

---

# Weight Priors for Learning Identity Relations

---

**Radha Kopparti**

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

**Tillman Weyde**

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

## Abstract

Learning abstract and systematic relations has been an open issue in neural network learning for over 30 years. It has been shown recently that neural networks do not learn relations based on identity and are unable to generalize well to unseen data. The Relation Based Pattern (RBP) approach has been proposed as a solution for this problem. In this work, we extend RBP by realizing it as a Bayesian prior on network weights to model the identity relations. This weight prior leads to a modified regularization term in otherwise standard network learning. In our experiments, we show that the Bayesian weight priors lead to perfect generalization when learning identity based relations and do not impede general neural network learning. We believe that the approach of creating an inductive bias with weight priors can be extended easily to other forms of relations and will be beneficial for many other learning tasks.

## 1 Introduction

Humans, including infants, are very effective at learning patterns of relations based on identity from sensory input and systematically applying them to new stimuli, even after a very brief exposure, while current neural networks in their standard form fail to detect these relations in unseen data [Marcus et al., 1999]. It has become evident that there are still relevant limitations to systematic generalization in current neural network architectures [Lake and Baroni, 2018, Marcus, 2018]. Neural networks are seen to be good at memorizing the numerical patterns seen in the training set but often not to extrapolate this representation outside the training set [Liska et al., 2018]. It was found that standard neural networks do not seem to learn identity relations, i.e. the equality of two vectors [Weyde and Kopparti, 2018], which are fundamental for many higher level tasks.

A well known study in this direction was conducted by Marcus et al. [1999], where a recurrent neural network failed to distinguish abstract patterns, based on equality or identity relations between the input stimuli, 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 in general, where results by Elman [1999], Altmann and Dienes [1999] and Shultz and Bale [2001] could not be reproduced by Vilcu and Hadley [2001, 2005], and Shultz and Bale [2006] disputed claims by Vilcu and Hadley [2005]. Other approaches, such as those by Shastri and Chang [1999] and Dominey and Ramus [2000], Alhama and Zuidema [2016], use specialized network architectures or different problem formulations or evaluation methods.

Recently, the Relation Based Patterns (RBP) approach has been introduced as a way to create a suitable inductive bias for the problem of learning identity relations on binary vectors Weyde and Kopparti [2018, 2019]. The task is to classify whether two halves of the input vectors are equal. i.e.  $u_i = u_{n+i} \forall i \in \{1, \dots, n\}$  for an input vector with  $2n$  dimensions. The RBP model introduced in Weyde and Kopparti [2018, 2019] is based on the comparison of input neurons that correspond to each other in a relation, e.g. the corresponding dimensions in a pair of binary vectors. For the

comparison they introduce Differentiator-Rectifier (DR) units, which calculate the absolute difference of two inputs:  $f(x, y) = |x - y|$ . For each dimension of the input vectors, a DR unit is introduced. This simplifies the learning problem because the summation of activations of the DR units is sufficient for generalisable identity detection. The results show no restriction on the learning of other tasks in practice.

There are different ways of integrating DR units into neural networks in RBP: *Early* and *Mid* Fusion. In *Early Fusion*, DR units are concatenated to input units, and in *Mid Fusion* they are concatenated to the hidden layer. In both cases, the existing input and hidden units are unchanged.

Adding units with hard-wired connections and unusual activation function limits the flexibility of the RBP approach. In this work, we introduce a modified RBP structure that can be formulated as a Bayesian prior to model the weight structure of Mid Fusion RBP in a standard feed forward network setting.

## 2 A Bayesian Approach to Relation Based Patterns

In the weight prior approach, we replace each DR units with two standard neurons and model the fixed weights with a default weight matrix  $D$ . This matrix contains the weights that enable a dimension-wise comparison of the inputs.

A diagram of the structure is given in Figure 1(a) where  $\alpha$  and  $\beta$  are the two input values being compared. A small example of this default matrix is shown in Figure 1(b). The incoming connections from two corresponding inputs to a neuron in the hidden layer to be compared have values of 1 and  $-1$ . For the same pair we use another hidden neuron with inverted signs of the weights, as in rows 1 and 2. With a ReLU activation function, this means that the value of one of the hidden neurons will be positive, if one input neurons have different activations. We therefore need at least  $n$  hidden units in the first hidden layer for a comparison of  $n$  two  $n/2$ -dimensional vectors. All other incoming connection weights are set to 0, including the bias. This ensures that in all cases where corresponding inputs are not equal, there will be a positive activation in one of the neurons.




Figure 1: a) For each input dimension there are two hidden layer neurons, e.g.  $h_1, h_2$  for dimension 1.  $\alpha$  and  $\beta$  indicate two input vectors of dimensionality  $n$ . b) The default weight matrix  $D$  for vector dim  $n=3$ . Each row corresponds to the incoming weights of a hidden neuron, e.g. the first row to  $h_1$  in figure 1 a). If there are more hidden neurons than pairs of input vectors, the additional rows contain only zeros.

The matrix  $D$  is then used to define default values of the network weights, i.e. we impose a loss based on the difference between  $D$  and the actual weight matrix  $W$ , that connects the input neurons to the first hidden layer. This loss term  $l_{RBP_{1/2}}$  for L1 or L2 loss are defined as:

$$l_{RBP_1} = \sum_{i=1}^k |W_k - D_k|, \quad l_{RBP_2} = \sum_{i=1}^k (W_k - D_k)^2, \quad (1)$$

where  $k$  is the number of elements in  $W$ . These loss functions correspond to Bayesian priors on the weights with the mean defined by the values of  $D$ . The  $L_2$  loss corresponds to a Gaussian and  $L_1$  loss to a Laplacian prior, such that backpropagation maximizes the posterior likelihood of the weights given the data Williams [1995]. The overall training loss  $l_t$  is defined as

$$l_t = l_c + \lambda \times l_{RBP} \quad (2)$$

where  $l_c$  is the cross entropy and  $\lambda$  is the regularization parameter, corresponding to the inverse of the variance of the prior, effectively regulating the strength of the RBP regularization. We call these methods of embedding the RBP into a standard network ERBP L1 and ERBP L2 respectively.

### 3 Generalizing and Learning of Identity Relations

For the task of learning identity relations, we generate synthetic data and use a standard feed-forward neural network. The input vector is binary and the target values are  $[0, 1]$  for unequal and  $[1, 0]$  for equal vector halves as described in section 1.

We use a grid search over hyper-parameters: the number of epochs  $[10, 20, 30]$ , the number of neurons per hidden layer was varied as  $[10, 20, 30]$ . For the Bayesian weight prior, we varied the regularization parameter  $\lambda$  with values  $[0.01, 0.03, 0.1, 0.3, 1, 3, 10, 30]$ . We ran a total of 10 simulations using the SGD and Adam optimizers, training for 20 epochs. We used a single hidden layer and a batch size of 1 unless indicated otherwise. The networks have been implemented using the PyTorch library<sup>1</sup>.

The train/test split was set to 75% / 25% for all tasks. We tested with vector dimensionalities  $n = 3, 10, 30$ . We generate all vectors with equal halves and take a random sample of all those with unequal halves to balance the classes. We downsample the size of the dataset to 1000 when it is greater.

#### 3.1 Identity Relations

Identity is an abstract relation in the sense that it is independent of the actual values of the individual arguments, it just depends on their combined configuration. In this task, pairs of vectors are presented to a feed-forward network and the task is to distinguish whether the two vectors are equal or not. We evaluated the performance of the network in different configurations on a held-out test set and the results are tabulated in Table 1 for different vector dimensions. As observed by Weyde and Kopparti [2018], standard networks do not improve much over random guessing, while ERBP L1 and ERBP L2, as well as Mid Fusion, almost always achieve perfect generalisation with sufficiently strong  $\lambda$  (see next section for details).

Table 1: Test set classification accuracy (in %) and standard deviation over 10 simulations (in brackets) using different models for Identity Learning (vector dimensions  $n = 3, 10, 30$ ). The networks were trained with the Adam optimizer for 20 epochs.

Type	Standard	Early Fusion	Mid Fusion	ERBP L1	ERBP L2
$n = 3$	55 (1.91)	65 (1.34)	100 (0.04)	100 (0.00)	100 (0.00)
$n = 10$	51 (1.67)	65 (1.32)	100 (0.08)	100 (0.04)	100 (0.02)
$n = 30$	50 (1.52)	65 (1.27)	100 (0.07)	100 (0.05)	100 (0.04)

#### 3.2 Parameter Variations

We study the effect of several parameters: number of hidden layers, choice of optimizer, regularization factor and weight initialization on identity relation learning tasks.

**Network Depth:** We tested identity learning with deeper neural network models, using  $h = 2, 3, 4, 5$  hidden layers. The results are tabulated in Table 2, showing only minor improvements in the network performance for deeper networks. However, ERBP L1 and L2 generalization is consistent and independent of the network depth.

**Optimiser:** We used both Stochastic Gradient Descent (SGD) and the Adam optimizer [Diederik P and Jimmy Lei, 2014] for training the ERBP L1 and L2. We observed faster convergence and greater improvement in the overall accuracy with the Adam compared to SGD. We observed similar results for both ERBP L1 and L2. Table 3 below summarizes the results of identity learning for both the optimizers with the regularization parameter  $\lambda$  set to 1. We observe that the SGD does not reach full generalization in this setting, however it does so at higher values of  $\lambda$ .

<sup>1</sup><http://pytorch.org>

Table 2: Test set classification accuracy (in %) with standard deviation (in brackets) for identity learning ( $n = 3$ ) using deeper networks. The networks were trained with the Adam optimizer for 20 epochs.

Hidden layers	No RBP	Early Fusion	Mid Fusion	ERBP L1	ERBP L2
h = 2	55 (1.65)	65 (1.26)	100 (0.02)	100 (0.00)	100 (0.00)
h = 3	55 (1.67)	67 (1.14)	100 (0.03)	100 (0.00)	100 (0.00)
h = 4	58 (1.63)	68 (1.25)	100 (0.02)	100 (0.00)	100 (0.00)
h = 5	59 (1.68)	72 (1.23)	100 (0.02)	100 (0.00)	100 (0.00)

Table 3: Accuracy (in %) and standard deviation of identity learning ( $n = 3$ ) using Adam and SGD optimizer for ERBP L1 and L2. SGD can also lead to 100% accuracy on the test set, but needs higher  $\lambda$  values.

Type	ERBP L1	ERBP L2
Adam	100 (0.00)	100 (0.00)
SGD	98 (0.06)	96 (0.05)

**Regularization Factor:** We varied the regularization factor  $\lambda$  in the loss function of the ERBP models. We observed that a factor of 3 or above reliably leads to perfect generalization in identity learning task. Figure 2, shows how the effect depends on the size the regularization factor  $\lambda$  using L1 and L2 loss functions for learning identity relations.




Figure 2: Test set classification accuracy (in %) of the network with ERBP L1 and L2 when varying the regularization parameter  $\lambda$  (shown in logarithmic scale) for identity learning with ( $n = 3$ ) and using the Adam optimizer.

### 3.3 Jointly learning identity relations and bit patterns

In order to test, whether the RBP and ERBP impedes other learning tasks, we also tested learning with some non-relational patterns that are based on values of specific neurons. The simplest form is that we have one bit that determines the target class. This case is learned with 100% generalization performance for a network with 2 additional outputs for the pattern classification.

Furthermore, we tested the learning of classification based on all even or odd elements in the vector being 0. In this case, we also get perfect generalization for  $\lambda$  up to 10. For  $\lambda = 30$  we see first deterioration of the accuracy, but there is a wide range of values where both the identity relations and the even/odd classification generalize perfectly.

## 4 Conclusions

Identity based relations are a fundamental form of relational learning. In this work, we re-visit the problem of learning identity relations using standard neural networks and show that creating a weight prior on the network weights leads to generalisable solutions of learning identity based relations. This also did not affect learning of non-relational patterns in our preliminary experiments, although more thorough testing remains to be done here. We believe that addressing these issues and coming up with effective solutions is necessary for more higher level relational learning tasks and also is relevant to address problems in general neural network learning. In future, we would like to extend this work towards learning other complex relational learning tasks and come with effective ways of creating an inductive bias using weight priors within the standard neural network architectures.

## References

- Gary Marcus, Sujith Vijayan, S. Bandi Rao, and Peter M. Vishton. Rule learning by seven-month-old infants. *Science*, 283, 5398:77–80, 1999.
- Brenden Lake and Marco Baroni. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *International Conference on Machine Learning*, pages 2879–2888, 2018.
- Gary Marcus. Deep learning : a critical appraisal. *arXiv:1801.00631*, 2018.
- Adam Liska, Germán Kruszewski, and Marco Baroni. Memorize or generalize? searching for a compositional RNN in a haystack. *CoRR*, abs/1802.06467, 2018. URL <http://arxiv.org/abs/1802.06467>.
- Tillman Weyde and Radha Kopparti. Feed-forward neural networks need inductive bias to learn equality relations. <https://arxiv.org/abs/1812.01662>, 2018.
- Jeffrey Elman. Generalization, rules, and neural networks: A simulation of marcus et. al. Technical report, 1999. URL <https://cr1.ucsd.edu/~elman/Papers/MVRVsimulation.html>.
- 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.
- 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, 2001.
- Marius Vilcu and Robert F Hadley. Generalization in simple recurrent networks. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 23:1072–1077, 2001.
- Marius Vilcu and Robert F Hadley. Two apparent ‘counterexamples’ to marcus: A closer look. *Minds and Machines*, 15(3-4):359–382, 2005.
- 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.
- Shastri and Chang. A spatiotemporal connectionist model of algebraic rule-learning. *International Computer Science Institute*, pages TR–99–011, 1999.
- Peter Dominey and Franck Ramus. Neural network processing of natural language: Sensitivity to serial, temporal and abstract structure of language in the infant. *Language and Cognitive Processes*, pages 15(1),87–127, 2000.
- 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.
- Tillman Weyde and Radha Kopparti. Modelling identity rules with neural networks. *Journal of Applied Logics*, 6(4):745–769, 2019.
- Peter M Williams. Bayesian regularization and pruning using a laplace prior. *Neural computation*, 7(1):117–143, 1995.
- Kingma Diederik P and Ba Jimmy Lei. Adam : A method for stochastic optimization. 2014. URL <https://arxiv.org/abs/1412.6980>.