

Approximate Fixed-Points in Recurrent Neural Networks

Zhengxiong Wang and Anton Ragni

Department of Computer Science, University of Sheffield, 211 Portobello, Sheffield S1 4DP, UK

donatello.wang@outlook.com, a.ragni@sheffield.ac.uk

Abstract

Recurrent neural networks are widely used in speech and language processing. Due to dependency on the past, standard algorithms for training these models, such as back-propagation through time (BPTT), cannot be efficiently parallelised. Furthermore, applying these models to more complex structures than sequences requires inference time approximations, which introduce inconsistency between inference and training. This paper shows that recurrent neural networks can be reformulated as fixed-points of non-linear equation systems. These fixed-points can be computed using an iterative algorithm exactly and in as many iterations as the length of any given sequence. Each iteration of this algorithm adds one additional Markovian-like order of dependencies such that upon termination all dependencies modelled by the recurrent neural networks have been incorporated. Although exact fixed-points inherit the same parallelization and inconsistency issues, this paper shows that approximate fixed-points can be computed in parallel and used consistently in training and inference including tasks such as lattice rescoring. Experimental validation is performed in two tasks, Penn Tree Bank and WikiText-2, and shows that approximate fixed-points yield competitive prediction performance to recurrent neural networks trained using the BPTT algorithm.

Index Terms: recurrent neural networks, fixed points

1. Introduction

Recurrent neural networks (RNN) are a popular choice for solving a variety of natural language and speech processing tasks including machine translation [1], language modelling [2] and acoustic modelling [3, 4, 5]. These powerful highly non-linear sequence models enable to model long-range dependencies impossible with many other model classes. Although RNNs have achieved excellent performance across many applications and tasks their practical use comes with a range of well-known issues [6]. As such, a variant of RNNs [7, 8] or a completely alternative sequence model [9] that can overcome those issues are of both practical and theoretical interest.

Unlike simpler n -gram and feed-forward neural network models used in language modelling [10], RNNs make predictions based on the complete rather than truncated history. The complete history representation used by RNNs is known to be unstable, which may cause optimisation issues and loss of information about long-term dependencies [11]. Although alternative history representations have been proposed to address these issues [7, 8], the robust and accurate modelling of long-term dependencies in these models remains a challenge. Long-term dependency modelling also causes issues [12] for efficient parallelization of training these models using algorithms such as back-propagation through time (BPTT) [13], which leads to increased development costs and time. Although computationally efficient training methodologies have been proposed [14] to replace BPTT, they do not offer easier parallelization. Furthermore, many important applications of RNNs involve rescoring

graph-like lattice structures, which, unlike linear chains or prefix trees, require inference-time approximations [15, 16].

This paper shows that the form of RNNs used in language modelling yields history representations that can be reformulated as fixed-points of a non-linear equation system. These fixed-points can be computed either exactly or approximately using an iterative algorithm with guaranteed convergence in as many steps as the lengths of underlying sequences. Each iteration in this algorithm enables to account for one additional order of dependencies until at termination the dependencies of all orders have been accounted for. Although approximate fixed points have fewer dependencies than provided by standard RNN history representations, they have a number of interesting properties. First, according to Banach theorem [17], approximate fixed points converge to exact fixed points exponentially fast which supports terminating the iterative algorithm after few iterations. Second, unlike standard RNN history representations, approximate fixed points can be parallelised. Third, given the ability to control the scope of dependencies (the number of iterations), the approximate fixed-points can be applied for lattice rescoring without making any inference time approximations.

The rest of this paper is organised as follows. Section 2 describes the form of RNNs examined in this paper. Section 3 shows how RNN history representations can be reformulated as fixed-points and presents an algorithm for computing them by iteratively refining approximate fixed-points. Experimental results comparing the RNN history representations, exact and approximate fixed points on two language modelling tasks are presented in Section 4. Finally, conclusions drawn from this work are given in Section 5.

2. Recurrent Neural Networks

There are many variants of recurrent neural networks (RNNs) [18, 7, 8, 19, 20]. This paper examines the class of RNNs commonly referred to as Elman networks [18]. An Elman network consists of an input layer, a hidden or history layer and an output layer. The hidden layer computes a history representation at time t using the following recursive process¹

$$\mathbf{h}_t = \phi(\mathbf{W}\mathbf{h}_{t-1} + \mathbf{V}\mathbf{x}_{t-1}) \quad (1)$$

where \mathbf{x}_{t-1} is a one-hot encoding or an embedding [21, 22] of a previous input word w_{t-1} , \mathbf{h}_{t-1} is a previous history representation, ϕ is an activation function, such as *tanh* and *sigmoid*, \mathbf{V} and \mathbf{W} are weight matrices associated with the input and hidden layer respectively. The history representation is used by the output layer to compute probability of predicting next input word, w_t , using *softmax* activation function

$$P(w_t = i | \mathbf{h}_t; \theta) = \frac{\exp(\mathbf{u}_i^\top \mathbf{h}_t)}{\sum_{j=1}^{|\mathcal{V}|} \exp(\mathbf{u}_j^\top \mathbf{h}_t)} \quad (2)$$

¹For simplicity all biases are omitted in the exposition.

where \mathbf{u}_i is the i -th row of output layer weight matrix \mathbf{U} , $|V|$ is the size of vocabulary and θ is the set of RNN parameters. The ‘quality’ of RNN predictions is assessed using perplexity

$$\text{PPL}(\mathcal{D}; \theta) = \exp \left(-\frac{1}{|\mathcal{D}|} \sum_{r=1}^R \log P(\mathbf{w}_{1:T_r}^{(r)}; \theta) \right) \quad (3)$$

where \mathcal{D} is a held-out set of R test word sequences and $|\mathcal{D}|$ denotes the number of words. The dependency on the complete past input sequence implied by eq. (1) makes it challenging to apply RNNs to structures other sequences. For instance, in order to apply RNN language models for rescoring lattices in speech recognition, a range of inference time approximations is commonly used [15, 16].

The RNN parameters θ can be estimated by minimising cross-entropy loss function

$$\mathcal{L}(\mathcal{D}; \theta) = -\frac{1}{R} \sum_{r=1}^R \log(P(\mathbf{w}_{1:T_r}^{(r)}; \theta)) \quad (4)$$

$$= -\frac{1}{R} \sum_{r=1}^R \sum_{t=1}^{T_r} \log(P(w_t^{(r)} | \mathbf{h}_t^{(r)}; \theta)) \quad (5)$$

where \mathcal{D} is a training set of R word sequences. For optimisation it is common to use stochastic gradient-based schemes [23]. In such cases \mathcal{D} is a sample (mini-batch) drawn from the training set. The gradient of the cross-entropy loss function with respect to the RNN parameters is

$$\frac{\partial \mathcal{L}(\mathcal{D}; \theta)}{\partial \theta} \propto \sum_{r=1}^R \sum_{t=1}^{T_r} \sum_{k=1}^t \frac{\partial \log(P(w_t^{(r)} | \mathbf{h}_t^{(r)}; \theta))}{\partial \mathbf{h}_t^{(r)}} \frac{\partial \mathbf{h}_t^{(r)}}{\partial \mathbf{h}_k^{(r)}} \frac{\partial \mathbf{h}_k^{(r)}}{\partial \theta} \quad (6)$$

The partial derivatives on the right hand side can be efficiently computed using the back-propagation through time (BPTT) algorithm [13]. The partial derivative in the middle, $\partial \mathbf{h}_t^{(r)} / \partial \mathbf{h}_k^{(r)}$, represents the dependencies between time steps t and k . Unfortunately, this partial derivative is known to exhibit stability issues as the difference between time steps increases [11]. Since BPTT development in the 1980s [24, 25, 26], it has inspired several advancements aimed at improving temporal modelling using neural networks and speeding up gradient computation [6, 7, 8, 9, 27, 28, 29]. However, limited work has been done on improving parallelization.

3. Fixed-Point Representation

Consider applying the Elman style RNN in eq. (1) to a sequence of length 3. Starting from the initial history vector \mathbf{h}_0 set to some value, the history states at each subsequent time can be computed as follows

$$\begin{aligned} \mathbf{h}_0 &= \mathbf{h}_0 \\ \mathbf{h}_1 &= \phi(\mathbf{W}\mathbf{h}_0 + \mathbf{V}\mathbf{x}_0) \\ \mathbf{h}_2 &= \phi(\mathbf{W}\mathbf{h}_1 + \mathbf{V}\mathbf{x}_1) \\ \mathbf{h}_3 &= \phi(\mathbf{W}\mathbf{h}_2 + \mathbf{V}\mathbf{x}_2) \end{aligned} \quad (7)$$

Noting that the history states $\mathbf{h}_0, \mathbf{h}_1, \mathbf{h}_2, \mathbf{h}_3$ appear on the left hand side *and* on the right hand side (ignoring for the moment the lack of \mathbf{h}_3 on the right hand side and the misalignment of indices) suggests that these equations can be written as a fixed-point equation $\vec{\mathbf{h}} = \mathbf{f}(\vec{\mathbf{h}}, \vec{\mathbf{x}})$. Indeed, it is possible to show that the above set of equations can be written in the following fixed-point form

$$\vec{\mathbf{h}} = \vec{\phi}(\vec{\mathbf{W}}\vec{\mathbf{h}} + \vec{\mathbf{V}}\vec{\mathbf{x}}) \quad (8)$$

where

$$\vec{\mathbf{x}} = [\mathbf{0}^\top \mathbf{x}_0^\top \dots \mathbf{x}_{T-1}^\top]^\top \quad (9)$$

$$\vec{\mathbf{h}} = [\mathbf{h}_0^\top \dots \mathbf{h}_T^\top]^\top \quad (10)$$

$$\vec{\phi}(\vec{\mathbf{z}}) = [\mathbf{z}_0^\top \phi(\mathbf{z}_1)^\top \dots \phi(\mathbf{z}_T)^\top]^\top \quad (11)$$

and

$$\vec{\mathbf{W}} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{W} & \mathbf{0} & \dots & \mathbf{0} \\ \vdots & \ddots & \ddots & \vdots \\ \mathbf{0} & \dots & \mathbf{W} & \mathbf{0} \end{bmatrix}, \vec{\mathbf{V}} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{V} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{V} \end{bmatrix} \quad (12)$$

For a sequence with three input words ($T = 3$) the fixed-point form in eq. (8) can be expressed as

$$\begin{bmatrix} \mathbf{h}_0 \\ \mathbf{h}_1 \\ \mathbf{h}_2 \\ \mathbf{h}_3 \end{bmatrix} = \vec{\phi} \left(\begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{W} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{W} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{W} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{h}_0 \\ \mathbf{h}_1 \\ \mathbf{h}_2 \\ \mathbf{h}_3 \end{bmatrix} + \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{V} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{V} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{V} \end{bmatrix} \begin{bmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} \right) \quad (13)$$

It is easy to verify that this equation simplifies to eq. (7). Thus, the sequence of RNN history states for any sequence is the fixed-point of the non-linear equation system in eq. (8). Non-linear equation systems appear in a range of machine learning approaches [17, 30, 31]. However, it is rarely possible to obtain exact fixed-points of those equations and closed-forms in particular. Instead, approximate fixed-points are commonly used.

3.1. Fixed-point iteration (FPI) algorithm

The Banach theorem [17] states that regardless of the initial starting point $\vec{\mathbf{h}}^{(0)}$ the following iterative process $\vec{\mathbf{h}}^{(n+1)} = \mathbf{f}(\vec{\mathbf{h}}^{(n)}, \vec{\mathbf{x}})$ will converge exponentially fast to the fixed-point of $\vec{\mathbf{h}} = \mathbf{f}(\vec{\mathbf{h}}, \vec{\mathbf{x}})$. For the fixed-point form in eq. (8) the corresponding iterative process can be written as

$$\vec{\mathbf{h}}^{(n+1)} = \vec{\phi}(\vec{\mathbf{W}}\vec{\mathbf{h}}^{(n)} + \vec{\mathbf{V}}\vec{\mathbf{x}}) \quad (14)$$

One interesting aspect of this process is that it enables *parallel* computation of all history states. Another interesting aspect is that it will converge to $\vec{\mathbf{h}} = [\mathbf{h}_0^\top \dots \mathbf{h}_T^\top]^\top$ (RNN history states) exactly in T steps. To illustrate this, consider the same example with the sequences of length 3 above. Starting with $\vec{\mathbf{h}}^{(0)} = [\mathbf{h}_0^{(0)} \mathbf{h}_1^{(0)} \mathbf{h}_2^{(0)} \mathbf{h}_3^{(0)}]^\top$, where $\mathbf{h}_0^{(0)}$ is set to the initial RNN state \mathbf{h}_0 and others set to arbitrary values, the first update based on equation (14) yields

$$\vec{\mathbf{h}}^{(1)} = \begin{bmatrix} \mathbf{h}_0^{(0)} \\ \phi(\mathbf{W}\mathbf{h}_0^{(0)} + \mathbf{V}\mathbf{x}_0) \\ \phi(\mathbf{W}\mathbf{h}_1^{(0)} + \mathbf{V}\mathbf{x}_1) \\ \phi(\mathbf{W}\mathbf{h}_2^{(0)} + \mathbf{V}\mathbf{x}_2) \end{bmatrix} \quad (15)$$

The second update will yield

$$\vec{\mathbf{h}}^{(2)} = \begin{bmatrix} \mathbf{h}_0^{(0)} \\ \phi(\mathbf{W}\mathbf{h}_0^{(0)} + \mathbf{V}\mathbf{x}_0) \\ \phi(\mathbf{W}\phi(\mathbf{W}\mathbf{h}_0^{(0)} + \mathbf{V}\mathbf{x}_0) + \mathbf{V}\mathbf{x}_1) \\ \phi(\mathbf{W}\phi(\mathbf{W}\mathbf{h}_1^{(0)} + \mathbf{V}\mathbf{x}_1) + \mathbf{V}\mathbf{x}_2) \end{bmatrix} \quad (16)$$

Finally, the third update will yield

$$\vec{h}^{(3)} = \begin{bmatrix} \phi(\mathbf{W} \underbrace{h_0^{(0)}}_{h_0} + \mathbf{V} \mathbf{x}_0) \\ \phi(\mathbf{W} \underbrace{\phi(\mathbf{W} h_0^{(0)} + \mathbf{V} \mathbf{x}_0)}_{h_1} + \mathbf{V} \mathbf{x}_1) \\ \phi(\mathbf{W} \underbrace{\phi(\mathbf{W} \phi(\mathbf{W} h_0^{(0)} + \mathbf{V} \mathbf{x}_0) + \mathbf{V} \mathbf{x}_1)}_{h_2} + \mathbf{V} \mathbf{x}_2) \end{bmatrix} \equiv \begin{bmatrix} h_0 \\ h_1 \\ h_2 \\ h_3 \end{bmatrix} \quad (17)$$

which is equivalent to the RNN history states. Note that all initial history states other than $h_0^{(0)}$ have been eliminated. Although the parallel forward pass with the fixed-point representation yields identical to the RNN history states values, the parallel backward pass would lead to some dependencies counted multiple times. For example, as shown in eq. (17) the first word \mathbf{x}_0 will make 3 contributions, the second word \mathbf{x}_1 will make 2 contributions and the final word \mathbf{x}_2 will make 1 contribution to the gradient. Thus, the information coming further from the past would receive more boosting than the more recent information.

3.2. Approximate fixed points (AFP)

The Banach theorem [17] states that the intermediate points $\vec{h}^{(n)}$ converge to the fixed-point \vec{h} exponentially fast, which lends support for approximating fixed-points by terminating the iterative process after a small number of iterations ρ . The section below illustrates a pseudo-code for updating the RNN parameters in the e -th training epoch (mini-batch size is 1 for simplicity), where $\vec{y} = [y_0^\top y_1^\top \dots y_T^\top]^\top$ is a sequence of output probability distributions, \vec{U} is a block-diagonal matrix with diagonal elements set to \mathbf{U} and $\vec{\sigma}$ is a block-softmax function. The simple example in eqs. (15) to (17) shows that each iter-

Algorithm 1: Approximate fixed-point training

```

1 for  $r$  in  $[1, R]$  do
2   for  $n$  in  $[1, \rho]$  do
3      $\vec{h}^{(r,n)} = \vec{\phi}(\vec{W} \vec{h}^{(r,n-1)} + \vec{V} \vec{x}^{(r)})$ 
4     // optionally skip propagation of dependencies
5      $\vec{h}^{(r,n)} = \vec{h}^{(r,n)}$   $\Big|_{\theta = \theta^{(e)}}$ 
6    $\vec{y}^{(r)} = \vec{\sigma}(\vec{U} \mathbf{h}^{(r,\rho)})$ 
7   Accumulate loss
8 Update parameters

```

ation adds dependencies on one more past word. Thus, the iteration number ρ is akin to the order of Markov assumption. The limited nature of dependencies possible with the approximate fixed-points make them suitable for lattice rescoring. By expanding lattices to order $\rho + 1$ enables fixed-points to be applied to all order $\rho + 1$ arcs in parallel. In contrast to RNNs, which require inference time approximations, the approximate fixed-points are consistent in both training and inference.

Although approximate fixed points (AFP) can be computed in parallel, as the example in eqs. (15) to (17) shows the overall amount of computation performed will be larger than in the standard RNN case. Table 1 compares complexities in RNN and FPI based forward passes, where T is the length of input sequence, H is the size of the history state, C is the number of computations (addition, multiplications, etc.) in the history

Table 1: Comparison of complexities in recurrent neural network and approximate fixed point based forward passes

Optimisation	Time complexity	Space complexity	Computational complexity
BPTT	$\mathcal{O}(T)$	$\mathcal{O}(HT)$	$\mathcal{O}(CT)$
FPI	$\mathcal{O}(\rho)$	$\mathcal{O}(2HT)$	$\mathcal{O}(CT\rho)$

state and ρ is the number of iterations. As mentioned above the overall amount of computation in the backward pass is also larger in the RNN case. To reduce the amount of computation in the backward pass it is possible to eliminate the propagation of dependencies from the past completely as shown in the line 5 of the Algorithm 1 above. Note that the lack of dependency propagation does not affect the computed history state values but will cause differences in the computed gradients.

4. Experiments

The experiments conducted in this work focused on assessing approximate fixed-points (AFP) in two language modelling tasks. The following Section 4.1 provides details about the configuration of baselines and AFPs. Sections 4.2 and 4.3 then report on their performance in each task respectively.

4.1. Experimental setup

The PennTree Bank (PTB) [32] and the WikiText-2 [33] used in this work are two relatively small scale datasets. Table 2 provides basic statistics including the percentage of out-of-vocabulary (OOV) words mapped to the special ‘word’ <unk>.

Table 2: Basic statistics of PTB and wikitext2 datasets

Dataset	Training words		OOV (%)
	Unique	Total	(train/dev/test)
PTB	10K	888K	5.1 / 5.0 / 6.1
WikiText-2	33K	2M	2.7 / 5.5 / 6.3

PTB: This dataset has been often used for conducting experiments in language modelling [2]. The original dataset consists of sentences extracted from the Wall Street Journal corpus [34]. A pre-processed version, which provides 888 thousand training words and 10 thousand word vocabulary, is used in this work. The development and test subsets contain 70 and 79 thousand words respectively. Compared with modern datasets, the vocabulary size in this dataset is small.

WikiText-2: This dataset consists of articles extracted from Wikipedia. A pre-processed version available as a part of PyTorch toolkit [35] is used in this work. Compared to the PTB, the training set and vocabulary sizes in this dataset are approximately 2 and 3 times larger respectively. Note that this dataset contains many foreign (e.g. Japanese) words which contribute to the increased vocabulary size.

A number of simple (feed-forward NN (FFNN)) and more complex (RNN, Long Short-Term Memory (LSTM)) baselines have been investigated. All NN LMs were implemented in PyTorch. The minibatch size used in stochastic optimisation was 20. The optimisation was performed using Adam [36] for either 20 (PTB) or 40 (WikiText-2) epochs. The FFNN baseline is a bigram LM, which lacks information about all previous words other than the most recent one when making predictions. The

initial history states of all recurrent LMs (including AFP) were set to zeros. Unless stated otherwise, the size of history states is 100. All other parameters were initialised randomly using the uniform distribution $\mathcal{U}(-0.05, 0.05)$. AFPs were configured identical to the RNN baselines. The estimation of AFPs benefited from the sparsity of \vec{W} , \vec{V} and \vec{U} matrices.

4.2. Penn Tree Bank

The first experiment examined learning approximate fixed points (AFP) using the fixed point iteration (FPI) algorithm in the PTB task. As discussed in Section 3.2, the FPI algorithm can be configured to either propagate dependencies or not ($\text{FPI}|_{\theta=\theta(\epsilon)}$) during the training. Figure 1 compares perplexities (PPL) of these two configurations against the number of iterations ρ used by the FPI algorithm. Note that in this task




Figure 1: PPL performance in the PTB task against the AFP iteration number ρ

the average number of words per sentence is 20. As Figure 1 shows, just two iterations of both FPI algorithms bring AFP performance close to that obtained with the BPTT-trained RNN that has access to the complete past information. Additional FPI iterations yield more gains for AFP. As expected, the removal of dependencies in $\text{FPI}|_{\theta=\theta(\epsilon)}$ has a negative effect on the final performance. Table 3 puts these results in the context of baselines other than the BPTT-trained RNN. The FFNN₂ (bigram)

Table 3: Summary of perplexities in the PTB task

LM	Optimisation	PPL
5-gram [37]	Kneser-Ney	141.2
FFNN ₂	BP	179.0
RNN [37]	TBPTT	142.1
RNN	BPTT	142.1
AFP	FPI	129.4
	FPI $ _{\theta=\theta(\epsilon)}$	136.4

baseline provides an interesting contrast. Similar to $\text{FPI}|_{\theta=\theta(\epsilon)}$ it does not extract statistics from words more than 1 time step in the past. However, the former benefits from the information accumulated within history states.

4.3. WikiText-2

Many high-performance recipes for WikiText-2 task involve the use of large, multi-layer, recurrent models. Due to a limited computational resource available, a simpler PTB-like configuration was created. Table 4 shows the impact of simplifications applied to the initial LSTM model with 2 hidden layers and 1024 units per layer trained using the truncated BPTT

[14] and dropout [38]. Compared to the more advanced base-

Table 4: Summary of perplexities in the WikiText-2 task

LM	Layers	Units	Optimisation	PPL
LSTM	2	1024	TBPTT+dropout	99.3
		200	TBPTT	131.18
	1	TBPTT	132.32	
RNN	1	200	TBPTT	168.31
BPTT			153.48	
FPI			149.35	
AFP			FPI $ _{\theta=\theta(\epsilon)}$	172.05

line, the final RNN model with single hidden layer and 200 units as expected shows a significantly worse performance yet is computationally more affordable. Figure 2 repeats the investigation conducted in the PTB task (see Figure 2). Note




Figure 2: PPL performance in the WikiText-2 task against the AFP iteration number ρ

that in this task the average number of words per sentence is 55, which is more than 2.5 times longer than in the PTB task. The AFP requires more iterations to match the performance of BPTT-trained RNNs, which suggests that this task requires significantly longer context for making accurate predictions. The latter would also provide an explanation why the simpler $\text{FPI}|_{\theta=\theta(\epsilon)}$ variant performs worse in this task. Overall these results illustrate that AFPs can offer competitive performance to BPTT-trained RNNs.

5. Conclusion and Future Work

This paper proposed a novel view on history representations obtained from recurrent neural networks (RNN) as fixed-points of non-linear equation systems. The novel view gives rise to a fixed-point iteration (FPI) algorithm that enables approximate but accurate history representations to be obtained in constant rather than linear with respect to sequence length time. The approximate fixed-points (AFP) enable efficient training (parallelization) and provide an opportunity for consistent inference on complex structures, such as lattices, which is impractical with RNNs trained using standard algorithms such as back-propagation through time (BPTT). Experimental validation was performed in two language modelling tasks where AFPs achieved competitive performance to BPTT-trained RNNs given a small number of FPI iterations. The future work will examine applying AFPs to lattice rescoring and will explore more advanced AFP optimisation.

6. References

- [1] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *NIPS*, 2014.
- [2] T. Mikolov, M. Karafiát, L. Burget, J. Černocký, and S. Khudanpur, "Recurrent neural network based language model," in *Interspeech*, 2010.
- [3] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv:1211.3711*, 2012.
- [4] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *ICASSP*, 2016.
- [5] V. Peddinti, Y. Wang, D. Povey, and S. Khudanpur, "Low latency acoustic modeling using temporal convolution and LSTMs," *IEEE Signal Processing Letters*, vol. 25, no. 3, pp. 373–377, 2017.
- [6] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *ICML*, 2013.
- [7] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, 1997.
- [8] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," *arXiv preprint arXiv:1406.1078*, 2014.
- [9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017.
- [10] H. Schwenk, "Continuous space translation models for phrase-based statistical machine translation," in *COLING*, 2012, pp. 1071–1080.
- [11] Y. Bengio, P. Simard, and P. Frasconi, "Learning long-term dependencies with gradient descent is difficult," *IEEE Transactions on Neural Networks*, vol. 5, no. 2, pp. 157–166, 1994.
- [12] X. Chen, Y. Wang, X. Liu, M. J. Gales, and P. C. Woodland, "Efficient GPU-based training of recurrent neural network language models using spliced sentence bunch," in *Interspeech*, 2014.
- [13] P. J. Werbos, "Backpropagation through time: what it does and how to do it," *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550–1560, 1990.
- [14] R. J. Williams and J. Peng, "An efficient gradient-based algorithm for on-line training of recurrent network trajectories," *Neural computation*, vol. 2, no. 4, pp. 490–501, 1990.
- [15] X. Liu, Y. Wang, X. Chen, M. J. F. Gales, and P. C. Woodland, "Efficient lattice rescoring using recurrent neural network language models," in *ICASSP*, 2014, pp. 4908–4912.
- [16] H. Xu, T. Chen, D. Gao, Y. Wang, K. Li, N. Goel, Y. Carmiel, D. Povey, and S. Khudanpur, "A pruned RNNLM lattice-rescoring algorithm for automatic speech recognition," in *ICASSP*, 2018.
- [17] J. Zhou, G. Cui, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun, "Graph neural networks: A review of methods and applications," *arXiv preprint arXiv:1812.08434*, 2018.
- [18] J. L. Elman, "Finding structure in time," *Cognitive science*, 1990.
- [19] M. I. Jordan, "Serial order: A parallel distributed processing approach," in *Advances in psychology*. Elsevier, 1997, vol. 121, pp. 471–495.
- [20] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Processing*, 1997.
- [21] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *EMNLP*, 2014, pp. 1532–1543.
- [22] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [23] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in *COMPSTAT*. Springer, 2010, pp. 177–186.
- [24] P. J. Werbos, "Backpropagation through time: what it does and how to do it," *IEEE*, 1990.
- [25] D. E. Rumelhart, G. E. H. GE, and R. J. Williams, "Learning internal representations by error propagation," California University San Diego, Tech. Rep., 1985.
- [26] A. J. Robinson and F. Fallside, "The utility driven dynamic error propagation network," Cambridge University, Tech. Rep., 1987.
- [27] R. J. Williams and J. Peng, "An efficient gradient-based algorithm for on-line training of recurrent network trajectories," *Neural Computation*, 1990.
- [28] R. J. Williams and D. Zipser, "A learning algorithm for continually running fully recurrent neural networks," *Neural Computation*, 1989.
- [29] A. Graves, G. Wayne, and I. Danihelka, "Neural Turing machines," *arXiv:1410.5401*, 2014.
- [30] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, "Neural ordinary differential equations," *arXiv preprint arXiv:1806.07366*, 2018.
- [31] S. Bai, J. Z. Kolter, and V. Koltun, "Deep equilibrium models," in *NIPS*, 2019.
- [32] M. P. Marcus, B. Santorini, and M. A. Marcinkiewicz, "Building a large annotated corpus of English: The Penn Treebank," *Computational Linguistics*, 1993.
- [33] S. Merity, C. Xiong, J. Bradbury, and R. Socher, "Pointer sentinel mixture models," *arXiv preprint arXiv:1609.07843*, 2016.
- [34] D. B. Paul and J. Baker, "The design for the Wall Street Journal-based CSR corpus," in *Speech and Natural Language*, Harriman, New York, 1992.
- [35] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic Differentiation in PyTorch," in *NIPS Workshop on Autodiff*, 2017.
- [36] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [37] T. Mikolov and G. Zweig, "Context dependent recurrent neural network language model," in *2012 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2012.
- [38] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The Journal of machine learning research*, 2014.