



City Research Online

City, University of London Institutional Repository

Citation: Weyde, T. & Kopparti, R. M. (2019). Modelling Identity Rules with Neural Networks. *Journal of Applied Logics*, 6(4), pp. 745-769.

This is the published version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/21601/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

MODELLING IDENTITY RULES WITH NEURAL NETWORKS

TILLMAN WEYDE

*Research Centre for Machine Learning, Department of Computer Science, City,
University of London, United Kingdom*

`t.e.veyde@city.ac.uk`

RADHA MANISHA KOPPARTI*

*Research Centre for Machine Learning, Department of Computer Science, City,
University of London, United Kingdom*

`radha.kopparti@city.ac.uk`

Abstract

In this paper, we show that standard feed-forward and recurrent neural networks fail to learn abstract patterns based on identity rules. We propose Relation Based Pattern (RBP) extensions to neural network structures that solve this problem and answer, as well as raise, questions about integrating structures for inductive bias into neural networks.

Examples of abstract patterns are the sequence patterns ABA and ABB where A or B can be any object. These were introduced by Marcus et al (1999) who also found that 7 month old infants recognise these patterns in sequences that use an unfamiliar vocabulary while simple recurrent neural networks do not. This result has been contested in the literature but it is confirmed by our experiments. We also show that the inability to generalise extends to different, previously untested, settings.

We propose a new approach to modify standard neural network architectures, called Relation Based Patterns (RBP) with different variants for classification and prediction. Our experiments show that neural networks with the appropriate RBP structure achieve perfect classification and prediction performance on synthetic data, including mixed concrete and abstract patterns. RBP also improves neural network performance in experiments with real-world sequence prediction tasks.

We discuss these finding in terms of challenges for neural network models and identify consequences from this result in terms of developing inductive biases for neural network learning.

We would like to thank the anonymous reviewers of this article for their valuable comments and suggestions that helped to improve this article.

*Funded by a PhD studentship from City, University of London

1 Introduction

Despite the impressive development of deep neural networks over recent years, there has been an increasing awareness that there are some tasks that still elude neural network learning or need unrealistic amounts of data. Humans, on the other hand, are remarkably quick at learning and abstracting from very few examples. Marcus [1] showed in an experiment that 7-month old infants already recognise sequences by identity rules, i.e. which elements are repeated, after just two minutes of familiarization. In that study a simple recurrent neural network model was also tested and it failed to generalise these identity rules to new data.

In this study, we re-visit this problem and evaluate the performance of frequently used standard neural network models in learning identity rules. More specifically, we find that feed-forward and recurrent neural networks (RNN) and their gated variants (LSTM and GRU) in standard set-ups clearly fail to learn general identity rules presented as classification and prediction tasks.

We tackle this problem by proposing *Relation Based Patterns* (RBP), which model identity relationships explicitly as extensions to neural networks for classification and prediction. We show experimentally that on synthetic data the networks with suitable RBP structures learn the relevant rules and generalise with perfect classification and prediction. We also show that this perfect performance extends to mixed rule-based and concrete patterns, and that RBP improves prediction on real-world language and music data.

Identity rules are clearly in the hypothesis space of the neural networks, but the networks fail to learn them by gradient descent. We identify that both the comparison of related input neurons and of input tokens needs to be predefined in the network to learn general rules from data. The RBP structures introduce this inductive bias in the neural networks and thus enable the learning of identity rules by standard neural networks.

Our contributions in this paper are specifically:

- we evaluate several common NN architectures: feed-forward networks, RNN, GRU, and LSTM, in novel settings, and find that they fail to learn general identity rules;
- we identify reasons that prevent the learning process from being successful in this context;
- we propose the Relation Based Patterns, a new method to enable the learning of identity rules within the regular network structure;

- we show in experiments that identity rules can be learnt with RBP structure on artificial data, including mixed rule-based and concrete patterns, and that they improve performance in real-world prediction tasks;

The remainder of this paper is structured as follows. Section 2 introduces related work on modelling identity rules. Section 3 presents results of our experiments with standard neural network architectures. Section 4 presents our RBP model and its different variants. Section 5 presents the results of experiments using RBP structures. Section 6 addresses the application of RBP to mixed patterns and real data. Section 7 discusses the implications of the presented experimental results and Section 8 concludes this paper.

2 Related work

Our task is the learning of rules from sequential data. This is often seen as grammar learning, on which there have been many studies in psychology. [2] made an early contribution on implicit learning and generalisation. Subsequently, [3, 4] studied specifically the knowledge acquired during artificial grammar learning tasks.]

The specific problem we are addressing in this study is the recognition of abstract patterns that are defined by the identity relation between tokens in a sequence. In the well-known experiments by [1], infants were exposed to sequences of one of the forms *ABA* or *ABB*, e.g. ‘la di la’ or ‘la 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 tu ba’ and ‘ba tu tu’) and they showed significantly different behaviour depending on whether the sequences exhibited the form they were familiarised with or not.

This type of pattern only depends on the equality between elements of the sequence and after successful learning it should be recognisable independently of the vocabulary used. However, [1] also showed that simple recurrent Elman networks were not able to perform this learning task. This finding sparked an exchange about whether human speech acquisition is based on rules or statistics and the proposal of several neural networks models that claimed to match the experimental results. [5] and [6] proposed a solution based on a distributed representation of the input and on pre-training where the network is first trained to recognise repeated items in a sequence. The network is subsequently trained on classifying *ABA* vs *ABB* patterns. Only [6] reports specific results and has only 4 test data points, but 100% accuracy. However, [7] reported that they could not recreate these results.

[8] and [9] suggested solutions which are based on modified network architectures and training methods. [10] could not replicate the results by [9] and found that the models by [8] do not generalise. The claims by [10] were again contested by [11]. A number of other methods were suggested that used specifically designed network architectures, data representations, and training methods, such as [12, 13, 14, 15]. More recent work by [16] suggests that prior experience or pre-defined context representation (“pre-training” or “pre-wiring”) is necessary for the network to learn general identity rules when using echo state networks. While these works are interesting and relevant, they do not answer our question whether more commonly used network architectures can learn general identity rules.

The discussion of this problem is part of a wider debate on the systematicity of language learning models, which started in the 1980s and 1990s [17, 18]. This debate, like the more specific one on identity rules, has been characterised by claims and counter-claims [19, 20, 21, 22, 23, 24], which, as stated by [25], often suffer from a lack of empirical grounding. Very recently, the work in [26] has defined a test of systematicity in a framework of translation, applied it to standard *seq2seq* neural network models [27]. They found that generalisation occurs in this setting, but it depends largely on the amount and type of data shown, and does not exhibit the extraction and systematic application of rules in the way a human learner would.

In most of the studies above, the evaluation has mostly been conducted by testing whether the output of the network shows a statistically significant difference between inputs that conform to a trained abstract pattern and those that do not. From a machine learning perspective, this criterion is not satisfactory as we, like [26], would expect that an identity rule should always be applied correctly once if it has been learned from examples, at least in cases of noise-free synthetic data. We are therefore interested in the question whether and how this general rule learning can be achieved with common neural network types for sequence classification.

This question also relates to recent discussions sparked by [28] about deep neural networks’ need for very large amounts of training data, lack of robustness and lack of transparency as also expressed, e.g., by [29, 30, 31]. We surmise that these issues relate to the lack of generalisation beyond the space covered by the input data, i.e. extrapolation, which is generally seen as requiring an inductive bias in the learning system, but there is no general agreement about the nature or implementation of inductive biases for neural networks, e.g. [32, 33]. In recent years, there was a trend to remove human designed features from neural networks, and leave everything to be learned from the data [34]. We follow here the inverse approach, to add a designed internal representation, as we find that for the given problem standard neural network methods consistently fail to learn any suitable internal representation from the data.

3 Experiment 1: standard neural networks

We test different network architectures to evaluate if and to what extent recurrent and feed-forward neural networks can learn and generalise abstract patterns based on identity rules.

3.1 Supervised learning of identity rules

The problem in the experiment by [1] is an unsupervised learning task, as the infants in the experiments were not given instructions or incentives. However, most common neural network architectures are designed for supervised learning and there are also natural formulations of abstract pattern recognition as supervised learning task in the form of classification or prediction.

In our case, abstract patterns are defined by identity relations. Expressed in logic, they can be described using the binary equality predicate $eq(\cdot, \cdot)$. For a sequence of three tokens α, β, γ the rule-based patterns ABA and ABB can be described by the following rules:

$$ABA : \neg eq(\alpha, \beta) \wedge eq(\alpha, \gamma) \tag{1}$$

$$ABB : \neg eq(\alpha, \beta) \wedge eq(\beta, \gamma). \tag{2}$$

These rules are independent of the actual values of α, β , and γ and also called abstract patterns. Concrete patterns, on the other hand, are defined in terms of values of from a vocabulary a, b, c, \dots . E.g., sequences $a**$, i.e. beginning with ‘ a ’, or $*bc$, ending with ‘ bc ’, can be formulated in logic as follows:

$$a** : \alpha = 'a' \tag{3}$$

$$*bc : \beta = 'b' \wedge \gamma = 'c'. \tag{4}$$

For the remainder of this article we use the informal notations ABA and $a**$ as far as they are unambiguous in their context.

For classification, the task is to assign a sequence to a class, i.e. ABA or ABB , after learning from labelled examples. For prediction, the task is to predict the next token given a sequence of two tokens after exposure to sequences of one of the classes (e.g. only ABA , or ABB respectively). These tasks are suitable for the most commonly used neural network architectures.

3.2 Experimental set-up

Network set-up We use the Feed-forward Neural Network (FFNN) (also called Multi-layer Perceptron) [35], the Simple Recurrent Neural Network (RNN, also

called Elman network [36]), the Gated Recurrent Unit (GRU) network [37], and the Long Short Term Memory (LSTM) network [38]. For Prediction we only use the RNN and its gated variants GRU and LSTM.

The input to the networks is a one-hot encoded vector representing each token with n neurons, where n is the size of the vocabulary. In the case of the FFNN, we encode the whole sequence of 3 tokens as a vector of size $3n$. For the recurrent models, we present the tokens sequentially, each as an n -dimensional vector. We set the number of neurons in each hidden layer to (10, 20, 30, 40, 50), using 1 or 2 hidden layers. We use Rectified Linear Units (ReLUs) for the hidden layers in all networks. The output layer uses the softmax activation function. The number of output units is 2 for classification and the size of the vocabulary for prediction. We train with the Adam optimisation method [39], using initial learning rates of 0.01, 0.1, 0.2, 0.4, and train with the synthetic datasets in one batch. We use regularisation with Dropout rates of 0.1, 0.2, 0.4 and set the number of epochs to 10, within which all trainings converged.

We conduct a full grid search over all hyperparameters using four-fold cross-validation to optimise the hyperparameters and determine test results. We run a total of 10 simulations for each evaluation and average the results. All experiments have been programmed in PyTorch and the code is publicly available.¹

Datasets For performing the rule learning experiments, we artificially generate data in the form of triples for each of the experiments. We consider our sample vocabulary as $a\dots l$ (12 letters) for both prediction and classification tasks. We generate triples in all five abstract patterns: AAA, AAB, ABA, ABB, and ABC for the experiments. The sequences are then divided differently for the different cases of classification. For all the experiments we use separate train, validation, and test sets with 50%, 25%, and 25% of the data, respectively. All sampling (train/test/validation split, downsampling) is done per simulation.

3.3 Classification

First we test three different classification tasks as listed below. We use half the vocabulary for training and the other half for testing and validation (randomly sampled). We divide the sequences into two classes as follows, always maintaining an equal size of both classes:

- 1) ABA/ABB vs other: In task a) class one contains only pattern ABA while the other contains all other possible patterns (AAA, AAB, ABB, ABC) downsam-

¹<https://github.com/radhamanisha1/RBP-architecture>

pled per pattern for class balance. The task is to detect whether $eq(\alpha, \gamma) \wedge \neg eq(\alpha, \beta)$ is true or false. Analogously, the task in b) ABB vs other is to detect $eq(\beta, \gamma) \wedge \neg eq(\alpha, \beta)$. This case corresponds to the experiment in [1], where only one rule-based pattern type is used for familiarisation.

- 2) ABA vs ABB: This task is like task 1 above, but only pattern ABB occurs in the second class, so that this task has less variance in the second class. We expected this task to be easier to learn because two equality predicates $eq(\alpha, \gamma), eq(\beta, \gamma)$ change their values between the classes and are each sufficient to indicate the class.
- 3) ABC vs other: In this case, class one (ABC) has no pair of equal tokens, while the *other* class has at least one of $eq(\alpha, \beta), eq(\alpha, \gamma), eq(\beta, \gamma)$ as *true*, i.e. detecting equalities without localising them is sufficient for correct classification.

In our experiments, the training converged quickly in all cases and the classification accuracy on the training data was 100%. The results on the test set are shown in Table 1. In all cases the baseline, corresponding to random guessing is 50%. This baseline is only exceeded for task 1) by the RNNs and their gated variants, and even then the accuracy is far from perfect at 55%.

Classification task	FFNN	RNN	GRU	LSTM
1a) ABA/other	50%	55%	55%	55%
1b) ABB/other	50%	55%	55%	55%
2) ABA/ABB	50%	50%	50%	50%
3) ABC/other	50%	50%	50%	50%

Table 1: Three classification tasks based on abstract patterns over 10 simulations. The numbers show test set accuracy after a grid search and cross validation as described in section 3.2. All values are rounded to the next percentage point.

3.4 Prediction

We performed prediction experiments on two tasks. In task 1) we train and test on ABA patterns and in task 2) on ABB. Training and test/validation set use different vocabularies. The training converged quickly in less than 10 epochs, and after training the classification accuracy on the training set is 100%.

The results on the test set are shown in Table 2. The baseline is 8.3...% as we have a vocabulary size of 12. We use again half the vocabulary (6 values) for training and half for validation/testing. The results show that the tested networks

fail completely to make correct predictions. They perform below the baseline at 0% accuracy, which is mostly because they predict only tokens that appear in the training set but not in the test set.

Prediction task	RNN	GRU	LSTM
1) ABA	0%	0%	0%
2) ABB	0%	0%	0%

Table 2: Prediction results for two different abstract patterns. The numbers show test set prediction accuracy after a grid search and cross validation as described in section 3.2.

3.5 Discussion

The results show clearly that FFNNs, RNNs, GRUs and LSTMs do not learn general abstract patterns based on identity rules. This agrees with the previously reported experiments by [1]. However, since there was some conflicting evidence in the literature, the clarity of the outcome was not expected.

Questions raised This result raises the question of why these neural networks do not learn to generalise abstract patterns from data. There are two aspects worth considering for an explanation: the capacity of the network and the necessary information for the network to solve the problem.

Regarding the capacity: the solution to the task is in the hypothesis space of the neural networks, since proofs exist of universal approximation properties for feed-forward networks with unbounded activation functions [40] and of Turing-completeness for recurrent networks [41]. We will present a constructive solution below, putting that result into practice, with a design of network instances that solve the problem.

The relevant question, as has been pointed out by [16], is therefore why learning with backpropagation does not lead to effective generalisation here. There are three different steps that are necessary to detect identity rules: a comparison of input neurons, a comparison of tokens, represented by multiple neurons, and a mapping of comparison results to classes or predictions.

Vocabulary hypothesis A possible reason for the failure of the networks to generalise what we call the vocabulary hypothesis. It is based on the separated vocabulary in one-hot encoded representation. This leads to some input neurons only being activated in the training set and some only in the validation and test sets.

In order to learn suitable weights for an input comparison, there would have to be a suitable gradient of the weights of the outgoing connections from these inputs. If parts of the vocabulary do not appear in the training data, i.e. the activation of the corresponding input neurons is always zero during training, the weights of their outgoing connections will not be adapted. We therefore expect that the separation of the vocabulary prevents generalisation from the training to the test set as the weights going out from neurons that are used during testing have not been adapted by the gradient descent. Based on this consideration we conducted another experiment with a shared vocabulary.

This experiment is called ABA-BAB vs other. We again represent our vocabulary as *a...l* (12 letters) for this task with train/validation/test split as 50%/25%/25%. Now we use the same vocabulary for training, validation, and testing, but we separate different sequences of the form ABA that use the same tokens between the training and validation/test sets. E.g., if *ded* is in the test set, then *ede* is in the training or validation set, so that there is no overlap in terms of actual sequences. Like in classification experiment 1), training converged quickly and resulted in perfect classification performance on the training set.

Classification task	FFNN	RNN	GRU	LSTM
ABA-BAB vs other	50%	50%	50%	50%

Table 3: Classification results on test sets with the same vocabulary used in test, validation and training set.

The results on the test set presented in Table 3 show performance at the baseline with no evidence of generalisation. This shows that activating all inputs by using a shared vocabulary is not sufficient to enable generalisation in the learning process.

Other explanations A second potential problem is which neurons should be compared. The FFNN has no prior information about neurons belonging to the same or different tokens or about which input neurons correspond to the same token values. In the RNN, one token is presented per time step, so that a comparison between the previous hidden state and the current input is possible as the same neurons are activated. However, with a full set of connections between the previous hidden layer and the current, there is no reason that relations between the same neurons at different time steps would be processed differently from different neurons.

On the other hand, if we had a representation that includes the information of which tokens are identical or different, then we would have all the information we need for a mapping, as these are the relations in which our defining rules are

formulated (e.g. ABA is defined as $eq(\alpha, \gamma) \wedge \neg eq(\alpha, \beta)$). This idea has led to the Relation Based Pattern (RBP) model that we introduce in the next section and then evaluate with respect to its effect on both abstract and concrete pattern learning.

4 Design of Relation Based Pattern models

To address the inability of neural networks to generalise rules in neural network learning, we developed the Relation Based Pattern (RBP) model as a constructive solution, where the comparisons between input neurons and between tokens and the mappings to outputs are added as a predefined structure to the network. The purpose of this structure is to enable standard neural networks to learn abstract patterns based on the identity rules over tokens while retaining other learning abilities.

In the RBP model there are two major steps. The first step is defining comparison units for detecting identity relations, called DR units, and the second step is adding the DR units to the neural network.

4.1 Comparison units

Comparing neurons We assume, as before, that input is a one-hot encoded vector of the current token along with the $n - 1$ previous vectors for a given context length n (in this study context length 3 for classification and 2 for prediction). We use comparison units, called DR units (differentiator-rectifier). As the name suggests, they apply a full wave rectification to the difference between two inputs: $f(x, y) = |x - y|$. The first level of DR units are DR_n units that are applied to every pair of corresponding input neurons (representing the same value) within a token representation, as shown in Figure 1.

Comparing tokens The next level of DR units are the DR_p units that sum the activations of the DR_n values that belong to one pair of tokens. Based on the sequence length n and vocabulary size a we create $k = a \times n(n - 1)/2$ DR_n units for all the possible pairs of tokens and i.e. in our classification example, we have a sequence of 3 tokens and a vocabulary size of 12, i.e. $12 \times 3(3 - 1)/2 = 36 \times 3$ DR_n units. All the DR_n units for a pair of tokens are then summed in a DR_p unit using connections with a fixed weight of +1. E.g. we have $5 \times (5 - 1)/2 = 10$ DR_p units for a context of length 5. Figure 2a shows the network structure with DR_n and DR_p units.

For the prediction case, we also use the same approach to represent the difference between each input token and the next token (i.e., the target network output). We create n DR_{pout} units that calculate the difference between each input in the

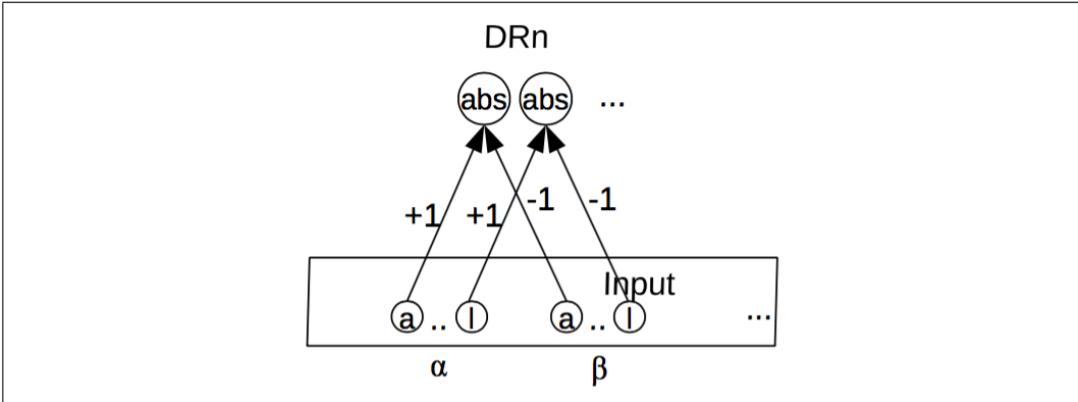


Figure 1: DR_n units comparing related inputs with an absolute of difference activation function. In one-hot encoding, there are k DR_n units for every pair of input tokens, where k is the vocabulary size.

given context and the next token. There are $k \times n$ $DR_{n,out}$ units that compare the corresponding neurons for each pair of input/output tokens, in the same way as for the pairs of input tokens. The overall network structure is shown in Figure 2b.

4.2 Neural network integration

We combine the DR units (DR_n and DR_p) with the neural network models in early, mid and late fusion approaches we call RBP1, RBP2 and RBP3, as outlined below. The weights that connect DR_n units to input and output, and the DR_n to DR_p units and the offset layer are fixed, all other weights that appear in the following models are trainable with backpropagation.

Early Fusion (RBP1n/p) In this approach, DR_n or DR_p units are added as additional inputs to the network, concatenated with the normal input. In Figure 3, the RBP1n/p structure is depicted. We use early fusion in both the prediction and classification tasks.

Mid Fusion (RBP2) The DR_p units are added to the hidden layer. Figure 4a shows the mid fusion structure for the feed-forward network and Figure 4b for the recurrent network respectively. The RBP2 approach is used for classification and prediction tasks.

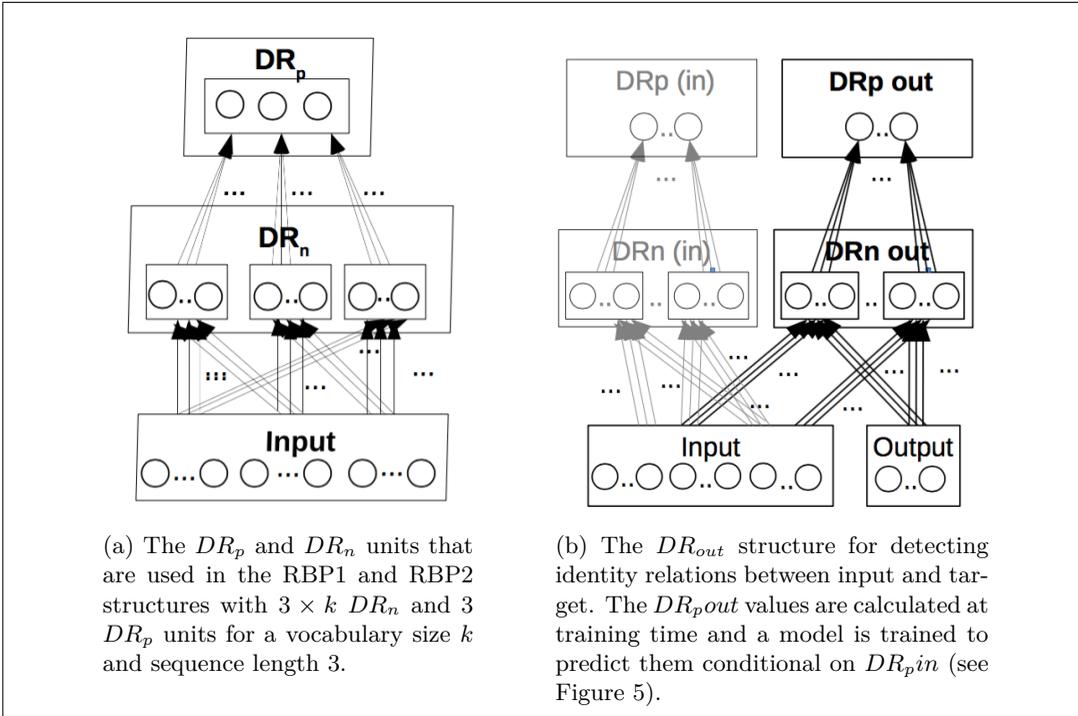


Figure 2: DR_n and DR_p units for inputs (all RBP) and outputs (RBP3).

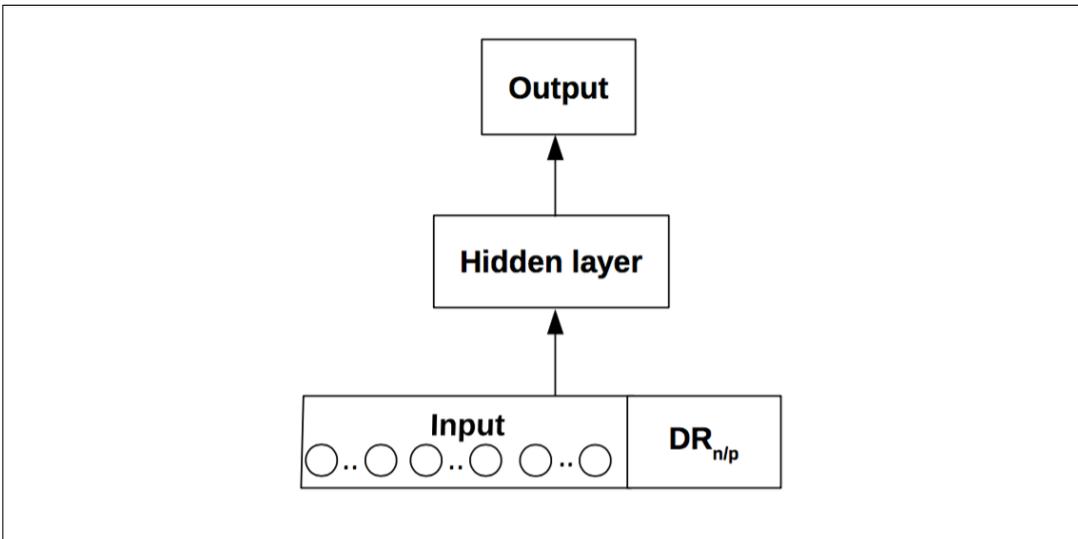


Figure 3: Overview of the RBP1n/RBP1p structure.

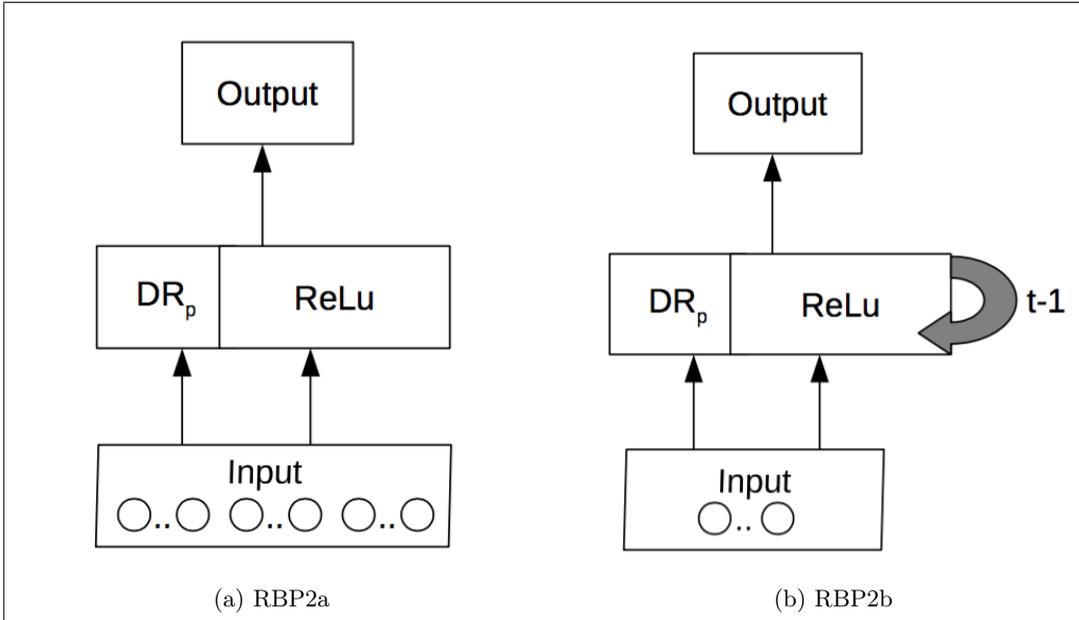


Figure 4: Overview of RBP2 approaches, where the $DR_{p,out}$ units are concatenated to the hidden layer.

Late Fusion (RBP3) In this approach, we use the same structure as in RBP2 (we call it $DR_{n,in}$ and $DR_{p,in}$ in this context), and in addition we estimate the probability of identity relations between the input and the output, i.e., that the token in the current context is repeated as the next token. We use a structure called $DR_{p,out}$ for this, and from there we project back to the vocabulary, to generate a probability offset for the tokens appearing in the context.

Figure 5 gives an overview of the RBP3 late fusion scheme. The $DR_{p,in}$ units detect identities between the input tokens in the current context as before. The $DR_{p,out}$ units model the identities between the context and the next token, as shown in the Figure 4b, where a repetition is encoded as 1, and a non-repeated token as a -1 . During training we use teacher-forcing, i.e., we set the values of the $DR_{p,out}$ units to the true values. We use a feed-forward neural network with one hidden layer to learn a mapping from the DR_{in} to the DR_{out} . This gives us an estimate of the DR_{out} units given the DR_{in} units. The DR_{out} values are then normalised subtracting the mean, and then mapped back to the output space (the one-hot vocabulary representation), using a zero value for the output values that don't appear in the input. These output offsets are then combined in a weighted sum (mixture of experts) with the output distribution estimated by the standard neural

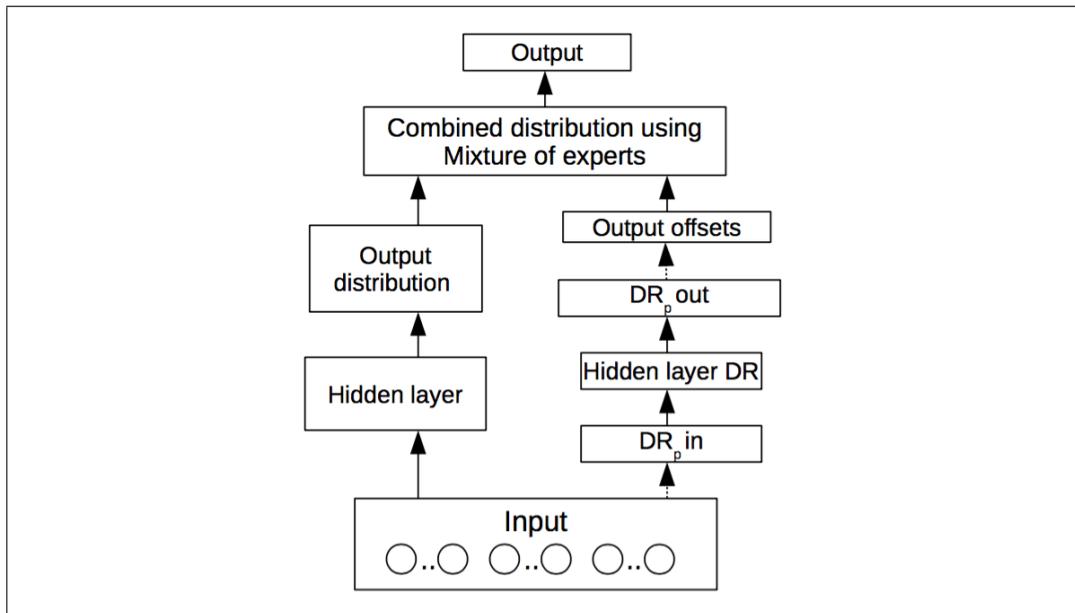


Figure 5: Overview of the RBP3 approach. The $DR_p in$ values are calculated as in RBP2. From there, we use a fully connected layer to predict $DR_p out$ (trained with teacher-forcing). The predicted $DR_p out$ values are mapped back to the vocabulary (based on the context tokens) and used as probability offsets in a mixture of experts with the standard neural network in the left part of the diagram. All connections are trainable except $Input$ to $DR_p in$ and $DR_p out$ to $Output offsets$ (dotted arrows).

network (on the left side in Figure 5). The weights in the mixture are trainable. The outputs from the combined distribution of mixture of experts are clipped between $[0,1]$ and renormalised. The final output distribution is a softmax layer providing the probability distribution over the vocabulary for the next token.

5 Experiment 2: neural networks with RBP structures

In the following we repeat the experiments from section 3 but also test networks with added RBP structures. For convenience we repeat the previous results in the tables in this section.

5.1 Classification experiments

This experiment is analogous to the first classification experiment. In the case of the feed-forward network, RBP2(a) was used in the mid fusion approach and for recurrent network, RBP2(b) was used. We trained again for 10 epochs and all networks converged to perfect classification on the training set. Table 4 provides the overall test accuracy for the three approaches.

Task	RBP	FFNN	RNN	GRU	LSTM
1a) ABA vs other	-	50%	55%	55%	55%
	RBP1n	50%	55%	55%	55%
	RBP1p	65%	70%	70%	70%
	RBP2	100%	100%	100%	100%
1b) ABB vs other	-	50%	55%	55%	55%
	RBP1n	50%	55%	55%	55%
	RBP1p	65%	70%	70%	70%
	RBP2	100%	100%	100%	100%
2) ABA vs ABB	-	50%	50%	50%	50%
	RBP1n	50%	60%	65%	65%
	RBP1p	75%	75%	75%	75%
	RBP2	100%	100%	100%	100%
3) ABC vs other	-	50%	50%	50%	50%
	RBP1n	55%	65%	65%	65%
	RBP1p	55%	70%	70%	70%
	RBP2	100%	100%	100%	100%
4) ABA-BAB vs other	-	50%	50%	50%	50%
	RBP1n	55%	72%	75%	75%
	RBP1p	69%	74%	75%	76%
	RBP2	100%	100%	100%	100%

Table 4: Classification experiments with RBP: test accuracy for the different models and tasks, as explained above. Results with ‘-’ in the RBP column are the same as in section 3 and shown here again for comparison.

The results with RBP1n models already show some improvement over the baseline in most configurations, but the results are only slightly above the standard networks, with RNNs, GRUs and LSTMs benefiting more than FFNN. This supports our hypothesis that learning to compare corresponding input neurons is a challenging task for neural networks. However, the results show that providing that comparison

is not sufficient for learning identity rules.

RBP1p structures also aggregate all the l DR_n neurons that belong to a pair of input tokens. The results show that providing that information leads to improved accuracy and provide evidence that this aggregation is another necessary step that the networks do not learn reliably from the data.

The RBP2 models enable the neural networks to make predictions and classifications that generalise according to identity rules that it learns from data. The RBP2 leads to perfect classification for all network types tested. This confirms the design consideration that comparing pairs of tokens provides the relevant information in the form required for classification, as the classes are defined by *equals* relations, so that the activations of the DRp units are directly correlated with the class labels.

A surprising result is the big difference between the generalisation using the RBP1p and the RBP2 structures. They both provide the same information, only in different layers of the network, but RBP1p only reaches at most 75% with a 50% baseline. We hypothesize that the additional expressive power provided by the nonlinearities in the hidden layer here hinders effective learning. This effect deserves further investigation.

5.2 Prediction experiments

Here we performed two experiments separately on ABA and ABB patterns as in experiment 1 on prediction. The tasks are the same as previously and we trained again for 10 epochs after which all networks had converged to perfect prediction accuracy on the training data. Table 5, summarises the accuracy for RNN, GRU and LSTM without RBP, and with RBP1n, 1p, 2, and 3.

Overall, we observe that only the LSTM benefits from RPB1n, RBP1p, and RBP2 structures, all other networks can apparently not make use of the information provided. The RBP3 model, on the other hand, leads to perfect classification on our synthetic dataset.

Our interpretation is that standard recurrent networks do not learn the more complex mapping that prediction requires, as not only recognition of a pattern but also selecting a prediction on the basis of that pattern is required. The somewhat better results of the LSTM networks are interesting. In the RBP3 model, the mapping between the identity patterns and back to the vocabulary adds considerable prior structure to the model and it is very effective in achieving the generalisation of rule-based patterns.

Pattern	RBP	RNN	GRU	LSTM
1) ABA	-	0%	0%	0%
	RBP1n	0%	0%	16%
	RBP1p	0%	0%	18%
	RBP2	0%	0%	20%
	RBP3	100%	100%	100%
2) ABB	-	0%	0%	0%
	RBP1n	0%	0%	17%
	RBP1p	0%	0%	20%
	RBP2	0%	0%	22%
	RBP3	100%	100%	100%

Table 5: Test set accuracy in prediction experiments for patterns ABA and ABB. As before, results are averaged over 10 simulations and rounded to the nearest decimal point. Results with ‘-’ in the RBP column are the same as in section 3 and shown here again for comparison.

6 Experiment 3: mixed tasks and real data

The results presented here were all obtained with synthetic data where classification was exclusively on rule-based abstract patterns. This raises the question whether the RBP will impede recognition of concrete patterns in a mixed situation. Furthermore, we would like to know whether RBP is effective with real data where the abstract and concrete patterns may interact.

6.1 Mixed abstract and concrete patterns

We conducted an experiment where the classes were defined by combinations of abstract and concrete patterns. Specifically we defined 4 classes based on the abstract patterns *ABA* and *ABB* combined with the concrete patterns *a*** and *b***. E.g., the class *ABA, a*** can be expressed logically as

$$eq(\alpha, \gamma) \wedge \neg eq(\alpha, \beta) \wedge \alpha = 'a'. \quad (5)$$

We use a vocabulary of 18 characters, out of which 12 are used for training and 6 are used for validation/testing in addition to ‘a’ and ‘b’, which need to appear in all sets because of the definition of the concrete patterns. For class 1/3 and class 2/4, abstract patterns ABA and ABB are used respectively. Class 1/2 and 3/4 start with tokens ‘a’ and ‘b’ respectively. The train, validation and test split is 50%, 25%, and 25% respectively. We trained the network for 10 epochs, leading to perfect

classification on the training set. A total of 10 simulations has been performed. We test a feed forward and a recurrent neural network without and with RBP1p and RBP2. The results are shown in Table 6.

RBP	FFNN	RNN
-	23%	42%
RBP1p	49%	57%
RBP2	100%	100%

Table 6: Test set accuracy for mixed abstract/concrete pattern classification.

As in the previous experiments, networks without RBP fail to generalise the abstract patterns. The results for RBP1p and RBP2 show, that the ability to learn and recognise the concrete patterns is not impeded by adding the RBP structures.

6.2 Language models with RBP

In order to test the capability of networks with RBP structure, we use them in two language modelling tasks. One is to predict characters in English text, and one is to predict the pitch of the next note in folk song melodies. We selected both tasks because of the prevalence of repetitions in the data, as notes in music and characters in English tend to be repeated more than words. Our RBP structures are designed to model identity-rules and we therefore expect them to be more effective on tasks with more repetitions.

Character prediction We conducted a character prediction experiment on a subset of the Gutenberg electronic book collection², consisting of text with the dataset size of 42252 words. We used 2 hidden layers with 50 neurons each. In the RBP2 model, the DRp units were concatenated with the first hidden layer. The learning rate is set to 0.01 and the network training converged after 30 epochs. Each character is predicted without and with the RBP variants using a context size of 5. The prediction results are summarized in Table 7.

Pitch prediction In another experiment we applied RBP to pitch prediction in melodies [42] taken from the Essen Folk Song Collection [43]. We performed a grid search for each context length for hyper parameter tuning, with [10,30,50,100] as the size of the hidden layer and [30,50] epochs with learning rate set to 0.01 with one hidden layer. The results for context length 5 are summarized in Table 8. RBP

²<https://www.gutenberg.org/>

RBP	RNN	GRU	LSTM
-	3.8281	3.8251	3.8211
RBP1p	4.4383	4.4368	4.4321
RBP2	3.7512	3.7463	3.7448
RBP3	3.4076	3.4034	3.4012

Table 7: Character prediction task. The numbers show the average cross entropy loss per character on the test set (lower is better, best values are set in bold), without and with RBP structures using context length 5.

improved the network performance for RNN, GRU, and LSTM. Overall, LSTM with late fusion produces the best result and also improves over the best reported performance in pitch prediction with a long-term model on this dataset with a cross-entropy of 2.547, which was achieved with a feature discovery approach by [44].

RBP	RNN	GRU	LSTM
-	2.6994	2.5702	2.5589
RBP1p	2.6992	2.5714	2.5584
RBP2	2.6837	2.5623	2.5483
RBP3	2.6588	2.5549	2.5242

Table 8: Pitch prediction task on the Essen Folk Song Collection. The numbers show the average cross entropy per note (lower is better, best values are set in bold), without and with RBP using context length 5.

Results In both character and pitch prediction, the addition of RBP3 structures improves the overall results consistently. RBP1n leads to a deterioration in character prediction and to inconsistent effect on pitch prediction, while RBP2 leads to a slight but consistent improvement in both tasks. This provides further evidence that the RBP structure enables the learning of relevant patterns in the data.

7 Discussion

7.1 Standard neural networks

The results of the experiments described above confirm the results of [1] and others that standard recurrent (and feed-forward) neural networks do not learn generalisable identity rules. From the tested models and settings of the task we can see that

the lack of activation of input neurons impedes learning, but avoiding this lack is not sufficient. The task assumes that the identity of input tokens is easy to recognise, classify and base predictions on, but the models we tested do not learn to generalise in this way. These results confirm our view that in order to generalise it is necessary to know which input neurons are related, similarly on the next level, which comparisons of input belong to a pair of tokens so that they can be aggregated per token. The structure of neural networks does not provide any prior preference for inputs that are related in this way over any other combinations of inputs. This makes it seem plausible that the solutions by [5, 6, 9] could not be replicated by [7, 10].

7.2 Constructive model with RBP

The RBP model addresses the learning of identity rules by adding neurons and connections with fixed weights. From the input neurons we add connections to a DR (differentiator-rectifier) unit from each pair of corresponding input neurons within any pair of tokens (represented in one-hot encoding). These DR units calculate the absolute of the difference of the activations of the two input neurons. They are followed by *DRp* units that aggregate by taking the sum of the DR unit activations for each pair of tokens. The fact that the DRp units relate to the difference between each pair of neurons makes the learning task for classification much simpler, as has been confirmed by our results. An open question in this context is why the RBP2 is so much more effective than RBP1p for classification, although the only difference is the layer in which the information is added into the network.

For prediction, we need a more complex structure, as beyond recognition of identity, also the selection of the token to predict is required, that depends on the tokens in the context and their similarity relations. The constructive RBP solution requires a transformation into a representation of identity relations in the input that is mapped to identities between input and output and that is mapped back to the token space by adding prediction probability to the tokens that are predicted to be identical between input and output. This created a complex predefined structure, but without it even the models that achieved perfect classification failed to make correct predictions with new data. Only the LSTM models could use the RBP1 and RBP2 information to make prediction above the baseline (22% vs 8.3%). We hypothesise that the gating structure of the LSTMs enables at least some effective mapping. The 100% correct predictions by all models using RBP3 shows the effectiveness of this structure.

7.3 Applications

Adding a bias into the network with a predefined structure such as RBP raises the question whether there is a negative effect on other learning abilities of the network and whether interactions between the abstract and concrete tasks can be learnt. In the mixed pattern experiment, RBP is still effective and showed no negative effect. In experiments with real language and music data we found that RBP3 has a positive effect on the prediction of characters in language and pitches in folk song melodies. The small negative effect of RBP1 on character prediction seems to indicate that there may be confounding effect where identity rules are less relevant. This effect did not appear in melody prediction, where repetition is more important.

7.4 Extrapolation and inductive bias

The results in this study confirm that an inductive bias is needed for extrapolation, in the terminology of [28], in order to generalise in some sense outside the space covered by the training data. This general challenge has recently attracted some attention. E.g., [45] provided several solutions to the related problem of learning equality of numbers (in binary representation), which does not generalise from even to odd numbers as pointed out already by [46]. As the authors point out in [45], an essential question is which biases are relevant to the domain and problem. The identity problem addressed here is in itself fundamental to learning about relations [47], as relations depend on object identity. This further raises the question what is needed to enable more complex concepts and rules to be learnt, such as more general logical concepts and rules.

The identity rules also point to the lower-level problem that the natural relations of position and belonging to objects are not naturally addressed in neural networks. Other tasks may require different structures, relating for example to arithmetics, geometry or physics [48]. We therefore see as an important task the definition or predefined structures in neural networks, so that they create useful inductive bias, but do not prevent learning of functions that do not conform to that bias.

8 Conclusions

Our experiments show that the observation by [1], that neural networks are unable to learn general identity rules, holds for standard feed-forward networks, recurrent neural networks, and networks of GRUs and LSTMs. The solution we propose here, the Relation Based Patterns (RBP), introduce an additional structure with fixed weights into the network. Our experiments confirm that the RBP structures

enable the learning of abstract patterns based on identity rules in classification and prediction as well as in mixed abstract and concrete patterns. We have further found that adding RBP structures improves performance in language and music prediction tasks.

Overall, we find that standard neural networks do not learn identity rules and that adding RBP structure creates an inductive bias which enables this extrapolation beyond training data with neural networks. This outcome raises the question on how to develop further inductive biases for neural networks to improve generalisation of learning on other tasks and more generally.

References

- [1] G. F. Marcus, S. Vijayan, S. Rao, P. Vishton, Rule learning by seven-month-old infants, *Science*, 283 283 (5398) (1999) 77–80.
- [2] P. C. Gordon, K. J. Holyoak, Implicit learning and generalization of the mere exposure effect., *Journal of Personality and Social Psychology*, 45, 492–500.
- [3] B. J. Knowlton, L. R. Squire, The information acquired during artificial grammar learning., *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 79–91.
- [4] B. J. Knowlton, L. R. Squire, Artificial grammar learning depends on implicit acquisition of both abstract and exemplar-specific information, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 169–181.
- [5] M. Seidenberg, J. Elman, Do infants learn grammar with algebra or statistics?, *Science* 284 (5413) (1999) 433–433.
- [6] J. Elman, Generalization, rules, and neural networks: A simulation of Marcus et. al, <https://crl1.ucsd.edu/~elman/Papers/MVRVsimulation.html>.
- [7] M. Vilcu, R. F. Hadley, Generalization in simple recurrent networks, in: *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 23, 2001, pp. 1072–1077.
- [8] T. R. Shultz, A. 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.
- [9] G. Altmann, Z. Dienes, Rule learning by seven-month-old infants and neural networks, In *Science* 284 (5416) (1999) 875–875.
- [10] M. Vilcu, R. F. Hadley, Two apparent ‘counterexamples’ to Marcus: A closer look, *Minds and Machines* 15 (3-4) (2005) 359–382.
- [11] T. R. Shultz, J.-P. Thivierge, D. Titone, Generalization in a model of infant sensitivity to syntactic variation, in: *Proceedings of the Annual Meeting of the Cognitive Science Society*, 2005, pp. 2009–20014.
- [12] L. Shastri, S. Chang, A spatiotemporal connectionist model of algebraic rule-learning, Tech. Rep. TR-99-011, Berkeley, California: International Computer Science Institute

- (1999).
- [13] M. Gasser, E. Colunga, Babies, variables, and connectionist networks, in: Proceedings of the 21st Annual Conference of the Cognitive Science Society, Lawrence Erlbaum, 1999, p. 794.
 - [14] P. F. Dominey, F. Ramus, Neural network processing of natural language: I. sensitivity to serial, temporal and abstract structure of language in the infant, *Language and Cognitive Processes* 15 (1) (2000) 87–127.
 - [15] M. H. Christiansen, C. M. Conway, S. Curtin, A connectionist single-mechanism account of rule-like behavior in infancy, in: Proceedings of the 22nd annual conference of the cognitive science society, 2000, pp. 83–88.
 - [16] R. G. Alhama, W. Zuidema, Pre-wiring and pre-training: What does a neural network need to learn truly general identity rules, *CoCo at NIPS*.
 - [17] J. A. Fodor, Z. W. Pylyshyn, Connectionism and cognitive architecture: A critical analysis, *Cognition* 28 (1-2) (1988) 3–71.
 - [18] R. F. Hadley, Systematicity in connectionist language learning, *Mind & Language* 9 (3) (1994) 247–272.
 - [19] P. Smolensky, The constituent structure of connectionist mental states: A reply to Fodor and Pylyshyn, *Southern Journal of Philosophy* 26 (Supplement) (1987) 137–161.
 - [20] J. Fodor, B. P. McLaughlin, Connectionism and the problem of systematicity: Why Smolensky’s solution doesn’t work, *Cognition* 35 (2) (1990) 183–204.
 - [21] D. Chalmers, Why Fodor and Pylyshyn were wrong: The simplest refutation, in: Proceedings of the Twelfth Annual Conference of the Cognitive Science Society, Cambridge, Mass, 1990, pp. 340–347.
 - [22] L. Niklasson, T. van Gelder, Systematicity and connectionist language learning, *Mind and Language* 9 (3) (1994) 28–302.
 - [23] M. H. Christiansen, N. Chater, Generalization and connectionist language learning, *Mind & Language* 9 (3) (1994) 273–287.
 - [24] R. F. Hadley, Systematicity revisited: reply to Christiansen and Chater and Niklasson and van Gelder, *Mind & Language* 9 (4) (1994) 431–444.
 - [25] S. L. Frank, Getting real about systematicity, in: P. Calvo, J. Symons (Eds.), *The architecture of cognition: Rethinking Fodor and Pylyshyn’s systematicity challenge*, MIT Press, 2014, pp. 147–164.
 - [26] B. Lake, M. Baroni, Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks, in: *International Conference on Machine Learning*, 2018, pp. 2879–2888.
 - [27] I. Sutskever, O. Vinyals, Q. V. Le, Sequence to sequence learning with neural networks, in: *Advances in neural information processing systems*, 2014, pp. 3104–3112.
 - [28] G. F. Marcus, Deep learning : a critical appraisal, arXiv:1801.00631.
 - [29] S. Sabour, N. Frosst, G. E. Hinton, Dynamic routing between capsules, in: *Advances in Neural Information Processing Systems*, 2017, pp. 3856–3866.
 - [30] K. Kansky, T. Silver, D. A. Mély, M. Eldawy, M. Lázaro-Gredilla, X. Lou, N. Dorfman,

- S. Sidor, S. Phoenix, D. George, Schema networks: Zero-shot transfer with a generative causal model of intuitive physics, arXiv preprint arXiv:1706.04317.
- [31] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, R. Fergus, Intriguing properties of neural networks, CoRR abs/1312.6199. arXiv:1312.6199. URL <http://arxiv.org/abs/1312.6199>
- [32] R. Feinman, B. M. Lake, Learning inductive biases with simple neural networks (2018). arXiv:1802.02745.
- [33] D. G. T. Barrett, F. Hill, A. Santoro, A. S. Morcos, T. Lillicrap, Measuring abstract reasoning in neural networks (2018). arXiv:1807.04225.
- [34] Y. Bengio, A. Courville, P. Vincent, Representation learning: A review and new perspectives, IEEE transactions on pattern analysis and machine intelligence 35 (8) (2013) 1798–1828.
- [35] D. E. Rumelhart, G. E. Hinton, R. J. Williams, Learning internal representations by error propagation, Tech. rep., California Univ San Diego La Jolla Inst for Cognitive Science (1985).
- [36] J. L. Elman, Finding structure in time, Cognitive science 14 (2) (1990) 179–211.
- [37] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using rnn encoder-decoder for statistical machine translation, arXiv preprint arXiv:1406.1078.
- [38] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural Computation 9 (8) (1997) 1735–1780. doi:10.1162/neco.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>
- [39] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv:1412.6980.
- [40] M. Leshno, V. Y. Lin, A. Pinkus, S. Schocken, Multilayer feedforward networks with a nonpolynomial activation function can approximate any function, Neural networks 6 (6) (1993) 861–867.
- [41] H. Siegelmann, E. Sontag, On the computational power of neural nets, Journal of Computer and System Sciences (1995) Volume 50, Issue 1, pp. 132–150.
- [42] R. M. Kopparti, T. Weyde, Evaluating repetition based melody prediction over different context lengths, ICML Joint Workshop on Machine Learning and Music, Stockholm, Sweden, July 14.
- [43] H. Schaffrath, The essen folksong collection in the humdrum kern format, Menlo Park, CA, Centre for Computer Assisted Research in the Humanities, 1995.
- [44] J. Langhabel, R. Lieck, M. Rohrmeier, Feature discovery for sequential prediction of monophonic music, International Society for Music Information Retrieval Conference (2017) 649–655.
- [45] J. Mitchell, P. Minervini, P. Stenetorp, S. Riedel, Extrapolation in NLP, arXiv:1805.06648.
- [46] G. F. Marcus, The algebraic mind: Integrating connectionism and cognitive science, Cambridge MIT Press.
- [47] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Ma-

- linowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, et al., Relational inductive biases, deep learning, and graph networks, arXiv preprint arXiv:1806.01261.
- [48] N. Cohen, A. Shashua, Inductive bias of deep convolutional networks through pooling geometry (2016). [arXiv:1605.06743](#).