

MODIFYING MEMORIES IN A RECURRENT NEURAL NETWORK UNIT

Anonymous authors

Paper under double-blind review

ABSTRACT

Long Short-Term Memory (LSTM) units have the ability to memorise and use long-term dependencies between inputs to generate predictions on time series data. We introduce the concept of modifying the cell state (memory) of LSTMs using rotation matrices parametrised by a new set of trainable weights. This addition shows significant increases of performance on some of the tasks from the bAbI dataset.

1 INTRODUCTION

In the recent years, Recurrent Neural Networks (RNNs) have been successfully used to tackle problems with data that can be represented in the shape of time series. Application domains include Natural Language Processing (NLP) (translation (Rosca & Breuel, 2016), summarisation (Nallapati et al., 2016), question answering and more), speech recognition (Hannun et al., 2014; Graves et al., 2013), text to speech systems (Arik et al., 2017), computer vision tasks (Stewart et al., 2016; Wu et al., 2017), and differentiable programming language interpreters (Riedel et al., 2016; Rocktäschel & Riedel, 2017).

An intuitive explanation for the success of RNNs in fields such as natural language understanding is that they allow words at the beginning of a sentence or paragraph to be memorised. This can be crucial to understanding the semantic content. Thus in the phrase "*The cat ate the fish*" it is important to memorise the subject (cat). However, often later words can change the meaning of a sentence in subtle ways. For example, "*The cat ate the fish, **didn't it***" changes a simple statement into a question. In this paper, we study a mechanism to enhance a standard RNN to enable it to modify its memory, with the hope that this will allow it to capture in the memory cells sequence information using a shorter and more robust representation.

One of the most used RNN units is the Long Short-Term Memory (LSTM) (Hochreiter & Schmidhuber, 1997). The core of the LSTM is that each unit has a *cell state* that is modified in a gated fashion at every time step. At a high level, the cell state has the role of providing the neural network with *memory* to hold long-term relationships between inputs. There are many small variations of LSTM units in the literature and most of them yield similar performance (Greff et al., 2017).

The memory (cell state) is expected to encode information necessary to make the next prediction. Currently the ability of the LSTMs to rotate and swap memory positions is limited to what can be achieved using the available gates. In this work we introduce a new operation on the memory that explicitly enables rotations and swaps of pairwise memory elements. Our preliminary tests show performance improvements on some of the bAbI tasks (Weston et al., 2015) compared with LSTM based architectures.

2 THE ROTATION GATE

In this section we introduce the idea of adding a new set of parameters for the RNN cell that enable rotation of the cell state. The following subsection shows how this is implemented in the LSTM unit.

One of the key innovations of LSTMs was the introduction of gated modified states so that if the gate neuron i is saturated then the memory $c_i(t - 1)$ would be unaltered. That is, $c_i(t - 1) \approx c_i(t)$

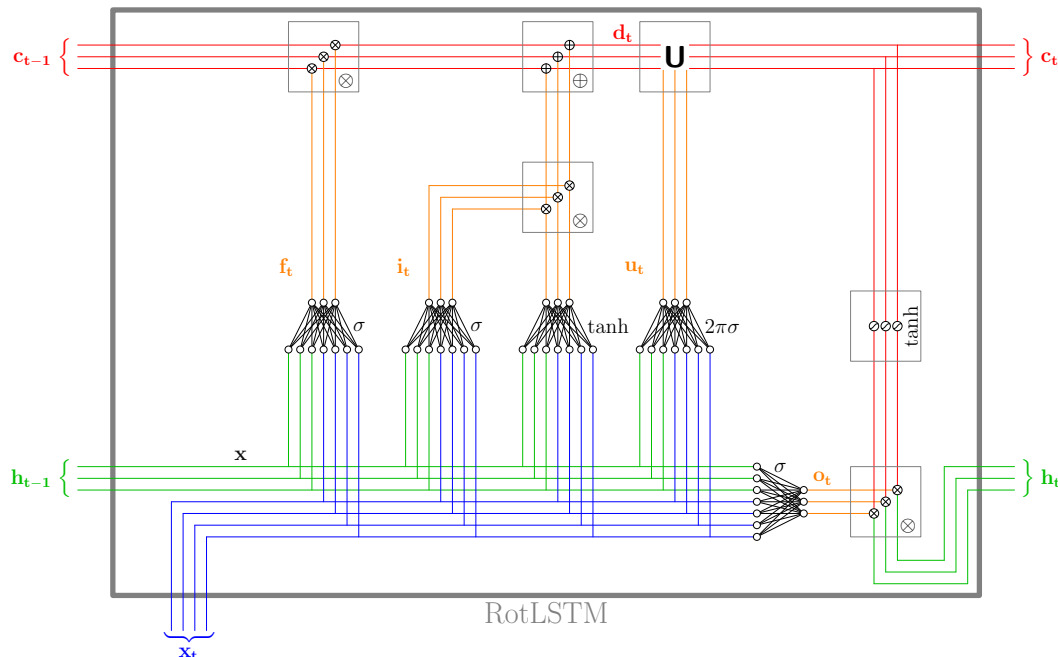


Figure 1: RotLSTM diagram. \mathbf{x} is the concatenation of \mathbf{h}_{t-1} and \mathbf{x}_t in the diagram (green and blue lines). Note that this differs from a regular LSTM by the introduction of the network producing angles \mathbf{u}_t and the rotation module marked \mathbf{U} . In the diagram input size is 4 and cell state size is 3.

John is in the kitchen. Where is John? A: Kitchen) and going to more difficult tasks of reasoning about size (example: *The football fits in the suitcase. The box is smaller than the football. Will the box fit in the suitcase? A: yes*) and path finding (example: *The bathroom is south of the office. The bathroom is north of the hallway. How do you go from the hallway to the office? A: north, north*). A summary of all tasks is available in Table 2. We are interested in evaluating the behaviour and performance of rotations on RNN units rather than beating state of the art.

We compare a model based on RotLSTM with the same model based on the traditional LSTM. All models are trained with the same hyperparameters and we do not perform any hyperparameter tuning apart from using the sensible defaults provided in the Keras library and example code (Chollet et al., 2015).

For the first experiment we train a LSTM and RotLSTM based model 10 times using a fixed cell state size of 50. In the second experiment we train the same models but vary the cell state size from 6 to 50 to assess whether the rotations help our models achieve good performance with smaller state sizes. We only choose even numbers for the cell state size to make all units go through rotations.

3.1 THE MODEL ARCHITECTURE

The model architecture, illustrated in Figure 2, is based on the Keras example implementation¹. This model architecture, empirically, shows better performance than the LSTM baseline published in Weston et al. (2015). The input question and sentences are passed through a word embedding layer (not shared, embeddings are different for questions and sentences). The question is fed into an RNN which produces a representation of the question. This representation is concatenated to every word vector from the story, which is then used as input to the second RNN. Intuitively, this helps the second RNN (Query) to *focus* on the important words to answer the question. The output of the second RNN is passed to a fully connected layer with a softmax activation of the size of the dictionary. The answer is the word with the highest activation.

¹Available at <https://goo.gl/9wfzr5>.




Figure 2: Model architecture for the bAbI dataset. The RNN is either LSTM or RotLSTM.

Table 1: Performance comparison on the bAbI dataset. Values are % average accuracy on test set \pm standard deviation taken from training each model 10 times. For each of the trained models the test accuracy is taken for the epoch with the best validation accuracy (picked from epoch numbers 1, 11, 21, 31, 40 since we only evaluated on the test set for those).

Task:	1	2	3	4	5	6	7	8	9	10
LSTM	49.8 \pm 3.3	26.5 \pm 3.8	21.1 \pm 1.0	63.7 \pm 9.4	33.6 \pm 6.2	49.2 \pm 0.8	62.7 \pm 15.5	68.8 \pm 8.3	63.9 \pm 0.2	45.2 \pm 2.0
RotLSTM	52.3 \pm 1.1	27.1 \pm 1.3	22.4 \pm 1.3	65.8 \pm 3.6	55.7 \pm 5.5	50.1 \pm 1.9	76.7 \pm 3.0	66.1 \pm 6.2	61.5 \pm 1.9	48.1 \pm 1.1
Task:	11	12	13	14	15	16	17	18	19	20
LSTM	73.6 \pm 2.2	74.3 \pm 1.7	94.4 \pm 0.2	20.7 \pm 3.3	21.4 \pm 0.5	48.0 \pm 2.2	48.0 \pm 0.0	70.3 \pm 20.8	8.5 \pm 0.6	87.7 \pm 4.2
RotLSTM	73.9 \pm 1.2	76.5 \pm 1.1	94.4 \pm 0.0	19.9 \pm 2.0	28.8 \pm 8.8	46.7 \pm 2.0	54.2 \pm 3.1	90.5 \pm 0.9	8.9 \pm 1.1	89.9 \pm 2.4

The categorical cross-entropy loss function was used for training. All dropout layers are dropping 30% of the nodes. The train-validation dataset split used was 95%-5%. The optimizer used was Adam with learning rate 0.001, no decay, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$. The training set was randomly shuffled before every epoch. All models were trained for 40 epochs. After every epoch the model performance was evaluated on the validation and training sets, and every 10 epochs on the test set. We set the random seeds to the same number for reproducibility and ran the experiments 10 times with 10 different random seeds. The source code is available at <https://goo.gl/Eopz2C>².

3.2 RESULTS

In this subsection we compare the the performance of models based on the LSTM and RotLSTM units on the bAbI dataset.

Applying rotations on the unit memory of the LSTM cell gives a slight improvement in performance overall, and significant improvements on specific tasks. Results are shown in Table 1. The most significant improvements are faster convergence, as shown in Figure 3, and requiring smaller state sizes, illustrated in Figure 4.

On tasks 1 (basic factoid), 11 (basic coreference), 12 (conjunction) and 13 (compound coreference) the RotLSTM model reaches top performance a couple of epochs before the LSTM model consistently. The RotLSTM model also needs a smaller cell state size, reaching top performance at state size 10 to 20 where the LSTM needs 20 to 30. The top performance is, however, similar for both models, with RotLSTM improving the accuracy with up to 2.5%.

The effect is observed on task 18 (reasoning about size) at a greater magnitude where the RotLSTM reaches top performance before epoch 20, after which it plateaus, while the LSTM model takes 40 epochs to fit the data. The training is more stable for RotLSTM and the final accuracy is improved by 20%. The RotLSTM reaches top performance using cell state 10 and the LSTM needs size 40. Similar performance increase for the RotLSTM (22.1%) is observed in task 5 (three argument relations), reaching top performance around epoch 25 and using a cell state of 50. Task 7 (counting) shows a similar behaviour with an accuracy increase of 14% for RotLSTM.

Tasks 4 (two argument relations) and 20 (agent motivation) show quicker learning (better performance in the early epochs) for the RotLSTM model but both models reach their top performance

²A GitHub repository will be made available after the double-blind review period ends.

after the same amount of training. On task 20 the RotLSTM performance reaches top accuracy using state size 10 while the LSTM incrementally improves until using state size 40 to 50.

Signs of overfitting for the RotLSTM model can be observed more prominently than for the LSTM model on tasks 15 (basic deduction) and 17 (positional reasoning).

Our models, both LSTM and RotLSTM, perform poorly on tasks 2 and 3 (factoid questions with 2 and 3 supporting facts, respectively) and 14 (time manipulation). These problem classes are solved very well using models that look over the input data many times and use an attention mechanism that allows the model to focus on the relevant input sentences to answer a question (Sukhbaatar et al., 2015; Kumar et al., 2016). Our models only look at the input data once and we do not filter out irrelevant information.

4 DISCUSSION AND FUTURE WORK

A limitation of the models in our experiments is only applying pairwise 2D rotations. Representations of past input can be larger groups of the cell state vector, thus 2D rotations might not fully exploit the benefits of transformations. In the future we hope to explore rotating groups of elements and multi-dimensional rotations. Rotating groups of elements of the cell state could potentially also force the models to learn a more structured representation of the world, similar to how forcing a model to learn specific representations of scenes, as presented in Higgins et al. (2017), yields semantic representations of the scene.

Rotations also need not be fully flexible. Introducing hard constraints on the rotations and what groups of parameters can be rotated might lead the model to learn richer memory representations. Future work could explore how adding such constraints impacts learning times and final performance on different datasets, but also look at what constraints can qualitatively improve the representation of long-term dependencies.

In this work we presented preliminary tests for adding rotations to simple models but we only used a toy dataset. The bAbI dataset has certain advantages such as being small thus easy to train many models on a single machine, not having noise as it is generated from a simulation, and having a wide range of tasks of various difficulties. However it is a toy dataset that has a very limited vocabulary and lacks the complexity of real world datasets (noise, inconsistencies, larger vocabularies, more complex language constructs, and so on). Another limitation of our evaluation is only using text, specifically question answering. To fully evaluate the idea of adding rotations to memory cells, in the future, we aim to look into incorporating our rotations on different domains and tasks including speech to text, translation, language generation, stock prices, and other common problems using real world datasets.

Tuning the hyperparameters of the rotation models might give better insights and performance increases and is something we aim to incorporate in our training pipeline in the future.

A brief exploration of the angles produced by \mathbf{u} and the weight matrix \mathbf{W}_{rot} show that \mathbf{u} does not saturate, thus rotations are in fact applied to our cell states and do not converge to 0 (or 360 degrees). A more in-depth qualitative analysis of the rotation gate is planned for future work. Peeking into the activations of our rotation gates could help understand the behaviour of rotations and to what extent they help better represent long-term memory.

A very successful and popular mutation of the LSTM is the Gated Recurrent Unit (GRU) unit (Cho et al., 2014). The GRU only has an output as opposed to both a cell state and an output and uses fewer gates. In the future we hope to explore adding rotations to GRU units and whether we can obtain similar results.

5 CONCLUSION

We have introduced a novel gating mechanism for RNN units that enables applying a parametrised transformation matrix to the cell state. We picked pairwise 2D rotations as the transformation and shown how this can be added to the popular LSTM units to create what we call RotLSTM.

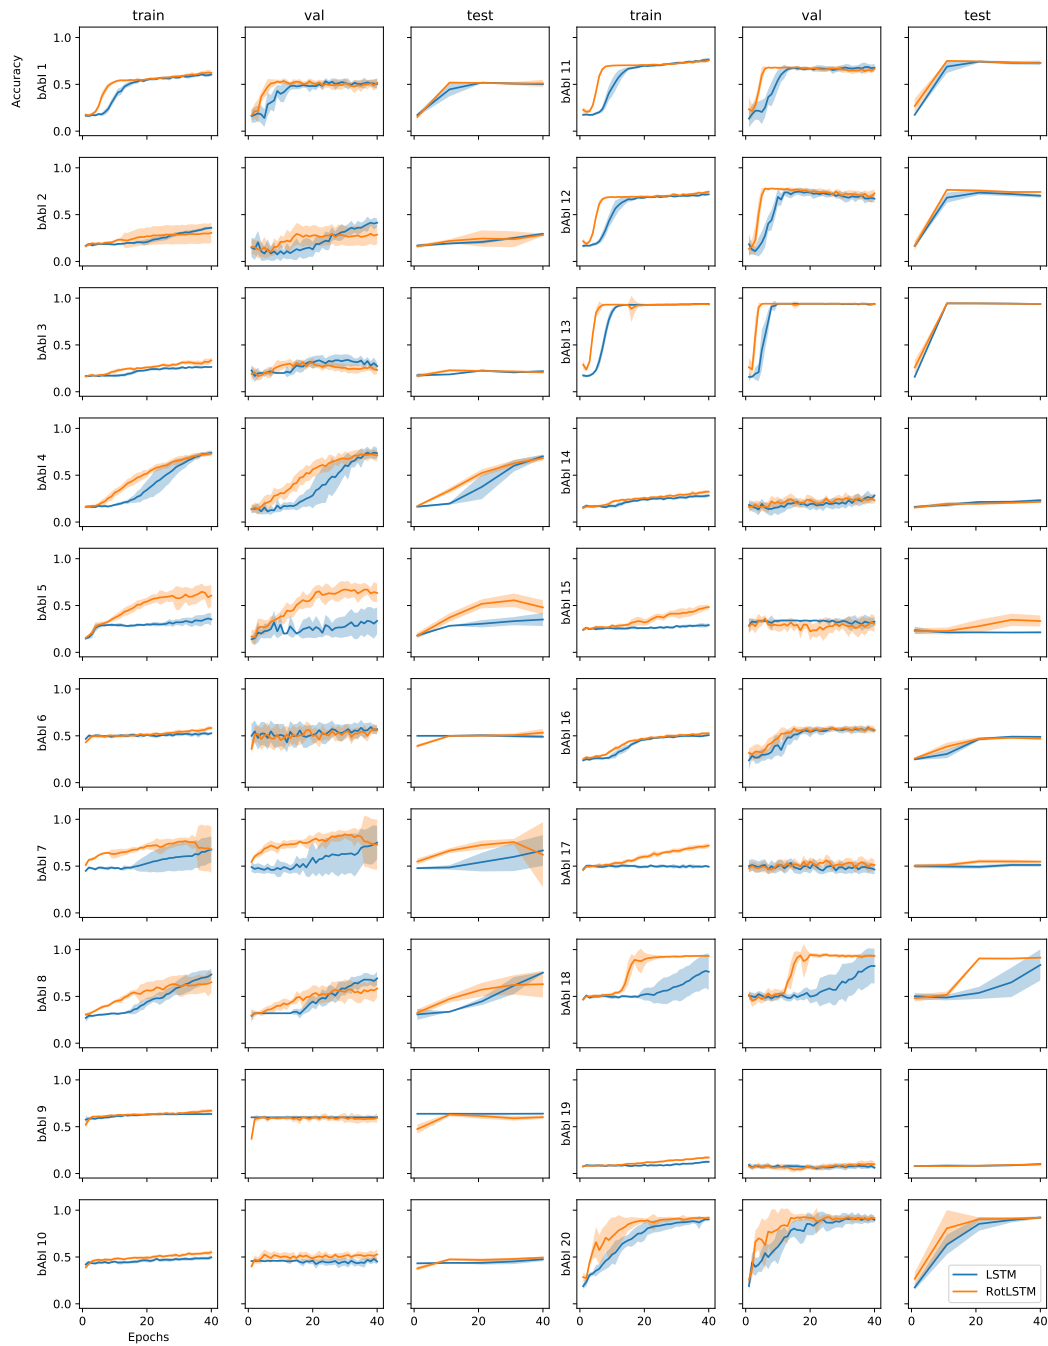


Figure 3: Accuracy comparison on training, validation (val) and test sets over 40 epochs for LSTM and RotLSTM models. The models were trained 10 times and shown is the average accuracy and in faded colour is the standard deviation. Test set accuracy was computed every 10 epochs.

We trained a simple model using RotLSTM units and compared them with the same model based on LSTM units. We show that for the LSTM-based architectures adding rotations has a positive impact on most bAbI tasks, making the training require fewer epochs to achieve similar or higher accuracy. On some tasks the RotLSTM model can use a lower dimensional cell state vector and maintain its performance.

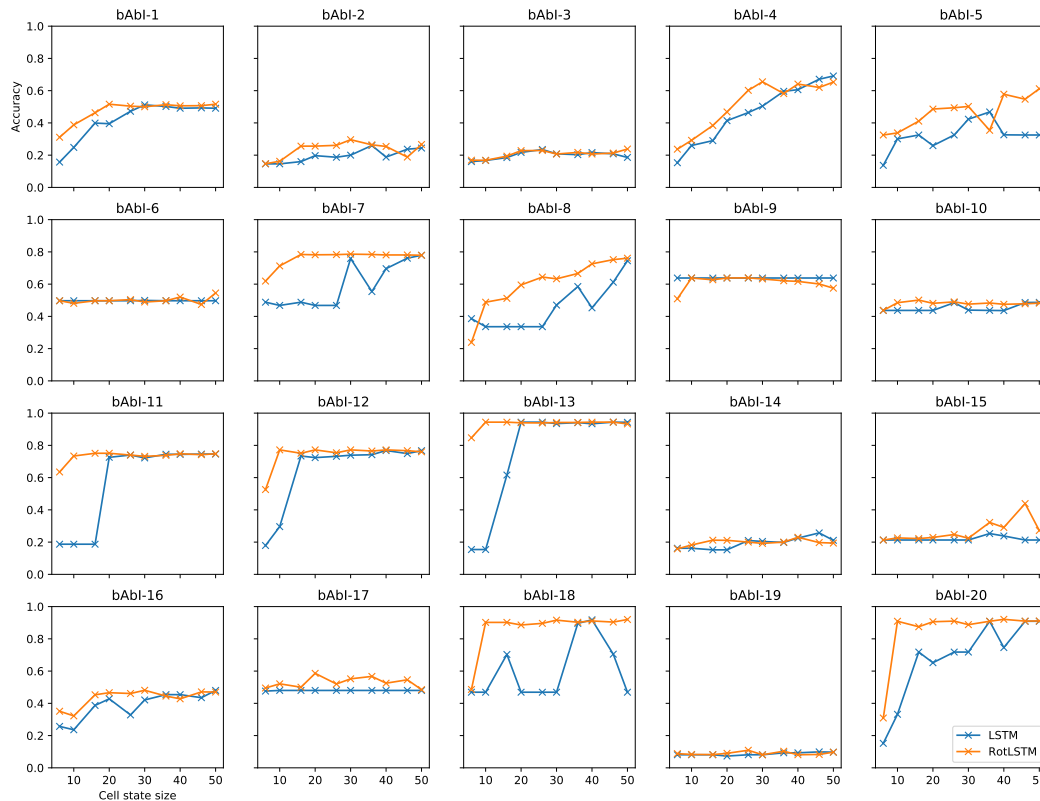


Figure 4: Accuracy on the test set for the LSTM and RotLSTM while varying the cell state size from 6 to 50. The shown numbers are for the epochs with best validation set accuracy.

Significant accuracy improvements of approximately 20% for the RotLSTM model over the LSTM model are visible on bAbl tasks 5 (three argument relations) and 18 (reasoning about size).

REFERENCES

- Sercan O Arik, Mike Chrzanowski, Adam Coates, Gregory Diamos, Andrew Gibiansky, Yongguo Kang, Xian Li, John Miller, Jonathan Raiman, Shubho Sengupta, et al. Deep voice: Real-time neural text-to-speech. *arXiv preprint arXiv:1702.07825*, 2017.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- François Chollet et al. Keras. <https://github.com/fchollet/keras>, 2015.
- Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *Acoustics, speech and signal processing (icassp), 2013 ieee international conference on*, pp. 6645–6649. IEEE, 2013.
- Klaus Greff, Rupesh K Srivastava, Jan Koutník, Bas R Steunebrink, and Jürgen Schmidhuber. Lstm: A search space odyssey. *IEEE transactions on neural networks and learning systems*, 2017.
- Awni Y. Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y. Ng. Deep speech: Scaling up end-to-end speech recognition. *CoRR*, abs/1412.5567, 2014. URL <http://arxiv.org/abs/1412.5567>.

- Irina Higgins, Nicolas Sonnerat, Loic Matthey, Arka Pal, Christopher P Burgess, Matthew Botvinick, Demis Hassabis, and Alexander Lerchner. Scan: Learning abstract hierarchical compositional visual concepts. *arXiv preprint arXiv:1707.03389*, 2017.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.
- Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain Paulus, and Richard Socher. Ask me anything: Dynamic memory networks for natural language processing. In *International Conference on Machine Learning*, pp. 1378–1387, 2016.
- Ramesh Nallapati, Bing Xiang, and Bowen Zhou. Sequence-to-sequence rnns for text summarization. *CoRR*, abs/1602.06023, 2016. URL <http://arxiv.org/abs/1602.06023>.
- Sebastian Riedel, Matko Bosnjak, and Tim Rocktäschel. Programming with a differentiable forth interpreter. *CoRR*, abs/1605.06640, 2016.
- Tim Rocktäschel and Sebastian Riedel. End-to-end differentiable proving. *arXiv preprint arXiv:1705.11040*, 2017.
- Mihaela Rosca and Thomas Breuel. Sequence-to-sequence neural network models for transliteration. *CoRR*, abs/1610.09565, 2016. URL <http://arxiv.org/abs/1610.09565>.
- Russell Stewart, Mykhaylo Andriluka, and Andrew Y Ng. End-to-end people detection in crowded scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2325–2333, 2016.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. End-to-end memory networks. In *Advances in neural information processing systems*, pp. 2440–2448, 2015.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. Towards ai-complete question answering: A set of prerequisite toy tasks. *CoRR*, abs/1502.05698, 2015. URL <http://arxiv.org/abs/1502.05698>.
- Qi Wu, Damien Teney, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Visual question answering: A survey of methods and datasets. *Computer Vision and Image Understanding*, 2017.

A BABI TASKS

A summary of all bAbI tasks is shown in Table 2 for reference.

Table 2: Descriptions of bAbI dataset tasks.

#	Description	#	Description
1	Factoid QA with one supporting fact	11	Basic coreference
2	Factoid QA with two supporting facts	12	Conjunction
3	Factoid QA with three supporting facts	13	Compound coreference
4	Two argument relations: subject vs. object	14	Time manipulation
5	Three argument relations	15	Basic deduction
6	Yes/No questions	16	Basic induction
7	Counting	17	Positional reasoning
8	Lists/Sets	18	Reasoning about size
9	Simple Negation	19	Path finding
10	Indefinite Knowledge	20	Reasoning about agent’s motivation