Comparing Morse-Smale Complexes via Co-Optimal Transport

Guangyu Meng, Mingzhe Li, Erin W. Chambers

University of Notre Dame

Foundations of Computational Geometry and Topology

ICERM, May 2026

Scalar Field in Science

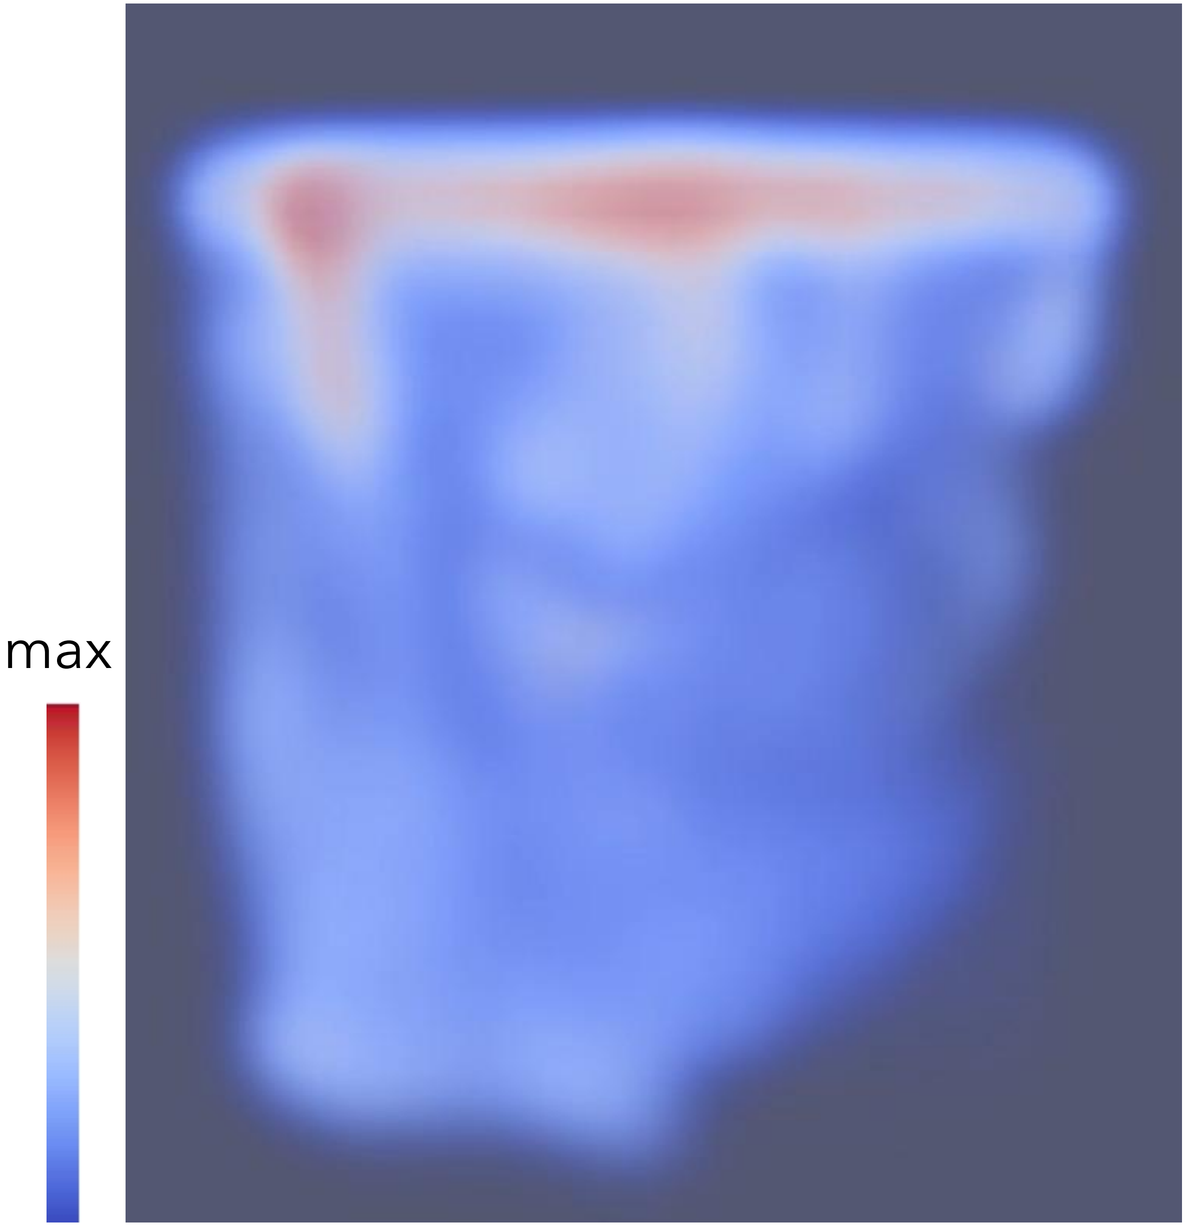

3D transient fluid flow simulation
(concentration field)

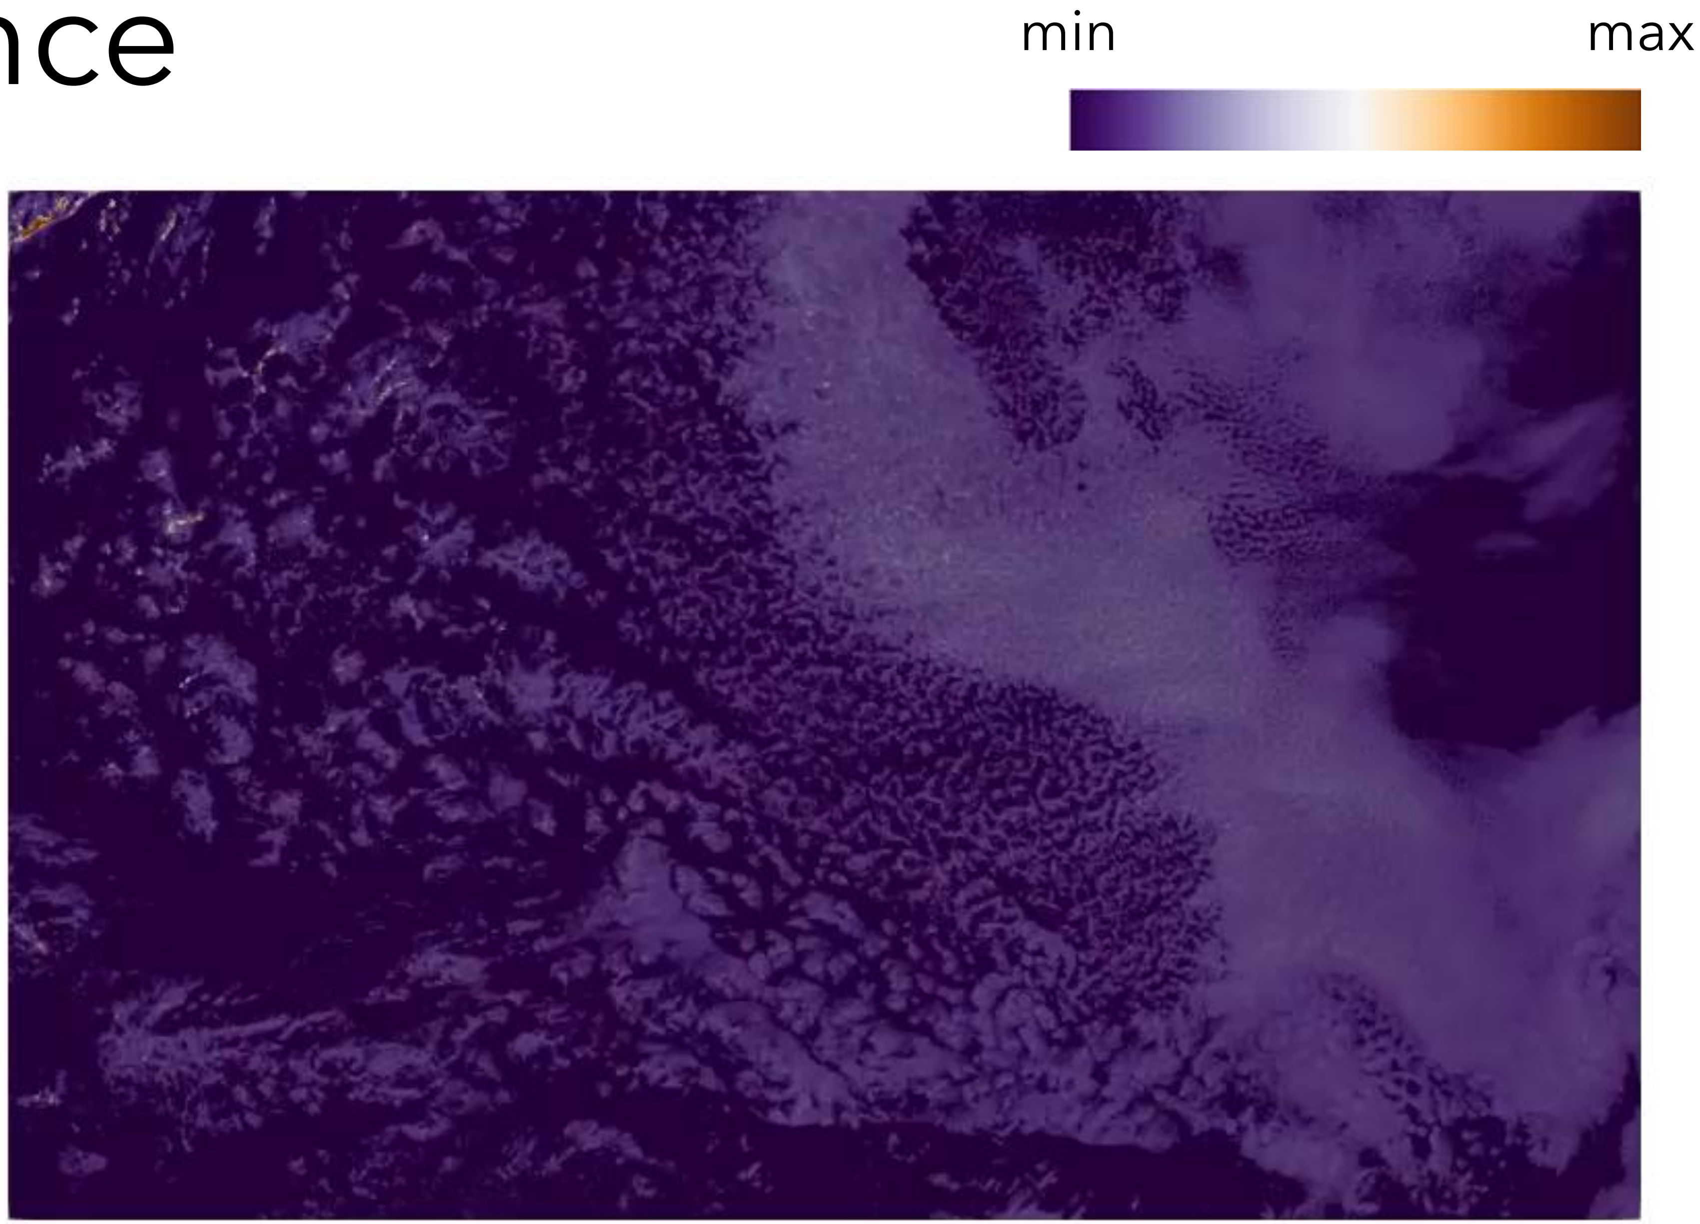

Marine Clouds: Time-varying cloud optical depth from satellite images of oceanic areas

Topological Descriptors

Topology-based descriptors capture the structure of scalar fields at multiple levels of detail:

Merge Tree

Hierarchical nesting of sub-level sets; tracks how components merge

Morse(-Smale) Complex

Decomposes domain by gradient flow into regions

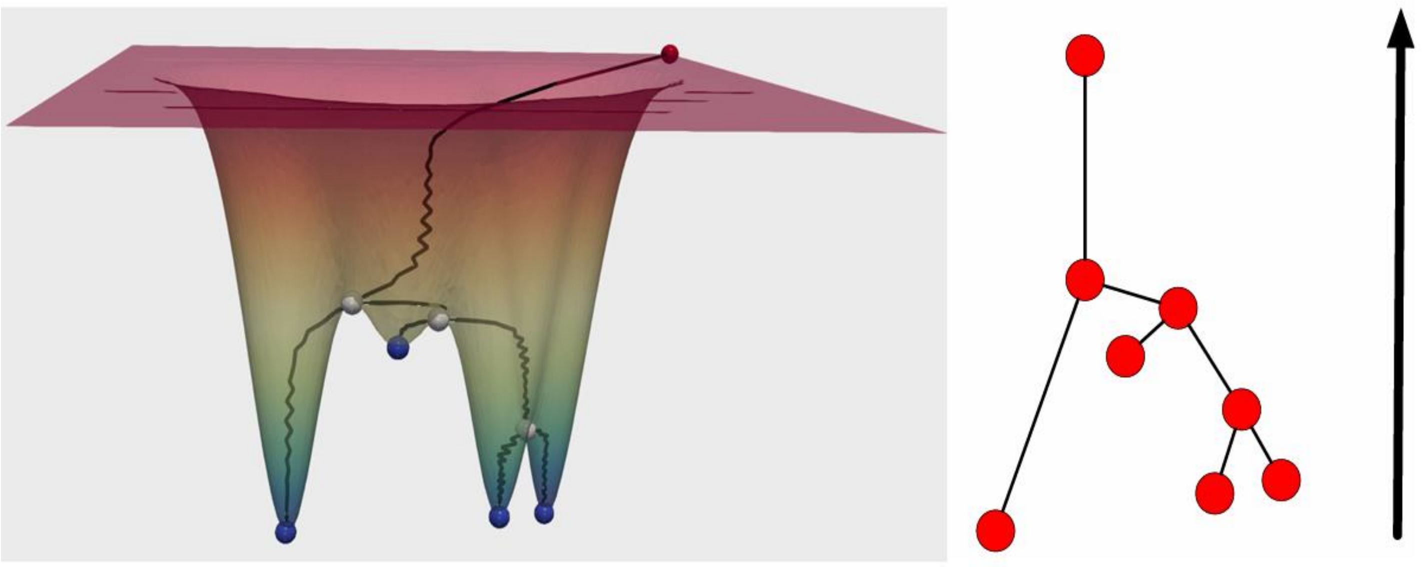

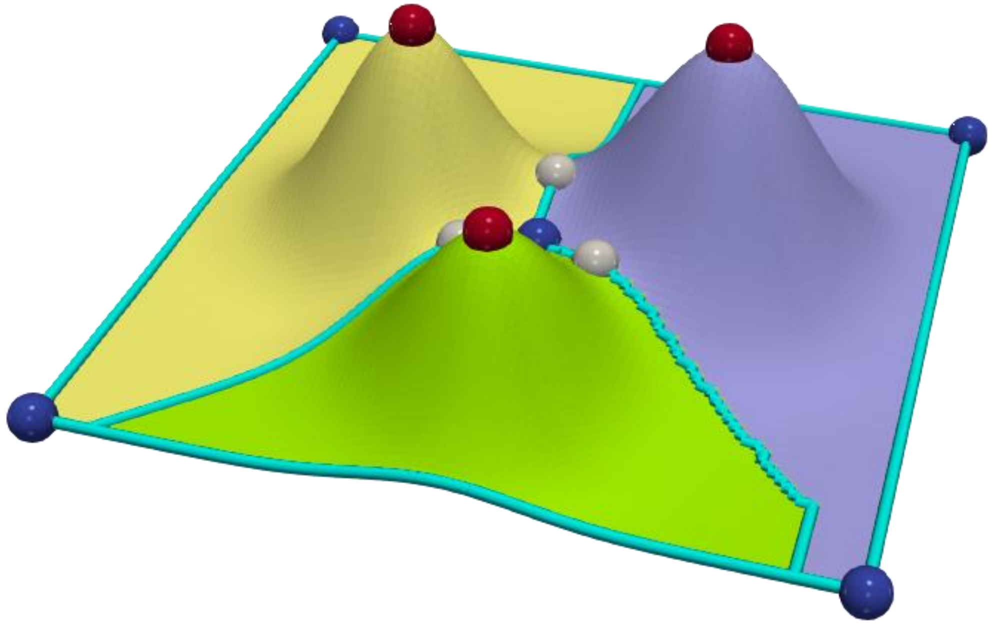

Optimal Transport (OT)

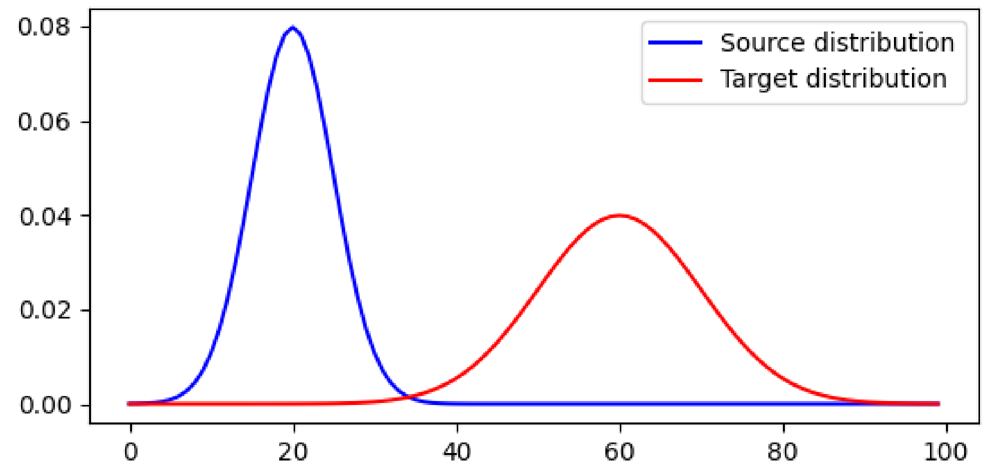

POT documentation
Flamary et al.

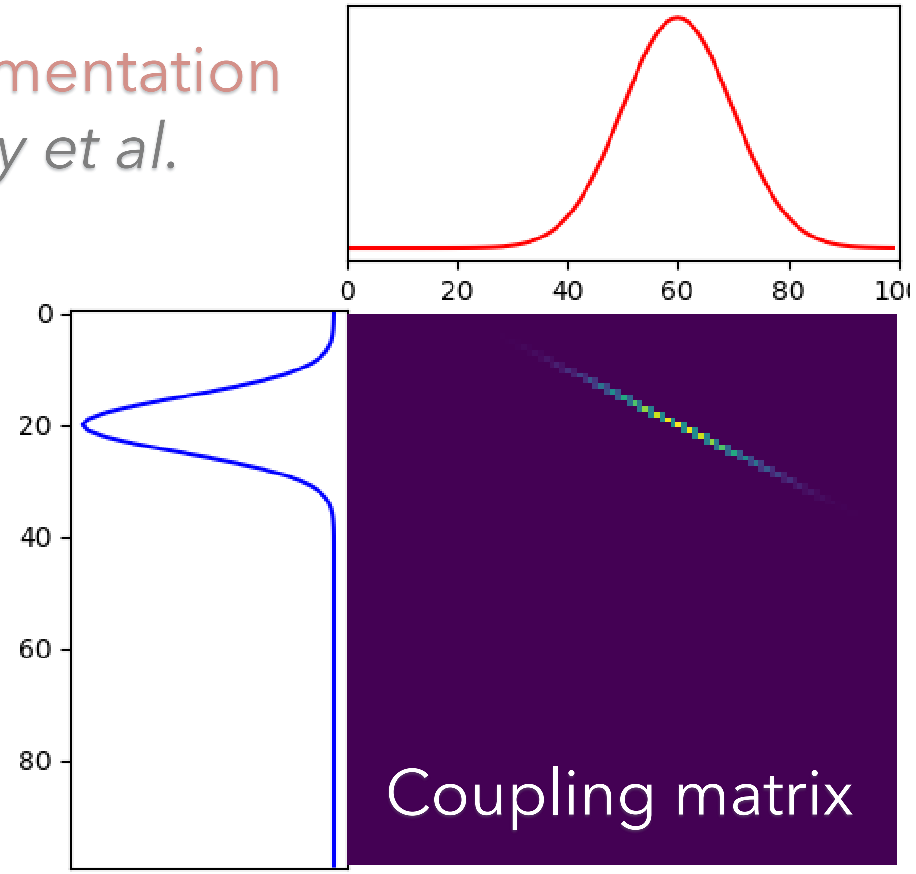

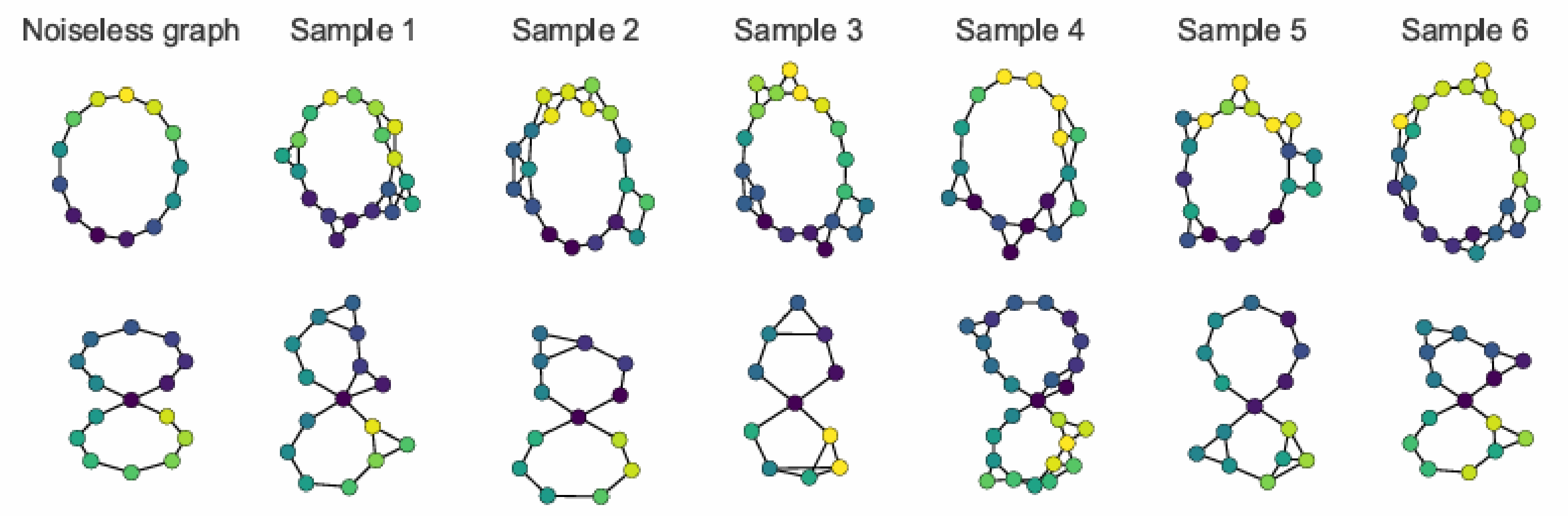

Graph matching and shape matching
Vayer et al. 2019

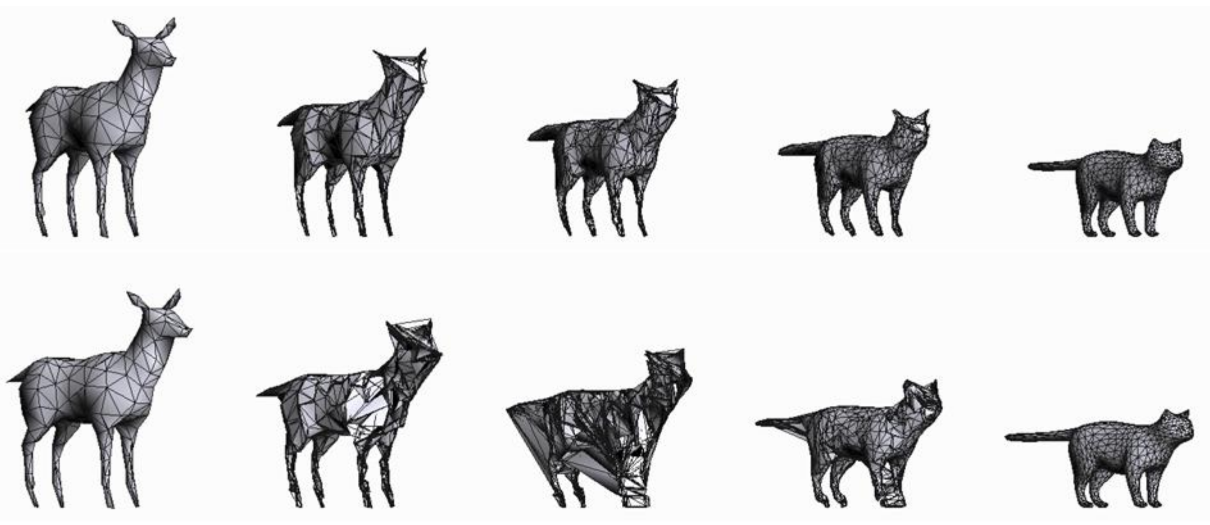

Line of Research:
Optimal Transport (OT) + Topological descriptors

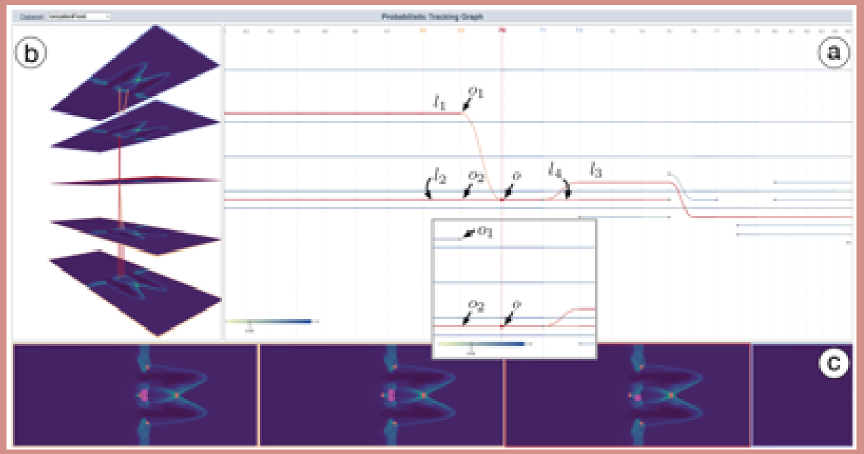

Flexible and Probabilistic Topology Tracking

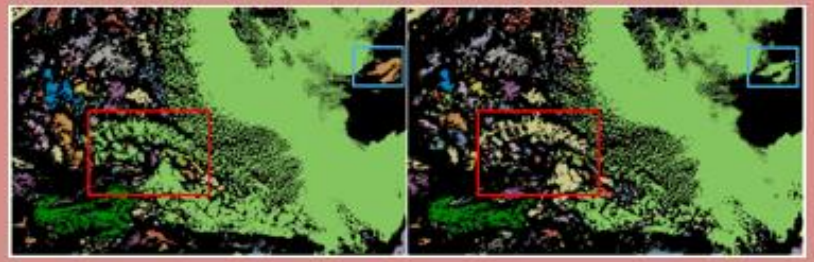

Tracking Low-Level Cloud Systems with Topology

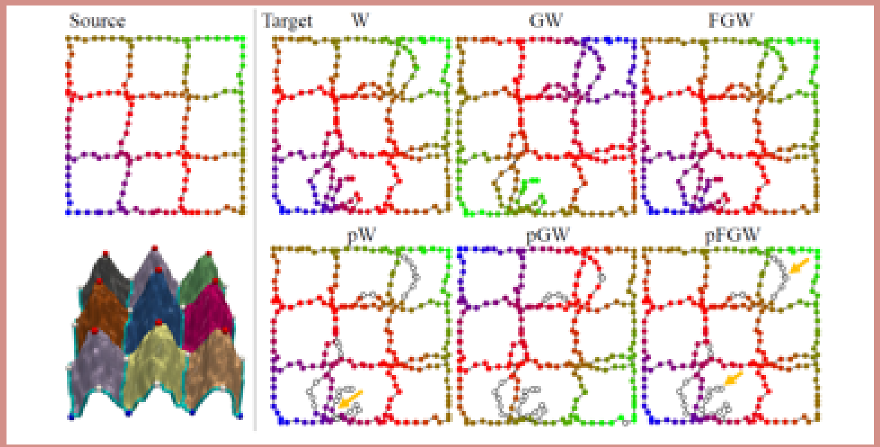

Comparing Morse Complexes using Optimal Transport

Line of Research: Optimal Transport (OT) + Topological descriptors

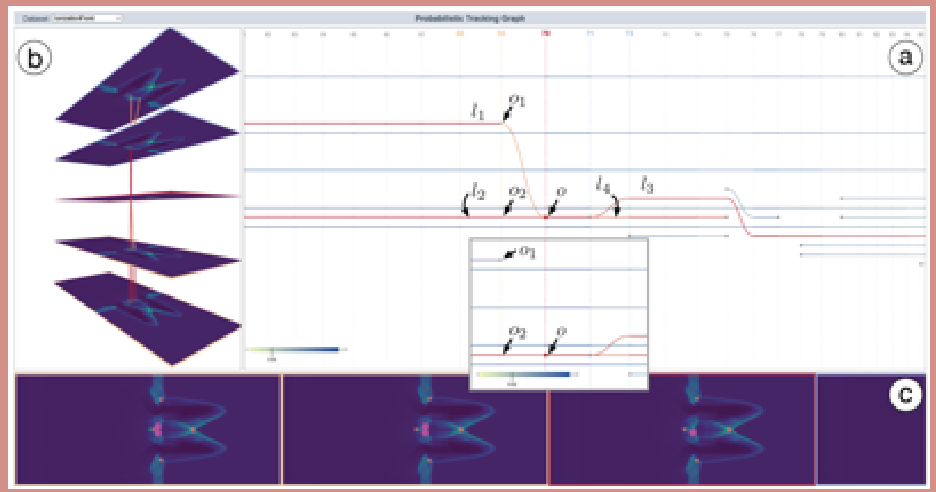

Flexible and Probabilistic Topology Tracking

Topological feature tracking using partial optimal transport

- Encoding spatial and topological information
- Partial matching allows ignoring noises
- Provide probabilistic tracking results

Mingzhe Li, Xinyuan Yan, Lin Yan, Tom Needham, Bei Wang, "Flexible and Probabilistic Topology Tracking with Partial Optimal Transport", IEEE Transactions on Visualization and Computer Graphics, 2025

Line of Research:
Optimal Transport (OT) + Topological descriptors

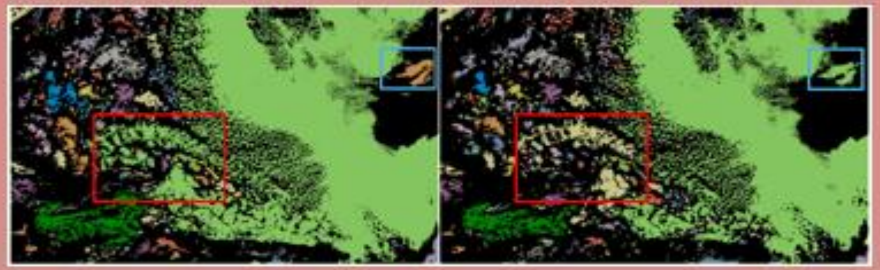

Tracking Low-Level Cloud Systems with Topology

Application: Cloud tracking

- Robust to complex cloud dynamics
- Robust for cloud systems with different characteristics

Mingzhe Li, Dwaipayan Chatterjee, Franziska Glassmeier, Fabian Senf, Bei Wang, "Tracking Low-Level Cloud Systems with Topology", In IEEE TopoInVis, 2025

Line of Research:
Optimal Transport (OT) + Topological descriptors

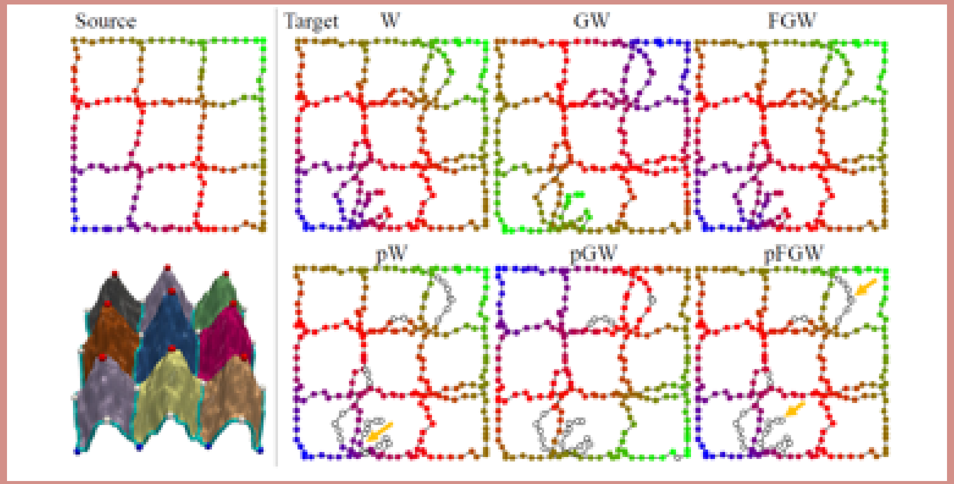

Comparing Morse Complexes using Optimal Transport

A class of comparative measures based on optimal transport for Morse complexes

- **Structural correspondences** between Morse graphs (1D-skeleton of Morse complexes)
- Supporting comparative analysis, including tasks such as classification or clustering

Mingzhe Li, Carson Storm, Austin Yang Li, Tom Needham, Bei Wang, "Tracking Low-Level Cloud Systems with Topology", In IEEE VIS Short Papers, 2023

MS-COOT: Comparing Morse-Smale Complexes via Co-Optimal Transport

Region-aware comparison of scalar-field topology

Guangyu Meng, Mingzhe Li, Erin W. Chambers

Under review

Morse-Smale (MS) Complex

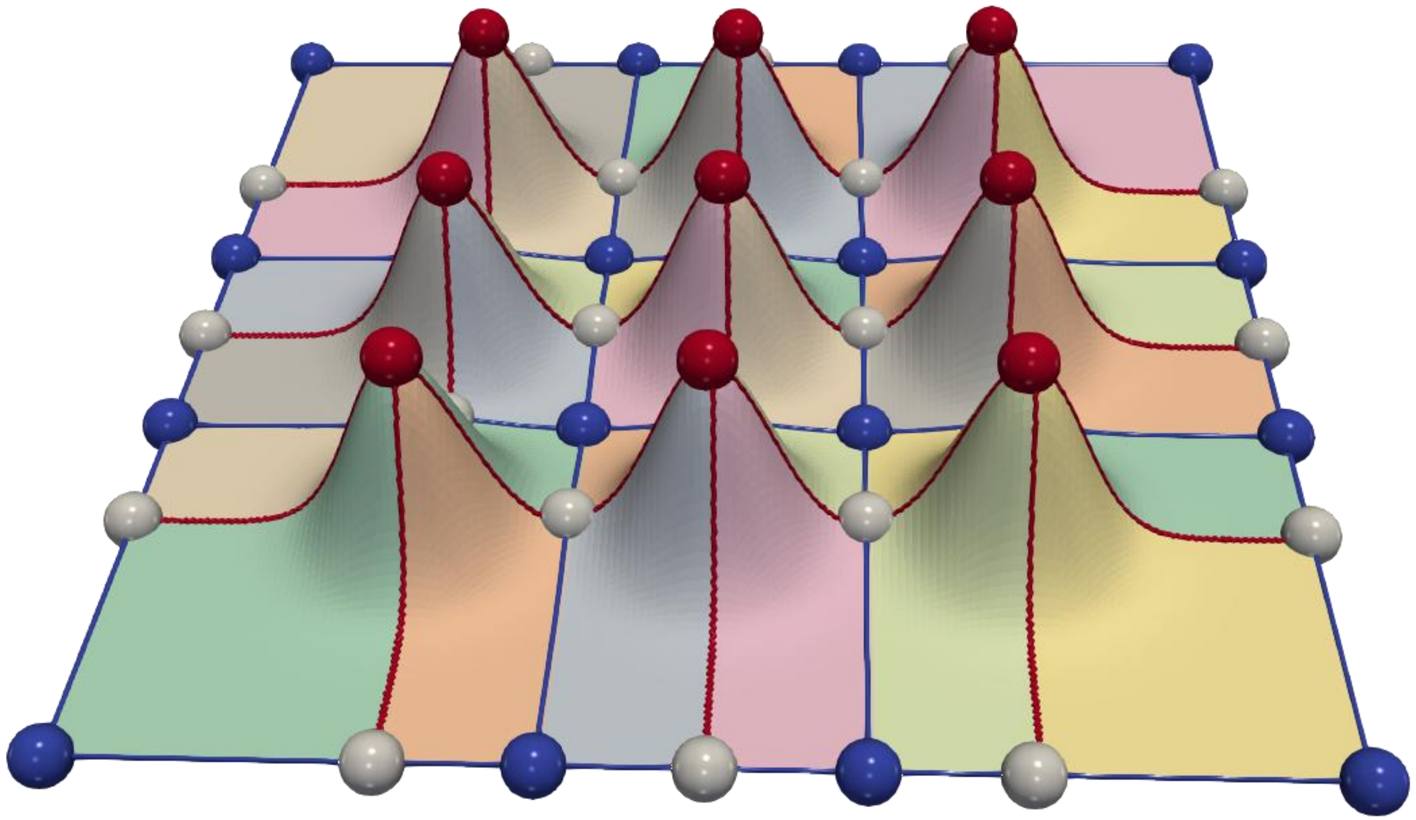

Terrain for 2D scalar field and its Morse-Smale complex

2D MS complexes

- Critical points: Minima, saddles, maxima
- 1D-skeletons: Morse graphs
- 2-cells: Segmentations/regions

Motivation: Gap in MS complex comparison

- Existing method: Optimal transport for Morse graph comparison
- Pros
  - Couplings support interpretation of correspondences
- Cons
  - Aware of graph structures only
  - No region-level correspondences

Our Approach: MS Complex as Hypergraph

Key insight: MS complex has natural hypergraph structure – critical points as nodes, regions as hyperedges.
Hyperedge: edge connecting two or more nodes.

Measure Hypernetwork H

V	Critical points (CPs)
E	Regions (hyperedges – incident to boundary CPs)
$\mu \in \Delta(V)$	Persistence-based probability measure on CPs
$\nu \in \Delta(E)$	Probability measure on regions (derived from μ)
$\omega: V \times E \to \mathbb{R}$	Hypernetwork function encoding CP-to-region proximity

Our Contributions

1. **MS-COOT distance**
   Co-optimal transport on hypergraph-structured MS complexes, jointly matching CPs and regions
2. **Region-level correspondence**
   Explicit region coupling ξ -- general distance-based formulation with region correspondences
3. **Domain-specific components**
   Hypernetwork function ω, persistence-based measures, and flexible sample cost term
4. **Comprehensive evaluation**
   Datasets across 2D, 3D surface, and 3D volumetric data

MS-COOT Distance

$$d_{MS-COOT}(H_f, H_g) = \min_{\pi,\xi} \sum_{i,j,k,l} |\omega_f[i,k] - \omega_g[j,l]|^2 \, \pi[i,j] \, \xi[k,l] + \alpha\langle C, \pi\rangle$$

$\pi \in \Pi(\mu_f, \mu_g)$ – CP coupling

$\xi \in \Pi(\nu_f, \nu_g)$ – Region coupling

Structure Term (quadratic)

- Compares relational values $\omega_f[i,k]$ and $\omega_g[j,l]$ under joint coupling (π, ξ)
- Encourages matched CPs and regions to have similar CP-to-region relationships
- Extends the standard COOT objective to MS hypernetworks

Sample Cost Term $\alpha\langle C, \pi\rangle$

- Default: type consistency
  $$C[i,j] = \mathbf{1}[type(CP_i) \neq type(CP_j)]$$
- For shape classification: scalar function value differences $C[i,j] = |f(CP_i) - f(CP_j)|$
- $\alpha \in [0,1]$ balances structural consistency vs. sample-level constraints

MS-COOT Pipeline Overview

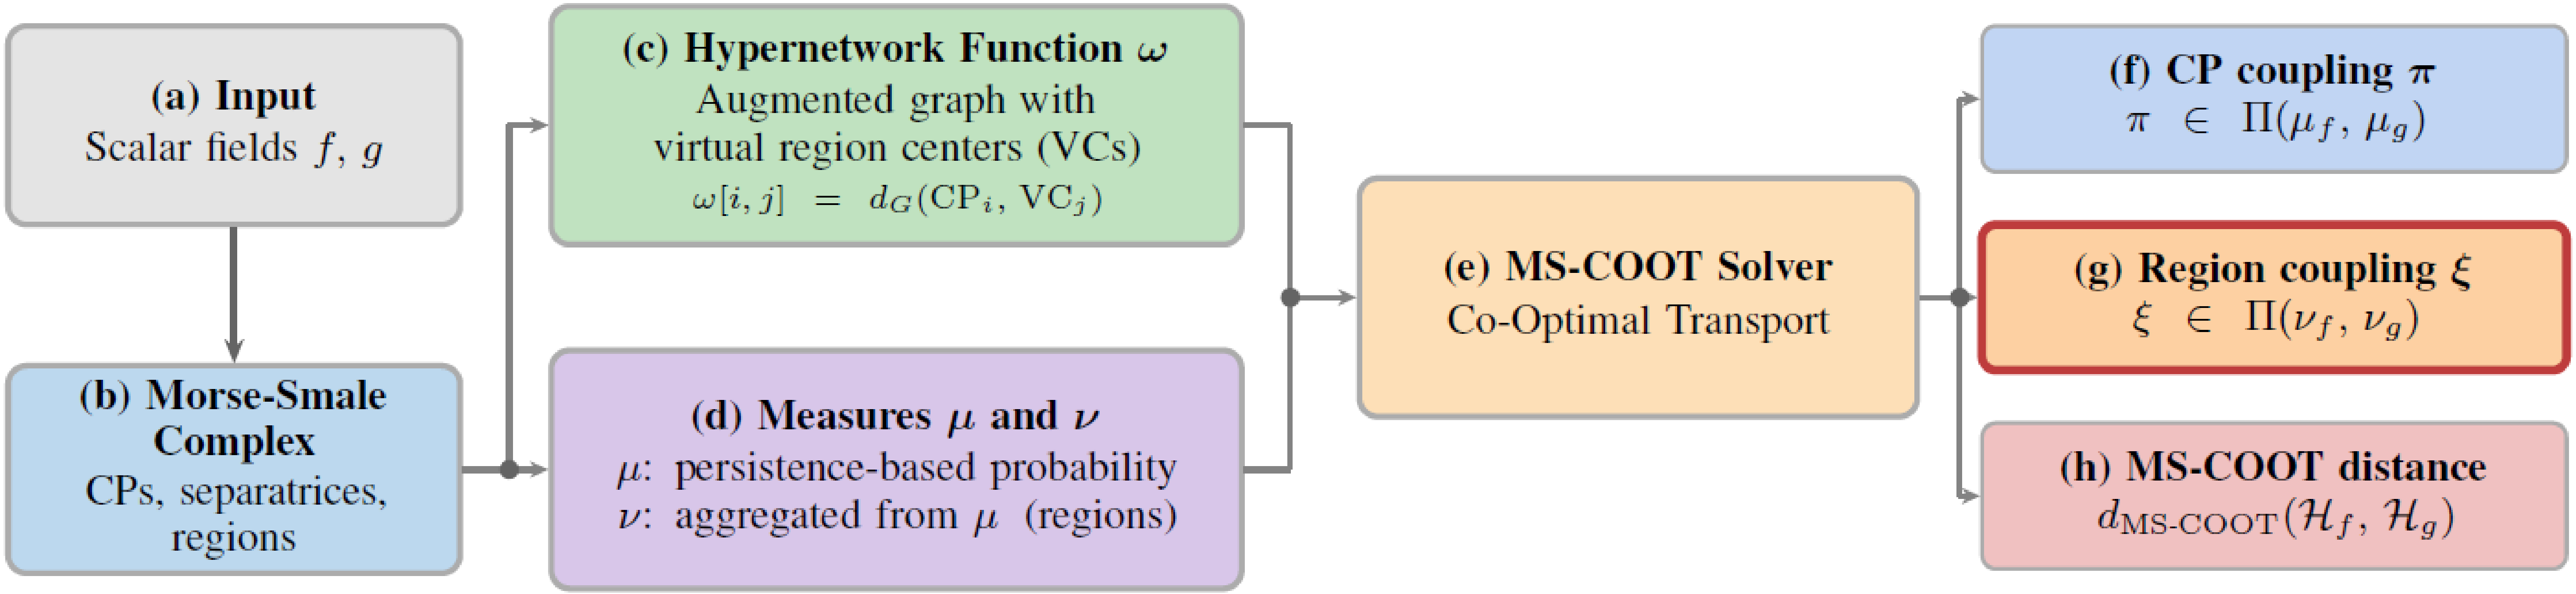

What MS-COOT Distance Provides

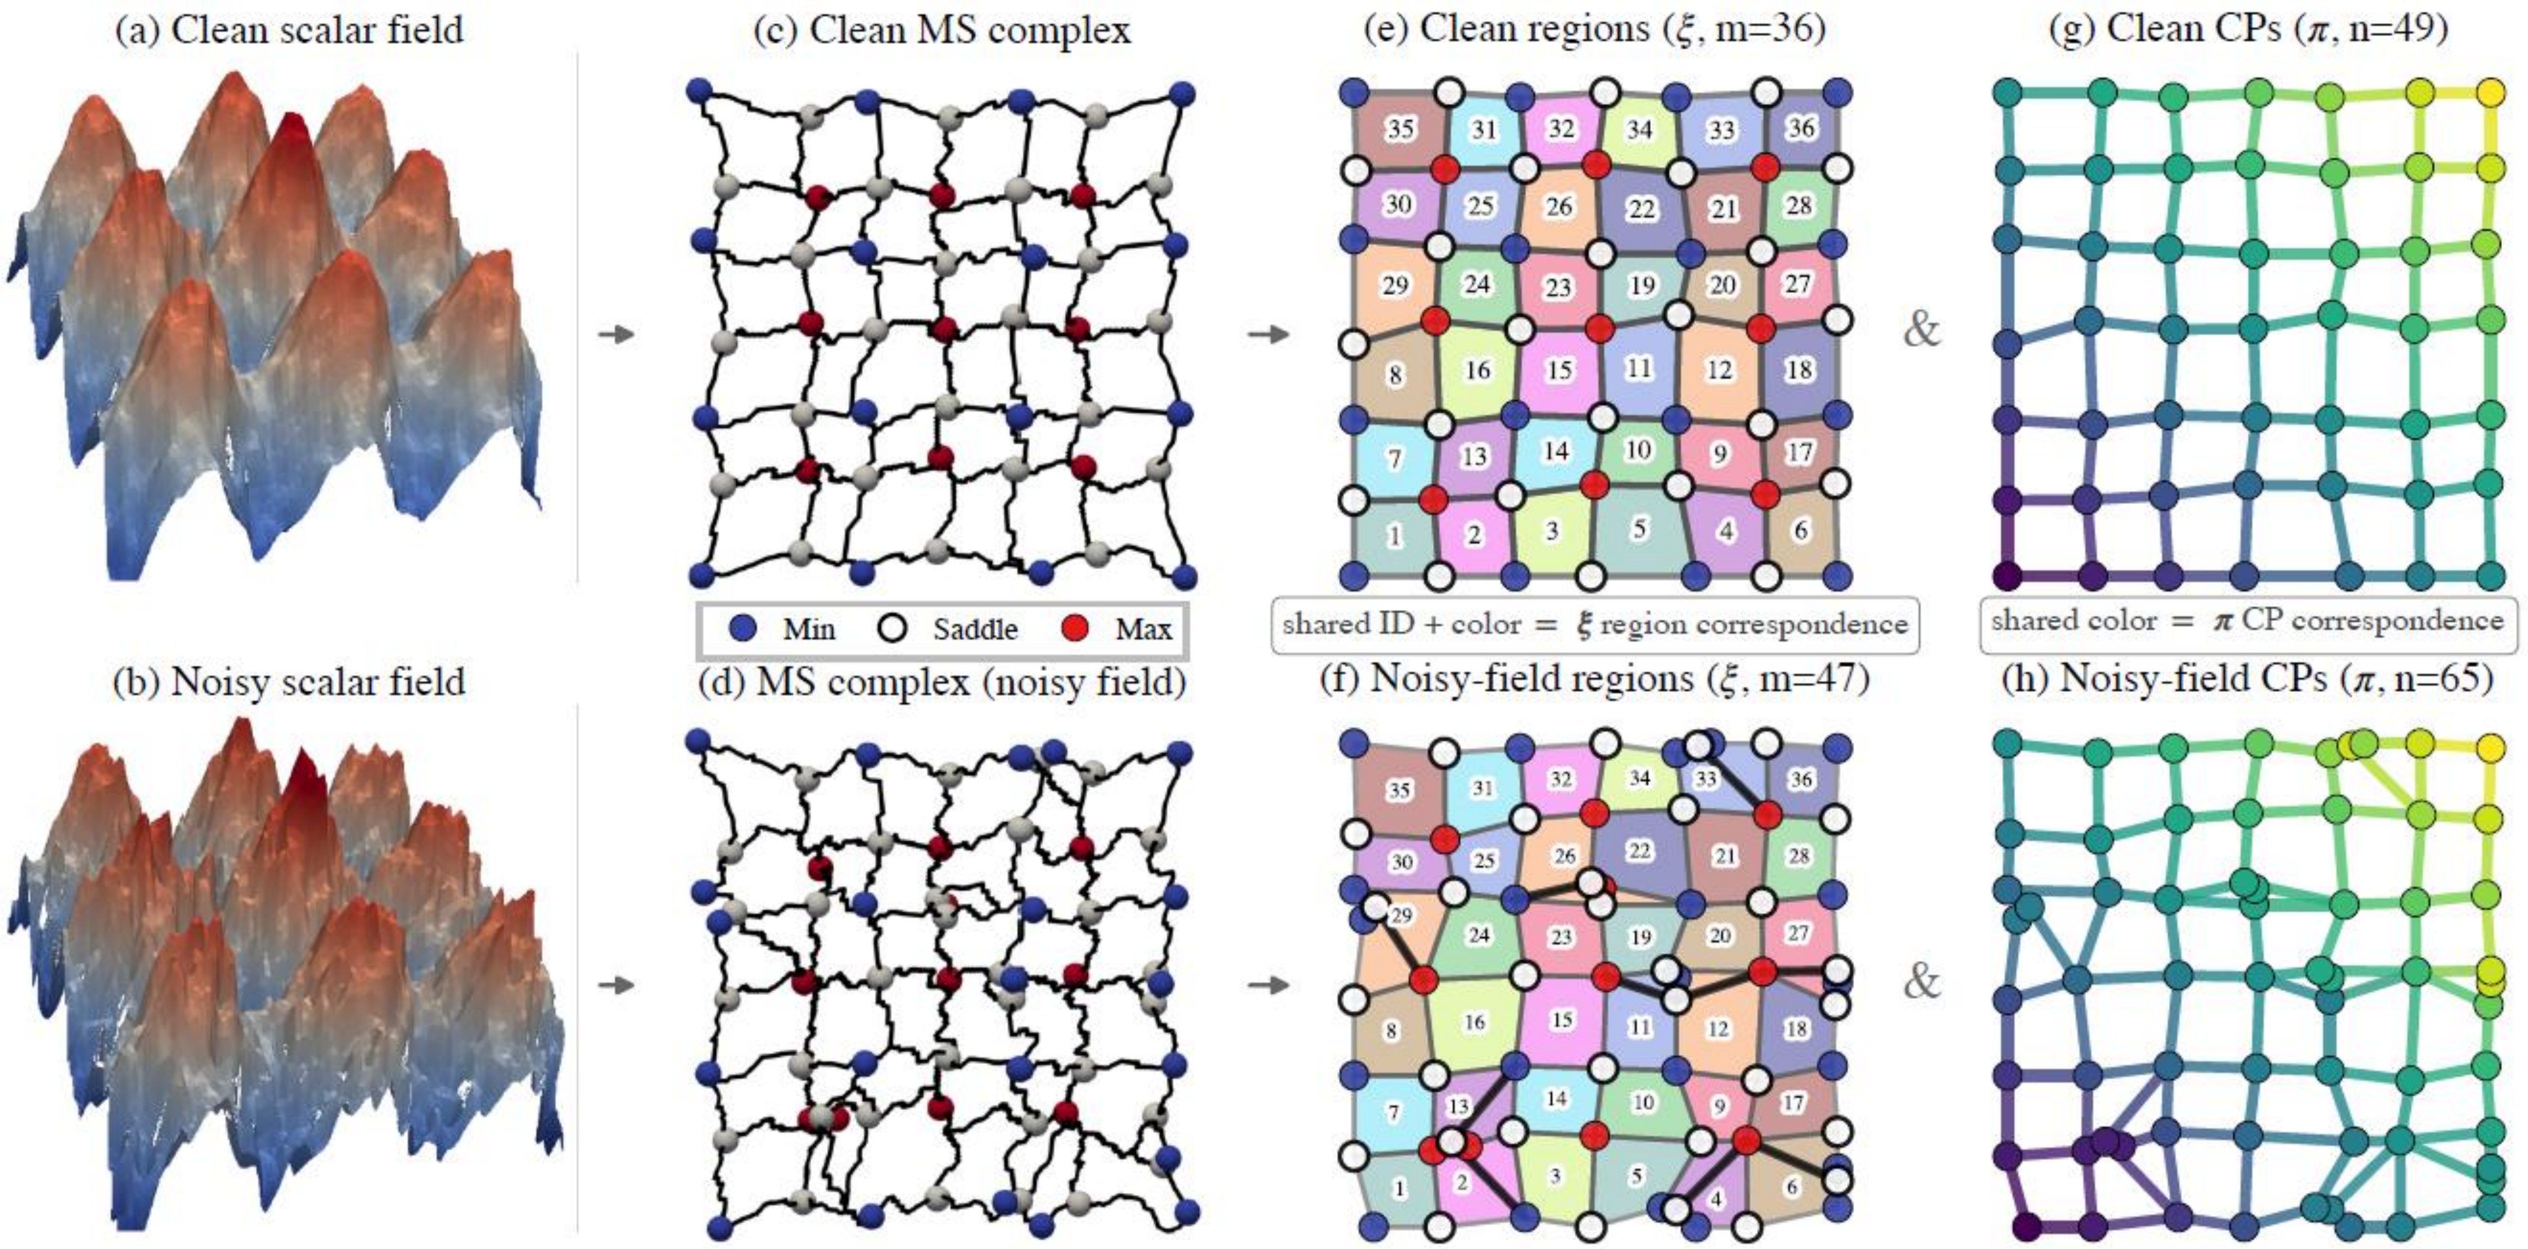

$\xi \in \Pi(\nu_f, \nu_g)$
Region coupling

$\pi \in \Pi(\mu_f, \mu_g)$
CP coupling

Hypernetwork Function ω — Shortest-Path Distances

$\omega[i,j] = d_G(CP_i, VC_j)$ — shortest-path distance on augmented graph G from each critical point to each virtual region center.

① **Virtual Region Centers (VCs)**

Each region R_j gets a representative point VC_j.
In 2D, four CPs on boundary: intersection of saddle-saddle and min-max connections (Quadrangle Lemma).
Otherwise, barycenter of boundary CPs.

② **Augmented Graph G**

Node set: $CPs \cup VCs$.
Edges: (a) CP-CP along separatrices, (b) CP-VC connecting region boundary CPs to VC, (c) VC-VC between adjacent regions.

③ **Shortest-Path Distance**

$$\omega[i,j] = d_G(CP_i, VC_j)$$

Captures global proximity: non-adjacent CPs and regions get meaningful graded distances, not just binary incidence.

④ **Joint Normalization**

Rescale both ω_f and ω_g by $Z = \max(\max \omega_f, \max \omega_g)$ so that both lie in [0, 1] before comparing.

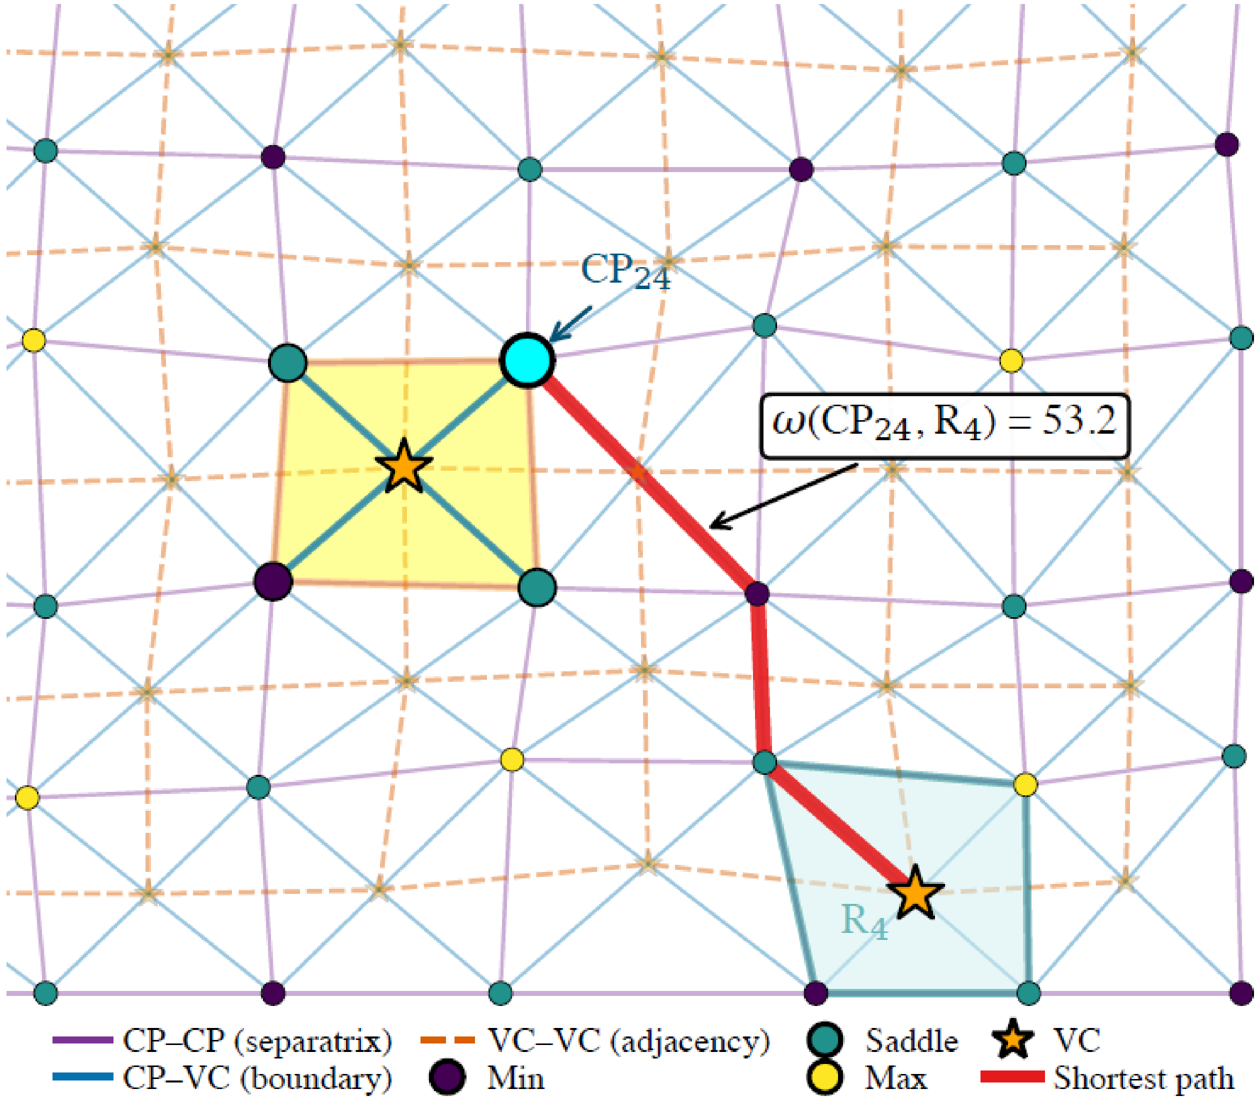

Persistence-Based Probability Measures

We weight CPs by topological persistence to emphasize structurally significant features.

Node Measure μ (on CPs)

Persistence Image (PI) representation of the persistence diagram.

- Place Gaussian kernel of bandwidth σ at each persistence pair (b_i, d_i), weighted by persistence $d_i - b_i$
- $\mu_i = PI_\sigma(b_i, d_i) \,/\, \sum_k PI_\sigma(b_k, d_k)$
- Smooth: similar-persistence CPs get similar weights; σ controls concentration
- Essential CPs without pairs with finite persistence: $0.1/n$

Hyperedge Measure ν (on Regions)

Regions' probability measures derive from boundary CPs:

$$\boldsymbol{\nu_j} = \frac{\sum_{\boldsymbol{CP_i} \in \boldsymbol{\partial R_j}} \boldsymbol{\mu_i}}{\sum_{\boldsymbol{l}} \sum_{\boldsymbol{CP_i} \in \boldsymbol{\partial R_l}} \boldsymbol{\mu_i}}$$

- Regions bounded by high-persistence CPs inherit larger weights
- Denominator accumulates per incidence (a CP on multiple regions contributes multiple times)
- Ensures structurally significant regions are well-represented in the region coupling ξ

Baselines – Graph-based OT-type Distances

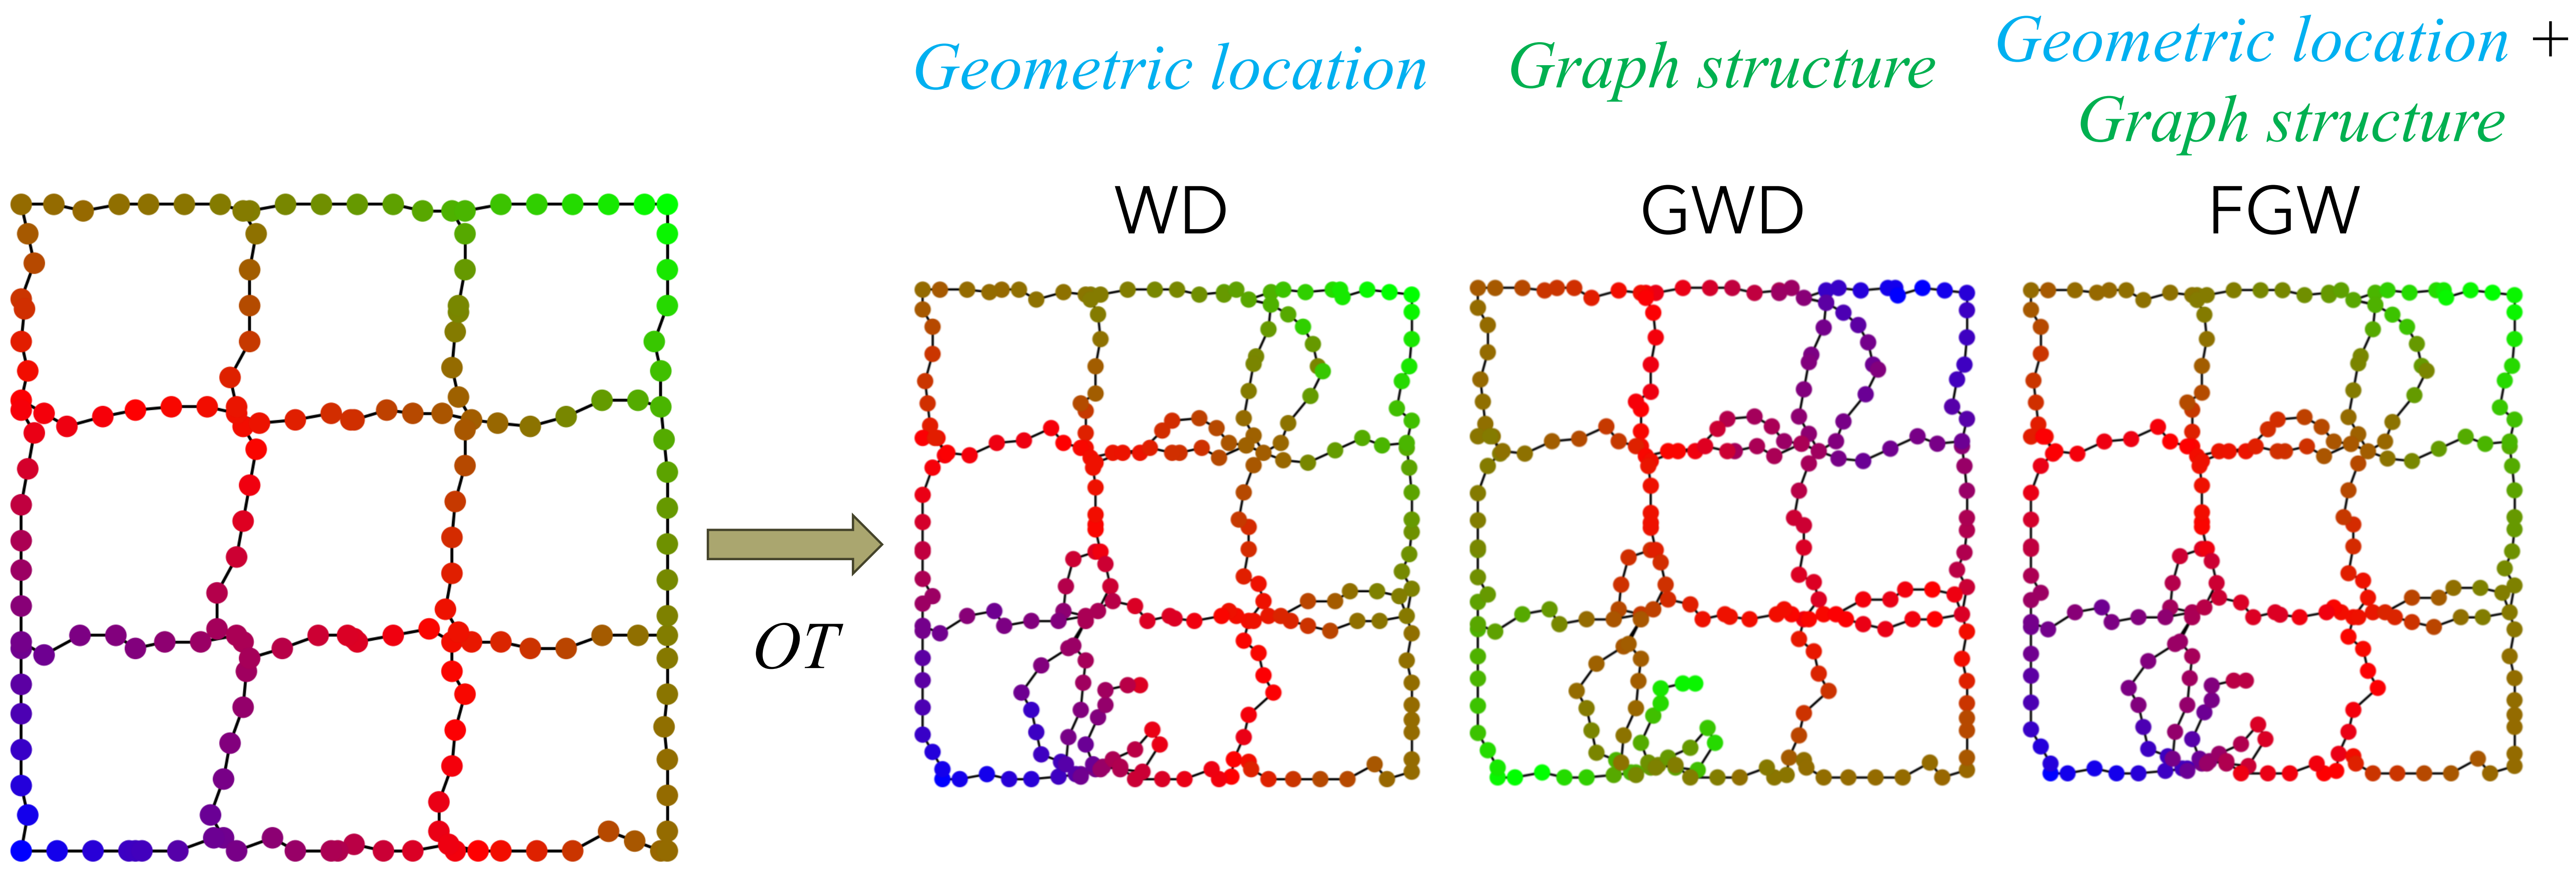

Results: Vortex Street — Periodic Dynamics

157 timesteps of 2D von Kármán vortex shedding.
Goal: capture periodicity and robustness to local perturbations.

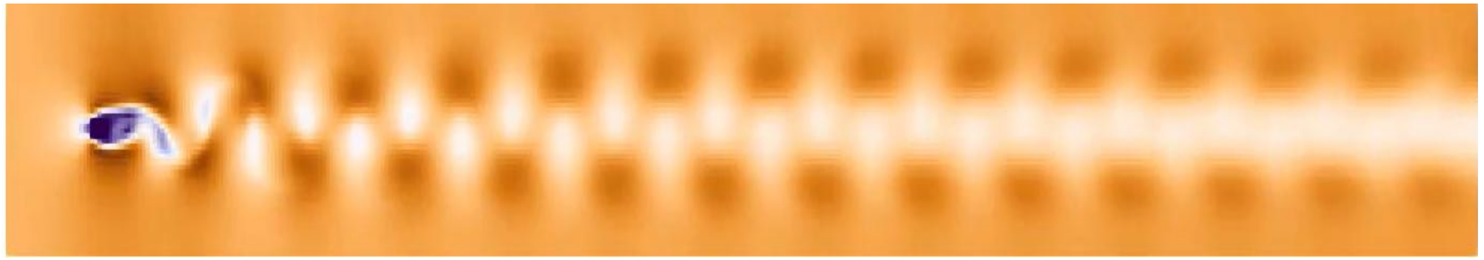

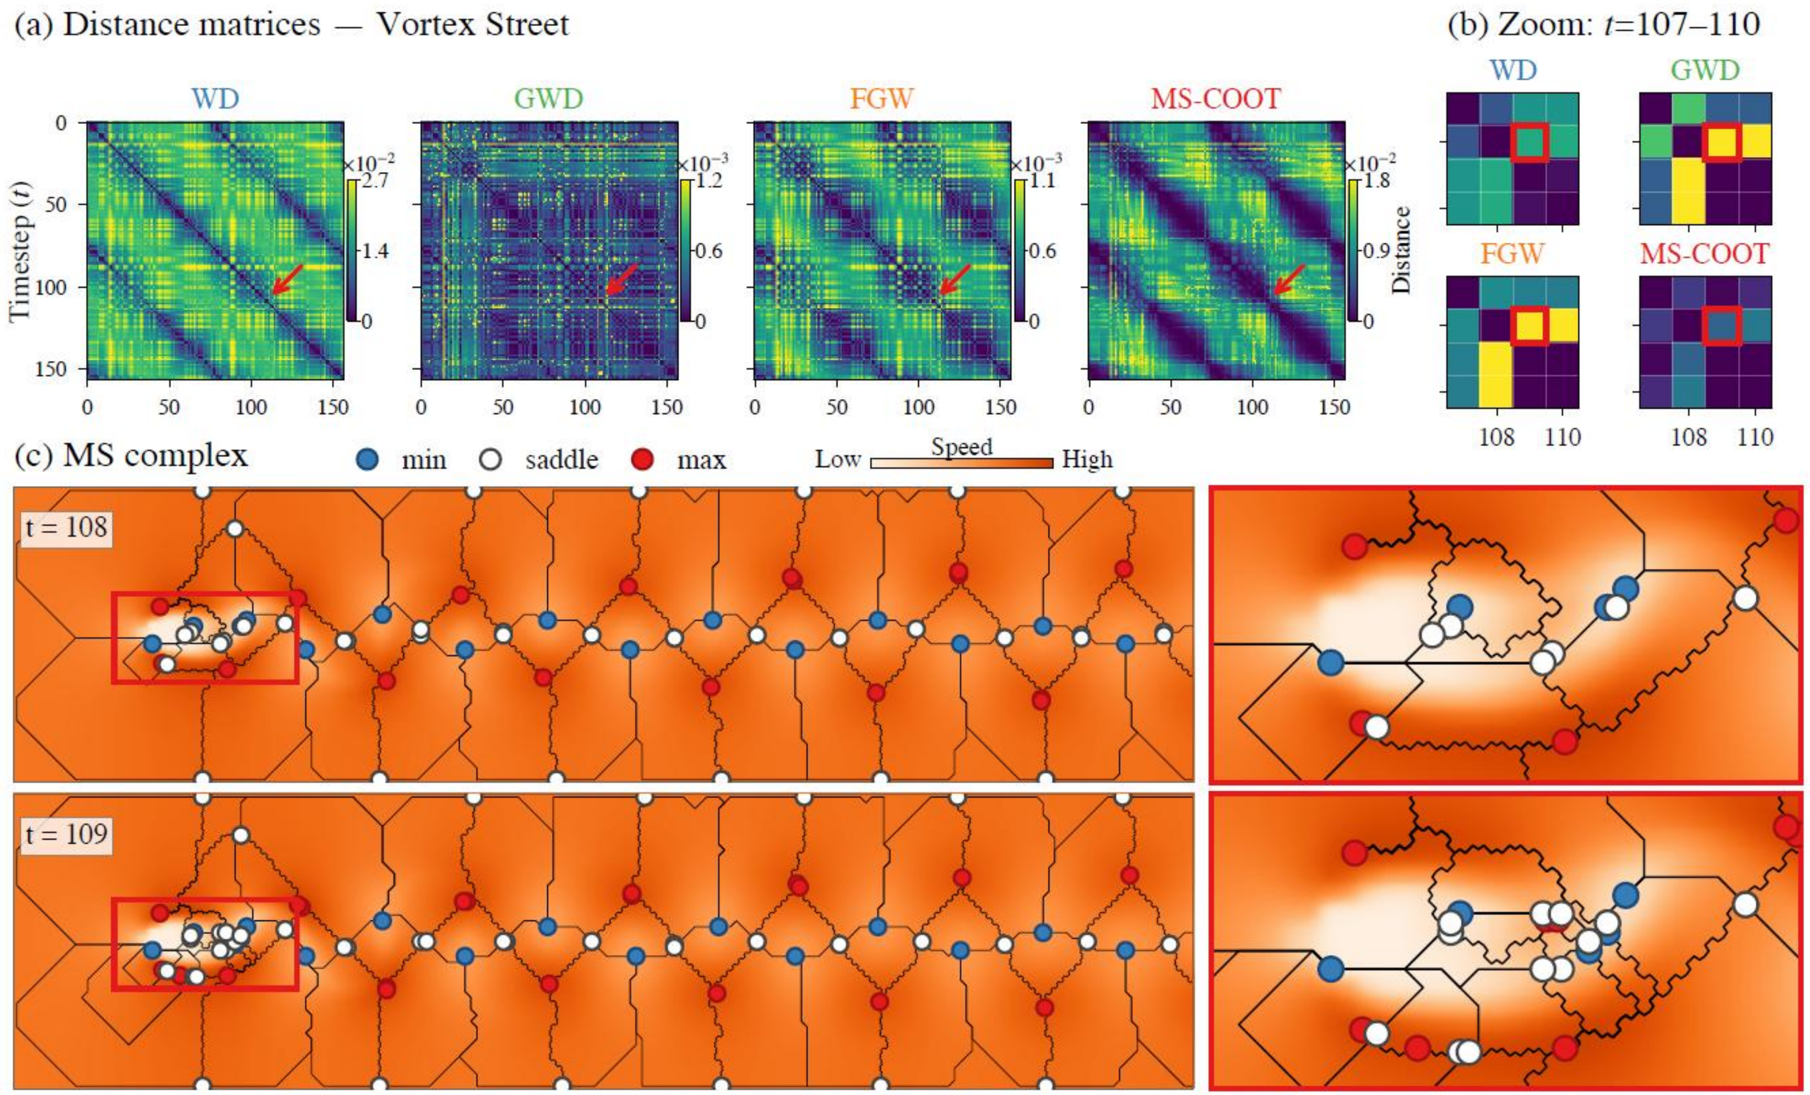

Off-diagonal banding

All methods capture shedding period a ≈ 74 timesteps; MS-COOT shows smooth bands

Smooth at t=108→109

Local CP perturbations → MS-COOT reports low distance; graph-based methods spike

Why MS-COOT is stable

Region adjacency in ω buffers minor separatrix changes; persistence measures suppress weak CPs

Results: TOSCA — Shape Classification

80 non-rigid 3D meshes in 9 categories; sample cost term C: average geodesic distance (AGD). k-NN leave-one-out classification.

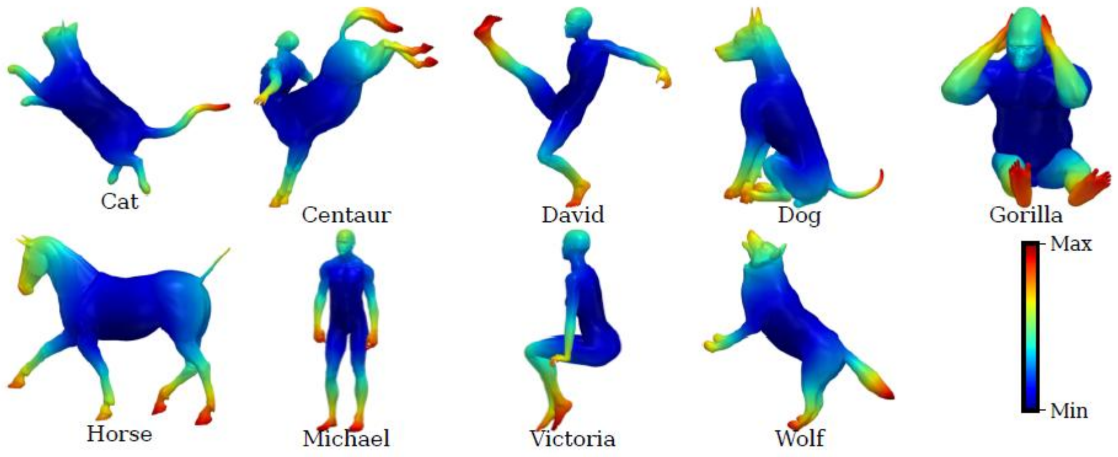

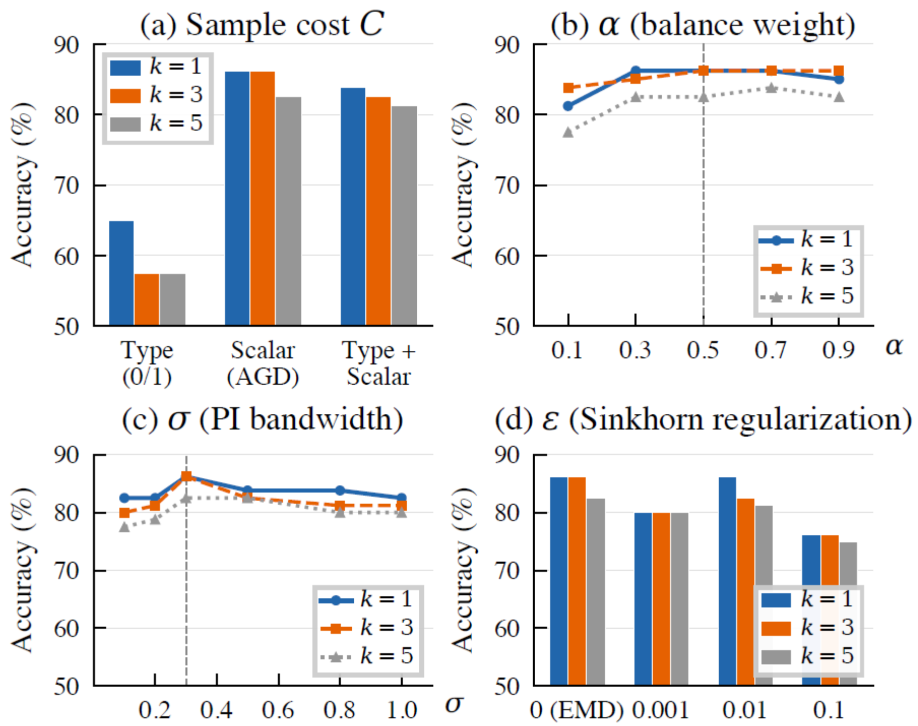

Table 1: **TOSCA shape classification.** Accuracy (%) by k-NN with $k=1$. N: number of meshes per class. Best per row in bold. The main confusion is David → Michael (see the Supplementary for details).

Class (N)	WD	GWD	FGW	MS-COOT
Cat (11)	90.91	63.64	90.91	**100.00**
Centaur (6)	83.33	**100.00**	83.33	**100.00**
David (7)	28.57	**57.14**	28.57	28.57
Dog (9)	**77.78**	66.67	**77.78**	**77.78**
Gorilla (4)	25.00	0.00	50.00	**75.00**
Horse (8)	75.00	**100.00**	**100.00**	**100.00**
Michael (20)	45.00	50.00	65.00	**85.00**
Victoria (12)	58.33	58.33	58.33	**100.00**
Wolf (3)	33.33	66.67	33.33	**100.00**
Overall (80)	60.00	62.50	68.75	**86.25**

- Higher detection accuracy than graph-based baselines
- Stable performance across parameter choices

Results: Viscous Finger — 3D Volumetric Data

IEEE SciVis Contest 2016 ensemble: salt dissolving in water (FPM simulation).
Three resolution levels (h=0.20/0.30/0.44), 5 runs each.

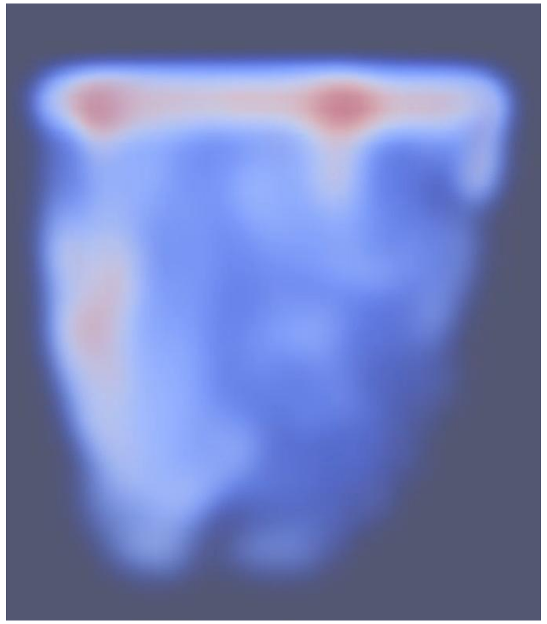

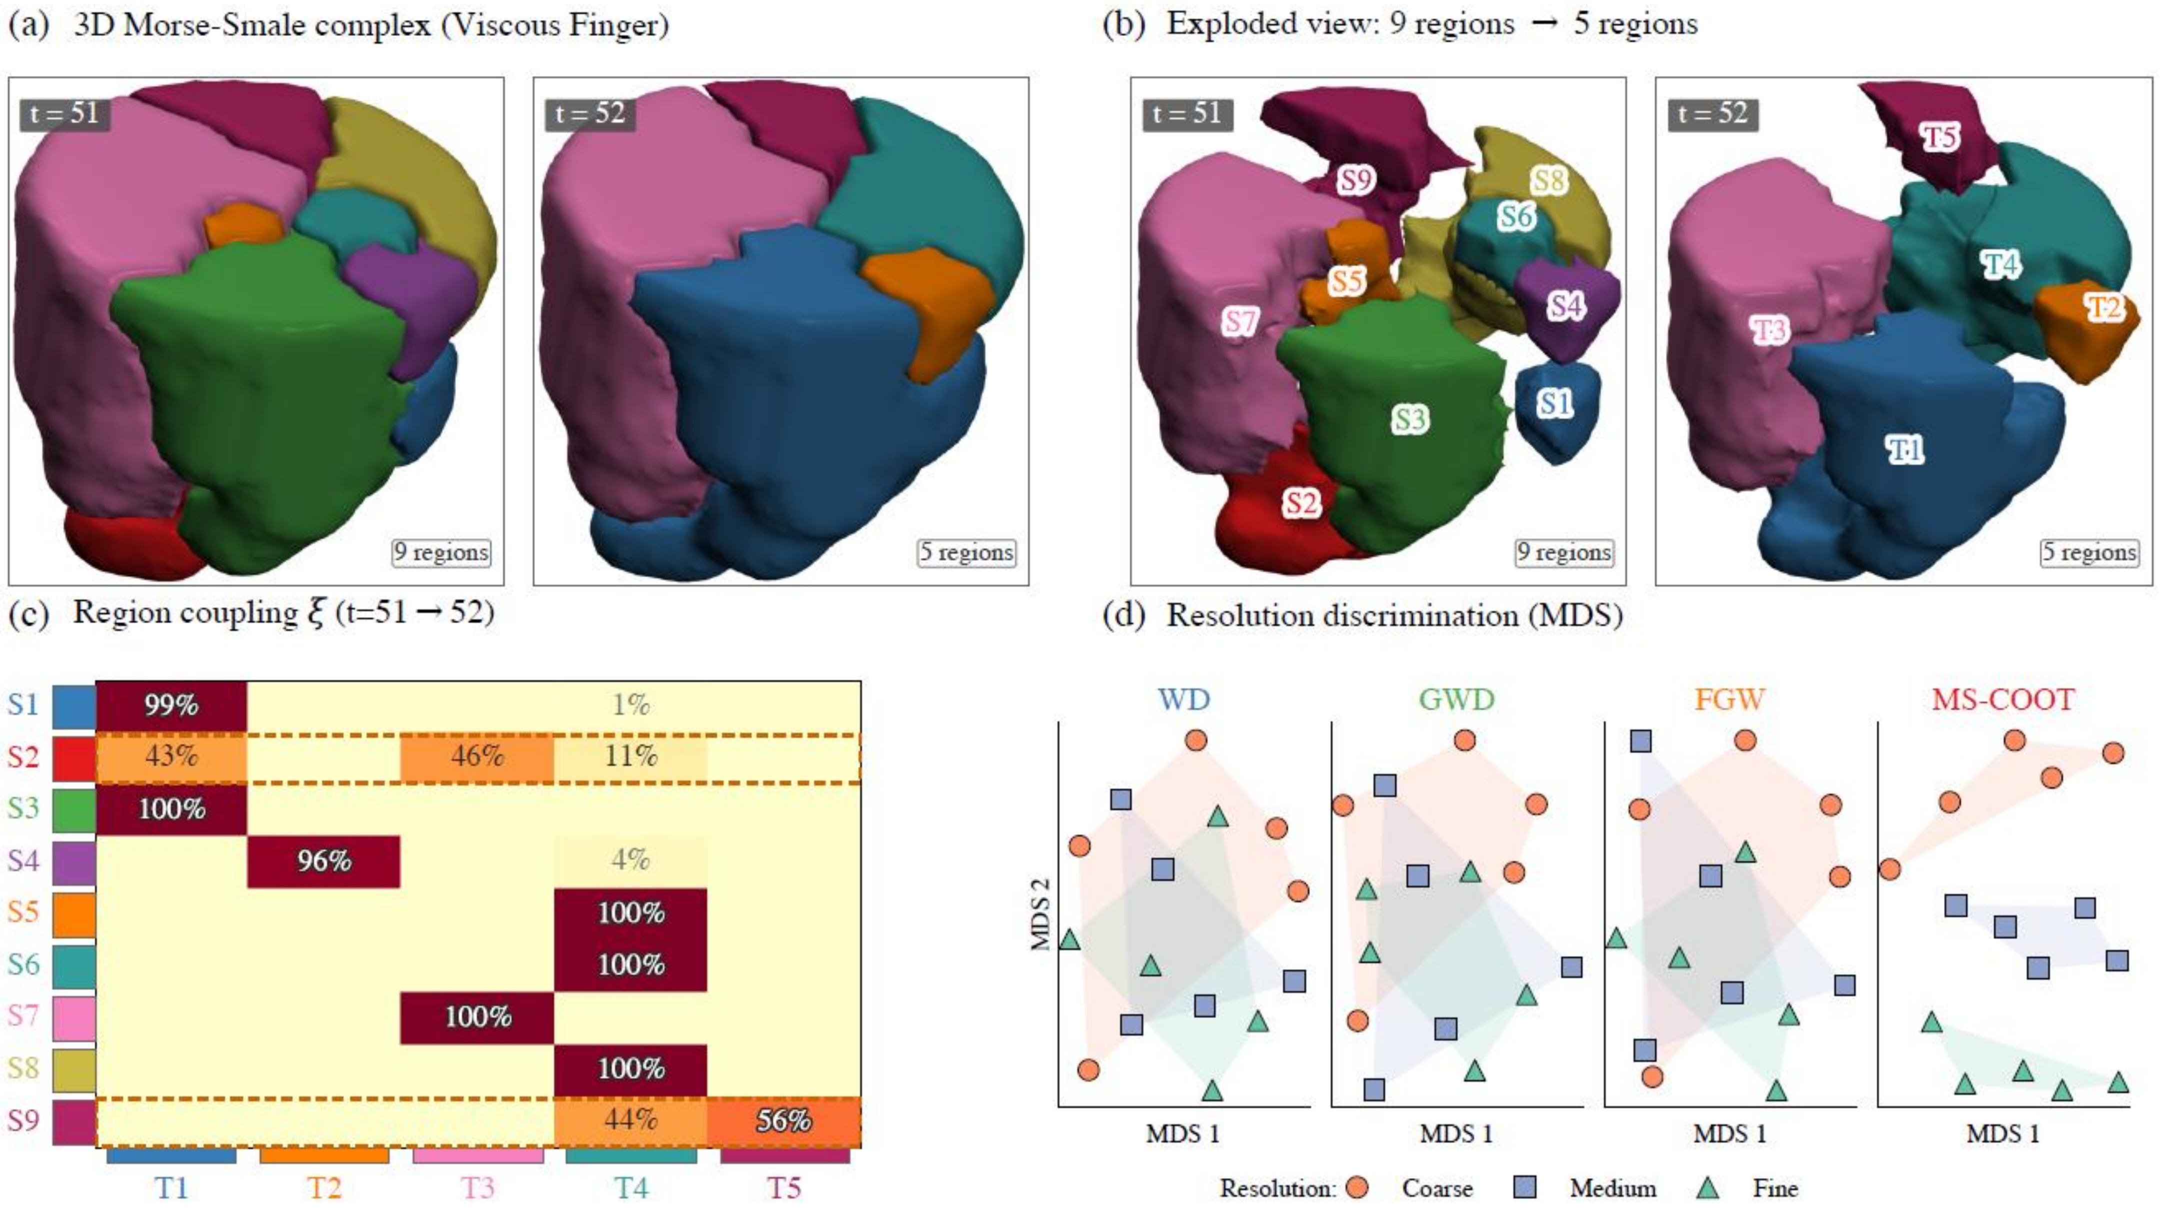

Resolution Classification (k-NN, N=15, 3 classes)		
Method	Accuracy	p-value
WD	53.3%	0.028*
GWD	46.7%	0.160
FGW	33.3%	0.427
MS-COOT	**80.0%**	**0.002****

Region coupling ξ reveals meaningful merge/split events:

- S1, S2, S3→T1
- S2, S7 → T3
- S5, S6, S8, S9 → T4
- S2→T1, T3, T4
- S9→T4, T5

Conclusion & Future Works

Contact:
- https://mingzhefluorite.github.io
- mli33@nd.edu

MS-COOT

✓ Co-optimal transport on hypergraph-structured MS complexes

✓ Jointly matches critical points (π) AND regions (ξ) simultaneously

✓ Explicit region-to-region correspondence

✓ Detects splits, merges, and large-scale topological transitions

✓ Strong classification and discrimination

Open Questions

- Non-convex objective → local optima; initialization sensitivity
- Scalability to large unsimplified complexes
- Geometric constraints for temporal tracking with ξ
- Metric properties when $\alpha > 0$
- Multi-parameter extensions

Thank you!

Optimization: Block Coordinate Descent + Sinkhorn

Initialize	Fix ξ → update π	Fix π → update ξ	Converged?	Output π, ξ, d
π from marginal OT problem on C ξ = uniform distribution	Entropic OT Solved via Sinkhorn	Entropic OT Solved via Sinkhorn	Check convergence or max BCD iterations reached	CP coupling, region coupling, and MS-COOT distance

Complexity:

$O(K(n_f n_g + m_f m_g))$ per BCD iter
n: # of CPs, m: # of regions
K Sinkhorn iters; one solve for π over CPs, one for ξ over regions

Initialization:

π from marginal OT on C; ξ uniform. Fallback to exact EMD ($\varepsilon = 0$) if Sinkhorn diverges.

Convergence:

BCD converges to a stationary point (Redko et al.); linear term $\alpha\langle C, \pi\rangle$ preserves this — convex in π-subproblem

Parameters:

ε = 0.001 (Sinkhorn reg.) α = 0.5 (balance) σ = 0.3 (PI bandwidth) — robust across tested datasets