

Abstract Rule Based Pattern Learning with Neural Networks

Radha Kopparti

Department of Computer Science, City, University of London

Abstract

In this research work, the problem of learning abstract rules using neural networks is studied and a solution called ‘Relation Based Patterns’ (RBP) which model abstract relationships based on equality is proposed.

The ability to learn abstractions and generalise is seen as the essence of human intelligence (Brooks 1991). Since 1950s, there have been efforts to build systems that learn and think like humans (Lake et al. 2016). It is observed that humans including infants tend to have good generalisation power when compared to the machine learning models in which hypothesis is usually approximated and may be prone to errors.

The examples proposed by Marcus (Marcus et al. 1999; Marcus 2001; 2018) such as the failure to generalise equality, distinguish between even to odd numbers or the recognition of ABA or ABB patterns of syllables have attracted a significant amount of attention in psychology, particularly in the study of human language learning, but they have not been addressed systematically as problems of machine learning and neural networks.

Generally, humans are very effective at extracting abstract relations (eg: as in Figure 1) from sensory input, often after very brief exposure.




Figure 1: How do humans perceive grammatical rules?

In abstract pattern learning, series of abstract patterns be it grammar like rules or sequences of data like music or language are used. One of the earliest work on abstract pattern learning was by (Marcus et al. 1999; Marcus 2001) where

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

the abstract grammar like rules are shown to six month old infants and their task was to distinguish the grammatical rules. The infants were exposed to sequences of one of the forms ABA or ABB, e.g. ‘le di le’ or ‘le di di’, for a few minutes in the familiarisation phase. In the test phase the infants were exposed to sequences with a different vocabulary (e.g. ‘ba po ba’ and ‘ba po po’) and the results from the experiments was that the infants were able to learn the grammatical patterns within minutes of habituation. The same task when reproduced using a recurrent neural network, failed to distinguish the grammar patterns.

The work by (Elman 1999; Seidenberg and Elman 1999; Christiansen and Curtin 1999; Shastri and Chang 1999) focused on using distributed representation of the inputs where the network was not able to learn the patterns as expected. Other works used localist representation (Altmann 2002; Altmann and Dienes 1999; Alhama and Zuidema 2016; Dominey and Ramus 2000) which suggests that an additional context or prior experience is necessary for the network to learn these identity rules. This has raised questions on whether the existing statistical and neural network models can generalize these abstract patterns or not.

An effective way of solving this problem is to introduce an inductive bias within the network structure to achieve better generalisation from fewer samples of training data. However, the question then arises about the type of inductive bias needed to improve the overall generalisation performance. To achieve better generalisation, the neural network should be able to learn the required knowledge and apply that to unseen circumstances or extend it beyond the scope of the actual problem.

In this research work, the problem of learning abstract rules using neural networks is studied and a solution called ‘Relation Based Patterns’ (RBP) which model abstract relationships based on equality is proposed. RBP creates an inductive bias in the neural networks that leads to the learning of generalisable solutions. It is observed that integration of RBP leads to almost perfect generalisation in abstract rule learning tasks with synthetic data and to improvements in neural language modelling on real-world data.

The primary focus of the research is on the following aspects :

- Why can't standard neural networks learn grammatical patterns for example say based on equality?
- How can we make the neural networks learn these grammatical abstractions?
- What does it mean to other real world tasks like language modelling where there are higher level of abstractions in data?

The RBP model is based on an input that consists of multiple items, where each can be represented by a vector of input neurons. DR units are used to compare corresponding neurons in different vectors. For the comparison differentiator-rectifier (DR) units are introduced, which calculate the absolute difference of two inputs: $f(x, y) = |x - y|$. One DR unit for every pair of corresponding input unit is created with the weights from the inputs to the DR units fixed at 1.

There are multiple vector comparisons that correspond to the different equality relations in patterns, e.g. equality of the pairs of vectors in positions (1,2), (1,3), (2,3), to recognise patterns of the forms ABA, ABB, AAB, ABC etc.

DR units are integrated as *Early Fusion* (added to the input layer), *Mid Fusion* (added to the hidden layer) and *Late Fusion* (added to the output layer) in standard neural networks.

RBP has been tested on various tasks like learning abstract relations of the form forms ABA, ABB, ABC, AAB and BAB in various orders to the network as a supervised formulation of (Marcus et al. 1999) with some variants. It was observed that without RBP networks never improve above chance level (50%), and Mid Fusion leads to almost perfect results.

RBP has also been tested on character prediction and melody prediction tasks, resulting in significant and consistent improvements, indicating that having a way of modelling lower level abstract relations in data improves the performance in real world tasks like language modelling.

Overall, through this study

- several neural network architectures like feed-forward networks and recurrent networks with their gated variants (LSTM and GRU) are evaluated and it was confirmed that these neural networks do not learn abstract grammar rules as expected.
- RBP (Relation Based Patterns) method as an inductive bias has been proposed to enable the learning of abstract grammatical patterns within the neural network structures.
- RBP can be integrated into standard neural network architectures in early, mid and late fusion settings.
- the networks with suitable RBP structure learned the abstract grammar patterns with 100% accuracy.
- the integration of RBP to neural network models improved the performance in neural language modelling tasks along with artificial grammar learning tasks which proves that RBP can be expanded to other sequential rule learning tasks as well.

In future this work can be extended towards improving the performance of neural networks in providing better ab-

stractions and generalizations for other forms of abstract relations on more complex tasks such as question answering or perception-based reasoning and other relational learning tasks.

References

- Alhama, R. G., and Zuidema, W. 2016. Pre-wiring and pre-training: What does a neural network need to learn truly general identity rules. *CoCo at NIPS*.
- Altmann, G., and Dienes, Z. 1999. Technical comment on rule learning by seven-month-old infants and neural networks. *In Science* 284(5416):875–875.
- Altmann, G. 2002. Learning and development in neural networks the importance of prior experience. *In Cognition* 85(2):B43–B50.
- Brooks, R. A. 1991. Intelligence without representation. *Artificial Intelligence* : [https://doi.org/10.1016/0004-3702\(91\)90053-M](https://doi.org/10.1016/0004-3702(91)90053-M) 47:139–159.
- Christiansen, M. H., and Curtin, S. L. 1999. The power of statistical learning: No need for algebraic rules. *Proceedings of the 21st annual conference of the Cognitive Science Society* 114:119.
- Dominey, P., and Ramus, F. 2000. Neural network processing of natural language: Isensitivity to serial, temporal and abstract structure of language in the infant. *Language and Cognitive Processes* 15(1),87–127.
- Elman, J. 1999. Generalization, rules, and neural networks: A simulation of marcus et. al. <https://crl.ucsd.edu/~elman/Papers/MVRVsimulation.html>.
- Lake, B. M.; Ullman, T. D.; Tenenbaum, J. B.; and Gershman, S. J. 2016. Building machines that learn and think like people. *CoRR* abs/1604.00289.
- Marcus, G. F.; Vijayan, S.; Rao, S.; and Vishton, P. 1999. Rule learning by seven-month-old infants. *Science*, 283 5398:77–80.
- Marcus, G. F. 2001. The algebraic mind: Integrating connectionism and cognitive science. *Cambridge MIT Press*.
- Marcus, G. F. 2018. Deep learning : a critical appraisal. *arXiv:1801.00631*.
- Seidenberg, M., and Elman, J. 1999. Do infants learn grammar with algebra or statistics? *Science* 284(5413):433–433.
- Shastri, and Chang. 1999. A spatiotemporal connectionist model of algebraic rule-learning. *International Computer Science Institute* TR–99–011.