Spatiotemporal Deformable Prototypes for Motion Anomaly Detection

Robert Bensch, Thomas Brox and Olaf Ronneberger

Department of Computer Science and
BIOSS Centre for Biological Signalling Studies
University of Freiburg, Germany

British Machine Vision Conference 2015
September 7-10, Swansea, UK

This study was supported by the Excellence Initiative of the German Federal and State Governments (EXC 294).

Motivation – Biological Motion Patterns

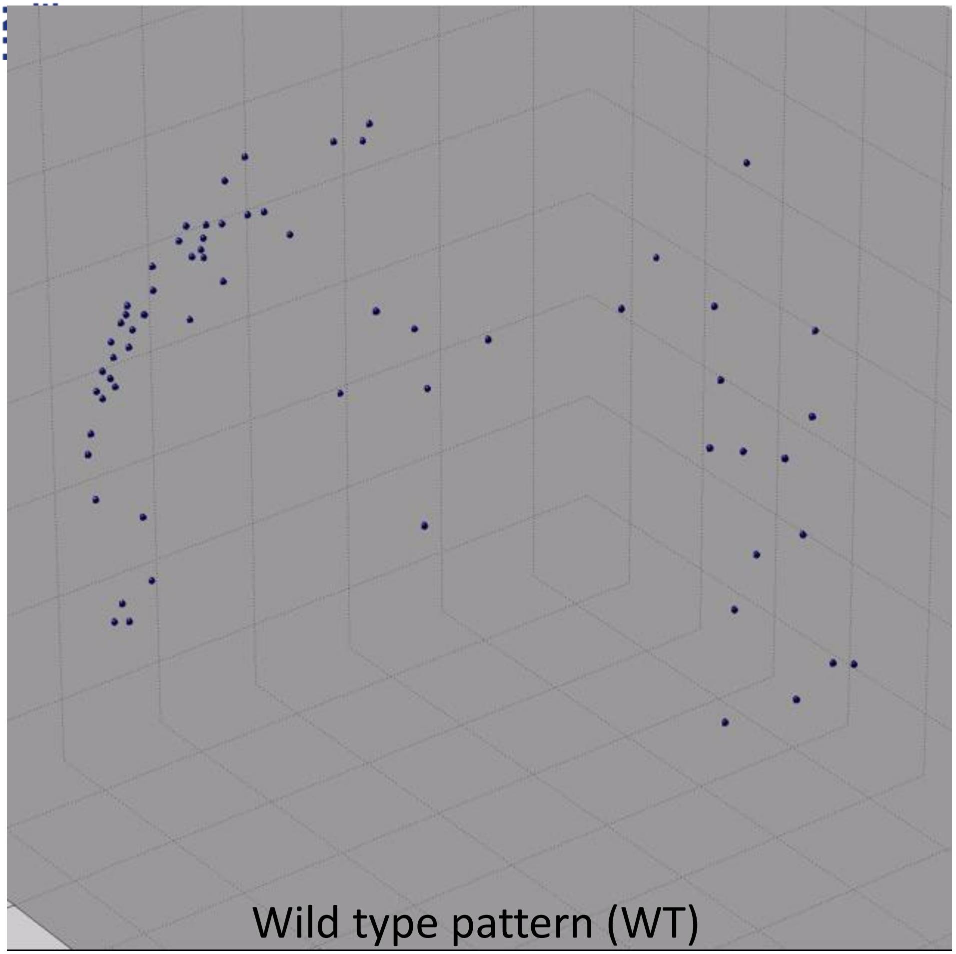

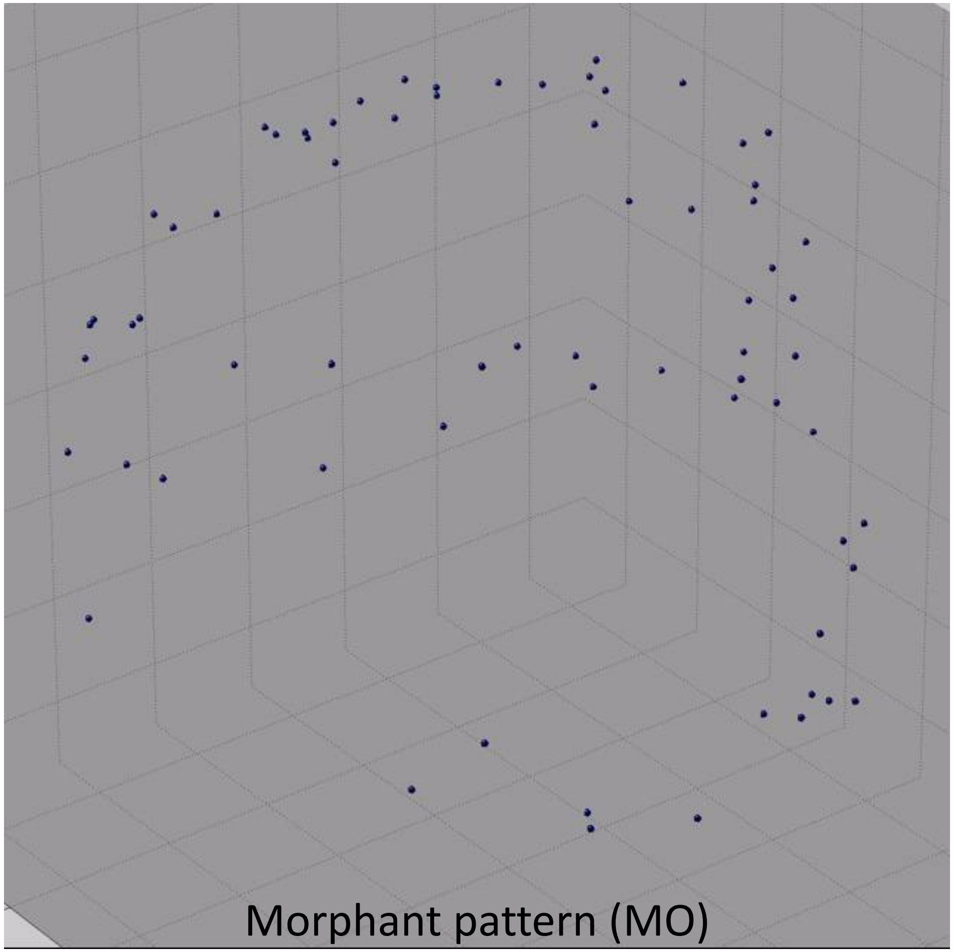

- **Goal:** Compare complex developmental growth patterns, identify differences to the wild type.
- Hard to get groundtruth in these biomedical settings

Motivation – Juggling Patterns

- 3D juggling dataset (recorded with Kinect camera)
- Normal pattern is a “3-ball cascade” juggling pattern
- **Goal:** Detect motion anomalies wrt. the normal pattern

Related Work and Contribution


Mahadevan et al., CVPR 2010

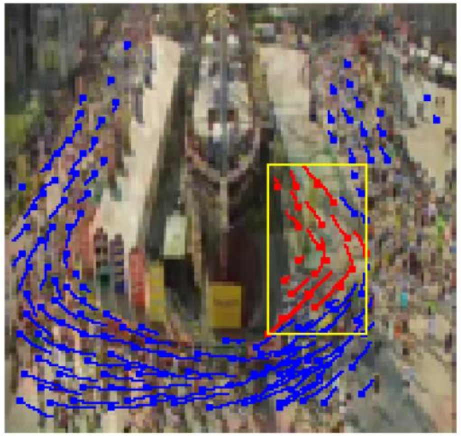
Wu et al., CVPR 2010

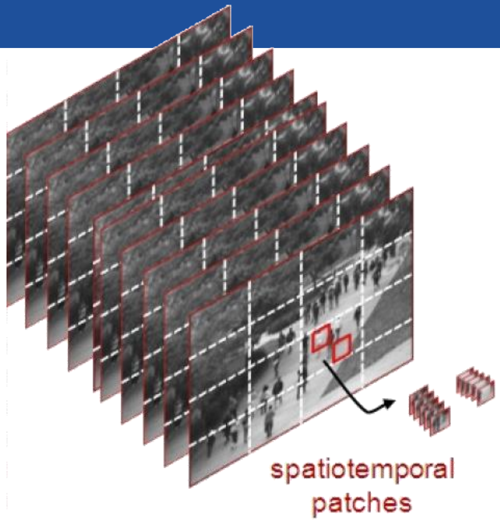

Mahadevan et al., CVPR 2010

- Surveillance scenarios and crowd analysis
  - Mahadevan et al., Anomaly detection in crowded scenes, CVPR 2010
  - Wu et al., Chaotic invariants of lagrangian particle trajectories for anomaly detection in crowded scenes, CVPR 2010
  - Cong et al., Sparse reconstruction cost for abnormal event detection, CVPR 2011
- Our contribution
  - Motion-based anomaly detection in a new setting and in *3D+time*
  - New method for elastic registration of motion patterns
  - New motion anomaly dataset (juggling patterns)

Approach Overview

- Input sequence

Approach Overview

- Compute dense point trajectories*

*Sundaram et al., Dense point trajectories by gpu-accelerated large displacement optical flow, ECCV 2010

Approach Overview


- Compute dense point trajectories*
- Kinect camera (RGB-D data, depth information)
- Compute *3D+time* trajectories

*Sundaram et al., Dense point trajectories by gpu-accelerated large displacement optical flow, ECCV 2010

Approach Overview

- Compute „supertrajectories“ in video (similar to „supervoxels“ in images)

Approach Overview

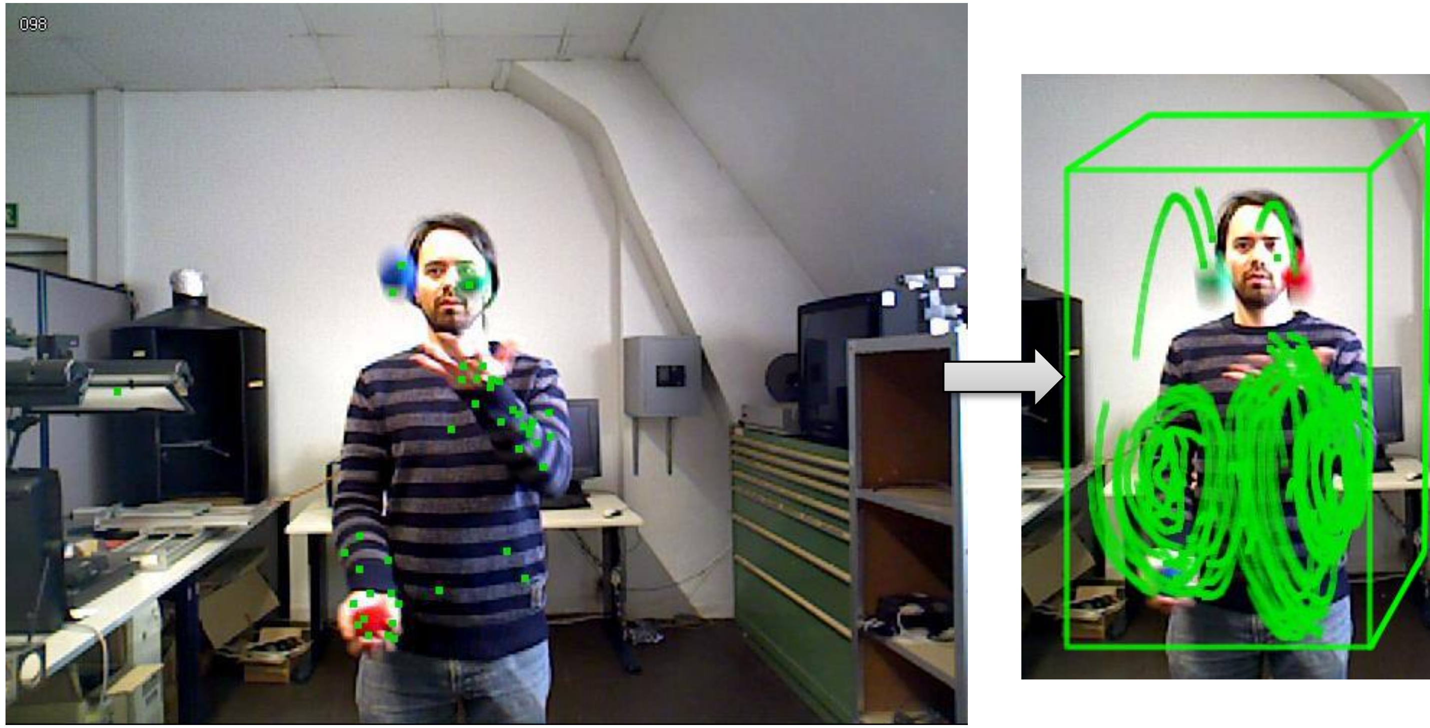

- Select a supertrajectory prototype (training phase)

Approach Overview

- Test sequence

Approach Overview

- Compute supertrajectories

Approach Overview

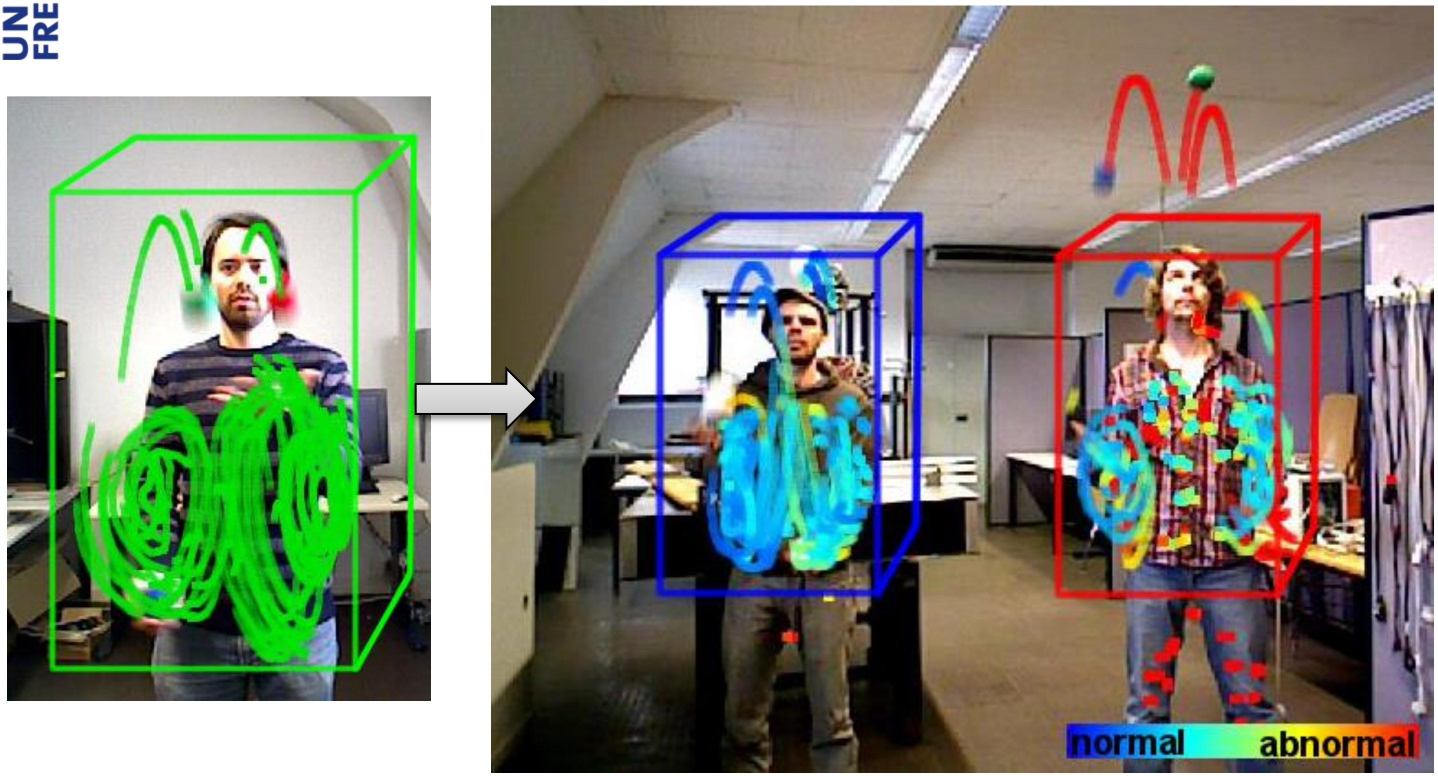

- Detect prototype instances (illustrated by bounding boxes)
- Rigid and elastic registration of prototype patterns (spatiotemporal)

Approach Overview

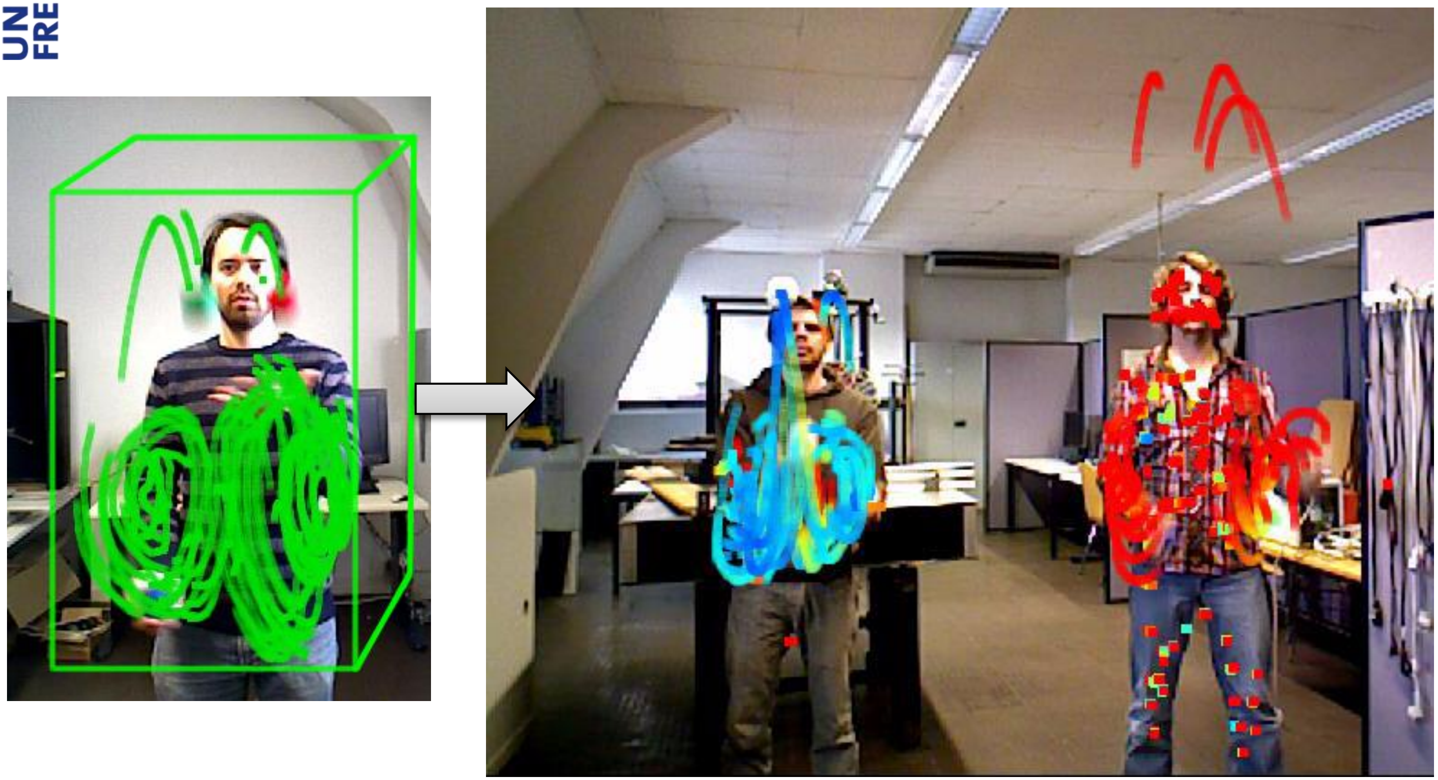

- Detect prototype instances (illustrated by bounding boxes)
- Rigid and elastic registration of prototype patterns (spatiotemporal)
- Reconstruct the whole test sequence by prototype placements
- **Training phase:** Learn prototype model (accepted variations)
- **Test phase:** Compute anomaly scores from the deviations to the aligned prototype patterns

Approach Details

- Supertrajectory representation
- Detection and elastic registration
- Learning the prototype model
- Detecting anomalies

Supertrajectory Representation

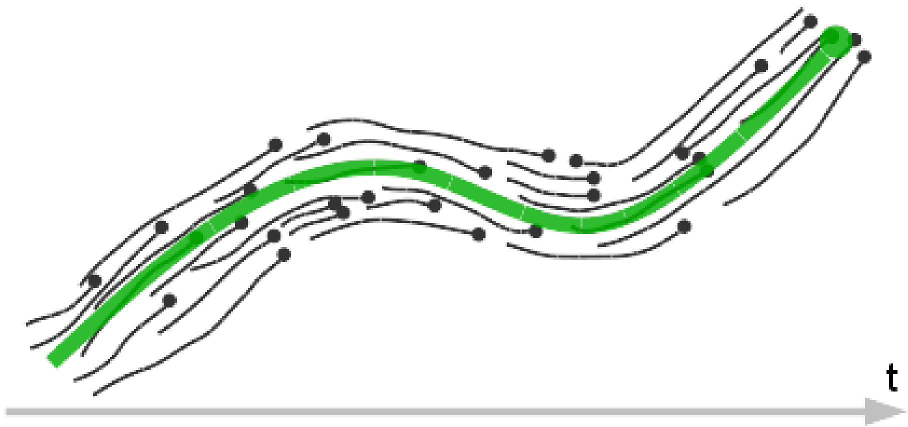

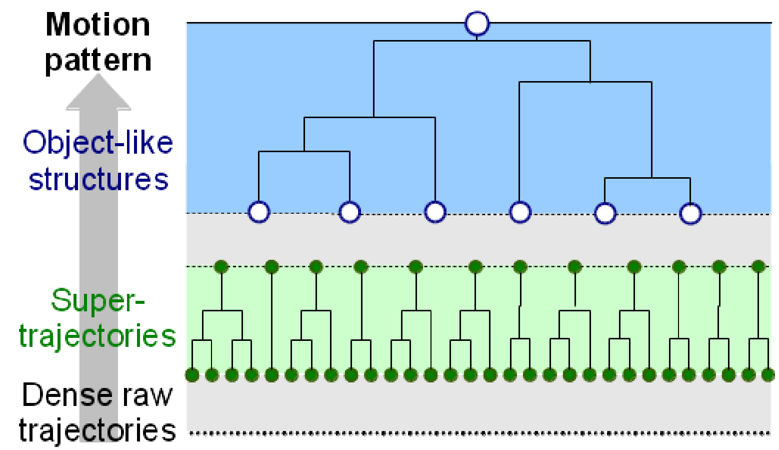

- „Supertrajectories“ serve as an efficient and robust representation
- Initialized by dense point trajectories*
- Hierarchical clustering groups trajectories in bottom-up manner
- Split the hierarchy at a certain intermediate level

*Sundaram et al., Dense point trajectories by gpu-accelerated large displacement optical flow, ECCV 2010

Supertrajectory Representation

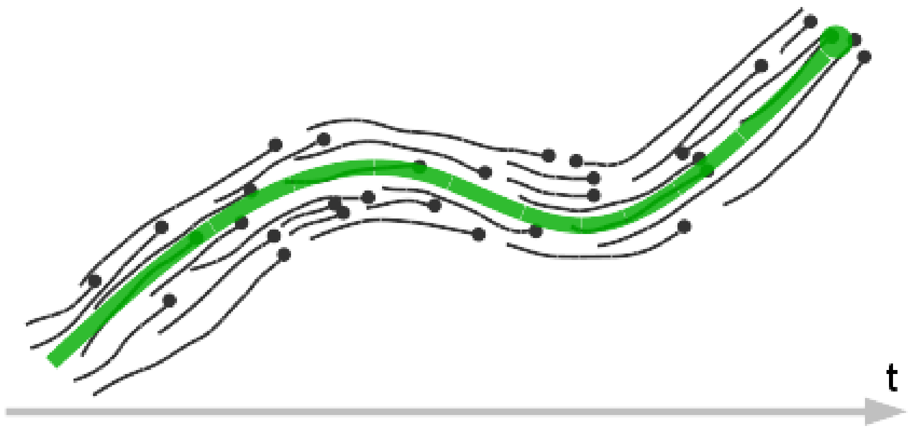

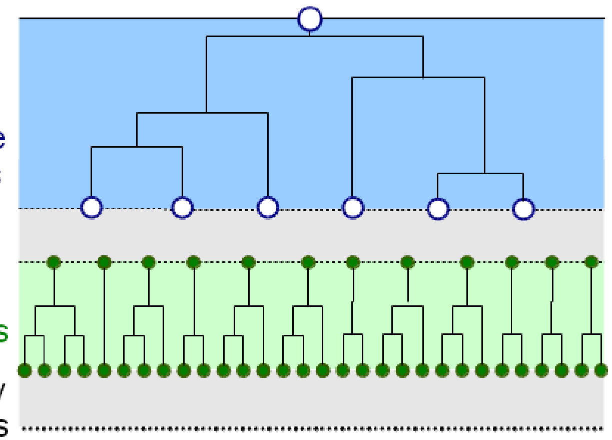

- Motion pattern definition

$$\mathbf{x}: \quad \Omega \to \mathbb{R}^3 \quad : (i,t) \mapsto \mathbf{x}(i,t), \quad i \in \{1,\ldots,N\} \subset \mathbb{N},\ t \in \mathbb{R}$$
$$w: \quad \Omega \to \{0,1\} \quad : (i,t) \mapsto w(i,t)$$

$\Omega \subset \mathbb{N} \times \mathbb{R}$ denotes domain of trajectories i and time t

Supertrajectory Representation

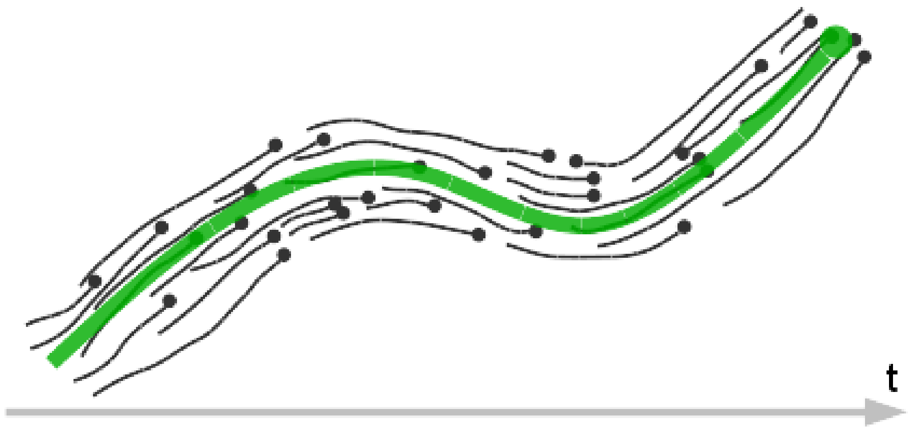

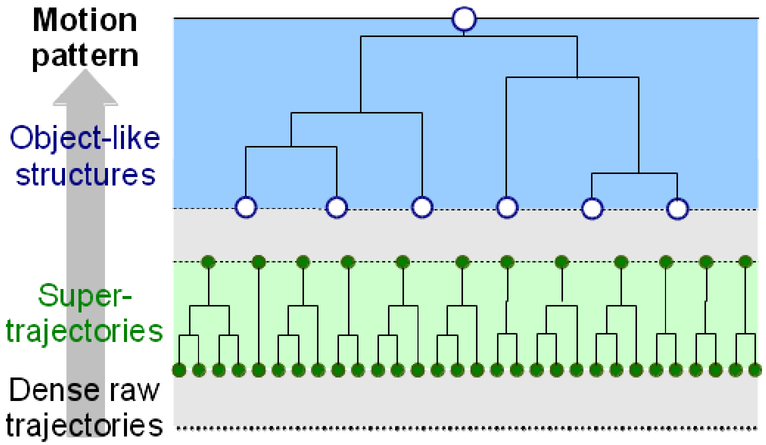

- Pairwise distances between trajectories

$$d(i,j) = \max_{t \in \mathbb{R}} (w(i,t) \cdot w(j,t) \cdot \|\mathbf{x}(i,t) - \mathbf{x}(j,t)\|)$$

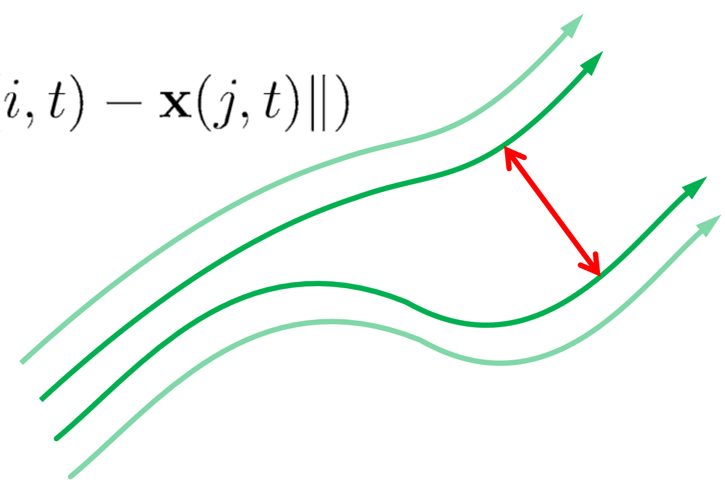

Supertrajectory Representation

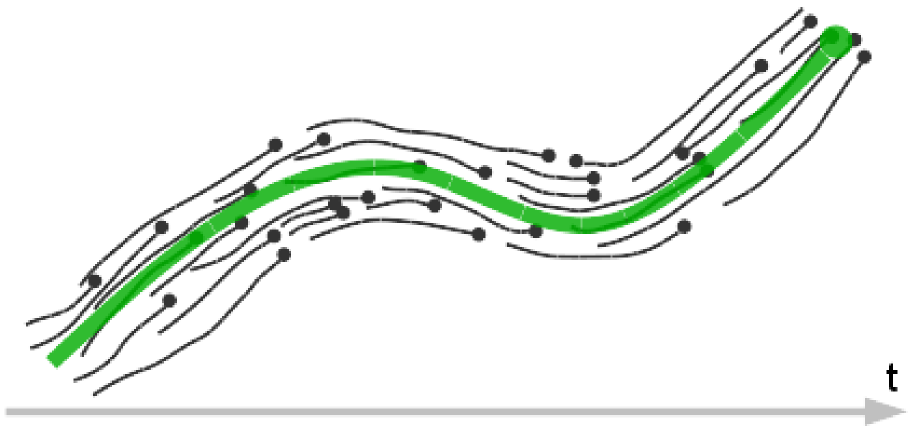

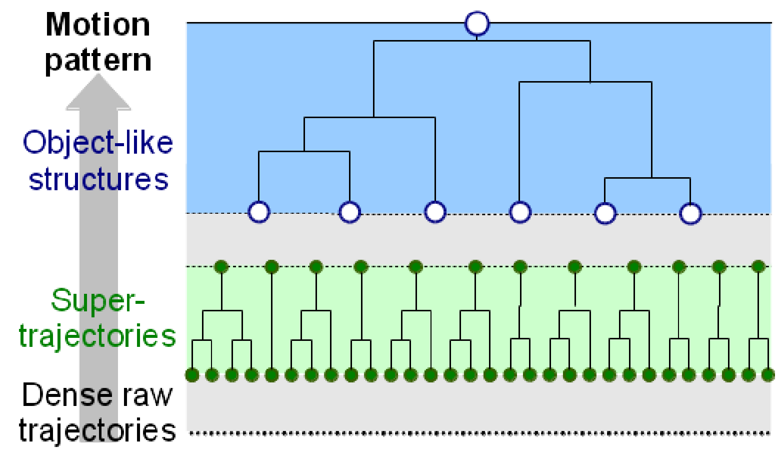

- Supertrajectory computation

$$\mathbf{x}(i,t) = \begin{cases} \frac{\sum_{i_{\text{raw}} \in \mathcal{X}_i} w_{\text{raw}}(i_{\text{raw}},t) \cdot \mathbf{x}_{\text{raw}}(i_{\text{raw}},t)}{\sum_{i_{\text{raw}} \in \mathcal{X}_i} w_{\text{raw}}(i_{\text{raw}},t)} & \text{if } w(i,t) = 1 \\ \mathbf{0} & \text{else,} \end{cases}$$

Detection and Elastic Registration

Prototype pattern
(Supertrajectories $\mathbf{x}^{a}$)

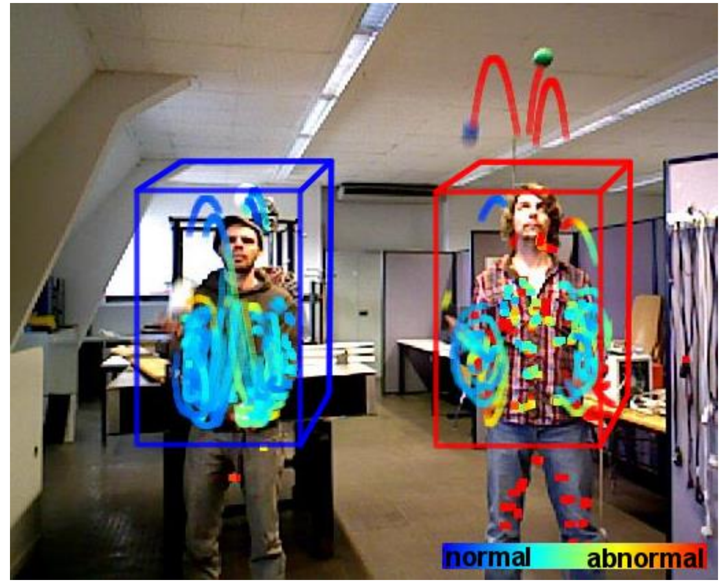

Test sequence (Supertrajectories $\mathbf{x}^{b}$)

Detection and Elastic Registration

Prototype pattern
(Supertrajectories $\mathbf{x}^a$)

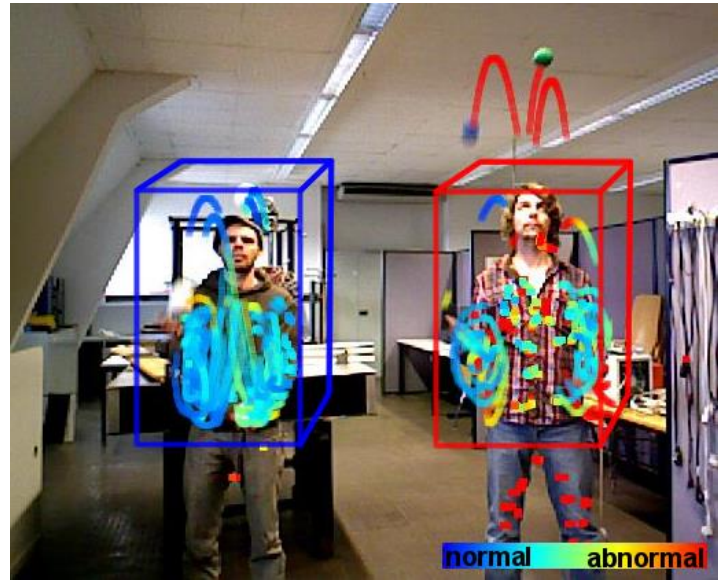

Test sequence (Supertrajectories $\mathbf{x}^b$)

- Compute detection hypotheses
  - Fast hashing approach*, extended for spatiotemporal setting
  - Temporal shift, 3D rotation and translation
  - Outputs rigid transformation hypotheses
    (temporal shift t_{shift}, spatial rigid transformation T)
- Refinement by rigid and elastic registration

*Winkelbach et al., Low-cost laser range scanner and fast surface registration approach, DAGM 2006

Rigid Pre-Alignment

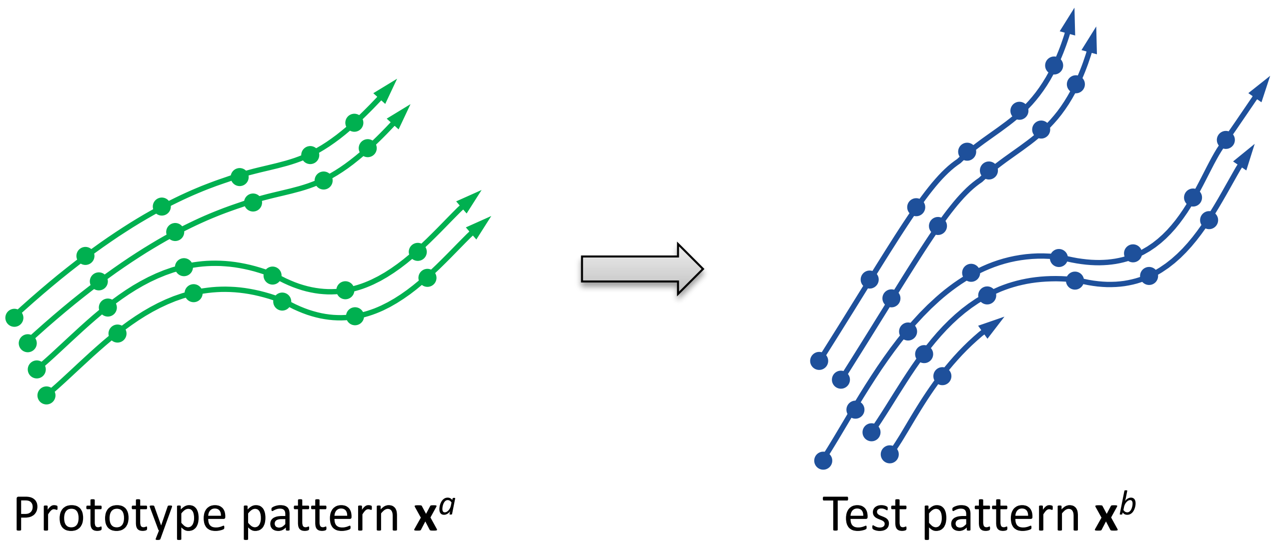

Rigid Pre-Alignment

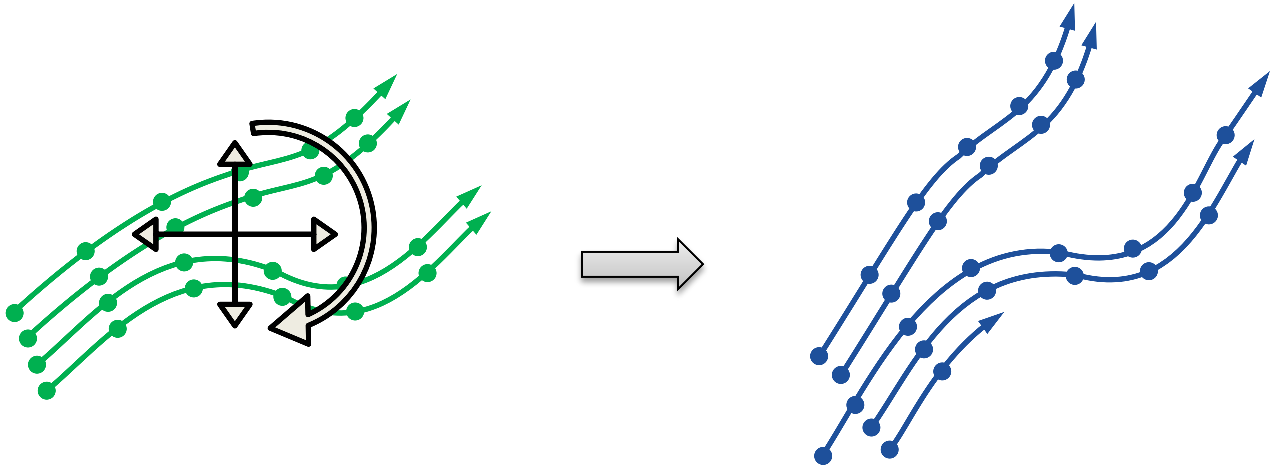

Prototype pattern $\mathbf{x}^a$

Test pattern $\mathbf{x}^b$

- Temporal shift, 3D rotation and translation

$$E_{\text{data}}(\mathbf{T}, t_{\text{shift}}, \sigma) = \sum_{\substack{(i_b,t_b)\in\Omega_b \\ w(i_b,t_b)=1 \\ \sigma(i_b,t_b)\neq 0}} \Psi\left(\|\mathbf{T}(\mathbf{x}^a(A[t_{\text{shift}}, \sigma](i_b, t_b))) - \mathbf{x}^b(i_b, t_b)\|^2\right) + \sum_{\substack{(i_b,t_b)\in\Omega_b \\ w(i_b,t_b)=1 \\ \sigma(i_b,t_b)=0}} d^2_{\text{undef}}$$

Rigid Pre-Alignment

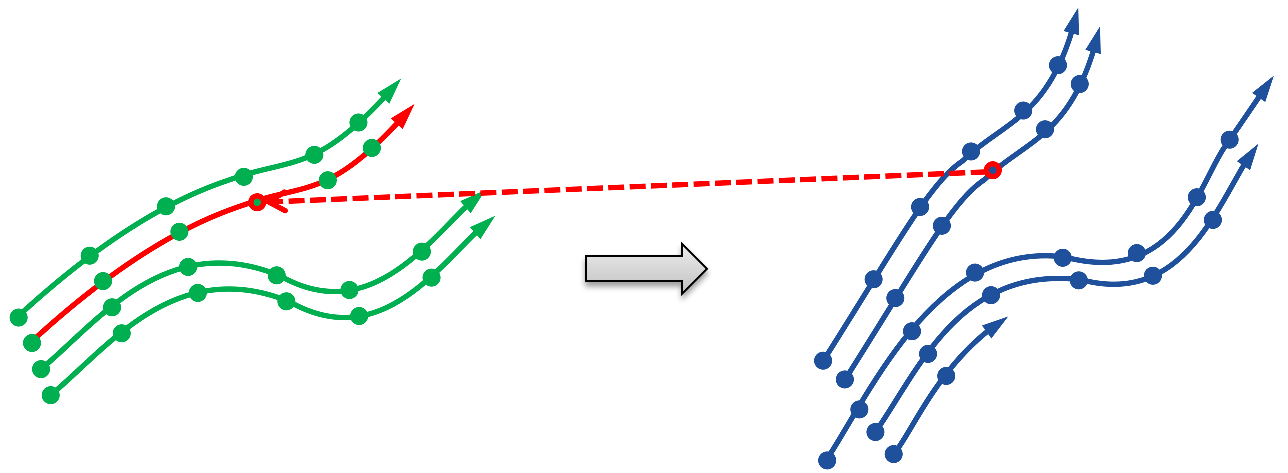

Prototype pattern $\mathbf{x}^a$ Test pattern $\mathbf{x}^b$

- Temporal shift, 3D rotation and translation
- Trajectory association function $\sigma : \Omega_b \rightarrow \mathbb{N} : (i_b, t_b) \rightarrow i_a$

$$E_{\text{data}}(\mathbf{T}, t_{\text{shift}}, \boxed{\sigma}) = \sum_{\substack{(i_b,t_b)\in\Omega_b \\ w(i_b,t_b)=1 \\ \sigma(i_b,t_b)\neq 0}} \Psi\left(\|\mathbf{T}(\mathbf{x}^a \boxed{A[t_{\text{shift}}, \sigma]}(i_b, t_b))) - \mathbf{x}^b(i_b, t_b)\|^2\right) + \sum_{\substack{(i_b,t_b)\in\Omega_b \\ w(i_b,t_b)=1 \\ \sigma(i_b,t_b)=0}} d^2_{\text{undef}}$$

Rigid Pre-Alignment

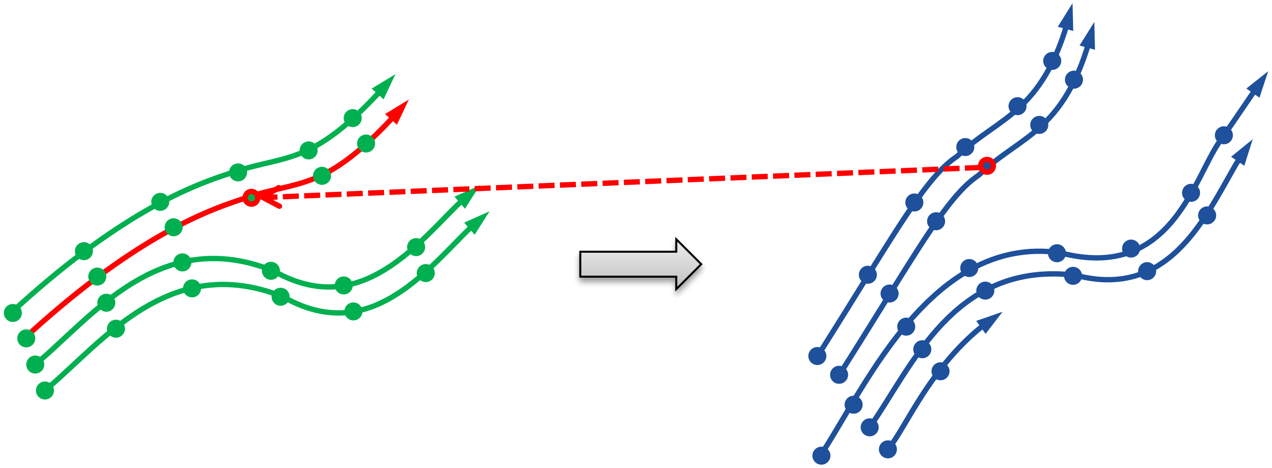

Prototype pattern $\mathbf{x}^a$ Test pattern $\mathbf{x}^b$

- Temporal shift, 3D rotation and translation
- Trajectory association function $\sigma : \Omega_b \rightarrow \mathbb{N} : (i_b, t_b) \rightarrow i_a$
- Minimize SSD of associated points

$$E_{\text{data}}(\mathbf{T}, t_{\text{shift}}, \sigma) = \sum_{\substack{(i_b,t_b)\in\Omega_b \\ w(i_b,t_b)=1 \\ \sigma(i_b,t_b)\neq 0}} \Psi\left(\left\|\mathbf{T}(\mathbf{x}^a(A[t_{\text{shift}}, \sigma](i_b, t_b))) - \mathbf{x}^b(i_b, t_b)\right\|^2\right) + \sum_{\substack{(i_b,t_b)\in\Omega_b \\ w(i_b,t_b)=1 \\ \sigma(i_b,t_b)=0}} d^2_{\text{undef}}$$

Rigid Pre-Alignment

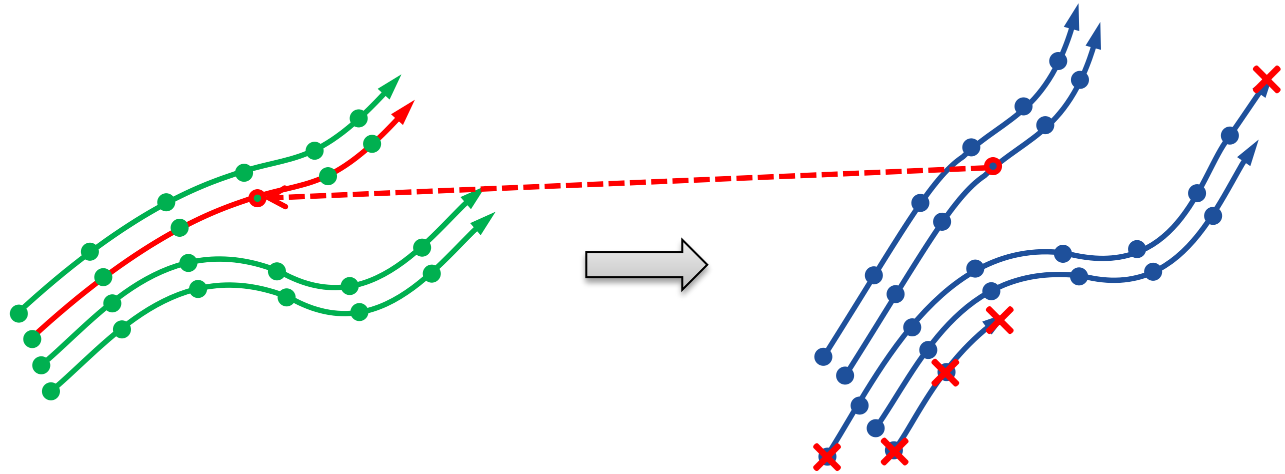

Prototype pattern $\mathbf{x}^a$ Test pattern $\mathbf{x}^b$

- Temporal shift, 3D rotation and translation
- Trajectory association function $\sigma : \Omega_b \rightarrow \mathbb{N} : (i_b, t_b) \rightarrow i_a$
- Minimize SSD of associated points
- Truncated squared norm (function Ψ)
- Penalize unassociated points (d_{undef})

$$E_{\text{data}}(\mathbf{T}, t_{\text{shift}}, \sigma) = \sum_{\substack{(i_b,t_b)\in\Omega_b \\ w(i_b,t_b)=1 \\ \sigma(i_b,t_b)\neq 0}} \boxed{\Psi}\left(\|\mathbf{T}(\mathbf{x}^a(A[t_{\text{shift}}, \sigma](i_b, t_b))) - \mathbf{x}^b(i_b, t_b)\|^2\right) + \sum_{\substack{(i_b,t_b)\in\Omega_b \\ w(i_b,t_b)=1 \\ \sigma(i_b,t_b)=0}} \boxed{d^2_{\text{undef}}}$$

Elastic Registration

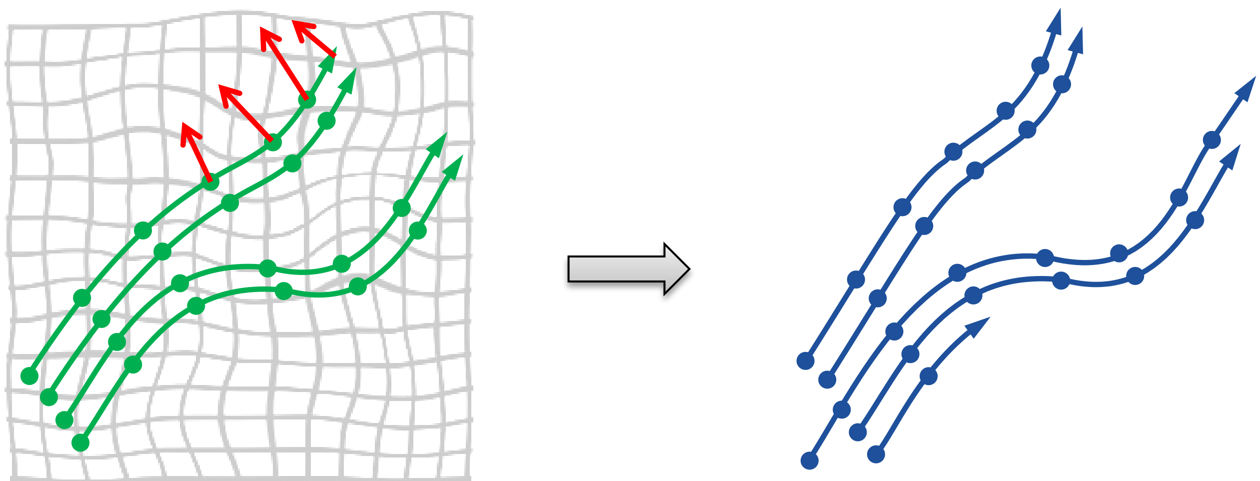

Prototype pattern $\mathbf{x}^a$ Test pattern $\mathbf{x}^b$

- Deformation, temporal warping

$$\mathbf{u}(i,t) \quad : \Omega_a \rightarrow \mathbb{R}^3$$
$$\tau(i,t) \quad : \Omega_a \rightarrow \mathbb{R}$$

Elastic Registration

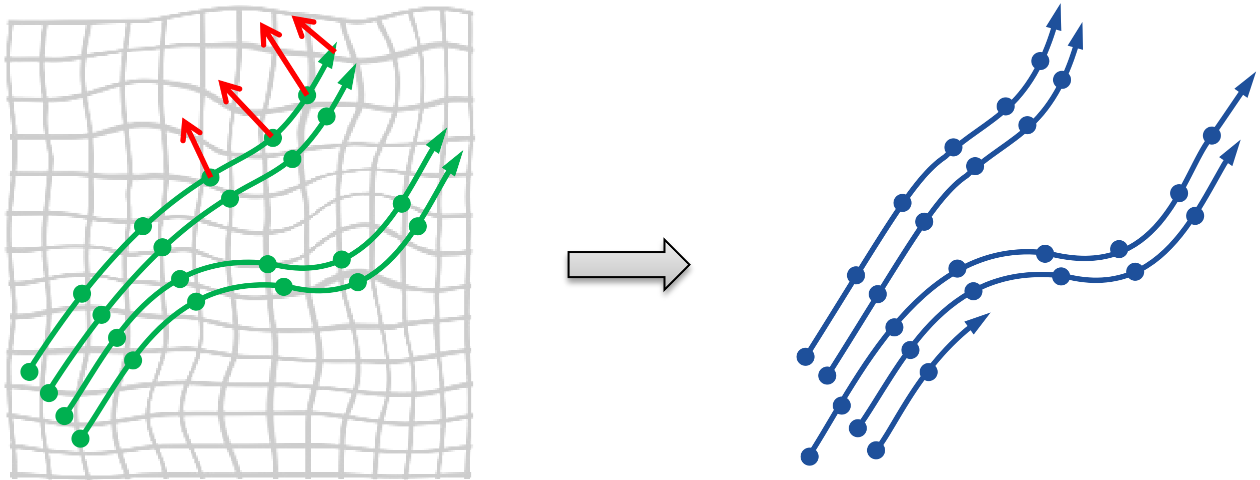

Prototype pattern $\mathbf{x}^a$ Test pattern $\mathbf{x}^b$

- Deformation, temporal warping — $\mathbf{u}(i,t) \quad : \Omega_a \rightarrow \mathbb{R}^3$
- **Data term** analogous to $E_{data}(\mathbf{T}, t_{shift}, \sigma)$ — $\tau(i,t) \quad : \Omega_a \rightarrow \mathbb{R}$

$$E(\mathbf{u},\tau,\sigma) = \boxed{E_{\text{data}}(\mathbf{u},\tau,\sigma)} + \alpha_{\text{spatial}} E_{\text{spatial}}(\mathbf{u},\tau) + \alpha_{\text{temp}} E_{\text{temp}}(\mathbf{u},\tau) + \alpha_{\text{assign}} E_{\text{assign}}(\sigma)$$

Elastic Registration

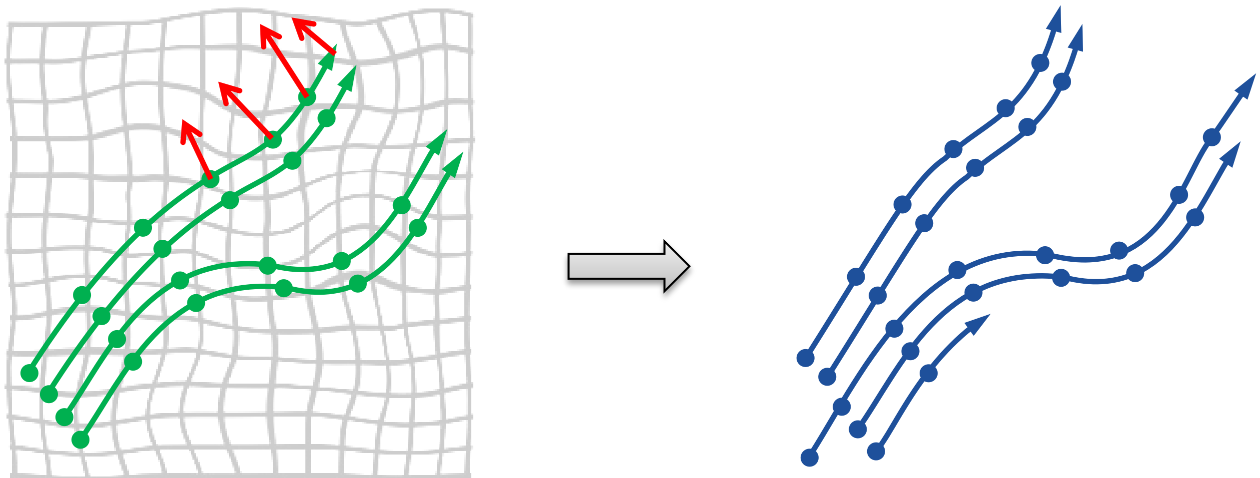

Prototype pattern $\mathbf{x}^a$ — Test pattern $\mathbf{x}^b$

- Deformation, temporal warping
- **Data term** analogous to $E_{\text{data}}(\mathbf{T}, t_{\text{shift}}, \sigma)$
- **Smoothness terms**

$$\mathbf{u}(i,t) \quad : \Omega_a \to \mathbb{R}^3$$
$$\tau(i,t) \quad : \Omega_a \to \mathbb{R}$$

$$E(\mathbf{u},\tau,\sigma) = E_{\text{data}}(\mathbf{u},\tau,\sigma) + \boxed{\alpha_{\text{spatial}} E_{\text{spatial}}(\mathbf{u},\tau) + \alpha_{\text{temp}} E_{\text{temp}}(\mathbf{u},\tau) + \alpha_{\text{assign}} E_{\text{assign}}(\sigma)}$$

Elastic Registration

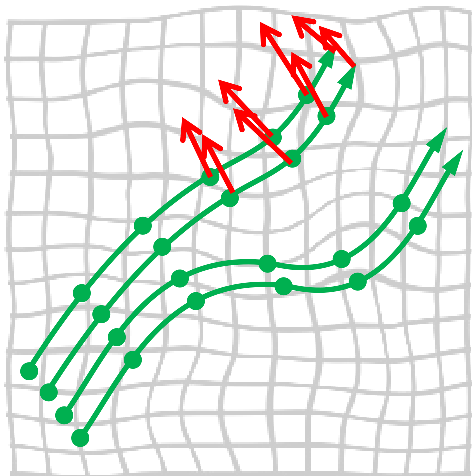

Prototype pattern $\mathbf{x}^a$

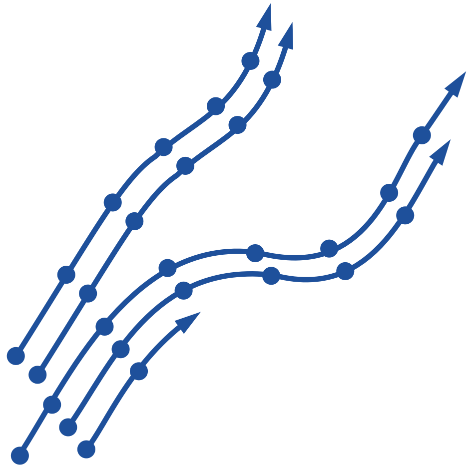

Test pattern $\mathbf{x}^b$

- Deformation, temporal warping $\quad \mathbf{u}(i,t) \quad : \Omega_a \to \mathbb{R}^3$
- **Data term** analogous to $E_{\text{data}}(\mathbf{T}, t_{\text{shift}}, \sigma)$ $\quad \tau(i,t) \quad : \Omega_a \to \mathbb{R}$
- **Spatial smoothness** (across trajectories)

$$E_{\text{spatial}}(\mathbf{u},\tau) = \sum_{\substack{i,j,t \\ (i,t)\in\Omega_a \wedge (j,t)\in\Omega_a \\ w(i,t)=1 \wedge w(j,t)=1}} C(i,j) \cdot \left(\|\mathbf{u}(i,t) - \mathbf{u}(j,t)\|^2 + \beta_{\text{temp}}(\tau(i,t) - \tau(j,t))^2 \right)$$

$$E(\mathbf{u},\tau,\sigma) = E_{\text{data}}(\mathbf{u},\tau,\sigma) + \alpha_{\text{spatial}} E_{\text{spatial}}(\mathbf{u},\tau) + \alpha_{\text{temp}} E_{\text{temp}}(\mathbf{u},\tau) + \alpha_{\text{assign}} E_{\text{assign}}(\sigma)$$

Elastic Registration

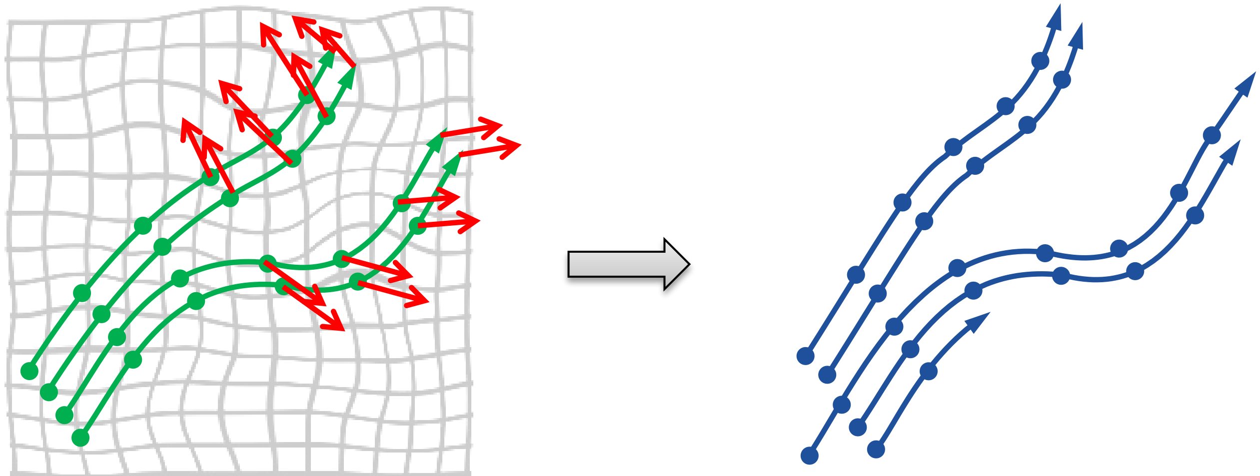

Prototype pattern $\mathbf{x}^a$ | Test pattern $\mathbf{x}^b$

- Deformation, temporal warping $\quad \mathbf{u}(i,t) \;:\Omega_a \to \mathbb{R}^3$
- **Data term** analogous to $E_{data}(\mathbf{T}, t_{shift}, \sigma)$ $\quad \tau(i,t) \;:\Omega_a \to \mathbb{R}$
- **Spatial smoothness** (across trajectories)

$$E_{\text{spatial}}(\mathbf{u},\tau) = \sum_{\substack{i,j,t \\ (i,t)\in\Omega_a \wedge (j,t)\in\Omega_a \\ w(i,t)=1 \wedge w(j,t)=1}} C(i,j) \cdot \left(\|\mathbf{u}(i,t) - \mathbf{u}(j,t)\|^2 + \beta_{\text{temp}}(\tau(i,t) - \tau(j,t))^2 \right)$$

$$C(i,j) = \exp(-d(i,j)^2/2r^2)$$

$$E(\mathbf{u},\tau,\sigma) = E_{\text{data}}(\mathbf{u},\tau,\sigma) + \alpha_{\text{spatial}} E_{\text{spatial}}(\mathbf{u},\tau) + \alpha_{\text{temp}} E_{\text{temp}}(\mathbf{u},\tau) + \alpha_{\text{assign}} E_{\text{assign}}(\sigma)$$

Energy Optimization

- Alternating optimization

  1) **Assignment function** σ computed by dynamic programming

  2) **Transformation**

     - Rigid: Procrustes algorithm
     - Elastic: L-BFGS optimization

  => Repeated until convergence

- Both parts solved globally optimal (convex energies)

Learning the Prototype Model

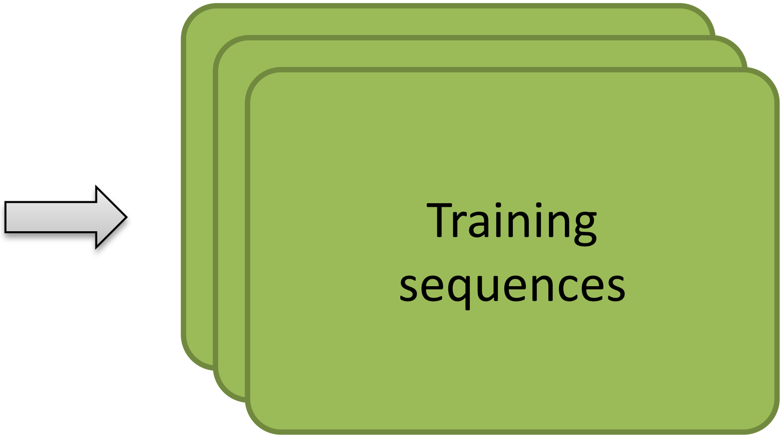

- Registration of the prototype to multiple instances in the training data
- Set of rigid and elastic transformation parameters
- **Statistical model**
  - **Global deformations** and data fitting costs

    => Defines bounds for validating prototype registrations
  - **Residual distances**

    => Locally aggregated for each prototype pattern point (residual model)

Detecting Anomalies

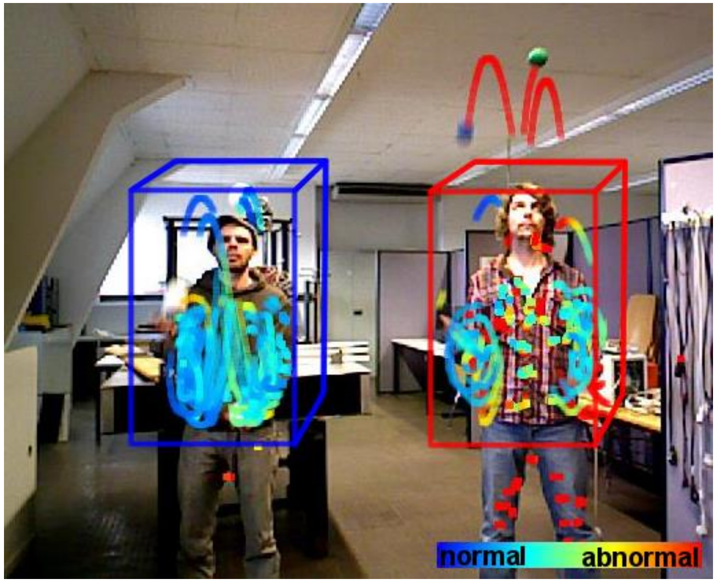

- Reconstruction by prototype placements (greedy search)
  - Iteratively find best placements of prototype patterns (candidates from detections)
  - Stop, if pattern is reconstructed completely (or if no candidates remain that can improve reconstruction)
- Pointwise and framewise anomaly score

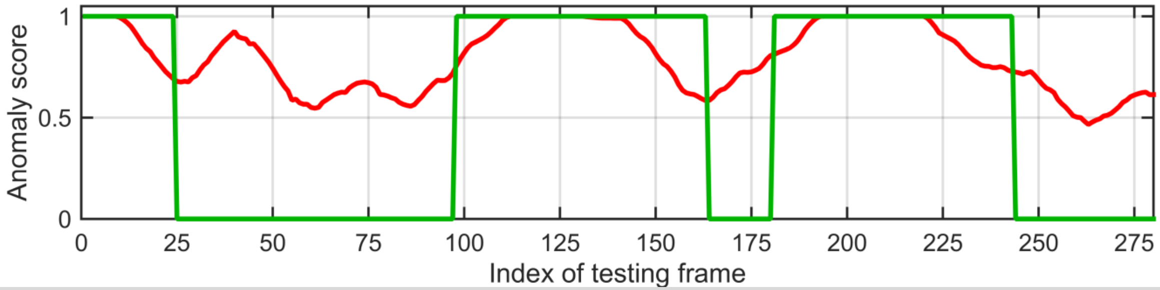

Experiments – Motion Anomaly Dataset

- Recorded juggling sequences (> 10.000 frames)
- Different jugglers, viewpoints, various anomalies, background motion

Anomaly Detection in Juggling Patterns



- Learn a protoype model from 3 training sequences only

Anomaly Detection in Juggling Patterns

Subset A (different viewpoints)

Subset B (similar viewpoints)

- Learn a protoype model from 3 training sequences only
- Split 29 test sequences in two subsets

Anomaly Detection in Juggling Patterns

Subset A (different viewpoints)

Subset B (similar viewpoints)

- Learn a protoype model from 3 training sequences only
- Split 29 test sequences in two subsets
- Compare against Chaotic invariants (CI) for anomaly detection in crowded scenes*

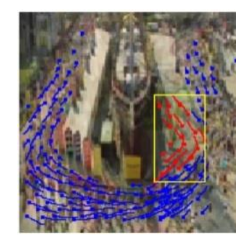

*Wu et al., Chaotic invariants of lagrangian particle trajectories for anomaly detection in crowded scenes, CVPR 2010

Anomaly Detection in Juggling Patterns

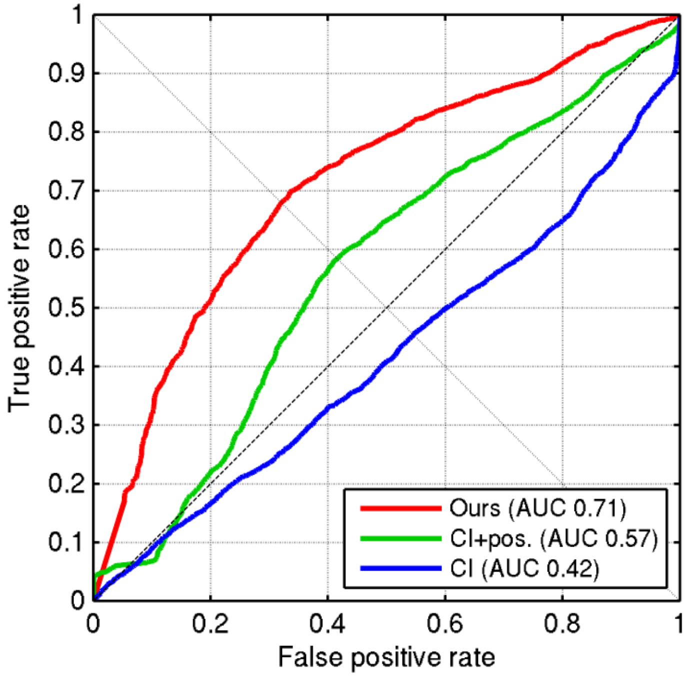

Subset A
(different viewpoints)

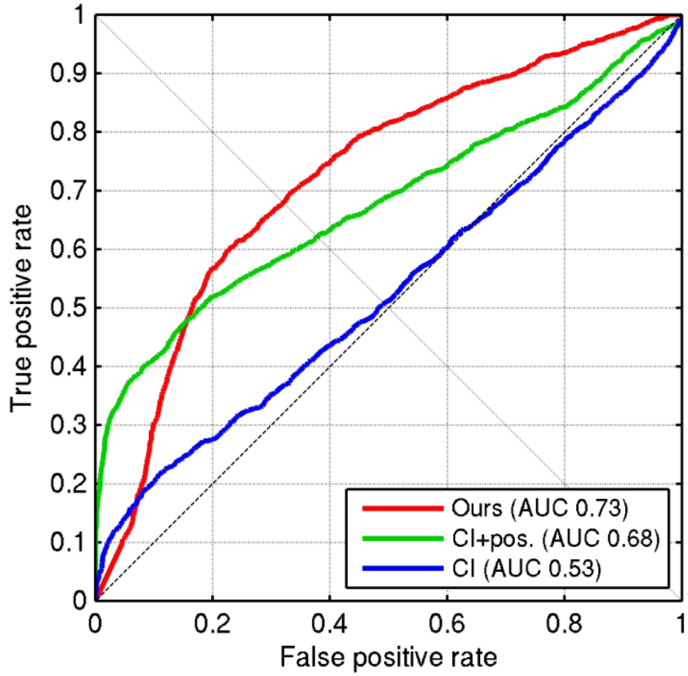

Subset B
(similar viewpoints)

- Evaluate framewise anomaly detection
- ROC curves from thresholding anomaly profile

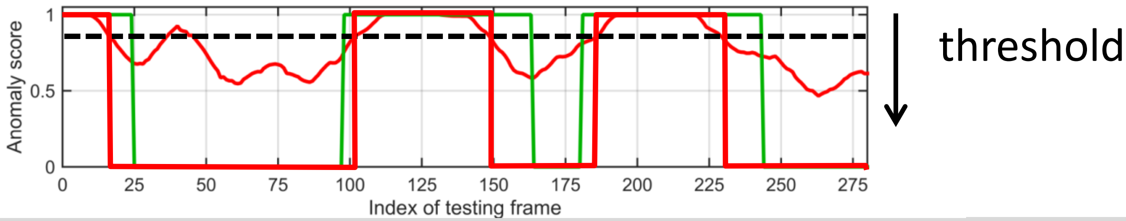

Anomaly Detection in Juggling Patterns

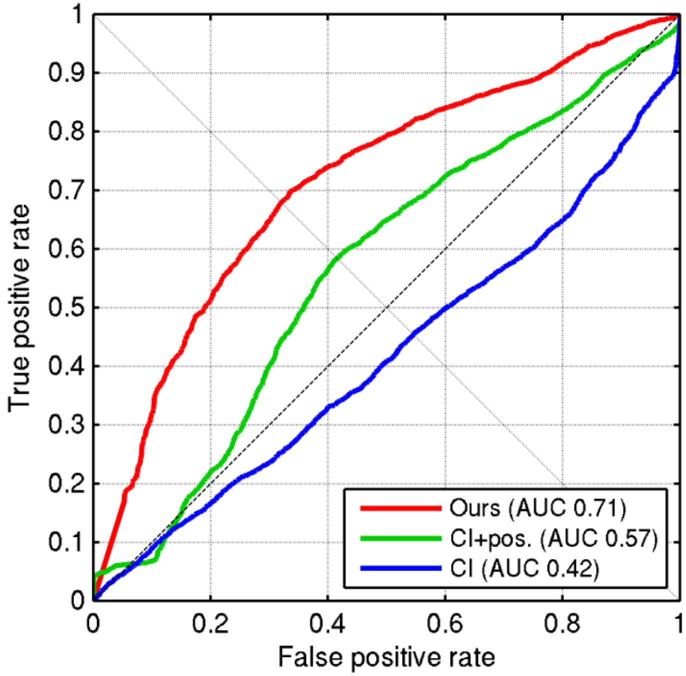

Subset A
(different viewpoints)

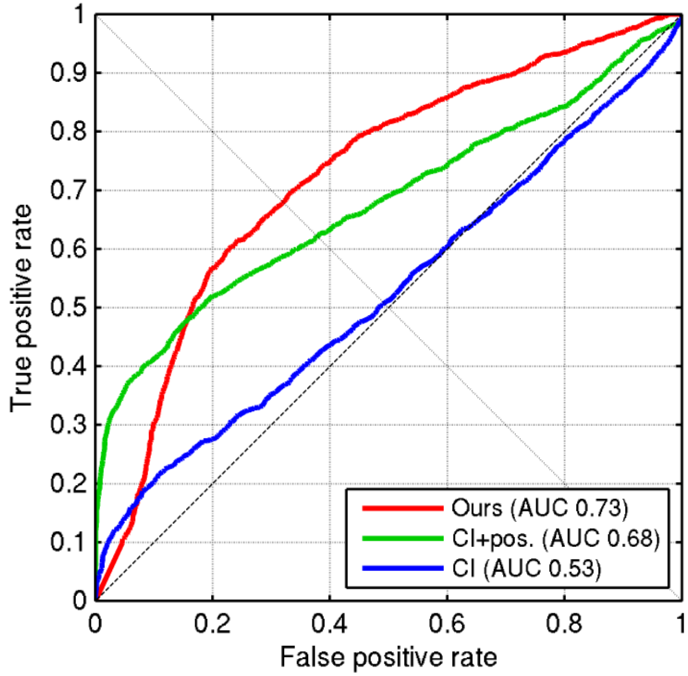

Subset B
(similar viewpoints)

Method	Subset A	Subset B	All sequences
CI (2D) [21]	0.42	0.53	0.46
CI+pos. (2D) [21]	0.57	0.68	0.62
CI (3D)	0.42	0.41	0.42
CI+pos. (3D)	0.51	0.55	0.53
Our method (3D)	**0.71**	**0.73**	**0.72**

Anomaly Detection in Biological Patterns*

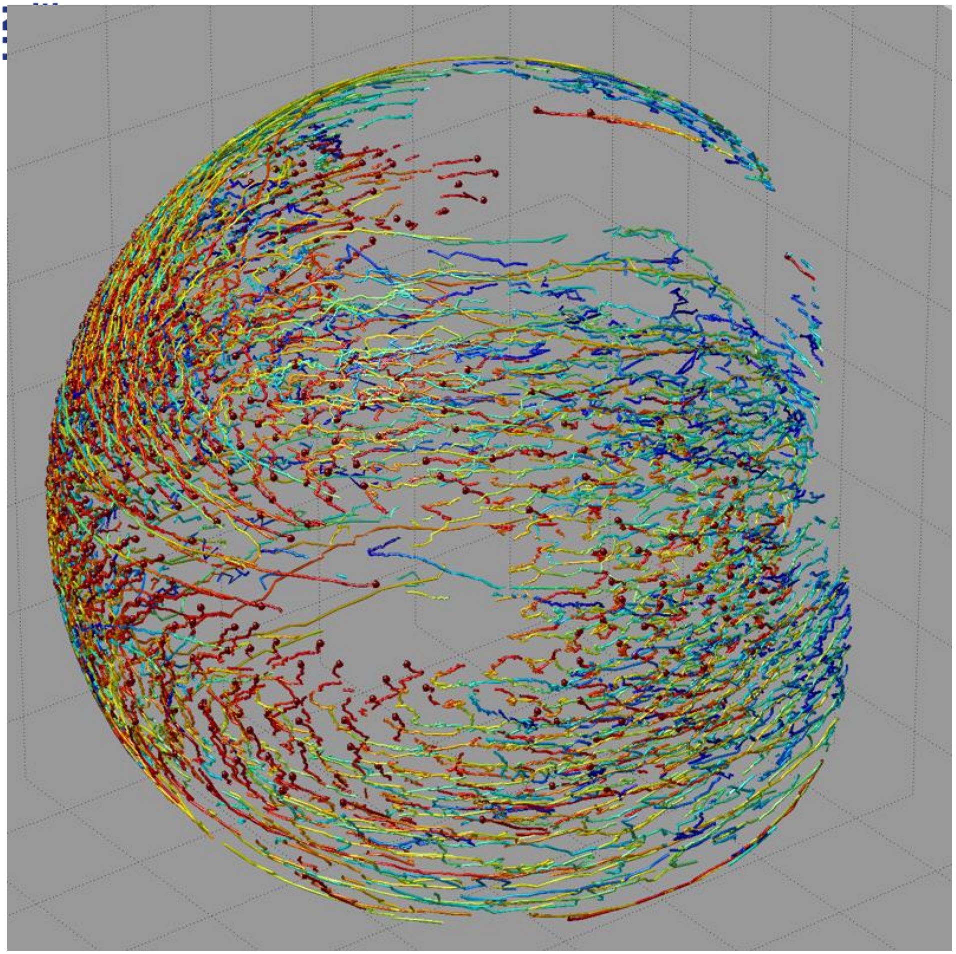

Wild type pattern (WT)

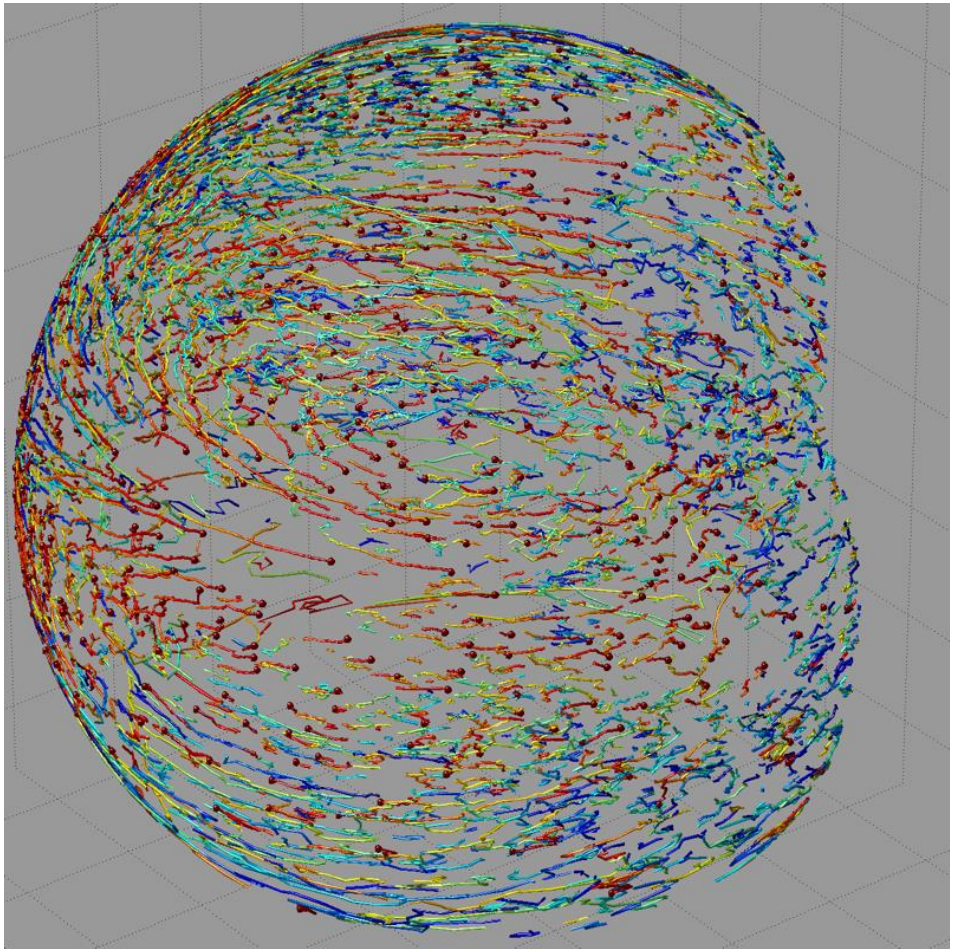

Morphant pattern (MO)

*__Trajectory data__ from** provided by Nico Scherf (Institute for Medical Informatics and Biometry, TU Dresden) and Jan Huisken (Max Planck Institudte of Molecular Cell Biology and Genetics (MPI-CBG), Dresden),
** **Schmid et al.**, High-speed panoramic light-sheet microscopy reveals global endodermal cell dynamics. *Nat Commun*, 2013

Anomaly Detection in Biological Patterns

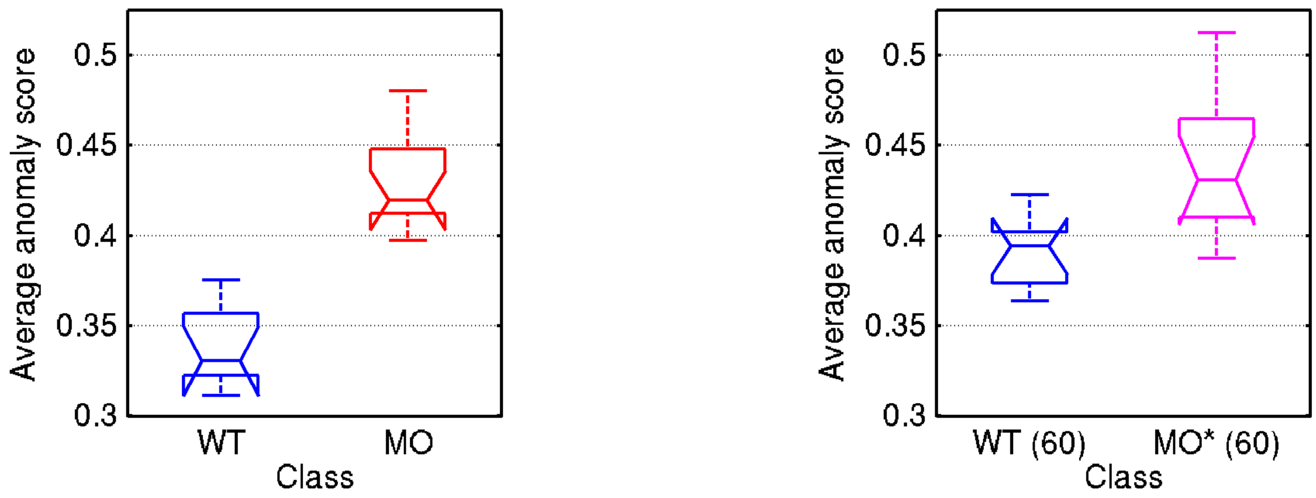

WT vs. MO   WT vs. MO*

- Learn prototype model (3 WT patterns)
- Test on remaining 9 WT and 12 MO patterns
- Compare global anomaly scores

Results:

1) Significant difference between WT and MO patterns
2) Time scaling (MO*) partially compensates the differences

Conclusion

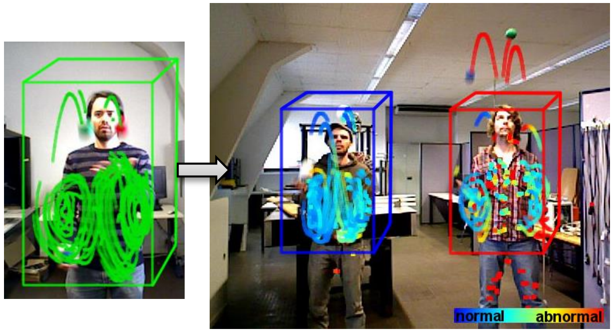

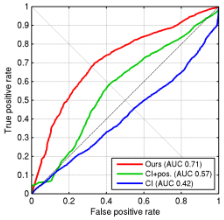

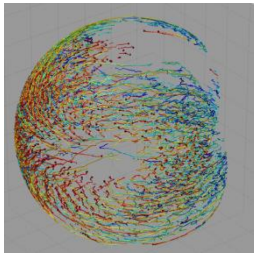

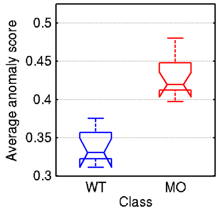

- New approach to motion anomaly detection
- Detect and localize subtle anomalies
- New motion anomaly dataset
- Important application area in biomedical image analysis

- **MATLAB code and datasets** (in preparation):
  http://lmb.informatik.uni-freiburg.de/resources/opensource/AnomalyDetection/

- **Acknowledgements**
  - Jugglers: J. Koch, A. Krämer, T. Paxian and D. Mai
  - Trajectory data: Nico Scherf (TU Dresden) and Jan Huisken (MPI-CBG, Dresden)
  - This study was supported by the Excellence Initiative of the German Federal and State Governments (EXC 294).

Thank you!