



HAL
open science

Improving Recurrent Neural Networks For Sequence Labelling

Marco Dinarelli, Isabelle Tellier

► **To cite this version:**

Marco Dinarelli, Isabelle Tellier. Improving Recurrent Neural Networks For Sequence Labelling. 2016.
hal-01489976

HAL Id: hal-01489976

<https://hal.science/hal-01489976>

Submitted on 14 Mar 2017

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.

Improving Recurrent Neural Networks For Sequence Labelling

Marco Dinarelli and Isabelle Tellier

LaTTiCe (UMR 8094), CNRS, ENS Paris, Université Sorbonne Nouvelle - Paris 3
PSL Research University, USPC (Université Sorbonne Paris Cité)
1 rue Maurice Arnoux, 92120 Montrouge, France
`marco.dinarelli@ens.fr, isabelle.tellier@univ-paris3.fr`

June 8, 2016

Abstract

In this paper we study different types of Recurrent Neural Networks (RNN) for sequence labeling tasks. We propose two new variants of RNNs integrating improvements for sequence labeling, and we compare them to the more traditional Elman and Jordan RNNs. We compare all models, either traditional or new, on four distinct tasks of sequence labeling: two on Spoken Language Understanding (ATIS and MEDIA); and two of POS tagging for the French Treebank (FTB) and the Penn Treebank (PTB) corpora. The results show that our new variants of RNNs are always more effective than the others.

1 Introduction

Recurrent Neural Networks (RNN) [1–3] are neural models able to take some context into account in their decision function. For this reason, they are particularly suitable for several NLP tasks, in particular sequential information prediction [4–8]. In RNNs, the contextual information is provided to the model by a loop connection in the network architecture. This connection allows to use at the current time step one or more pieces of information predicted at previous time steps. This architecture seems particularly effective for neural networks since it allows to combine the power of distributional representations (or *embeddings*) with the effectiveness of contextual information.

In the literature about RNNs for NLP, two main variants have been proposed, also called “simple” RNNs: the Elman [2] and the Jordan [1] RNN models. The difference between these models lies in the position of the loop connection giving the *recurrent character* to the network: in the Elman RNN, it is put in the hidden layer whereas in

the Jordan RNN it connects the output layer to the hidden layer. In this last case, the recurrent connection allows to use, at the current time step, the information predicted at previous time steps. In the last few years, these two types of RNNs have been very successful for language modeling [9, 10], and for some sequence labeling tasks [6–8, 11, 12].

The intuition at the origin of this article is that embeddings allow a fine and effective modeling not only of words, but also of *labels* and their dependencies, which are very important for sequence labeling. In this paper, we define two new variants of RNN to achieve this more effective modeling.

In the first variant, the recurrent connection is between the output and the input layers. In other words, this variant gives labels predicted at previous positions in a sequence as input to the network. Such contextual information is added to the usual input context made of words, and both are used to predict the label at the current position in the sequence. Moreover we modified the hidden layer activity computation with respect to Elman and Jordan RNN, so that to take the different information provided by words and labels into account. From our intuition, thanks to label embeddings and to features-learning abilities of the hidden layer, this variant models in a more effective way label dependencies. The second variant we propose combines an Elman RNN and our first variant. This variant can thus exploit both contextual information provided by the previous states of the hidden layer, and the labels predicted at previous positions of a sequence.

A high-level schema of the Elman’s, Jordan’s and our first variant of RNN are shown in Figure 1. The schema of our second variant can be obtained by adding the recursion of the first one to the Elman architecture. In this figure, w is the input word, y is the predicted label, E , H , O and R are the parameter matrices between each pair of layers: they will be described in details in the next section. Before, it is worth discussing the advantages our variants can bring with respect to traditional RNN architectures, and explaining why they are expected to provide better modeling abilities.

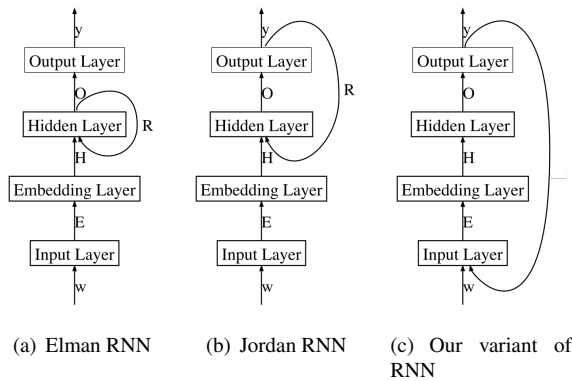


Figure 1: High Level Schema of the Main RNNs Studied in this Work.

First, since the output at previous positions is given as input to the network at the

current position, the contextual information flows across the whole network, affecting each layer’s input and output, at both forward and backward phases. In contrast, in Elman and Jordan RNNs, not all layers are affected at both forward and backward phases.

A second advantage of our variants is given by label embeddings. Indeed, the first layer of our RNNs is just a *look-up* table mapping sparse “one-hot” representations into distributional representations.¹ Since in our variants the output of the network at previous steps is given as input at the current step, the mapping from sparse representations to embeddings involves both words and labels. Label embeddings can be pre-trained from data as it is usually done for words. Pre-trained word embeddings, e.g. with *word2vec*, have already shown their ability to capture very attractive syntactic and semantic properties [13, 14]. Using label embeddings, the same properties can be learned also for labels. More importantly, using several predicted labels as embeddings provide a more effective modeling of label dependencies via the internal state of the network, which is the hidden layer’s output.

Another advantage coming from the use of label embeddings and different previous labels as context, is an increased robustness of the model to prediction mistakes. This effect comes from the syntactic and semantic properties that embeddings can encode [14].

All these advantages are also supported by our second variant, which uses both a label embedding context, like the first variant, and the loop connection at the hidden layer, like an Elman RNN.

All RNNs in this article are studied in their *forward*, *backward* and *bidirectional* versions [3]. In order to have a fair and straightforward comparison, we give the results of our new variants of RNNs together with those obtained with our implementation of Elman and Jordan RNNs. These implementations are very close to state-of-the-art, even if we did not implement every optimization feature.

All models are evaluated on four tasks. Two are Spoken Language Understanding (SLU) tasks [15]: ATIS [16] and MEDIA [17], which can be both modeled as sequence labeling problems. Two are POS-tagging tasks, one on the French Treebank (FTB) [18, 19] and one on the Penn Treebank (PTB) [20]. The results we obtain on these tasks with our implementations, despite they are not always better than the state-of-the-art, provide a stable ranking of different RNN architectures: at least one of our variants, most of the time and surprisingly the simpler one, is always better than Jordan and Elman RNNs.

In the remainder of the paper, we introduce RNNs and we describe in more details the variants proposed in this work (section 2). In section 3, we describe the corpora used for evaluation and all the results obtained, in comparison with state-of-the-art models. In section 4, we draw our conclusions.

¹The “one-hot” representation of an element at position i in a dictionary V is a vector of size $|V|$ where the i -th component has the value 1 whereas all the others are 0.

2 Improving Recurrent Neural Networks

The RNNs we consider in this work have the same architecture also used for Feedforward Neural Network Language Models (NNLM), described in [21]. In this architecture, we have four layers: input, embedding, hidden and output. Words are given as input to the network as indexes, corresponding to their position in a dictionary V .

The index of a word is used to select its embedding (or distributional representation) in a real-valued matrix $E \in R^{|V| \times N}$, $|V|$ being the size of the dictionary and N the dimensionality of the embeddings (which is a parameter to be chosen). We name $E(v(w_t))$ the embedding of the word w given as input at the position t of a sequence. $v(w_t) = i$ is the index of the word w_t in the dictionary, and it can be seen alternatively as a “one-hot” vector representation (the vector is zero everywhere except at position $v(w)$, where it is 1).

In contrast to NNLM, RNNs have one more connection, the recursive connection, between two layers, depending on the type of RNNs. As mentioned previously, Elman RNNs have a recursion loop in the hidden layer. Since this layer encodes the internal representation of the input to the network, the recurrent connection of an Elman network allows to keep “in memory” words used as input at previous positions in the sequence. Jordan RNNs have instead a recursion between the output and the hidden layer. This means that a Jordan RNN can take previous predicted labels into account to predict the label at the current position in a sequence. For every type of RNN, we call R the matrix of parameters of the recursion connection.

Our implementation of Jordan and Elman RNNs is like in the literature. [1–3].²

2.1 RNN Learning

Learning the described RNNs consists in learning the parameters $\Theta = (E, H, O, R)$ between each pair of layers (see Figure 1), and we omit biases to keep notations lighter. We use a cross-entropy cost function between the expected label c_t and the predicted label y_t at the position t in the sequence, plus a $L2$ regularization term [22]:

$$C = -c_t \cdot \log(y_t) + \frac{\lambda}{2} |\Theta|^2 \quad (1)$$

λ is an hyper-parameter of the model. Since y_t is a probability distribution over output labels, we can also view the output of a RNN as the probability of the predicted label y_t : $P(y_t|I, \Theta)$. I is the input given to the network plus the contextual information provided by the recurrent connection. For Elman RNN $I_{Elman} = w_{t-w} \dots w_t \dots w_{t+w}, h_{t-1}$, that is the word input context and the output of the hidden layer at the previous position in the sequence. For Jordan RNN $I_{Jordan} = w_{t-w} \dots w_t \dots w_{t+w}, y_{t-1}$, that is the same word input context as Elman RNN, and the label predicted at the previous position. We associate the following decision function to predict the label at position t in a sequence:

$$\tilde{t}_t = \underset{j \in \{1, \dots, |L|\}}{\operatorname{argmax}} P(y_t^j | I, \Theta) \quad (2)$$

²We also find [7] quite easy to understand, even for readers not familiar with RNNs.

\tilde{l}_t is a particular discrete label.

We use the back-propagation algorithm and the stochastic gradient descent with momentum [22] for learning the weights Θ .

2.2 Learning Variants

An important choice for learning RNN models concern the back-propagation algorithm. Indeed, because of the recurrent nature of their architecture, in order to properly learn RNNs the Back-Propagation Through Time algorithm (BPTT) should be used [23]. The BPTT algorithm consists basically in unfolding the recurrent architecture for a chosen number of steps, and then learning the network as a standard Feed-Forward network. This is supposed to allow RNNs to learn arbitrarily long past context. However, [10] has shown that RNNs for language modelling learn best with just 5 time steps in the past. This may be due to the fact that, at least in NLP tasks, the past information kept “in memory” by the network via the recurrent architecture, actually fades away after some time steps. Moreover, in many NLP tasks, using an arbitrarily-long context on either input or output side, doesn’t guarantee better performances, as increasing the context size also increases the noise. Since BPTT is quite more expensive than the traditional back-propagation algorithm, [7] has preferred to use explicit output context in Jordan RNNs and to learn the model with the traditional back-propagation algorithm, not surprisingly without losing performance.

In this work we use the same variant as [7]. When using an explicit context of output labels from previous time steps, the hidden layer activity of a Jordan RNN is computed as:

$$h_t = \Sigma(I_t \cdot H + [y_{t-c+1}y_{t-c+2}\dots y_{t-1}] \cdot R) \quad (3)$$

where c is the size of the history of previous labels that we want to explicitly use as context to predict next label. $[\cdot]$ indicates the concatenation of vectors.

All the modifications applied to the Jordan RNN so far can be applied in a similar way to the Elman RNN.

2.3 New Variants of RNN

As mentioned in the introduction, the new variants of RNN proposed in this work present two differences with respect to traditional RNNs: i) the recurrent connection is from the output to the input layer, meaning that predicted labels are converted into embeddings in the same way as words; ii) the hidden layer activities are computed in a slightly different way. Indeed, in Elman and Jordan RNNs the contextual information provided by the recurrent connection is summed to the input information (see equation 3 above). In our variants of RNNs instead, word and label embeddings are concatenated and provided to the hidden layer as different inputs.

The most interesting consequence of modifications in our RNN variants, is the fact that output labels are mapped into distributional representations, as it is usually done for input items. Indeed, the first layer of our network, is just a mapping from sparse “one-hot” representations to distributional representations. Such mapping results in fine

features and attractive syntactic and semantic properties, as shown by *word2vec* and similar works [13]. Such representations can be learnt from data the same way as for words. In the simplest case, this can be done by using sequences of output labels. When structured information is available, like syntactic parse trees or structured semantic labels such as named entities or entity relations, more sophisticated embeddings can be learnt. In this work, we learn label embeddings using sequences of output labels associated to word sequences in annotated data. It is worth noting that the idea of using label embeddings has been introduced by [24] in the context of dependency parsing. In this paper, we focus on the use of several label embeddings as context, thus encoding label dependencies, which are very important in sequence labeling tasks.

Using the same notation as above, we name E_w the embedding matrix for words, and E_l the embedding matrix for labels. We name

$$I_t = [E_w(v(w_{t-w})) \dots E_w(v(w_t)) \dots E_w(v(w_{t+w}))]$$

the concatenation of the vectors representing the input words when processing the position t in a sequence, while

$$L_t = [E_l(v(y_{t-c+1})) E_l(v(y_{t-c+2})) \dots E_l(v(y_{t-1}))]$$

is the concatenation of the vectors representing the output labels predicted at the previous c steps. The hidden layer activities are computed as:

$$h_t = \Sigma([I_t L_t] \cdot H)$$

Σ is the sigmoid activation function [22], $[\cdot]$ means the concatenation of the two matrices and we omit biases to keep notations lighter. The remainder of the layer activities, as well as the error computation and back-propagation are computed the same way as in traditional RNNs.

Note that in this variant of RNN there is no R matrix at the recurrent connection. The recurrent connection here means that the output is given back as input to the network and it is thus converted explicitly from probability distribution given by the softmax into a label index, which is used in turn to select a label embedding from the matrix E_l . Basically the role played from matrix R in Elman and Jordan RNN, is played from matrix E_l in our variant.

Another important interest of having the recursion between output and input layers is robustness. This is a direct consequence of using embeddings for output labels. Since we use several predicted labels as context at each position t (see L_t above), at least in the later stages of learning (when the model is close to the final optimum), it is unlikely to have several mistakes in the same context. Even then, thanks to the properties of distributed representations [14], wrong labels have very similar representations to the correct ones. Taking an example cited in [14]: if we use *Paris* instead of *Rome*, it has no effect for many NLP tasks, as they are both proper nouns for POS-tagging, locations for named entity recognition etc. Distributed representations for labels provide the same robustness on the output side.³

Jordan RNNs cannot provide in general the same robustness. We can interpret hidden activity computation in a Jordan RNN in two ways.

On the one hand, if we interpret a Jordan RNN as using sparse label representations as input to the hidden layer, such representations are either “one-hot” representations of

³Sometimes in POS-tagging, models mistake verbs and nouns. They make such errors because some particular verbs occur in the same context of nouns (e.g. “the sleep is important”), and so have similar representations.

labels or probability distributions given as output by the softmax at the output layer. In the first case it is clear that a mistake may have more effect than in the second one, as the only value that is not zero is in a wrong position. But when probability distributions are used, we have found that most of the probability mass is “picked” on one or few labels, which thus does not provide much more softness than a “one-hot” representation. In this interpretation, sparse labels are an additional input for the hidden layer, the matrix R on the recurrent connection plays the same role as H does for input words. The two matrices are then summed to compute the total input of the hidden layer.

On the other hand, we can see the multiplication of sparse representation of labels in equation 3 as the selection of an embedding for labels from matrix R .⁴ Even in this interpretation there is a substantial difference between the Jordan RNN and our variant of RNN, I-RNN henceforth (I stands for *Improved*). In order to understand in detail this difference, we focus on the equations for computing the hidden activities. For Jordan RNN we have:

$$h_t = \Sigma(I_t \cdot H + [y_{t-c+1}y_{t-c+2}\dots y_{t-1}] \cdot R)$$

For I-RNN we have:

$$h_t = \Sigma([I_t L_t] \cdot H)$$

where L_t is $[E_i(v(y_{t-c+1}))E_i(v(y_{t-c+2}))\dots E_i(v(y_{t-1}))]$, that is the concatenation of c previous label embeddings.

In this second interpretation, in Jordan RNN labels are not directly concerned in the computation of the total input for the hidden layer, since they are not directly multiplied by matrix H . Only input context I_t is multiplied by H . The result of this multiplication is summed to the result of label embedding selection, performed as $y[t-i] \cdot R, i = 1 \dots c-1$. Finally the hidden non-linear function Σ is applied. Mixing input processing $I_t \cdot H$ and label processing $y[t-i] \cdot R$ with a sum, can make sense for the tasks for which the Jordan network was designed [1], as output units were of the same nature as input units (speech signal). However we believe that it doesn’t express sufficiently well words and labels interactions in NLP tasks. Also, since $y[t-i] \cdot R$ is an embedding selection for labels, labels and words are not processed in the same way: I_t is already made of word embeddings, which are further transformed by $I_t \cdot H$. This further transformation is not applied to label embeddings.

In I-RNN in contrast, sparse labels are first converted into embeddings with $E_i(v(y_{t-i}))$, $i = 1 \dots c$, and their concatenation results in L_t . This matrix is further concatenated to the input context I_t . The result of the concatenation is multiplied by matrix H to compute the total input to the hidden layer.⁵ Finally the hidden non-linear function Σ is applied. This means that information provided by the input context I_t is not mixed with L_t with a sum like in Jordan RNN. These two data are given to the hidden layer as separated inputs. More in particular, the concatenation of I_t and L_t is performed neuron-wise, that is each hidden neuron receives as input all context words and all context labels, encoding them as a network internal feature. Thus we let the hidden layer itself to learn labels interactions, and words-labels interactions. This is indeed in agreement with the “philosophy” of neural networks, where features designing is turned into features learning. Since words and labels have different nature in sequence labeling

⁴Multiplying a “one-hot” representation by a matrix is equivalent to selecting one row of the matrix.

⁵We remind the reader that we are omitting biases to keep notation lighter

tasks, we believe that modeling interactions in this way is more effective. Also, with the \mathbb{I} -RNN architecture, words and labels are processed in the same way, as both are first converted into embeddings, and then “transformed” again via multiplication by H matrix. In order to make results obtained with different RNNs comparable, we used the same number of hidden neurons for all RNNs. In \mathbb{I} -RNN thus, each hidden neuron receives as input much more information than Jordan and Elman hidden neurons.

In order to make the explanation more clear, \mathbb{I} -RNN architecture is detailed in figure 2. Symbols have the same meaning as equations in the paper, the only exception is that labels are indicated in the figure with upper-case L . Matrix H is replicated at the hidden layer computation meaning that all neurons receive the whole $[I_t L_t]$ input, which is made of $2 \cdot w + 1 + c$ D -dimensional embeddings: $2w + 1$ word embeddings and c label embeddings. *CAT* is the concatenation operator. Please note that the concatenations of embeddings are performed at two different steps just for the sake of clearness and to be coherent with equations in the paper, all concatenations can be performed in one step.

Our second variant of RNN combines the characteristics of an Elman RNN and of our first variant. In this variant, the only difference with the first one is the computation of the hidden layer activities, where we use the concatenation of the c previous hidden layer states in addition to the information already used in the first variant:

$$h_t = \Sigma([I_t L_t] \cdot H + [h_{t-c+1} h_{t-c+2} \dots h_{t-1}] \cdot R)$$

2.4 Forward, Backward and Bidirectional RNNs

All RNNs described in this work are studied in their *forward*, *backward* and *bidirectional* versions [3]. Forward RNNs work as described so far. Backward RNNs have exactly the same architecture, the only difference being that they process sequences in the reverse order, from the end to the begin. Backward RNNs can be used to predict future labels in a sequence.

Bidirectional RNNs use both past and future information to predict the next label, which is both words and labels in our variants and in Jordan RNNs, or both words and hidden layers in Elman RNNs. When labeling sequences with bidirectional RNNs, a backward network is first used to predict labels backward. The bidirectional RNN then processes sequences in forward direction using past contextual information as usual, and future contextual information provided by the states and labels predicted by the backward network. The hidden layer of the bidirectional version of our first variant of RNNs is thus computed as:

$$h_t = \Sigma([I_t L_t^p L_t^f] \cdot H)$$

where $L_t^p = L_t$ introduced above, while

$$L_t^f = [E_l(v(y_{t+1})) \dots E_l(v(y_{t+c-1})) E_l(v(y_{t+c}))]$$

is the concatenation of the vectors representing the c future labels predicted by the backward model. It is very similar for our second variant, refer to [3] for details.

2.5 Recurrent Neural Network Complexity

We provide here an analysis of the complexity in terms of the number of parameters involved in each model. In Jordan RNN we count:

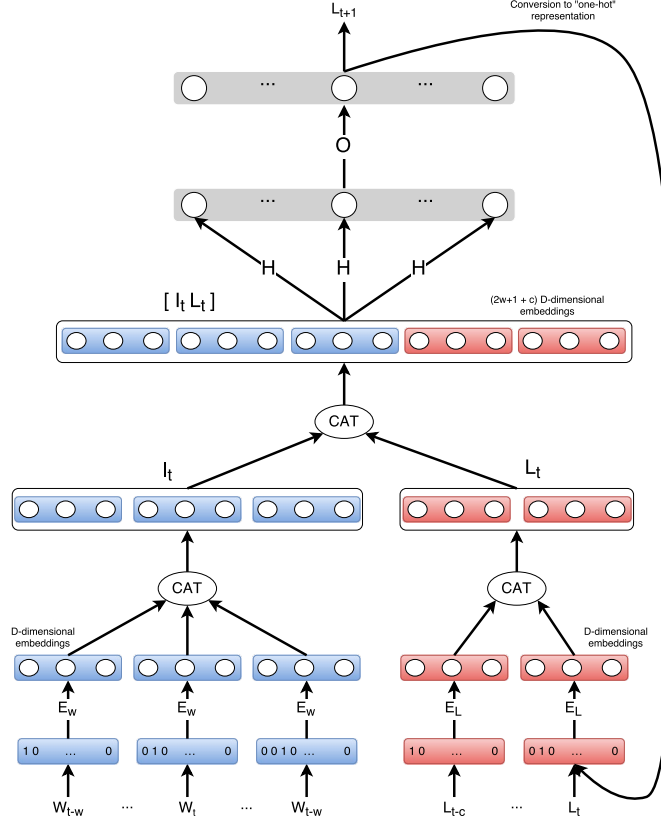


Figure 2: Detailed architecture of I-RNN. Symbols have the same meaning as equations in the paper, the only exception is that labels are indicated in the figure with upper-case L . Please note that matrix H is replicated at the hidden layer computation meaning that all neurons receive the whole $[l_t L_t]$ input, which is made of $2 \cdot w + 1 + c$ D -dimensional embeddings: $2w + 1$ word embeddings and c label embeddings. CAT is the concatenation operator.

$$|V| \times D + ((2w + 1)D + c|O|) \times |H| + |H| \times |O|$$

where $|V|$ is the size of the input dictionary, D the dimensionality of the embeddings, $|H|$ and $|O|$ are the size of the hidden and output layers, respectively, w the size of the window of words used as context on the input side⁶ and c is the size of the context of labels, which is multiplied by the dimensionality of the output label dictionary $|O|$.

With the same symbols, in an Elman RNN and in our first variant we have, respectively:

$$|V| \times D + ((2w + 1)D + c|H|) \times |H| + |H| \times |O|$$

and

$$|V| \times D + |O| \times D + ((2w + 1)D + cD) \times |H| + |H| \times |O|$$

⁶The word input context is thus made of w words on the left and w on the right of the word at a given position t , plus the word at t itself, which gives a total of $2w + 1$ input words

The only difference between the Jordan and Elman RNNs lies in the factors $c|O|$ and $c|H|$. Their difference in complexity depends thus on the size of the output layer (Jordan) with respect to the size of the hidden layer (Elman). Since in sequence labeling tasks the hidden layer is often bigger, Elman RNN is more complex than Jordan RNN. The difference between the Jordan RNN and our first variant is in the factors $|O| \times D$ and cD . The first is due to the label embeddings⁷, the second is due to the use of such embeddings as input to the hidden layer. Since often D and O have sizes in the same order of magnitude, and thanks to the use of vectorized operations on matrices, we didn't find a noticeable difference in terms of training and testing time between the Jordan RNN and our first variant. This simple analysis also shows that our first variant roughly needs the same number of connections in the hidden layer as a Jordan RNN. Our first variant is thus architecturally equivalent to a Jordan RNN.

In contrast, for the second variant we have:

$$|V| \times D + |O| \times D + ((2w + 1)D + cD + c|H|) \times |H| + |H| \times |O|$$

The additional term $c|H|$, is due to the same recurrent connection as in an Elman RNN. Using vectorized operations for matrix calculations, we found the second variant slower for both training and testing time by a factor 1.15 with respect to the other RNNs. The same complexity analysis holds for backward RNNs. But bidirectional RNNs are even more complex. Without deriving any new formula, we note that they are slower with respect to their corresponding forward/backward models by a factor of roughly 1.5.

3 Evaluation

3.1 Corpora

We used four distinct corpora:

The Air Travel Information System (ATIS) task [16] has been designed to automatically provide flight information in SLU systems. The semantic representation is frame-based and the goal is to find the correct frame and the corresponding semantic slots. For example, for the sentence “*I want the flights from Boston to Philadelphia today*”, the correct frame is FLIGHT and the words *Boston*, *Philadelphia* and *today* must be annotated with the concepts DEPARTURE.CITY, ARRIVAL.CITY and DEPARTURE.DATE, respectively.

ATIS is a relatively simple task dating from 1993. The training set is made of 4978 sentences taken from the “context independent” data in the ATIS-2 and ATIS-3 corpora. The test set is made of 893 sentences, taken from the ATIS-3 NOV93 and DEC94 datasets. There are no official development data provided with this corpus, we have thus taken a part of the training data at random to play this role.⁸

The French corpus MEDIA [17] has been created to evaluate SLU systems providing tourist information, in particular hotel information in France. It is composed of 1250 dialogues acquired with a Wizard-of-OZ protocol where 250 speakers have applied 5 hotel reservation scenarios. The dialogues have been annotated following a rich

⁷We use embeddings of the same size D for words and labels.

⁸Please see [16] for more details on the ATIS corpus.

	training		development		test	
# sentences	12,908		1,259		3,005	
	words	concepts	words	concepts	words	concepts
# words	94,466	43,078	10,849	4,705	25,606	11,383
# vocab.	2,210	99	838	66	1,276	78
# OOV%	–	–	1.33	0.02	1.39	0.04

Table 1: Statistics on the corpus MEDIA

Section	# sentences	# tokens	# unk. tokens
FTB-train	9881	278083	
FTB-dev	1235	36508	1790 (4.9%)
FTB-test	1235	36340	1701 (4.7%)

Table 2: Statistics on the FTB corpus used for POS-tagging

semantic ontology. Semantic components can be combined to create complex semantic labels.⁹ In addition to this rich annotation, another difficulty lies in the coreferences. Some words cannot be correctly annotated without information about previous dialog turns. For example in the sentence “*Yes, the one at less than fifty euros per night*”, *the one* refers to an hotel previously introduced in the dialog. Statistics on training, development and test data from this corpus are shown in table 1.

Both ATIS and MEDIA can be modelled as sequence labeling tasks using the *BIO* chunking notation [25]. Several different works compared on ATIS [6–8, 26]. [8] is the only work providing results on MEDIA with RNNs, it also provides results obtained with CRFs [27], allowing an interesting comparison.

The French Treebank (FTB) corpus is presented in [18]. The version we use for POS-tagging is exactly the same as in [19]. But, in contrast to them, who reach the best result on this task with an external lexicon, we don’t use any external resource here¹⁰. Statistics on the FTB corpus are shown in table 2.

The Penn Treebank (FTB) corpus is presented in [20]. In order to have a direct comparison with previous works [28] [29] [5], we split the data as they do: sections 0 – 18 are used for training, 19 – 21 for validation and 22 – 24 for testing. Statistics on the PTB corpus are shown in table 3.

3.2 RNNs Parameters Settings

In order to compare with some published works on the ATIS and MEDIA tasks, we use the same dimensionality settings used by [6], [7] and [8], that is embeddings have 200 dimensions, hidden layer has 100 dimensions. We also use the same context size for words, that is $w = 3$, and we use $c = 6$ as labels context size in our variants and in Jordan RNN. We use the same tokenization, basically consisting of words lower-casing.

In contrast, in our models the sigmoid activation function at the hidden layer and the $L2$ regularization. While [6], [7], [8] and [26] use the rectified linear activation function and the dropout regularization [22] [30].

⁹For example the label *localisation* can be combined with *city*, *relative-distance*, *general-relative-place*, *street* etc.

¹⁰[19] also provides results without using the external lexicon

Data set	Sections	Sentences	Tokens	Unknown
Training	0-18	38,219	912,344	0
Development	19-21	5,527	131,768	4,467
Test	22-24	5,462	129,654	3,649

Table 3: Statistics of the training, development and test data texts of the Penn Treebank corpus

For the FTB POS-tagging task we have used 200-dimensional embeddings, 300-dimensional hidden layer, again $w = 3$ for the context on the input side, and 6 context labels on the output side. The bigger hidden layer gave better results during validation experiments, due to the larger word dictionary in this task with respect to the others, roughly 25000 for the FTB against 2000 for MEDIA and 1300 for ATIS. In contrast to [19], which has used several features of words (prefixes, suffixes, capitalisation information etc.), we only performed a simple tokenization for reducing the size of the input dictionary: all numbers have been mapped to a conventional symbol (NUM), and nouns not corresponding to proper names and starting with a capital letter have been converted to lowercase. We preferred this simple tokenisation without using rich features, because our goal in this work is not obtaining the best results ever, it is to compare Jordan and Elman RNNs with our variants of RNN and show that our variants works better for sequence labeling. Adding many features and/or building sophisticated models would make the message less clear, as results would be probably better but the improvements could be attributed to rich and sophisticated models, instead of to the model itself.

For the PTB POS-tagging task we use exactly the same settings and pre-processing as for the FTB task, except that we used 400 hidden neurons. During validation we found that this works better, again due to the size of the dictionary which is 45000 for this task (after pre-processing).

We trained all RNNs with exactly the same protocol: i) we first train neural language models to obtain word and label embeddings. This language model is like the one in [21], except it uses both words/labels in the past and in the future to predict next word/label. ii) we train all RNNs using the same embeddings trained at previous step. We train the RNN for word embeddings for 20 epochs, the RNN for label embeddings for 10 epochs, and we train the RNNs for sequence labeling for 20 epochs. The number of epochs has been roughly optimized on development data. At the end of training we keep the model which gave the best tagging accuracy on development data. Also we roughly optimized on development data the learning rate and the parameter λ for regularization, the best values found are 0.5 and 0.003, respectively.

3.3 Training and Tagging Time

Since implementations of RNNs use in this work are prototypes¹¹, it does not make sense to compare them to state-of-the-art in terms of training and tagging time. However it is worth providing training times at least to have an idea and to have a comparison among different RNNs.

¹¹Our implementations are basically written in Octave <https://www.gnu.org/software/octave/>

As explained in section 2.5, our first variant I -RNN and the Jordan RNN have the same complexity. Also, since the size of the hidden layer is in the same order of magnitude as the size of the output layer (i.e. the number of labels), also Elman RNN has roughly the same complexity as the Jordan RNN. This is reflected in the training time.

The training time for label embeddings is always relatively short, as the size of the output layer, that is the number of labels, is always relatively small. This training time thus can vary from few minutes for ATIS and MEDIA, to less than 1 hour for the FTB corpus.

Training word embeddings is also very fast on ATIS and MEDIA, taking less than 30 minutes for ATIS and roughly 40 minutes for MEDIA. In contrast training word embeddings on the FTB corpus takes roughly 5 days, training time goes to roughly 6 days for the PTB word embeddings.

Training the RNN taggers takes roughly the same time as training the word embeddings, as the size of the word dictionary is the dimension that most affects the computational complexity at softmax used to predict next label.

Concerning the second variant of RNN proposed in this work, $I+E$ -RNN, this is slower as it is more complex in terms of number of parameters. Training $I+E$ -RNN on ATIS and MEDIA takes roughly 45 minutes and 1 hour, respectively. In contrast training $I+E$ -RNN on the FTB and PTB corpora takes roughly 6 days and 7 days, respectively.

We didn't keep track of tagging time, however it is always negligible with respect to training time, and it is always measured in few minutes. All times provided here are with a 1.7 GHz CPU, single process.

3.4 Sequence Labeling Results

The evaluation of all models described here is shown in table 4 5 6 and 7, in terms of F1 measure for ATIS and MEDIA, *Accuracy* on the FTB and PTB. Our implementations of Elman and Jordan RNN are indicated in tables as E -RNN and J -RNN. Our new variants are indicated as I -RNN and $I+E$ -RNN.

Table 4 shows results on the ATIS corpus. In the higher part of the table we show results obtained using only words as input. In the lower part, results are obtained using both words and word-classes available for this task. Such classes concern city names, airports and time expressions. They allow the models to generalize from specific words triggering concepts.¹²

Note that our results on the ATIS corpus are not always comparable with those published in the literature because: *i*) Models published in the literature use a *rectified linear* activation function at the hidden layer¹³ and the *dropout* regularization. Our models use the sigmoid activation function and the $L2$ regularization. *ii*) For experiments on ATIS we have used roughly 18% of the training data for development, we thus

¹²For example the cities of *Boston* and *Philadelphia* in the example above are mapped to the class CITY-NAME. If a model has never seen *Boston* during the training phase, but has seen at least one city name, it will possibly annotate *Boston* as a departure city thanks to some discriminative context, such as the preposition *from*.

¹³ $f(x) = \max(0, x)$.

used a smaller training set. *iii*) The works we compare with do not always give details on how classes available for the task have been integrated into their models. *iv*) Layer dimensionality and hyper-parameter settings do not always match those of published works. In fact, to avoid running too much experiments, we have based our settings on known works, but this doesn't allow a straightforward comparison with other published works.

Despite this, the message we want to claim in this work is still true for two reasons: *i*) Some of the results obtained with Elman and Jordan RNNs are close, or even better, than state-of-the-art. Thus they are not weak baselines. *ii*) We provide a fair comparison of our variants with traditional Elman and Jordan RNNs.

The results in the higher part of table 4 show that the best model on the ATIS task, with these settings, is the Elman RNN of [26]. Note that it is not clear how the improvements of [26] with respect to [7] (in part due to same authors) have been obtained. Indeed, in [7] authors obtain the best result with a Jordan RNN, while in [26] an Elman RNN gets the best performance. During our experiments, using the same experimentation protocol as [26], we could not reach the same performances. We conclude that the difference between our results and those in [26] are due to reasons mentioned above. Beyond this, we note that our Elman and Jordan RNN implementations are equivalent to those of [7]. Also, our first variant of RNNs, \mathbb{I} -RNN, obtains the second best result (93.84 in bold), which is the best result we could reach. Our second variant is roughly equivalent to a Jordan RNN on this task.

The results in the lower part of the table 4 (*Classes*), obtained with both words and classes as input, are quite better than those obtained with words only. They roughly follow the same behavior, except that in this case our Jordan RNN is slightly better than our second variant. The first variant \mathbb{I} -RNN obtains again the best result among our implementations (95.21 in bold). In this case also, we attribute the differences with respect to published results to the different settings mentioned above. For comparison, in this part of the table we show also results obtained using CRF.

On the ATIS task, using either words or both words and classes as inputs, we can see that results are always quite high. This task is relatively simple. Label dependencies can be easily modeled, as there is basically no segmentation of concepts over different consecutive words (one concept corresponds to one word). In this settings, the potential of our new variants of RNNs cannot be fully exploited. This limitation is confirmed by the results obtained by [8] and [26] using CRFs. Indeed, RNNs don't take the whole sequence of labels into account in their decision function. In contrast, CRFs use a global decision function taking all the possible labelings of a given input sequence into account to predict the best sequence of labels. The fact that CRFs are less effective than RNN on ATIS is a clear sign that label dependency modeling is relatively simple in this task.

Table 5 shows the results on the corpus MEDIA. As already mentioned, this task is more difficult because of its richer semantic annotation, and also because of the coreferences and of the segmentation of concepts over several words. This creates relatively long label dependencies, which cannot be taken into account by simple models. The difficulty of this task is confirmed by the magnitude of results (10 F1 points lower than on the ATIS task).

As can be seen in Table 5, CRFs are in general much more effective than Jordan

and Elman RNNs on MEDIA. This outcome could be expected as RNNs use a local decision function not able to take long label dependencies into account. We can also see that our implementations of Elman and Jordan RNNs are comparable, even better in the case of Elman RNN, with state-of-the-art RNNs of [8].

More importantly, results on the MEDIA task shows that in this particular experimental settings where taking label dependencies into account is crucial, our new variants are remarkably more effective than both our implementations and state-of-the-art implementations [8] of Elman and Jordan RNNs. This holds for forward, backward and bidirectional RNNs. Moreover, the bidirectional version of our variants of RNNs outperforms even CRF. We attribute this effectiveness to a better modeling of label dependencies, due to label embeddings.

Table 6 shows the results obtained on the POS-tagging task of the FTB corpus. On this task, we compare our RNN implementations to the state-of-the-art results obtained in [19] with the model $MElt_{fr}^0$. We would like to underline that $MElt_{fr}^0$, when it does not use external resources like in the model obtaining the best absolute result in [19], nevertheless uses several features associated with words that provide an advantage over features used in our RNNs. As can be expected thus, $MElt_{fr}^0$ outperforms the RNNs. Results in table 6 shows that forward and backward RNNs are quite close to each other on this task, I-RNN, providing a little improvement. In contrast, the bidirectional version of I-RNN and I+E-RNN provide a significant improvement over Jordan and Elman RNNs.

Table 7 shows the results obtained on the WSJ POS tagging task. We also provide the results of [28] and [29] for comparison with the state-of-the-art. This is just for a matter of comparison, as the results shown in those works were achieved with rich features and more sophisticated models. Instead, we can roughly compare our results with those of [5], which were also obtained with neural networks. In particular, we can compare with the model called NN+SLL. Note that this is a rough comparison as the model of [5], though not a RNN, integrates capitalisation features and uses a convolution and a *max-over-time* layer to encode large context information.

Results for POS tagging on the PTB corpus roughly confirm the conclusions reached for FTB results, with the only difference that in this case I+E-RNN is slightly better than I-RNN. Our RNNs don't improve the state-of-the-art, but are all more effective than the model of [5]. This result is particularly important, as it shows that RNNs, even without using a sophisticated encoding of the context like the model NN+SLL in [5], are intrinsically a better model for sequence labeling. This claim is enforced also by the fact that NN+SLL of [5] implements a global probability computation strategy similar to CRF (SLL stands for *Sentence Level Likelihood*), while all RNNs presented here use a local decision function (see equation 2). Again, the ranking of RNN models on the PTB POS-tagging task is stable, the variants of RNNs proposed in this work being more effective than traditional Elman and Jordan RNNs.

We would like to notify that we have also performed some experiments on ATIS and MEDIA without pre-training label embeddings. The results reached are not substantially different from those obtained with pre-training label embeddings. Indeed, on relatively small data, it is not rare to have similar or even better results without pre-training. This is due to the fact that learning effective embeddings requires a relatively large amount of data. [8] also shows results obtained with embeddings pre-trained with

word2vec and without embedding pre-training. His conclusions are indeed very similar. More generally, reaching roughly the same results on a difficult task like MEDIA without label embedding pre-training is a clear sign that our variants of RNNs are superior to traditional RNNs because they use a context made of label embeddings. As a matter of fact, the gain with respect to Elman and Jordan RNNs cannot be attributed in this case to the use of pre-trained embeddings. On relatively larger corpora like FTB and PTB, label embedding pre-training seems to provide a slight improvement.

Finally, we have run experiments also modifying Jordan and Elman RNNs so that to model words and labels interactions more like \mathbb{I} -RNN does, that is word and label embeddings (or hidden states in Elman RNN) are not summed together, instead they are concatenated. Results obtained were not substantially different from those obtained with “traditional” Jordan and Elman RNN and in any case \mathbb{I} -RNN was always performing best, still keeping a large gain over Elman and Jordan RNN on the MEDIA task. The explanation of this outcome is that keeping word and label embeddings separated, and then multiplying both by matrix H to compute hidden activities, as we do in \mathbb{I} -RNN, is more effective than concatenating $I_t \cdot H$ and $y[t - 1] \cdot R$, as we did for Jordan RNN, and analogously with previous hidden layer for Elman RNN. This is not surprising, as \mathbb{I} -RNN also in this case is applying one transformation more when multiplying label embeddings by H to compute the total input to the hidden layer.

It is somewhat surprising that \mathbb{I} -RNN systematically outperforms $\mathbb{I}+\mathbb{E}$ -RNN, the latter model integrates more information at the hidden layer and thus should be able to take advantage of both Elman and \mathbb{I} -RNN characteristics. While an analysis to explain this outcome is not trivial, our interpretation is that using two recursions in a RNN gives actually redundant information to the model. Indeed, the output of the hidden layer keeps the internal state of the network, which is the internal (distributed) representation of the input n-gram of words around position t and the previous c labels. The recursion at the hidden layer allows to keep this information “in memory” and to use it at the next step $t + 1$. However using the recursion of \mathbb{I} -RNN the previous c labels are also given explicitly as input to the hidden layer. This may be redundant, and constrain the model to learn an increased amount of noise. A similar idea of hybrid RNN model has been tested in [26] without showing a clear advantage on Elman and Jordan RNNs.

What can be said in general from the results obtained on all the presented tasks, is that RNN architectures using label embedding context can model label dependencies in a more effective way, even when these dependencies are relatively simple (like in ATIS and POS tagging tasks). The two variant of RNNs proposed in this work, in particular the \mathbb{I} -RNN variant, are for this reason more effective than Elman and Jordan RNNs on sequence labeling tasks.

3.5 Comparison of Jordan-RNN and \mathbb{I} -RNN label representations

We compare Jordan RNN and \mathbb{I} -RNN label representations under the interpretation where Jordan RNN hidden activity computation uses sparse labels as an additional input to the hidden layer, as explained in section 2.3. As explained also in the same section, under the other interpretation \mathbb{I} -RNN provides the advantage of performing an

Model	F1 measure		
	<i>forward</i>	<i>backward</i>	bidirectional
Words			
[7] E-RNN	93.65%	92.12%	–
[7] J-RNN	93.77%	93.31%	93.98%
[26] E-RNN	94.98%	–	94.73%
E-RNN	93.41%	93.18%	93.41%
J-RNN	93.61%	93.55%	93.74%
I-RNN	93.84%	93.56%	93.73%
I+E-RNN	93.74%	93.44%	93.60%
Classes			
[8] E-RNN	96.16%	–	–
[8] CRF	–	–	95.23%
[6] E-RNN	96.04%	–	–
[26] E-RNN	96.24%	–	96.29%
[26] CRF	–	–	95.16%
E-RNN	94.73%	93.61%	94.71%
J-RNN	94.94%	94.80%	94.89%
I-RNN	95.21%	94.64%	94.75%
I+E-RNN	94.84%	94.61%	94.79%

Table 4: Results of SLU on the ATIS corpus

Model	F1 measure		
	<i>forward</i>	<i>backward</i>	bidirectional
[8] E-RNN	81.94%	–	–
[8] J-RNN	83.25%	–	–
[8] CRF	–	–	86.00%
E-RNN	82.64%	82.61%	83.13%
J-RNN	83.06%	83.74%	84.29%
I-RNN	84.91%	86.28%	86.71%
I+E-RNN	84.58%	85.84%	86.21%

Table 5: Results of SLU on the MEDIA corpus

additional transformation on labels, and gives words and labels embeddings as separated inputs to the hidden layer.

Under the first interpretation, the advantage of using label embeddings in I-RNN instead of “one-hot” or probability distribution representations like in Jordan RNNs, is an increased amount of *signal* flowing across the network. A semantic interpretation of the interaction of these two representations with the network is not trivial. Indeed, in the probability representation output by the softmax in a Jordan RNN, the different dimensions are just probabilities associated to different labels. In contrast, in label embeddings used in I-RNN, the dimensions are different distributional features, related to how a particular label is used in particular label contexts. A comparison between these two representation is thus not really meaningful.

Instead, we performed a simple analysis of the magnitude of values found in the probability distribution used as representations in Jordan RNNs, when using development data of the MEDIA corpus. We summarize this analysis as follows:

1. 7843 out of 11051 (71%) of the time, the maximum value is greater than 0.9
2. 9928 out of 11051 (90%) of the time, the sum of the 3 highest probabilities is greater than 0.9
3. Excluding the 3 highest probabilities, the remaining values in the distribution

Model	Accuracy		
	<i>forward</i>	<i>backward</i>	bidirectional
[19] $MElt_{fr}^0$	–	–	97.00%
E-RNN	96.31%	96.30%	96.32%
J-RNN	96.31%	96.27%	96.32%
I-RNN	96.36%	96.37%	96.51%
I+E-RNN	96.28%	96.33%	96.48%

Table 6: Results of POS-tagging on the FTB

Model	Accuracy		
	<i>forward</i>	<i>backward</i>	bidirectional
[28]	–	–	97.24%
[29]	–	–	97.33%
[5] NN+SLL	–	–	96.37%
E-RNN	96.75%	96.76%	96.75%
J-RNN	96.71%	96.69%	96.77%
I-RNN	96.75%	96.72%	96.90%
I+E-RNN	96.73%	96.71%	96.93%

Table 7: Results of POS-tagging on the PTB

have very small values (less than 0.001)

This simple analysis shows that the probability distributions used for label representations in Jordan RNNs do not provide much more information to the network than a “one-hot” representation, and not much signal into the network. This problem is somewhat similar to the “vanishing gradient” problem [31]: as the network learns, the probability gets concentrated on few dimensions and all the other values get very small, limiting network learning. This problem is all the more obvious that label dependency modeling is more important for the task. On an absolute scale however, this is less serious than the vanishing gradient problem, as Jordan RNNs still reach competitive performances.

4 Conclusions

In this paper we have studied different architectures of Recurrent Neural Networks for sequence labeling tasks. We have proposed two new variants of RNNs to better model label dependencies, and we have compared these variants to the traditional architectures of Elman and Jordan RNNs. We have explained the advantages provided by the proposed variants with respect to previous RNNs. We have evaluated all RNNs, either new or traditional, on 4 different tasks: two of Spoken Language Understanding and two of POS-tagging. The results show that, even though RNNs don’t always improve the state-of-the-art, our new variants of RNNs always outperform the traditional Elman and Jordan RNNs.

References

- [1] Jordan, M.I.: Serial order: A parallel, distributed processing approach. In Elman, J.L., Rumelhart, D.E., eds.: Advances in Connectionist Theory: Speech.

Erlbaum, Hillsdale, NJ (1989)

- [2] Elman, J.L.: Finding structure in time. *COGNITIVE SCIENCE* **14** (1990) 179–211
- [3] Schuster, M., Paliwal, K.: Bidirectional recurrent neural networks. *Trans. Sig. Proc.* **45** (1997) 2673–2681
- [4] Collobert, R., Weston, J.: A unified architecture for natural language processing: Deep neural networks with multitask learning. In: *Proceedings of the 25th International Conference on Machine Learning. ICML '08, New York, NY, USA, ACM (2008)* 160–167
- [5] Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., Kuksa, P.: Natural language processing (almost) from scratch. *J. Mach. Learn. Res.* **12** (2011) 2493–2537
- [6] Yao, K., Zweig, G., Hwang, M.Y., Shi, Y., Yu, D.: Recurrent neural networks for language understanding, *Interspeech* (2013)
- [7] Mesnil, G., He, X., Deng, L., Bengio, Y.: Investigation of recurrent-neural-network architectures and learning methods for spoken language understanding. In: *Interspeech 2013*. (2013)
- [8] Vukotic, V., Raymond, C., Gravier, G.: Is it time to switch to word embedding and recurrent neural networks for spoken language understanding? In: *InterSpeech, Dresde, Germany* (2015)
- [9] Mikolov, T., Karafiát, M., Burget, L., Cernocký, J., Khudanpur, S.: Recurrent neural network based language model. In: *INTERSPEECH 2010, 11th Annual Conference of the International Speech Communication Association, Makuhari, Chiba, Japan, September 26-30, 2010*. (2010) 1045–1048
- [10] Mikolov, T., Kombrink, S., Burget, L., Cernock, J., Khudanpur, S.: Extensions of recurrent neural network language model. In: *ICASSP, IEEE (2011)* 5528–5531
- [11] Xu, W., Auli, M., Clark, S.: CCG supertagging with a recurrent neural network. In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 2: Short Papers*. (2015) 250–255
- [12] Zennaki, O., Semmar, N., Besacier, L.: Unsupervised and lightly supervised part-of-speech tagging using recurrent neural networks. In: *Proceedings of the 29th Pacific Asia Conference on Language, Information and Computation, PACLIC 29, Shanghai, China, October 30 - November 1, 2015*. (2015)
- [13] Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. *CoRR* **abs/1301.3781** (2013)

- [14] Mikolov, T., Yih, W., Zweig, G.: Linguistic regularities in continuous space word representations. In: *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics*. (2013) 746–751
- [15] De Mori, R., Bechet, F., Hakkani-Tur, D., McTear, M., Riccardi, G., Tur, G.: Spoken language understanding: A survey. *IEEE Signal Processing Magazine* **25** (2008) 50–58
- [16] Dahl, D.A., Bates, M., Brown, M., Fisher, W., Hunicke-Smith, K., Pallett, D., Pao, C., Rudnicky, A., Shriberg, E.: Expanding the scope of the atis task: The atis-3 corpus. In: *Proceedings of the Workshop on Human Language Technology. HLT '94*, Stroudsburg, PA, USA, Association for Computational Linguistics (1994) 43–48
- [17] Bonneau-Maynard, H., Ayache, C., Bechet, F., Denis, A., Kuhn, A., Lefèvre, F., Mostefa, D., Qugnard, M., Rosset, S., Servan, S., Vilaneau, J.: Results of the french evalda-media evaluation campaign for literal understanding. In: *LREC*, Genoa, Italy (2006) 2054–2059
- [18] Abeillé, A., Clément, L., Toussanel, F.: Building a Treebank for French. In: *Treebanks : Building and Using Parsed Corpora*. Springer (2003) 165–188
- [19] Denis, P., Sagot, B.: Coupling an annotated corpus and a lexicon for state-of-the-art pos tagging. *Lang. Resour. Eval.* **46** (2012) 721–736
- [20] Marcus, M.P., Santorini, B., Marcinkiewicz, M.A.: Building a large annotated corpus of english: The penn treebank. *COMPUTATIONAL LINGUISTICS* **19** (1993) 313–330
- [21] Bengio, Y., Ducharme, R., Vincent, P., Jauvin, C.: A neural probabilistic language model. *JOURNAL OF MACHINE LEARNING RESEARCH* **3** (2003) 1137–1155
- [22] Bengio, Y.: Practical recommendations for gradient-based training of deep architectures. *CoRR* **abs/1206.5533** (2012)
- [23] Werbos, P.: Backpropagation through time: what does it do and how to do it. In: *Proceedings of IEEE*. Volume 78. (1990) 1550–1560
- [24] Chen, D., Manning, C.: A fast and accurate dependency parser using neural networks. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, Association for Computational Linguistics (2014) 740–750
- [25] Ramshaw, L., Marcus, M.: Text chunking using transformation-based learning. In: *Proceedings of the 3rd Workshop on Very Large Corpora*, Cambridge, MA, USA (1995) 84–94

- [26] Mesnil, G., Dauphin, Y., Yao, K., Bengio, Y., Deng, L., Hakkani-Tur, D., He, X., Heck, L., Tur, G., Yu, D., Zweig, G.: Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* (2015)
- [27] Lafferty, J., McCallum, A., Pereira, F.: Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: *Proceedings of the Eighteenth International Conference on Machine Learning (ICML)*, Williamstown, MA, USA (2001) 282–289
- [28] Toutanova, K., Klein, D., Manning, C.D., Singer, Y.: Feature-rich part-of-speech tagging with a cyclic dependency network. In: *NAACL '03: Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology*, Morristown, NJ, USA, Association for Computational Linguistics (2003) 173–180
- [29] Shen, L., Satta, G., Joshi, A.: Guided learning for bidirectional sequence classification. In: *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, Prague, Czech Republic, Association for Computational Linguistics (2007) 760–767
- [30] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* **15** (2014) 1929–1958
- [31] Hochreiter, S., Bengio, Y., Frasconi, P., Schmidhuber, J.: Gradient flow in recurrent nets: the difficulty of learning long-term. In Kremer, Kolen, eds.: *A Field Guide to Dynamical Recurrent Neural Networks*. IEEE Press (2001)