

Multi-layer Relation Networks for relational reasoning

Marius Jahrens

Institute for Neuro- and Bioinformatics
Lübeck, Germany
jahrens@inb.uni-luebeck.de

Thomas Martinetz

Institute for Neuro- and Bioinformatics
Lübeck, Germany
martinetz@inb.uni-luebeck.de

ABSTRACT

Relational Networks (RN) as introduced by Santoro et al. in 2017 have demonstrated strong relational reasoning capabilities with a rather shallow architecture. Its single-layer design, however, only considers pairs of information objects, making it unsuitable for problems requiring reasoning across a higher number of facts. To overcome this limitation, we propose a multi-layer relation network architecture which enables successive refinements of relational information through multiple layers. We show that the increased depth allows for more complex relational reasoning by applying it to the bAbI 20 QA dataset, solving all 20 tasks with joint training and surpassing the state-of-the-art results.

CCS CONCEPTS

• Computing methodologies → Machine learning algorithms;

ACM Reference Format:

Marius Jahrens and Thomas Martinetz. 2019. Multi-layer Relation Networks for relational reasoning. In *2nd International Conference on Applications of Intelligent Systems (APPIS 2019), January 7–9, 2019, Las Palmas de Gran Canaria, Spain*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3309772.3309782>

1 INTRODUCTION

While Neural Networks have been very successful in classification and regression, their ability to cope with relational tasks has remained rather limited. In order to overcome this limitation, Neural Networks need a way to perform reasoning on different aspects of an input signal, so that more complex relations can be deduced through logical chains.

Relation Networks (RNs) as introduced by Santoro et al. [8] were a milestone in this direction. Their architecture is illustrated in Figure 1. RNs expect object vectors as input from previous layers or from samples, representing bits of information from the same domain. These objects can be as different as outputs of recurrent layers encoding the information content of sentences in natural language processing or feature vectors representing abstract image features in the output of stacked convolutional layers. The objects' relations to one another with respect to an input question q is evaluated by an evaluation module g for all combinatorial pairs of objects. The objects are expected to contain all information

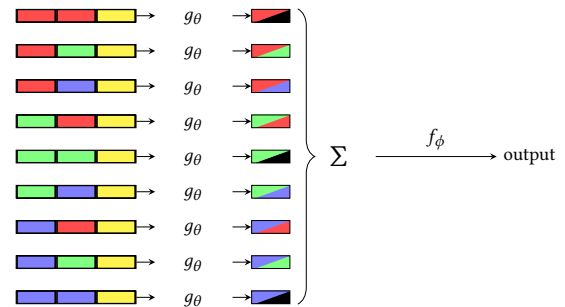


Figure 1: In the RN architecture pairs of objects (o_i, o_j) (in red, green, blue) with input question q (yellow) are handled independently by the relation evaluation function g_θ . The resulting vectors are summed up and the sum vector is transformed by f_ϕ to form the final network output.

necessary to reason about their relation, most importantly also including their relative or absolute spatial or, in case of sequences, temporal position. The extracted relational information is then combined by summing the output vectors of g , which keeps the result invariant to the order of the objects and the order in which the pairs are evaluated. Finally, the combined vector is fed to an output function f , which leads to

$$RN(O, q) = f_\phi \left(\sum_{i,j} g_\theta(o_i, o_j, q) \right) \quad (1)$$

as the final output of the RN. f and g are typically chosen as MLPs with weights ϕ and θ , respectively. Since Relation Networks are differentiable if the internal functions are modeled accordingly, e.g. by using MLPs, the combination of object generating layers and Relation Network components can be trained end-to-end, so that suitable object representations can be learned inside the network.

A number of extensions of the original RN have been introduced. In a recent paper by Palm et al. [6] the network has one recurrent cell for each input object which is repeatedly fed its previous state, its input object, as well as transformed previous outputs from the other objects' recurrent cells. It can be shown that their Recurrent Relation Network is a special case of our approach, if our relation evaluation functions are made to share weights across all layers. However, doing so forces the inputs of each layer to reside in the same domain, thereby limiting layers' output representations to have a similar level of abstraction as the inputs. In another approach to enhance Relation Networks, Moon et al. [5] proposed a memory enhanced variation with an attention mechanism which, while improving upon the RN results, doesn't solve all bAbI tasks. Other related architectures like RelNet by Bansal et al. [1] were

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

APPIS 2019, January 7–9, 2019, Las Palmas de Gran Canaria, Spain

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6085-2/19/01.

<https://doi.org/10.1145/3309772.3309782>

benchmarked on each task separately, yet our approach achieves similar or better results despite being trained on all tasks jointly.

We extend the RN simply by further layers, however, in a way that it scales linearly to arbitrary numbers of such layers. We stack these layers such that more complex relations can be inferred, now being able to solve tasks that proved to be difficult in the past due to the number of related facts. The main advantage of our approach over other attempts to enhance RNs is its simplicity and its scalability, all while showing improved performance on the benchmark dataset.

2 MULTI-LAYER RELATION NETWORKS

The single-layer architecture discussed above only allows the relation evaluation function g_θ to consider information from two objects, making it unsuitable for problems requiring reasoning across more than two facts. While expanding g to evaluate triplets would be possible, the computational complexity would increase from $O(n^2)$ to $O(n^3)$ with n being the number of objects. Hence, simply scaling up the combinatorial layer is not feasible from a performance point of view.

Instead we sum up all the relations which reason about an object o_i , so that we end up with one sum vector per object. We then use these vectors as input objects for another single-layer Relation Network, thus forming a double-layer RN in total. The architecture of this Double-Layer Relation-Network (DLRN) is shown in Figure 2 and described by

$$DLRN(O, q) = f_\phi \left(\sum_{i,k} h_\psi \left(\sum_j g_\theta(o_i, o_j, q), \sum_l g_\theta(o_k, o_l, q), q \right) \right). \quad (2)$$

This way the relation evaluation function h_ψ reasons about objects which already contain binary relational knowledge. The construction process can be repeated for arbitrary layer stack counts m , with $m = 0$ being equivalent to the single-layer RN. In a recursive formulation we obtain

$$MLRN(O, q) = f \left(\sum_{i,k} h_{m,i,k} \right) \quad (3)$$

$$h_{d,i,k} = h_d \left(\sum_j h_{d-1,i,j}, \sum_l h_{d-1,k,l}, q \right) \quad d = 1, \dots, m \quad (4)$$

$$h_{0,i,k} = g(o_i, o_k, q). \quad (5)$$

Note that the Multi-Layer Relation Network (MLRN) has a computational complexity of $O(n^2m)$, making it computationally suitable for more complex reasoning tasks.

3 EXPERIMENTS

3.1 The bAbI dataset

Published by Weston et al. (2015) [9], the bAbI 20 QA tasks are a dataset of algorithmically generated sentences and related questions and answers formulated in natural language. Each task represents a type of reasoning as denoted in Table 1. The dataset contains 1000 or 10000 training/validation samples per task and 100 or 1000 test samples on the small and the large variant, respectively. For our work and all comparisons, we used the 10k variant with joint training on all tasks and predefined separation into 9k training samples,

1k validation samples and 1k test samples. A task is considered to be solved successfully if the test accuracy exceeds 95%.

Under the same conditions, Differentiable Neural Computers (DNC) have managed to solve 18 of these 20 tasks [8], Sparse DNCs 19/20 [8], Relation Networks 18/20 [8], Recurrent RNs 19/20 [6], Relation Memory Networks 19/20 [5] and Working Memory Networks 19/20 [7]. While Adaptive Memory Networks [4] reported 20/20 successful tasks, their experiment settings mention different model training parameters for different types of tasks, hinting that the training was probably not performed on all tasks jointly.

Examples of the most difficult tasks 2, 3, and 16, are given in the Appendix in Table 2. Note that for task 16 about 6% of the samples are not disambiguous as the context provides multiple possible answers and the answer provided in the label is not always the one with the most supporting facts.

3.2 Comparison of single- and double-layer RN

Since choosing layer count $m=0$ yields an architecture which is equivalent to the original single-layer RN, we decided to reproduce the original results first and use the determined hyperparameters as a starting point for the multi-layer experiments.

As in the original RN paper we reduced or padded each sample’s context to contain exactly 20 sentences. While this preprocessing step might have an influence on the difficulty of the tasks since fewer unimportant sentences in the context equate to less input noise, preliminary tests of the MLRN indicated that this limitation does not have a significant impact on the network’s performance. To produce the input objects for the relation network, we use a word embedding matrix and feed the embedded word vectors, for each sentence separately, into an LSTM with peepholes [3]. The final LSTM output is then concatenated with a position encoding vector which is a one-hot vector, encoding the position of the sentence inside the current sample’s context, with a randomized padding length per sample to reduce overfitting (as proposed by Adam Santoro, one of the authors of the original RN paper, in personal communication). The word embedding matrix and the LSTM weights were learned together with the RN weights during the end-to-end training.

We chose a word embedding dimension of 256, an LSTM cell with 32 units, a position encoding vector size of 40, $f = MLP(256, 512, dict_size)$ with ReLU activations for the inner layers and a final linear layer with softmax for the output. For the relation evaluation function, we chose $g = MLP(256, 256, 256, 256)$ with ReLU activations, added L2 weight regularization terms with a penalty of $2e-5$ for the weights in f and g and used a cross-entropy loss function. For the gradient descent we used the adam optimizer with a fixed learning rate of $5e-5$ and a batch size of 32. About 4 million training steps were necessary to finish the training, taking about 4 days on a single GTX 1080 Ti.

For the double-layer case ($m = 1$), we used the same parameters as before, except switching to a weight penalty factor of $1e-4$ and choosing $h_1 = MLP(256, 256, 256)$. The training took 6 million steps, equating to about 8 days on a single GTX 1080 Ti.

For the sake of comparison tasks 2 and 3 are of most interest because they represent the failure cases of the single-layer RNs. Task 2 consists of questions with two supporting facts and task 3

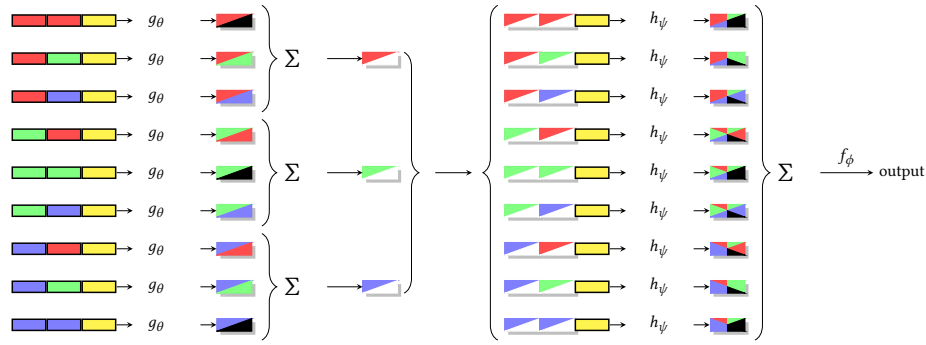


Figure 2: Double-layer RN architecture: The output vectors of the relation evaluation function g_θ are grouped and summed by the first index i of its input pairs (o_i, o_j) (in red, green, blue) to form new object vectors (red/white, green/white, blue/white). The new objects are used as inputs for a single-layer RN (right half) with relation evaluation function h_ψ .

| | Original RN results | Layer count $m = 0$ | Layer count $m = 1$ |
|---------------------------------------|---------------------|---------------------|---------------------|
| Task 1: Single Supporting Fact | > 95 | 100 | 100 |
| Task 2: Two Supporting Facts | 91.9 | 81.3 | 99.8 |
| Task 3: Three Supporting Facts | 83.5 | 90.5 | 99.3 |
| Task 4: Two Argument Relations | > 95 | 100 | 100 |
| Task 5: Three Argument Relations | > 95 | 99.1 | 99.9 |
| Task 6: Yes/No Questions | > 95 | 100 | 100 |
| Task 7: Counting | > 95 | 100 | 100 |
| Task 8: Lists/Sets | > 95 | 100 | 100 |
| Task 9: Simple Negation | > 95 | 100 | 100 |
| Task 10: Indefinite Knowledge | > 95 | 99.4 | 100 |
| Task 11: Basic Coreference | > 95 | 99.1 | 100 |
| Task 12: Conjunction | > 95 | 100 | 100 |
| Task 13: Compound Coreference | > 95 | 100 | 100 |
| Task 14: Time Reasoning | > 95 | 99.9 | 100 |
| Task 15: Basic Deduction | > 95 | 100 | 100 |
| Task 16: Basic Induction | 97.9 | 97.0 | 97.4 |
| Task 17: Positional Reasoning | > 95 | 97.8 | 98.3 |
| Task 18: Size Reasoning | > 95 | 99.5 | 99.6 |
| Task 19: Path Finding | > 95 | 99.9 | 100 |
| Task 20: Agent's Motivations | > 95 | 100 | 100 |
| Tasks succeeded (> 95%) | 18/20 | 18/20 | 20/20 |
| Mean error | unknown | 1.825 | 0.285 |

Table 1: Task accuracies on the test set in % for a single-layer RN as reported by the original authors, a single-layer RN with our generalized architecture ($m = 0$) and our double-layer RN ($m = 1$). Bold are the two tasks which significantly improved due to higher order relations.

represents questions with three supporting facts. Note that while our results for tasks 2 and 3 seem to be interchanged when compared to the results from the original paper, no evidence for a fault could be found. The results in Table 1 show that the increased layer count does indeed allow the network to learn more complex relations.

The MLRN not only learns more complex relations, but also learns more quickly. In Figure 3 we show the learning curves on the validation sets of Tasks 2 and 3 and of Tasks 16 and 19. As we can see, for reaching the smaller error in Task 2 and 3 no more training is necessary. The final error is reached very quickly. The faster learning is even more evident in Task 19. The single- and the double-layer RN both reach the same small error in this less complex task, but the MLRN reaches it with almost an order of magnitude less training steps. For Task 16 also our MLRN needs

many training steps for learning, even more than the single-layer RN. Task 16 is the reason why the overall training on all tasks takes about twice as much time as for the single-layer RN. Also the single-layer RN has difficulties learning this task. We expect that these difficulties are due to the ambiguity of the training data, as shown in the Appendix. It is well known that the kind of symmetry breaking, which has to take place in such cases, is usually very time consuming.

4 CONCLUSION

We have proposed a generalization of the Relation Network which can be scaled easily to multiple layers for higher task complexities. We have shown that this Multi-Layer Relation Network is capable of solving all of the bAbI QA tasks with joint training on all tasks, making it to our knowledge the first model to achieve this, all

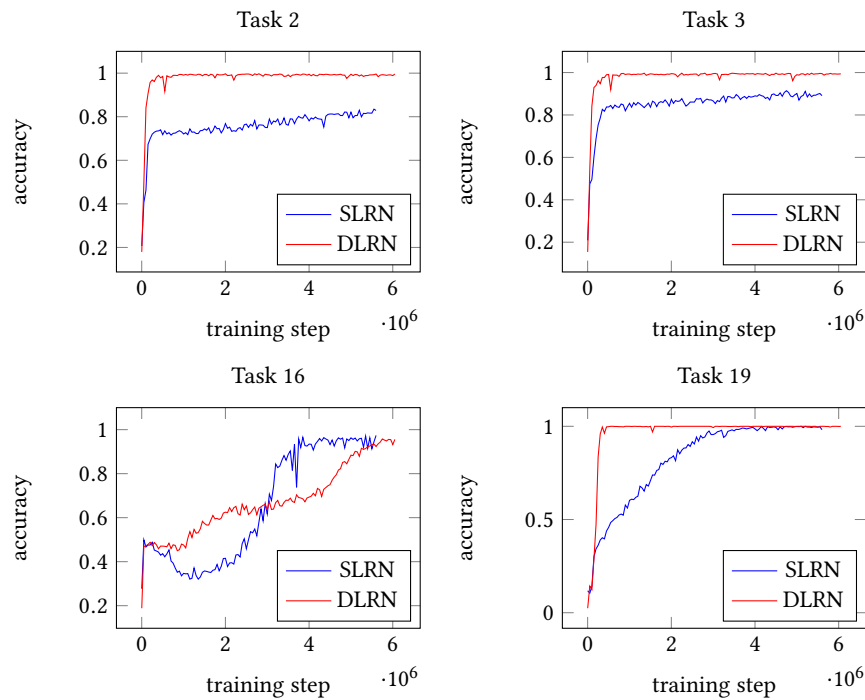


Figure 3: Accuracy on the validation set for a single-layer RN and a double-layer RN periodically measured during training.

without resorting to using ensembles or taking the best of multiple runs. Its total test error rate of 0.285% even matches the state-of-the-art of models trained on each task separately [1]. In most cases the MLRN also learns its tasks significantly faster than the single-layer RN, despite more parameters to train. This is a hint that its structure is better suited for these tasks.

When compared to other proposed Relation Network enhancements, we attribute its advantage over Recurrent RNs [6] to the fact that the shared weights in the recurrent architecture don't allow them to learn more complex relations in deeper layers because the weight-sharing forces all the layer outputs to be in the same domain, and therefore a similar level of abstraction, as the first layer's input.

With its simple construction the MLRN is suitable for a large variety of problems. In a next step we will evaluate our multi-layer architecture on further relational reasoning tasks like the relation of objects in images as in [8] or even challenges for abstract reasoning inspired by human IQ tests as proposed by [2].

REFERENCES

- [1] Trapit Bansal, Arvind Neelakantan, and Andrew McCallum. 2017. RelNet: End-to-End Modeling of Entities & Relations. arXiv:1706.07179. arXiv:1706.07179
- [2] David G.T. Barrett, Felix Hill, Adam Santoro, Ari S. Morcos, and Timothy Lillicrap. 2018. Measuring abstract reasoning in neural networks. arXiv:1807.04225. arXiv:1807.04225
- [3] Felix A. Gers and Juergen Schmidhuber. 2000. Recurrent nets that time and count. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, Vol. 3. IEEE, 189–194 vol.3. <https://doi.org/10.1109/ijcnn.2000.861302>
- [4] Daniel Li and Asim Kadav. 2018. Adaptive Memory Networks. arXiv:1802.00510. arXiv:1802.00510
- [5] Jihyung Moon, Hyochang Yang, and Sungzoon Cho. 2018. Finding ReMO (Related Memory Object): A simple neural architecture for text based reasoning. arXiv:1801.08459. arXiv:1801.08459
- [6] Rasmus Berg Palm, Ulrich Paquet, and Ole Winther. 2017. Recurrent Relational Networks for complex relational reasoning. arXiv:1711.08028. arXiv:1711.08028
- [7] Juan Pavez and Héctor Allende. 2018. Working Memory Networks: Augmenting Memory Networks with a Relational Reasoning Module. arXiv:1805.09354. arXiv:1805.09354
- [8] Adam Santoro, David Raposo, David G.T. Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. 2017. A simple neural network module for relational reasoning. arXiv:1706.01427. arXiv:1706.01427
- [9] Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M. Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. 2015. Towards ai-complete question answering: A set of prerequisite toy tasks. arXiv:1502.05698. arXiv:1502.05698

A BABI EXAMPLES

| | |
|------------------------|---|
| Task 2 | <p>Context:</p> <ol style="list-style-type: none"> 1. Mary got the milk there. 2. John moved to the bedroom. 3. Sandra went back to the kitchen. 4. Mary travelled to the hallway. <p>Question:</p> <p>Where is the milk?</p> <p>Answer:</p> <p>hallway</p> |
| Task 3 | <p>Context:</p> <ol style="list-style-type: none"> 1. John moved to the bedroom. 2. John grabbed the apple there. 3. Sandra moved to the hallway. 4. John went to the office. 5. Sandra went back to the bedroom. 6. Sandra took the milk. 7. John journeyed to the bathroom. 8. John travelled to the office. 9. Sandra left the milk. 10. Mary went to the bedroom. 11. Mary moved to the office. 12. John travelled to the hallway. 13. Sandra moved to the garden. 14. Mary moved to the kitchen. 15. Daniel took the football. 16. Mary journeyed to the bedroom. 17. Mary grabbed the milk there. 18. Mary discarded the milk. 19. John went to the garden. 20. John discarded the apple there. <p>Question:</p> <p>Where was the apple before the bathroom?</p> <p>Answer:</p> <p>office</p> |
| Task 16 (ambiguous) | <p>Context:</p> <ol style="list-style-type: none"> 1. Greg is a frog. 2. Bernhard is a swan. 3. Julius is a frog. 4. Bernhard is white. 5. Julius is green. 6. Lily is a frog. 7. Brian is a frog. 8. Lily is gray. 9. Brian is gray. <p>Question:</p> <p>What color is Greg?</p> <p>Answer:</p> <p>gray</p> |
| Task 16 (disambiguous) | <p>Context:</p> <ol style="list-style-type: none"> 1. Lily is a swan. 2. Bernhard is a lion. 3. Greg is a swan. 4. Bernhard is white. 5. Brian is a lion. 6. Lily is gray. 7. Julius is a rhino. 8. Julius is gray. 9. Greg is gray. <p>Question:</p> <p>What color is Brian?</p> <p>Answer:</p> <p>white</p> |

Table 2: Samples from the most difficult tasks 2, 3 and 16 (one sample with and one without ambiguity).