Factorised neural relational inference for multi-interaction systems

Ezra Webb*[1], Ben Day*[2], Helena Andres-Terre[2] and Pietro Lió[2]
1. Department of Physics, 2. Department of Computer Science & Technology, University of Cambridge
Cambridge, United Kingdom.
ben.day@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

1. Introduction

Many complex natural and cultural phenomena are well modelled by systems of simple interactions between particles. A number of architectures have been developed to articulate this kind of structure, both implicitly and explicitly. We consider an unsupervised explicit model, the NRI model (Kipf et al., 2018), and make a series of **representational adaptations** and **physically motivated changes**.

Most notably, the inferred latent interaction graph is **factorised into a multiplex network**, allowing each layer-graph to encode a different interaction-type.

This fNRI model has fewer parameters and **significantly outperforms the original in both edge and trajectory prediction**, establishing a new state-of-the-art.

fNRI vs. sfNRI

The original formulation interaction graph is a single-layer graph with many edge types. By factorising we change to having many graphs but a choice remains as to whether to continue to use many edge types or to have a single edge type (i.e. present/absent only). We refer to these two models as the fNRI and sfNRI, respectively, as the **s**fNRI uses a sigmoid activation function over the edges.

Example trajectories

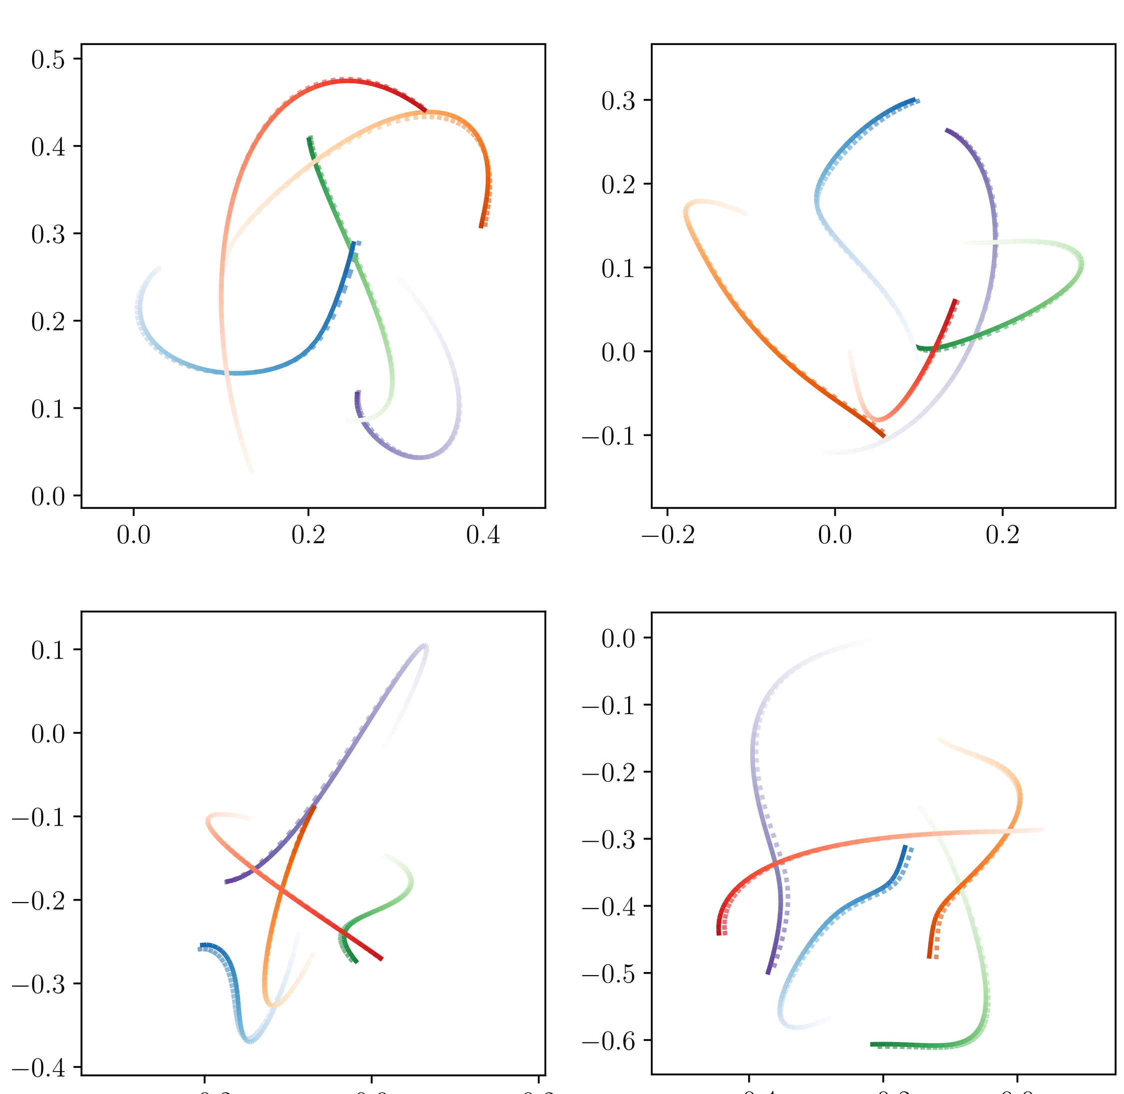

Figure 3: In these examples the interaction graph has been inferred correctly by the fNRI model. Dotted lines are ground truth trajectories, solid fading lines are the model predictions with lines becoming fainter into the past. To make these predictions the model only receives a starting configuration and infers all 50 time steps shown here without further access to the ground truth (i.e. there is no forcing).

References

Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. Neural Relational Inference for Interacting Systems. 2 2018. URL http://arxiv.org/abs/1802.04687.

2. Factorisation

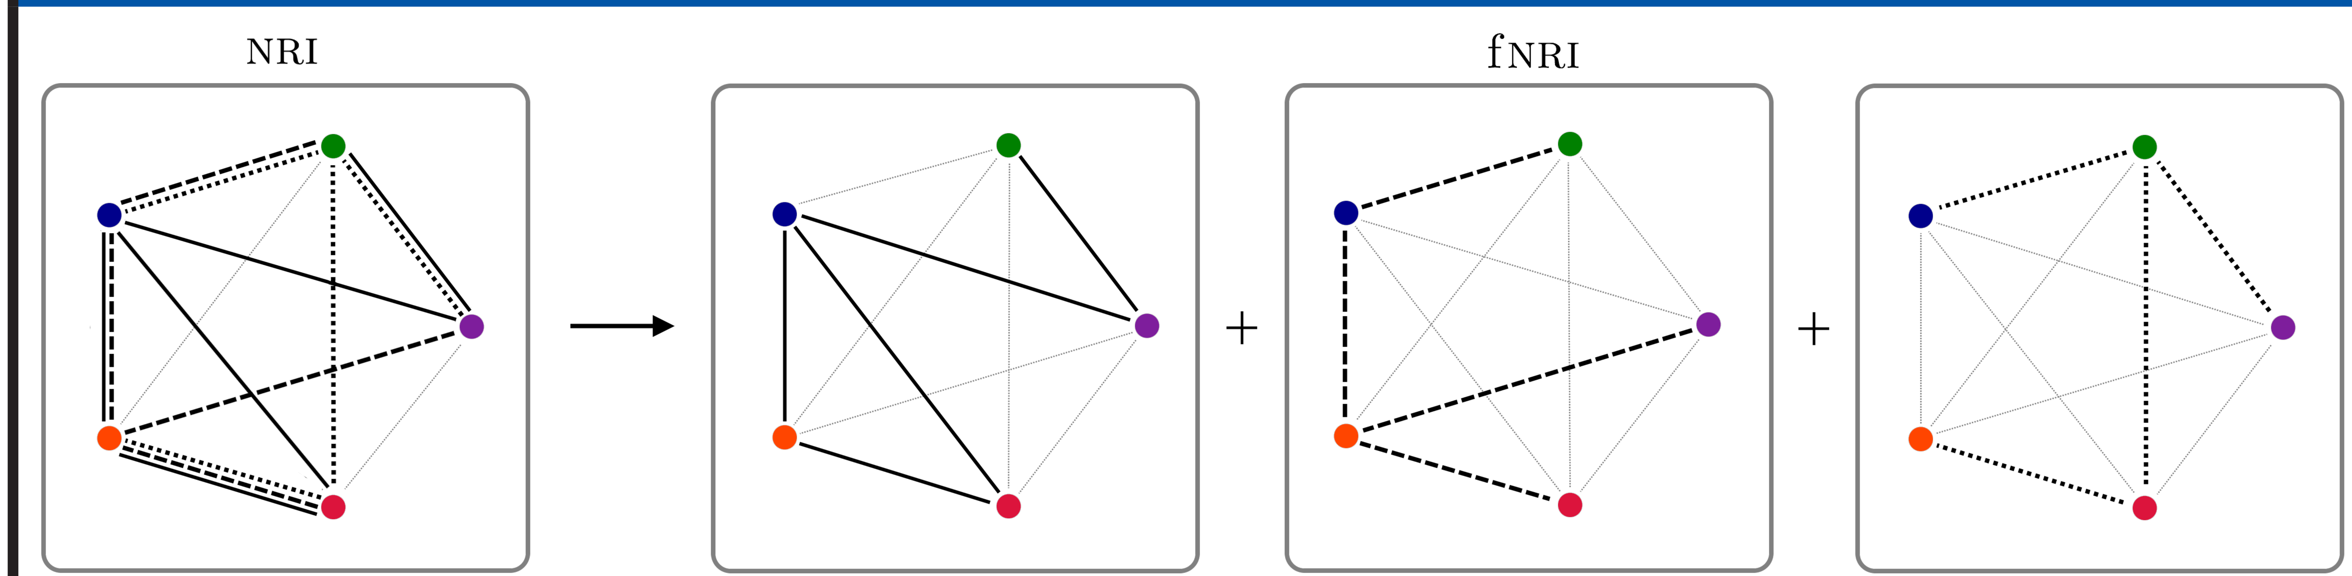

Figure 1: Schematic showing the representational change in the interaction graph between the NRI and fNRI models when there are three independent interaction types represented by solid, dashed and dotted lines, in addition to no interaction, represented by thin grey lines. In the NRI model, the possible combinations of interactions require eight ($= 2^3$) edge-types.

3. Geometric interpretation

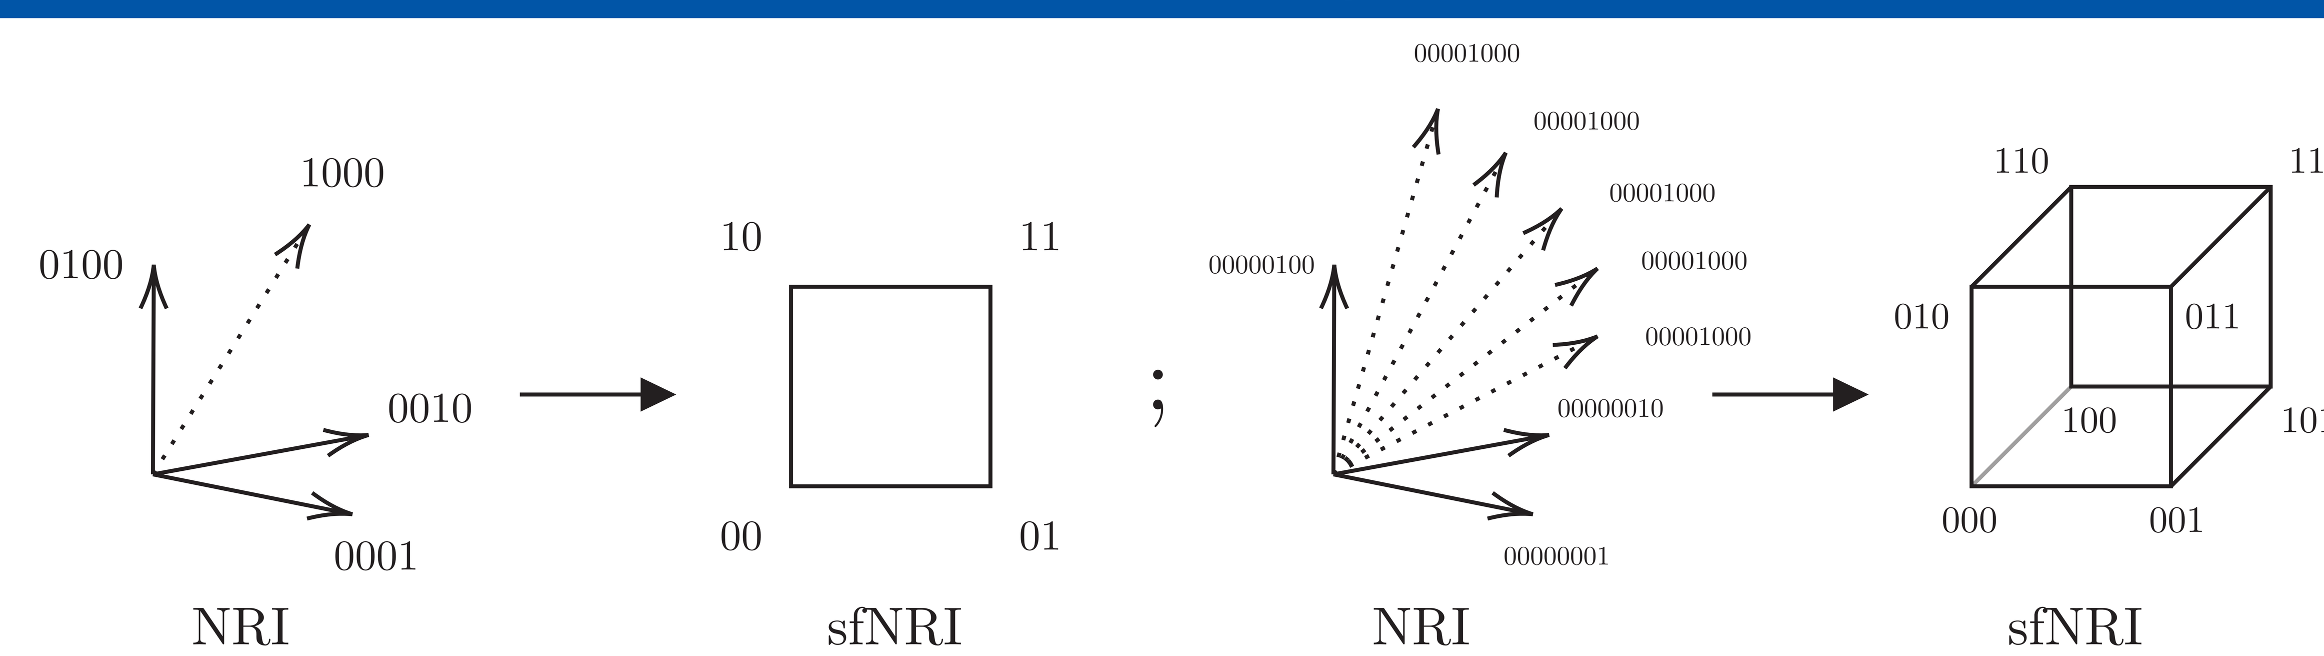

Figure 2: In factorising the graph we shift from a categorical representation (single graph, many edge-types) to a multicategorical representation (many graphs, single edge-type). On the left we show the different representations of the combination of two interaction types and on the right the combination of three types, and how the implicit geometries of these representations differ between the original and our formulation. The NRI implies each combination ($\propto 2^N$) is orthogonal and equidistant from the others. The sfNRI only requires a new dimension for each interaction ($\propto N$), representing the combinations as points on a square, cube or, more generally, an N-dimensional hypercube.

4. Results

We test the models in trajectory prediction and interaction graph recovery in physics simulations. There are three interaction types: idealised springs (I-springs) that apply a linear attractive force ($\propto x$); like-charges that produce an inverse-square repulsion ($\propto x^{-2}$); and springs of fixed length (F-springs) that are repulsive at short range and attractive at long range ($\propto |x - k|$). The rows labelled (supervised) or (true graph) are provided access to ground truth interaction graphs and provide an upper bound. **Our model improves upon the original in all cases.**

Table 1: Accuracy (%) in recovering the ground truth interaction graph. Higher is better.

	I-Springs+Charges			**I-Springs+Charges+F-springs**			
Accuracy	**Combined**	I-Springs	Charges	**Combined**	I-Springs	Charges	F-Springs
Random	25.0	50.0	50.0	12.5	50.0	50.0	50.0
NRI (learned)	89.1 ± 0.4	97.9 ± 0.0	91.0 ± 0.4	57.9 ± 6.1	88.5 ± 0.9	87.3 ± 6.2	70.7 ± 2.3
fNRI (learned)	**94.0 ± 1.4**	98.0 ± 0.1	95.8 ± 1.3	**63.3 ± 6.5**	86.9 ± 2.7	97.7 ± 0.7	69.2 ± 5.5
sfNRI (learned)	88.8 ± 0.8	97.6 ± 0.1	91.1 ± 0.8	45.1 ± 5.1	90.0 ± 2.3	98.2 ± 0.8	52.4 ± 2.7
NRI (supervised)	**98.3 ± 0.0**	98.6 ± 0.0	99.7 ± 0.0	80.9 ± 0.7	92.4 ± 0.3	99.0 ± 0.1	84.4 ± 0.4
fNRI (supervised)	**98.3 ± 0.0**	98.8 ± 0.4	99.4 ± 0.4	**81.8 ± 0.1**	93.3 ± 0.1	99.3 ± 0.0	85.8 ± 0.1
sfNRI (supervised)	98.0 ± 0.0	98.3 ± 0.0	99.6 ± 0.0	81.0 ± 0.3	92.9 ± 0.1	99.2 ± 0.0	85.2 ± 0.2

Table 2: Mean squared error (MSE) / 10^{-5} in trajectory prediction. Lower is better.

	I-Springs+Charges			**I-Springs+Charges+F-Springs**		
Predictions Steps	1	10	20	1	10	20
Static	19.4	283	783	12.8	274	782
NRI (learned)	0.88 ± 0.06	4.05 ± 0.22	11.5 ± 0.5	0.95 ± 0.05	8.67 ± 0.45	29.1 ± 1.4
fNRI (learned)	**0.80 ± 0.04**	3.54 ± 0.09	9.93 ± 0.29	0.81 ± 0.05	7.78 ± 0.20	26.8 ± 0.8
sfNRI (learned)	1.03 ± 0.09	**3.32 ± 0.23**	**9.68 ± 0.74**	**0.77 ± 0.03**	**5.69 ± 0.21**	**19.3 ± 0.8**
NRI (true graph)	0.85 ± 0.04	1.59 ± 0.26	3.20 ± 0.15	0.75 ± 0.02	1.55 ± 0.07	3.43 ± 0.21
fNRI (true graph)	**0.70 ± 0.03**	**1.30 ± 0.06**	**2.52 ± 0.11**	**0.51 ± 0.05**	0.97 ± 0.08	2.44 ± 0.28
sfNRI (true graph)	0.86 ± 0.09	1.32 ± 0.06	2.77 ± 0.07	0.56 ± 0.04	**0.89 ± 0.06**	**2.28 ± 0.15**