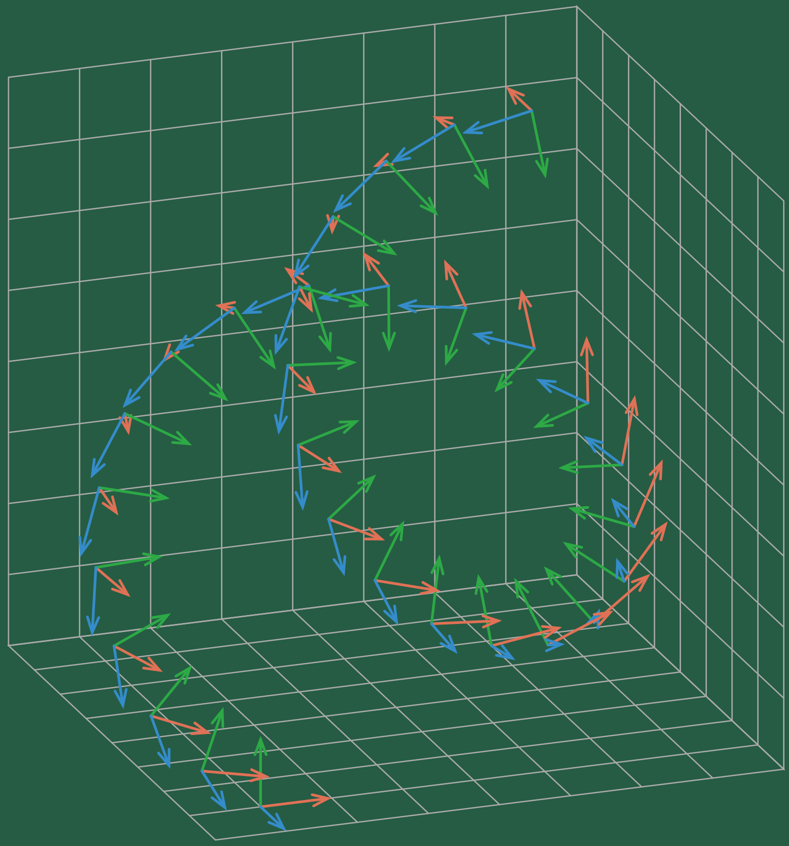

Euclidean Equivariant Processing and Machine Learning on Position-Orientation Space

Gijs Bellaard

Introduction

Introduction

This thesis explores the interplay between:

- Euclidean equivariant processing.
- Position-orientation space.
- Machine learning.

Euclidean Equivariant Processing.

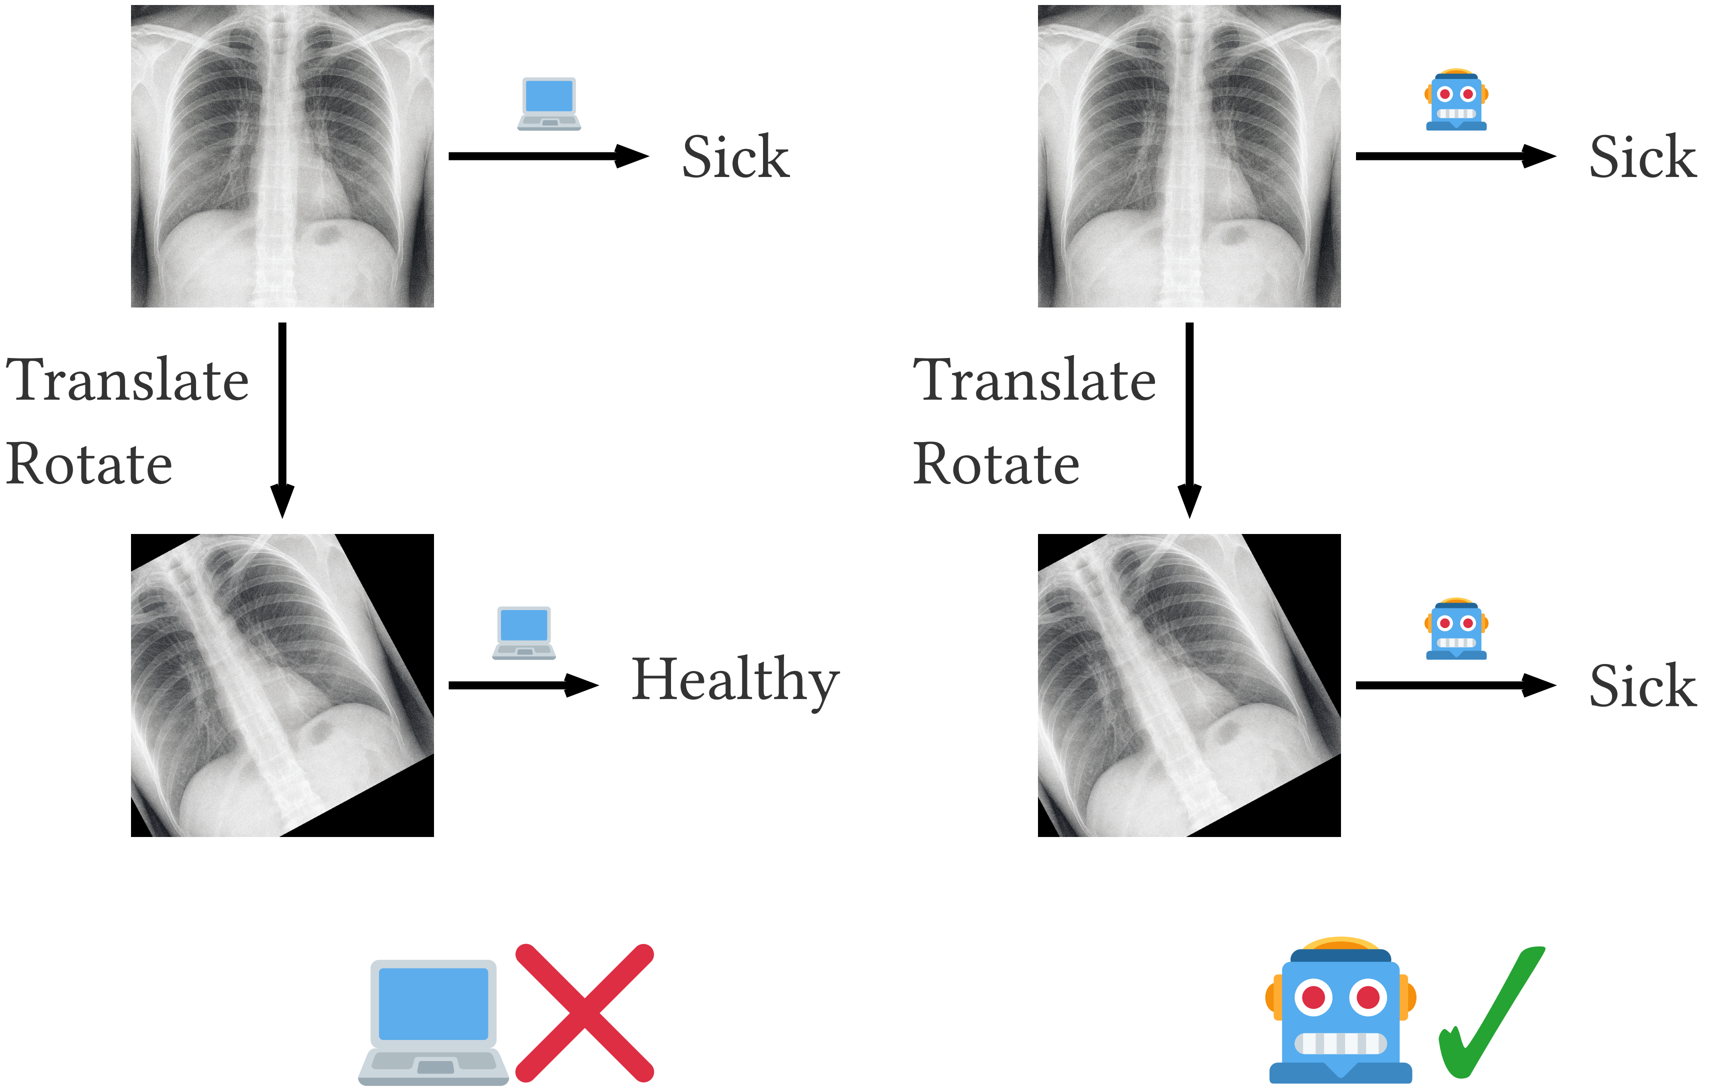

Euclidean Equivariant Processing.

- We say a model or process is *equivariant* if it respects a desired set of transformations.
- The transformations that we consider are rotations and translations.
- These are called *Euclidean transformations.*
- In short, we consider *Euclidean equivariant processing*.

Position-Orientation Space.

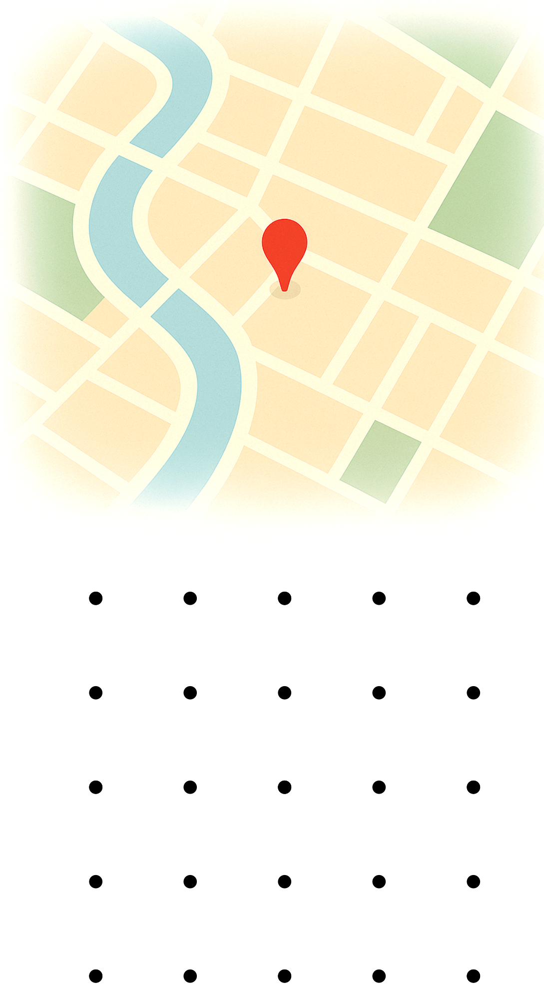

Space of Positions

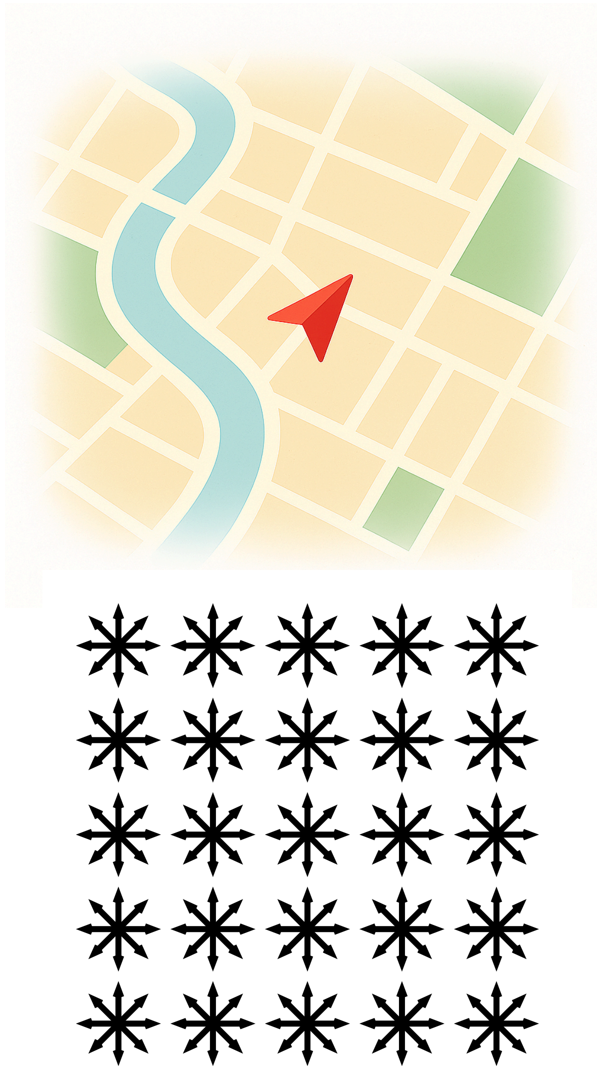

Space of Position-Orientations

Machine Learning.

- More and more problems are being tackled on computers using *machine learning*.
- Machine learning is about the development and study of models that can “learn” from existing data and generalize to unseen data.
- Machine learning has led to major breakthroughs in tasks previously considered intractable by classical methods.

Supervised Machine Learning

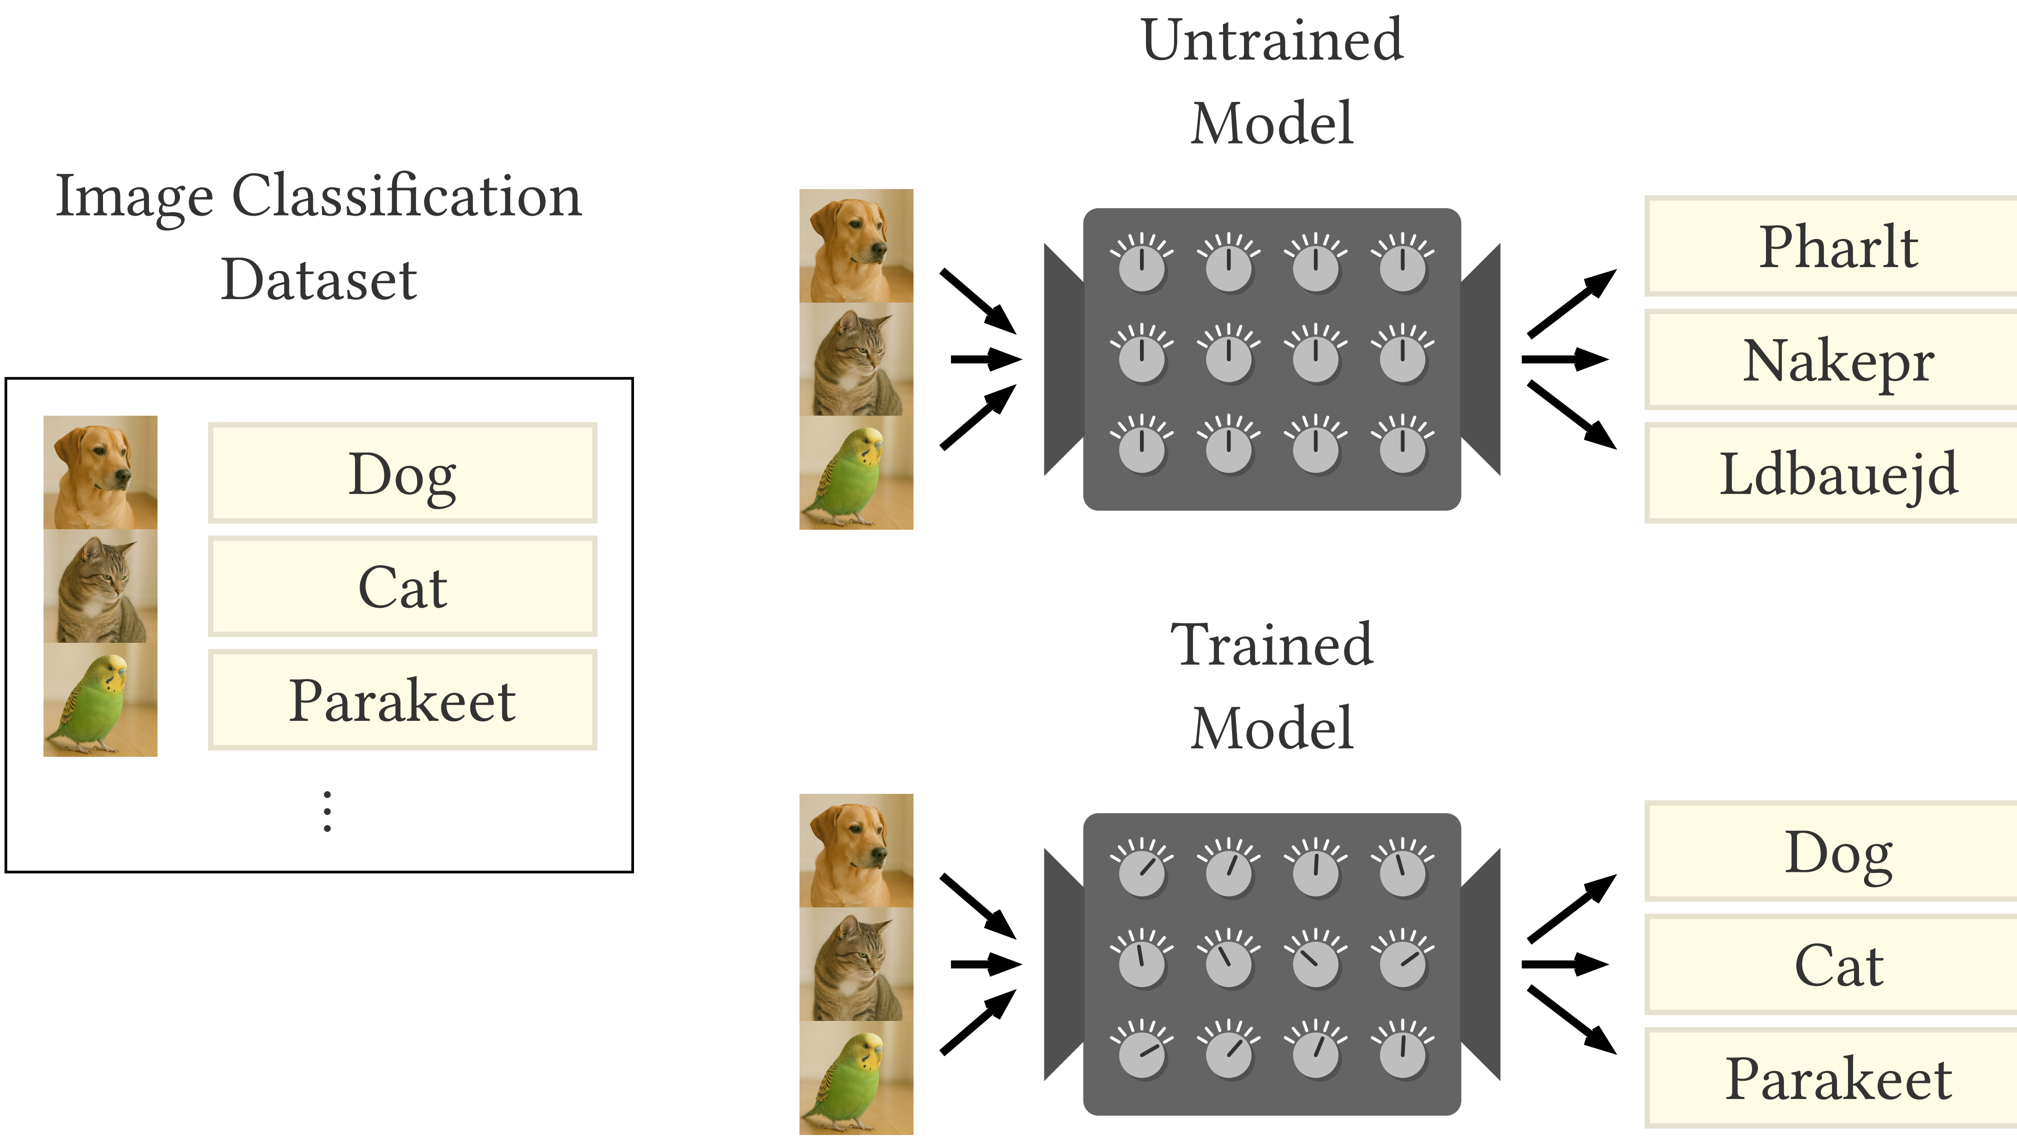

Concrete Applications

PDE-based Neural Networks

- In Duits, Smets, Bekkers and Portegies (2021) a machine learning framework called *PDE-based neural networks* is introduced.
- This framework uses position-orientation space and is Euclidean equivariant.
- It can be used to segment vessels in retinal images and coronary arteries in X-ray angiograms.

Retinal vessels segmentation

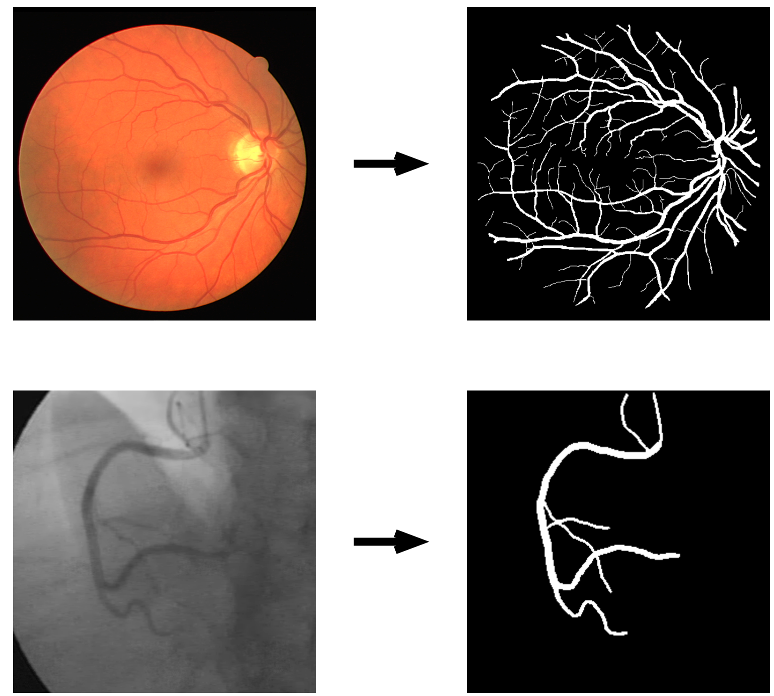

Coronary artery segmentation

PONITA

- In Bekkers, Vadgama, Hesselink, van der Linden and Romero (2024) a machine learning model *PONITA* is introduced.
- This model uses position-orientation space and is Euclidean equivariant.
- It can be used to predict the chemical properties of molecules.

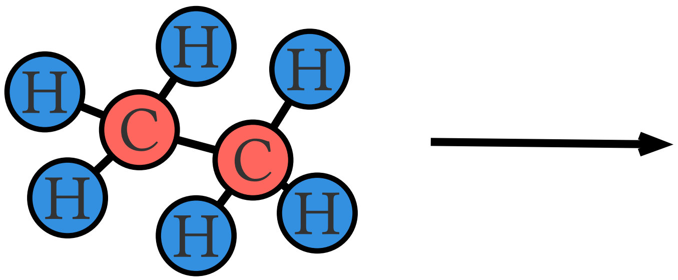

Dipole moment:	0.05	D
Isotropic polarizability:	9.13	a_0^3
Internal energy at 0K:	-2.18	MeV

Contributions

Contributions

- PDE-based neural networks use the *distance* between position-orientations. However, this exact distance is *expensive.*
- In Chapter 3 we introduce and analyze *efficient approximations* of the distance. We investigate how these approximations impact the accuracy of the networks.

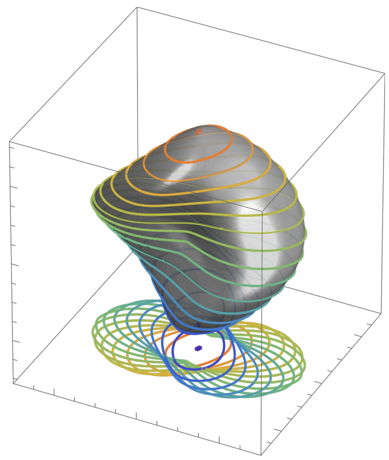

- The distance between position-orientations depends on the *Euclidean invariant Riemannian metric* that is chosen.
- In Chapter 4 we classify and parametrize *all* such metrics, and describe the *mav distance* as an efficient alternative to the distance.

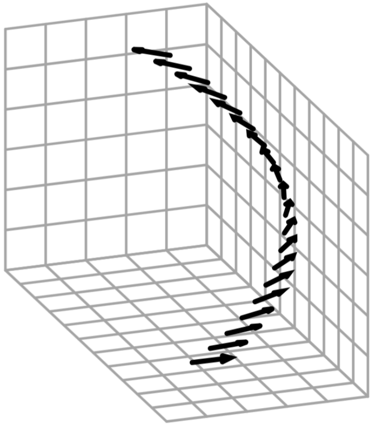

Contributions

- The PONITA network builds upon *Euclidean invariants* between position-orientations.
- In Chapter 5 we describe an *independent* and *universal* collection of invariants. We show that this collection of invariants can improve the accuracy of PONITA.

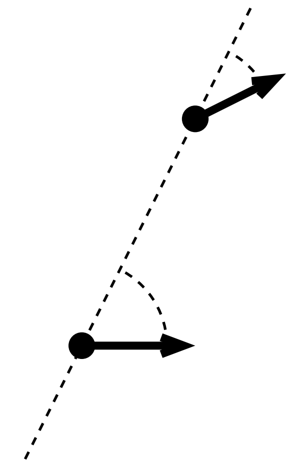

- Data usually "lives" on position space which we need to transfer to position-orientation space, a process called *lifting*.
- In Chapter 6 we show that lifting using *cake wavelets* can perform just as well as a *trained lifting* within PDE-based neural networks.

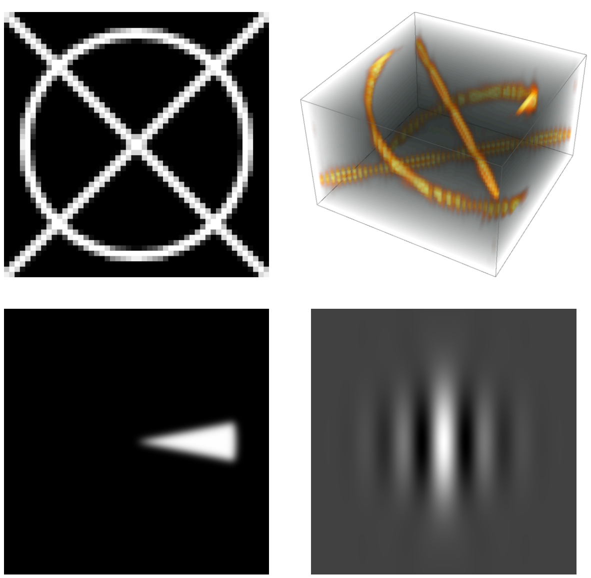

Contributions

- PDE-based neural networks use PDEs that come from *scale-space theory*.
- In Chapter 7 we generalize this theory to *semifields*. Our axiomatic setup reveals new PDEs to use within the networks, and we show their impact on accuracy.

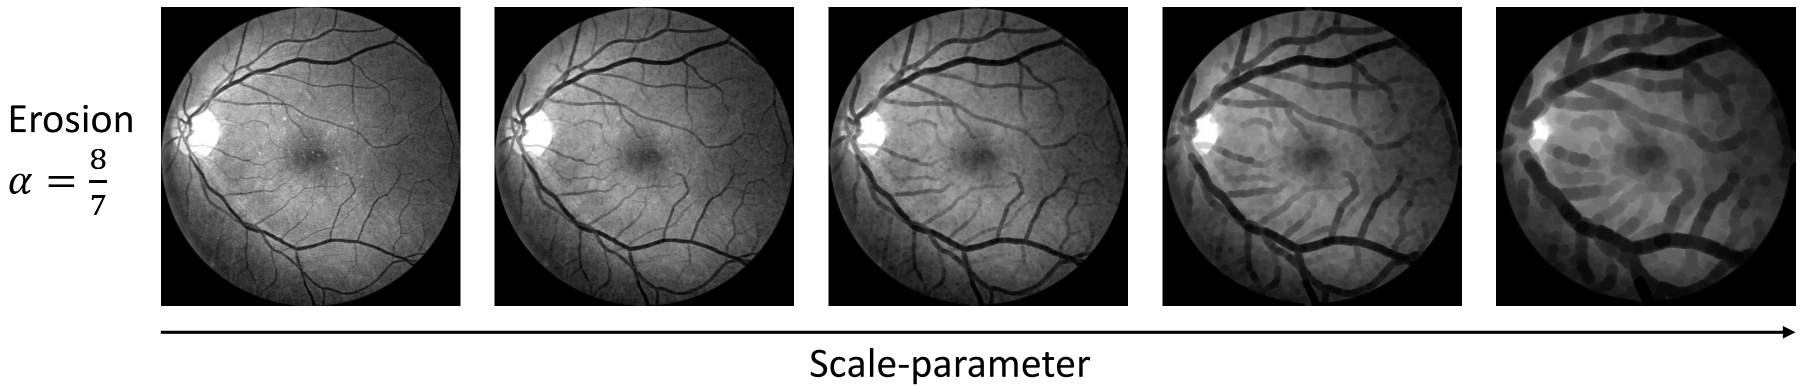

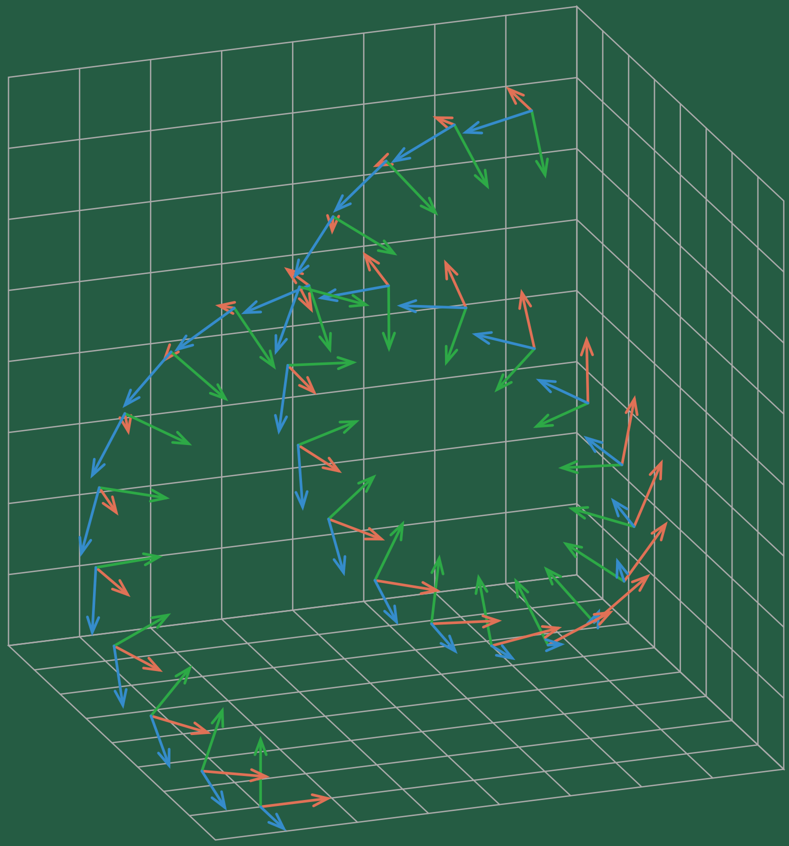

Euclidean Equivariant Processing and Machine Learning on Position-Orientation Space

Gijs Bellaard