

pix2rule: End-to-end Neuro-symbolic Rule Learning

Nuri Cingillioglu
nuric@imperial.ac.uk

Alessandra Russo
a.russo@imperial.ac.uk

Imperial College
London



Abstract

Humans have the ability to seamlessly combine low-level visual input with high-level symbolic reasoning often in the form of recognising objects, learning relations between them and applying rules. Neuro-symbolic systems aim to bring a unifying approach to connectionist and logic-based principles for visual processing and abstract reasoning respectively. This paper presents a complete neuro-symbolic method for processing images into objects, learning relations and logical rules in an end-to-end fashion. The main contribution is a differentiable layer in a deep learning architecture from which symbolic relations and rules can be extracted by pruning and thresholding. We evaluate our model using two datasets: subgraph isomorphism task for symbolic rule learning and an image classification domain with compound relations for learning objects, relations and rules. We demonstrate that our model scales beyond state-of-the-art symbolic learners and outperforms deep relational neural network architectures.


Semi-symbolic Layer

If we utilise t-norms to implement fuzzy logic, as the number of inputs increases, the operation suffers from vanishing gradients and becomes unviable for downstream layers such as CNNs. This phenomenon occurs due to the 1 out of n failure or success characteristic of conjunction and disjunction respectively. Hence, we are interested in an operation that does not starve gradients and **eventually** converges to the desired semantics.

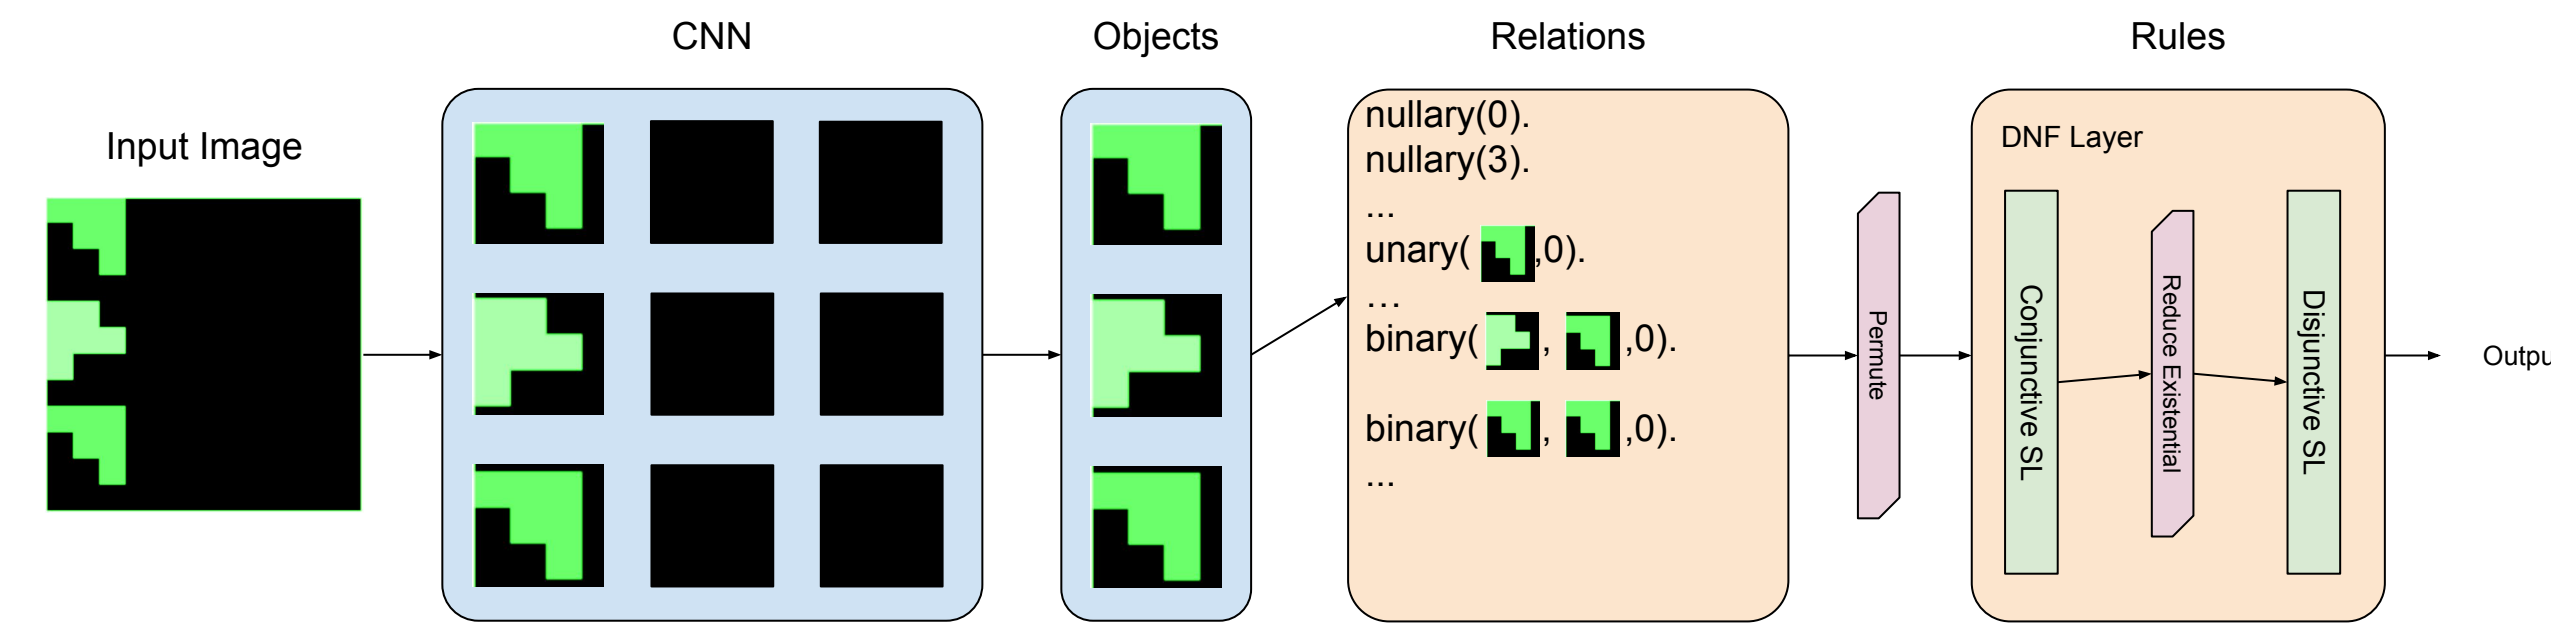
$$y = f\left(\sum_i w_i x_i + \beta\right) \quad (1)$$

$$\beta = \delta(\max_i |w_i| - \sum_i |w_i|) \quad (2)$$

Delta is the semantic gate selector ranging from 0 to 1 for conjunction and -1 for disjunction. Negation naturally is implemented as multiplicative inverse and we select f to be tanh. The semi-symbolic layer thus can **gradually switch** between a linear layer and a logical gate.



We construct a disjunctive normal form layer (DNF) by stacking two semi-symbolic layers, one conjunctive and one disjunctive. In order to learn first order rules with variables, we curate all permutations of object to variable binding, grounding the logic program. We use the **max** operator to reduce existential variables.



For the full neuro-symbolic model, we combine the DNF layer with an upstream Convolutional Neural Network, object selection and relational layers. The entire model is **end-to-end differentiable** and the object representations, their relations and rules are learnt together. For symbolic tasks we just use the DNF layer as a standalone model.

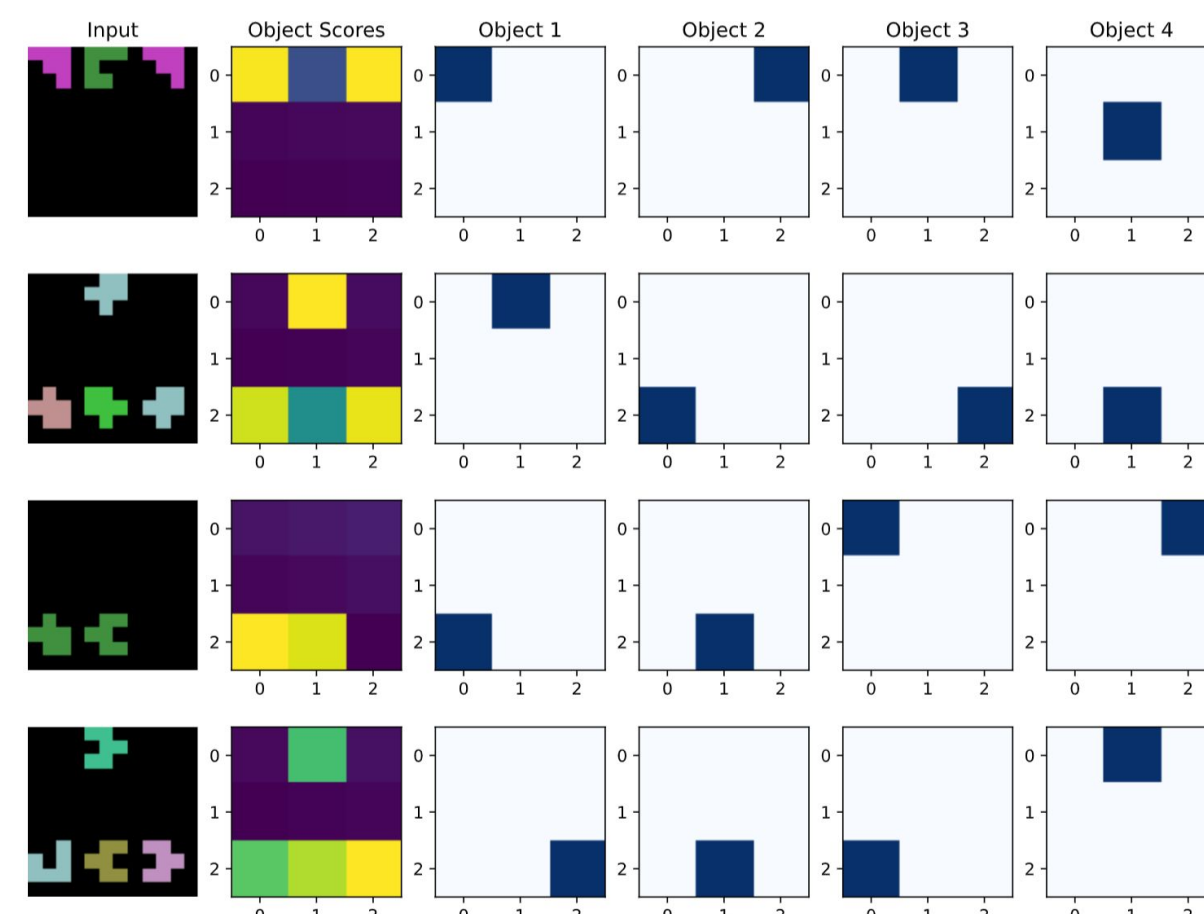
Experiments

Difficulty	Test Accuracy			Training Time		
	Easy	Medium	Hard	Easy	Medium	Hard
DNF	1.0±0.0	1.00±0.0	1.00±0.00	127.13± 4.17	136.90±10.00	129.56± 5.77
DNF+t	1.0±0.0	0.99±0.0	0.99±0.01	125.02± 6.53	135.67± 8.56	143.67±22.60
FastLASv3	1.0±0.0			29.85± 0.69		
ILASP-2i	1.0±0.0			3336.00±994.45		

For the subgraph set isomorphism problem, we create three difficulties in which the rule lengths increase. We compare against two state-of-the-art rule learners: ILASP [1] and FastLAS [2]. Our DNF model scales better to longer rules. DNF+t is with thresholded weights giving symbolic rules.

Set	Task Model	All			Between			Occurs			Same			XOccurs		
		100	1000	5000	100	1000	5000	100	1000	5000	100	1000	5000	100	1000	5000
Hex.	DNF	0.94	0.97	0.98	0.90	0.99	0.99	0.56	0.99	0.99	0.94	1.00	1.00	0.49	0.80	0.93
	DNF-h	0.98	0.99	0.99	0.95	1.00	0.99	0.62	0.99	0.99	0.97	1.00	1.00	0.50	0.98	0.96
	DNF-h+t	0.51	0.55	0.92	0.91	0.97	0.98	0.50	0.79	0.96	0.53	1.00	0.98	0.48	0.51	0.51
	DNF-hi	0.98	0.99	1.00	0.97	0.99	1.00	0.69	0.99	0.99	0.97	1.00	1.00	0.51	0.99	0.99
	DNF-r	0.94	0.98	0.98	0.84	1.00	0.99	0.63	0.99	0.99	0.96	1.00	1.00	0.51	0.94	0.51
	PrediNet	0.85	0.95	0.96	0.66	0.99	0.99	0.57	0.95	0.97	0.99	1.00	1.00	0.50	0.58	0.95

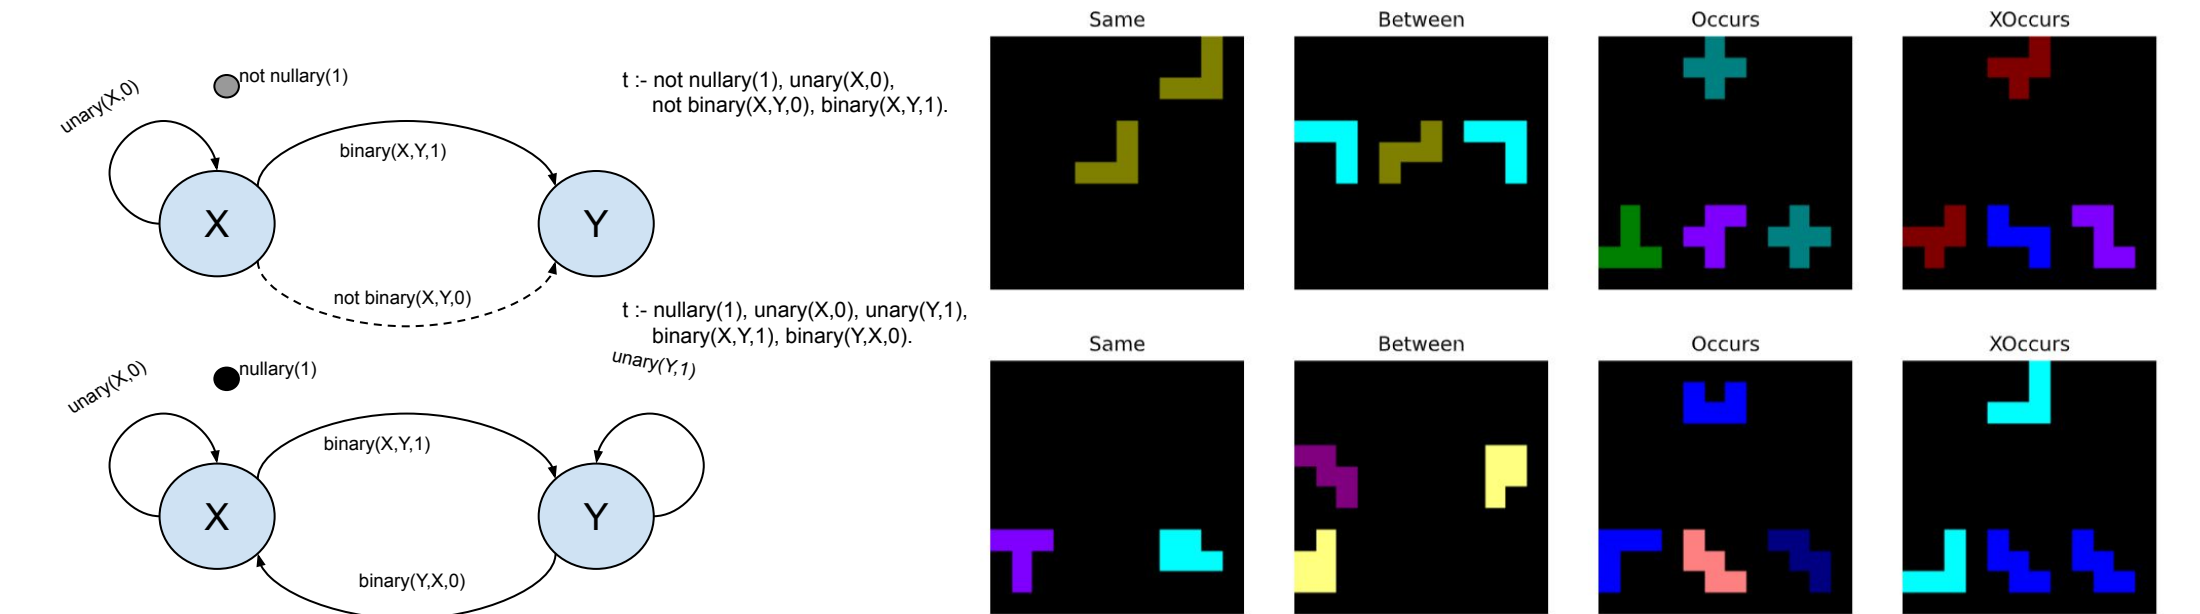
For the full neuro-symbolic pipeline, we compare against PrediNet [3], a state-of-the-art deep relational neural network. We demonstrate that our model is more data efficient if not competitive against PrediNet.



When we plot the selected objects from the relations game dataset, we observe that the model learns to select image patches that correspond to shapes. If there are extra slots, then blank patches are selected.


Datasets

We use two datasets: (i) **Subgraph set isomorphism** requires a model to decide whether any subgraph of a given graph is isomorphic to a set of other graphs and (ii) **Relations Game dataset** consists of an input image with different shapes and colours exhibiting compound relations.




Analysis

Trained on the between task with 1k training examples, the following symbolic rule can be used to solve it. The last argument is the predicate ID. This rule achieves 97% test accuracy.



We remove one atom at a time and compute the drop in accuracy. We then plot the truth cases for the most important atom. For the example rule, **binary(X,Y,6) drops accuracy by 18%**. Object arguments (top and bottom rows) that make binary(X,Y,6) true and false exhibit **no common pattern** to principal concepts of the dataset.



Limitations

- The rule learning requires a ground logic program which may not scale well to large number of constants.
- The model may overfit and under-utilise explicit object, relation and rule separation, e.g. learn a rule with just a single relation.

References

- [1] Mark Law et al. The ILASP system for Inductive Learning of Answer Set Programs. ArXiv, 2020. arXiv:2005.00904.
- [2] Mark Law et al. FastLAS: Scalable Inductive Logic Programming Incorporating Domain-Specific Optimisation Criteria. AAAI, 2020.
- [3] Murray Shanahan et al. An Explicitly Relational Neural Network Architecture. ICML, 2020.

