

---

# Feed-Forward Neural Networks Need Inductive Bias to Learn Equality Relations

---

**Tillman Weyde**

Department of Computer Science  
City, University of London  
London, United Kingdom  
t.e.weyde@city.ac.uk

**Radha Manisha Kopparti**

Department of Computer Science  
City, University of London  
London, United Kingdom  
radha.kopparti@city.ac.uk

## Abstract

Basic binary relations such as equality and inequality are fundamental to relational data structures. Neural networks should learn such relations and generalise to new unseen data. We show in this study, however, that this generalisation fails with standard feed-forward networks on binary vectors. Even when trained with maximal training data, standard networks do not reliably detect equality.

We introduce differential rectifier (DR) units that we add to the network in different configurations. The DR units create an inductive bias in the networks, so that they do learn to generalise, even from small numbers of examples and we have not found any negative effect of their inclusion in the network. Given the fundamental nature of these relations, we hypothesize that feed-forward neural network learning benefits from inductive bias in other relations as well. Consequently, the further development of suitable inductive biases will be beneficial to many tasks in relational learning with neural networks.

## 1 Introduction

Basic relations such as equality are fundamental to relational data structures. One goal of applying neural networks to relational data is that the networks learn to infer these relational structure from data. Although equality is typically not learned from data, equality or approximate equality may be embedded as part of other tasks. The modelling of equality is clearly in the hypothesis space of feed-forward neural networks (FFNNs) [1], but [2, 3] already highlighted that learning of identity relationships with neural networks may not generalise to unseen data. Therefore, we see learning to recognise equality as relevant from a theoretical and practical perspective.

In this study we test whether feed-forward networks learn equality as well as a numeric comparison, thresholded digit sum, and digit reversal of pairs of binary vectors and then generalise this to new data in different settings regarding the task, the amount of data provided, and the depth of the network. We find that the recognition of binary relations is not generalised reliably by feed-forward networks.

To address this problem, we introduce an inductive bias with additional predefined network structures, that we call differential rectifier (DR) units. We find in our experiments that DR units induce reliable perfect generalisation for equality and all other tasks except in digit reversal.

We see two questions that these results raise: First, which other relations neural networks do not learn and what that means for more complex tasks. Second, what kinds of inductive biases to design and how to implement them.

The remainder of this paper is organised as follows: Section 2 reviews related literature, Section 3 introduces the task of learning vector equality and our DR units for inductive bias. Section 4 presents the experimental results and in Section 5 follow the conclusions of this paper.

## 2 Related work

In relational learning, equality is often not learned from the data, with the exception of the work by [4] who learn to detect equality attributed of objects from images. Learning equality could be interesting in the context of constraint learning [5] to learn when equality constraints should be regarded as satisfied. Another relevant area is rule learning and application, where soft unification like in [6] could be replaced with a learnt model.

Since neural networks are currently by far the most popular machine learning method, it seems of interest whether they can learn equality. There have been a number of theoretical contributions showing that feed-forward networks are universal approximators, most generally to our knowledge by [1]. Presumably because of these results there was relatively little interest in the question which functions neural networks can not learn. One of the few studies in this direction was undertaken in [2] in 1999, where a recurrent neural network failed to distinguish abstract patterns, based on equality relations between sequence elements, although seven-month-old infants showed the ability to distinguish them after a few minutes of exposure. This was followed by an lively exchange on rule learning by neural networks and in human language acquisition, where results by [7–9] could not be reproduced by [10, 11] and [12] disputed claims by [11]. Other approaches, such as [13–15], use different network architectures, problem formulations or evaluation methods.

A more specific problem of learning equality relations was posed in [3] by showing that learning of equality on even numbers does not transfer to odd numbers in binary representation. This relates to the input neuron for the least significant bit not being set to 1 during training. Recently, [16] addressed this specific problem with different approaches as an example for extrapolation and inductive biases for machine learning in natural language processing. However, they did not address the general question of learning equality with neural networks.

If standard neural networks do not generalise equality relations despite the solution being in their hypothesis space, as we will show for FFNNs below, then the question is how we can enable the learning of solutions that do generalise. Inductive biases as a solution can be realised in a number of ways and have been of increased interest recently [17, 18].

## 3 Equality relation learning

The studies listed above motivated the approach taken here to study a reduced problem outside common contexts such as image analysis or cognitive modelling: whether feed-forward neural networks trained with back-propagation generally have the ability to learn equality relations and generalise to unseen data.

The general task is to learn the relation between pairs of binary vectors. This leads to a binary classification of the pairs according to the equality or otherwise of its element vectors. We use a standard FFNN as sketched in Figure 1a). This network has  $2n$  input neurons, where  $n$  is the vector dimensionality. The hidden layer has 10 neurons with ReLu activation. The output layer has two neurons representing the two classes (equal/unequal), which use softmax activation. The training uses the Adam optimiser[19] with cross-entropy loss.

The data we train and test the network with is synthetically generated and we vary the type and the distribution of the data in the experiments below. We are interested in how many training examples are needed until the network learns to correctly classify pairs of equal vs. unequal vectors. This network, like the following ones have been implemented in Python using PyTorch (<http://pytorch.org>).

**Inductive bias creation with DR units** In our model, we use differential rectifier (DR) units that compare input values by calculating the absolute difference:  $f(x, y) = |x - y|$ . We create one DR unit for every vector dimension with weights from the inputs to the DR units fixed at 1, thus learning the suitable summation weights for the DRs is sufficient for creating a generalisable equality detector.

We use two ways of integrating DR units into the neural networks: *Early Fusion* and *Mid Fusion*. In Early Fusion, DR units are concatenated to input units 1b), and in Mid Fusion they are added to the hidden layer 1c). In both cases the existing input and hidden units are unchanged.




Figure 1: Network architectures: a) standard feed-forward network, DR integration with b) early fusion and c) mid fusion. The DR units receive their input in both cases from *vec 1* and *vec 2*.

Vector Dimensions	Plain FFNN	Early Fusion	Mid Fusion
n=2	52%	82%	100%
n=3	55%	75%	100%
n=5	37%	67%	100%
n=10	52%	75%	100%
n=30	65%	75%	100%
n=100	75%	100%	100%

Table 1: Accuracy of the different network types on pairs of vectors of different dimensions. The joint train and test data covers all possible equal vector pairs up to 1000, and a random selection where there are more. Only the Mid Fusion architecture reaches reliable equality detection.

## 4 Experiments and Results

We performed different sets of experiments using binary vectors for estimating vector equality in relation to vector dimensionality, data size and dataset structure. We also use two additional tasks to test the effect of DR units in different contexts.

**Effect of network architecture and vector dimensionality** We generate pairs of random binary vectors with dimensionality  $n$  between 2 and 100 as shown in table 1. We use all the possible binary vectors to generate equal pairs, i.e.  $2^n$  pairs, for  $n < 10$  and a random selection of 1000 vectors otherwise. We also generate the same number of randomly selected unequal vector pairs. Then we use stratified sampling to split the data 75:25 into train and test set. The network is then trained for 20 epochs, which led to convergence in all cases. We run 10 simulations for each configuration. The average results are shown in Table 1.

We see that the standard FFNNs never fully generalise, and in many cases barely exceed chance level (50%). The early fusion model improves results, but only reaches full performance for 100 dimensions. The Mid Fusion reaches perfect test performance in all cases.

For the plain FFNN, it looks like there is a trend towards better performance at higher dimensionality, but with the observed variation that may be coincidental. We did not perform an exhaustive grid search over all hyperparameters, but tested higher numbers of hidden layers (2,3), and larger hidden layers (20,30 neurons) without observing a significant change in the results.

**Effect of training data size** We study here how much the performance depends on the training data size. For this, we vary only the training data size and keep the test set and all other parameters constant. We use training data sizes of 1% to 50% (in relation to the totally available data as defined above) and the accuracy achieved in various conditions is plotted in Figure 2. It is worth noting, that the Mid Fusion network reaches 100% accuracy from 10% data size on while the FFNN shows only small learning effects.

**Effect of vector coverage** A possible hypothesis for the results of the FFNN is that the coverage of the vectors in the training set plays a role. To share vectors in equal pairs between training and test set would mean to train on the test data, but we created a training set that contains all vectors that




Figure 2: Accuracy of FFNN for 10 dimensional binary vectors after varying the distributions of training data from (1%-50%) keeping the testing data fixed

Type	a)	b)	c)	d)	e)
1) Plain FFNN	50%	52%	75%	77%	50%
2) Early Fusion	75%	87%	92%	82%	55%
3) Mid Fusion	100%	100%	100%	100%	58%

Table 2: Test set accuracy of FFNN without and with DR units for different vector coverage (a,b, see text for details) and for classification by c) numeric comparison ( $\geq$ ), d) digit sum  $\geq 3$ , and e) inversion of digits.

appear in the test set in the unequal pairs for  $n = 10$ . The results are shown in column a) of Table 2. We also created a training set where each vector appeared as above, but in both position 1 and 2. The results are shown in column b) of table 2. The results in both cases are similar to those without this additional coverage in Table 1.

**Other classification tasks** We evaluate here whether the DR units have a negative effect on other learning tasks (using  $n = 3$ ). We evaluated the networks on the classification by comparing the two vectors in the pair as binary numbers with results shown in column c) of Table 2. We also tested a task that is not a comparison of the two vectors in the pair, by calculating the digit sum. We classify by checking if the digit sum is  $\geq 3$ . In both c) and d) we see, that the performance is actually not hindered but helped by the DR units. We finally tested the task of recognising digit reversal (swapping least with most significant bits), which DR units are not designed for, as they compare corresponding digits. As we can see in column e), DR units do not deliver a perfect solution here, but still lead to somewhat better results than a plain FFNN.

## 5 Conclusions

In this study we examined the learning behaviour of feed-forward neural networks in vector equality detection and observed that the networks do not generalise well to unseen data. We also had similar results in other tasks like numeric inequality and sum of bits of binary vectors. We therefore introduced a simple modification to the network with differential rectifier (DR) units and noticed substantial improvements on unseen test data. This improvement is largely independent of vector dimension, data size and other parameters.

The question why standard FFNNs do not learn vector equality relations in a generalisable way is a relevant one, and deserves further theoretical and empirical study. It is also important to investigate the design of further measures for creating and controlling inductive biases in neural network learning, as we find that even relatively simple tasks like generalising equality require them.

## References

- [1] Moshe Leshno, Vladimir Ya Lin, Allan Pinkus, and Shimon Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural networks*, 6(6):861–867, 1993.
- [2] G. F. Marcus, S. Vijayan, S.B. Rao, and P.M. Vishton. Rule learning by seven-month-old infants. *Science*, 283, (5398):77–80, 1999.
- [3] G. F. Marcus. The algebraic mind: Integrating connectionism and cognitive science. *Cambridge MIT Press*, 2001.
- [4] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Tim Lillicrap. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, pages 4967–4976, 2017.
- [5] Luc De Raedt, Andrea Passerini, and Stefano Teso. Learning constraints from examples. In *Proceedings in Thirty-Second AAI Conference on Artificial Intelligence*, 2018.
- [6] Andres Campero, Aldo Pareja, Tim Klinger, Josh Tenenbaum, and Sebastian Riedel. Logical rule induction and theory learning using neural theorem proving, 2018.
- [7] Jeffrey Elman. Generalization, rules, and neural networks: A simulation of marcus et. al. <https://crl.ucsd.edu/elman/Papers/MVRVsimulation.html>, 1999.
- [8] Gerry Altmann and Zoltan Dienes. Technical comment on rule learning by seven-month-old infants and neural networks. In *Science*, (284(5416)):875–875, 1999.
- [9] Thomas R. Shultz and Alan C. Bale. Neural network simulation of infant familiarization to artificial sentences: Rule-like behavior without explicit rules and variables. *Infancy*, 2:4, 501-536, DOI: 10.1207/S15327078IN020407, 2001.
- [10] Marlus Vilcu and Robert F Hadley. Generalization in simple recurrent networks. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 23, 2001.
- [11] Marius Vilcu and Robert F Hadley. Two apparent ‘counterexamples’ to marcus: A closer look. *Minds and Machines*, 15(3-4):359–382, 2005.
- [12] Thomas R Shultz and Alan C Bale. Neural networks discover a near-identity relation to distinguish simple syntactic forms. *Minds and Machines*, 16(2):107–139, 2006.
- [13] Shastri and Chang. A spatiotemporal connectionist model of algebraic rule-learning. *International Computer Science Institute*, pages TR–99–011, 1999.
- [14] P. Dominey and F. Ramus. Neural network processing of natural language: Isensitivity to serial, temporal and abstract structure of language in the infant. *Language and Cognitive Processes*, pages 15(1),87–127, 2000.
- [15] Raquel G. Alhama and Willem Zuidema. Pre-wiring and pre-training: What does a neural network need to learn truly general identity rules. *CoCo at NIPS*, 2016.
- [16] Jeff Mitchell, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Extrapolation in nlp. *arXiv:1805.06648*, 2018.
- [17] Jessica B Hamrick, Kelsey R Allen, Victor Bapst, Tina Zhu, Kevin R McKee, Joshua B Tenenbaum, and Peter W Battaglia. Relational inductive bias for physical construction in humans and machines. *arXiv preprint arXiv:1806.01203*, 2018.
- [18] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.
- [19] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014.