

Continuous Learning in a Hierarchical Multiscale Neural Network

Thomas Wolf, Julien Chaumond & Clement Delangue

Hugging Face Inc.

81 Prospect St.

Brooklyn, New York 11201, USA

{thomas, julien, clement}@huggingface.co

Abstract

We reformulate the problem of encoding a multi-scale representation of a sequence in a language model by casting it in a continuous learning framework. We propose a hierarchical multi-scale language model in which short time-scale dependencies are encoded in the hidden state of a lower-level recurrent neural network while longer time-scale dependencies are encoded in the dynamic of the lower-level network by having a meta-learner update the weights of the lower-level neural network in an online meta-learning fashion. We use elastic weights consolidation as a higher-level to prevent catastrophic forgetting in our continuous learning framework.

1 Introduction

Language models are a major class of natural language processing (NLP) models whose development has led to major progress in many areas like translation, speech recognition or summarization (Schwenk, 2012; Arisoy et al., 2012; Rush et al., 2015; Nallapati et al., 2016). Recently, the task of language modeling has been shown to be an adequate proxy for learning unsupervised representations of high-quality in tasks like text classification (Howard and Ruder, 2018), sentiment detection (Radford et al., 2017) or word vector learning (Peters et al., 2018).

More generally, language modeling is an example of online/sequential prediction task, in which a model tries to predict the next observation given a sequence of past observations. The development of better models for sequential prediction is believed to be beneficial for a wide range of applications like model-based planning or reinforcement learning as these models have to encode some

form of memory or causal model of the world to accurately predict a future event given past events.

One of the main issues limiting the performance of language models (LMs) is the problem of capturing long-term dependencies within a sequence.

Neural network based language models (Hochreiter and Schmidhuber, 1997; Cho et al., 2014) learn to implicitly store dependencies in a vector of hidden activities (Mikolov et al., 2010). They can be extended by attention mechanisms, memories or caches (Bahdanau et al., 2014; Tran et al., 2016; Graves et al., 2014) to capture long-range connections more explicitly. Unfortunately, the very local context is often so highly informative that LMs typically end up using their memories mostly to store short term context (Daniluk et al., 2016).

In this work, we study the possibility of combining short-term representations, stored in neural activations (hidden state), with medium-term representations encoded in a set of dynamical weights of the language model. Our work extends a series of recent experiments on networks with dynamically evolving weights (Ba et al., 2016; Ha et al., 2016; Krause et al., 2017; Moniz and Krueger, 2018) which show improvements in sequential prediction tasks. We build upon these works by formulating the task as a hierarchical online meta-learning task as detailed below.

The motivation behind this work stems from two observations.

On the one hand, there is evidence from a physiological point of view that time-coherent processes like working memory can involve differing mechanisms at differing time-scales. Biological neural activations typically have a 10 ms coherence timescale, while short-term synaptic plasticity can temporarily modulate the dynamic of the neural network itself on timescales of 100 ms to minutes. Longer time-scales (a few minutes

to several hours) see long-term learning kicks in with permanent modifications to neural excitability (Tsodyks et al., 1998; Abbott and Regehr, 2004; Barak and Tsodyks, 2007; Ba et al., 2016). Interestingly, these psychological observations are paralleled, on the computational side, by a series of recent works on recurrent networks with dynamically evolving weights that show benefits from dynamically updating the weights of a network during a sequential task (Ba et al., 2016; Ha et al., 2016; Krause et al., 2017; Moniz and Krueger, 2018).

In parallel to that, it has also been shown that temporal data with multiple time-scales dependencies can naturally be encoded in a hierarchical representation where higher-level features are changing slowly to store long time-scale dependencies and lower-level features are changing faster to encode short time-scale dependencies and local timing (Schmidhuber, 1992; El Hahi and Bengio, 1995; Koutnk et al., 2014; Chung et al., 2016).

As a consequence, we would like our model to encode information in a multi-scale hierarchical representation where

1. *short time-scale dependencies* can be encoded in fast-updated neural activations (hidden state),
2. *medium time-scale dependencies* can be encoded in the dynamic of the network by using dynamic weights updated more slowly, and
3. *a long time-scale memory* can be encoded in a static set of parameters of the model.

In the present work, we take as dynamic weights the full set of weights of a RNN language model (usually word embeddings plus recurrent, input and output weights of each recurrent layer).

2 Dynamical Language Modeling

Given a sequence of T discrete symbols $S = (w_1, w_2, \dots, w_T)$, the language modeling task consists in assigning a probability to the sequence $P(S) = p(w_1, \dots, w_T)$ which can be written, using the chain-rule, as

$$P(S | \theta) = \prod_{t=1}^T P(w_t | w_{t-1}, \dots, w_0, \theta) P(w_0 | \theta). \quad (1)$$

where θ is a set of parameters of the language model.

In the case of a neural-network-based language model, the conditional probability $P(w_t | w_{t-1}, \dots, w_0, \theta)$ is typically parametrized using an autoregressive neural network as

$$P(w_t | w_{t-1}, \dots, w_0, \theta) = f_{\theta}(w_{t-1}, \dots, w_0) \quad (2)$$

where θ are the parameters of the neural network.

In a dynamical language modeling framework, the parameters θ of the language model are not tied over the sequence S but are allowed to evolve. Thus, prior to computing the probability of a future token w_t , a set of parameters θ_t is estimated from the past parameters and tokens as $\theta_t = \operatorname{argmax}_{\theta} P(\theta | w_{t-1}, \dots, w_0, \theta_{t-1} \dots \theta_0)$ and the updated parameters θ_t are used to compute the probability of the next token w_t .

In our hierarchical neural network language model, the updated parameters θ_t are estimated by a higher level neural network g parametrized by a set of (static) parameters ϕ :

$$\theta_t = g_{\phi}(w_{t-1}, \dots, w_0, \theta_{t-1} \dots \theta_0) \quad (3)$$

2.1 Online meta-learning formulation

The function computed by the higher level network g , estimating θ_t from an history of parameters $\theta_{<t}$ and data points $w_{<t}$, can be seen as an online meta-learning task in which a high-level meta-learner network is trained to update the weights of a low-level network from the loss of the low-level network on a previous batch of data.

Such a meta-learner can be trained (Andrychowicz et al., 2016) to reduce the loss of the low-level network with the idea that it will generalize a gradient descent rule

$$\theta_t = \theta_{t-1} - \alpha_t \nabla_{\theta_{t-1}} \mathcal{L}_t \quad (4)$$

where α_t is a learning rate at time t and $\nabla_{\theta_{t-1}} \mathcal{L}_t^{LM}$ is the gradient of the loss \mathcal{L}_t^{LM} of the language model on the t -th dataset with respect to previous parameters θ_{t-1} .

Ravi and Larochelle (2016) made the observation that such a gradient descent rule bears similarities with the update rule for LSTM cell-states

$$c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t \quad (5)$$

when $c_t \rightarrow \theta_t$, $i_t \rightarrow \alpha_t$ and $\tilde{c}_t \rightarrow -\nabla_{\theta_{t-1}} \mathcal{L}_t$

We extend this analogy to the case of a multi-scale hierarchical recurrent model illustrated on figure 1 and composed of:

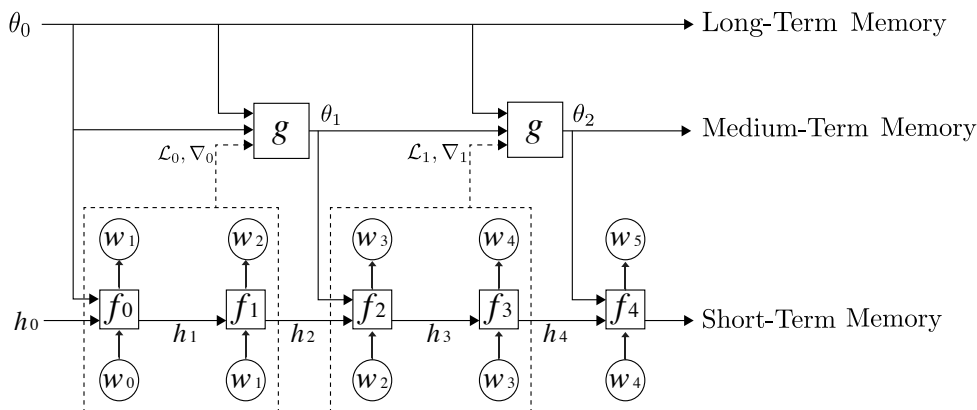


Figure 1: A diagram of the Dynamical Language Model. The lower-level neural network f (short-term memory) is a conventional word-level language model where w_0, \dots, w_5 are words tokens. The medium-level language model g is a feed-forward or recurrent neural network while the higher-level memory is formed by a static set of consolidated pre-trained weights (see text).

1. *Lower-level / short time-scale*: a RNN-based language model f encoding representations in the activations of a hidden state,
2. *Middle-level / medium time-scale*: a meta-learner g updating the set of weights of the language model to store medium-term representations, and
3. *Higher-level / long time-scale*: a static long-term memory of the dynamic of the RNN-based language model (see below).

The meta-learner g is trained to update the lower-level network f by computing $f_t, i_t, z_t = g_\phi(\theta_{t-1}, \mathcal{L}_t^{LM}, \nabla_{\theta_{t-1}} \mathcal{L}_t^{LM}, \theta_0)$ and updating the set of weights as

$$\theta_t = f_t \odot \theta_{t-1} + i_t \odot \nabla_{\theta_{t-1}} \mathcal{L}_t^{LM} + z_t \odot \theta_0 \quad (6)$$

This hierarchical network could be seen as an analog of the hierarchical recurrent neural networks (Chung et al., 2016) where the gates f_t, i_t and z_t can be seen as controlling a set of COPY, FLUSH and UPDATE operations:

1. COPY (f_t): part of the state copied from the previous state θ_{t-1} ,
2. UPDATE (i_t): part of the state updated by the loss gradients on the previous batch, and
3. FLUSH (z_t): part of the state reset from a static long term memory θ_0 .

One difference with the work of (Chung et al., 2016) is that the memory was confined to the hidden in the later while the memory of our hierarchical network is split between the weights of the lower-level network and its hidden-state.

The meta-learner can be a feed-forward or a RNN network. In our experiments, simple linear feed-forward networks lead to the lower perplexities, probably because it was easier to regularize and optimize. The meta-learner implements coordinate-sharing as described in (Andrychowicz et al., 2016; Ravi and Larochelle, 2016) and takes as input the loss \mathcal{L}_t^{LM} and loss-gradients $\nabla_{\theta_{t-1}} \mathcal{L}_t^{LM}$ over a previous batch B_i (a sequence of M tokens w_0, \dots, w_M as illustrated on figure 1). The size M of the batch adjusts the trade-off between the noise of the loss/gradients and updating frequency of the medium-term memory, smaller batches leading to faster updates with higher noise.

2.2 Continual learning

The interaction between the meta-learner and the language model implements a form of continual-learning and the language model thus faces a phenomenon known as catastrophic forgetting (French, 1999). In our case, this correspond to the lower-level network over-specializing to the lexical field of a particular topic after several updates of the meta-learner (e.g. while processing a long article on a specific topic).

To mitigate this effect we use a higher-level static memory initialized using “elastic weight consolidation” (EWC) introduced by Kirkpatrick

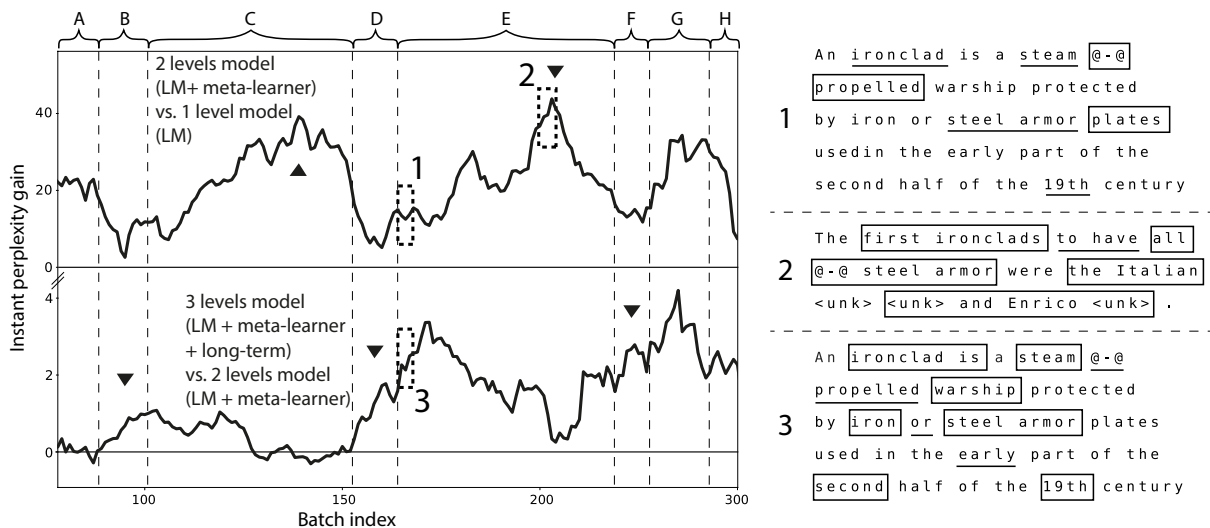


Figure 2: Medium and long-term memory effects on a sample of Wikitext-2 test set with a sequence of Wikipedia articles (letters *A* – *H*). (Left) Instantaneous perplexity gain: difference in batch perplexity between models. Higher values means the first model has locally a lower perplexity than the second model. (Top curve) Comparing a two-levels model (LM + meta-learner) with a one-level model (LM). (Bottom curve) Comparing a three-levels model (LM + meta-learner + long-term) with a two-levels model. (Right) Token loss difference on three batch samples indicated on the left curves. A squared (resp. underlined) word means the first model has a lower (resp. higher) loss on that word than the second model. We emphasize only words associated with a significant difference in loss by setting a threshold at 10 percent of the maximum absolute loss of each sample.

et al. (2017) to reduce forgetting in multi-task reinforcement learning.

Casting our task in the EWC framework, we define a task *A* which is the language modeling task (prediction of next token) when no context is stored in the weights of the lower-level network. The solution of task *A* is a set of weights toward which the model could advantageously come back when the context stored in the weights become irrelevant (for example when switching between paragraphs on different topics). To obtain a set of weights for task *A*, we train the lower-level network (RNN) alone on the training dataset and obtain a set of weights that would perform well on average, i.e. when no specific context has been provided by a context-dependent weight update performed by the meta-learner.

We then define a task *B* which is a language modeling task when a context has been stored in the weights of the lower-level network by an update of the meta-learner. The aim of EWC is to learn task *B* while retaining some performance on task *A*.

Empirical results suggest that many weights configurations result in similar performances (Sussmann, 1992) and there is thus likely a solu-

tion for task *B* close to a solution for task *A*. The idea behind EWC is to learn task *B* while protecting the performance in task *A* by constraining the parameters to stay around the solution found for task *A*.

This constraint is implemented as a quadratic penalty, similarly to spring anchoring the parameters, hence the name elastic. The stiffness of the springs should be greater for parameters that most affect performance in task *A*. We can formally write this constrain by using Bayes rule to express the conditional log probability of the parameters when the training dataset \mathcal{D} is split between the training dataset for task *A* (\mathcal{D}_A) and the training dataset for task *B* (\mathcal{D}_B):

$$\log p(\theta | \mathcal{D}) = \log p(\mathcal{D}_B | \theta) + \log p(\theta | \mathcal{D}_A) - \log p(\mathcal{D}_B) \quad (7)$$

The true posterior probability on task *A* $p(\theta | \mathcal{D}_A)$ is intractable, so we approximate the posterior as a Gaussian distribution with mean given by the parameters and a diagonal precision given by the diagonal of the Fisher information matrix F which is equivalent to the second derivative of the loss near a minimum and can be computed from first-order derivatives alone.

3 Related work

Several works have been devoted to dynamically updating the weights of neural networks during inference. A few recent architectures are the Fast-Weights of Ba et al. (2016), the Hypernetworks of Ha et al. (2016) and the Nested LSTM of Moniz and Krueger (2018). The weights update rules of these models use as inputs one or several of (i) a previous hidden state of a RNN network or higher level network and/or (ii) the current or previous inputs to the network. However, these models do not use the predictions of the network on the previous tokens (i.e. the loss and gradient of the loss of the model) as in the present work. The architecture that is most related to the present work is the study on dynamical evaluation of Krause et al. (2017) in which a loss function similar to the loss function of the present work is obtained empirically and optimized using a large hyper-parameter search on the parameters of the SGD-like rule.

4 Experiments

4.1 Architecture and hyper-parameters

As mentioned in 2.2, a set of pre-trained weights of the RNN language model is first obtained by training the lower-level network f and computing the diagonal of the Fisher matrix around the final weights.

Then, the meta-learner g is trained in an online meta-learning fashion on the validation dataset (alternatively, a sub-set of the training dataset could be used). A training sequence S is split in a sequence of mini-batches B_i , each batch B_i containing M inputs tokens ($w_{i \times M}, \dots, w_{i \times M + M}$) and M associated targets ($w_{i \times M + 1}, \dots, w_{i \times M + M + 1}$). In our experiments we varied M between 5 and 20.

The meta-learner is trained as described in (Andrychowicz et al., 2016; Li and Malik, 2016) by minimizing the sum over the sequence of LM losses: $\mathcal{L}_{meta} = \sum_{i>0} \mathcal{L}_i^{LM}$. The meta-learner is trained by truncated back-propagation through time and is unrolled over at least 40 steps as the reward from the medium-term memory is relatively sparse (Li and Malik, 2016).

To be able to unroll the model over a sufficient number of steps while using a state-of-the-art language model with over than 30 millions parameters, we use a memory-efficient version of back propagation through time based on gradi-

ent checkpointing as described by Grusly et al. (2016).

4.2 Experiments

We performed a series of experiments on the Wikitext-2 dataset (Merity et al., 2016) using an AWD-LSTM language model (Merity et al., 2017) and a feed-forward and RNN meta-learner.

The test perplexity are similar to perplexities obtained using dynamical evaluation (Krause et al., 2017), reaching 46.9 with a linear feed-forward meta-learner when starting from a one-level language model with test perplexity of 64.8.

In our experiments, the perplexity could not be improved by using a RNN meta-learner or a deeper meta-learner. We hypothesis that this may be caused by several reasons. First, storing a hidden state in the meta-learner might be less important in an online meta-learning setup than it is in a standard meta-learning setup (Andrychowicz et al., 2016) as the target distribution of the weights is non-stationary. Second, the size of the hidden state cannot be increased significantly without reducing the number of steps along which the meta-learner is unrolled during meta-training which may be detrimental.

Some quantitative experiments are shown on Figure 2 using a linear feed-forward network to illustrate the effect of the various layers in the hierarchical model. The curves shows differences in batch perplexity between model variants.

The top curve compares a one-level model (language model) with a two-levels model (language model + meta-learner). The meta-learner is able to learn medium-term representations to progressively reduce perplexity along articles (see e.g. articles C and E). Right sample 1 (resp. 2) details sentences at the begging (resp. middle) of article E related to a warship called “Ironclad”. The addition of the meta-learner reduces the loss on a number of expression related to the warship like “ironclad” or “steel armor”.

Bottom curve compares a three-levels model (language model + meta-learner + long-term memory) with the two-levels model. The local loss is reduced at topics changes and beginning of new topics (see e.g. articles B, D and F). The right sample 3 can be contrasted with sample 1 to illustrate how the hierarchical model is able to better recover a good parameter space following a change in topic.

References

- L. F. Abbott and Wade G. Regehr. 2004. [Synaptic computation](#). *Nature*, 431(7010):796–803.
- Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W. Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando de Freitas. 2016. [Learning to learn by gradient descent by gradient descent](#). *arXiv:1606.04474 [cs]*. ArXiv: 1606.04474.
- Ebru Arisoy, Tara N. Sainath, Brian Kingsbury, and Bhuvana Ramabhadran. 2012. Deep neural network language models. In *Proceedings of the NAACL-HLT 2012 Workshop: Will We Ever Really Replace the N-gram Model? On the Future of Language Modeling for HLT*, pages 20–28. Association for Computational Linguistics.
- Jimmy Ba, Geoffrey Hinton, Volodymyr Mnih, Joel Z. Leibo, and Catalin Ionescu. 2016. [Using Fast Weights to Attend to the Recent Past](#). *arXiv:1610.06258 [cs, stat]*. ArXiv: 1610.06258.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. [Neural machine translation by jointly learning to align and translate](#). *arXiv preprint arXiv:1409.0473*.
- Omri Barak and Misha Tsodyks. 2007. [Persistent activity in neural networks with dynamic synapses](#). *PLoS computational biology*, 3(2):e35.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation](#). *arXiv:1406.1078 [cs, stat]*. ArXiv: 1406.1078.
- Junyoung Chung, Sungjin Ahn, and Yoshua Bengio. 2016. [Hierarchical Multiscale Recurrent Neural Networks](#). *arXiv:1609.01704 [cs]*. ArXiv: 1609.01704.
- Micha Daniluk, Tim Rocktschel, Johannes Welbl, and Sebastian Riedel. 2016. [Frustratingly Short Attention Spans in Neural Language Modeling](#).
- Salah El Hahi and Yoshua Bengio. 1995. [Hierarchical Recurrent Neural Networks for Long-term Dependencies](#). In *Proceedings of the 8th International Conference on Neural Information Processing Systems, NIPS’95*, pages 493–499, Cambridge, MA, USA. MIT Press.
- Robert M. French. 1999. [Catastrophic forgetting in connectionist networks](#). *Trends in Cognitive Sciences*, 3(4):128–135.
- Alex Graves, Greg Wayne, and Ivo Danihelka. 2014. [Neural Turing machines](#). *arXiv preprint arXiv:1410.5401*.
- Audrnas Gruslys, Remi Munos, Ivo Danihelka, Marc Lanctot, and Alex Graves. 2016. [Memory-Efficient Backpropagation Through Time](#). *arXiv:1606.03401 [cs]*. ArXiv: 1606.03401.
- David Ha, Andrew Dai, and Quoc V. Le. 2016. [HyperNetworks](#). *arXiv:1609.09106 [cs]*. ArXiv: 1609.09106.
- Sepp Hochreiter and Jrgen Schmidhuber. 1997. [Long Short-Term Memory](#). *Neural Comput.*, 9(8):1735–1780.
- Jeremy Howard and Sebastian Ruder. 2018. [Fine-tuned Language Models for Text Classification](#). *arXiv:1801.06146 [cs, stat]*. ArXiv: 1801.06146.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2017. [Overcoming catastrophic forgetting in neural networks](#). *Proceedings of the National Academy of Sciences*, 114(13):3521–3526.
- Jan Koutnk, Klaus Greff, Faustino Gomez, and Jrgen Schmidhuber. 2014. [A Clockwork RNN](#). In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML’14*, pages II–1863–II–1871, Beijing, China. JMLR.org.
- Ben Krause, Emmanuel Kahembwe, Iain Murray, and Steve Renals. 2017. [Dynamic Evaluation of Neural Sequence Models](#). *arXiv:1709.07432 [cs]*. ArXiv: 1709.07432.
- Ke Li and Jitendra Malik. 2016. [Learning to Optimize](#). *arXiv:1606.01885 [cs, math, stat]*. ArXiv: 1606.01885.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2017. [Regularizing and Optimizing LSTM Language Models](#). *arXiv:1708.02182 [cs]*. ArXiv: 1708.02182.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. [Pointer Sentinel Mixture Models](#). *arXiv:1609.07843 [cs]*. ArXiv: 1609.07843.
- Tomas Mikolov, Martin Karafit, Lukas Burget, Jan Cernock, and Sanjeev Khudanpur. 2010. *Recurrent neural network based language model*, volume 2.
- Joel Ruben Antony Moniz and David Krueger. 2018. [Nested LSTMs](#). *arXiv:1801.10308 [cs]*. ArXiv: 1801.10308.
- Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence rnns and beyond](#). *arXiv preprint arXiv:1602.06023*.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). *arXiv:1802.05365 [cs]*. ArXiv: 1802.05365.

- Alec Radford, Rafal Jozefowicz, and Ilya Sutskever. 2017. [Learning to Generate Reviews and Discovering Sentiment](#). *arXiv:1704.01444 [cs]*. ArXiv: 1704.01444.
- Sachin Ravi and Hugo Larochelle. 2016. [Optimization as a Model for Few-Shot Learning](#).
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. *arXiv preprint arXiv:1509.00685*.
- J. Schmidhuber. 1992. [Learning Complex, Extended Sequences Using the Principle of History Compression](#). *Neural Computation*, 4(2):234–242.
- Holger Schwenk. 2012. Continuous space translation models for phrase-based statistical machine translation. *Proceedings of COLING 2012: Posters*, pages 1071–1080.
- Hctor J. Sussmann. 1992. [Uniqueness of the weights for minimal feedforward nets with a given input-output map](#). *Neural Networks*, 5(4):589–593.
- Ke Tran, Arianna Bisazza, and Christof Monz. 2016. [Recurrent Memory Networks for Language Modeling](#). *arXiv:1601.01272 [cs]*. ArXiv: 1601.01272.
- Misha Tsodyks, Klaus Pawelzik, and Henry Markram. 1998. *Neural Networks with Dynamic Synapses*.