



HAL
open science

Hybrid Neural Models For Sequence Modelling: The Best Of Three Worlds

Marco Dinarelli, Loïc Grobol

► **To cite this version:**

Marco Dinarelli, Loïc Grobol. Hybrid Neural Models For Sequence Modelling: The Best Of Three Worlds. 2019. hal-02416404

HAL Id: hal-02416404

<https://hal.science/hal-02416404>

Submitted on 17 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Hybrid Neural Models For Sequence Modelling: The Best Of Three Worlds

Marco Dinarelli¹ Loïc Grobol^{2,3}

(1) LIG, Bâtiment IMAG - 700 avenue Centrale - Domaine Universitaire de Saint-Martin-d'Hères

(2) Lattice CNRS, 1 rue Maurice Arnoux, 92120 Montrouge, France

(3) ALMAAnCH Inria, 2 rue Simone Iff, 75589 Paris, France

marco.dinarelli@univ-grenoble-alpes.fr, loic.grobol@gmail.com

ABSTRACT

We propose a neural architecture with the main characteristics of the most successful neural models of the last years : bidirectional RNNs, *encoder-decoder*, and the *Transformer* model. Evaluation on three sequence labelling tasks yields results that are close to the state-of-the-art for all tasks and better than it for some of them, showing the pertinence of this hybrid architecture for this kind of tasks.

KEYWORDS : Neural Networks, sequence modelling, MEDIA, WSJ, TIGER

1 Introduction

Sequence labelling is an important NLP task, many NLP problems can be modeled indeed as sequence labelling. The most traditional labelling problems are Part-Of-Speech (POS) tagging, syntactic chunking, Named Entity Recognition (NER) (Collobert *et al.*, 2011), and also Spoken Language Understanding (SLU) in the context of human-machine dialog systems (De Mori *et al.*, 2008).

Moreover there are other NLP problems that can be decomposed into several processing steps, the first of which can be often modeled as a sequence labelling problem.

We can place among such problems the syntactic constituency parsing, which can be decomposed into POS tagging and constituents analysis (Collins, 1997); the coreference chain resolution (Soon *et al.*, 2001; Ng & Cardie, 2002), composed of mention recognition from raw text and coreferent mention detection; and also the structured named entity recognition (Grouin *et al.*, 2011; Dinarelli & Rosset, 2012a,b).

More widely, machine translation and syntactic constituency parsing can be modeled as end-to-end sequence prediction problems (Sutskever *et al.*, 2014; Bahdanau *et al.*, 2014; Vaswani *et al.*, 2017; Vinyals *et al.*, 2015), as well as a wide class of language understanding tasks (Devlin *et al.*, 2018). It would be possible thus, at least in principle, to design a model for sequence prediction in a multi-task learning framework, in order to deal with most of the NLP tasks, which shows in turn the interest in the research for sequence modelling problems.

The NLP problems that can be decomposed into several processing steps can be solved also

with a single model. For example the syntactic constituency parsing can be performed end-to-end with a model performing POS tagging and constituent analysis at the same time, e.g. (Rush *et al.*, 2012). However, even in such cases, pre-training representations with neural networks and fine-tuning them afterwards on simpler tasks, like sequence labelling, can still improve the prediction performance (Peters *et al.*, 2018). We can also consider this training procedure similar to embedding pre-training, which effectiveness has already been proved (Lample *et al.*, 2016; Ma & Hovy, 2016).

In this paper we focus on neural architectures for sequence modelling in the sense this is meant in problems such as POS tagging, morpho-syntactic analysis, and semantic chunking and tagging. The latter is performed for SLU tasks in the context of human-machine dialog systems (De Mori *et al.*, 2008).

We inspired our work from (Chen *et al.*, 2018), from which we also inspired the title of our paper, and which proposes neural hybrid models combining the *Encoder-Decoder* (or *Sequence-to-Sequence*) and *Transformer* architectures. We propose an architecture with the main characteristics of the most successful neural models of the last few years : the bidirectional RNNs (Lample *et al.*, 2016; Ma & Hovy, 2016), the *Sequence-to-Sequence* model (Sutskever *et al.*, 2014; Bahdanau *et al.*, 2014), and the *Transformer* (Vaswani *et al.*, 2017).

We evaluate our architecture on three traditional tasks of sequence labelling : Spoken Language Understanding (SLU) in French (MEDIA corpus) (Bonneau-Maynard *et al.*, 2006), POS tagging in English (WSJ corpus) (Marcus *et al.*, 1993), and morpho-syntactic tagging in German (TIGER corpus) (Brants *et al.*, 2004). The results outperform often the state-of-the-art, and allow in any case to conclude that our architecture finds its place among neural architectures.

2 Neural Architectures

The proposed neural models are inspired from the *LSTM+CRF* architectures (Lample *et al.*, 2016; Ma & Hovy, 2016) in the way input encoding (tokens, characters, etc.) is performed. They are inspired from the *Sequence-to-Sequence* architecture for the overall architecture, and from the models proposed in our previous work (Dinarelli & Tellier, 2016a,b; Dinarelli *et al.*, 2017; Dupont *et al.*, 2017; Dinarelli & Grobol, 2019) for making predictions based on a bidirectional context on the output side (labels). Afterwards we added to this architecture some of the characteristics of the *Transformer* model.

This work constitutes a step ahead with respect to (Dinarelli & Grobol, 2019). While we have not run full experiments to show the improvements with respect to our previous work, and results are equivalent in some cases, some outcomes suggest that the model proposed here is more effective. In any case the model described here is more general, as the architecture has been slightly modified so that it can be applied also to tasks like machine translation and syntactic parsing, though we do not evaluate on these tasks in this work. Moreover in other cases results are state-of-the-art, though they are obtained on data (TIGER) on which we did not evaluate previously.

2.1 Encoder

The encoder of our network is made of a bidirectional GRU (Cho *et al.*, 2014), which takes as input a sequence S^{lex} representing input tokens as the concatenation of word embeddings ($E_w(w_i)$) and representations computed from their characters $h_c(w_i)$. The latter is computed by an *ad-hoc* recurrent layer, in a fashion similar to those used in the literature Ma & Hovy (2016), with the difference that we use a GRU recurrent layer instead of a convolution layer.

The character-level representation of a word w is computed like in (Dinarelli & Grobol, 2019) :

$$\begin{aligned} S^c(w) &= (E_c(c_{w,1}), \dots, E_c(c_{w,n})) \\ (h_c(c_{w,1}), \dots, h_c(c_{w,n})) &= \text{GRU}_c(S^c(w), h_c^0) \\ h_c(w) &= \text{FFNN}(\text{Sum}(h_c(c_{w,1}), \dots, h_c(c_{w,n}))) \end{aligned} \quad (1)$$

where E_c is the matrix of character embeddings, $c_{w,i}$ is the i -th character of the word w , $S^c(w)$ is the sequence of character embeddings for the word w , GRU_c is the GRU layer for characters, $h_c(c_{w,i})$ is the hidden state associated to the i -th character of the word w . GRU_c , like GRU_w , is a bidirectional GRU layer.

The full representation of the i -th word w_i in a sequence is computed as :

$$\begin{aligned} S^{lex} &= ([E_w(w_1), h_c(w_1)], \dots, [E_w(w_N), h_c(w_N)]) \\ (h_{w_1}, \dots, h_{w_N}) &= \text{GRU}_w(S^{lex}, h_w^0) \end{aligned} \quad (2)$$

Since GRU_w scans the sequence forward and backward, h_{w_i} depends on the word w_i and its context. When additional features are available at token level, they are embedded in the same way as words and concatenated to word embeddings in the sequence S^{lex} .

2.2 Decoders

Our model uses a representation of both left and right context at label level like proposed in Dinarelli & Tellier (2016a); Dinarelli *et al.* (2017); Dupont *et al.* (2017); Dinarelli & Grobol (2019).

We use a $\overleftarrow{\text{GRU}}_e$ *backward* layer for encoding the right context, and a $\overrightarrow{\text{GRU}}_e$ *forward* layer for the left context with respect to a given labelling position i in the input sequence. These layers take as input the lexical level representation computed by the encoder and the label embeddings $E_e(e_i)$, which makes them similar to the decoder used in the original *Sequence-to-Sequence* architecture for machine translation Sutskever *et al.* (2014); Bahdanau *et al.* (2014). An evolution with respect to such architecture is the use of two decoders, one for the left label-level context and one for the right label-level context. This idea of using both left and right context at the output item's level has attracted some interest in the last few years, and has been applied in different variants to machine translation (Xia *et al.*, 2017; Zhang *et al.*, 2018a).

The right label-level context is computed by the *backward* $\overleftarrow{\text{GRU}}_e$ decoder as follows :

$$\overleftarrow{h}_{e_i} = \overleftarrow{\text{GRU}}_e([h_{w_i}, E_e(e_{i+1})], \overleftarrow{h}_{e_{i+1}}) \quad (3)$$

The left label-level context \overrightarrow{h}_{e_i} is computed in a similar way by the *forward* $\overrightarrow{\text{GRU}}_e$ decoder.

2.3 Output Layer

In order to allow the model to make global decisions when predicting the current label, we add an output layer on top of the backward decoder, made of an affine transformation followed by a *log-softmax* activation. This computes log-probabilities of backward predictions :

$$\begin{aligned} \log\text{-P}(\overleftarrow{e}_i) &= \log\text{-softmax}(W_{bw}[h_{w_i}, \overleftarrow{h}_{e_i}] + b_{bw}) \\ \overleftarrow{e}_i &= \text{argmax}(\log\text{-P}(\overleftarrow{e}_i)) \end{aligned} \quad (4)$$

The forward decoder predicts labels using as input the label-level context representations \overrightarrow{h}_{e_i} and \overleftarrow{h}_{e_i} , as well as the lexical level information h_{w_i} computed by the encoder :

$$\begin{aligned} \log\text{-P}(\overrightarrow{e}_i) &= \log\text{-softmax}(W_o[\overrightarrow{h}_{e_i}, h_{w_i}, \overleftarrow{h}_{e_i}] + b_o) \\ \overrightarrow{e}_i &= \text{argmax}(\log\text{-P}(\overrightarrow{e}_i)) \end{aligned} \quad (5)$$

In order to strengthen the global character of the decision, the final output of the model is computed as the arithmetic mean (in the log space, which gives a geometric mean of probabilities) of the *forward* and *backward* log-probabilities : $\frac{1}{2}(\log\text{-P}(\overrightarrow{e}_i) + \log\text{-P}(\overleftarrow{e}_i))$. This *teaches* the model to compute real good predictions also in the backward phase, instead of basing decisions on raw heuristics, or even on the forward (left) context only. Indeed, the backward decoder has only a partial view of the output level information, while the forward decoder uses both left and right context.

2.4 Learning

All models described in this paper are learned minimizing the minus log-likelihood \mathcal{LL} of the training data. More precisely :

$$-\mathcal{LL}(\Theta|D) = -\sum_{d=1}^{|D|} \sum_{i=1}^{N_d} \frac{1}{2} (\log\text{-P}(\overrightarrow{e}_i) + \log\text{-P}(\overleftarrow{e}_i)) + \frac{\lambda}{2} |\Theta|^2 \quad (6)$$

The first sum scans the data D of size $|D|$, while the second sum is over training examples S_d , of size N_d .

Since data used in this work for evaluation are relatively small, while our models are relatively complex, we add a Gaussian regularization term L_2 with a coefficient λ .

2.5 *The best of three worlds*

The neural architecture described so far uses the same ideas introduced in (Dinarelli & Teller, 2016a; Dinarelli *et al.*, 2017; Dupont *et al.*, 2017; Dinarelli & Grobol, 2019) for predicting labels using both representations of left (forward) and right (backward) contexts, and for both input-level information (words, characters, etc.) and labels. Our model integrates also the characteristics of bidirectional RNNs and of Encoder-Decoder architectures. Moreover we use two decoders, instead of one like in the original architecture.

Starting from this neural architecture we added also some of the characteristics of the *Transformer* model.

We note that the paper from which we inspired our work (Chen *et al.*, 2018) investigated the combination of two Encoder-Decoder architectures, one based on recurrent layers (Sutskever *et al.*, 2014; Bahdanau *et al.*, 2014), while the other is based on the *Transformer* model, which is not a recurrent network. This work shows that such combination is possible and effective compared to the two combined architectures.

In this work we combined three different architectures with the goal of putting together their strengths, and trying to overcome some limitations.

First of all, in the tasks of sequence labelling used for evaluation in this work, there is a *one-to-one* correspondence between input and output items (words and labels, respectively). Our preliminary experiments on mention detection for coreference resolution using a *Transformer*, as well as some experiments in the literature (Guo *et al.*, 2019), suggest that ignoring such correspondence leads to a remarkable performance drop.¹ For this reason the encoder and the decoders used in our architecture exploit the one-to-one correspondence of sequence labelling tasks. We are aware though this information could not be used in *real* sequence-to-sequence problems such like machine translation or syntactic parsing.

Second, we chose an Encoder-Decoder overall architecture since this proved to be the most effective for sequence-to-sequence problems (including sequence labelling). Moreover, we showed that keeping a label-level context into account, in some sequence labelling problems, using a decoder is more effective than doing the same with a CRF neural layer (Dinarelli *et al.*, 2017), we will show the same also in this work. As we already mentioned, in our architecture we use two decoders. This solution allows for an even more effective modelling of a label-level context. This solution is not possible, at least in principle, with the original *Transformer* model.

Third, we decided to integrate some of the characteristics of the *Transformer* in order to overcome limitations of RNNs concerning the learning phase. As discussed in (Vaswani *et al.*, 2017), RNNs are difficult to train because of the length of paths the learning signal has to back-propagate through. The length of such paths is the consequence of the recurrent nature of the networks.

The *Transformer* model (Vaswani *et al.*, 2017) proposed an alternative, non-recurrent Encoder-Decoder architecture based on a multi-head attention mechanism. Recent works (Dehghani *et al.*, 2018; Dai *et al.*, 2019) suggest however that using some kind of recurrent connections *Transformer*'s performance can be improved.

1. F1 measure drops more than 3 points using the *Transformer* in our experiments

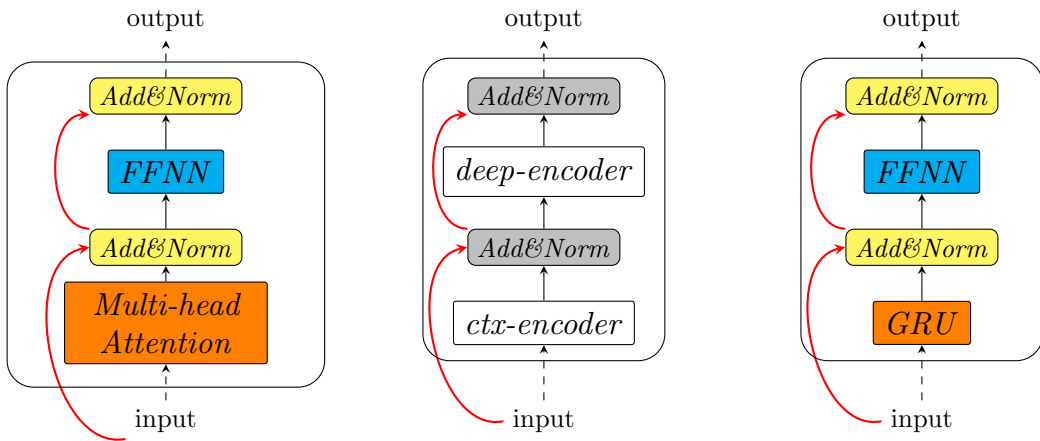


FIGURE 1 – High-level architecture, from a conceptual point-of-view, of each block of a *Transformer* model (in the middle). On the left we show its instantiation in the *Transformer*, on the right we show its instantiation in our architecture, where we use a bidirectional GRU layer to encode contextual features

Another interesting characteristic of the *Transformer* is the use of residual connections, or skip-connections (Bengio *et al.*, 2003). This characteristic allows for alleviating the vanishing gradient problem of deep networks. Indeed, the use of gates in recurrent networks reduces the amplitude of this problem, however the literature and our own experiments suggest that for long sequences, gates are not sufficient for an effective learning.

The attention mechanism does not suffer from this problem. Each element of the input sequence is related to a small, fixed number of layers in the attention mechanism. The back-propagation path of the learning signal is thus much shorter than that of RNNs. Moreover the use of skip-connections contributes in keeping a higher learning signal amplitude, allowing to skip a whole layer, over which the learning signal would be weakened (because of gradient multiplications).

We can interpret each block of a *Transformer* architecture, both encoder and decoder, as the combination of a sub-module which “*encodes*” contextual features, we call this sub-module *ctx-encoder*, and a *feed-forward* sub-module which “*maps*” contextual features into a deeper feature space, we call the second sub-module *deep-encoder*. The output of the two sub-modules is added to their respective input (skip-connection) and normalized with a *layer normalization*. This interpretation is shown in the figure 1 in the middle. Skip-connections are highlighted in red.

The *Transformer* proposed in Vaswani *et al.* (2017) instantiate the generic architecture described above using as *ctx-encoder* the multi-head attention mechanism, while it uses a two-layer feed-forward network as *deep-encoder*. A single block of this architecture is shown in the figure 1 on the left.

We modified our neural architecture for instantiating the generic architecture. In our solution we use a bidirectional GRU layer as *ctx-encoder*, while we use exactly the same sub-module




FIGURE 2 – Overall architecture of our model, this is the same as the model presented in our previous work Dinarelli & Grobol (2019)

as the *Transformer* as *deep-encoder*. A block of our architecture is shown in the figure 1 on the right.

The characteristics of the *Transformer* model we integrated in our architecture are thus the skip-connections and layer normalization (*Add&Norm* in the figure 1), as well as the two-layer feed-forward sub-module to re-encode contextual features once the skip-connection has been applied (*FFNN* in the figure 1). Following the computation chain suggested by Chen *et al.* (2018), the output of our encoder is computed as follows :

$$\begin{aligned}
 \hat{S}_i^{lex} &= \text{Norm}(S_i^{lex}) \\
 \hat{h}_{w_i} &= \text{GRU}_w(\hat{S}_i^{lex}, h_{w_{i-1}}) \\
 h_{w_i} &= \text{FFNN}(\text{Norm}(\text{Dropout}(\hat{h}_{w_i}) + \hat{S}_i^{lex}))
 \end{aligned}
 \tag{7}$$

where we use *Norm* for the layer normalization, and *Dropout* for a dropout regularization (Srivastava *et al.*, 2014). The other GRU layers of our architecture, introduced in previous sections, are modified in a similar way.

In this work we don't use any attention mechanism like in the *Transformer*, context is granted by the use of GRU layers (cf. section 2.1). Our choice is motivated by the work described in Levy *et al.* (2018), where it is shown that RNNs, such like GRU, can encode long contexts as a weighted mean of the inputs, in a way that can be considered similar to an attention mechanism. Moreover, using GRU layers avoid the use of positional encoding, the sequential structure of an input sentence is encoded implicitly by the RNNs.²

The overall architecture of our model is the same as our previous work (Dinarelli & Grobol, 2019) (cf figure 2), the difference is in the computation of layer's output described in previous sections.

2. We don't exclude however that using positional encoding would improve our architecture

| | Training | | Dev | | Test | |
|-------------|----------|----------|--------|----------|--------|----------|
| # sentences | 12 908 | | 1 259 | | 3 005 | |
| | Words | Concepts | Words | Concepts | Words | Concepts |
| # tokens | 94 466 | 43 078 | 10 849 | 4 705 | 25 606 | 11 383 |
| dictionary | 2 210 | 99 | 838 | 66 | 1 276 | 78 |
| OOV% | – | – | 1,33 | 0,02 | 1,39 | 0,04 |

TABLE 1 – Statistics of the French corpus MEDIA

| | Training | | Dev | | Test | |
|-------------|----------|--------|---------|--------|---------|--------|
| # sentences | 38 219 | | 5 527 | | 5 462 | |
| | Words | Labels | Words | Labels | Words | Labels |
| # tokens | 912 344 | – | 131 768 | – | 129 654 | – |
| dictionary | 43 210 | 45 | 15 081 | 45 | 13 968 | 45 |
| OOV% | – | – | 22,61 | 0 | 20,00 | 0 |

TABLE 2 – Statistics of the English corpus WSJ

3 Evaluation

We evaluate our models on three tasks of sequence labelling :

The French corpus MEDIA (Bonneau-Maynard *et al.*, 2006) was created for evaluating spoken dialog systems in the context of hotel information and reservation in France. Data have been annotated following a rich concept ontology. Simple semantic components can be combined to create more precise and complex semantic units.³ Concepts span in general more than one word, the goal is thus to tag words and then reconstruct semantic units from the tags. Statistics of the corpus MEDIA are shown in figure 1.

This task is modelled as sequence labelling splitting semantic units over multiple words using the *BIO* formalism (Ramshaw & Marcus, 1995). We already evaluated on this task in the past, using a large variety of probabilistic models (Dinarelli *et al.*, 2009a,b; Quarteroni *et al.*, 2009; Hahn *et al.*, 2010; Dinarelli, 2010; Dinarelli & Rosset, 2011; Dinarelli *et al.*, 2011).

In the MEDIA corpus some word classes (clusters) are provided with the annotation. Such classes allow for a better generalization on words belonging to categories for which lists can be easily retrieved, or constructed, in the context of a real human-machine dialog system. Examples of such classes are the name of cities in France (class *CITY*), the hotel brands (class *HOTEL*), currency amounts (*AMOUNTS*), etc. We used these classes in some of our experiments, this is indicated with *FEAT* in the tables of the following section.

The English Penn TreeBank corpus (Marcus *et al.*, 1993), named WSJ henceforth, is one of the most used corpus for evaluating models for sequence labelling. The task consists in assigning each word its POS tag. For our experiments we use the most common data split : sections 0-18 are used as training data, sections 19-21 as development data, and sections 22-24 are used for testing. Statistics of this corpus are given in the table 2.

The German corpus TIGER (Brants *et al.*, 2004) is annotated with rich morpho-syntactic information. These include POS tags like in the WSJ corpus, but also gender, number, cases, and other inflection information, as well as conjugation information for verbs. The

3. For example the component *localization* (we translate from French) can be combined with *city*, *relative-distance*, *general-relative-localization*, *street*, etc.

| | Training | | Dev | | Test | |
|-------------|----------|--------|--------|--------|--------|--------|
| # sentences | 40 472 | | 5 000 | | 5 000 | |
| | Words | Labels | Words | Labels | Words | Labels |
| # tokens | 719 530 | – | 76 704 | – | 92 004 | – |
| dictionary | 77 220 | 681 | 15 852 | 501 | 20 149 | 537 |
| OOV% | – | – | 30,90 | 0,01 | 37,18 | 0,015 |

TABLE 3 – Statistics of the German corpus TIGER

combination of all these pieces of information constitutes the output labels. This task is similar to the POS tagging from a sequence labelling point of view, however it is made more difficult on one side by the language (German has more complex syntactic structures than English), on the other side by the relatively large number of labels to be predicted (694 in total, while there are 138 labels in MEDIA and 45 in the WSJ corpus). We use the same data split used in (Lavergne & Yvon, 2017). Statistics of this corpus are given in the table 3. As we can see from tables, the size of dictionaries, for both words and labels, is relatively large in the WSJ and TIGER corpora.

We note that the Out-Of-Vocabulary (OOV) rate in the MEDIA corpus is quite small. Indeed there are only two OOV words in both Dev and Test data, moreover they are function words. There is basically no unknown word in this task. In contrast the OOV rate in both the WSJ and the TIGER corpora is quite high : roughly one word out of five in the first and one out of three in the second.

3.1 Settings

We used the corpus MEDIA for fine-tuning the hyper-parameters of our models in (Dinarelli & Grobol, 2019). MEDIA is the smallest corpus among the three used, and allows thus a faster optimization. The settings used in this work are the same as those used in (Dinarelli *et al.*, 2017; Dinarelli & Grobol, 2019), where the same data were used for evaluation (except for the TIGER corpus). The settings for the WSJ corpus are the same, except for the word embeddings (300 dimensions) and the learning rate ($2,5 \times 10^{-4}$). We use the same settings for WSJ and TIGER. Since the latter corpus is the largest, we ran few experiments to choose the correct hidden layer size (300).

Like we already explained in (Dinarelli & Grobol, 2019), at first our experiments were not fitting in the memory of our GPUs. In order to solve this problem we found two solutions. In the first we organize the training data as a single stream of items. This is then split into chunks of fixed size, shifting each chunk by one token with respect to the previous (thus consecutive chunks overlap). The second solution is more intuitive and consists in clustering sentences of the same length. This way, very long (rare) sentences are assigned to very small clusters, thus fitting in memory, and small and average clusters form bigger clusters, which still fit in memory. In our experiments we found out that the first solution works far better on the MEDIA corpus. On the WSJ and TIGER corpora the two solutions are equivalent, we thus prefer the second for these two tasks, which is more intuitive and is expected to work better on bigger corpora.

Our neural network system is coded with pytorch⁴, experiments are run under Linux Debian

4. <https://pytorch.org/docs/stable/index.html>

| Model | Accuracy | F1 | CER |
|---|--------------|--------------|-------------|
| MEDIA DEV | | | |
| Seq2Biseq (Dinarelli & Grobol, 2019) | 89.97 | 86.57 | 10.42 |
| Seq2Biseq _{2-opt} (Dinarelli & Grobol, 2019) | 90.22 | 86.88 | 9.97 |
| Seq2Biseq-Transformer _{2-opt} | 90.08 | 86.93 | 9.87 |
| Seq2Biseq-Transformer+FEAT _{2-opt} | 90.14 | 87.05 | 9.54 |
| MEDIA TEST | | | |
| BiGRU+CRF (Dinarelli <i>et al.</i> , 2017) | – | 86.69 | 10.13 |
| LD-RNN _{deep} (Dinarelli <i>et al.</i> , 2017) | – | 87.36 | 9.8 |
| Seq2Biseq (Dinarelli & Grobol, 2019) | 89.57 | 87.50 | 10.26 |
| Seq2Biseq _{2-opt} (Dinarelli & Grobol, 2019) | 89.79 | 87.69 | 9.93 |
| Seq2Biseq-Transformer _{2-opt} | 89.98 | 87.81 | 9.67 |
| Seq2Biseq-Transformer+FEAT _{2-opt} | 90.12 | 87.94 | 9.48 |

TABLE 4 – Results of different variants of our model for the semantic tagging task of the corpus MEDIA, compared to the state-of-the-art

distribution with GeForce GTX 1080 GPUs.

3.2 Results

Results for the MEDIA task are means over 10 experiments, the parameters are initialized at random⁵ for each experiment. On this task we evaluate our models in terms of accuracy and, since the goal is to re-construct semantic units spanning over multiple words, also with F1 measure and *Concept Error Rate* (CER). CER is computed by aligning the model prediction with the gold annotation with the Levenshtein distance algorithm, then the sum of insertions, deletions and substitutions is divided by the length of the gold annotation. On the other two tasks, since there is no label splitting over multiple words, we evaluate the models only in terms of accuracy.

Since the model presented in this paper is an evolution of the *Seq2Biseq* model proposed in (Dinarelli & Grobol, 2019), the model proposed here is named *Seq2Biseq-Transformer*, as it adds characteristics of the *Transformer* on top of the *Seq2Biseq* model.

Results on Dev and Test data of the MEDIA task are shown in table 4. For comparison we show also results presented in our previous works (Dinarelli *et al.*, 2017; Dinarelli & Grobol, 2019), which are state-of-the-art on this task.

As described in (Dinarelli & Grobol, 2019), the Seq2Biseq model gives better results when trained with 2 optimizers (see (Dinarelli & Grobol, 2019) for details), one for each decoder. The Seq2Biseq-Transformer model is thus always trained with 2 optimizers.

As we can see in table 4, results obtained on MEDIA with the Seq2Biseq-Transformer model are slightly worse than Seq2Biseq’s results in terms of accuracy, and they are slightly better in terms of F1 and CER.⁶ Adding word classes available for the MEDIA task (cf. the begin

5. Using default initialization of the *Pytorch* library

6. We did not perform significance tests

| Model | Accuracy | |
|---|--------------|--------------|
| | WSJ DEV | WSJ TEST |
| LD-RNN _{deep} (Dinarelli <i>et al.</i> , 2017) | 96.90 | 96.91 |
| LSTM+CRF (Ma & Hovy, 2016) | – | 97.13 |
| Seq2Biseq (Dinarelli & Grobol, 2019) | 97.13 | 97.20 |
| Seq2Biseq _{2-opt} (Dinarelli & Grobol, 2019) | 97.33 | 97.35 |
| Seq2Biseq-Transformer _{2-opt} | 97.33 | 97.36 |
| Models using pre-trained embeddings ⁸ | | |
| Transformer+Char (Guo <i>et al.</i> , 2019) | – | 97.14 |
| LSTM+CRF + Glove (Ma & Hovy, 2016) | 97.46 | 97.55 |
| LSTM+LD-RNN + Glove (Zhang <i>et al.</i> , 2018b) | – | 97.59 |
| Star-Transformer+Char (Guo <i>et al.</i> , 2019) | – | 97.64 |

TABLE 5 – Results of different variants of our model for the POS tagging task on the WSJ corpus, compared to the state-of-the-art

of section 3, data description) as input to the model, results improve with all evaluation metrics and constitute the new state-of-the-art on this task with manual transcription.⁷

Results of POS tagging on the WSJ corpus are presented in table 5. As we can see in this case, adding functionalities of the *Transformer* to the Seq2Biseq model give exactly the same result as the Seq2Biseq model on the Dev data, and slightly better on the Test data (on a single experiment).

Keeping both these results and those on the MEDIA task into account, we could conclude that the Seq2Biseq-Transformer model does not really improve our previous model. However we would like to point out that the main goal of this work was to integrate characteristics of the *Transformer* model, and to come up with a generic architecture, which can be applied not only to traditional sequence labelling tasks, but also to tasks such like machine translation and syntactic parsing. Indeed the Seq2Biseq-Transformer architecture has been modified, with respect to our previous model, in order to perform all these tasks. We consider that obtaining similar results (which are still state-of-the-art) with a more versatile model is still a positive result. Additionally, observing the training procedure of the Seq2Biseq-Transformer model, we found out that while final results are similar, the loss values on both training and Dev data are better than those obtained with the Seq2Biseq model. This could actually mean the model is over-fitting, and we speculate that instead of using the same settings as in (Dinarelli & Grobol, 2019), since the Seq2Biseq-Transformer is more complex (at least in terms of parameters, and see below), a fine-tuning of parameters could lead to better performances.

In table 5 we can see that, whether we use one or two decoders for training, our models improve the *LSTM+CRF* model (Ma & Hovy, 2016) without using the embeddings pre-trained with *GloVe* (Pennington *et al.*, 2014). For a matter of comparison, we show also the best results on this task, obtained with pre-trained embeddings. Though our results are lower than state-of-the-art in terms of accuracy, they are not too far, proving that our hybrid neural

⁷. SLU is performed on automatic transcription, in this work however we evaluate our models against sequence labelling problems

| Model | Accuracy | |
|--------------------------------|----------------------|----------------------|
| | TIGER DEV | TIGER TEST |
| Seq2Biseq _{2-opt} | 93.90 (98.30) | 91.86 (97.74) |
| VO-CRF (Lavergne & Yvon, 2017) | – | 88.78 |

TABLE 6 – Results of different variants of our neural architecture for the morpho-syntactic tagging task on the TIGER corpus, compared with state-of-the-art

architecture is quite effective. In particular, we compare our results with those reported in (Guo *et al.*, 2019) obtained with the *Transformer* and *Star-Transformer* models.⁹ Whether the *Transformer* model has been trained with or without pre-trained embeddings, our models outperform it by a significant margin. As we mentioned in section 2.5, we have experienced a drop of 3 points in the F1 measure using a *Transformer* for mention detection in preliminary experiments (not reported in this work). We hypothesize that the one-to-one correspondence between input and output units in sequence labelling tasks is a very useful information, and the lack of performance of the *Transformer* on sequence labelling tasks can be explained by the fact that this model doesn’t use it.

In table 6 we show results obtained on the morpho-syntactic tagging task on the German corpus TIGER. To the best of our knowledge, the best results in the literature for this task are reported in Lavergne & Yvon (2017), obtained using a *Variable Order* CRF (VO-CRF). We can see thus that our model improves the state-of-the-art on this task. It is important to note however that the TIGER corpus has been less used for the evaluation of neural models for sequence labelling. Results reported in Lavergne & Yvon (2017) are indeed obtained with a discrete-representation based, though quite sophisticated, CRF model. In table 6 we report also, between parenthesis, POS tagging results. These are obtained from morpho-syntactic tagging results extracting the POS tag from the complete label. That is no model training has been performed for the POS tagging task alone.

4 Conclusions

In this paper we proposed a neural model for sequence labelling integrating characteristics of the most successful models of the last couple of years : bidirectional RNNs, *Sequence-to-Sequence* model, and the *Transformer* model. An evaluation on three tasks of sequence labelling shows that the proposed model is very effective for this kind of tasks. Indeed it achieves results comparable to the state-of-the-art, but keeping a generic architecture that allows to use the same model on other kind of tasks like machine translation and syntactic parsing. Moreover in some cases our model outperforms state-of-the-art models, proving that our hybrid architecture has its place among current neural architectures.

9. In (Guo *et al.*, 2019) it is actually not clear if the *Transformer* and *Star-Transformer* models have been trained with or without pre-trained embeddings. Authors report “*Advanced Techniques* except widely-used pre-trained embeddings” in the caption of table 5. However results are compared with literature models using pre-trained embeddings, and this is not mentioned in the text.

Acknowledgments

This research is part of the program “Investissements d’Avenir” funded by the French National Research Agency (ANR) ANR-10-LABX-0083 (Labex EFL).

This work has been partially funded by the French ANR project DEMOCRAT (“*Description et modélisation des chaînes de référence : outils pour l’annotation de corpus et le traitement automatique*”) project ANR-15-CE38-0008.

References

- BAHDANAU D., CHO K. & BENGIO Y. (2014). Neural machine translation by jointly learning to align and translate. *CoRR*, **abs/1409.0473**.
- BENGIO Y., DUCHARME R., VINCENT P. & JAUVIN C. (2003). A neural probabilistic language model. *JOURNAL OF MACHINE LEARNING RESEARCH*, **3**, 1137–1155.
- BONNEAU-MAYNARD H., AYACHE C., BECHET F., DENIS A., KUHN A., LEFÈVRE F., MOSTEFA D., QUGNARD M., ROSSET S. & SERVAN, S. VILANEAU J. (2006). Results of the french evalda-media evaluation campaign for literal understanding. In *LREC*, p. 2054–2059, Genoa, Italy.
- BRANTS S., DIPPER S., EISENBERG P., HANSEN-SCHIRRA S., KONIG E., LEZIUS W., ROHRER C., SMITH G. & USZKOREIT H. (2004). TIGER : Linguistic interpretation of a german corpus. *Research on Language and Computation*, **2**(4), 597–620.
- CHEN M. X., FIRAT O., BAPNA A., JOHNSON M., MACHEREY W., FOSTER G., JONES L., SCHUSTER M., SHAZEER N., PARMAR N., VASWANI A., USZKOREIT J., KAISER L., CHEN Z., WU Y. & HUGHES M. (2018). The Best of Both Worlds : Combining Recent Advances in Neural Machine Translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, volume 1, p. 76–86 : Association for Computational Linguistics.
- CHO K., VAN MERRIENBOER B., GULCEHRE C., BAHDANAU D., BOUGARES F., SCHWENK H. & BENGIO Y. (2014). Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, p. 1724–1734 : Association for Computational Linguistics.
- COLLINS M. (1997). Three generative, lexicalised models for statistical parsing. In *Proceedings of ACL*, p. 16–23, Stroudsburg, PA, USA : Association for Computational Linguistics.
- COLLOBERT R., WESTON J., BOTTOU L., KARLEN M., KAVUKCUOGLU K. & KUKSA P. (2011). Natural language processing (almost) from scratch. *J. Mach. Learn. Res.*, **12**.
- DAI Z., YANG Z., YANG Y., CARBONELL J., LE Q. V. & SALAKHUTDINOV R. (2019). Transformer-XL : Attentive Language Models Beyond a Fixed-Length Context. *arXiv preprint 1901.02860*.
- DE MORI R., BECHET F., HAKKANI-TUR D., MCTEAR M., RICCARDI G. & TUR G. (2008). Spoken language understanding : A survey. *IEEE Signal Processing Magazine*, **25**, 50–58.

DEGHANI M., GOUWS S., VINYALS O., USZKOREIT J. & KAISER L. (2018). Universal transformers. *CoRR*, **abs/1807.03819**.

DEVLIN J., CHANG M.-W., LEE K. & TOUTANOVA K. (2018). BERT : Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint 1810.04805*.

DINARELLI M. (2010). *Spoken Language Understanding : from Spoken Utterances to Semantic Structures*. PhD thesis, International Doctoral School in Information and Communication Technology, Dipartimento di Ingegneria e Scienza dell' Informazione, via Sommarive 14, 38100 Povo di Trento (TN), Italy.

DINARELLI M. & GROBOL L. (2019). Seq2biseq : Bidirectional output-wise recurrent neural networks for sequence modelling. *CoRR*, **abs/1904.04733**.

DINARELLI M., MOSCHITTI A. & RICCARDI G. (2009a). Concept segmentation and labeling for conversational speech. In *Proceedings of the International Conference of the Speech Communication Assosiation (Interspeech)*, Brighton, U.K.

DINARELLI M., MOSCHITTI A. & RICCARDI G. (2009b). Re-ranking models based on small training data for spoken language understanding. In *Conference of Empirical Methods for Natural Language Processing*, p. 11–18, Singapore.

DINARELLI M., MOSCHITTI A. & RICCARDI G. (2011). Discriminative reranking for spoken language understanding. *IEEE TASLP*, **20**, 526–539.

DINARELLI M. & ROSSET S. (2011). Hypotheses selection criteria in a reranking framework for spoken language understanding. In *Conference of Empirical Methods for Natural Language Processing*, p. 1104–1115, Edinburgh, U.K.

DINARELLI M. & ROSSET S. (2012a). Tree representations in probabilistic models for extended named entity detection. In *European Chapter of the Association for Computational Linguistics (EACL)*, p. 174–184, Avignon, France.

DINARELLI M. & ROSSET S. (2012b). Tree-structured named entity recognition on ocr data : Analysis, processing and results. In N. C. C. CHAIR), K. CHOUKRI, T. DECLERCK, M. U. DOGAN, B. MAEGAARD, J. MARIANI, J. ODIJK & S. PIPERIDIS, Eds., *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey : European Language Resources Association (ELRA).

DINARELLI M. & TELLIER I. (2016a). Improving recurrent neural networks for sequence labelling. *CoRR*, **abs/1606.02555**.

DINARELLI M. & TELLIER I. (2016b). New recurrent neural network variants for sequence labeling. In *Proceedings of the 17th International Conference on Intelligent Text Processing and Computational Linguistics*, Konya, Turkey : Lecture Notes in Computer Science (Springer).

DINARELLI M., VUKOTIC V. & RAYMOND C. (2017). Label-dependency coding in Simple Recurrent Networks for Spoken Language Understanding. In *Interspeech*, Stockholm, Sweden.

DUPONT Y., DINARELLI M. & TELLIER I. (2017). Label-dependencies aware recurrent neural networks. In *Proceedings of CICling*, Budapest, Hungary : LNCS, Springer.

GROUIN C., DINARELLI M., ROSSET S., WISNIEWSKI G. & ZWEIGENBAUM P. (2011). Coreference resolution in clinical reports. the limsi participation in the i2b2/va 2011 challenge. In *In Proceedings of i2b2/VA 2011 Coreference Resolution Workshop*.

- GUO Q., QIU X., LIU P., SHAO Y., XUE X. & ZHANG Z. (2019). Star-transformer. *CoRR*, [abs/1902.09113](#).
- HAHN S., DINARELLI M., RAYMOND C., LEFÈVRE F., LEHEN P., DE MORI R., MOSCHITTI A., NEY H. & RICCARDI G. (2010). Comparing stochastic approaches to spoken language understanding in multiple languages. *IEEE TASLP*, **99**.
- LAMPLE G., BALLESTEROS M., SUBRAMANIAN S., KAWAKAMI K. & DYER C. (2016). Neural Architectures for Named Entity Recognition. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics : Human Language Technologies*, p. 260–270 : Association for Computational Linguistics.
- LAVERGNE T. & YVON F. (2017). Learning the structure of variable-order crfs : a finite-state perspective. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, p. 433–439 : Association for Computational Linguistics.
- LEVY O., LEE K., FITZGERALD N. & ZETTLEMOYER L. (2018). Long short-term memory as a dynamically computed element-wise weighted sum. In *Proceedings of ACL*, p. 732–739 : ACL.
- MA X. & HOVY E. (2016). End-to-end sequence labeling via bi-directional lstm-cnns-crf. In *Proceedings of ACL*.
- MARCUS M. P., SANTORINI B. & MARCINKIEWICZ M. A. (1993). Building a large annotated corpus of english : The penn treebank. *COMPUTATIONAL LINGUISTICS*, **19**(2).
- NG V. & CARDIE C. (2002). Improving Machine Learning Approaches to Coreference Resolution. In *Proceedings of ACL'02*, p. 104–111.
- PENNINGTON J., SOCHER R. & MANNING C. D. (2014). Glove : Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, p. 1532–1543.
- PETERS M., NEUMANN M., IYYER M., GARDNER M., CLARK C., LEE K. & ZETTLEMOYER L. (2018). Deep Contextualized Word Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics : Human Language Technologies*, volume 1, p. 2227–2237 : Association for Computational Linguistics.
- QUARTERONI S., RICCARDI G. & DINARELLI M. (2009). What’s in an ontology for spoken language understanding. In *Proceedings of the International Conference of the Speech Communication Assosiation (Interspeech)*, Brighton, U.K.
- RAMSHAW L. & MARCUS M. (1995). Text chunking using transformation-based learning. In *Proceedings of the 3rd Workshop on Very Large Corpora*, p. 84–94, Cambridge, MA, USA.
- RUSH A. M., REICHART R., COLLINS M. & GLOBERSON A. (2012). Improved parsing and pos tagging using inter-sentence consistency constraints. In *Proceedings of EMNLP-CoNLL*, Stroudsburg, PA, USA.
- SOON W. M., NG H. T. & LIM D. C. Y. (2001). A Machine Learning Approach to Coreference Resolution of Noun Phrases. *Computational Linguistics*, **27**(4), 521–544.
- SRIVASTAVA N., HINTON G., KRIZHEVSKY A., SUTSKEVER I. & SALAKHUTDINOV R. (2014). Dropout : A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, **15**, 1929–1958.

- SUTSKEVER I., VINYALS O. & LE Q. V. (2014). Sequence to sequence learning with neural networks. In *Proceedings of NIPS*, Cambridge, MA, USA : MIT Press.
- VASWANI A., SHAZEER N., PARMAR N., USZKOREIT J., JONES L., GOMEZ A. N., KAISER L. & POLOSUKHIN I. (2017). Attention is All you Need. In I. GUYON, U. V. LUXBURG, S. BENGIO, H. WALLACH, R. FERGUS, S. VISHWANATHAN & R. GARNETT, Eds., *Advances in Neural Information Processing Systems 30*, p. 5998–6008. Curran Associates, Inc.
- VINYALS O., KAISER L., KOO T., PETROV S., SUTSKEVER I. & HINTON G. (2015). Grammar As a Foreign Language. In *Proceedings of the 28th International Conference on Neural Information Processing*, volume 2 of *NIPS'15*, p. 2773–2781, Cambridge, MA, USA : MIT Press.
- XIA Y., TIAN F., WU L., LIN J., QIN T., YU N. & LIU T.-Y. (2017). Deliberation networks : Sequence generation beyond one-pass decoding. In I. GUYON, U. V. LUXBURG, S. BENGIO, H. WALLACH, R. FERGUS, S. VISHWANATHAN & R. GARNETT, Eds., *Advances in Neural Information Processing Systems 30*, p. 1784–1794. Curran Associates, Inc.
- ZHANG X., SU J., QIN Y., LIU Y., JI R. & WANG H. (2018a). Asynchronous bidirectional decoding for neural machine translation. *CoRR*, **abs/1801.05122**.
- ZHANG Y., CHEN H., ZHAO Y., LIU Q. & YIN D. (2018b). Learning tag dependencies for sequence tagging. In *International Joint Conference on Artificial Intelligence (IJCAI)*.