Corinna Coupette
Max Planck Institute for Informatics

Sebastian Dalleiger
CISPA

Bastian Rieck
Helmholtz Munich

Ollivier-Ricci Curvature for Hypergraphs

A Unified Framework

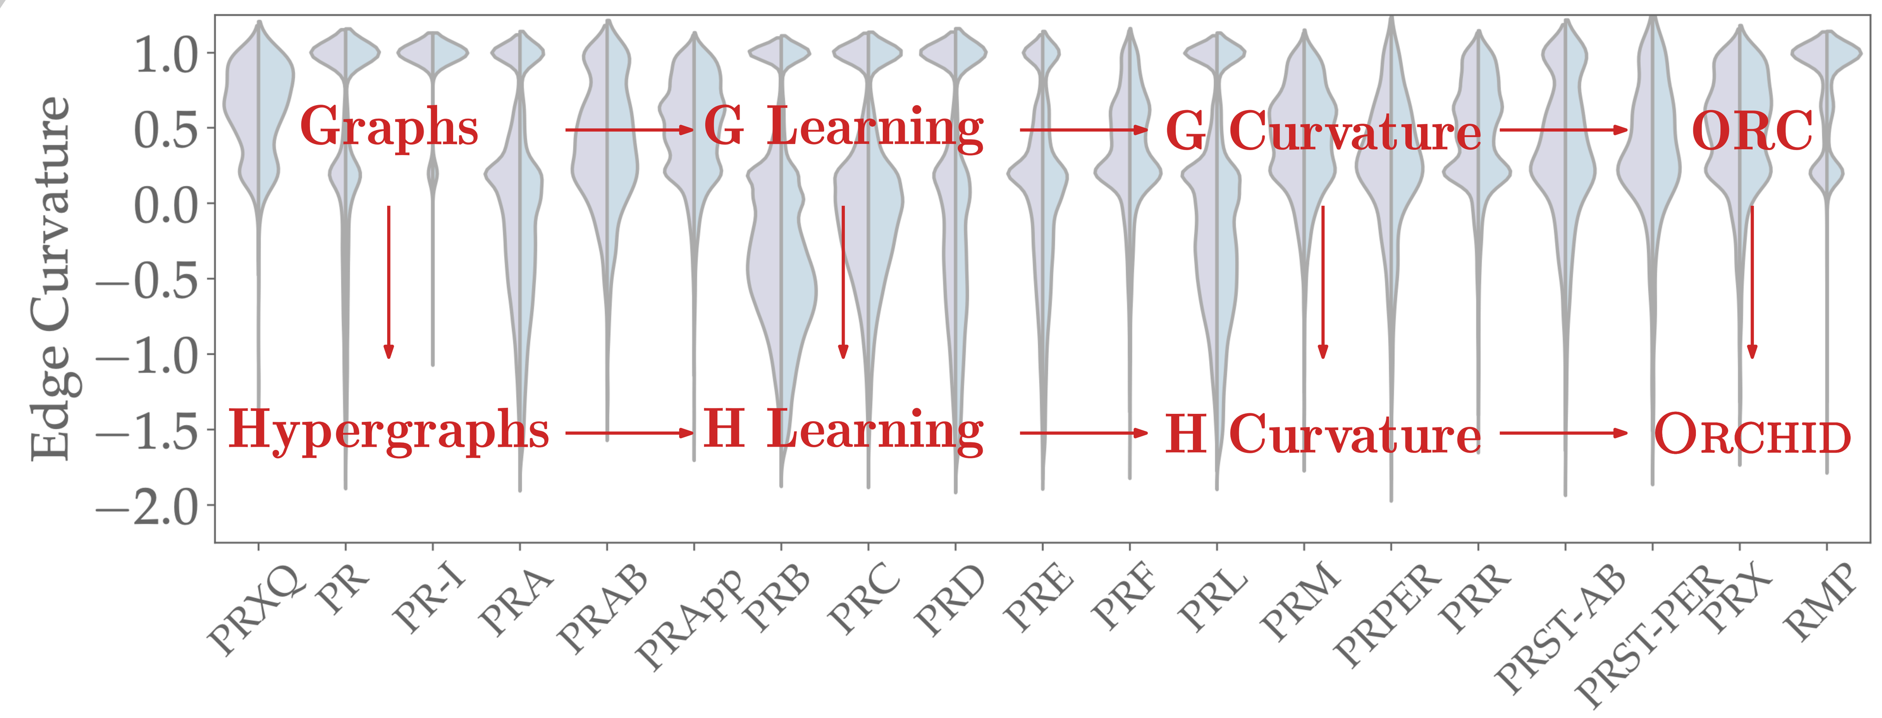

Ollivier-Ricci Curvature

- Discretization of Ricci curvature
- Value in $[-2, 1]$ for *local* geometry around nodes and edges
  $\leftrightarrow$ diffusion behavior of random walks
- Ingredients:
  - $d(i, j)$: Shortest-path distance between nodes i and j
  - μ_i: Probability measure associated with node i
    $\leftrightarrow$ α-lazy random walk (α: smoothing parameter)
  - W_1: Wasserstein distance between two probability measures
- *Edge* curvature: $\kappa(i, j) = 1 - \frac{W_1(\mu_i, \mu_j)}{d(i,j)}$
- *Node* curvature: Mean of incident edge curvatures

Curvature in Graph Learning

- Comparing real-world networks
- Identifying bottlenecks in real-world networks
- Alleviating oversquashing in graph neural networks

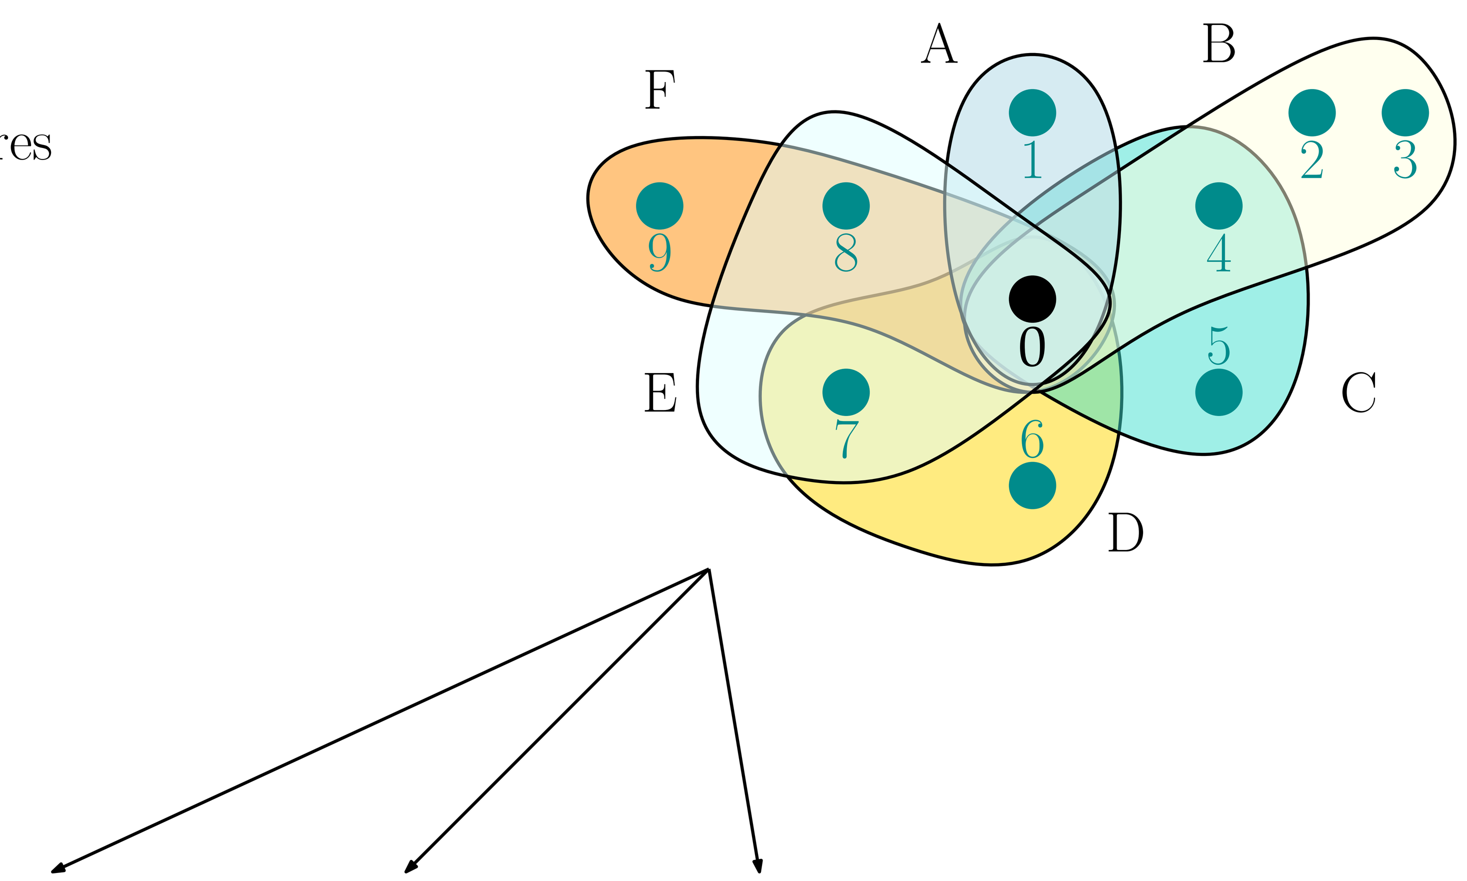

Random Walks on Hypergraphs

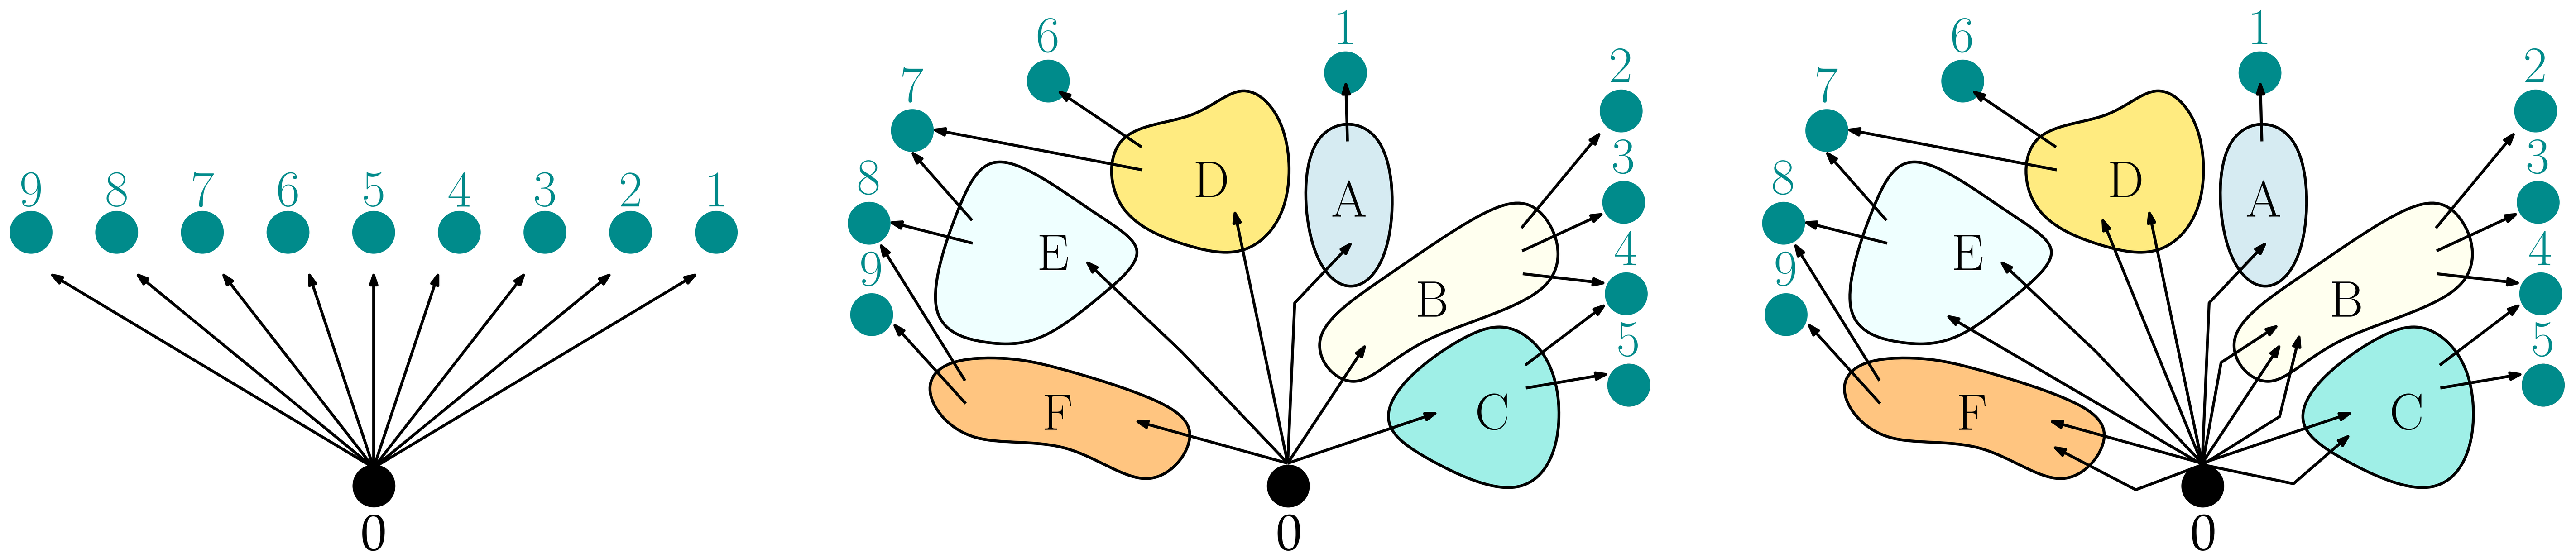

EN: Equal-Nodes Random Walk
$\updownarrow$
Unweighted Clique Expansion

EE: Equal-Edges Random Walk
$\updownarrow$
Unweighted Star Expansion

WE: Weighted-Edges Random Walk
$\updownarrow$
Weighted Clique Expansion

The ORCHID Framework for Hypergraph Curvature

- Probability measure μ_i for node i: α-lazy { EN, EE, WE } random walk
- *Directional* curvature along direction $i - j$: $\kappa(i, j) = 1 - \frac{W_1(\mu_i, \mu_j)}{d(i,j)}$
- *Edge* curvature for edge e: $1 - \text{AGG}(e)$, for permutation-invariant AGG function:
  - $\text{AGG}_A(e) = \frac{2}{|e|(|e|-1)} \sum_{\{i,j\} \subseteq e} W_1(\mu_i, \mu_j)$
  - $\text{AGG}_M(e) = \max \{W_1(\mu_i, \mu_j) \mid \{i, j\} \subseteq e\}$
  - $\text{AGG}_B(e) = \frac{1}{|e|-1} \sum_{i \in e} W_1(\mu_i, \bar{\mu})$, where $\bar{\mu}$ is the barycenter of $\{\mu_i \mid i \in e\}$
- *Node* curvature for node i:
  - Mean of *directional* curvatures $\{\kappa(i, j) \mid i \sim j\}$
  - Mean of *edge* curvatures $\{\kappa(e) \mid i \in e\}$
- Provably generalizes Ollivier-Ricci curvature to hypergraphs

ORCHID Curvatures in the Wild

Setup:
- 4 individual hypergraphs
  + 9 hypergraph collections
- Up to >3M nodes and >6M edges
- ~20–2 000 hypergraphs per collection
- Varying densities + distributions

Results:
- Distinct curvature profiles
- Discriminative test statistics
- Interpretable embeddings
- Cohesive clusterings