improvement of 0.2%)⁶.

Compared to NMT, output size is drastically smaller ($\mathcal{Y}_{\text{SwDA}} = 42$ while $\mathcal{Y}_{\text{MRDA}} = 5$) for DA classification. When considering alternative paths with small output space in imbalanced datasets the beam search is more likely to consider very unlikely sequences as alternatives (eg. “s s s s”).

Fine tuning with a sequence loss: As previously mentioned, using beam search during inference only leads to a limited improvement in accuracy. We finetune a seq2seq composed with a HGRU encoder and a decoder with *hard guided attention* (this model has been selected in the previous steps) with the introduced sequence level loss describes in Section 4. Table 5 shows that this fine tuning steps improves the performances of 1% on SwDA (84% vs 85%) and 1.2% on RMDA (90.4% vs 91.6%).

seq2seq_{BEST}: Our *seq2seq_{BEST}* model is composed of a HGRU encoder and a decoder with *hard guided attention* finetuned with $B_{\text{train}} = 2$ and $B_{\text{inf}} = 5$ for SwDA and $B_{\text{train}} = 5$ and $B_{\text{inf}} = 1$ for SwDA.

Experiment 5: Comparison with state-of-the-art models

In this section, we compare the performances of *seq2seq_{BEST}* with other state of the art models and analyse the performances of the models. Table 6 shows the performances of best performing model seq2seq_{BEST} on the test set. seq2seq_{BEST} achieves an accuracy of 85% on the SwDA corpora. This model outperforms Chen et al. (2018) and Raheja and Tetreault (2019) which achieve

⁶The considered beam size are small compared to other applications (Li et al. 2016a). While increasing the beam size, we see that the beam search become very conservative (Gimpel et al. 2013) and tends to output labels highly represented in the training set (e.g., *sd* for SwDA).

| | | SwDA | | RMDA | |
|------------------|--------------------|-------------|------|------|------|
| B_{inf} | B_{train} | 2 | 5 | 2 | 5 |
| | 1 | | 84.8 | 84.7 | 91.3 |
| 2 | | 84.9 | 84.8 | 91.3 | 91.6 |
| 5 | | 85.0 | 84.9 | 91.5 | 91.6 |

Table 5: Accuracy on the dev set of seq2seq model trained with sequence level loss. B_{train} stands for the beam size during training, B_{inf} for the one during inference⁸. For SwDA, Wilcoxon test (10 runs with different seeds) has been performed for $B_{\text{train}} = 2$ and $B_{\text{inf}} = 2$ against all other models. For RMDA, Wilcoxon test has been performed (10 runs with different seeds) for $B_{\text{train}} = 5$ and $B_{\text{inf}} = 1$ against all model with $B_{\text{train}} = 2$.

an accuracy of 82.9%. On MRDA, our best performing model reaches an accuracy of 91.6% where current state-of-the-art systems, Chen et al.; Kumar et al. (2018; 2018) achieve respectively 92.2% and 91.7%.

| Models | SwDa | MRDA |
|-----------------------------|-------------|-------------|
| Li et al. (2018a) | 82.9 | 92.2 |
| Chen et al. (2018) | 81.3 | 91.7 |
| Kumar et al. (2018) | 79.2 | 90.9 |
| Raheja and Tetreault (2019) | 82.9 | 91.1 |
| seq2seq _{BEST} | 85.0 | 91.6 |

Table 6: Accuracy of our best models (seq2seq) and Baseline_{CRF} on SwDA and MRDA test sets.

7 Conclusion

In this work, we have presented a novel approach to the DA classification problem. We have shown that our seq2seq model, using a newly devised *guided attention* mechanisms, achieves state-of-the-art results thanks to its ability to better model global dependencies.

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *CoRR* abs/1409.0473.
- Bojanowski, P.; Grave, E.; Joulin, A.; and Mikolov, T. 2017. Enriching word vectors with subword information. *Transactions of ACL* 5:135–146.
- Bothe, C.; Weber, C.; Magg, S.; and Wermter, S. 2018. A context-based approach for dialogue act recognition using simple recurrent neural networks. *CoRR* abs/1805.06280.
- Chen, Z.; Yang, R.; Zhao, Z.; Cai, D.; and He, X. 2018. Dialogue act recognition via crf-attentive structured network. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 225–234. ACM.
- Cho, K.; van Merriënboer, B.; Gülçehre, Ç.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *CoRR* abs/1406.1078.
- Gimpel, K.; Batra, D.; Dyer, C.; and Shakhnarovich, G. 2013. A systematic exploration of diversity in machine translation. In *Proceedings of EMNLP 2013*, 1100–1111.
- Janin, A.; Baron, D.; Edwards, J.; Ellis, D.; Gelbart, D.; Morgan, N.; Peskin, B.; Pfau, T.; Shriberg, E.; Stolcke, A.; et al. 2003. The icisi meeting corpus. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03).*, volume 1, I–I. IEEE.
- Jozefowicz, R.; Zaremba, W.; and Sutskever, I. 2015. An empirical exploration of recurrent network architectures. In *International Conference on Machine Learning*, 2342–2350.
- Keizer, S.; op den Akker, R.; and Nijholt, A. 2002. Dialogue act recognition with bayesian networks for dutch dialogues. In *Proceedings of the Third SIGdial Workshop on Discourse and Dialogue*.
- Khanpour, H.; Guntakandla, N.; and Nielsen, R. 2016. Dialogue act classification in domain-independent conversations using a deep recurrent neural network. In *COLING*.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kumar, H.; Agarwal, A.; Dasgupta, R.; and Joshi, S. 2018. Dialogue act sequence labeling using hierarchical encoder with crf. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *nature* 521(7553):436.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, W. B. 2016a. A diversity-promoting objective function for neural conversation models. In *HLT-NAACL*.
- Li, J.; Galley, M.; Brockett, C.; Spithourakis, G.; Gao, J.; and Dolan, B. 2016b. A persona-based neural conversation model. In *Proceedings of the 54th Annual Meeting of ACL (Volume 1: Long Papers)*, 994–1003. Association for Computational Linguistics.
- Li, R.; Lin, C.; Collinson, M.; Li, X.; and Chen, G. 2018a. A dual-attention hierarchical recurrent neural network for dialogue act classification. *CoRR*.
- Li, Z.; Cai, J.; He, S.; and Zhao, H. 2018b. Seq2seq dependency parsing. In *Proceedings of the 27th International Conference on Computational Linguistics*, 3203–3214.
- Loshchilov, I., and Hutter, F. 2017. Fixing weight decay regularization in adam. *arXiv preprint arXiv:1711.05101*.
- Luong, T.; Pham, H.; and Manning, C. D. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of EMNLP 2015*, 1412–1421. Lisbon, Portugal: Association for Computational Linguistics.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *NeurIPS*, 3111–3119.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *Proceedings of EMNLP 2014*, 1532–1543.
- Raheja, and Tetreault. 2019. Dialogue act classification with context-aware self-attention. *CoRR* abs/1904.02594.
- Shriberg, E.; Dhillon, R.; Bhagat, S.; Ang, J.; and Carvey, H. 2004. The icisi meeting recorder dialog act (mrda) corpus. In *Proceedings of the 5th SIGdial Workshop on Discourse and Dialogue at HLT-NAACL 2004*.
- Sordani, A.; Bengio, Y.; Vahabi, H.; Lioma, C.; Grue Simonsen, J.; and Nie, J.-Y. 2015. A hierarchical recurrent encoder-decoder for generative context-aware query suggestion. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, 553–562. ACM.
- Stolcke, A.; Shriberg, E.; Bates, R.; Coccaro, N.; Jurafsky, D.; Martin, R.; Meteer, M.; Ries, K.; Taylor, P.; Van Ess-Dykema, C.; et al. 1998. Dialog act modeling for conversational speech. In *AAAI Spring Symposium on Applying Machine Learning to Discourse Processing*, 98–105.
- Stolcke, A.; Ries, K.; Coccaro, N.; Shriberg, E.; Bates, R.; Jurafsky, D.; Taylor, P.; Martin, R.; Ess-Dykema, C. V.; and Meteer, M. 2000. Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational linguistics* 26(3):339–373.
- Surendran, D., and Levow, G.-A. 2006. Dialog act tagging with support vector machines and hidden markov models. In *Ninth International Conference on Spoken Language Processing*.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *NeurIPS*, 3104–3112.
- Wiseman, S., and Rush, A. M. 2016. Sequence-to-sequence learning as beam-search optimization. In *EMNLP*.
- Wu, Y.; Schuster, M.; Chen, Z.; Le, Q. V.; Norouzi, M.; Macherey, W.; Krikun, M.; Cao, Y.; Gao, Q.; Macherey, K.; et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.