

# Roto-translated Local Coordinate Frames For Interacting Dynamical Systems

---

Miltiadis Kofinas<sup>1</sup>, Naveen Shankar Nagaraja<sup>2</sup>, Efstratios Gavves<sup>1</sup>


<sup>1</sup>VIS Lab  
University of Amsterdam  
Amsterdam, Netherlands

<sup>2</sup>Department of Autonomous Driving  
BMW Group  
Munich, Germany


Learning on Graphs and Geometry Reading Group  
1 February 2022

# Interacting systems are everywhere

- Colliding particles
- N-body systems
- Molecules
- Motion capture
- Traffic scenes



3D Charged particles [6]




Traffic scene [8]

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. “Neural relational inference for interacting systems”. In: *ICML*. 2018


[8] Tim Salzmann<sup>†</sup>, Boris Ivanovic<sup>†</sup>, et al. “Trajectron++: Dynamically-Feasible Trajectory Forecasting With Heterogeneous Data”. In: *ECCV*. 2020

# Future forecasting


Past



Future



## Related work - Neural Relational Inference [6]




- Explicitly infer graph structure over latent edge types
- Simultaneously learn the dynamical system

[5] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *ICLR*. 2014

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018


## Related work - Dynamic Neural Relational Inference [4]



- Dynamic relations through time
- Sequential approximate posterior based on past states

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: CVPR. 2020

# Related work - E(n) Equivariant Graph Networks [9]




- Leverage rotation equivariant relative positions and invariant euclidean distances

[9] Víctor García Satorras, Emiel Hoogetboom, and Max Welling. "E(n) Equivariant Graph Neural Networks". In: *ICML*. 2021


# Motivation

What happens when we rotate/translate the inputs?


Past



Future

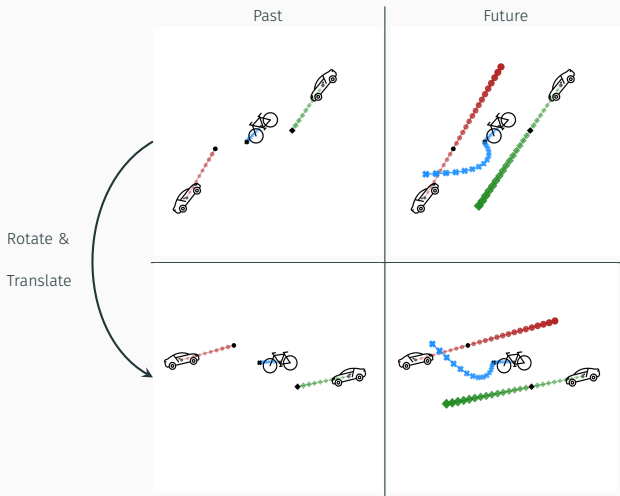


# Motivation






# Motivation



**Roto-translation equivariance**

Dynamics do not change under rotations and translations

# Motivation




## Ego-centric perspective

Objects operate in ego-centric and asymmetric views of the world

## Global coordinate frames

Graphs embedded in arbitrary global coordinate frames

# Motivation




## Ego-centric perspective

Objects operate in ego-centric and asymmetric views of the world


## Global coordinate frames

Graphs embedded in arbitrary global coordinate frames


## Local coordinate frames




## Local coordinate frames




## Local coordinate frames



# Local coordinate frames




# Local coordinate frames





# Local coordinate frames



# Local coordinate frames as Neural relational inference models






[5] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *ICLR*. 2014

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020

# Local coordinate frames as Neural relational inference models




[5] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *ICLR*. 2014

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020

# Local coordinate frames as Neural relational inference models




[5] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *ICLR*. 2014

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020

# Local coordinate frames as Neural relational inference models

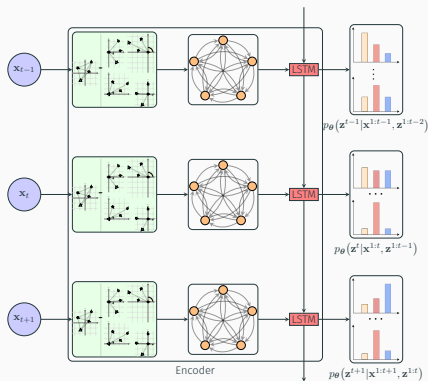


[5] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *ICLR*. 2014

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020

# Local coordinate frames as Neural relational inference models

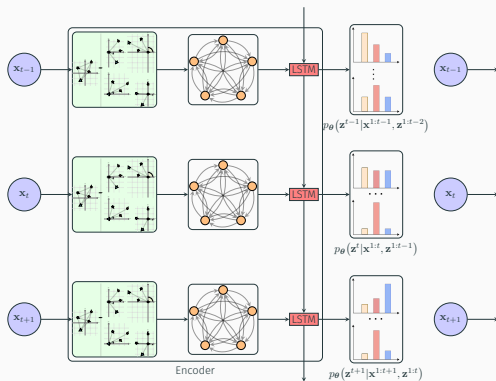


[5] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *ICLR*. 2014

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020

# Local coordinate frames as Neural relational inference models

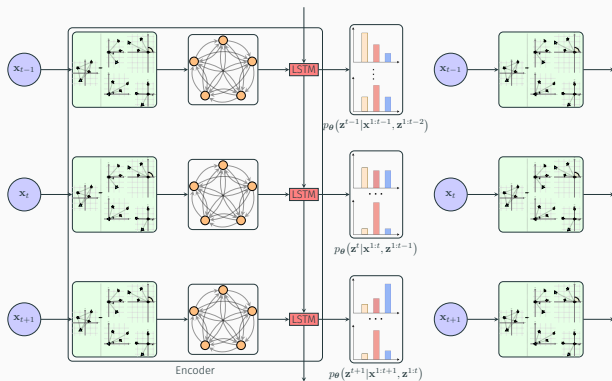


[5] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *ICLR*. 2014

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020

# Local coordinate frames as Neural relational inference models



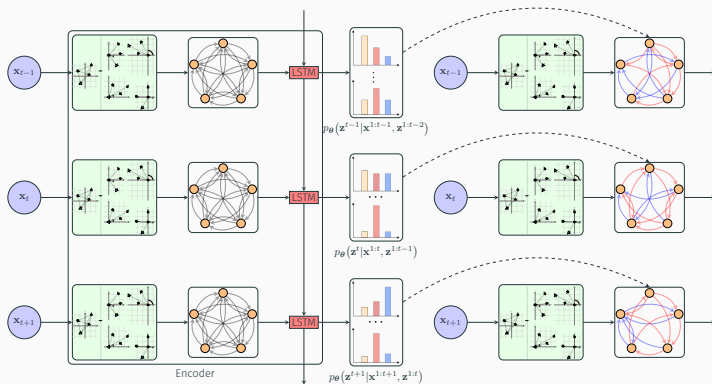
[5] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *ICLR*. 2014

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020



# Local coordinate frames as Neural relational inference models

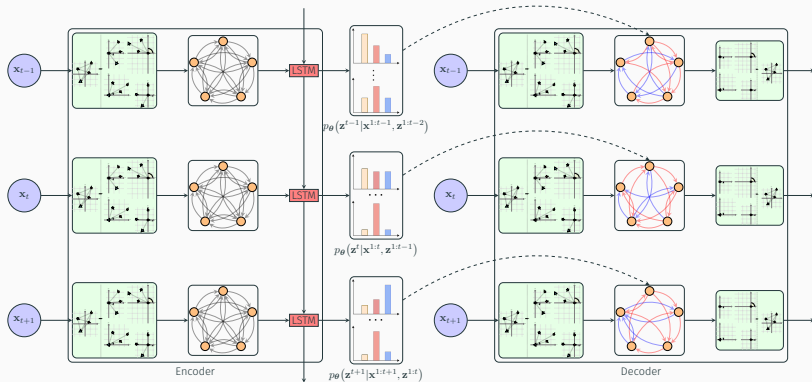


[5] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *ICLR*. 2014

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020

# Local coordinate frames as Neural relational inference models

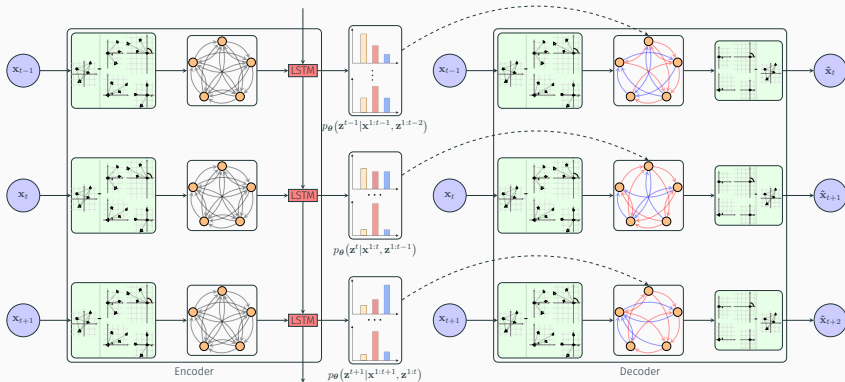


[5] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *ICLR*. 2014

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020

# Local coordinate frames as Neural relational inference models





[5] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes". In: *ICLR*. 2014

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020

# Anisotropic continuous filtering in local coordinate frames

Directionality in graphs  $\implies$  Anisotropic filtering



# Experiments

Synthetic [1]

- 2D, repulsive forces

InD [1]

- Traffic scenes, 2D, social interactions


Charged particles [6]

- 3D, electrostatic forces

CMU Motion capture [3]

- 3D, subject #35, walking trajectories

# Results




[4] Colin Graber and Alexander G Schwing. “Dynamic Neural Relational Inference”. In: *CVPR*. 2020

[1] Julian Bock et al. “The inD dataset: A drone dataset of naturalistic road user trajectories at german intersections”. In: *2020 IEEE Intelligent Vehicles Symposium (IV)*. 2020

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. “Neural relational inference for interacting systems”. In: *ICML*. 2018

[3] CMU. *Carnegie-Mellon Motion Capture Database*. 2003. URL: <http://mocap.cs.cmu.edu>

# Results - Synthetic Dataset



Relation prediction F1 score on synthetic dataset


Method	NRI	dNRI	LoCS
F1	26.5	60.8	88.9

[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018


[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020

# Qualitative results - charged particles


Groundtruth




LoCS (Ours)




dNRI [4]



NRI [6]



EGNN [9]



[6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, et al. "Neural relational inference for interacting systems". In: *ICML*. 2018

[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020


[9] Víctor García Satorras, Emiel Hoogetboom, and Max Welling. "E(n) Equivariant Graph Neural Networks". In: *ICML*.

2021




# Qualitative results - inD


Groundtruth




LoCS (Ours)



dNRI [4]



EGNN [9]



[4] Colin Graber and Alexander G Schwing. "Dynamic Neural Relational Inference". In: *CVPR*. 2020

[9] Víctor García Satorras, Emiel Hoogetboom, and Max Welling. "E(n) Equivariant Graph Neural Networks". In: *ICML*.

2021

# Conclusion

- Local coordinate frames for all objects
- Invariance/equivariance to global roto-translations
- Anisotropic continuous filters in local coordinate frames
- Demonstrate effectiveness on a range of 2D/3D settings
- Source code: <https://github.com/mkofinas/locs>
- Paper: <https://arxiv.org/abs/2110.14961>



- [1] Julian Bock et al. “The inD dataset: A drone dataset of naturalistic road user trajectories at german intersections”. In: *2020 IEEE Intelligent Vehicles Symposium (IV)*. 2020.
- [2] Kyunghyun Cho et al. “Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2014.
- [3] CMU. *Carnegie-Mellon Motion Capture Database*. 2003. URL: <http://mocap.cs.cmu.edu>.
- [4] Colin Graber and Alexander G Schwing. “Dynamic Neural Relational Inference”. In: *CVPR*. 2020.

- [5] Diederik P Kingma and Max Welling. “Auto-encoding variational bayes”. In: *ICLR*. 2014.
- [6] Thomas Kipf<sup>†</sup>, Ethan Fetaya<sup>†</sup>, Kuan-Chieh Wang, Max Welling, and Richard Zemel. “Neural relational inference for interacting systems”. In: *ICML*. 2018.
- [7] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. “Stochastic backpropagation and approximate inference in deep generative models”. In: *ICML*. 2014.
- [8] Tim Salzmann<sup>†</sup>, Boris Ivanovic<sup>†</sup>, Punarjay Chakravarty, and Marco Pavone. “Trajectron++: Dynamically-Feasible Trajectory Forecasting With Heterogeneous Data”. In: *ECCV*. 2020.

- [9] Víctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. “E(n) Equivariant Graph Neural Networks”. In: *ICML*. 2021.
- [10] Martin Simonovsky and Nikos Komodakis. “Dynamic edge-conditioned filters in convolutional neural networks on graphs”. In: *CVPR*. 2017.