KTH ROYAL INSTITUTE
OF TECHNOLOGY

Learning and Evaluating the Geometric Structure of Representation Spaces

PhD Thesis Defense, June 13th 2022

Petra Poklukar

Outline

- **Data Representations: Learning and Evaluation**
  - two main challenges

- **Geometric Aspects of Evaluation and Representation Learning**
  - Paper A - ICML: Geometric Component Analysis
  - Paper B - ICLR: Delaunay Component Analysis
  - Paper D - ICML: Geometric Multimodal Contrastive Learning

Theoretical contributions

- **Applications**
  - Paper E - T-RO: Latent Space Roadmap
  - Paper C - preprint: GraphDCA

- **Conclusion, takeaways, future directions**

Data Representations: Learning and Evaluation

Representation Learning: what and why

What:

Automatic discovery of useful feature patterns in the observed data.

Why:

An AI system should be able to *identify* and *structure* the underlying semantic information hidden in the observed data, and *leverage* that information for subsequent reasoning.

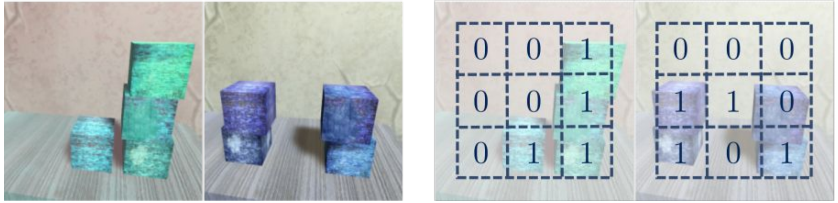

Learning and Evaluating Representations

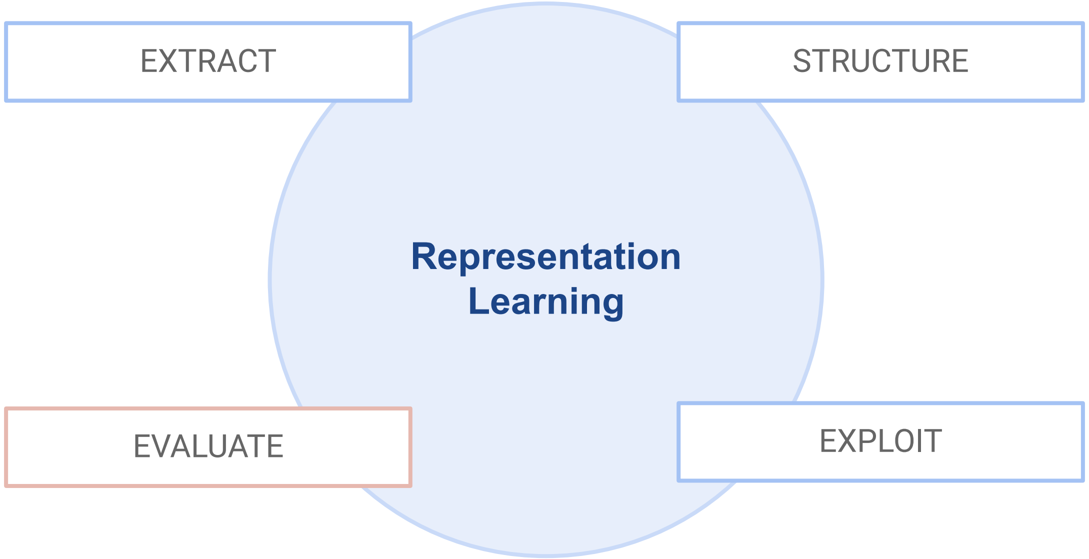

Learning and Evaluating Representations

Locatello et al., *A Sober Look at the Unsupervised Learning of Disentangled Representations and their Evaluation*, JMLR 2020

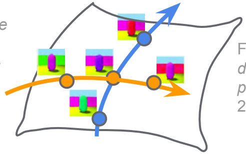

Fumero et al., *Learning disentangled representations via product manifold projection*, ICML 2021

Law et al., *Ultrahyperbolic representa learning*, NeurIPS 2020

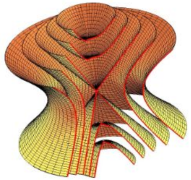

Chen et al., *A Simple Framework for Contrastive Learning of Visual Representations*, ICML 2020

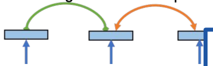

EXTRACT

STRUCTURE

Korman, Self-supervised representation learning on manifolds, ICLR 2021 Workshop on Geometrical and Topological Representation Learning

Gidaris et. al. *Unsupervised Representation Learning by Predicting Image Rotations*, ICLR 2018.

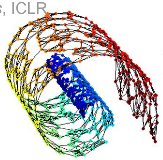

McInnes et al., (2018). *UMAP: Uniform Manifold Approximation and Projection.* Journal of Open Source Software, 3(29), 861

Moor et al., *Topological Autoencoders*, ICML 2020

Laturnus et al, *MorphVAE: Generating Neural Morphologies from 3D-Walks using a Variational Autoencoder with Spherical Latent Space*, ICML 2021

Representation Learning

Sodhani et al, *Multi-Task Reinforcement Learning with Context-based Representations*, ICML 2021

Whitney et al., *Evaluating representations by the complexity of learning low-loss predictors, arXiv:2009.07368.*

Barannikov et al., *Manifold Topology Divergence: a Framework for Comparing Data Manifolds*, NeurIPS 2021

EVALUATE

EXPLOIT

Collins et al., *Exploiting Shared Representations for Personalized Federated Learning*, ICML 2021

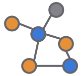

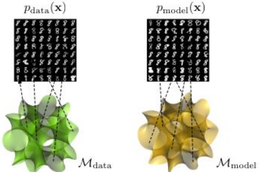

Kynkäänniemi et al., *Improved precision and recall metric for assessing generative models*, NeurIPS 2019

Kirichenko et. al., *Why Normalizing Flows Fail to Detect Out-of-Distribution Data*, NeurIPS 2020

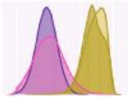

Bevilacqua et al., *Size-Invariant Graph Representations for Graph Classification Extrapolations*, ICML 2021

Khrulkov et al., *Geometry Score: A Method For Comparing Generative Adversarial Networks*, ICML 2018

Learning and Evaluating Representations

Locatello et al., *A Sober Look at the Unsupervised Learning of Disentangled Representations and their Evaluation*, JMLR 2020

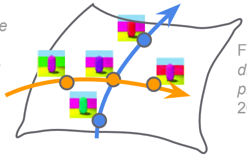

Fumero et al., *Learning disentangled representations via product manifold projection*, ICML 2021

Law et al., *Ultrahyperbolic representa* *learning*, NeurIPS 2020

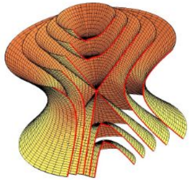

Chen et al., *A Simple Framework for Contrastive Learning of Visual Representations*, ICML 2020

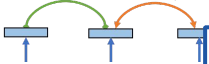

EXTRACT

STRUCTURE

Korman, Self-supervised representation learning on manifolds, ICLR 2021 Workshop on Geometrical and Topological Representation Learning

Gidaris et. al. *Unsupervised Representation Learning by Predicting Image Rotations*, ICLR 2018.

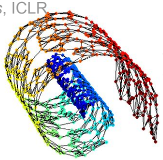

McInnes et al., (2018). *UMAP: Uniform Manifold Approximation and Projection.* Journal of Open Source Software, 3(29), 861

Moor et al., *Topological Autoencoders*, ICML 2020

Laturnus et al, *MorphVAE: Generating Neural Morphologies from 3D-Walks using a Variational Autoencoder with Spherical Latent Space*, ICML 2021

Representation Learning

Sodhani et al, *Multi-Task Reinforcement Learning with Context-based Representations*, ICML 2021

Whitney et al., *Evaluating representations by the complexity of learning low-loss predictors, arXiv:2009.07368.*

Barannikov et al., *Manifold Topology Divergence: a Framework for Comparing Data Manifolds*, NeurIPS 2021

EVALUATE

EXPLOIT

Collins et al., *Exploiting Shared Representations for Personalized Federated Learning*, ICML 2021

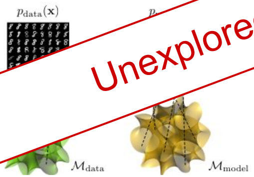

Unexplored

anniemi et al., *Improved precision and recall metric for assessing generative models*, NeurIPS 2019

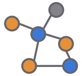

Kirichenko et. al., *Why Normalizing Flows Fail to Detect Out-of-Distribution Data*, NeurIPS 2020

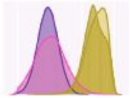

Bevilacqua et al., *Size-Invariant Graph Representations for Graph Classification Extrapolations*, ICML 2021

Khrulkov et al., *Geometry Score: A Method For Comparing Generative Adversarial Networks*, ICML 2018

Representation Learning: main challenges

To learn useful data representations, we need to consider the following two problems:

1. How to **design** representation **learning models** that identify semantically useful information and encode it into structured low dimensional representations?

- Paper D - ICML: Geometric Multimodal Contrastive Learning
- Paper E - T-RO: Latent Space Roadmap

2. How to **design** reliable **evaluation frameworks** for assessing the quality of the resulting representations?

- Paper A - ICML: Geometric Component Analysis
- Paper B - ICLR: Delaunay Component Analysis
- Paper C - preprint: GraphDCA

Learning and Evaluating Representations: Geometric and Topological Aspects

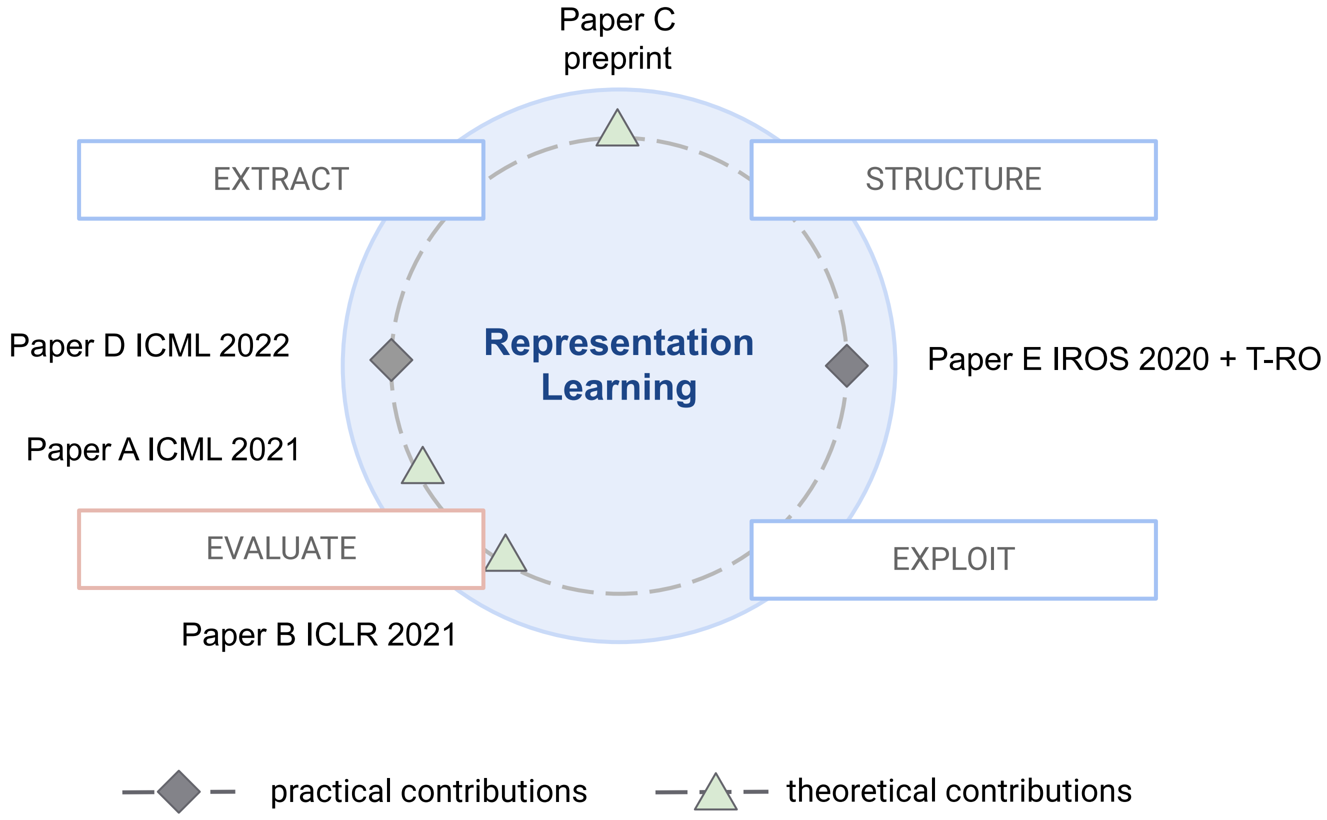

Evaluating *Learned* Data Representations

Evaluation of learned data representations

Typically on pre-designed downstream tasks that:

- either rely on labelled data
    - classification tasks [1, 2, 3]
    - prediction tasks [4]

- or are handcrafted
    - robotics task [5, 6]
    - performance of the policy in reinforcement learning [7, 8]

Limitations of such evaluation:

- time consuming
- data hungry
- often no potential downstream tasks
- too many potential downstream tasks bias the evaluation
- handcrafted tasks bias the learning

[1] Chen et al. "Big self-supervised models are strong semi-supervised learners", *Advances in Neural Information Processing Systems 2020.*
[2] Ermolov et al. "Whitening for self-supervised representation learning", *International Conference on Machine Learning 2021.*
[3] Bevilacqua et al. "Size-invariant graph representations for graph classification extrapolation", *International Conference on Machine Learning 2021.*
[4] Li et al. "Learning object-centric representations of multi-object scenes from multiple views", *Advances in Neural Information Processing Systems 2020.*
[5] Chamzas et al. "Comparing reconstruction-and contrastive-based models for visual task planning", *arXiv preprint arXiv:2109.06737, 2021.*
[6] Lippi et al. "Latent space roadmap for visual action planning of deformable and rigid object manipulation", *International Conference on Intelligent Robots and Systems 2020.*
[7] Ghadirzadeh et. al. "Data-efficient visuomotor policy training using reinforcement learning and generative models", *arXiv preprint arXiv:2007.13134, 2020.*
[8] Laskin et al " CURL: Contrastive unsupervised representations for reinforcement learning", *International Conference on Machine Learning 2020.*

Geometric evaluation of representations

Idea: compare topological and geometrical properties of two sets of representations [9, 10, 11, 12]

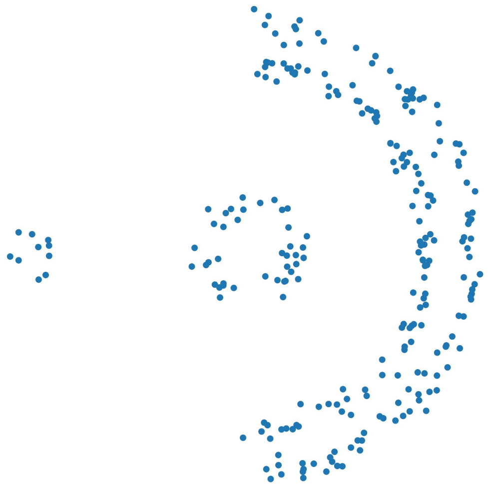

[9] Poklukar et al. "GeomCA: Geometric Evaluation of Data Representations," *2021 28th International Conference on Machine Learning (ICML)*.
[10] Poklukar et al. "Delaunay Component Analysis for Evaluation of Data Representations," *2022 International Conference on Learning Representations (ICLR)*.
[11] Kynkäänniemi et al. "Improved Precision and Recall Metric for Assessing Generative Models," *2019 Conference on Neural Information Processing Systems (NeurIPS)*.
[12] Khrulkov et al. "Geometry Score: {A} Method For Comparing Generative Adversarial Networks," *2018 28th International Conference on Machine Learning (ICML)*.

Geometric evaluation of representations

Idea: compare topological and geometrical properties of <u>two sets</u> of representations [9, 10, 11, 12]

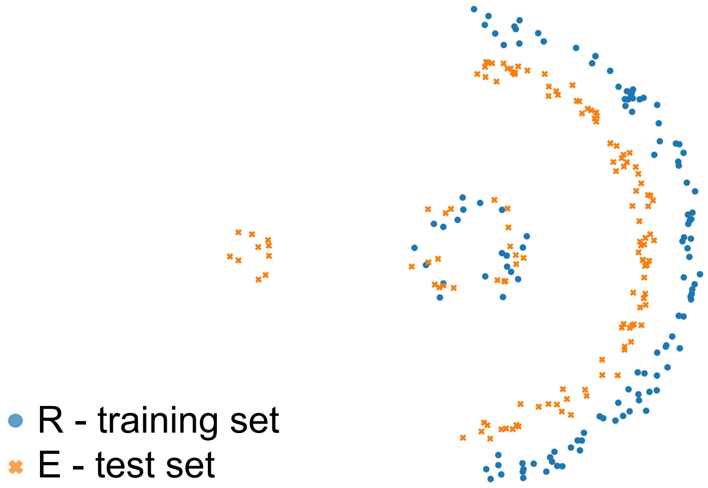

[9] Poklukar et al. "GeomCA: Geometric Evaluation of Data Representations," *2021 28th International Conference on Machine Learning (ICML)*.
[10] Poklukar et al. "Delaunay Component Analysis for Evaluation of Data Representations," *2022 International Conference on Learning Representations (ICLR)*.
[11] Kynkäänniemi et al. "Improved Precision and Recall Metric for Assessing Generative Models," *2019 Conference on Neural Information Processing Systems (NeurIPS)*.
[12] Khrulkov et al. "Geometry Score: {A} Method For Comparing Generative Adversarial Networks," *2018 28th International Conference on Machine Learning (ICML)*.

Geometric evaluation of representations

Idea: compare topological and geometrical properties of two sets of representations [9, 10, 11, 12]

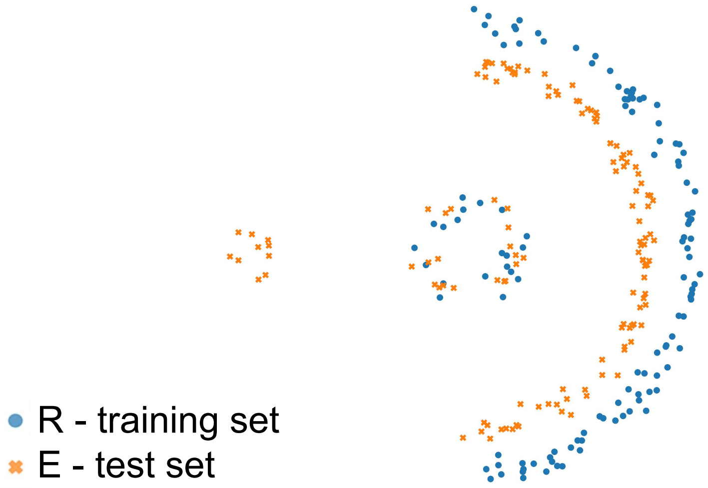

Needed:

1. “give volume to points”

2. “compare orange volume with the blue one”

[9] Poklukar et al. "GeomCA: Geometric Evaluation of Data Representations," *2021 28th International Conference on Machine Learning (ICML)*.
[10] Poklukar et al. "Delaunay Component Analysis for Evaluation of Data Representations," *2022 International Conference on Learning Representations (ICLR)*.
[11] Kynkäänniemi et al. "Improved Precision and Recall Metric for Assessing Generative Models," *2019 Conference on Neural Information Processing Systems (NeurIPS)*.
[12] Khrulkov et al. "Geometry Score: {A} Method For Comparing Generative Adversarial Networks," *2018 28th International Conference on Machine Learning (ICML)*.

Geometric evaluation of representations

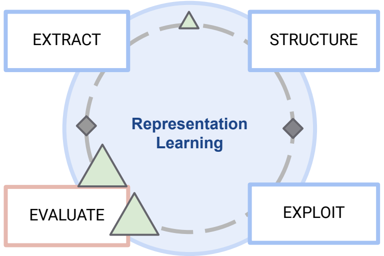

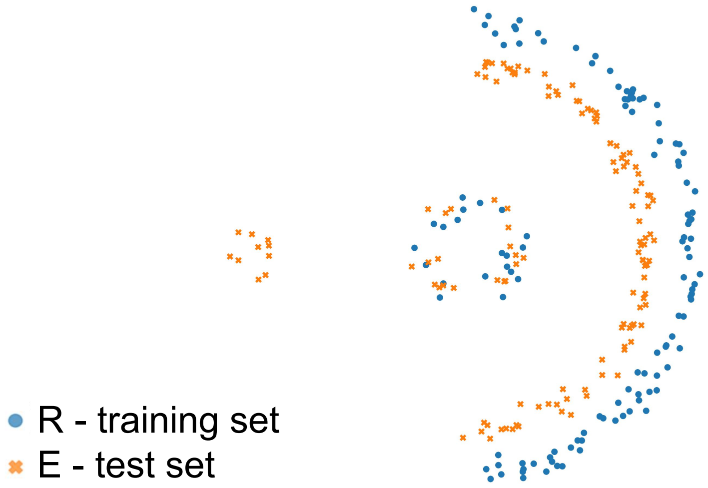

Paper B
Delaunay Component Analysis for Evaluation of Data Representations, Petra Poklukar, Vladislav Polianskii, Anastasiia Varava, Florian T. Pokorny, Danica Kragic, *ICLR 2022*

Needed:

1. “give volume to points”
2. “compare orange volume with the blue one”

Paper A
GeomCA: Geometric Evaluation of Data Representations, Petra Poklukar, Anastasia Varava, Danica Kragic, *ICML 2021*

Geometric Component Analysis (GeomCA)

Idea: compare topological and geometrical properties of two sets of representations

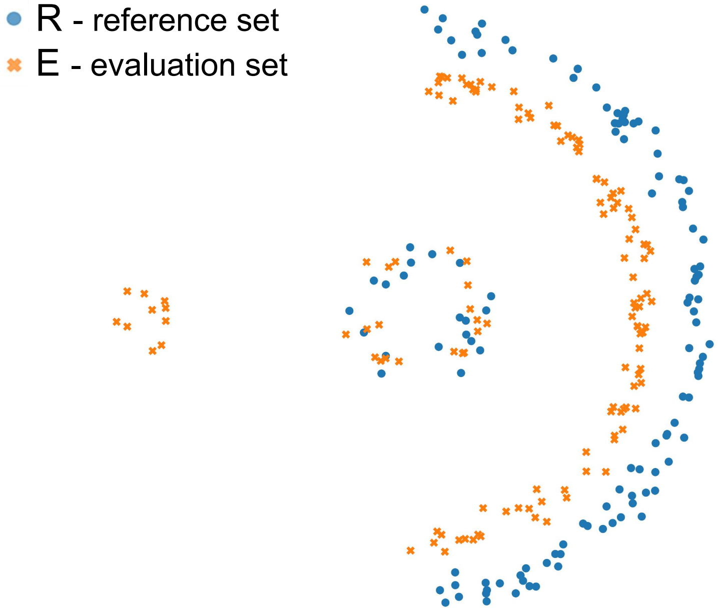

Geometric Component Analysis (GeomCA)

Idea: compare topological and geometrical properties of two sets of representations

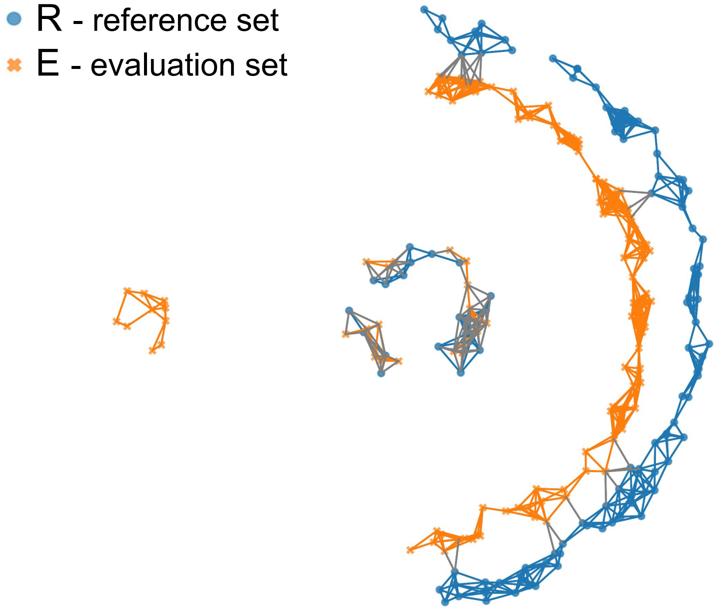

Simple approach: build a graph and derive <u>GeomCA scores</u> by analysing its connected components

Geometric Component Analysis (GeomCA)

Idea: compare topological and geometrical properties of two sets of representations

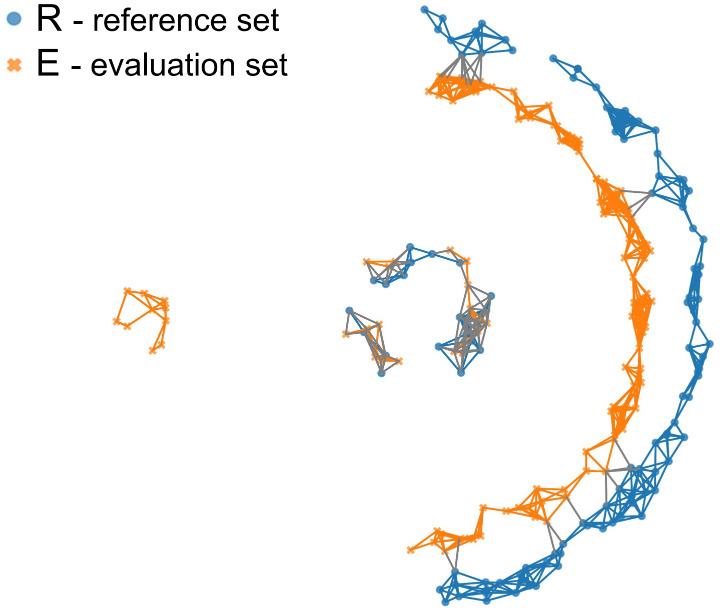

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

- diversity of *points* from both sets

Geometric Component Analysis (GeomCA)

Idea: compare topological and geometrical properties of two sets of representations

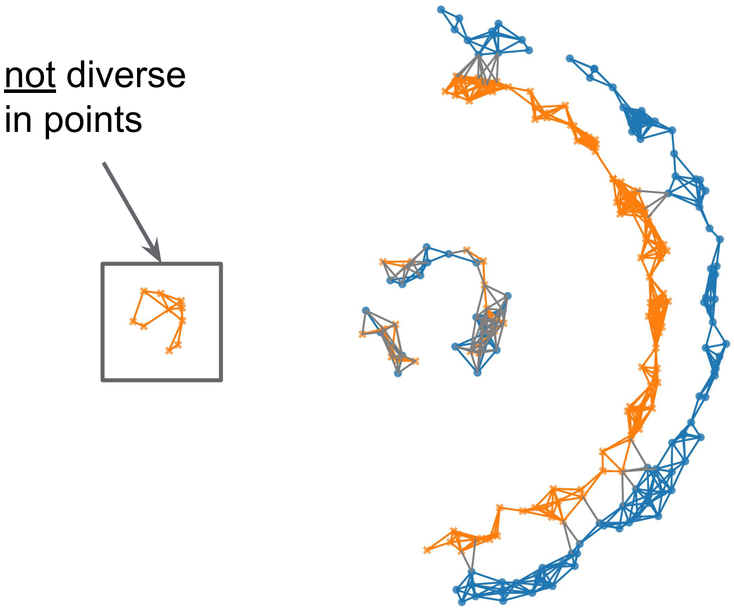

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

- diversity of *points* from both sets

Geometric Component Analysis (GeomCA)

Idea: compare topological and geometrical properties of two sets of representations

not diverse
in points

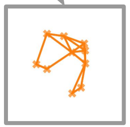

diverse in
points

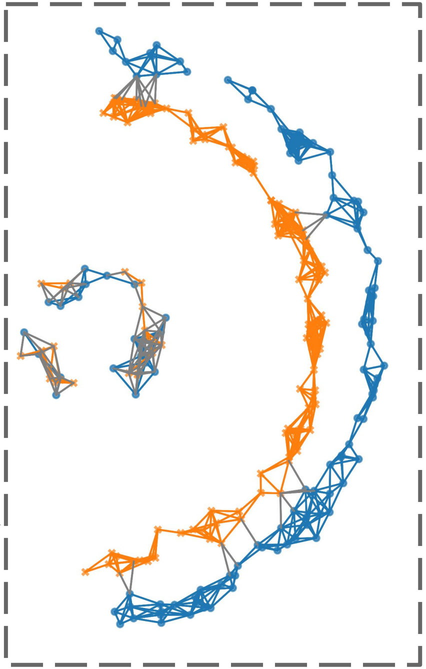

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

- diversity of *points* from both sets

Geometric Component Analysis (GeomCA)

Idea: compare topological and geometrical properties of two sets of representations

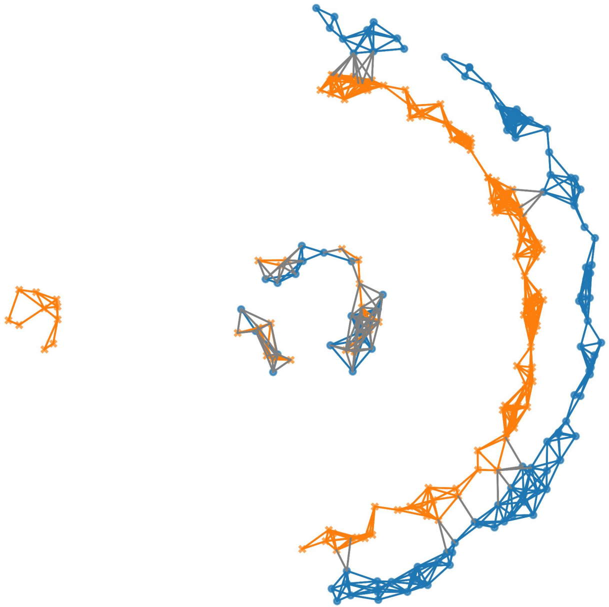

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

- diversity of *points* from both sets
- diversity of *edges* among points from each of the sets

Geometric Component Analysis (GeomCA)

Idea: compare topological and geometrical properties of two sets of representations

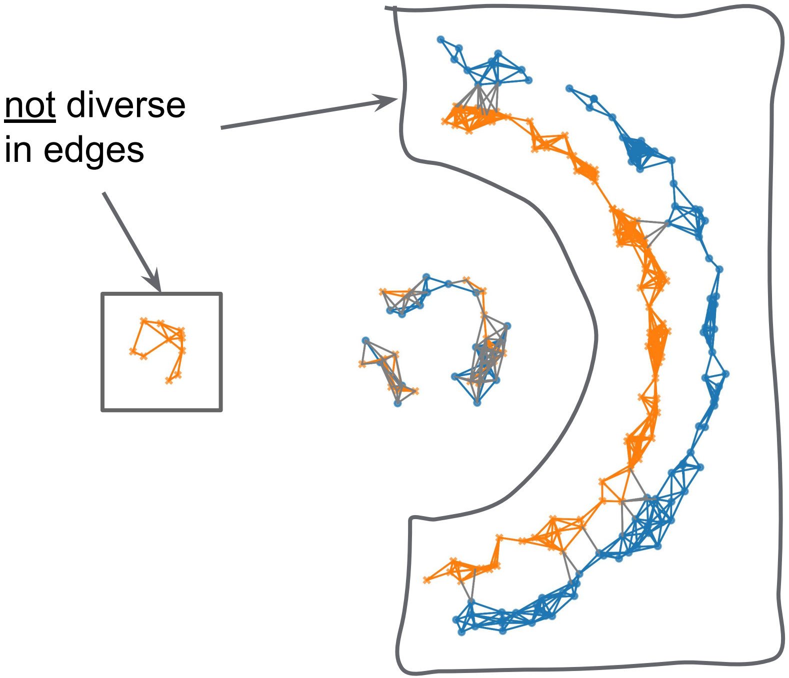

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

- diversity of *points* from both sets
- diversity of *edges* among points from each of the sets

Geometric Component Analysis (GeomCA)

Idea: compare topological and geometrical properties of two sets of representations

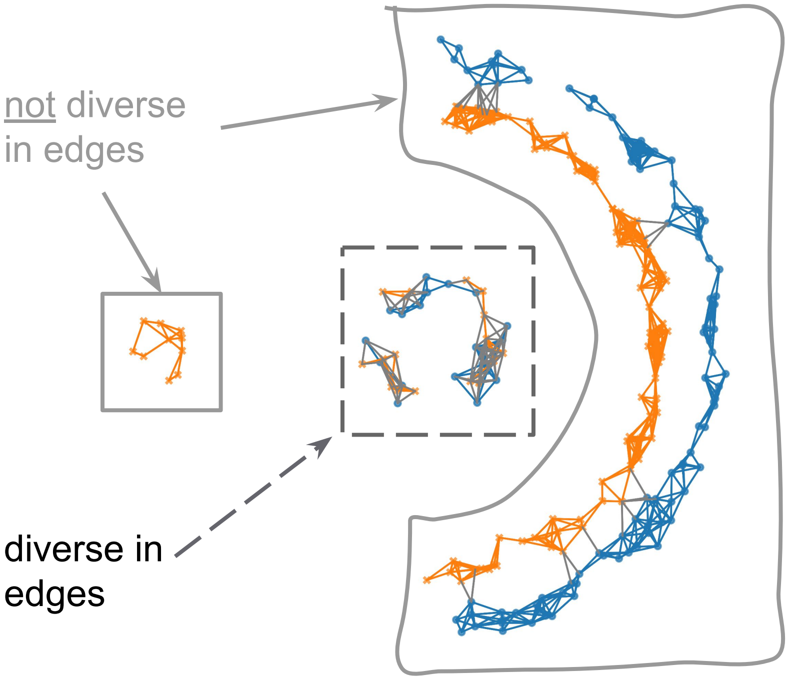

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

- diversity of *points* from both sets
- diversity of *edges* among points from each of the sets

GeomCA: limitation

Relies on ε-proximity graphs

Limitation #2

- no functionality to evaluate the *quality of a single query representation*

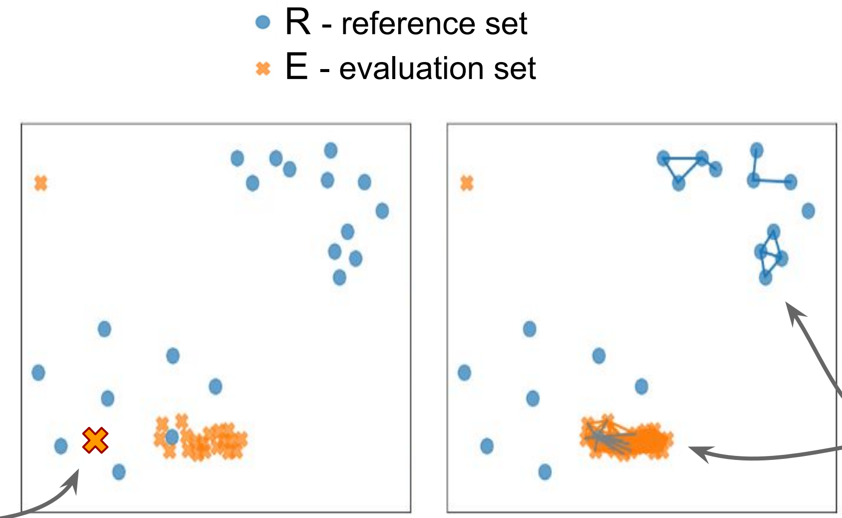

original points

ε-proximity graph

Limitation #1

- *single value of* ε does not always capture clusters of different shape and density

Paper B
Delaunay Component Analysis for Evaluation of Data Representations

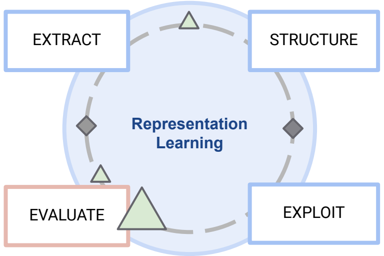

Petra Poklukar, Vladislav Polianskii, Anastasiia Varava, Florian T. Pokorny, Danica Kragic

International Conference on Learning Representations 2022

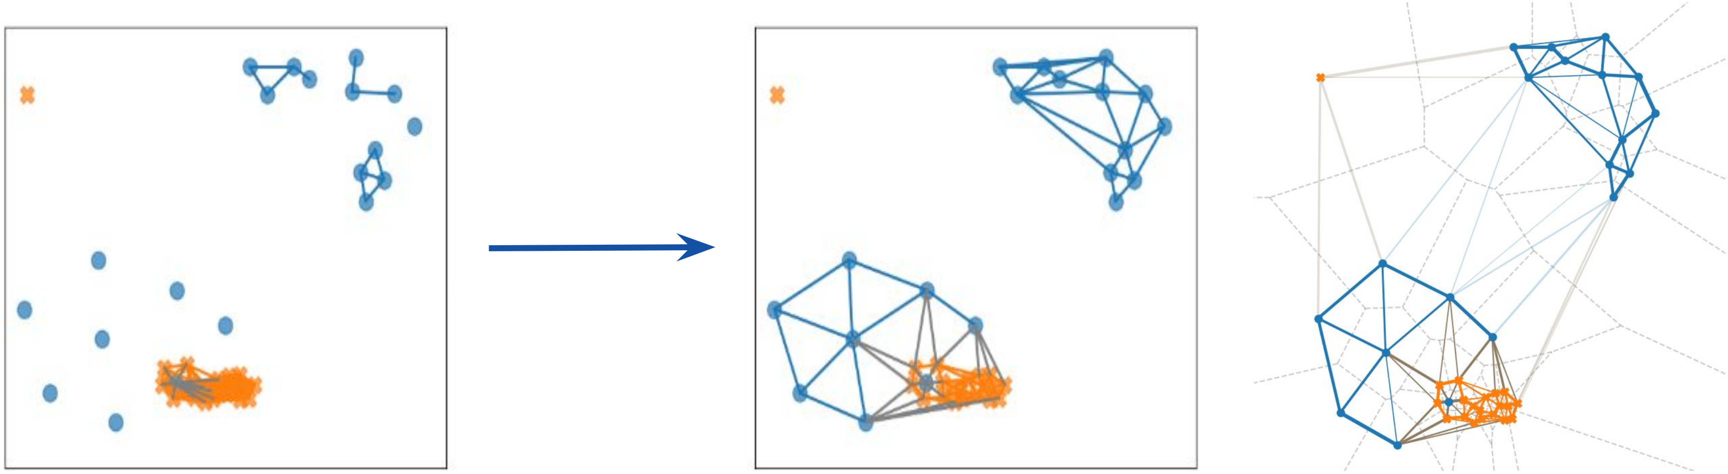

Applications of GeomCA and DCA

1. Evaluation of generative models [9, 10]

2. Analysis of semantic similarity in representation spaces [9, 10]

3. Evaluation of contrastive representations of multimodal data [13] (Paper D)

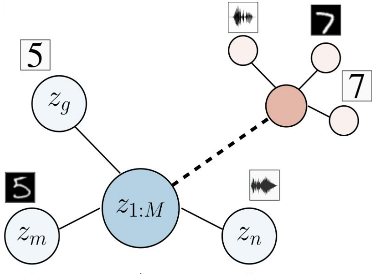

4. Comparing large input graphs [14] (Paper C)

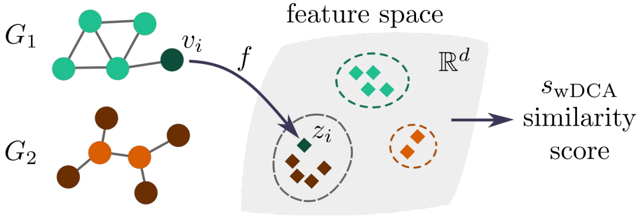

[9] Poklukar et al. "GeomCA: Geometric Evaluation of Data Representations," *2021 International Conference on Machine Learning (ICML)*.
[10] Poklukar et al. "Delaunay Component Analysis for Evaluation of Data Representations," *2022 International Conference on Learning Representations (ICLR)*.
[13] Poklukar et al. "GMC: Geometric Multimodal Contrastive Representation Learning," *2022 International Conference on Machine Learning (ICML)*.
[14] Ceylan et al. "GraphDCA: A Framework for Node Distribution Comparison in Real and Synthetic Graphs," *(under review)*.

Paper D
GMC: Geometric Multimodal Contrastive Representation Learning

Petra Poklukar*, Miguel Vasco*, Hang Yin, Francisco S. Melo, Ana Paiva, Danica Kragic

International Conference on Machine Learning 2022

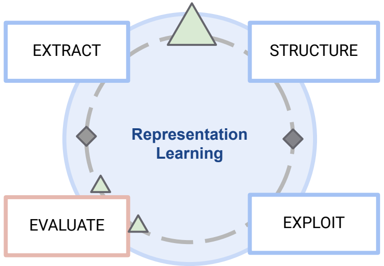

Goal: learn representations of multimodal data that are (informative and) robust to missing modalities at test time

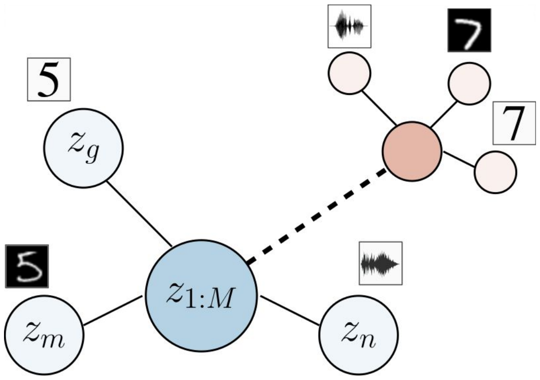

Paper D

GMC: Geometric Multimodal Contrastive Representation Learning

Petra Poklukar*, Miguel Vasco*, Hang Yin, Francisco S. Melo, Ana Paiva, Danica Kragic

International Conference on Machine Learning 2022

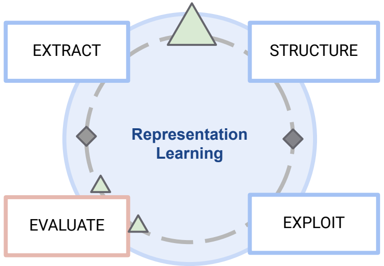

Goal: learn representations of multimodal data that are (informative and) robust to missing modalities at test time

(a) MVAE (Wu & Goodman, 2018)

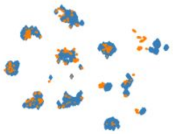

(b) MMVAE (Shi et al., 2019)

(c) Nexus (Vasco et al., 2022)

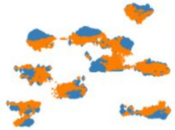

(d) MUSE (Vasco et al., 2021)

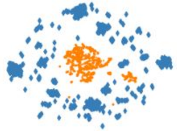

(e) MFM (Tsai et al., 2018)

(f) **GMC (Ours)**

Quantify the geometric alignment with DCA (Paper B)

- R - complete representations
- E - image modality representations

Applications

Representation Learning: main challenges

To learn useful data representations, we need to consider the following two problems:

1. How to **design** representation **learning models** that identify semantically useful information and encode it into structured low dimensional representations?

   - Paper D - ICML: Geometric Multimodal Contrastive Learning
   - Paper E - T-RO: Latent Space Roadmap

2. How to **design** reliable **evaluation frameworks** for assessing the quality of the resulting representations?

   - Paper A - ICML: Geometric Component Analysis
   - Paper B - ICLR: Delaunay Component Analysis
   - Paper C - preprint: GraphDCA

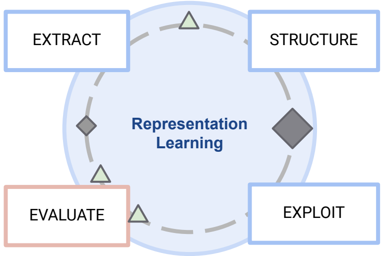

Paper E
Enabling Visual Action Planning for Object Manipulation through Latent Space Roadmap

Martina Lippi*, Petra Poklukar*, Michael C. Welle*, Anastasiia Varava, Hang Yin, Alessandro Marino and Danica Kragic

Conditionally accepted to IEEE Transactions of Robotics

Goal: visual action planning of complex manipulation tasks with high-dimensional state spaces such as deformable objects

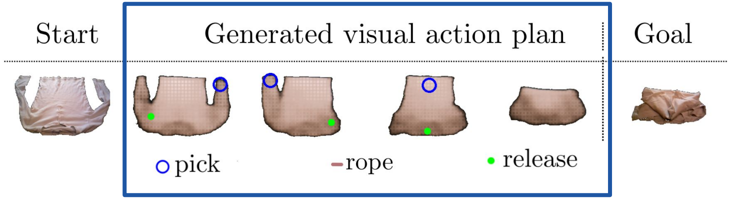

Approach: use representations learned by a VAE

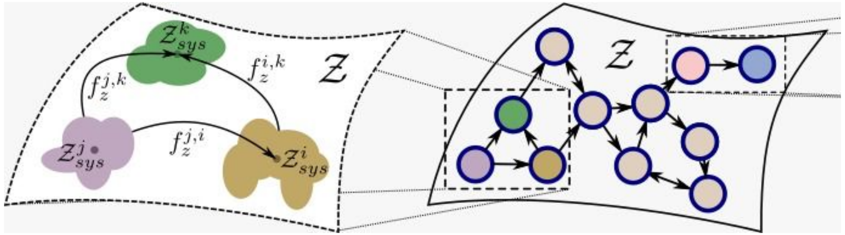

Paper C
GraphDCA: a Framework for Node Distribution Comparison in Real and Synthetic Graphs

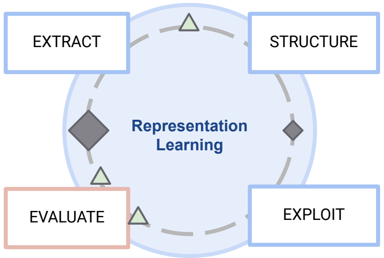

Ciwan Ceylan*, Petra Poklukar*, Hanna Hultin, Alexander Kravberg, Anastasia Varava, Danica Kragic

Preprint

Goal: develop an evaluation procedure for comparing input graphs in terms of their node structural features

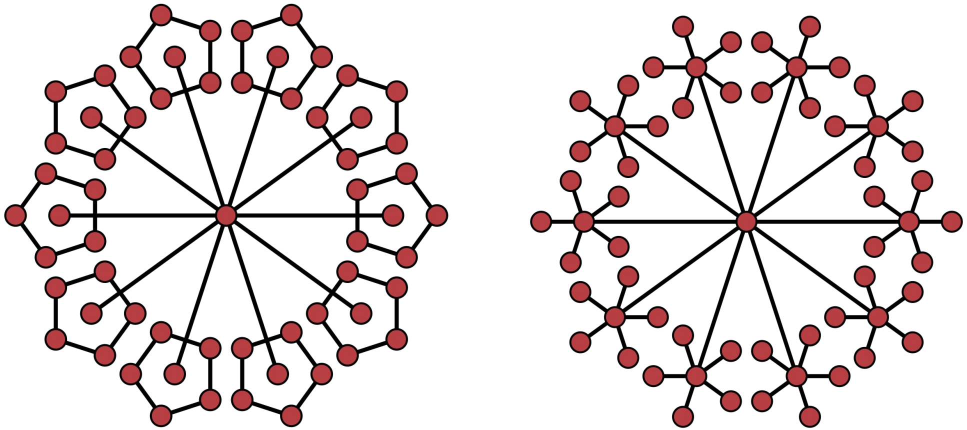

Paper C
GraphDCA: a Framework for Node Distribution Comparison in Real and Synthetic Graphs

Ciwan Ceylan*, Petra Poklukar*, Hanna Hultin, Alexander Kravberg, Anastasia Varava, Danica Kragic

Preprint

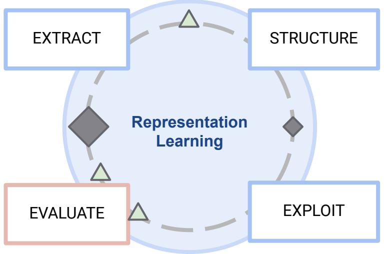

Goal: develop an evaluation procedure for comparing input graphs in terms of their node structural features

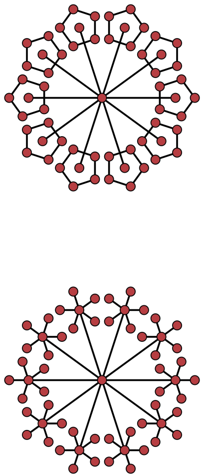

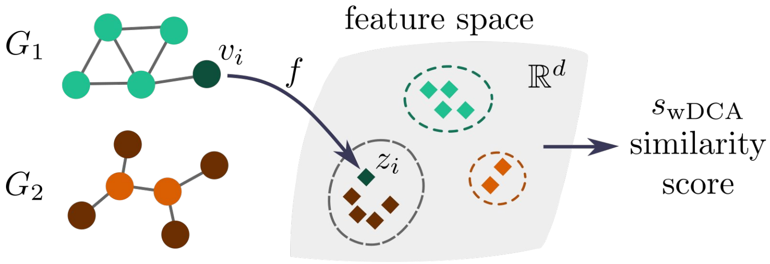

1. Extract local structural properties as node representations
2. Analyze their alignment with DCA (Paper B)

Future directions

- Geometric regularization of deep learning models
  - Extension of GMC to a subset of modalities during test time

- Applications of the geometric evaluation frameworks to
  - other data domains, e.g., biology
  - non-Euclidean representation spaces

- Develop an approximate DCA that scales to very large sets

Key takeaways

1. **Well-structured** representation spaces improve the performance of downstream tasks

2. **Geometry**-based evaluation of representation spaces **can offer** valuable **insights into** semantic similarities of **representations**

Learning and Evaluating the Geometric Structure of Representation Spaces

PhD Thesis Defense

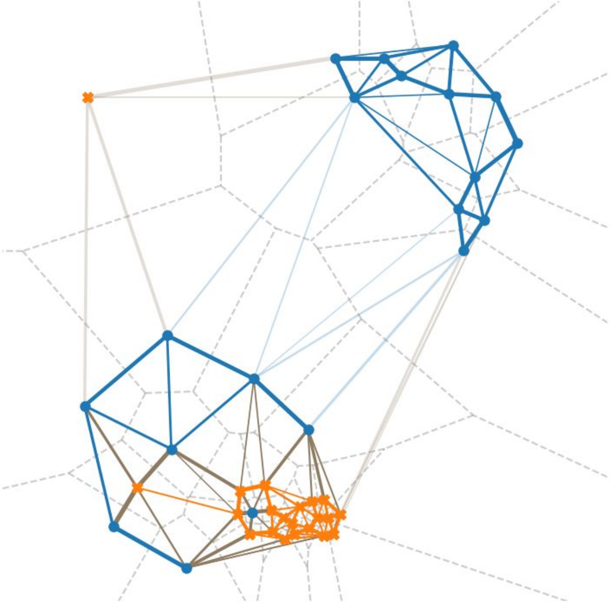

Petra Poklukar
poklukar@kth.se
https://people.kth.se/~poklukar/