

Log-Linear RNNs : Towards Recurrent Neural Networks with Flexible Prior Knowledge

Marc Dymetman Chunyang Xiao
Xerox Research Centre Europe, Grenoble, France
{marc.dymetman, chunyang.xiao}@xrce.xerox.com

November 2016

Abstract

We introduce *LL-RNNs* (Log-Linear RNNs), an extension of Recurrent Neural Networks that replaces the softmax output layer by a log-linear output layer, of which the softmax is a special case. This conceptually simple move has two main advantages. First, it allows the learner to combat training data sparsity by allowing it to model words (or more generally, output symbols) as complex combinations of attributes without requiring that each combination is directly observed in the training data (as the softmax does). Second, it permits the inclusion of flexible prior knowledge in the form of *a priori* specified modular features, where the neural network component learns to dynamically control the weights of a log-linear distribution exploiting these features.

We conduct experiments in the domain of language modelling of French, that exploit morphological prior knowledge and show an important decrease in perplexity relative to a baseline RNN.

We provide other motivating illustrations, and finally argue that the log-linear and the neural-network components contribute complementary strengths to the LL-RNN: the LL aspect allows the model to incorporate rich prior knowledge, while the NN aspect, according to the “representation learning” paradigm, allows the model to discover novel combination of characteristics.

This is an updated version of the e-print arXiv:1607.02467, in particular now including experiments.

1 Introduction

Recurrent Neural Networks (Goodfellow et al., 2016, Chapter 10) have recently shown remarkable success in sequential data prediction and have been applied to such NLP tasks as Language Modelling (Mikolov et al., 2010), Machine Translation (Sutskever et al., 2014; Bahdanau et al., 2015), Parsing (Vinyals et al., 2014), Natural Language Generation (Wen et al., 2015) and Dialogue (Vinyals and Le, 2015), to name only a few. Specially popular RNN architectures in these applications have been models able to exploit long-distance correlations, such as LSTMs (Hochreiter and Schmidhuber, 1997; Gers et al., 2000) and GRUs (Cho et al., 2014), which have led to groundbreaking performances.

RNNs (or more generally, Neural Networks), at the core, are machines that take as input a real vector and output a real vector, through a combination of linear and non-linear operations.

When working with symbolic data, some conversion from these real vectors from and to discrete values, for instance words in a certain vocabulary, becomes necessary. However most RNNs have taken an oversimplified view of this mapping. In particular, for converting output vectors into distributions over symbolic values, the mapping has mostly been done through a *softmax* operation, which assumes that the RNN is able to compute a real value for each individual member of the vocabulary, and then converts this value into a probability through a direct exponentiation followed by a normalization.

This rather crude “softmax approach”, which implies that the output vector has the same dimensionality as the vocabulary, has had some serious consequences.

To focus on only one symptomatic defect of this approach, consider the following. When using words as symbols, even large vocabularies cannot account for all the actual words found either in training or in test, and the models need to resort to a catch-all “unknown” symbol *unk*, which provides a poor support for prediction and requires to be supplemented by diverse pre- and post-processing steps (Luong et al., 2014; Jean et al., 2015). Even for words inside the vocabulary, unless they have been witnessed many times in the training data, prediction tends to be poor because each word is an “island”, completely distinct from and without relation to other words, which needs to be predicted individually.

One partial solution to the above problem consists in changing the granularity by moving from word to character symbols (Sutskever et al., 2011; Ling et al., 2015). This has the benefit that the vocabulary becomes much

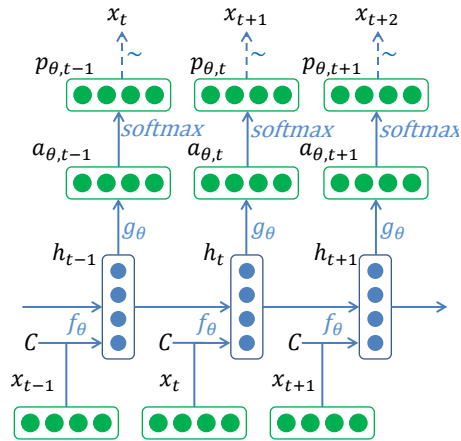


Figure 1: A generic RNN.

smaller, and that all the characters can be observed many times in the training data. While character-based RNNs have thus some advantages over word-based ones, they also tend to produce non-words and to necessitate longer prediction chains than words, so the jury is still out, with emerging hybrid architectures that attempt to capitalize on both levels (Luong and Manning, 2016).

Here, we propose a different approach, which removes the constraint that the dimensionality of the RNN output vector has to be equal to the size of the vocabulary and allows generalization across related words. However, its crucial benefit is that it introduces *a principled and powerful way of incorporating prior knowledge inside the models*.

The approach involves a very direct and natural extension of the softmax, by considering it as a special case of an *conditional exponential family*, a class of models better known as *log-linear* models and widely used in “pre-NN” NLP. We argue that this simple extension of the softmax allows the resulting “log-linear RNN” to compound the aptitude of log-linear models for *exploiting* prior knowledge and predefined features with the aptitude of RNNs for *discovering* complex new combinations of predictive traits.

2 Log-Linear RNNs

2.1 Generic RNNs

Let us first recap briefly the generic notion of RNN, abstracting away from different styles of implementation (LSTM (Hochreiter and Schmidhuber, 1997; Graves, 2012), GRU (Cho et al., 2014), attention models (Bahdanau et al., 2015), different number of layers, etc.).

An RNN is a generative process for predicting a sequence of symbols $x_1, x_2, \dots, x_t, \dots$, where the symbols are taken in some vocabulary V , and where the prediction can be conditioned by a certain observed context C . This generative process can be written as:

$$p_\theta(x_{t+1}|C, x_1, x_2, \dots, x_t),$$

where θ is a real-valued parameter vector.¹ Generically, this conditional probability is computed according to:

$$h_t = f_\theta(C; x_t, h_{t-1}), \tag{1}$$

$$a_{\theta,t} = g_\theta(h_t), \tag{2}$$

$$p_{\theta,t} = \text{softmax}(a_{\theta,t}), \tag{3}$$

$$x_{t+1} \sim p_{\theta,t}(\cdot). \tag{4}$$

Here h_{t-1} is the hidden state at the previous step $t-1$, x_t is the output symbol produced at that step and f_θ is a neural-network based function (e.g. a LSTM network) that computes the next hidden state h_t based on C , x_t , and h_{t-1} . The function g_θ ,² is then typically computed through an MLP, which returns a real-valued vector $a_{\theta,t}$ of dimension $|V|$. This vector is then normalized into a probability distribution over V through the softmax transformation:

$$\text{softmax}(a_{\theta,t})(x) = 1/Z \exp(a_{\theta,t}(x)),$$

with the normalization factor:

$$Z = \sum_{x' \in V} \exp(a_{\theta,t}(x')),$$

¹We will sometimes write this as $p_\theta(x_{t+1}|C; x_1, x_2, \dots, x_t)$ to stress the difference between the “context” C and the prefix x_1, x_2, \dots, x_t . Note that some RNNs are “non-conditional”, i.e. do not exploit a context C .

²We do not distinguish between the parameters for f and for g , and write θ for both.

and finally the next symbol x_{t+1} is sampled from this distribution. See Figure 1.

Training of such a model is typically done through back-propagation of the cross-entropy loss:

$$-\log p_{\theta}(\bar{x}_{t+1}|x_1, x_2, \dots, x_t; C),$$

where \bar{x}_{t+1} is the actual symbol observed in the training set.

2.2 Log-Linear models

Definition

Log-linear models play a considerable role in statistics and machine learning; special classes are often known through different names depending on the application domains and on various details: *exponential families* (typically for unconditional versions of the models) (Nielsen and Garcia, 2009) *maximum entropy* models (Berger et al., 1996; Jaynes, 1957), *conditional random fields* (Lafferty et al., 2001), binomial and multinomial *logistic regression* (Hastie et al., 2001, Chapter 4). These models have been especially popular in NLP, for example in Language Modelling (Rosenfeld, 1996), in sequence labelling (Lafferty et al., 2001), in machine translation (Berger et al., 1996; Och and Ney, 2002), to name only a few.

Here we follow the exposition (Jebara, 2013), which is useful for its broad applicability, and which defines a *conditional log-linear model* — which we could also call a *conditional exponential family* — as a model of the form (in our own notation):

$$p(x|K, a) = \frac{1}{Z(K, a)} b(K, x) \exp\left(a^{\top} \phi(K, x)\right). \quad (5)$$

Let us describe the notation:

- x is a variable in a set V , which we will take here to be discrete (i.e. countable), and sometimes finite.³ We will use the terms *domain* or *vocabulary* for this set.
- K is the *conditioning variable* (also called *condition*).

³The model is applicable over continuous (measurable) spaces, but to simplify the exposition we will concentrate on the discrete case, which permits to use sums instead of integrals.

- a is a parameter vector in \mathbb{R}^d , which (for reasons that will appear later) we will call the *adaptor* vector.⁴
- ϕ is a *feature function* $(K, x) \rightarrow \mathbb{R}^d$; note that we sometimes write $(x; K)$ or $(K; x)$ instead of (K, x) to stress the fact that K is a condition.
- b is a nonnegative function $(K, x) \rightarrow \mathbb{R}^+$; we will call it the *background* function of the model.⁵
- $Z(K, a)$, called the *partition* function, is a normalization factor:

$$Z(K, a) = \sum_x b(K, x) \exp\left(a^\top \phi(K, x)\right).$$

When the context is unambiguous, we will sometimes leave the condition K as well as the parameter vector a implicit, and also simply write Z instead of $Z(K, a)$; thus we will write:

$$p(x) = \frac{1}{Z} b(x) \exp\left(a^\top \phi(x)\right), \quad (6)$$

or more compactly:

$$p(x) \propto b(x) \exp\left(a^\top \phi(x)\right). \quad (7)$$

The background as a “prior”

If in equation (7) the background function is actually a normalized probability distribution over V (that is, $\sum_x b(x) = 1$) and if the parameter vector a is null, then the distribution p is identical to b .

Suppose that we have an initial belief that the parameter vector a should be close to a_0 , then by reparametrizing equation (7) in the form:

$$p(x) \propto b'(x) \exp\left(a'^\top \phi(x)\right), \quad (8)$$

with $b'(x) = b(x) \exp(a_0^\top \phi(x))$ and $a' = a - a_0$, then our initial belief is represented by taking $a' = 0$. In other words, we can always assume that our initial belief is represented by the background probability b' along with a null parameter vector $a' = 0$. Deviations from this initial belief are then representation by variations of the parameter vector away from 0 and a

⁴In the NLP literature, this parameter vector is often denoted by λ .

⁵Jebara (2013) calls it the *prior* of the family.

simple form of regularization can be obtained by penalizing some p -norm $\|a'\|_p$ of this parameter vector.⁶

Gradient of cross-entropy loss

An important property of log-linear models is that they enjoy an extremely intuitive form for the gradient of their log-likelihood (aka cross-entropy loss).

If \bar{x} is a training instance observed under condition K , and if the current model is $p(x|a, K)$ according to equation (5), its likelihood loss at \bar{x} is defined as: $-\log L = -\log p(\bar{x}|a, K)$. Then a simple calculation shows that the gradient $\frac{\partial \log L}{\partial a}$ (also called the ‘‘Fisher score’’ at \bar{x}) is given by:

$$\frac{\partial \log L}{\partial a} = \phi(\bar{x}; K) - \sum_{x \in V} p(x|a, K) \phi(x; K). \quad (9)$$

In other words, the gradient is minus the difference between the model expectation of the feature vector and its actual value at \bar{x} .⁷

2.3 Log-Linear RNNs

We can now define what we mean by a *log-linear RNN*. The model, illustrated in Figure 2, is similar to a standard RNN up to two differences:

The first difference is that we allow a more general form of input to the network at each time step; namely, instead of allowing only the latest symbol x_t to be used as input, along with the condition C , we now allow an arbitrary feature vector $\psi(C, x_1, \dots, x_t)$ to be used as input; this feature vector is of fixed dimensionality $|\psi|$, and we allow it to be computed in

⁶Contrarily to the generality of the presentation by Jebara (2013), many presentations of log-linear models in the NLP context do not make an explicit reference to b , which is then implicitly taken to be uniform. However, the more statistically oriented presentations (Jordan, 20XX; Nielsen and Garcia, 2009) of the strongly related (unconditional) exponential family models do, which makes the mathematics neater and is necessary in presence of non-finite or continuous spaces. One advantage of the explicit introduction of b , even for finite spaces, is that it makes it easier to speak about the prior knowledge we have about the overall process.

⁷More generally, if we have a training set consisting of N pairs of the form $(\bar{x}_n; K_n)$, then the gradient of the log-likelihood for this training set is given by:

$$\frac{\partial \log L}{\partial a} = \sum_{n=1}^N \left(\phi(\bar{x}_n; K_n) - \sum_{x \in V} p(x|a, K_n) \phi(x; K_n) \right).$$

In other words, this gradient is *the difference between the feature vectors at the true labels minus the expected feature vectors under the current distribution* (Jebara, 2013).

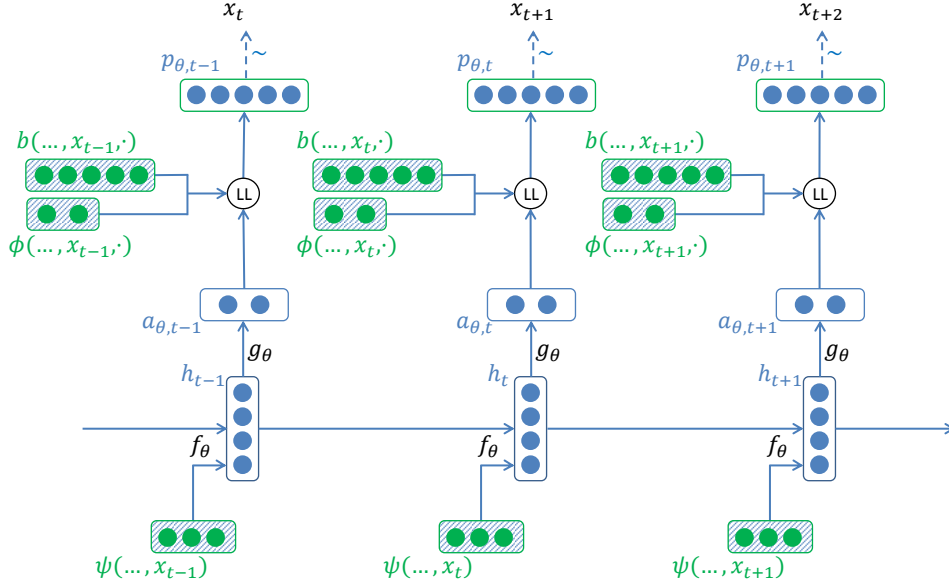


Figure 2: A Log-Linear RNN.

an *arbitrary* (but deterministic) way from the combination of the currently known prefix x_1, \dots, x_{t-1}, x_t and the context C . This is a relatively minor change, but one that usefully expands the expressive power of the network. We will sometimes call the ψ features the *input features*.

The second, major, difference is the following. We do compute $a_{\theta, t}$ in the same way as previously from h_t , however, after this point, rather than applying a softmax to obtain a distribution over V , we now apply a log-linear model. While for the standard RNN we had:

$$p_{\theta, t}(x_{t+1}) = \text{softmax}(a_{\theta, t})(x_{t+1}),$$

in the LL-RNN, we define:

$$p_{\theta, t}(x_{t+1}) \propto b(C, x_1, \dots, x_t, x_{t+1}) \exp\left(a_{\theta, t}^\top \phi(C, x_1, \dots, x_t, x_{t+1})\right). \quad (10)$$

In other words, we assume that we have *a priori* fixed a certain background function $b(K, x)$, where the condition K is given by $K = (C, x_1, \dots, x_t)$, and also defined M features defining a feature vector $\phi(K, x_{t+1})$, of fixed dimensionality $|\phi| = M$. We will sometimes call these features the *output features*. Note that both the background and the features have access to the context $K = (C, x_1, \dots, x_t)$.

In Figure 2, we have indicated with $\textcircled{\text{LL}}$ (LogLinear) the operation (10) that combines $a_{\theta,t}$ with the feature vector $\phi(C, x_1, \dots, x_t, x_{t+1})$ and the background $b(C, x_1, \dots, x_t, x_{t+1})$ to produce the probability distribution $p_{\theta,t}(x_{t+1})$ over V . We note that, here, $a_{\theta,t}$ is a vector of size $|\phi|$, which may or may not be equal to the size $|V|$ of the vocabulary, by contrast to the case of the softmax of Figure 1.

Overall, the LL-RNN is then computed through the following equations:

$$h_t = f_{\theta}(\psi(C, x_1, \dots, x_t), h_{t-1}), \quad (11)$$

$$a_{\theta,t} = g_{\theta}(h_t), \quad (12)$$

$$p_{\theta,t}(x) \propto b(C, x_1, \dots, x_t, x) \cdot \exp\left(a_{\theta,t}^{\top} \phi(C, x_1, \dots, x_t, x)\right), \quad (13)$$

$$x_{t+1} \sim p_{\theta,t}(\cdot). \quad (14)$$

For prediction, we now use the combined process p_{θ} , and we train this process, similarly to the RNN case, according to its cross-entropy loss relative to the actually observed symbol \bar{x} :

$$-\log p_{\theta}(\bar{x}_{t+1} | C, x_1, x_2, \dots, x_t). \quad (15)$$

At training time, in order to use this loss for backpropagation in the RNN, we have to be able to compute its gradient relative to the previous layer, namely $a_{\theta,t}$. From equation (9), we see that this gradient is given by:

$$\left(\sum_{x \in V} p(x | a_{\theta,t}, K) \phi(K; x) \right) - \phi(K; \bar{x}_{t+1}), \quad (16)$$

with $K = C, x_1, x_2, \dots, x_t$.

This equation provides a particularly intuitive formula for the gradient, namely, as the difference between the expectation of $\phi(K; x)$ according to the log-linear model with parameters $a_{\theta,t}$ and the observed value $\phi(K; \bar{x}_{t+1})$. However, this expectation can be difficult to compute. For a finite (and not too large) vocabulary V , the simplest approach is to simply evaluate the right-hand side of equation (13) for each $x \in V$, to normalize by the sum to obtain $p_{\theta,t}(x)$, and to weight each $\phi(K; x)$ accordingly. For standard RNNs (which are special cases of LL-RNNs, see below), this is actually what the simpler approaches to computing the softmax gradient do, but more sophisticated approaches have been proposed, such as employing a ‘‘hierarchical softmax’’ (Morin and Bengio, 2005). In the general case (large or infinite V),

the expectation term in (19) needs to be approximated, and different techniques may be employed, some specific to log-linear models (Elkan, 2008; Jebara, 2013), some more generic, such as contrastive divergence (Hinton, 2002) or Importance Sampling; a recent introduction to these generic methods is provided in (Goodfellow et al., 2016, Chapter 18), but, despite its practical importance, we will not pursue this topic further here.

2.4 LL-RNNs generalize RNNs

It is easy to see that LL-RNNs generalize RNNs. Consider a finite vocabulary V , and the $|V|$ -dim “one-hot” representation of $x \in V$, relative to a certain fixed ordering of the elements of V :

$$\text{oneHot}(x) = [0, 0, \dots, \underset{\substack{\uparrow \\ x}}{1}, \dots, 0].$$

We assume (as we implicitly did in the discussion of standard RNNs) that C is coded through some fixed-vector and we then define:

$$\psi(C, x_1, \dots, x_t) = C \oplus \text{oneHot}(x_t), \quad (17)$$

where \oplus denotes vector concatenation; thus we “forget” about the initial portion x_1, \dots, x_{t-1} of the prefix, and only take into account C and x_t , encoded in a similar way as in the case of RNNs.

We then define $b(x)$ to be uniformly 1 for all $x \in V$ (“uniform background”), and ϕ to be:

$$\phi(C, x_1, \dots, x_t, x_{t+1}) = \text{oneHot}(x_{t+1}).$$

Neither b nor ϕ depend on C, x_1, \dots, x_t , and we have:

$$p_{\theta,t}(x_{t+1}) \propto b(x_{t+1}) \exp\left(a_{\theta,t}^\top \phi(x_{t+1})\right) = \exp a_{\theta,t}(x_{t+1}),$$

in other words:

$$p_{\theta,t} = \text{softmax}(a_{\theta,t}).$$

Thus, we are back to the definition of RNNs in equations (1-4). As for the gradient computation of equation (19):

$$\left(\sum_{x \in V} p(x|a_{\theta,t}, K) \phi(K; x)\right) - \phi(K; \bar{x}_{t+1}), \quad (18)$$

it takes the simple form:

$$\left(\sum_{x \in V} p_{\theta,t}(x) \text{oneHot}(x) \right) - \text{oneHot}(\bar{x}_{t+1}), \quad (19)$$

in other words this gradient is the vector ∇ of dimension $|V|$, with coordinates $i \in 1, \dots, |V|$ corresponding to the different elements $x_{(i)}$ of V , where:

$$\nabla_i = \begin{cases} p_{\theta,t}(x_{(i)}) - 1 & \text{if } x_{(i)} = \bar{x}_{t+1}, \\ p_{\theta,t}(x_{(i)}) & \text{for the other } x_{(i)}\text{'s.} \end{cases} \quad (20a)$$

$$(20b)$$

This corresponds to the computation in the usual softmax case.

3 A motivating illustration: rare words

We now come back to the our starting point in the introduction: the problem of unknown or rare words, and indicate a way to handle this problem with LL-RNNs, which may also help building intuition about these models.

Let us consider some moderately-sized corpus of English sentences, tokenized at the word level, and then consider the vocabulary V_1 , of size 10K, consisting of the 9999 most frequent words to occur in this corpus plus one special symbol UNK used for tokens not among those words (“unknown words”).

After replacing the unknown words in the corpus by UNK, we can train a language model for the corpus by training a standard RNN, say of the LSTM type. Note that if translated into a LL-RNN according to section 2.4, this model has 10K features (9999 features for identity with a specific frequent word, the last one for identity with the symbol UNK), along with a uniform background b .

This model however has some serious shortcomings, in particular:

- Suppose that none of the two tokens *Grenoble* and *37* belong to V_1 (i.e. to the 9999 most frequent words of the corpus), then the learnt model cannot distinguish the probability of the two test sentences: *the cost was 37 euros* / *the cost was Grenoble euros*.
- Suppose that several sentences of the form *the cost was NN euros* appear in the corpus, with NN taking (say) values 9, 13, 21, all belonging to V_1 , and that on the other hand 15 also belongs to V_1 , but appears in non-cost contexts; then the learnt model cannot give a reasonable

probability to *the cost was 15 euros*, because it is unable to notice the similarity between 15 and the tokens 9, 13, 21.

Let's see how we can improve the situation by moving to a LL-RNN.

We start by extending V_1 to a much larger finite set of words V_2 , in particular one that includes all the words in the union of the training and test corpora,⁸ and we keep b uniform over V_2 . Concerning the ψ (input) features, for now we keep them at their standard RNN values (namely as in (17)). Concerning the ϕ features, we keep the 9999 word-identity features that we had, but not the UNK-identity one; however, we do add some new features (say $\phi_{10000} - \phi_{10020}$):

- A binary feature $\phi_{10000}(x) = \phi_{\text{number}}(x)$ that tells us whether the token x can be a number;
- A binary feature $\phi_{10001}(x) = \phi_{\text{location}}(x)$ that tells us whether the token x can be a location, such as a city or a country;
- A few binary features $\phi_{\text{noun}}(x)$, $\phi_{\text{adj}}(x)$, ..., covering the main POS's for English tokens. Note that a single word may have simultaneously several such features firing, for instance *flies* is both a noun and a verb.⁹
- Some other features, covering other important classes of words.

Each of the $\phi_1, \dots, \phi_{10020}$ features has a corresponding weight that we index in a similar way a_1, \dots, a_{10020} .

Note again that we *do allow* the features to overlap freely, nothing preventing a word to be both a location and an adjective, for example (e.g. *Nice* in *We visited Nice / Nice flowers were seen everywhere*), and to also appear in the 9999 most frequent words. For exposition reasons (ie in order to simplify the explanations below) we will suppose that a number N will always fire the feature ϕ_{number} , but no other feature, apart from the case where it also belongs to V_1 , in which case it will also fire the word-identity feature that corresponds to it, which we will denote by $\phi_{\tilde{N}}$, with $\tilde{N} \leq 9999$.

Why is this model superior to the standard RNN one?

To answer this question, let's consider the encoding of N in ϕ feature space, when N is a number. There are two slightly different cases to look at:

⁸We will see later that the restriction that V is finite can be lifted.

⁹Rather than using the notation ϕ_{10000}, \dots , we sometimes use the notation $\phi_{\text{number}}, \dots$, for obvious reasons of clarity.

1. N does not belong to V_1 . Then we have $\phi_{10000} = \phi_{\text{number}} = 1$, and $\phi_i = 0$ for other i 's.
2. N belongs to V_1 . Then we have $\phi_{10000} = \phi_{\text{number}} = 1$, $\phi_{\tilde{N}} = 1$ and $\phi_i = 0$ for other i 's.

Let us now consider the behavior of the LL-RNN during training, when at a certain point, let's say after having observed the prefix *the cost was*, it is now coming to the prediction of the next item $x_{t+1} = x$, which we assume is actually a number $\bar{x} = N$ in the training sample.

We start by assuming that N does not belong to V_1 .

Let us consider the current value $a = a_{\theta,t}$ of the weight vector calculated by the network at this point. According to equation (9), the gradient is:

$$\frac{\partial \log L}{\partial a} = \phi(N) - \sum_x p(x|a) \phi(x),$$

where L is the cross-entropy loss and p is the probability distribution associated with the log-linear weights a .

In our case the first term is a vector that is null everywhere but on coordinate ϕ_{number} , on which it is equal to 1. As for the second term, it can be seen as the model average of the feature vector $\phi(x)$ when x is sampled according to $p(x|a)$. One can see that this vector has all its coordinates in the interval $[0, 1]$, and in fact strictly between 0 and 1.¹⁰ As a consequence, the gradient $\frac{\partial \log L}{\partial a}$ is strictly positive on the coordinate ϕ_{number} and strictly negative on all the other coordinates. In other words, the backpropagation signal sent to the neural network at this point is that it should modify its parameters θ in such a way as to increase the a_{number} weight, and decrease all the other weights in a .

A slightly different situation occurs if we assume now that N belongs to V_1 . In that case $\phi(N)$ is null everywhere but on its two coordinates ϕ_{number} and $\phi_{\tilde{N}}$, on which it is equal to 1. By the same reasoning as before we see that the gradient $\frac{\partial \log L}{\partial a}$ is then strictly positive on the two corresponding coordinates, and strictly negative everywhere else. Thus, the signal sent to the network is to modify its parameter towards increasing the a_{number} and $a_{\tilde{N}}$ weights, and decrease them everywhere else.

Overall, on each occurrence of a number in the training set, the network is then learning to increase the weights corresponding to the features (either

¹⁰This last fact is because, for a vector a with finite coordinates, $p(x|a)$ can never be 0, and also because we are making the mild assumption that for any feature ϕ_i , there exist x and x' such that $\phi_i(x) = 0, \phi_i(x') = 1$; the strict inequalities follow immediately.

both a_{number} and $a_{\tilde{N}}$ or only a_{number} , depending on whether N is in V_1 or not) firing on this number, and to decrease the weights for all the other features. This contrasts with the behavior of the previous RNN model where only in the case of $N \in V_1$ did the weight $a_{\tilde{N}}$ change. This means that at the end of training, when predicting the word x_{t+1} that follows the prefix *The cost was*, the LL-RNN network will have a tendency to produce a weight vector $a_{\theta,t}$ with especially high weight on a_{number} , some positive weights on those $a_{\tilde{N}}$ for which N has appeared in similar contexts, and negative weights on features not firing in similar contexts.¹¹

Now, to come back to our initial example, let us compare the situation with the two next-word predictions *The cost was 37* and *The cost was Grenoble*. The LL-RNN model predicts the next word x_{t+1} with probability:

$$p_{\theta,t}(x_{t+1}) \propto \exp\left(a_{\theta,t}^\top \phi(x_{t+1})\right).$$

While the prediction $x_{t+1} = 37$ fires the feature ϕ_{number} , the prediction $x_{t+1} = \textit{Grenoble}$ does not fire any of the features that tend to be active in the context of the prefix *The cost was*, and therefore $p_{\theta,t}(37) \gg p_{\theta,t}(\textit{Grenoble})$. This is in stark contrast to the behavior of the original RNN, for which both *37* and *Grenoble* were undistinguishable unknown words.

We note that, while the model is able to capitalize on the generic notion of number through its feature ϕ_{number} , it is also able to learn to privilege certain specific numbers belonging to V_1 if they tend to appear more frequently in certain contexts. A log-linear model has the important advantage of being able to handle redundant features¹² such as ϕ_{number} and $\phi_{\bar{3}}$ which both fire on 3. Depending on prior expectations about typical texts in the domain being handled, it may then be useful to introduce features for distinguishing between different classes of numbers, for instance “small numbers” or “year-like numbers”, allowing the LL-RNN to make useful generalizations based on these features. Such features need not be binary, for example a small-number feature could take values decreasing from 1 to 0, with the higher values reserved for the smaller numbers.

While our example focussed on the case of numbers, it is clear that our observations equally apply to other features that we mentioned, such as

¹¹If only numbers appeared in the context *The cost was*, then this would mean *all* “non-numeric” features, but such words as *high*, *expensive*, *etc.* may of course also appear, and their associated features would also receive positive increments.

¹²This property of log-linear models was what permitted a fundamental advance in Statistical Machine Translation beyond the initial limited noisy-channel models, by allowing a freer combination of different assessments of translation quality, without having to bother about overlapping assessments (Berger et al., 1996; Och and Ney, 2002).

$\phi_{\text{location}}(x)$, which can serve to generalize predictions in such contexts as *We are travelling to*.

In principle, generally speaking, any features that can support generalization, such as features representing semantic classes (e.g. nodes in the Wordnet hierarchy), morphosyntactic classes (lemma, gender, number, etc.) or the like, can be useful.

4 Some potential applications

The extension from softmax to log-linear outputs, while formally simple, opens a significant range of potential applications other than the handling of rare words. We now briefly sketch a few directions.

***A priori* constrained sequences** For some applications, sequences to be generated may have to respect certain *a priori* constraints. One such case is the approach to semantic parsing of (Xiao et al., 2016), where starting from a natural language question an RNN decoder produces a sequential encoding of a logical form, which has to conform to a certain grammar. The model used is implicitly a simple case of LL-RNN, where (in our present terminology) the output feature vector ϕ remains the usual oneHot, but the background b is not uniform anymore, but constrains the generated sequence to conform to the grammar.

Language model adaptation We saw earlier that by taking b to be uniform and ϕ to be a oneHot, an LL-RNN is just a standard RNN. The opposite extreme case is obtained by supposing that we already *know* the exact generative process for producing x_{t+1} from the context $K = C, x_1, x_2, \dots, x_t$. If we define $b(K; \cdot) = b(K; x)$ to be identical to this true underlying process, then in order to have the best performance in test, it is sufficient for the adaptor vector $a_{\theta,t}$ to be equal to the null vector, because then, according to (13), $p_{\theta,t}(x) \propto b(K; x)$ is equal to the underlying process. The task for the RNN to learn a θ such that $a_{\theta,t}$ is null or close to null is an easy one (just take the higher level parameter matrices to be null or close to null), and in this case the adaptor has actually nothing to adapt to.

A more interesting, intermediary, case is when $b(K; x)$ is not too far from the true process. For example, b could be a word-based language model (n-gram type, LSTM type, etc.) trained on some large monolingual corpus, while the current focus is on modeling a specific domain for which much less data is available. Then training the RNN-based adaptor a_{θ} on the specific domain data would still be able to rely on b for test words not seen in the

specific data, but learn to upweight the prediction of words often seen in these specific data.¹³

Input features In a standard RNN, a word x_t is vector-encoded through a one-hot representation both when it is produced as the current output of the network but also when it is used as the next input to the network. In section 3, we saw the interest of defining the “output” features ϕ to go beyond word-identity features — i.e. beyond the identification $\phi(x) = \text{oneHot}(x)$ —, but we kept the “input” features as in standard RNNs, namely we kept $\psi(x) = \text{oneHot}(x)$. However, let us note an issue there. This usual encoding of the input x means that if $x = 37$ has rarely (or not at all) been seen in the training data, then the network will have few clues to distinguish this word from another rarely observed word (for example the adjective *preposterous*) when computing f_θ in equation 11. The network, in the context of the prefix *the cost was*, is able to give a reasonable probability to 37 thanks to ϕ . However, when assessing the probability of *euros* in the context of the prefix *the cost was 37*, this is not distinguished by the network from the prefix *the cost was preposterous*, which would not allow *euros* as the next word. A promising way to solve this problem here is to take $\psi = \phi$, namely to encode the input x using the same features as the output x . This allows the network to “see” that 37 is a number and that *preposterous* is an adjective, and to compute its hidden state based on this information. We should note, however, that there is no requirement that ψ be equal to ϕ in general; the point is that we can include in ψ features which can help the network predict the next word.

Infinite domains In the example of section 3, the vocabulary V_2 was large, but finite. This is quite artificial, especially if we want to account for words representing numbers, or words taken in some open-ended set, such as entity names. Let us go back to the equation (5) defining log-linear models, and let us ignore the context K for simplicity: $p(x|a) = \frac{1}{Z(a)} b(x) \exp(a^\top \phi(x))$, with $Z(a) = \sum_{x \in V} b(x) \exp(a^\top \phi(x))$. When V is finite, then the normalization factor $Z(a)$ is also finite, and therefore the probability $p(x|a)$ is well defined; in particular, it is well-defined when $b(x) = 1$ uniformly. However, when V is (countably) infinite, then this is unfortunately not true anymore. For instance, with $b(x) = 1$ uniformly, and with $a = 0$, then $Z(a)$ is infinite and the probability is undefined. By contrast, let’s assume that the background

¹³For instance, focussing on the simple case of an adaptor over a oneHot ϕ , as soon as $a_{\theta,t}(K; x)$ is positive on a certain word x , then the probability of this word is increased relative to what the background indicates.

function b is in $L_1(V)$, i.e. $\sum_{x \in V} b(x) < \infty$. Let’s also suppose that the feature vector ϕ is uniformly bounded (that is, all its coordinates ϕ_i are such that $\forall x \in V, \phi_i(x) \in [\alpha_i, \beta_i]$, for $\alpha_i, \beta_i \in \mathbb{R}$). Then, for any a , $Z(a)$ is finite, and therefore $p(x|a)$ is well-defined.

Thus, the standard RNNs, which have (implicitly) a uniform background b , have no way to handle infinite vocabularies, while LL-RNNs, by using a finite-mass b , can. One simple way to ensure that property on tokens representing numbers, for example, is to associate them with a geometric background distribution, decaying fast with their length, and a similar treatment can be done for named entities.

Condition-based priming Many applications of RNNs, such as machine translation (Sutskever et al., 2014) or natural language generation (Wen et al., 2015), etc., depend on a condition C (source sentence, semantic representation, etc.). When translated into LL-RNNs, this condition is taken into account through the input feature vector $\psi(C, x_1, \dots, x_t) = C \oplus \text{oneHot}(x_t)$, see (17), but does not appear in $b(C, x_1, \dots, x_t; x_{t+1}) = b(x_{t+1}) = 1$ or $\phi(C, x_1, \dots, x_t; x_{t+1}) = \text{oneHot}(x_{t+1})$.

However, there is opportunity for exploiting the condition inside b or ϕ . To sketch a simple example, in NLG, one may be able to predefine some weak unigram language model for the realization that depends on the semantic input C , for example by constraining named entities that appear in the realization to have some evidence in the input. Such a language model can be usefully represented through the background process $b(C, x_1, \dots, x_t; x_{t+1}) = b(C; x_{t+1})$, providing a form of “priming” for the combined LL-RNN, helping it to avoid irrelevant tokens.

A similar approach was recently exploited in Goyal et al. (2016), in the context of a character-based seq2seq LSTM for generating utterances from input “dialog acts”. In this approach, the background b , formulated as a weighted finite-state automaton over characters, is used both for encouraging the system to generate character strings that correspond to possible dictionary words, as well as to allow it to generate strings corresponding to such non-dictionary tokens as named-entities, numbers, addresses, and the like, but only when such strings have evidence in the input dialog act.

5 Experiments: French language model using morphological features

5.1 Datasets

Our datasets are based on the annotated French corpora¹⁴ provided by the *Universal Dependencies* initiative¹⁵. These corpora are tagged at the POS level as well as at the dependency level. In our experiments, we only exploit the POS annotations, and we use lowercased versions of the corpora.

Table 2 shows the sentence sizes of our different datasets, and Table 2 overall statistics in terms of word-tokens and word-types.

Training	Validation	Test1	Test2
13,826	728	298	1596

Table 1: Number of sentences in the different datasets. The UD (Universal Dependency) Training set (14554 sents.) is the union of our Training and Validation sets. The UD Validation set is our Test2 set, while the UD Test set is our Test1 set.

Total sentences	Tokens	Avg. sent. length	Types	Token-type ratio
16448	463069	28.15	42894	10.80

Table 2: Data statistics.

5.2 Features

The corpora provide POS and Morphological tags for each word token in the context of the sentence in which it appears. Table 3 shows the 52 tags that we use, which we treat as binary features. In addition, we select the M most frequent word types appearing in the entire corpus, and we use $M + 1$ additional binary features which identify whether a given word is identical to one of the M most frequent words, or whether it is outside this set. In total, we then use $M + 53$ binary features.

We collect all the word types appearing in the entire corpus and we associate with each a binary vector of size $M + 53$ which is the boolean union of the binary vectors associated with all the tokens for that type.

¹⁴<http://universaldependencies.org/#fr>.

¹⁵<http://universaldependencies.org> (Version 1).

POS:ADJ	Case:Abl	PronType:Dem
POS:ADP	Case:Acc	PronType:In
POS:ADV	Case:Nom	PronType:Int
POS:AUX	Case:Voc	PronType:Neg
POS:CONJ	Definite:Def	PronType:Prs
POS:DET	Definite:Ind	PronType:Rel
POS:INTJ	Degree:Cmp	Reflex:Yes
POS:NIL	Gender:Fem	Tense:Fut
POS:NOUN	Gender:Masc	Tense:Imp
POS:NUM	Gender:Neut	Tense:Past
POS:PART	Mood:Cnd	Tense:Pres
POS:PRON	Mood:Imp	VerbForm:Fin
POS:PROPN	Mood:IndN	VerbForm:Inf
POS:PUNC	Mood:SubT	VerbForm:Part
POS:SCON	Number:PlurJ	
POS:SYM	Number:Sing	
POS:VERB	Person:1	
POS:X	Person:2	
POS:_	Person:3	

Table 3: POS and Morphological features.

In case of an ambiguous word, the binary vector may have ones on several POS simultaneously.¹⁶ Thus, here, we basically use the corpus as a proxy for a morphological analyser of French, and we do not use the contextual information provided by the token-level tags.

5.3 Models

In these experiments, we use a finite word vocabulary V consisting of the 42894 types found in the entire corpus (including the validation and test sets). We then compare our LL-RNN with a vanilla RNN, both over this vocabulary V . Thus none of the models has unknown words. Both models are implemented in Keras (Chollet, 2015) over a Theano (Theano Development Team, 2016) backend.

The baseline RNN is using one-hot encodings for the words in V , and consists of an embedding layer of dimension 256 followed by two LSTM

¹⁶Thus for instance, currently, the vector associated with the word *le* has ones not only on the DET (determiner) and the PRON (pronoun) features, but also on the PROPN (proper noun) feature, due to the appearance of the name *Le* in the corpus...

(Hochreiter and Schmidhuber, 1997) layers of dimension 256, followed by a dense layer and finally a softmax layer both of dimension $|V|$. The LSTM sequence length for predicting the next word is fixed at 8 words. SGD is done through rmsprop and the learning rate is fixed at 0.001.

The LL-RNN has the same architecture and parameters, but for the following differences. First, the direct embedding of the input words is replaced by an embedding of dimension 256 of the representation of the words in the space of the $M + 53$ features (that is the input feature vector ψ is of dimension $M + 53$). This is followed by the same two LSTMs as before, both of dimension 256. This is now followed by a dense layer of output dimension $M + 53$ (the a_θ weights over the output feature vector ϕ , here identical to ψ). This layer is then transformed in a deterministic way into a probability distribution over V , after incorporation of a fixed background probability distribution b over V . This background b has been precomputed as the unigram probability distribution of the word types over the entire corpus.¹⁷

5.4 Results

Table 4 shows the perplexity results we obtain for different cases of the LL-RNNs as compared to the baseline RNN. We use a notation such as LL-RNN (2500) to indicate a value $M = 2500$ for the number of frequent word types considered as features. For each model, we stopped the training after the validation loss (not shown) did not improve for three epochs.

¹⁷We thus use also the test corpora to estimate these unigram probabilities. This is because the background requires to have some estimate of the probability of all the words it may encounter not only in training but also in test. However, this method is only a proxy to a proper estimate of the background for all possible words, which we leave to future development. We note that, similarly, for the baseline RNN, we need to know before hand all words it may encounter, otherwise we have to resort to the UNK category, which we did not want to do in order to be able to do a direct comparison of perplexities.

	Training	Test1	Test2
RNN	4.17	6.07	6.15
LL-RNN (10000)	4.68	5.17	5.17
LL-RNN (5000)	4.33	5.13	5.12
LL-RNN (3000)	4.65	5.11	5.09
LL-RNN (2500)	4.74	5.08	5.07
LL-RNN (2000)	4.77	5.11	5.10
LL-RNN (1000)	4.84	5.13	5.13
LL-RNN (500)	4.93	5.20	5.18
LL-RNN (10)	5.30	5.45	5.44

Table 4: Log-perplexities per word (base e) of different models. The perplexity per word corresponding to a log-perplexity of 5.08 (resp. 6.07) is 161 (resp. 433).

We observe a considerable improvement of perplexity between the baseline and all the LL-RNN models, the largest one being for $M = 2500$ — where the perplexity is divided by a factor of $433/161 \simeq 2.7$ — with some tendency of the models to degrade when M becomes either very large or very small.

An initial, informal, qualitative look at the sentences generated by the RNN model on the one hand and by the best LL-RNN model on the other hand, seems to indicate a much better ability of the LL-RNN to account for agreement in gender and number at moderate distances (see Table 5), but a proper evaluation has not yet been performed.

elle	TOPFORM:elle, POS:PRON, POS:PROPN, Gender:Fem, Number:Sing, Person:3, PronType:Prs
est	TOPFORM:est, POS:ADJ, POS:AUX, POS:NOUN, POS:PROPN, POS:SCONJ, POS:VERB, POS:X, Gender:Fem, Gender:Masc, Mood:Ind, Number:Sing, Person:3, Tense:Pres, VerbForm:Fin
très	TOPFORM:très, POS:ADV
souvent	TOPFORM:souvent, POS:ADV
représentée	TOPFORM:@notTop, POS:VERB, Gender:Fem, Number:Sing, Tense:Past, VerbForm:Part
en	TOPFORM:en, POS:ADP, POS:ADV, POS:PRON, Person:3
réaction	TOPFORM: réaction, POS:NOUN, Gender:Fem, Number:Sing
à	TOPFORM:à, POS:ADP, POS:AUX, POS:NOUN, POS:VERB, Mood:Ind, Number:Sing, Person:3, Tense:Pres, VerbForm:Fin
l'	TOPFORM:l', POS:DET, POS:PART, POS:PRON, POS:PROPN, Definite:Def, Gender:Fem, Gender:Masc, Number:Sing, Person:3, PronType:Prs
image	TOPFORM:image, POS:NOUN, Gender:Fem, Number:Sing
de	TOPFORM:de, POS:ADP, POS:DET, POS:PROPN, POS:X, Definite:Ind, Gender:Fem, Gender:Masc, Number:Plur, Number:Sing, PronType:Dem
la	TOPFORM:la, POS:ADP, POS:DET, POS:NOUN, POS:PRON, POS:PROPN, POS:X, Definite:Def, Gender:Fem, Gender:Masc, Number:Sing, Person:3, PronType:Prs
république	TOPFORM:république, POS:NOUN, POS:PROPN, Gender:Fem, Number:Sing
en	TOPFORM:en, POS:ADP, POS:ADV, POS:PRON, Person:3
1999	TOPFORM:1999, POS:NUM
.	TOPFORM:., POS:PUNCT

Table 5: Example of a sentence generated by LL-RNN (2500). The right column shows the non-null features for each word. Note that *représentée*, which is not among the most frequent 2500 words (TOPFORM:@notTop), has proper agreement (Gender:Fem, Number:Sing) with the distant pronoun *elle*.¹⁹

¹⁹As a side remark, we observe some flawed features, due to a small number of gold-annotation errors, such as the fact that *à* appears both with the correct POS:ADP (adposition - a generic term covering prepositions), but also some impossible POS's (AUX, NOUN, VERB). We have not attempted to filter out these (relatively rare) gold-annotation mistakes, but doing so could only improve the results.

6 Discussion

LL-RNNs simply extend RNNs by replacing the softmax parametrization of the output with a log-linear one, but this elementary move has two major consequences.

The first consequence is that two elements $x, x' \in V$, rather than being individuals without connections, can now share attributes. This is a fundamental property for linguistics, where classical approaches represent words as combination of “linguistic features”, such as POS, lemma, number, gender, case, tense, aspect, register, etc. With the standard RNN softmax approach, two words that are different on even a single dimension have to be predicted independently, which can only be done effectively in presence of large training sets. In the LL-RNN approach, by associating different ϕ features to the different linguistic “features”, the model can learn to predict a plural number based on observation of plural numbers, an accusative based on the observation of accusatives, and so on, and then predict word forms that are combinations that have never been observed in the training data. We saw an example of this phenomenon in the experiments of section 5.²⁰ If the linguistic features encompass semantic classes (possibly provided by Wordnet, or else by semantically-oriented embeddings) then generalizations become possible over these semantic classes also. By contrast, in the softmax case, not only the models are deficient in presence of sparsity of training data for word forms, but they also require to waste capacity of the RNN parameters θ to make them able to map to the large a_θ vectors that are required to discriminate between the many elements of V ; with LL-based RNNs, the parametrization a_θ can in principle be smaller, because fewer ϕ features need to be specified to obtain word level predictions.

The second consequence is that we can exploit rich prior knowledge through the input features ψ , the background b , and the output features ϕ . We already gave some illustrations of incorporating prior knowledge in this way, but there are many other possibilities. For example, in a dialogue application that requires some answer utterances to contain numerical data that can only be obtained by access to a knowledge base, a certain binary “expert feature” $\phi_e(K; x)$ could take the value 1 if and only if x is *either* a non-number word *or* a specific number n obtained by some (more or less complex) process exploiting the context K in conjunction with the knowledge base. In combination with a background b and other features in ϕ ,

²⁰Similar observations have been done, in the quite different “factored” model recently proposed by García-Martínez et al. (2016).

who would be responsible for the linguistic quality of the answer utterance, the ϕ_e feature, when activated, would ensure that *if* a number is produced at this point, it is equal to n , but would not try to decide at exactly *which* point a number should be produced (this is better left to the “language specialists”: b and the other features). *Whether* the feature ϕ_e is activated would be decided by the RNN: a large value of the e coordinate of $a_{\theta,t}$ would activate the feature, a small (close to null) value deactivate it.²¹

We conclude by a remark concerning the complementarity of the log-linear component and the neural network component in the LL-RNN approach. On its own, as has been amply demonstrated in recent years, a standard softmax-based RNN is already quite powerful. On *its* own, a stand-alone log-linear model is *also* quite powerful, as older research also demonstrated. Roughly, the difference between a log-linear model and a LL-RNN model is that in the first, the log-linear weights (in our notation, a) are fixed after training, while in the LL-RNN they dynamically vary under the control of the neural network component.²² However, the strengths of the two classes of models lie in different areas. The log-linear model is very good at exploiting prior knowledge in the form of complex features, but it has no ability to discover new combinations of features. On the other hand, the RNN is very good at discovering which combinations of characteristics of its input are predictive of the output (representation learning), but is ill-equipped for exploiting prior knowledge. We argue that the LL-RNN approach is a way to capitalize on these complementary qualities.

²¹The idea is reminiscent of the approach of Le et al. (2016), who use LSTM-based mixtures of experts for a similar purpose; the big difference is that here, instead of using a linear mixture, we use a “log-linear mixture”, i.e. our features are combined multiplicatively rather than additively, with exponents given by the RNN, that is they are “collaborating”, while in their approach the experts are “competing”: their expert corresponding to ϕ_e needs to decide on its own at which exact point it should produce the number, rather than relying on the linguistic specialist to do it.

This “multiplicative” aspect of the LL-RNNs can be related to the *product of experts* introduced by Hinton (2002). However, in his case, the focus is on learning the individual experts, which are then combined through a direct product, not involving exponentiations, and therefore not in the log-linear class. In our case, the focus is on exploiting predefined experts (or features), but on letting a “controlling” RNN decide about their exponents.

²²Note how a standard log-linear model with oneHot features over V would not make sense: with a fixed, it would always predict the same distribution for the next word. By contrast, a LL-RNN over the same features does make sense: it is a standard RNN. Standard log-linear models have to employ more interesting features.

Acknowledgments

We thank Salah Ait-Mokhtar, Matthias Gallé, Claire Gardent, Éric Gaussier, Raghav Goyal and Florent Perronnin for discussions at various stages of this research.

References

- Bahdanau, D., Cho, K., and Bengio, Y. (2015). Neural Machine Translation by Jointly Learning to Align and Translate. *ICLR*, pages 1–15.
- Berger, A. L., Della Pietra, S. A., and Della Pietra, V. J. (1996). A maximum entropy approach to natural language processing. *Computational linguistics*, 22(1):39–71.
- Cho, K., van Merriënboer, B., Bahdanau, D., and Bengio, Y. (2014). On the Properties of Neural Machine Translation: Encoder-Decoder Approaches. *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111.
- Chollet, F. (2015). Keras. <https://github.com/fchollet/keras>.
- Elkan, C. (2008). Log-linear models and conditional random fields. *Tutorial notes at CIKM*, 8:1–12.
- García-Martínez, M., Barrault, L., and Bougares, F. (2016). Factored Neural Machine Translation. *arXiv :1609.04621*.
- Gers, F. A., Schmidhuber, J. A., and Cummins, F. A. (2000). Learning to forget: Continual prediction with lstm. *Neural Comput.*, 12(10):2451–2471.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). Deep learning. Book in preparation for MIT Press.
- Goyal, R., Dymetman, M., and Gaussier, E. (2016). Natural Language Generation through Character-based RNNs with Finite-State Prior Knowledge. In *Proc. COLING*, Osaka, Japan.
- Graves, A. (2012). *Supervised sequence labelling with recurrent neural networks*, volume 385. Springer.

- Hastie, T., Tibshirani, R., and Friedman, J. (2001). *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York Inc., New York, NY, USA.
- Hinton, G. E. (2002). Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Phys. Rev.*, 106:620–630.
- Jean, S., Cho, K., Memisevic, R., and Bengio, Y. (2015). On Using Very Large Target Vocabulary for Neural Machine Translation. *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 000:1–10.
- Jebara, T. (2013). Log-Linear Models, Logistic Regression and Conditional Random Fields. Lecture notes: www.cs.columbia.edu/~jebara/6772/notes/notes4.pdf.
- Jordan, M. (20XX). The exponential family: Basics. Lecture notes: <http://people.eecs.berkeley.edu/~jordan/courses/260-spring10/other-readings/chapter8.pdf>.
- Lafferty, J. D., McCallum, A., and Pereira, F. C. N. (2001). Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01*, pages 282–289, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Le, P., Dymetman, M., and Renders, J.-M. (2016). LSTM-based mixture-of-experts for knowledge-aware dialogues. In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pages 94–99, Berlin, Germany. Association for Computational Linguistics.
- Ling, W., Trancoso, I., Dyer, C., and Black, A. W. (2015). Character-based Neural Machine Translation. *ICLR'16*, pages 1–11.
- Luong, M.-T. and Manning, C. D. (2016). Achieving Open Vocabulary Neural Machine Translation with Hybrid Word-Character Models. *arXiv:1604.00788v2*.

- Luong, M.-T., Sutskever, I., Le, Q. V., Vinyals, O., and Zaremba, W. (2014). Addressing the Rare Word Problem in Neural Machine Translation. *arXiv:1410.8206v3*.
- Mikolov, T., Karafiat, M., Burget, L., Cernocky, J., and Khudanpur, S. (2010). Recurrent Neural Network based Language Model. *Interspeech*, (September):1045–1048.
- Morin, F. and Bengio, Y. (2005). Hierarchical probabilistic neural network language model. *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics*, pages 246–252.
- Nielsen, F. and Garcia, V. (2009). Statistical exponential families: A digest with flash cards. *arXiv:0911.4863*.
- Och, F. J. and Ney, H. (2002). Discriminative training and maximum entropy models for statistical machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL '02*, pages 295–302, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Rosenfeld, R. (1996). A maximum entropy approach to adaptive statistical language modelling. *Computer Speech & Language*, 10(3):187–228.
- Sutskever, I., Martens, J., and Hinton, G. (2011). Generating Text with Recurrent Neural Networks. *Neural Networks*, 131(1):1017–1024.
- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Theano Development Team (2016). Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688.
- Vinyals, O., Kaiser, L., Koo, T., Petrov, S., Sutskever, I., and Hinton, G. (2014). Grammar as a Foreign Language. *arXiv: 1412.7449v3*.
- Vinyals, O. and Le, Q. (2015). A neural conversational model. *arXiv:1506.05869*.
- Wen, T., Gasic, M., Mrksic, N., Su, P., Vandyke, D., and Young, S. J. (2015). Semantically conditioned lstm-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on*

Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015, pages 1711–1721.

Xiao, C., Dymetman, M., and Gardent, C. (2016). Sequence-based structured prediction for semantic parsing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1341–1350, Berlin, Germany. Association for Computational Linguistics.