Relational representation learning with spiking neural networks

Dominik Dold[1], Victor C. Chian[2], Josep Soler Garrido[3], Marcel Hildebrandt[2], Thomas Runkler[2]

[1]European Space Agency (ESA), European Space Research and Technology Centre (ESTEC), Advanced Concepts Team (ACT), Noordwijk, Netherlands.
[2]Siemens AG, Siemens Technology, Munich, Germany. [3]European Commission, Joint Research Centre (JRC), Seville, Spain.

Knowledge graphs

Flexible data structure: molecules, social networks, projects, …

Represented via triples:

C. Fisher	*plays*	*L. Organa*
L. Organa	*appearsIn*	*Star Wars*
	…	

or as a **heterogeneous graph:**

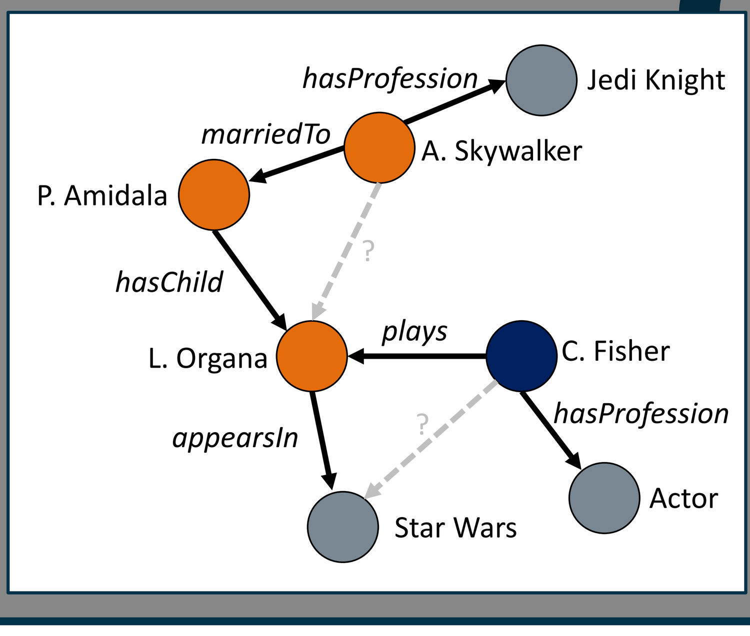

Inference: predict missing links

C. Fisher	*appearsIn*	*Star Wars?*
A. Skywalker	*hasChild*	*L. Organa?*
	…	

Graph embedding

„Equivalent" **vector space** representation learned via **gradient descent**.

graph $\xrightarrow{\text{embed}}$ $\vec{e}_n, \vec{r}_p$ $\xrightarrow{f(\vec{e}_s, \vec{r}_p, \vec{e}_o)}$ link existence

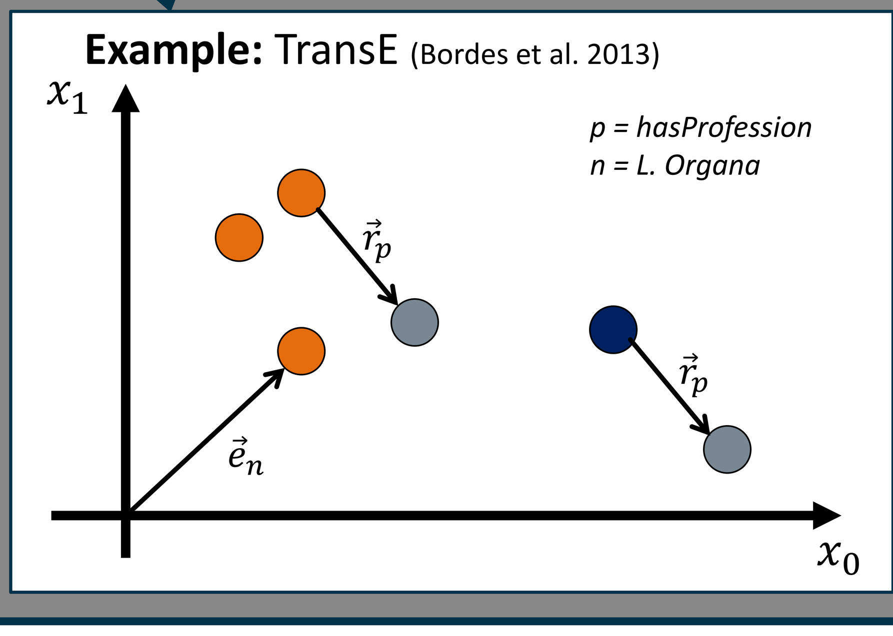

Decoder function f**:** reconstructs graph property

Inference: evaluate model

$f(\vec{e}_{C.F.}, \vec{r}_{appearsIn}, \vec{e}_{S.W.})$

$f(\vec{e}_{A.S.}, \vec{r}_{hasChild}, \vec{e}_{L.O.})$

3. Spike-based graph neural networks

Extension of **convolution operator** to graphs:

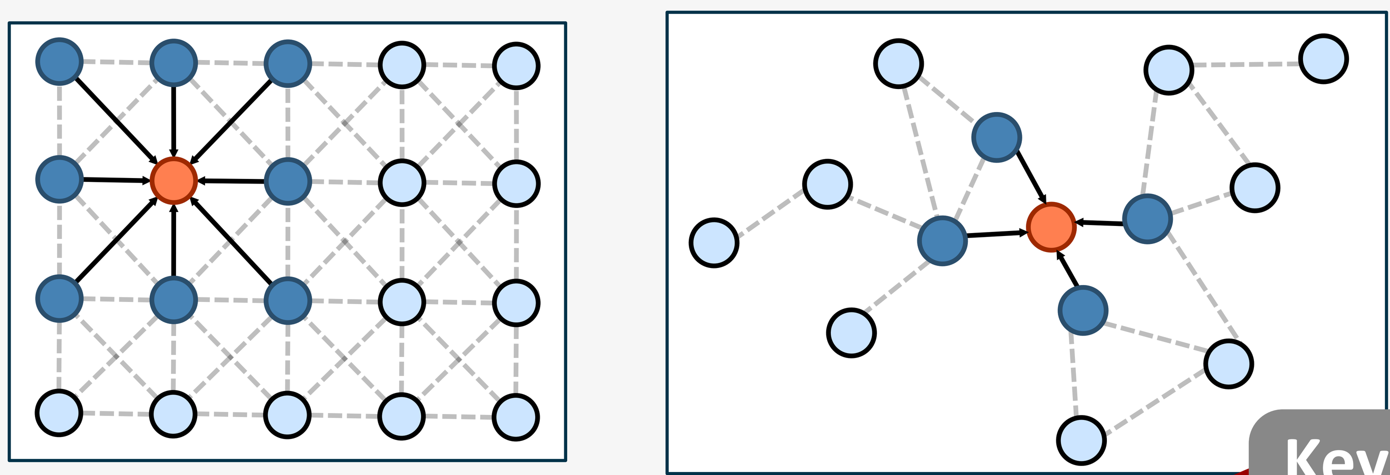

Idea: update embeddings using local information.

New:
- irregular structure,
- permutation invariant w.r.t. nodes,
- concepts not seen during training can be embedded!

1) Weight sharing: freeze weights

Dataset	Non-frozen MRR	Frozen MRR
FB15k-237	0.23	**0.26**
UMLS	0.58	**0.80**

Embeddings **align** to static and random weights.

2) Building spike-based neural networks for graph inference

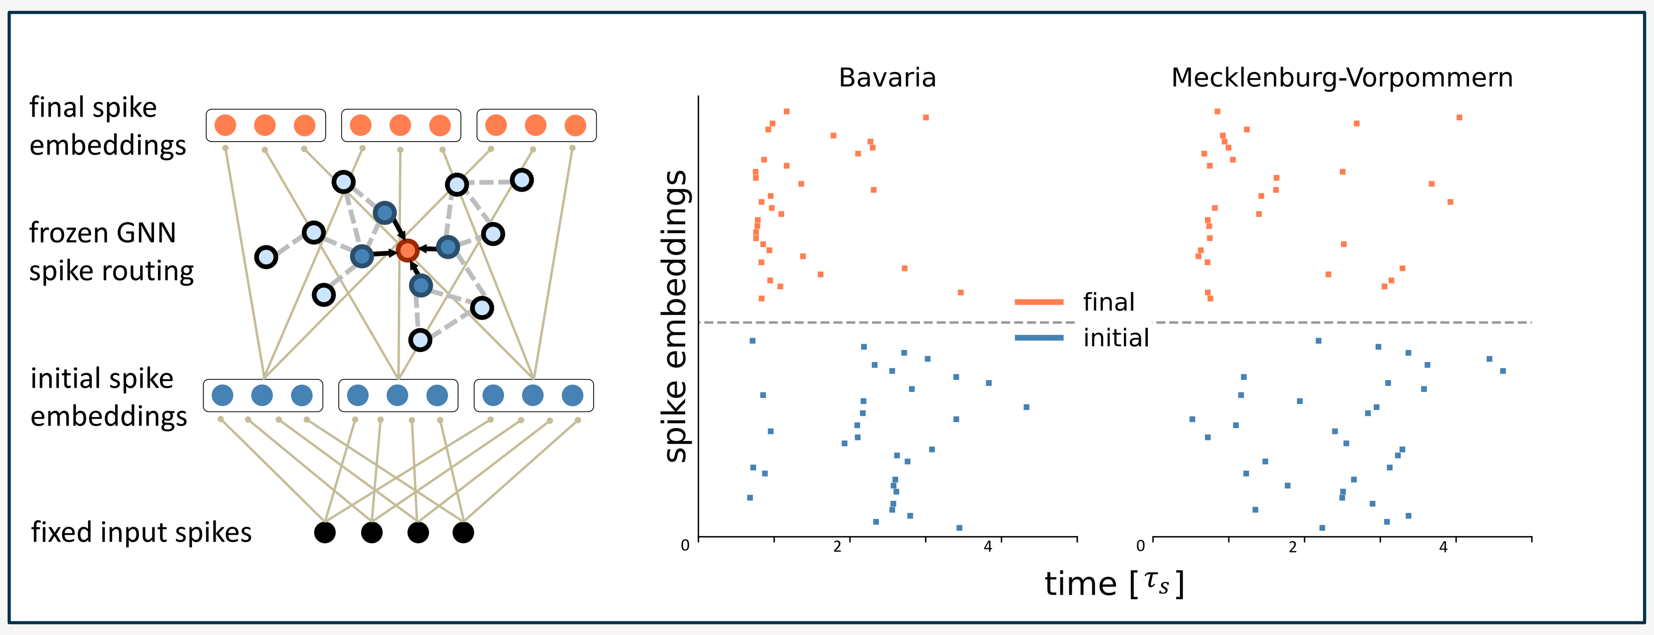

$$\tau_s \, \dot{u}_s(t) \propto \sum_{\text{relations } r} \sum_{\text{neighbors } n} W_r \, \boldsymbol{\kappa}(t, \vec{e}_n)$$

Dataset	MRR SGNN	Reference (TransE)
German States	**0.56**	0.69
Starcraft	**0.67**	0.71

Simultaneous and **sparse** (20-30%) computation!

Key message

Given a concept like *actor*, **how** can it be represented using spikes? **How** are such representations used to infer that, e.g., a person is an actor?

We **propose** to represent such **concepts** ("person", "actor") and **relationships** ("*is an*") in the **spike domain** using graph embedding.

1. Time to first spike embeddings

Node embedding: time to first spike of population.

Relation embedding: spike time differences.

Decoder: $\sum_j ||d(\vec{e}_s, \vec{e}_o) - \vec{r}_p||_j$

with $d(\vec{e}_s, \vec{e}_o) = ||\vec{e}_s - \vec{e}_o||$

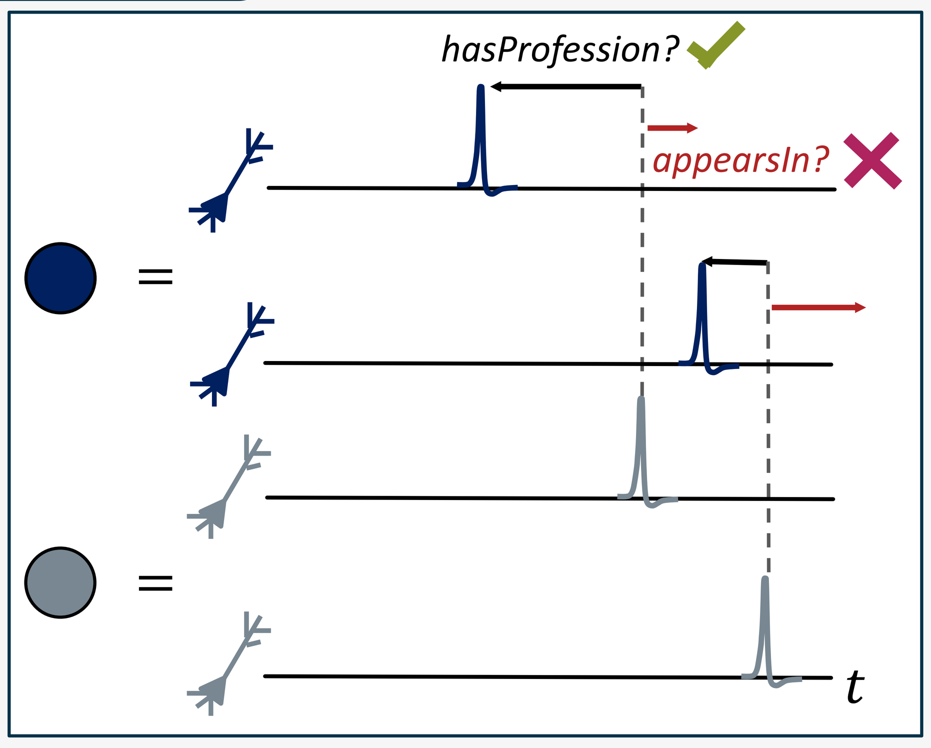

Neuron model: here: I&F, only requires calculatable **gradient w.r.t. spike times**.

2. Spike train embeddings

Node embedding: spike train.

Relation embedding: spike time differences.

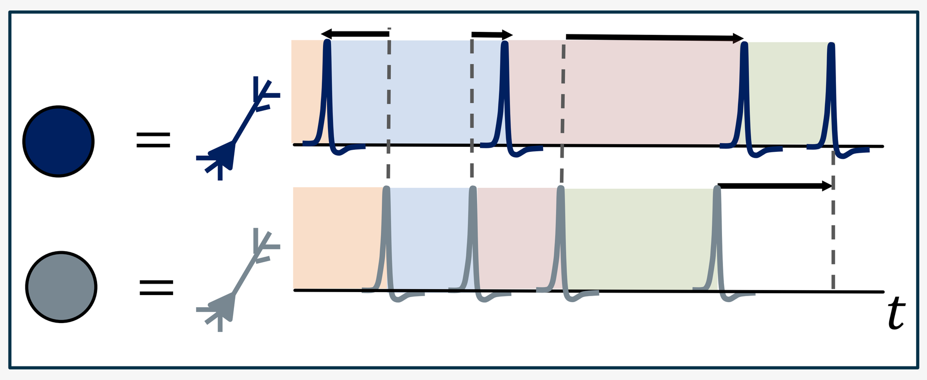

Challenge: spike order has to be conserved.

Solution: updates via inter-spike intervals.

$$\begin{pmatrix} \mathrm{ISI}_0 \\ \mathrm{ISI}_0 + \mathrm{ISI}_1 \\ \mathrm{ISI}_0 + \mathrm{ISI}_1 + \mathrm{ISI}_2 \\ \mathrm{ISI}_0 + \mathrm{ISI}_1 + \mathrm{ISI}_2 + \mathrm{ISI}_3 \end{pmatrix}$$

1) Zachary Karate Club

Karate Club splits into two groups led by person **1** and **34**.

Given: social graph
Task: how do they split?

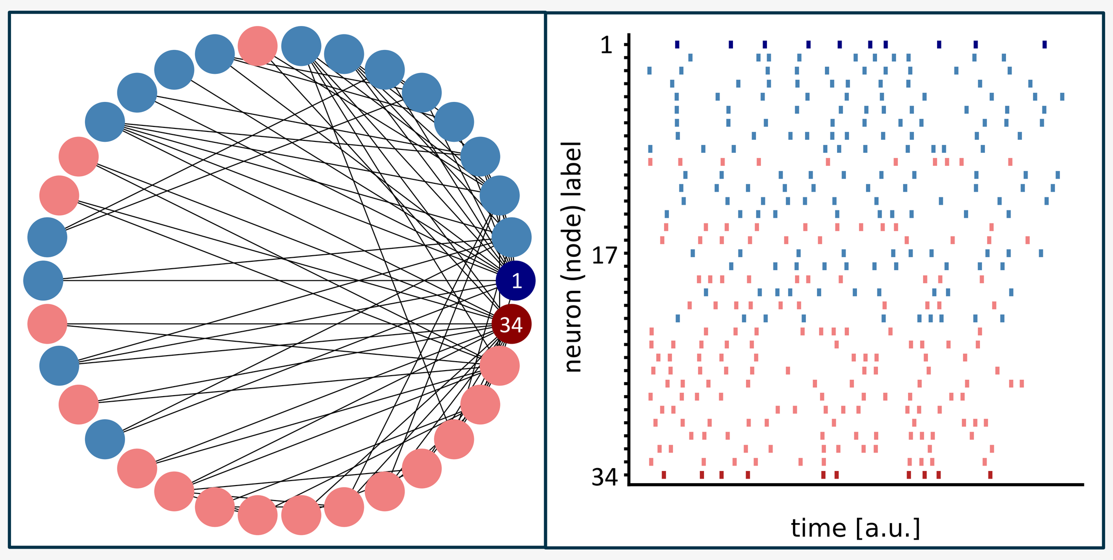

2) Link prediction benchmarks

Mean Reciprocal Rank (MRR) measures link prediction performance.

Data set	MRR ours	TransE	RESCAL
FB15k-237	**0.21**	0.21	0.28
CoDEx-S	**0.30**	0.35	0.40
IAD	**0.66**	0.66	0.61
UMLS	**0.81**	0.81	0.88
Kinships	**0.47**	0.48	0.81

To boldly go…

Spiking neural networks for Earth observation

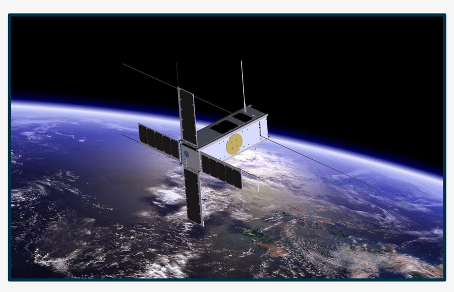

Differentiable graphs for lattice inverse design

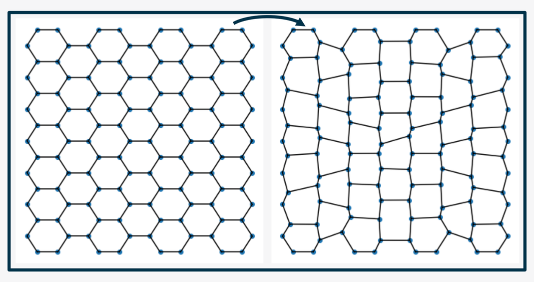

Self-configuring cube ensembles

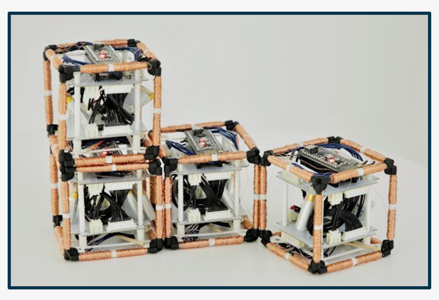

Check out our projects!

Publications:

SpikE: spike-based embeddings for multi-relational graph data, arXiv:2104.13398.

Learning through structure: towards deep neuromorphic knowledge graph embeddings, arXiv:2109.10376.

Relational representation learning with spike trains, arXiv:2205.09140.

Neuro-symbolic computing with spiking neural networks, arXiv:2208.02576.

We thank Serghei Mogoreanu, Alexander Hadjiivanov and Gabriele Meoni for helpful discussions. We further thank our colleagues at the Siemens Semantics and Reasoning Research Group, the Siemens AI Lab and ESA's Advanced Concepts Team for their support. This work was partially funded by the Federal Ministry for Economic Affairs and Energy of Germany (IIP-Ecosphere Project) and by the German Federal Ministry for Education and Research ("MLWin", grant 01IS18050). DD acknowledges support through the European Space Agency fellowship programme.

Email contact: dominik.dold@esa.int