

Capturing Long-range Contextual Dependencies with Memory-enhanced Conditional Random Fields

Fei Liu Timothy Baldwin Trevor Cohn

School of Computing and Information Systems

The University of Melbourne

Victoria, Australia

fliu3@student.unimelb.edu.au

tb@ldwin.net t.cohn@unimelb.edu.au

Abstract

Despite successful applications across a broad range of NLP tasks, conditional random fields (“CRFs”), in particular the linear-chain variant, are only able to model local features. While this has important benefits in terms of inference tractability, it limits the ability of the model to capture long-range dependencies between items. Attempts to extend CRFs to capture long-range dependencies have largely come at the cost of computational complexity and approximate inference. In this work, we propose an extension to CRFs by integrating external memory, taking inspiration from memory networks, thereby allowing CRFs to incorporate information far beyond neighbouring steps. Experiments across two tasks show substantial improvements over strong CRF and LSTM baselines.

1 Introduction

While long-range contextual dependencies are prevalent in natural language, for tractability reasons, most statistical models capture only local features (Finkel et al., 2005). Take the sentence in Figure 1, for example. Here, while it is easy to determine that *Interfax* in the second sentence is a named entity, it is hard to determine its semantic class, as there is little context information. The usage in the first sentence, on the other hand, can be reliably disambiguated due to the post-modifying phrase *news agency*. Ideally we would like to be able to share such contexts across all usages (and variants) of a given named entity for reliable and consistent identification and disambiguation.

A related example is forum thread discourse analysis. Previous work has largely focused on

linear-chain Conditional Random Fields (CRFs) (Wang et al., 2011; Zhang et al., 2017), framing the task as one of sequence tagging. Although CRFs are adept at capturing local structure, the problem does not naturally suit a linear sequential structure, i.e., a post may be a reply to either a neighbouring post or one posted far earlier within the same thread. In both cases, contextual dependencies can be long-range, necessitating the ability to capture dependencies between arbitrarily distant items. Identifying this key limitation, Sutton and McCallum (2004) and Finkel et al. (2005) proposed the use of CRFs with skip connections to incorporate long-range dependencies. In both cases the graph structure must be supplied a priori, rather than learned, and both techniques incur the need for costly approximate inference.

Recurrent neural networks (RNNs) have been proposed as an alternative technique for encoding sequential inputs, however plain RNNs are unable to capture long-range dependencies (Bengio et al., 1994; Hochreiter et al., 2001) and variants such as LSTMs (Hochreiter and Schmidhuber, 1997), although more capable of capturing non-local patterns, still exhibit a significant locality bias in practice (Lai et al., 2015; Linzen et al., 2016).

In this paper, taking inspiration from the work of Weston et al. (2015) on memory networks (MEMNETS), we propose to extend CRFs by integrating external memory mechanisms, thereby enabling the model to look beyond localised features and have access to the entire sequence. This is achieved with attention over every entry in the memory. Experiments on named entity recognition and forum thread parsing, both of which involve long-range contextual dependencies, demonstrate the effectiveness of the proposed model, achieving state-of-the-art performance on the former, and outperforming a number of strong baselines in the case of the latter. A full imple-




Figure 1: A NER example with long-range contextual dependencies. The vertical dash line indicates a sentence boundary.

mentation of the model is available at: <https://github.com/liufly/mecrf>.

The paper is organised as follows: after reviewing previous studies on capturing long range contextual dependencies and related models in Section 2, we detail the elements of the proposed model in Section 3. Section 4 and 5 present the experimental results on two different datasets: one for thread discourse structure prediction and the other named entity recognition (NER), with analyses and visualisation in their respective sections. Lastly, Section 6 concludes the paper.

2 Related Work

In this section, we review the different families of models that are relevant to this work, in capturing long-range contextual dependencies in different ways.

Conditional Random Fields (CRFs). CRFs (Lafferty et al., 2001), in particular linear-chain CRFs, have been widely adopted and applied to sequence labelling tasks in NLP, but have the critical limitation that they only capture local structure (Sutton and McCallum, 2004; Finkel et al., 2005), despite non-local structure being common in structured language classification tasks. In the context of named entity recognition (“NER”), Sutton and McCallum (2004) proposed skip-chain CRFs as a means of alleviating this shortcoming, wherein distant items are connected in a sequence based on a heuristic such as string identity (to achieve label consistency across all instances of the same string). The idea of label consistency and exploiting non-local features has also been explored in the work of Finkel et al. (2005), who take long-range structure into account while maintaining tractable inference with Gibbs sampling (Geman and Geman, 1984), by performing approximate inference over factored probabilistic models. While both of these lines of work report impressive results on information extraction tasks, they come at the price of high computational cost and

incompatibility with exact inference.

Similar ideas have also been explored by Krishnan and Manning (2006) for NER, where they apply two CRFs, the first of which makes predictions based on local information, and the second combines named entities identified by the first CRF in a single cluster, thereby enforcing label consistency and enabling the use of a richer set of features to capture non-local dependencies. Liao and Grishman (2010) make a strong case for going beyond sentence boundaries and leveraging document-level information for event extraction.

While we take inspiration from these earlier studies, we do not enforce label consistency as a hard constraint, and additionally do not sacrifice inference tractability: our model is capable of incorporating non-local features, and is compatible with exact inference methods.

Recurrent Neural Networks (RNNs). Recently, the broad adoption of deep learning methods in NLP has given rise to the prevalent use of RNNs. Long short-term memories (“LSTMs”: Hochreiter and Schmidhuber (1997)), a particular variant of RNN, have become particularly popular, and been successfully applied to a large number of tasks: speech recognition (Graves et al., 2013), sequence tagging (Huang et al., 2015), document categorisation (Yang et al., 2016), and machine translation (Cho et al., 2014; Bahdanau et al., 2014). However, as pointed out by Lai et al. (2015) and Linzen et al. (2016), RNNs — including LSTMs — are biased towards immediately preceding (or neighbouring, in the case of bi-directional RNNs) items, and perform poorly in contexts which involve long-range contextual dependencies, despite the inclusion of memory cells. This is further evidenced by the work of Cho et al. (2014), who show that the performance of a basic encoder–decoder deteriorates as the length of the input sentence increases.

Memory networks (MEMNETs). More recently, Weston et al. (2015) proposed memory

networks and showed that the augmentation of memory is crucial to performing inference requiring long-range dependencies, especially when document-level reasoning between multiple supporting facts is required. Of particular interest to our work are so-called “memory hops” in memory networks, which are guided by an attention mechanism based on the relevance between a question and each supporting context sentence in the memory hop. Governed by the attention mechanism, the ability to access the entire sequence is similar to the soft alignment idea proposed by Bahdanau et al. (2014) for neural machine translation. In this work, we borrow the concept of memory hops and integrate it into CRFs, thereby enabling the model to look beyond localised features and have access to the whole sequence via an attention mechanism.

3 Methodology

In the context of sequential tagging, we assume the input is in the form of sequence pairs: $\mathcal{D} = \{\mathbf{x}^{(n)}, \mathbf{y}^{(n)}\}_{n=1}^N$ where $\mathbf{x}^{(n)}$ is the input of the n -th example in dataset \mathcal{D} and consists of a sequence: $\{x_1^{(n)}, x_2^{(n)}, \dots, x_T^{(n)}\}$. Similarly, $\mathbf{y}^{(n)}$ is of the same length as $\mathbf{x}^{(n)}$ and consists of the corresponding labels $\{y_1^{(n)}, y_2^{(n)}, \dots, y_T^{(n)}\}$. For notational convenience, hereinafter we omit the superscript denoting the n -th example.

In the case of NER, each x_t is a word in a sentence with y_t being the corresponding NER label. For forum thread discourse analysis, x_t represents the text of an entire post whereas y_t is the dialogue act label for the t -th post.

The proposed model extends CRFs by integrating external memory and is therefore named a **Memory-Enhanced Conditional Random Field** (“ME-CRF”). We take inspiration from Memory Networks (“MEMNETs”: Weston et al. (2015)) and incorporate so-called memory hops into CRFs, thereby allowing the model to have unrestricted access to the whole sequence rather than localised features as in RNNs (Lai et al., 2015; Linzen et al., 2016).

As illustrated in Figure 2, ME-CRF can be divided into two major parts: (1) the memory layer; and (2) the CRF layer. The memory layer can be further broken down into three main components: (a) the input memory $\mathbf{m}_{1:t}$; (b) the output memory $\mathbf{c}_{1:t}$; and (c) the current input \mathbf{u}_t , which represents the current step (also known as the “question” in the context of MEMNET). The input and

output memory representations are connected via an attention mechanism whose weights are determined by measuring the similarity between the input memory and the current input. The CRF layer, on the other hand, takes the output of the memory layer as input. In the remainder of this section, we detail the elements of ME-CRF.

3.1 Memory Layer

3.1.1 Input memory

Every element (word/post) in a sequence \mathbf{x} is encoded with $\mathbf{x}_t = \Phi(x_t)$, where $\Phi(\cdot)$ can be any encoding function mapping the input x_t into a vector $\mathbf{x}_t \in \mathbb{R}^d$. This results in the sequence $\{\mathbf{x}_1, \dots, \mathbf{x}_T\}$. While this new sequence can be seen as the memory in the context of MEMNETs, one major drawback of this approach, as pointed out by Seo et al. (2017), is the insensitivity to temporal information between memory cells. We therefore follow Xiong et al. (2016) in injecting temporal signal into the memory using a bi-directional GRU encoding (Cho et al., 2014):

$$\vec{\mathbf{m}}_t = \overrightarrow{\text{GRU}}(\mathbf{x}_t, \vec{\mathbf{m}}_{t-1}) \quad (1)$$

$$\overleftarrow{\mathbf{m}}_t = \overleftarrow{\text{GRU}}(\mathbf{x}_t, \overleftarrow{\mathbf{m}}_{t+1}) \quad (2)$$

$$\mathbf{m}_t = \tanh(\overrightarrow{\mathbf{W}}_m \vec{\mathbf{m}}_t + \overleftarrow{\mathbf{W}}_m \overleftarrow{\mathbf{m}}_t + \mathbf{b}_m) \quad (3)$$

where $\overrightarrow{\mathbf{W}}_m$, $\overleftarrow{\mathbf{W}}_m$ and \mathbf{b}_m are learnable parameters.

3.1.2 Current input

This is used to represent the current step x_t , be it a word or a post. As in MEMNETs, we want to enforce the current input to be in the same space as the input memory so that we can determine the attention weight of each element in the memory by measuring the relevance between the two. We denote the current input by $\mathbf{u}_t = \mathbf{m}_t$.

3.1.3 Attention

To determine the attention value of each element in the memory, we measure the relevance between the current step \mathbf{u}_t and \mathbf{m}_i for $i \in [1, t]$ with a softmax function:

$$p_{t,i} = \text{softmax}(\mathbf{u}_t^\top \mathbf{m}_i) \quad (4)$$

where $\text{softmax}(a_i) = \frac{e^{a_i}}{\sum_j e^{a_j}}$.

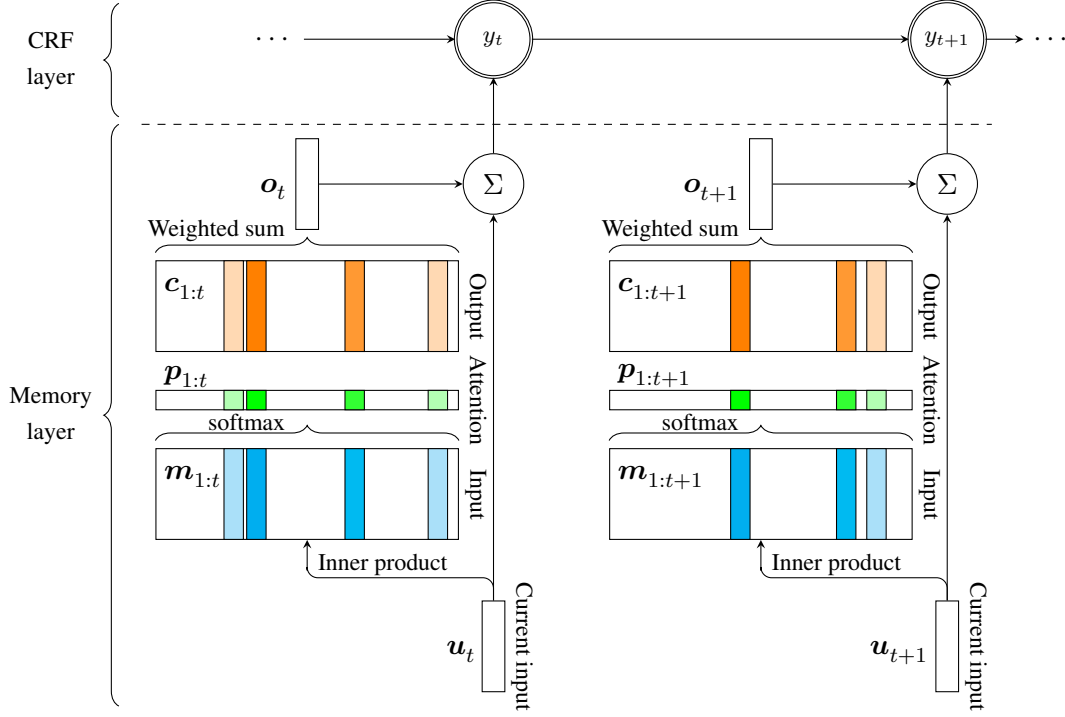


Figure 2: Illustration of ME-CRFs with a single memory hop, showing the network architecture at time step t and $t + 1$.

3.1.4 Output memory

Similar to m_t , c_t is the output memory, and is calculated analogously but with a different set of parameters in the GRUs and tanh layers of Equations (1), (2) and (3). The output memory is used to generate the final output of the memory layer and fed as input to the CRF layer.

3.1.5 Memory layer output

Once the attention weights have been computed, the memory access controller receives the response o in the form of a weighted sum over the output memory representations:

$$o_t = \sum_i p_{t,i} c_i \quad (5)$$

This allows the model to have unrestricted access to elements in previous steps as opposed to a single vector h_t in RNNs, thereby enabling ME-CRFs to detect and effectively incorporate long-range dependencies.

3.1.6 Extension

For more challenging tasks requiring complex reasoning capabilities with multiple supporting facts from the memory, the model can be further extended by stacking multiple memory hops, in

which case the output of the k -th hop is taken as input to the $(k + 1)$ -th hop:

$$u_t^{k+1} = o_t^k + u_t^k \quad (6)$$

where u_t^{k+1} encodes not only information at the current step (u_t^k) but also relevant knowledge from the memory (o_t^k). In the scope of this work, we limit the number of hops to 1.

3.2 CRF Layer

Once the representation of the current step u_t^{K+1} is computed — incorporating relevant information from the memory (assuming the total number of memory hops is K) — it is then fed into a CRF layer:

$$s(\mathbf{x}, \mathbf{y}) = \sum_{t=0}^T \mathbf{A}_{y_t, y_{t+1}} + \sum_{t=1}^T \mathbf{P}_{t, y_t} \quad (7)$$

where $\mathbf{A} \in \mathbb{R}^{|\mathcal{Y}| \times |\mathcal{Y}|}$ is the CRF transition matrix, $|\mathcal{Y}|$ is the size of the label set, and $\mathbf{P} \in \mathbb{R}^{T \times |\mathcal{Y}|}$ is a linearly transformed matrix from u_t^{K+1} such that $\mathbf{P}_{t,:}^\top = \mathbf{W}_s u_t^{K+1}$ where $\mathbf{W}_s \in \mathbb{R}^{|\mathcal{Y}| \times h}$ with h being the size of m_t . Here, $A_{i,j}$ represents the score of the transition from the i -th tag to the j -th tag whereas $P_{i,j}$ is the score of the j -th tag at time i .

Using the scoring function in Equation (7), we calculate the score of the sequence \mathbf{y} normalised by the sum of scores of all possible sequences $\tilde{\mathbf{y}}$, and this becomes the probability of the true sequence:

$$p(\mathbf{y}|\mathbf{x}) = \frac{\exp(s(\mathbf{x}, \mathbf{y}))}{\sum_{\tilde{\mathbf{y}} \in \mathbf{Y}_{\mathbf{x}}} \exp(s(\mathbf{x}, \tilde{\mathbf{y}}))} \quad (8)$$

We train the model to maximise the probability of the gold label sequence with the following loss function:

$$\mathcal{L} = \sum_{n=1}^N \log p(\mathbf{y}^{(n)}|\mathbf{x}^{(n)}) \quad (9)$$

where $p(\mathbf{y}^{(n)}|\mathbf{x}^{(n)})$ is calculated using the forward-backward algorithm. Note that the model is fully end-to-end differentiable.

At test time, the model predicts the output sequence with maximum a posteriori probability:

$$\mathbf{y}^* = \arg \max_{\tilde{\mathbf{y}} \in \mathbf{Y}_{\mathbf{x}}} p(\tilde{\mathbf{y}}|\mathbf{x}) \quad (10)$$

Since we are only modelling bigram interactions, we adopt the Viterbi algorithm for decoding.

4 Thread Discourse Structure Prediction

In this section, we describe how ME-CRFs can be applied to the task of thread discourse structure prediction, wherein we attempt to predict which post(s) a given post directly responds to, and in what way(s) (as captured by dialogue acts). This is a novel approach to this problem and capable of natively handling both tasks within the same network architecture.

4.1 Dataset and Task

In this work, we adopt the dataset of Kim et al. (2010),¹ which consists of 315 threads and 1,332 posts, collected from the Operating System, Software, Hardware and Web Development subforums of CNET.² Every post has been manually linked to preceding post(s) in the thread that it is a direct response to (in the form of “links”), and the nature of the response for each link (in the form of “dialogue acts”, or “DAs”). In this dataset, it is not uncommon to see messages respond to posts which occur much earlier in the thread (based on the chronological ordering of posts). In fact, 18%

¹<http://people.eng.unimelb.edu.au/tbaldwin/resources/con112010-thread/>

²<http://forums.cnet.com/>

of the posts link to posts other than their immediately preceding post.

The task is defined as follows: given a list of preceding posts x_1, \dots, x_{t-1} and the current post x_t , to classify which posts it links to (l_t) and the dialogue act (y_t) of each such link. In the scope of this work, ME-CRFs are capable of modelling both tasks natively, and therefore a natural fit for this problem.

4.2 Experimental Setup

In this dataset, in addition to the body of text, each post is also associated with a title. We therefore use two encoders, $\Phi_{title}(\cdot)$ and $\Phi_{text}(\cdot)$, to process them separately and then concatenate $\mathbf{x}_t = [\Phi_{title}(x_t); \Phi_{text}(x_t)]^\top$. Here, $\Phi_{title}(\cdot)$ and $\Phi_{text}(\cdot)$ take word embeddings as input, processing each post at the word level, as opposed to the post-level bi-directional GRU in Equations (1) and (2), and the representation of a post \mathbf{x}_t (either title or text) is obtained by transforming the last and first hidden states of the forward and backward word-level GRU, similar to Equation (3). Note that $\Phi_{title}(\cdot)$ and $\Phi_{text}(\cdot)$ do not share parameters. As in Tang et al. (2016), we further restrict $\mathbf{m}_i^k = \mathbf{c}_i^k$ to curb overfitting.

In keeping with Wang (2014), we complement the textual representations with hand-crafted structural features $\Phi_s(x_t) \in \mathbb{R}^2$:

- initiator: a binary feature indicating whether the author of the current post is the same as the initiator of the thread,
- position: the relative position of the current post, as a ratio over the total number of posts in the thread;

Also drawing on Wang (2014), we incorporate punctuation-based features $\Phi_p(x_t) \in \mathbb{R}^3$: the number of question marks, exclamation marks and URLs in the current post. The resultant feature vectors are projected into an embedding space by \mathbf{W}_s and \mathbf{W}_p and concatenated with \mathbf{x}_i , resulting in the new \mathbf{x}'_i . Subsequently, \mathbf{x}'_i is fed into the bi-directional GRUs to obtain \mathbf{m}_i .

For link prediction, we generate a supervision signal from the annotated links, guiding the attention to focus on the correct memory position:

$$\mathcal{L}_{\text{LNK}} = \sum_{n=1}^N \sum_{t=1}^T \text{CrossEntropy}(\mathbf{l}_t^{(n)}, \mathbf{p}_t^{(n)}) \quad (11)$$

where $\mathbf{l}_t^{(n)}$ is a one-hot vector indicating where the link points to for the t -th post in the n -th thread,

and $\mathbf{p}_t^{(n)} = \{p_{t,1}, \dots, p_{t,t}\}$ is the predicted distribution of attention over the t posts in the memory. To accommodate the first post in a thread, as it points to a virtual “head” post, we set a dummy post, $\mathbf{m}_0 = \mathbf{0}$, of the same size as \mathbf{m}_i . While the dataset contains multi-headed posts (posts with more than one outgoing link), following Wang (2014), we only include the most recent linked post during training, but evaluate over the full set of labelled links.

For this task, ME-CRF is jointly trained to predict both the link and dialogue act with the following loss function:

$$\mathcal{L}' = \alpha \mathcal{L}_{\text{DA}} + (1 - \alpha) \mathcal{L}_{\text{LNK}} \quad (12)$$

where \mathcal{L}_{DA} is the CRF likelihood defined in Equation (9), and α is a hyper-parameter for balancing the emphasis between the two tasks.

Training is carried out with Adam (Kingma and Ba, 2015) over 50 epochs with a batch size of 32. We use the following hyper-parameter settings: word embedding size of 20, $\mathbf{W}_p \in \mathbb{R}^{100 \times 3}$, $\mathbf{W}_s \in \mathbb{R}^{50 \times 2}$, $\alpha = 0.5$, hidden size of Φ_{title} and Φ_{text} is 20, hidden size of GRU and GRU is 50. Dropout is applied to all GRU recurrent units on the input and output connections with a keep rate of 0.7.

Lastly, we also explore the idea of curriculum learning (Bengio et al., 2009), by fixing the CRF transition matrix $\mathbf{A} = \mathbf{0}$ for the first $e = 20$ epochs, after which we train the parameters for the remainder of the run. This allows the ME-CRF to learn a good strategy for DA and link prediction, as independent “maxent” type classifier, before attempting to learn sequence dynamics. We refer to this variant as “ME-CRF+”.

4.3 Evaluation

Following Wang (2014), we evaluate based on post-level micro-averaged F-score. All experiments were carried out with 10-fold cross validation, stratifying at the thread level.

We benchmark against the following previous work: the feature-rich CRF-based approach of Kim et al. (2010), where the authors trained independent models for each of link and DA classification (“CRF_{KIM}”); the feature-rich CRF-based approach of Wang (2014), where the author further extended the feature set of Kim et al. (2010) and jointly trained a CRF over the link and DA prediction tasks (“CRF_{WANG}”); and the dependency

Model	Link	DA	Joint
CRF _{KIM}	86.3	75.1	—
CRF _{WANG}	82.3	73.4	66.5
DEPPARSER	85.0	75.7	70.6
MEMNET	85.8	76.0	69.5
ME-CRF	86.4	77.5	70.9
ME-CRF+	86.3	77.4	71.2

Table 1: Post-level Link and DA F-scores. Performance for ME-CRF and ME-CRF+ is macro-averaged over 5 runs.

parser-based approach of Wang (2014), where the author treated the discourse structure prediction task as a constrained dependency parsing problem, with posts as nodes in the dependency graph, and the constraint that links must connect to preceding posts in the thread (“DEPPARSER”).³ In addition to the CRF/parser-based systems, we also build a MEMNET-based baseline (named MEMNET) where MEMNET shares the architecture of the memory layer in ME-CRF but excludes the use of the CRF layer. Instead, MEMNET, following the work of Sukhbaatar et al. (2015), predicts the final answer by:

$$\hat{\mathbf{y}} = \text{softmax}(\mathbf{W}_{\text{DA}}(\mathbf{u}^{K+1})) \quad (13)$$

where $\hat{\mathbf{y}}$ is the predicted DA distribution, $\mathbf{W}_{\text{DA}} \in \mathbb{R}^{|\mathcal{Y}| \times d}$ is a parameter matrix for the model to learn, and $K = 1$ is the total number of hops. This is equivalent to classifying link and DA independently at each time step t without taking transitions between DA labels into account.

4.4 Results

The experimental results are presented in Table 1, wherein the first three rows are the three baseline systems.

State-of-the-art post-level results. ME-CRFs achieve state of the art results in terms of joint post-level F-score, substantially better than the baselines. While ME-CRF slightly outperforms the current state-of-the-art (DEPPARSER), ME-CRF+ improves the performance and achieves a further 0.3% absolute gain.

³Note that a mistake was found in the results in the original paper (Wang et al., 2011), and we use the corrected results from Wang (2014).




Figure 3: Breakdown of post-level Joint F-scores by post depth, where e.g. “[1, 2]” is the joint F-score over posts of depth 1–2, i.e. the first or second post in the thread. Note that we take the reported performance of CRF_{WANG} and DEPPARSER from Wang (2014).

Curriculum learning improves joint prediction.

Despite the slight performance drop on the DA and link prediction tasks, ME-CRF+, with the CRF transition matrix frozen for the first 20 epochs, achieves a $\sim 0.3\%$ absolute gain in joint F-score over ME-CRF. This suggests that the sequence dynamics between posts, while difficult to capture, are beneficial to the overall task (resulting in more coherent DA and link predictions) if trained with proper initialisation.

MEMNET vs. ME-CRFs. We see consistent gains across all three tasks when the CRF layer is added. Although not presented in Table 1, the difference is most notable at the thread-level (i.e. a thread is correct iff all posts are tagged correctly), highlighting the importance of sequential transitional information between posts.

CRF vs. ME-CRFs. Note that CRF_{KIM} is not trained jointly on the two component tasks, but individually on each task. Without additional data, jointly training on the two tasks generates results that are comparable or substantially better over the individual tasks. This highlights the effectiveness of ME-CRF, especially with the link prediction performance comparable to that of a single-task model CRF, and surpassing it in the case of DA.




Figure 4: Breakdown of post-level Link and DA F-scores by post depth.

4.5 Analysis

We break the performance down by the depth of each post in a given thread, and present the results in Figure 3. Although below the baselines for the interval [1, 2], ME-CRF+ consistently outperforms CRF_{WANG} from depth 3 onwards, and is superior to DEPPARSER for depths [7,). Breaking down the performance further to the individual tasks of Link and DA prediction, as displayed in Figure 4, we observe a similar trend. In line with the findings in the work of Wang (2014), this confirms that prediction becomes progressively more difficult as threads grow longer, which is largely due to the increased variability in discourse structure. Despite the escalated difficulty, ME-CRF+ is substantially superior to the baselines when classifying deeper posts.

Between the CRF-based models, it is worth noting that despite the lower performance for [1, 2], ME-CRF+ benefits from having global access to the entire sequence, and consistently outperforms CRF_{WANG} for depths [3,), highlighting the effectiveness of the memory mechanism. Overall, these results validate our hypothesis that having unrestricted access to the whole sequence is beneficial, especially for long-range dependencies, offering further evidence of the power of

ME-CRFs.

5 Named Entity Recognition

In this section, we present experiments in a second setting: named entity recognition, over the CoNLL 2003 English NER shared task dataset (Tjong Kim Sang and De Meulder, 2003). Our interest here is in evaluating the ability of ME-CRF to capture document context, to aid in the identification and disambiguation of NEs.

5.1 Dataset and Task

The CoNLL 2003 NER shared task dataset consists of 14,041/3,250/3,453 sentences in the training/development/test set, resp., extracted from 946/216/231 Reuters news articles from the period 1996–97. The goal is to identify individual token occurrences of NEs, and tag each with its class (e.g. LOCATION or ORGANISATION). Here, we use the IOB tagging scheme. In terms of tagging schemes, while some have shown improvements with a more expressive IOBES marginally (Ratinov and Roth, 2009; Dai et al., 2015), we stick to the BIO scheme for simplicity and the observation of little improvement between these schemes by Lample et al. (2016).

5.2 Experimental Setup

We choose $\Phi(x_t)$ to be a lookup function, returning the corresponding embedding x_t of the word x_t . In addition to the word features, we employ a subset of the lexical features described in Huang et al. (2015), based on whether the word:

- starts with a capital letter;
- is composed of all capital letters;
- is composed of all lower case letters;
- contains non initial capital letters;
- contains both letters and digits;
- contains punctuation.

These features are all binary and referred to as $\Phi_l(x_t)$. Similar to the thread structure prediction experiments, we concatenate $\Phi_l(x_t)$ with x_t to generate the new input x' to the bi-directional GRUs in Equations (1) and (2).

In order to incorporate information in the document beyond sentence boundaries, we encode every word sequentially in a document with $\overrightarrow{\Phi}$ and $\overleftarrow{\text{GRU}}$, and store them in the memory m_i and c_i for $i \in [1, t']$, where t' is the index of the current word t in the document.

Training is carried out with Adam, over 100 epochs with a batch size of 32. We use the following hyper-parameter settings: word embedding size = 50; hidden size of $\overrightarrow{\text{GRU}}$ and $\overleftarrow{\text{GRU}}$ = 50; \overrightarrow{W}_m and $\overleftarrow{W}_m \in \mathbb{R}^{50 \times 50}$; and $b_m \in \mathbb{R}^{50}$. Dropout is applied to all GRU recurrent units on the input and output connections, with a keep rate of 0.8. We initialise ME-CRF with pre-trained word embeddings and keep them fixed during training. While we report results only on the test set, we use early stopping based on the development set.

5.3 Evaluation

Evaluation is based on span-level NE F-score, based on the official CoNLL evaluation script.⁴

We compare against the following baselines:

1. a CRF over hand-tuned lexical features (“CRF”: Huang et al. (2015))
2. an LSTM and bi-directional LSTM (“LSTM” and “BI-LSTM”, resp.: Huang et al. (2015))
3. a CRF taking features from a convolutional neural network as input (“CONV-CRF”: Collobert et al. (2011))
4. a CRF over the output of either a simple LSTM or bidirectional LSTM (“LSTM-CRF” and “BI-LSTM-CRF”, resp.: Huang et al. (2015))

Note that for our word embeddings, while we observe better performance with GLOVE (Pennington et al., 2014), for fair comparison purposes, we adopt the same SENNA embeddings (Collobert et al., 2011) as are used in the baseline methods.⁵

5.4 Results

The experimental results are presented in Table 2. Results for the baseline methods are based on the published results of Huang et al. (2015) and Collobert et al. (2011). Note that none of the systems in Table 2 use external gazetteers, to make the comparison fair. As can be observed, ME-CRF achieves the best performance, beating all the baselines.

To gain a better understanding of what the model has learned, Table 3 presents two examples

⁴<http://www.cnts.ua.ac.be/conll2000/chunking/conlleval.txt>

⁵Lample et al. (2016) report a higher result of 90.9 using a BI-LSTM-CRF architecture, but augmented with skip n -grams (Ling et al., 2015) and character embeddings. Due to the differing underlying representation, we exclude it from the comparison.

Model	F-score
CRF	86.1
LSTM	83.7
Bi-LSTM	85.2
CONV-CRF	88.7
LSTM-CRF	88.4
Bi-LSTM-CRF	88.8
ME-CRF	89.5

Table 2: NER performance on the CoNLL 2003 English NER shared task dataset.

where ME-CRF focuses on words beyond the current sentence boundaries. In the example on the left, where the target word is *Juventus* (an Italian soccer team), ME-CRF directs the attention mainly to the occurrence of the same word in a previous sentence and a small fraction to *Manchester* (a UK soccer team, in this context). Note that it does not attend to the other NE (*Europe*) in that sentence, which is of a different NE class. In the example on the right, on the other hand, ME-CRF allocates attention to the same words as the target word in the current sentence. Note that the second occurrence of *Interfax* in the memory is the same occurrence as the first word in the current sentence. While more weight is placed on the second *Interfax*, close to one third of the attention is also assigned to the first occurrence. Given that the memory, m_i and c_i , is encoded with bi-directional GRUs, the first *Interfax* should, to some degree, capture the succeeding neighbouring elements: *news agency*.

This is reminiscent of label consistency in the works of Sutton and McCallum (2004) and Finkel et al. (2005), but differs in that the consistency constraint is soft as opposed to hard in previous studies, and automatically learned without the use of heuristics.

6 Conclusion

In this paper, we have presented ME-CRF, a model extending linear-chain CRFs by including external memory. This allows the model to look beyond neighbouring items and access long-range context. Experimental results demonstrate the effectiveness of the proposed method over two tasks: forum thread discourse analysis, and named entity recognition.

Memory	$p_{t,i}$	Memory	$p_{t,i}$
...		...	
Manchester	0.23	,	0.00
United	0.00	Interfax	0.32
face	0.00	news	0.00
Juventus	0.65	agency	0.00
in	0.00	said	0.00
Europe	0.00	.	0.00
...		<u>Interfax</u>	0.68
European champions		<u>Interfax</u> quoted	
<u>Juventus</u> ...		Russian ...	

Table 3: An NER example showing learned attention to long-range contextual dependencies. The last row is the current sentence. The underlined words indicate the target word x_t , and the dashed line indicates a sentence boundary.

Acknowledgments

We thank the anonymous reviewers for their valuable feedback, and gratefully acknowledge the support of Australian Government Research Training Program Scholarship and National Computational Infrastructure (NCI Australia). This work was also supported in part by the Australian Research Council.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015)*. San Diego, USA.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning (ICML 2009)*. Montreal, Canada, pages 41–48.
- Yoshua Bengio, Patrice Simard, and Paolo Frasconi. 1994. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks* 5(2):157–166.
- Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*. Doha, Qatar, pages 103–111.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from

- scratch. *Journal of Machine Learning Research* 12:2493–2537.
- Hong-Jie Dai, Po-Ting Lai, Yung-Chun Chang, and Richard Tzong-Han Tsai. 2015. Enhancing of chemical compound and drug name recognition using representative tag scheme and fine-grained tokenization. *Journal of Cheminformatics* 7(1):S14.
- Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL 2005)*. Ann Arbor, USA, pages 363–370.
- Stuart Geman and Donald Geman. 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6:721–741.
- Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech recognition with deep recurrent neural networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2013)*. Vancouver, Canada, pages 6645–6649.
- Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, and Jürgen Schmidhuber. 2001. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*.
- Su Nam Kim, Li Wang, and Timothy Baldwin. 2010. Tagging and linking web forum posts. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning (CoNLL 2010)*. Association for Computational Linguistics, Uppsala, Sweden, pages 192–202.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the 3th International Conference on Learning Representations (ICLR 2015)*. San Diego, USA.
- Vijay Krishnan and Christopher D. Manning. 2006. An effective two-stage model for exploiting non-local dependencies in named entity recognition. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics (ACL 2006)*. Sydney, Australia, pages 1121–1128.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001)*. San Francisco, USA, pages 282–289.
- Siwei Lai, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Recurrent convolutional neural networks for text classification. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI 2015)*. Austin, USA, volume 333, pages 2267–2273.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2016)*. San Diego, USA, pages 260–270.
- Shasha Liao and Ralph Grishman. 2010. Using document level cross-event inference to improve event extraction. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics (ACL 2010)*. Uppsala, Sweden, pages 789–797.
- Wang Ling, Yulia Tsvetkov, Silvio Amir, Ramon Fernandez, Chris Dyer, Alan W Black, Isabel Trancoso, and Chu-Cheng Lin. 2015. Not all contexts are created equal: Better word representations with variable attention. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP 2015)*. Lisbon, Portugal, pages 1367–1372.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. Assessing the ability of lstms to learn syntax-sensitive dependencies. *Transactions of the Association for Computational Linguistics (TACL 2016)* 4:521–535.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*. Doha, Qatar, pages 1532–1543.
- Lev Ratinov and Dan Roth. 2009. Design challenges and misconceptions in named entity recognition. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL 2009)*. Boulder, USA, pages 147–155.
- Minjoon Seo, Sewon Min, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Query-reduction networks for question answering. In *Proceedings of the 5th International Conference on Learning Representations (ICLR 2017)*. Toulon, France.
- Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. 2015. End-to-end memory networks. In *Proceedings of Advances in Neural Information Processing Systems (NIPS 2015)*. Montréal, Canada, pages 2440–2448.
- Charles Sutton and Andrew McCallum. 2004. Collective segmentation and labeling of distant entities in information extraction. Technical report, Massachusetts Univ Amherst Dept of Computer Science.

- Duyu Tang, Bing Qin, and Ting Liu. 2016. Aspect level sentiment classification with deep memory network. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)*. Austin, USA, pages 214–224.
- Erik F Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In *Proceedings of the seventh conference on North American Chapter of the Association for Computational Linguistics (NAACL 2003)*. Edmonton, Canada, pages 142–147.
- Li Wang. 2014. *Knowledge discovery and extraction of domain-specific web data*. Ph.D. thesis, The University of Melbourne.
- Li Wang, Marco Lui, Su Nam Kim, Joakim Nivre, and Timothy Baldwin. 2011. Predicting thread discourse structure over technical web forums. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing (EMNLP 2011)*. Edinburgh, UK, pages 13–25.
- Jason Weston, Sumit Chopra, and Antoine Bordes. 2015. Memory networks. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015)*. San Diego, USA.
- Caiming Xiong, Stephen Merity, and Richard Socher. 2016. Dynamic memory networks for visual and textual question answering. In *Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)*. New York, USA, pages 2397–2406.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics — Human Language Technologies (NAACL HLT 2016)*. San Diego, USA, pages 1480–1489.
- Amy Zhang, Bryan Culbertson, and Praveen Paritosh. 2017. Characterizing online discussion using coarse discourse sequences. In *Proceedings of the 11th AAAI International Conference on Web and Social Media (ICWSM 2017)*. Palo Alto, USA, pages 357–366.