KTH ROYAL INSTITUTE
OF TECHNOLOGY

Learning and Evaluating Representations: Geometric and Topological Aspects

Petra Poklukar

Outline

- **Data Representations: Learning and Evaluation**

- **Geometric Aspects of Evaluation and Representation Learning**
  - Paper A - ICML: Geometric Component Analysis
  - Paper B - ICLR: Delaunay Component Analysis
  - *Work In Progress*: Geometric Multimodal Contrastive Learning
  - *Work In Progress*: Improved Training of Generative Models

- **Applications**
  - Paper C & D - IROS, T-RO: Latent Space Roadmap
  - Paper E - JMLR: Data-Efficient Visuomotor Policy Training using Reinforcement Learning and Generative Models
  - *Work In Progress*: Evaluation of Graph Generative Models

Theoretical contributions

Practical contributions

Data Representations: Learning and Evaluation

Representation Learning: what and why

What:

Learning more *general* and *abstract* representations of the data that make it easier to extract useful information when building *deep learning models*

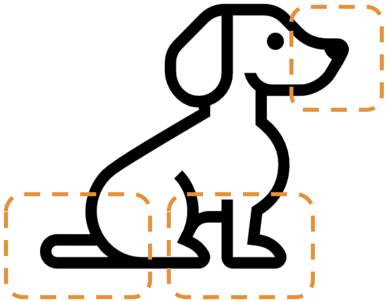

Why:

An AI system should learn to *identify* and *structure* the underlying semantic information hidden in the observed data

Learning and Evaluating Representations

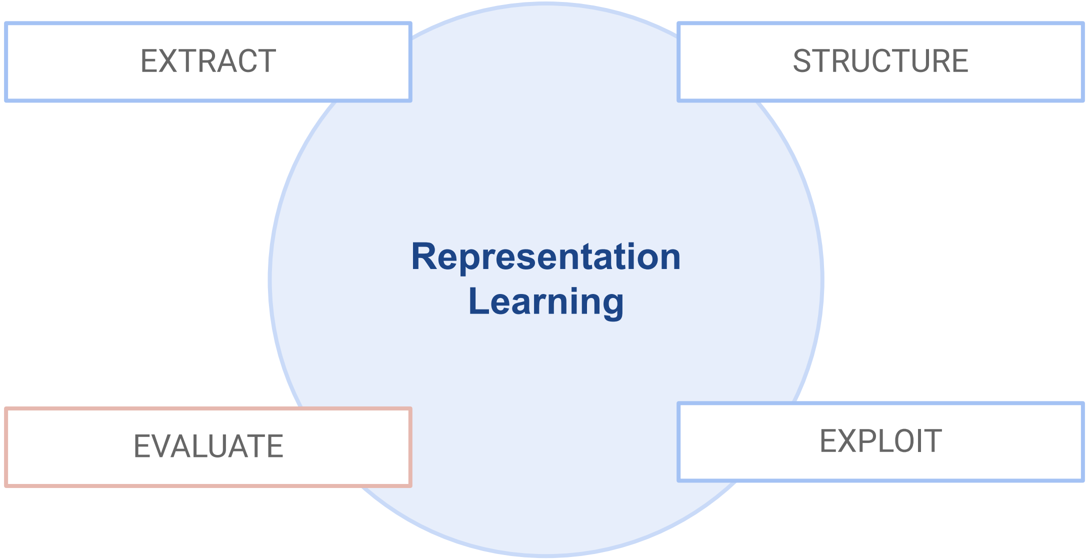

Learning and Evaluating Representations

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

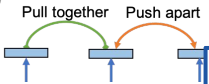

Locatello et al., *A Sober Look at the Unsupervised Learning of Disentangled Representations and their Evaluation*, JMLR 2020

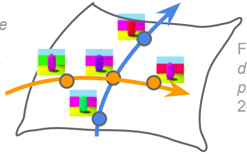

Fumero et al., *Learning disentangled representations via product manifold projection*, ICML 2021

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

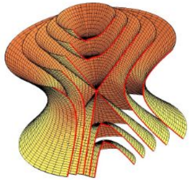

Korman, Self-supervised representation learning on manifolds, ICLR 2021 Workshop on Geometrical and Topological Representation Learning

EXTRACT

STRUCTURE

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

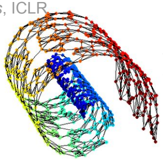

McInnes et al., (2018). *UMAP: Uniform Manifold Approximation and Projection.* Journal of Open Source Software, 3(29), 861

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

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

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

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

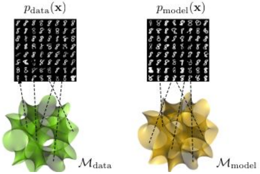

Kynkäänniemi et al., *Improved precision and recall metric for assessing generative models*, NeurIPS 2019

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

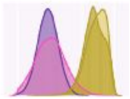

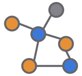

Bevilacqua et al., *Size-Invariant Graph Representations for Graph Classification Extrapolations*, ICML 2021

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

Learning and Evaluating Representations

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

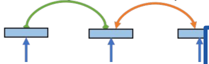

Locatello et al., *A Sober Look at the Unsupervised Learning of Disentangled Representations and their Evaluation*, JMLR 2020

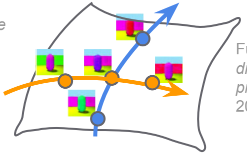

Fumero et al., *Learning disentangled representations via product manifold projection*, ICML 2021

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

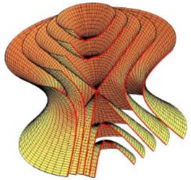

EXTRACT

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.

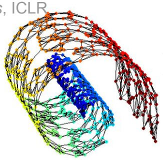

McInnes et al., (2018). *UMAP: Uniform Manifold Approximation and Projection.* Journal of Open Source Software, 3(29), 861

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

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

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

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

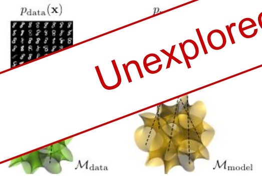

Unexplored

...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

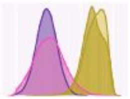

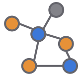

Bevilacqua et al., *Size-Invariant Graph Representations for Graph Classification Extrapolations*, ICML 2021

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

Learning and Evaluating Representations: Geometric and Topological Aspects

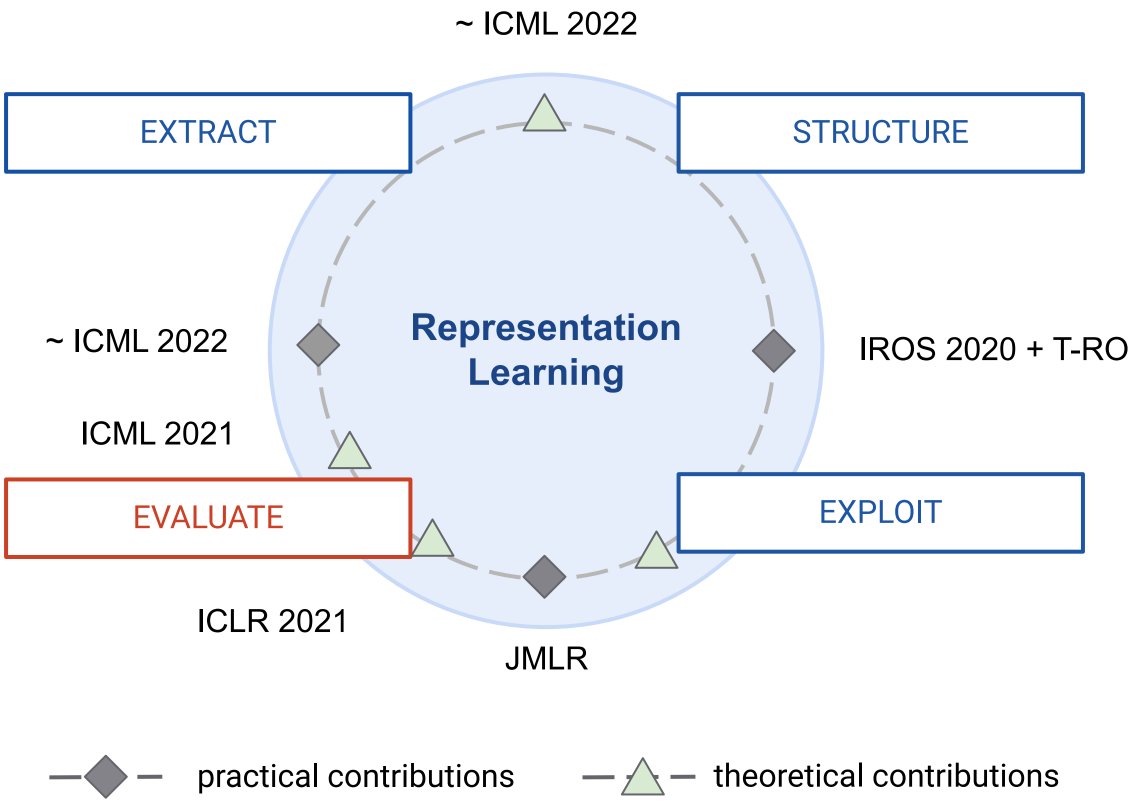

My work so far: statistics

conference publications: 3

workshops: 5

journal submissions: 2

collaborators: 15+

submissions in pipeline: 2

Personal website: https://people.kth.se/~poklukar/

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

[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.*

What if we cannot define an evaluation task?

Example: generative models

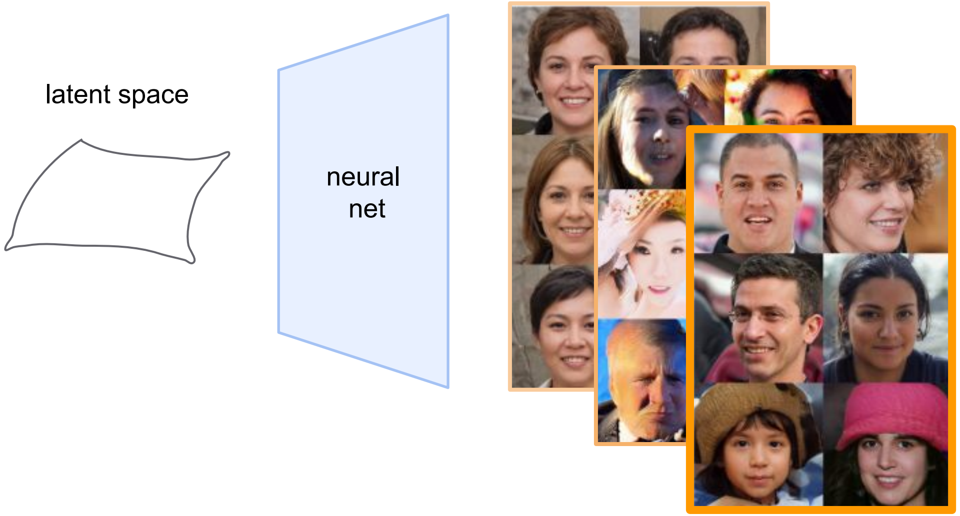

How to evaluate the quality of the generated images?

Image source: Kynkäänniemi et. al., *Improved Precision and Recall Metric for Assessing Generative Models* in Advances in Neural Information Processing Systems (NeurIPS) 2019

What if we cannot define an evaluation task?

Example: generative models

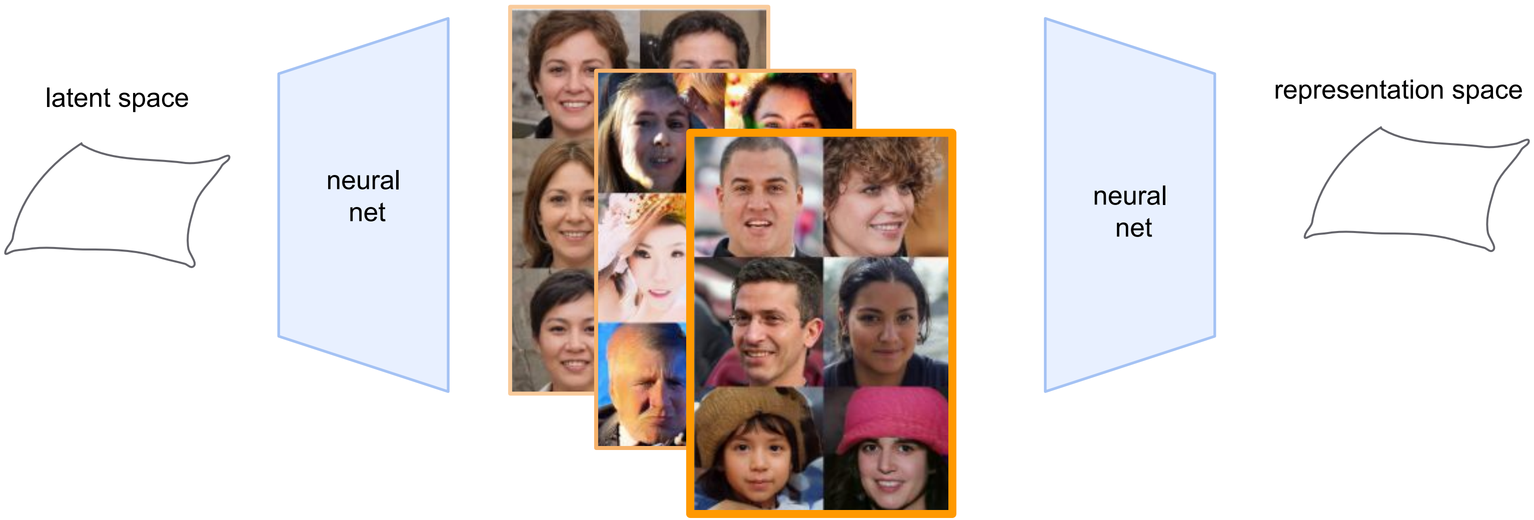

Image source: Kynkäänniemi et. al., *Improved Precision and Recall Metric for Assessing Generative Models* in Advances in Neural Information Processing Systems (NeurIPS) 2019

Geometry to the Rescue

Geometric evaluation frameworks

Idea: compare topological and geometrical properties of <u>two sets</u> of representations

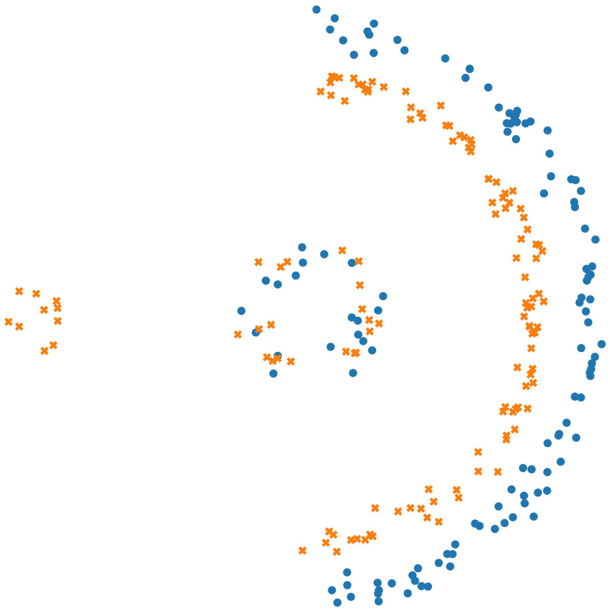

training

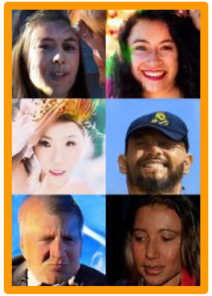

generated

Needed:

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

Image source: Kynkäänniemi et. al., *Improved Precision and Recall Metric for Assessing Generative Models* in Advances in Neural Information Processing Systems (NeurIPS) 2019

Ways to approximate data manifolds

graphs

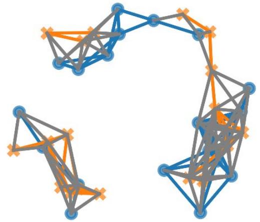

simplicial complexes

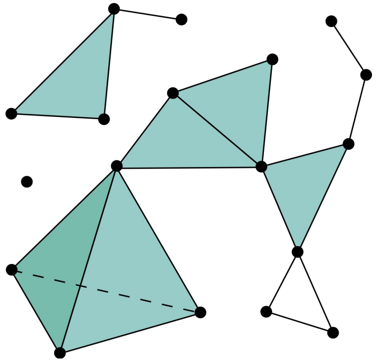

Khrulkov & Oseledets *Geometry Score: A Method For Comparing Generative Adversarial Networks*. ICML 2018

spherical approximations

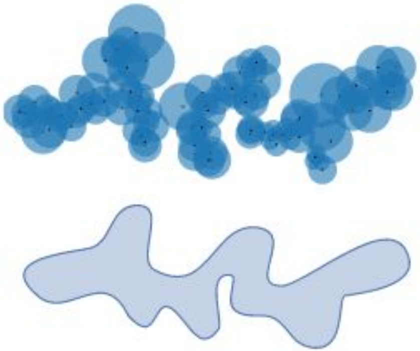

Kynkäänniemi et al., *Improved precision and recall metric for assessing generative models*, NeurIPS 2019

Once we have a manifold...

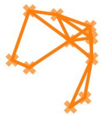

vs

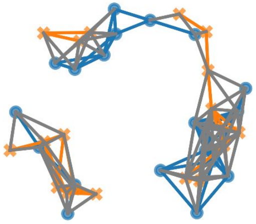

What to compare?

- Number of connected components (0-dim holes)
- Number of 1-dim (or n-dim) holes
- Alignment of points
- Statistics on edges
- ….

Topology
properties preserved under continuous deformations, such as stretching, twisting, crumpling, and bending

Geometry
properties related to distance, shape, size and relative position

Geometric evaluation frameworks: related work

simplicial complexes

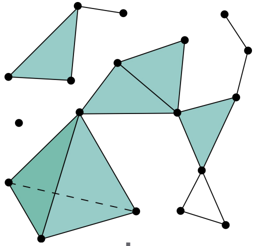

Number of n-dim holes

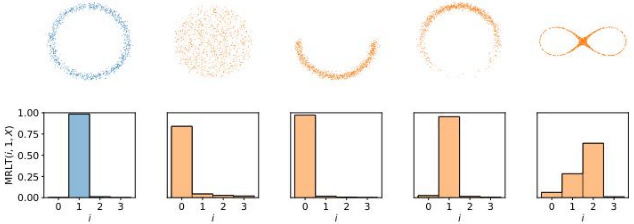

Khrulkov & Oseledets ***Geometry Score****: A Method For Comparing Generative Adversarial Networks*. ICML 2018

spherical approximations

Alignment of points

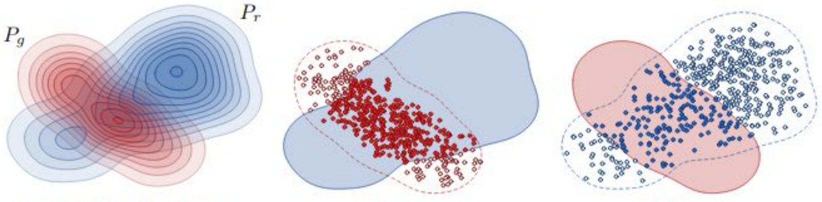

(a) Example distributions (b) Precision (c) Recall

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

GeomCA: Geometric Evaluation of Data Representations

Petra Poklukar, Anastasia Varava, Danica Kragic

Proceedings of the 38th International Conference on Machine Learning

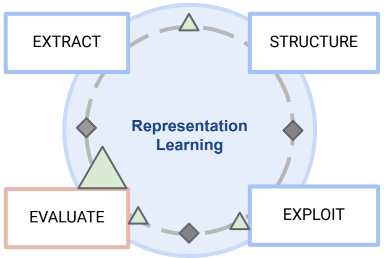

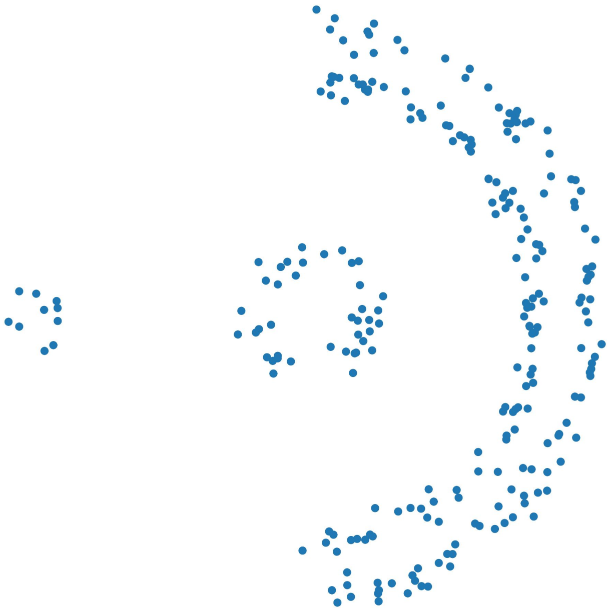

How to evaluate the quality of *any* data representations without relying on downstream tasks?

How to evaluate the quality of learned data representations?

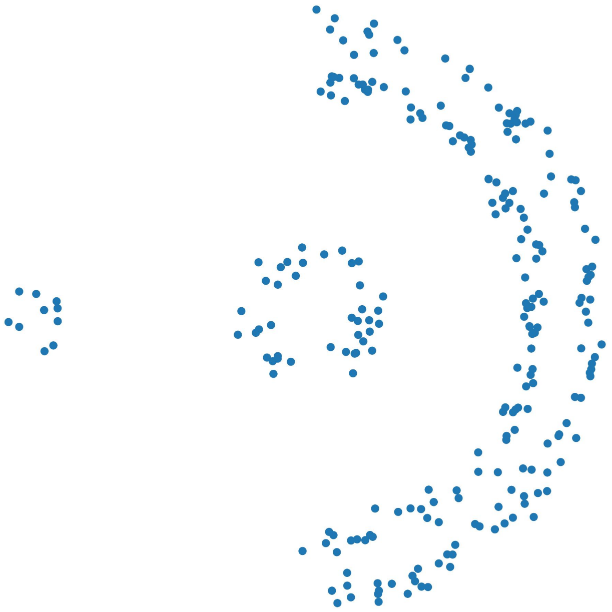

An evaluation framework should be

- **General** and independent of**:**
  - *downstream evaluation task*
  - *the model* generating representations
  - *dimensionality* of representations
- **Informative** and provide insights into**:**
  - the *global structure* of learned representation space
  - *local areas* where potential failures arise

Geometric Component Analysis (GeomCA)

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

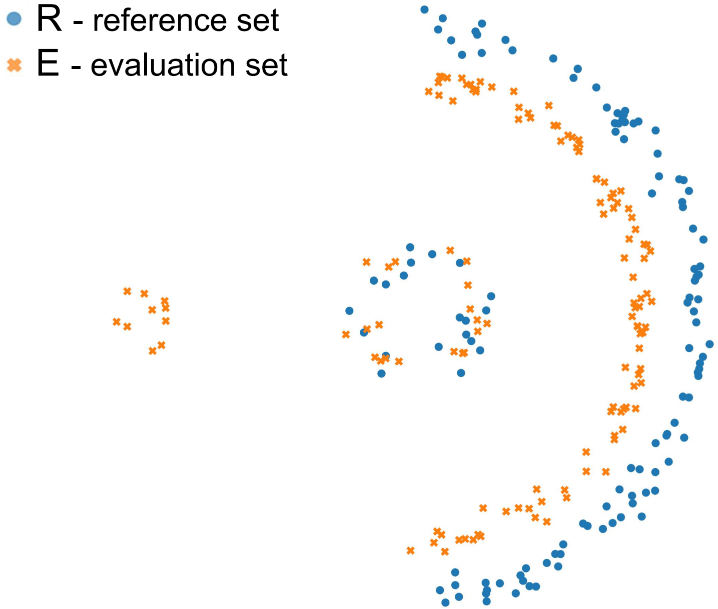

Geometric Component Analysis (GeomCA)

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

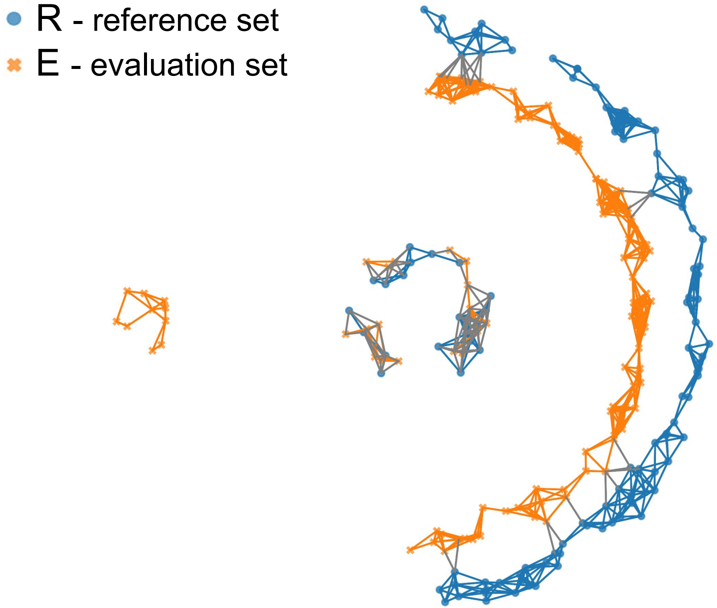

Simple approach: build a graph and derive GeomCA scores by analysing its connected components

Geometric Component Analysis (GeomCA)

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

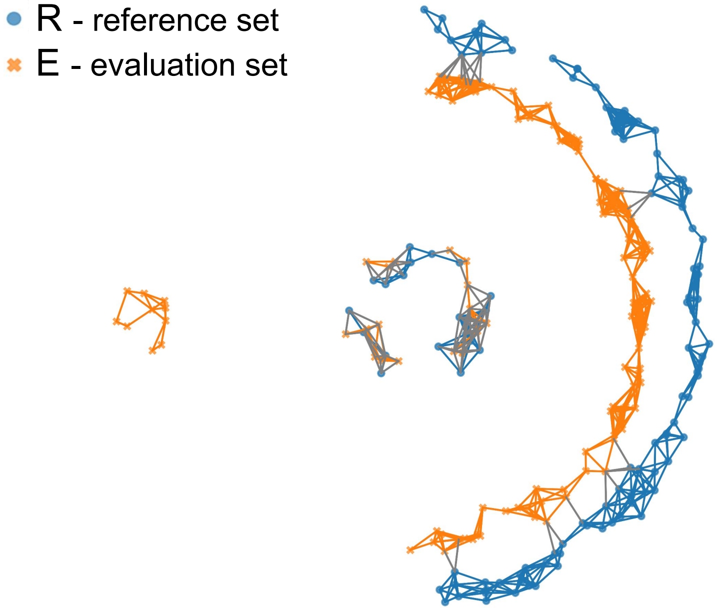

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

- diversity of *points* from both sets

Geometric Component Analysis (GeomCA)

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

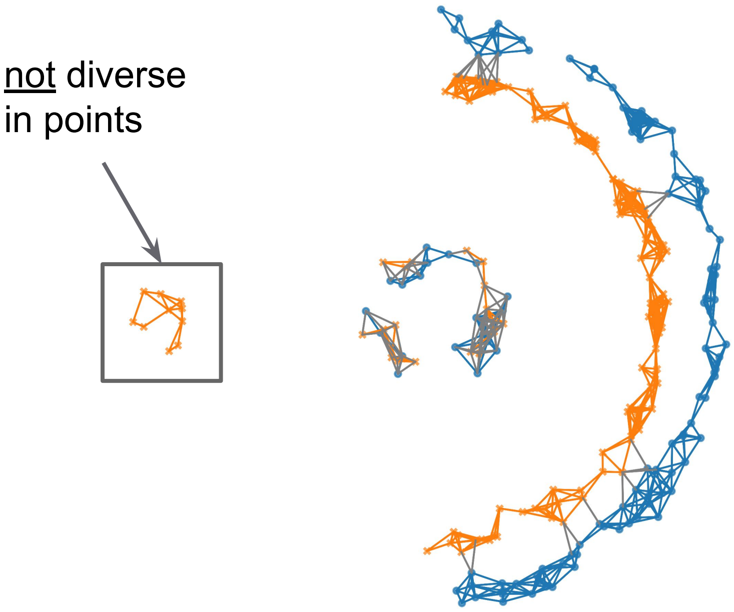

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

- diversity of *points* from both sets

Geometric Component Analysis (GeomCA)

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

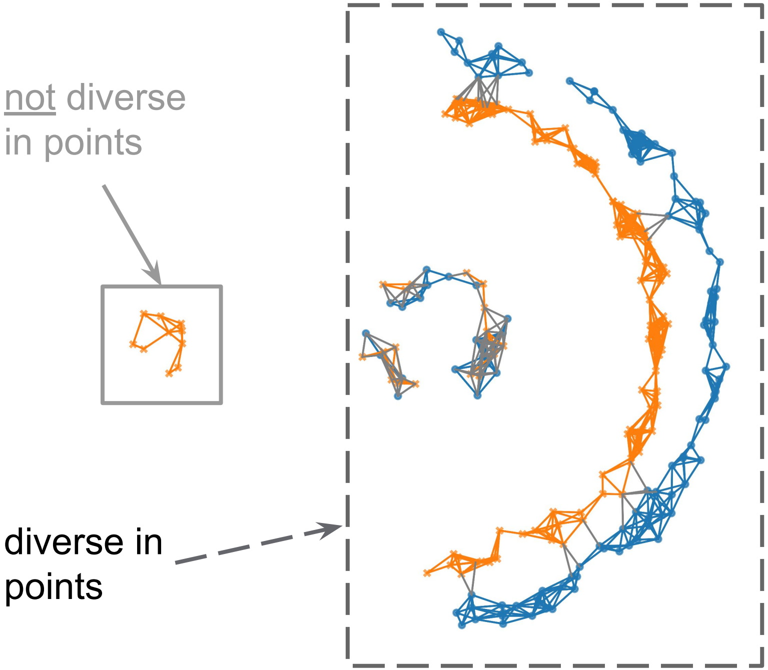

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

- diversity of *points* from both sets

Geometric Component Analysis (GeomCA)

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

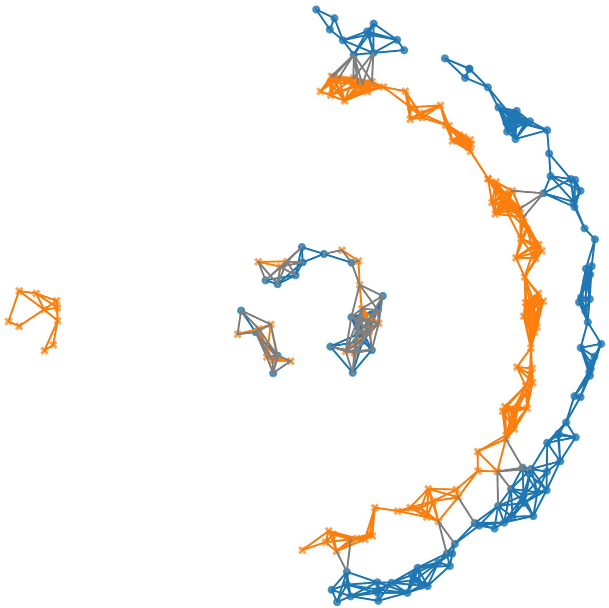

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

- 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

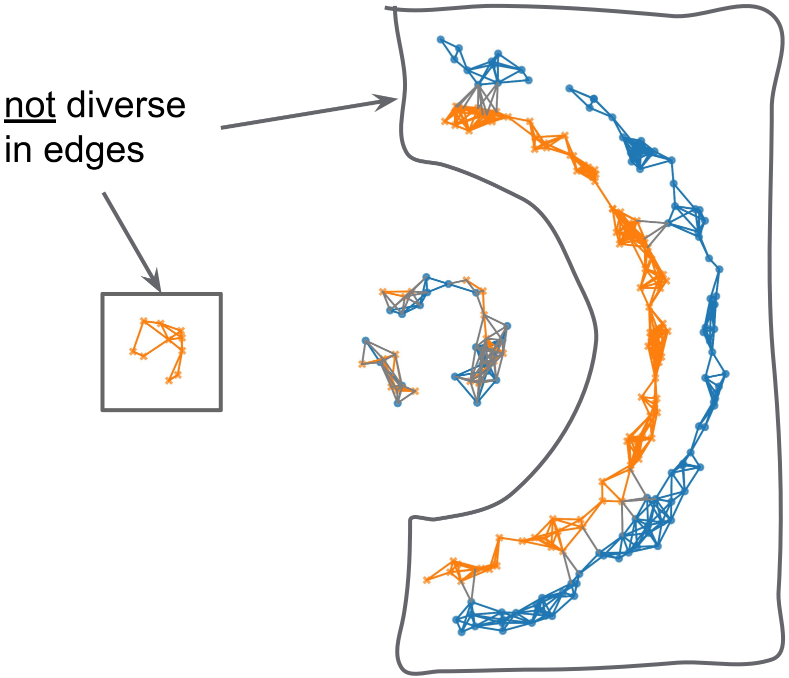

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

- 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

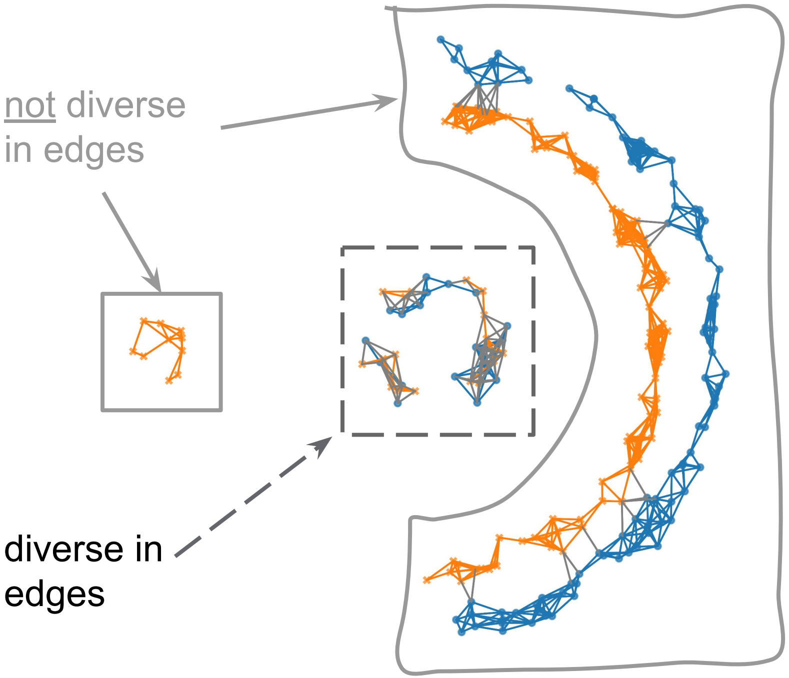

Simple approach: build a graph and derive GeomCA scores by analysing its connected components in terms of

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

GeomCA scores can be used for both ...

global evaluation ...

for example, to investigate mode collapse

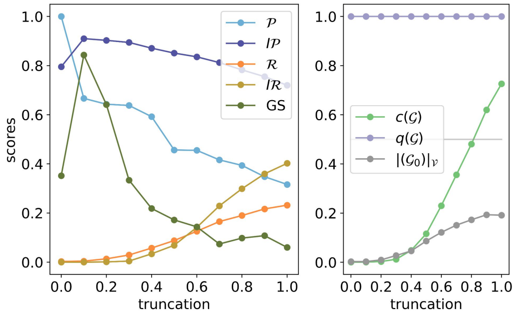

Investigating mode collapse in StyleGAN trained on the FFHQ dataset [1]

… and *local* evaluation

for example, to detect outliers and identify points in specific components

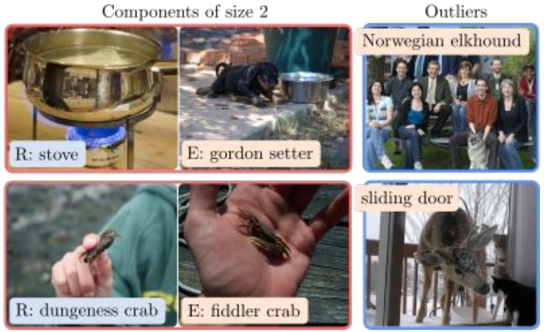

Investigating individual ImageNet points in VGG16 representation space [2, 3]

[1] Karras et. al., "A style-based generator architecture for generative adversarial networks." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019.
[2] Liu et al. "Very deep convolutional neural network based image classification using small training sample size," *2015 3rd IAPR Asian Conference on Pattern Recognition (ACPR)*.
[3] Deng et al. "ImageNet: A large-scale hierarchical image database," *2009 IEEE Conference on Computer Vision and Pattern Recognition*.

GeomCA: graph construction

Relies on ε-proximity graphs

- two points connected by an edge if they are at most ε distance apart

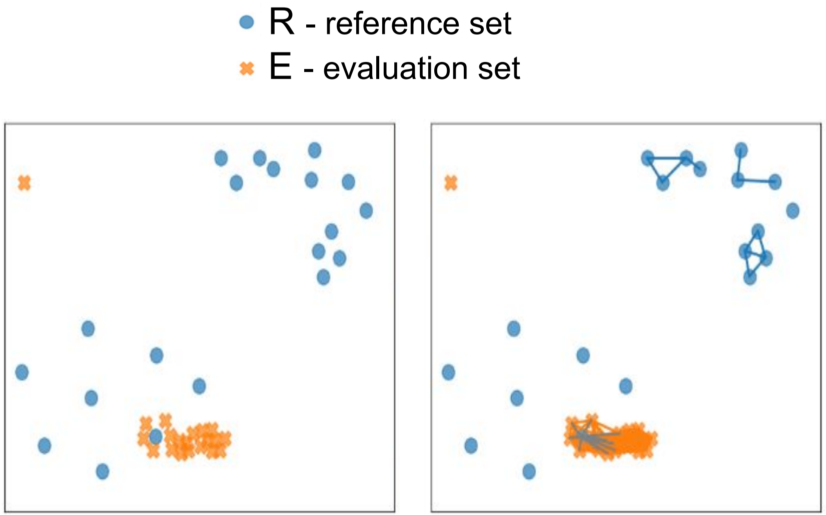

original points

ε-proximity graph

GeomCA: limitation

Relies on ε-proximity graphs

- two points connected by an edge if they are at most ε distance apart

Limitation #2

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

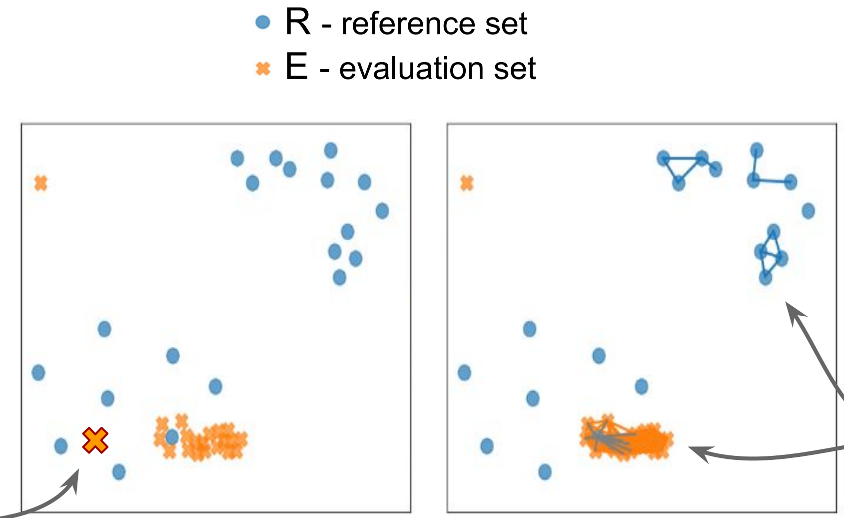

original points ε-proximity graph

Limitation #1

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

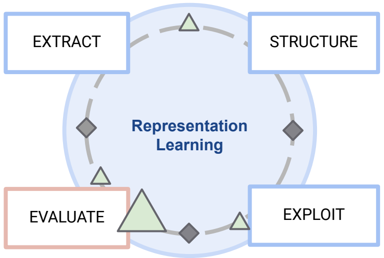

Delaunay Component Analysis for Evaluation of Data Representations

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

International Conference on Learning Representations 2022

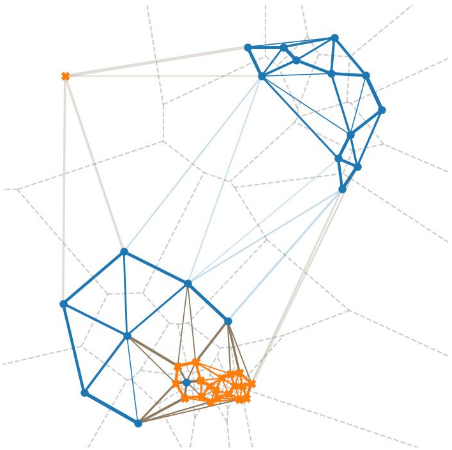

Delaunay Component Analysis (DCA)

Relies on natural neighbourhood graphs called Delaunay graphs

- two points connected by an edge if they are “spatial” neighbours, i.e., if their Voronoi cells intersect

● R - reference set
✖ E - evaluation set

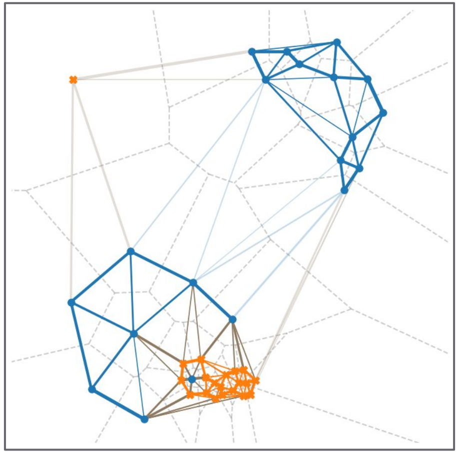

Delaunay graph

Delaunay Component Analysis (DCA): contributions

Relies on natural neighbourhood graphs called Delaunay graphs

- two points connected by an edge if they are “spatial” neighbours, i.e., if their Voronoi cells intersect

easy evaluation of the **quality of a single query representation**

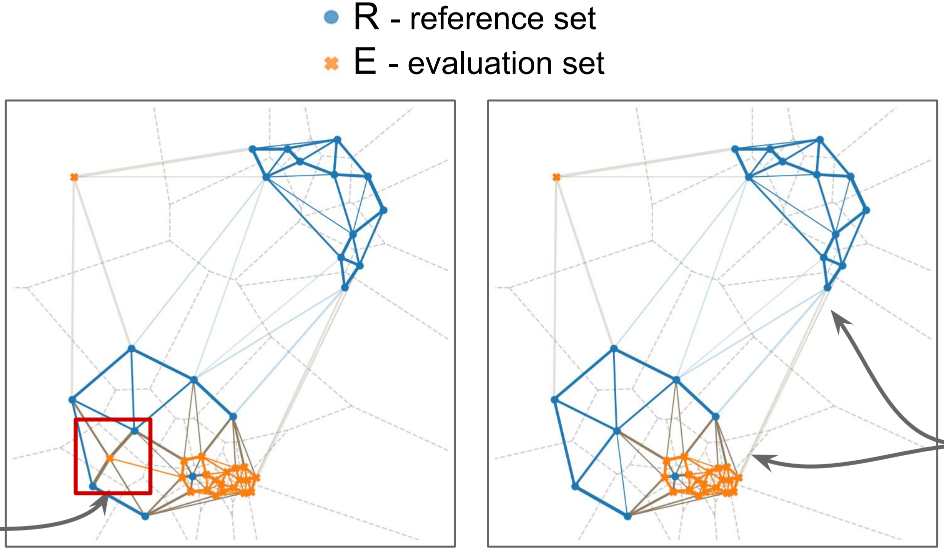

Delaunay graph

naturally capture **local changes in the density** of the representation space

DCA framework

1. Approximate Delaunay graph

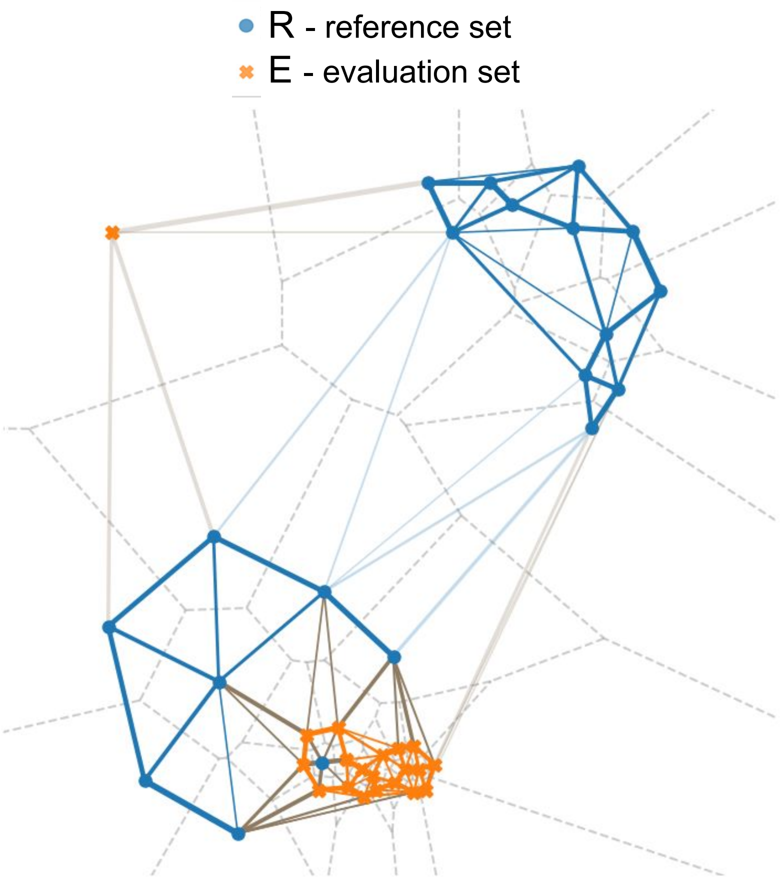

DCA framework

1. Approximate Delaunay graph

2. Distill Delaunay graph into connected components

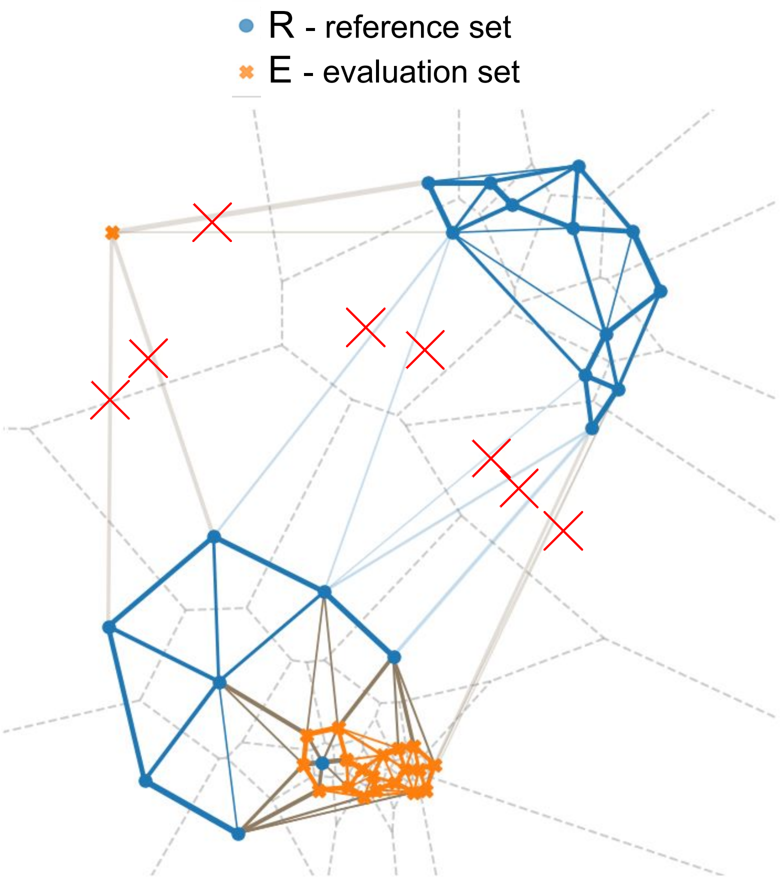

DCA framework

1. Approximate Delaunay graph

2. Distill Delaunay graph into connected components

3. Apply GeomCA scores on the obtained connected components
   - diversity of nodes
   - diversity of edges

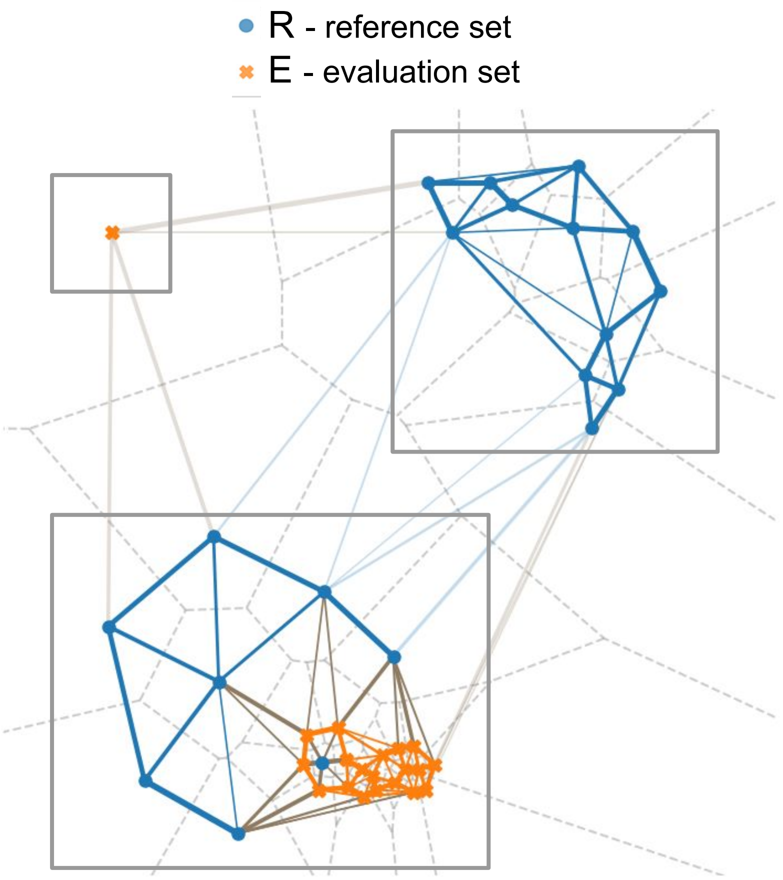

DCA framework

1. Approximate Delaunay graph

2. Distill Delaunay graph into connected components

3. Apply GeomCA scores on the obtained connected components
   - diversity of nodes
   - diversity of edges

4. Evaluate graph neighbourhood of the query representation

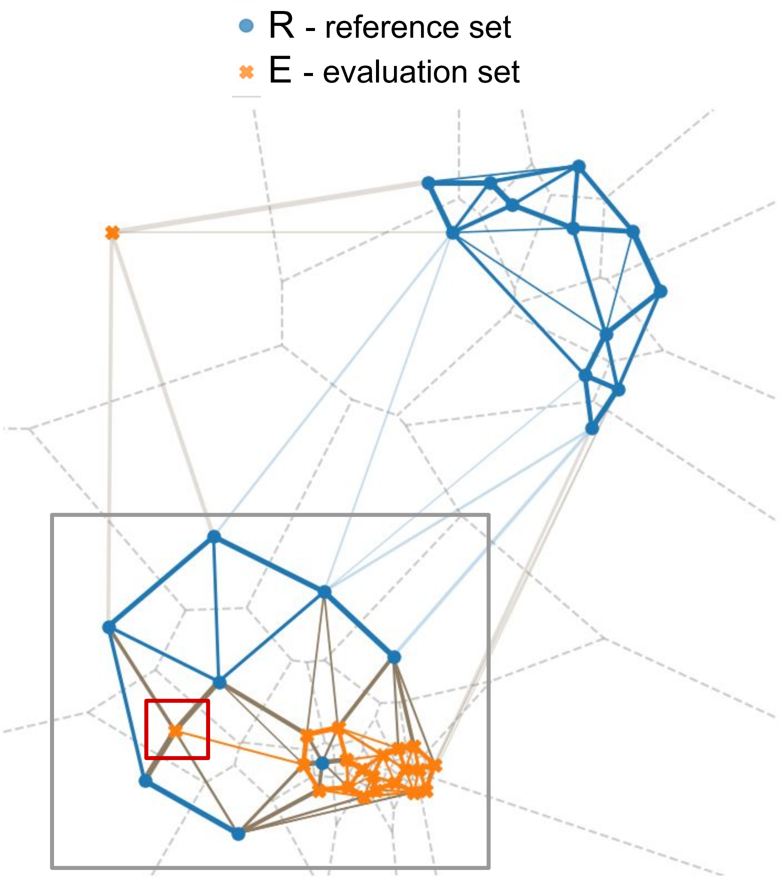

DCA: evaluation

DCA

- is more **stable and reliable** in presence of outliers and components of varying density compared to benchmark methods [4, 5, 6]
- detects peculiar geometric arrangements of points in high dimensional spaces

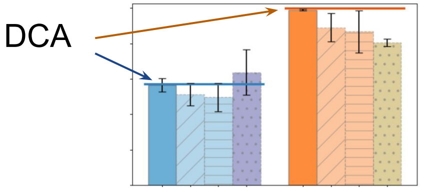

Query point DCA

- various use cases, e.g. evaluation of single GAN generated image, cluster assignments or semantic similarity in the representation space

[4] Poklukar et al. "GeomCA: Geometric evaluation of data representations," *2021 28th International Conference on Machine Learning (ICML)*.
[5] Kynkäänniemi et al. "Improved Precision and Recall Metric for Assessing Generative Models," *2019 Conference on Neural Information Processing Systems (NeurIPS)*.
[5] Khrulkov et al. "Geometry Score: {A} Method For Comparing Generative Adversarial Networks," *2018 28th International Conference on Machine Learning (ICML)*.

Work in progress
Geometric Multimodal Contrastive Learning

Miguel Vasco, Petra Poklukar

Plan: International Conference on Machine Learning 2022

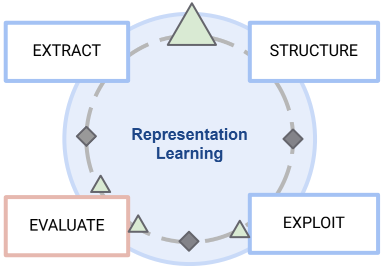

Goal: learn representations of complete and partial observations that are robust to missing modalities

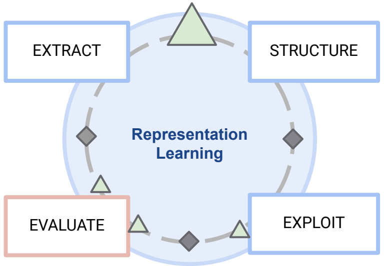

Geometric Multimodal Contrastive Learning

Goal: learn representations of complete and partial observations that are robust to missing modalities

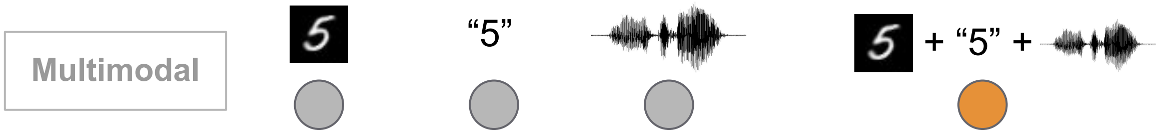

Geometric Multimodal Contrastive Learning

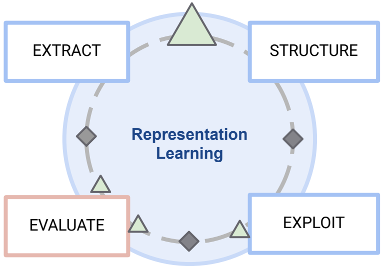

Goal: learn representations of complete and partial observations that are robust to missing modalities

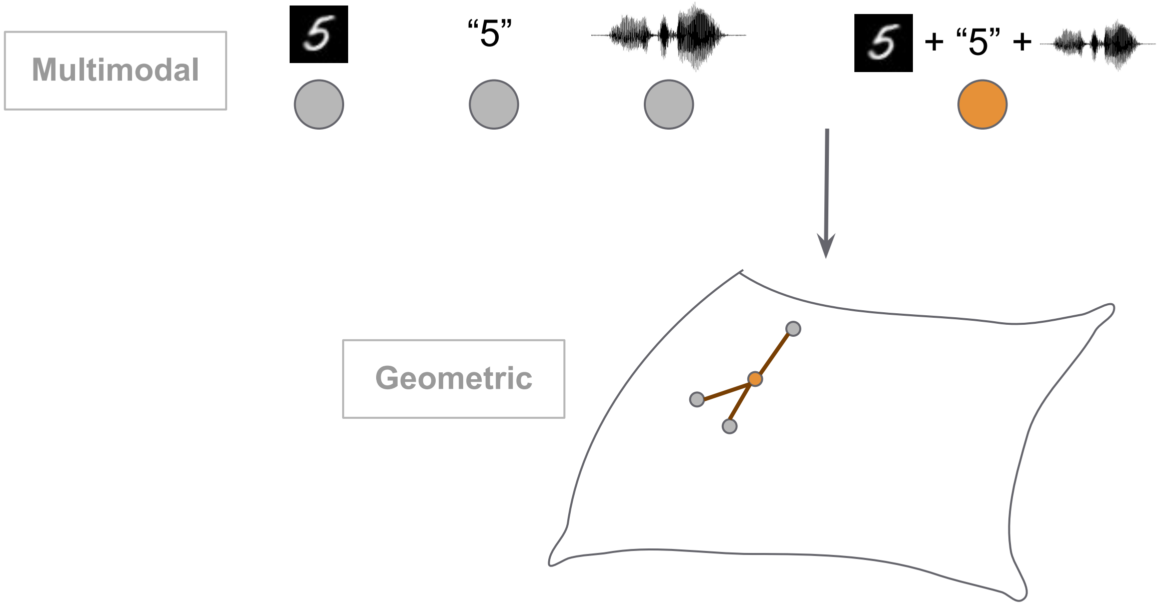

Geometric Multimodal Contrastive Learning

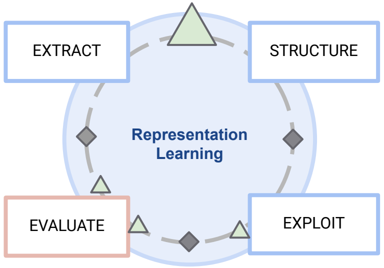

Goal: learn representations of complete and partial observations that are robust to missing modalities

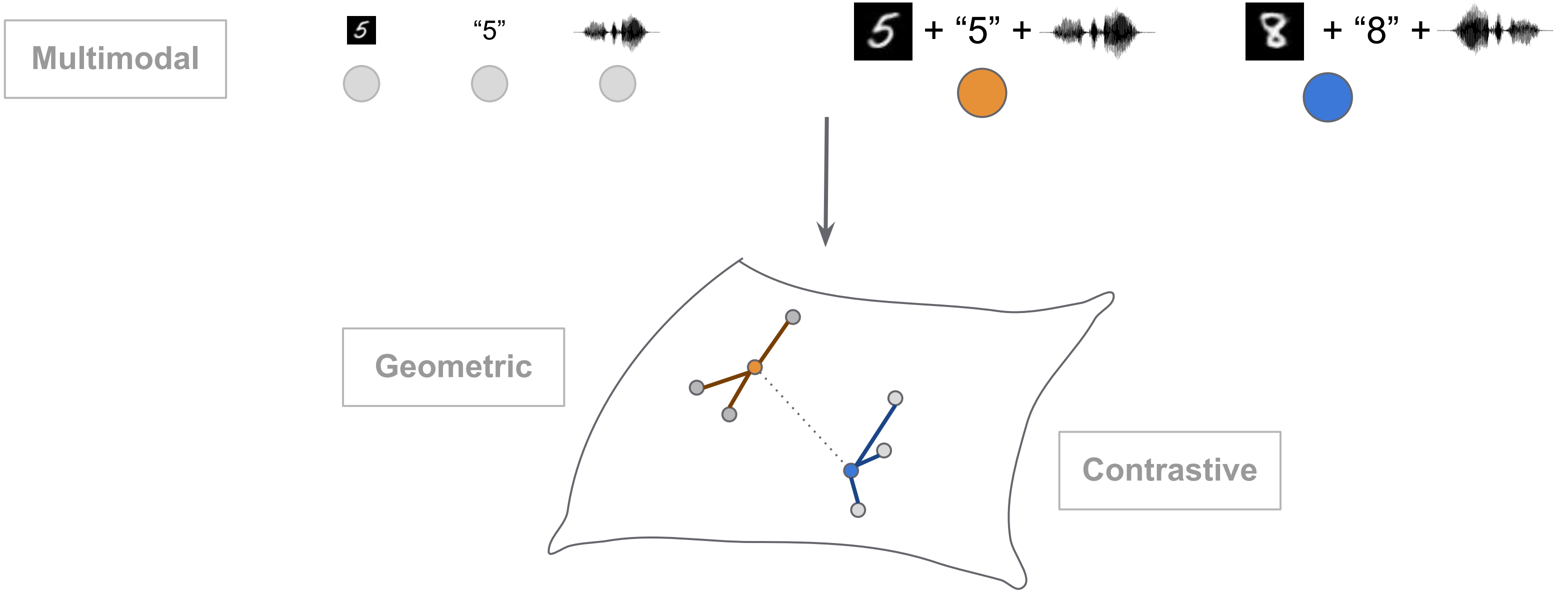

Work in Progress
Improved Training of Generative Models

Yonk Shi, Michael Welle, Petra Poklukar

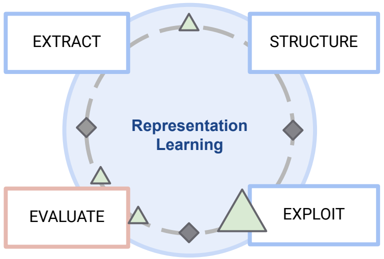

Goal: exploit geometric regularization to improve training of generative models

Improved Training of Generative Models

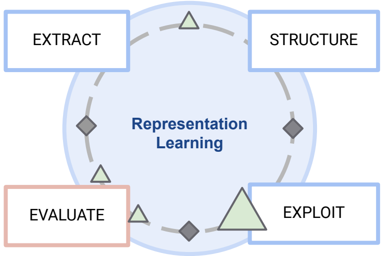

Goal: exploit geometric regularization to improve training of generative models

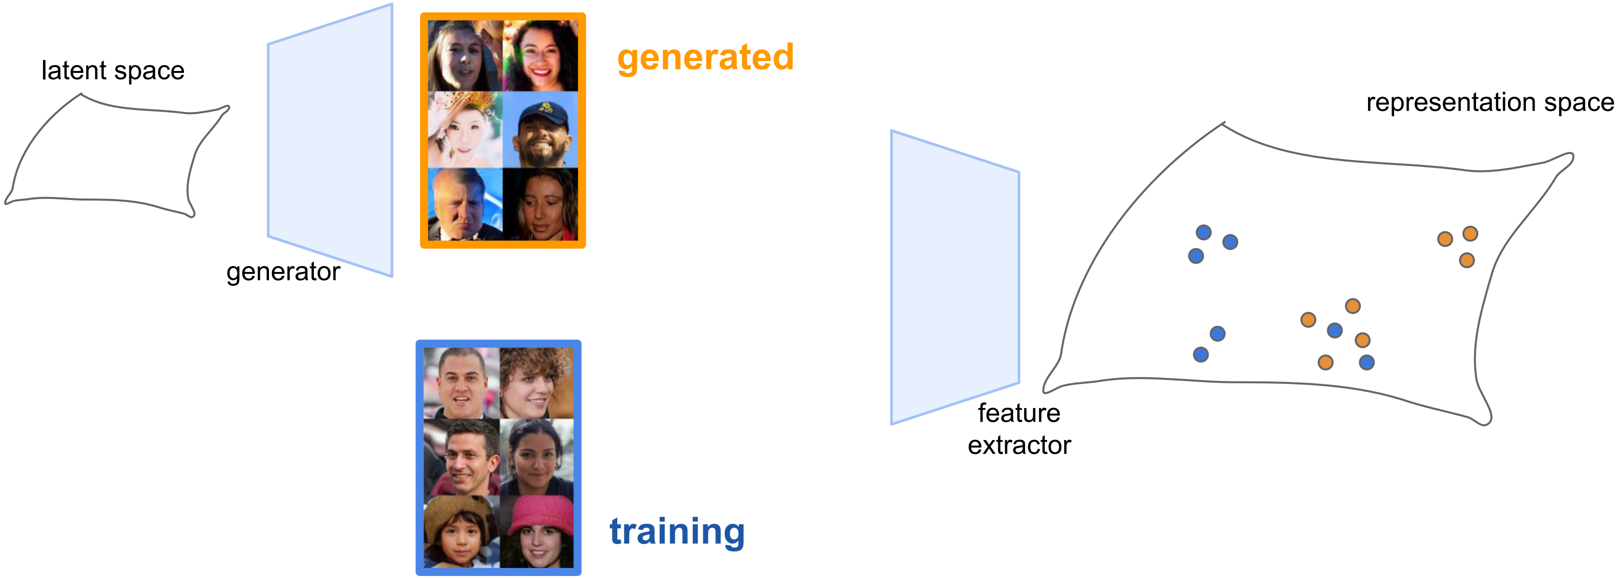

Image source: Kynkäänniemi et. al., *Improved Precision and Recall Metric for Assessing Generative Models* in Advances in Neural Information Processing Systems (NeurIPS) 2019

Improved Training of Generative Models

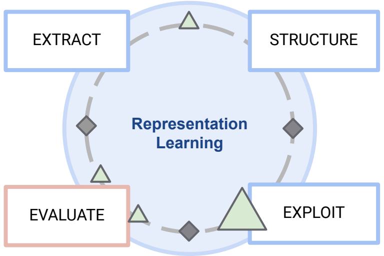

Goal: exploit geometric regularization to improve training of generative models

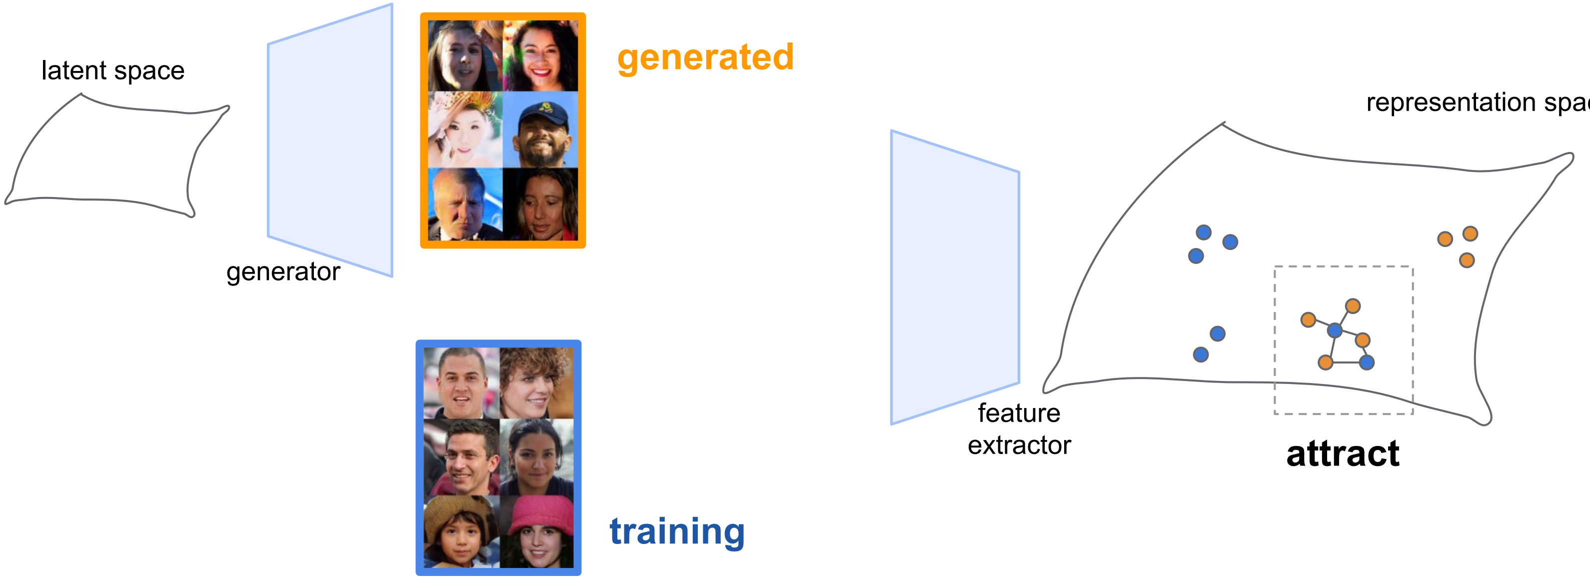

Image source: Kynkäänniemi et. al., *Improved Precision and Recall Metric for Assessing Generative Models* in Advances in Neural Information Processing Systems (NeurIPS) 2019

Improved Training of Generative Models

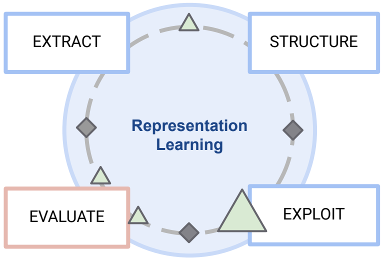

Goal: exploit geometric regularization to improve training of generative models

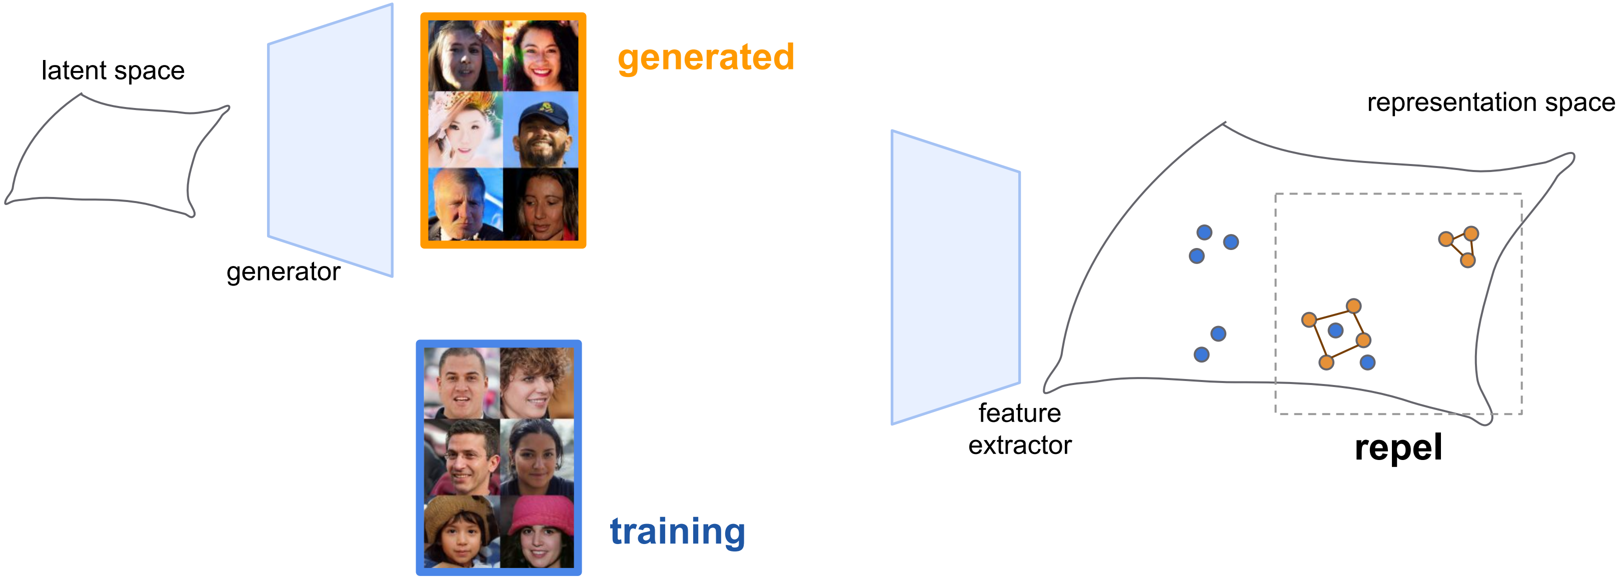

Applications

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

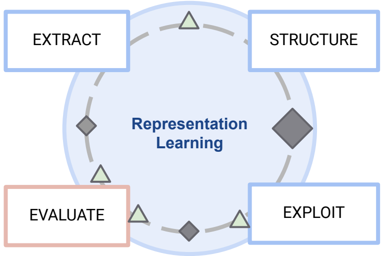

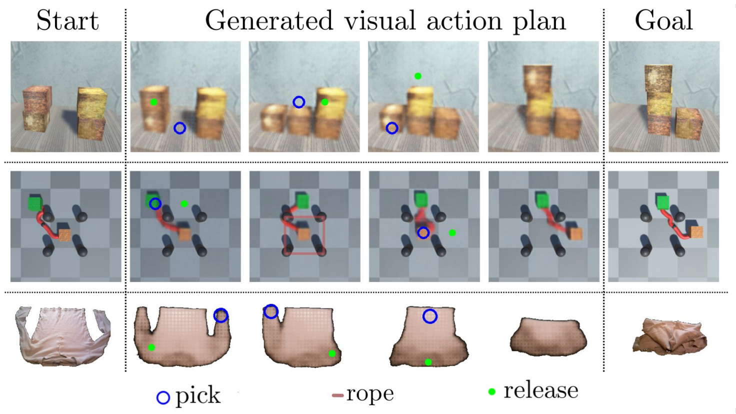

Goal: visual action planning of complex manipulation tasks with high-dimensional state spaces such as deformable objects

Image source: Lippi et al. *"Enabling Visual Action Planning for Object Manipulation through Latent Space Roadmap", arXiv preprint arXiv:2103.02554.*

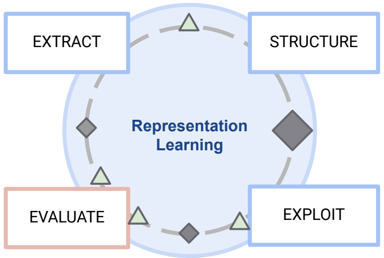

Enabling Visual Action Planning for Object Manipulation through Latent Space Roadmap

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

Conditionally accepted to IEEE Transactions of Robotics

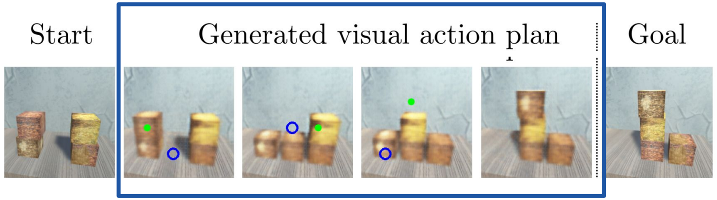

using representations learned by a VAE

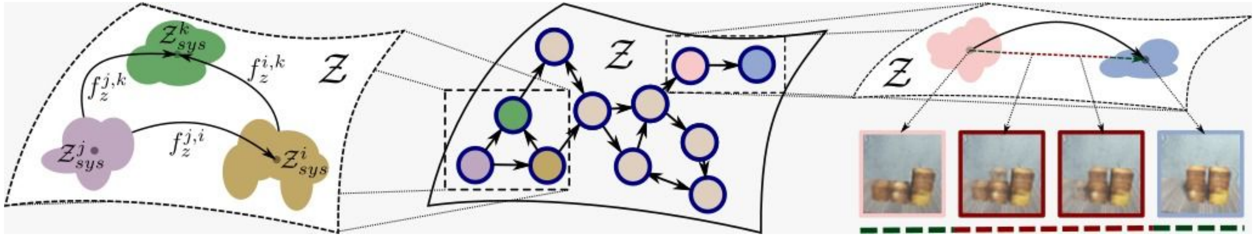

Image source: Lippi et al. *“Enabling Visual Action Planning for Object Manipulation through Latent Space Roadmap”, arXiv preprint arXiv:2103.02554.*

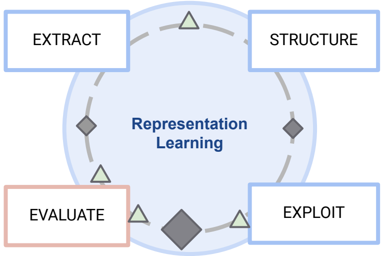

Data-Efficient Visuomotor Policy Training using Reinforcement Learning and Generative Models

Ali Ghadirzadeh*, Petra Poklukar*, Karol Arndt, Chelsea Finn, Ville Kyrki, Danica Kragic and Mårten Björkman

In revision for Journal of Machine Learning Research

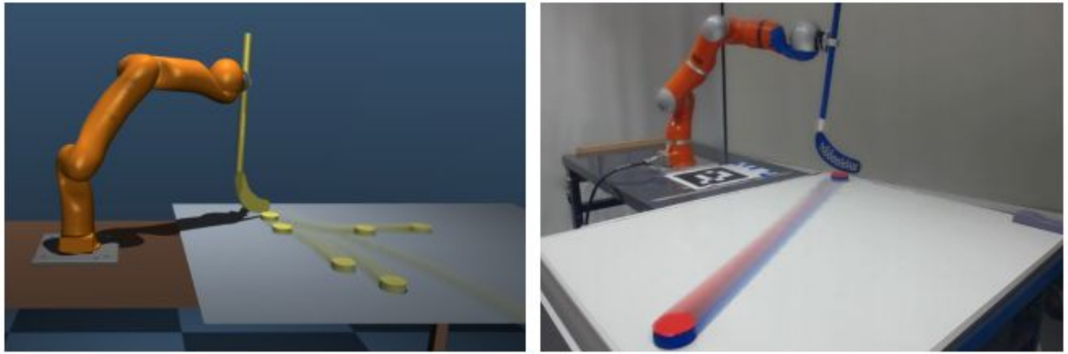

Arndt, Karol, et al. "Meta reinforcement learning for sim-to-real domain adaptation." IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2020.

Goal: use deep generative models to reduce the complexity of the problem

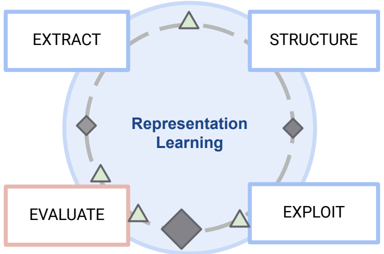

Data-Efficient Visuomotor Policy Training using Reinforcement Learning and Generative Models

Ali Ghadirzadeh*, Petra Poklukar*, Karol Arndt, Chelsea Finn, Ville Kyrki, Danica Kragic and Mårten Björkman

In revision for Journal of Machine Learning Research

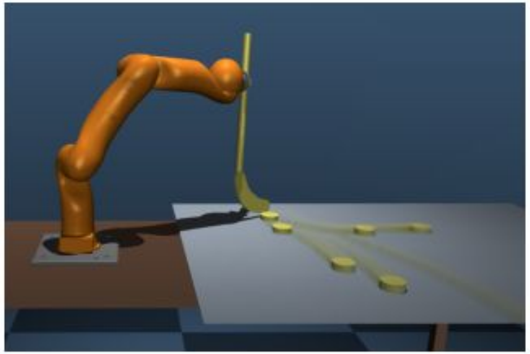

Arndt, Karol, et al. "Meta reinforcement learning for sim-to-real domain adaptation." IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2020.

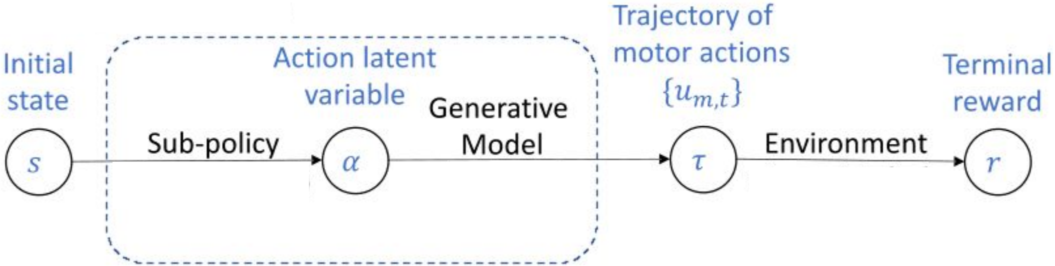

Relate the policy performance with characteristics of representations learned by deep generative models

Work in progress
Evaluation of Graph Representations and Graph Generative Models

Petra Poklukar, Ciwan Ceylan, Hanna Hultin

Plan ICML 2022

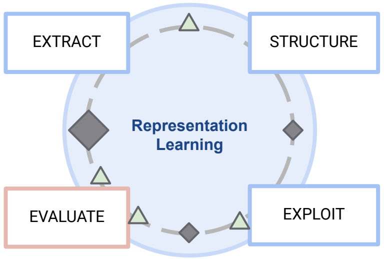

Goal: develop an evaluation procedure for assessing the quality of graph representations and graph generation models (in collaboration with SEB)

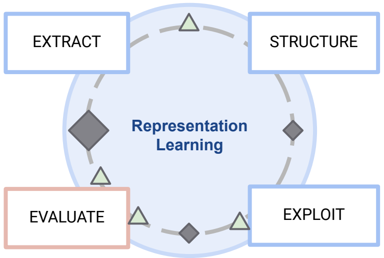

Evaluation of Graph Representations and Graph Generative Models

Goal: develop an evaluation procedure for assessing the quality of graph representations and graph generation models (in collaboration with SEB)

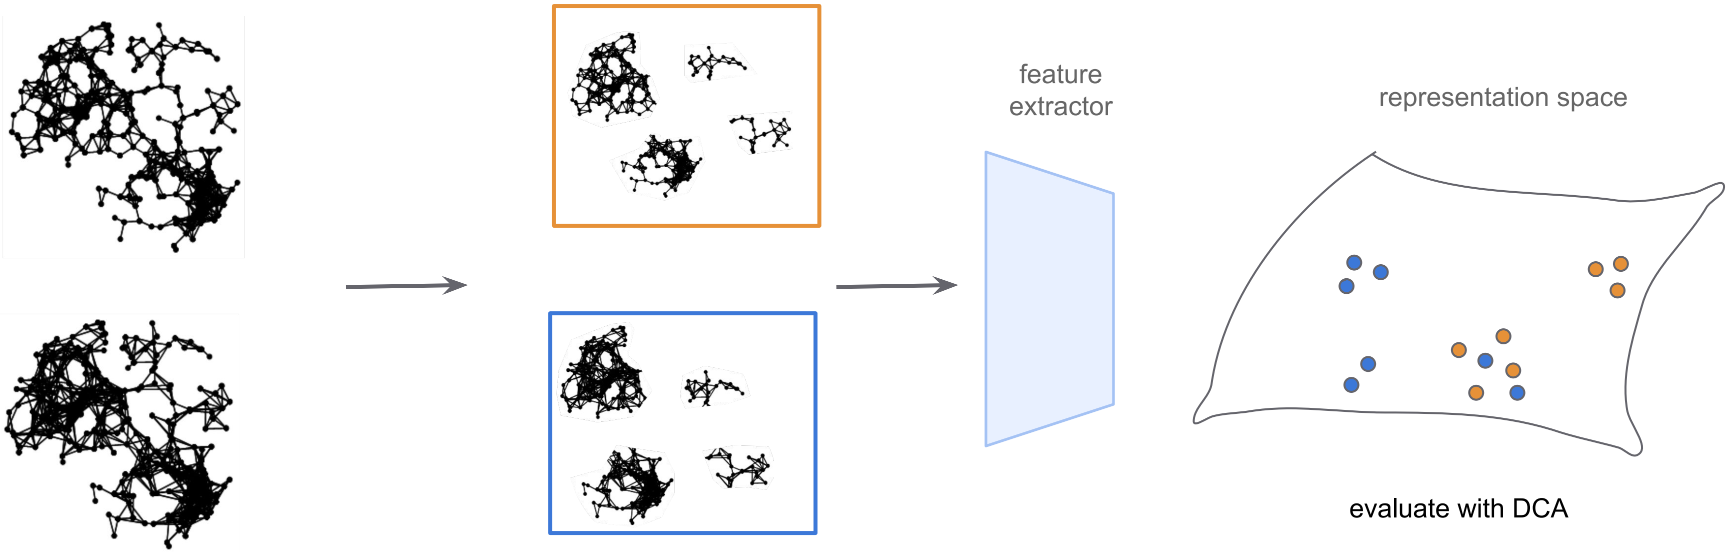

Image source: Rendsburg et al. "*NetGAN without GAN: From Random Walks to Low-Rank Approximations*", International Conference on Machine Learning (ICML) 2020.

Evaluation of Graph Representations and Graph Generative Models

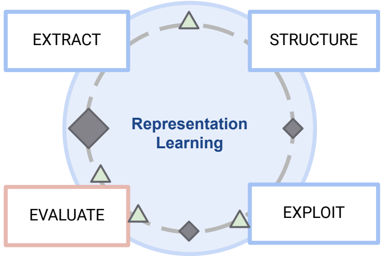

Goal: develop an evaluation procedure for assessing the quality of graph representations and graph generation models (in collaboration with SEB)

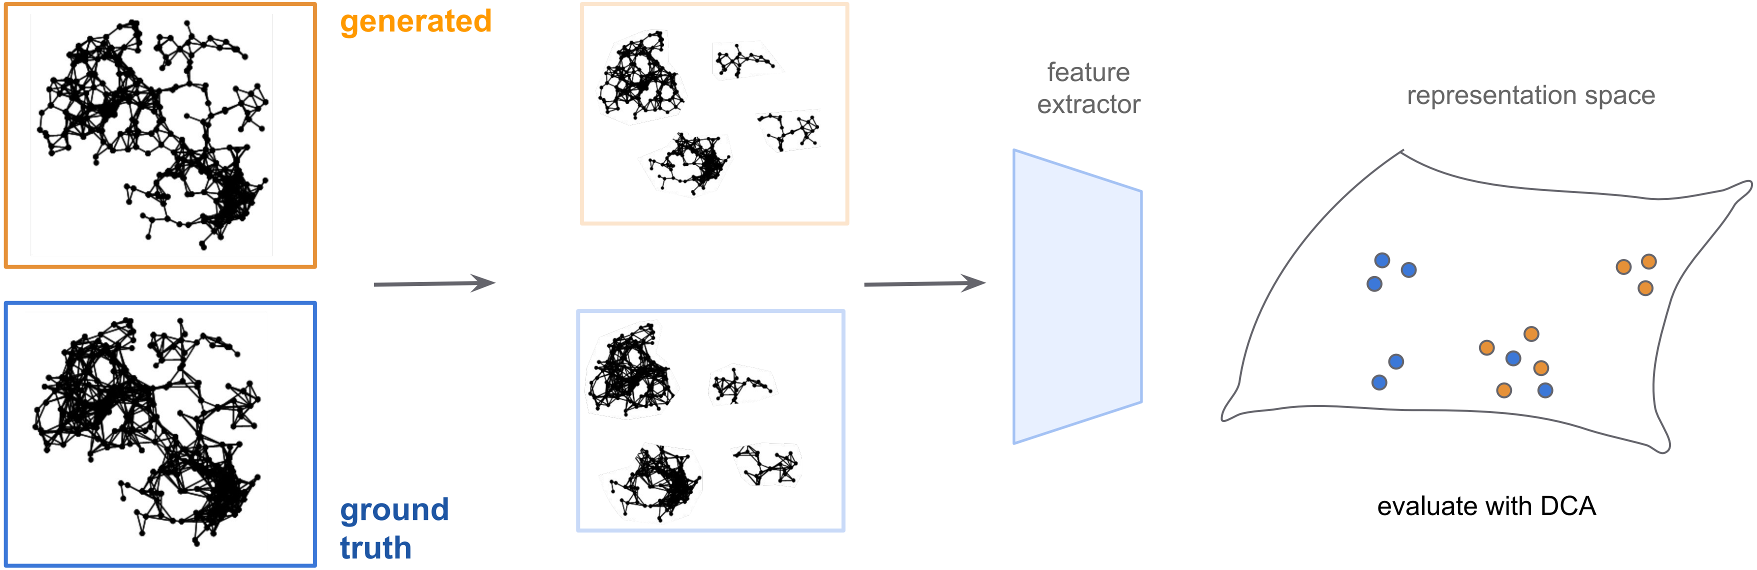

Image source: Rendsburg et al. "*NetGAN without GAN: From Random Walks to Low-Rank Approximations*", International Conference on Machine Learning (ICML) 2020.

Learning and Evaluating Representations: Geometric and Topological Aspects

80% seminar

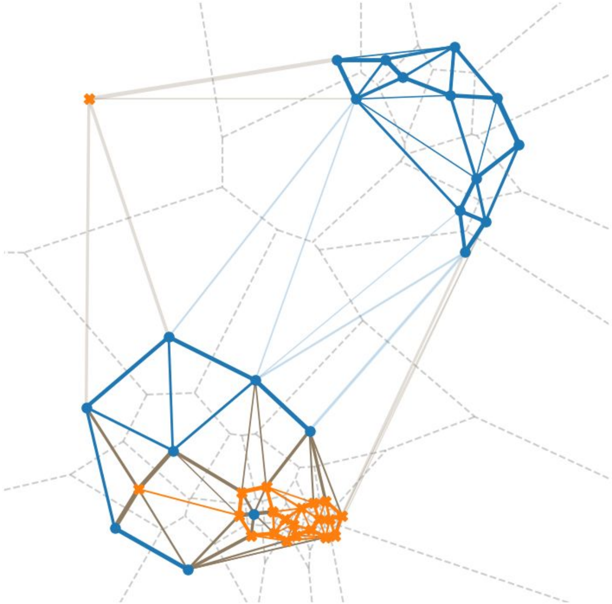

Petra Poklukar
poklukar@kth.se
https://people.kth.se/~poklukar/