Learning and Evaluating Data Representations

Petra Poklukar

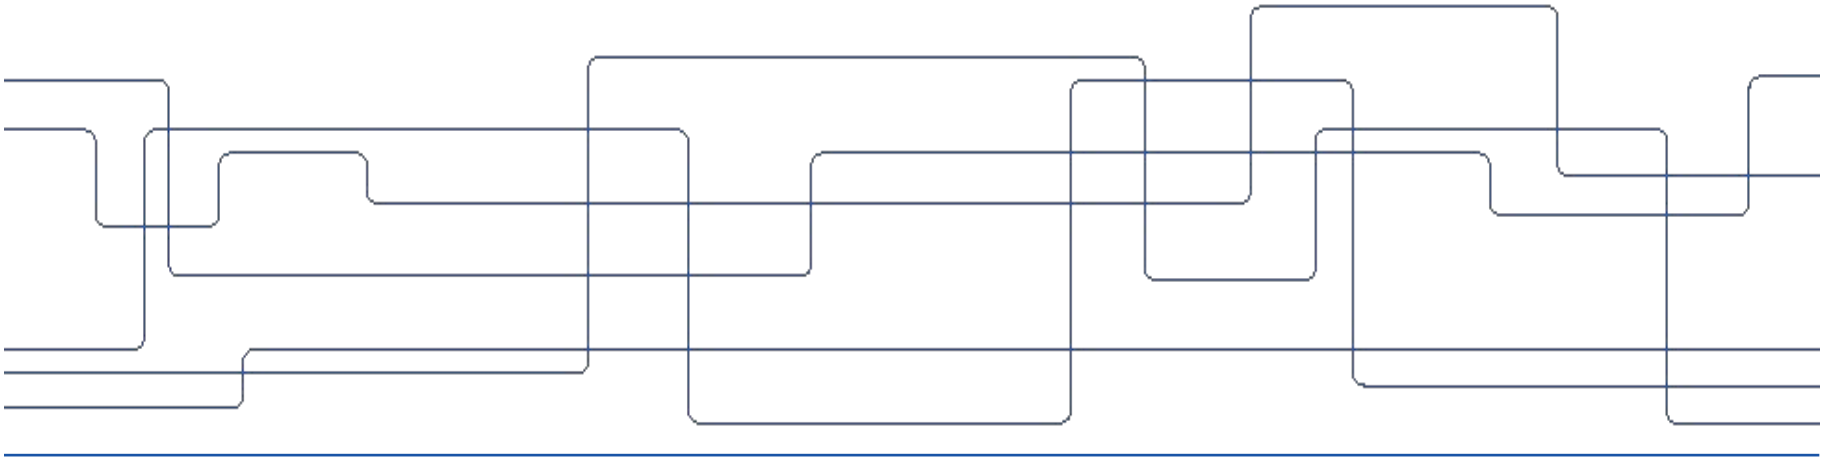

Outline

- About me

- Representation Learning:
  - What and why
  - How
    - *example from robotics*
    - *general*
  - Evaluation
    - *example from generative modelling*

Briefly about myself

- Bachelor's & Masters in Mathematics at University of Ljubljana [2011-2014 & 2014-2016]
  - Topology, Geometry and Programming

- Work experience:
  - approximately 1 year as a PhD in Mathematics at Stockholm University [2016-2018]
  - Data Scientist at King [2018]
  - PhD in Machine Learning at KTH Royal Institute of Technology, Stockholm [2018-2022/2023]

Representation Learning: what and why

Example: computer vision *before* deep learning

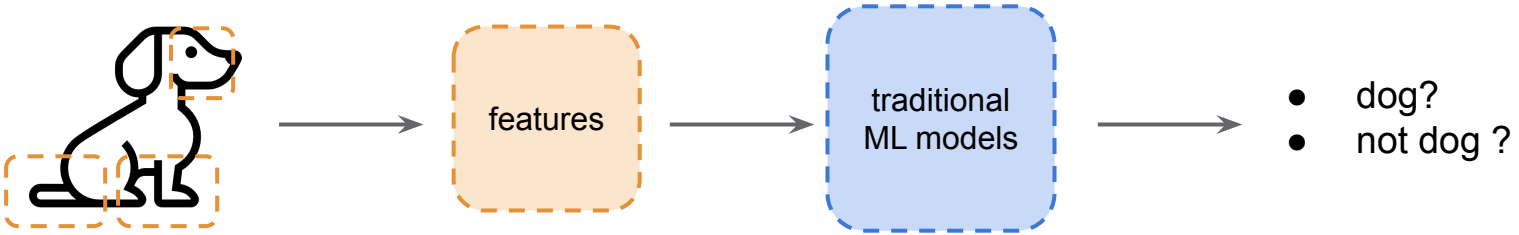

Representation Learning: what and why

Example: computer vision *after* deep learning

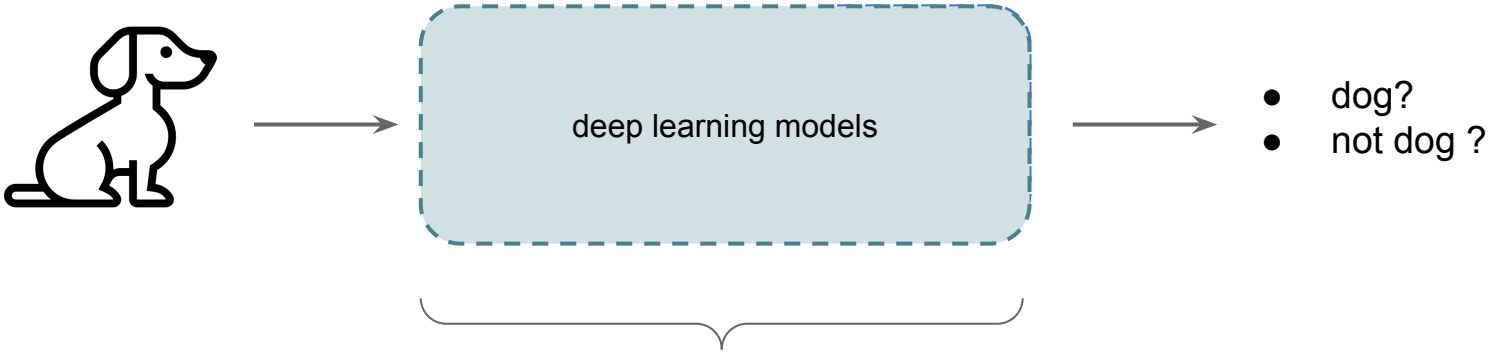

How to learn data representations that make it easier to extract useful information for such classifier?

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*

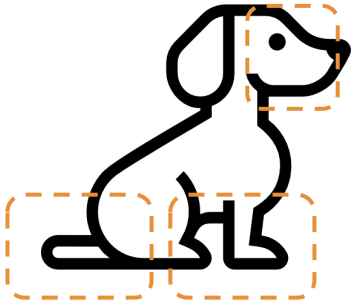

Why:

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

Representation Learning: how

Example from robotics:

Latent Space Roadmap for Visual Action Planning of Deformable and Rigid Object Manipulation,

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

International Conference on Intelligent Robots and Systems (IROS) 2020

Latent Space Roadmap for Visual Action Planning of Deformable and Rigid Object Manipulation

Given: start and goal state
Goal: visual action planning

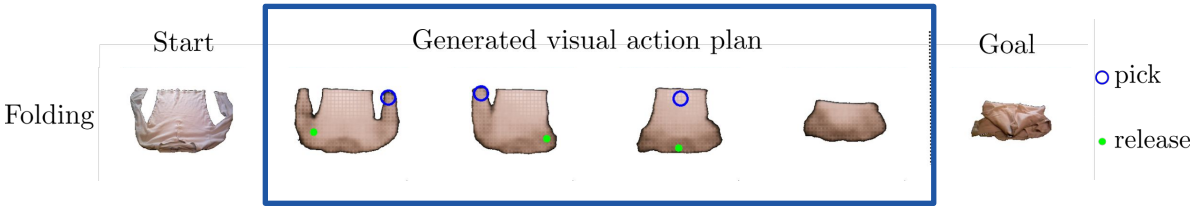

Problem: image space
Solution: using learned low-dimensional representations

What is a good representation for visual action planning?

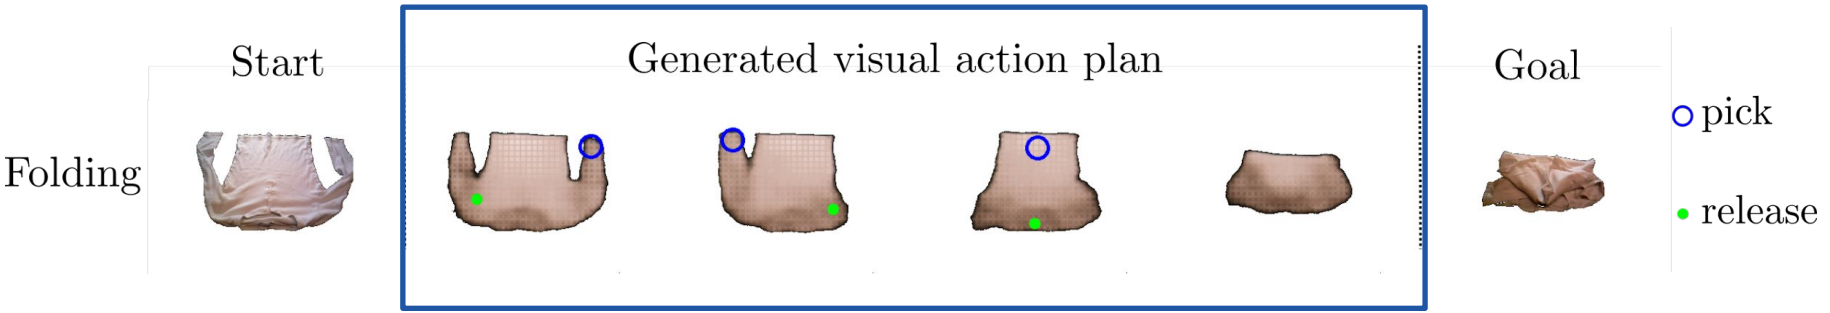

- low-dimensional
- extracts features representing each state
- cluster the extracted features

How to learn a good representation for visual action planning?

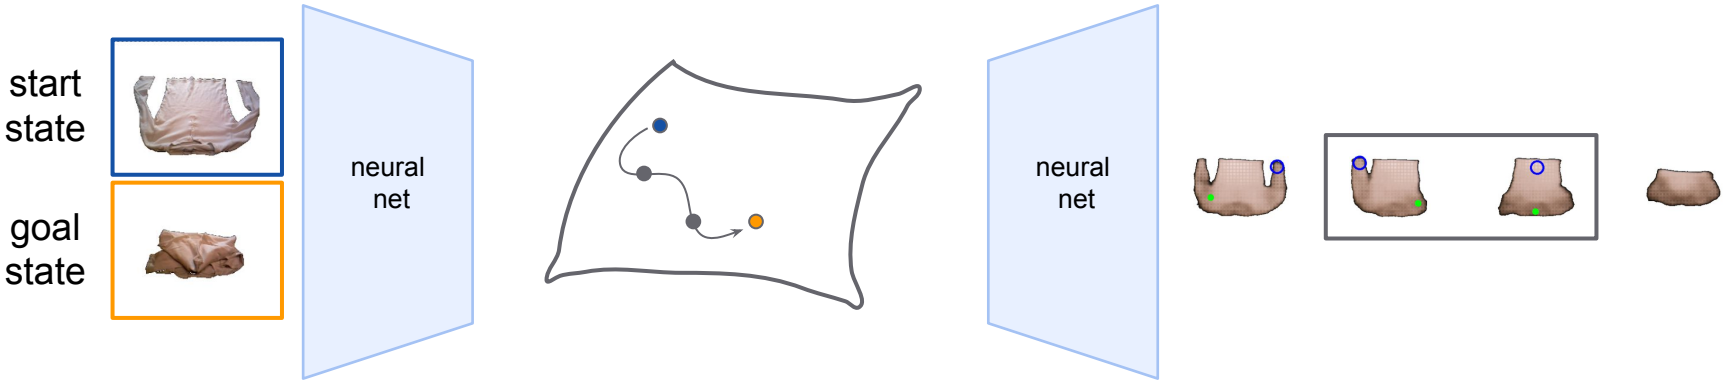

- extracted features represent state
- similar extracted features are clustered together

How to learn a good representation for visual action planning?

start state

goal state

neural net

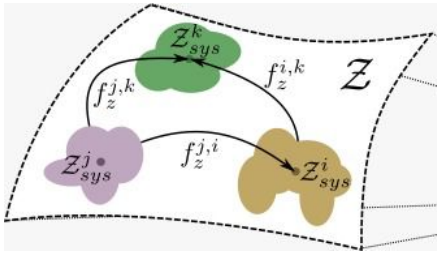

neural net

contrastive loss:

- encode similar states together

- encode dissimilar states far apart

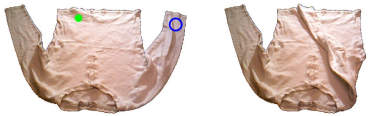

Representation Learning: how more generally

More examples:

- Contrastive learning [1] (previous example)

[1] Le-Khac et. al., *Contrastive Representation Learning: A Framework and Review* in IEEE Access, vol. 8, pp. 193907-193934, 2020.

Representation Learning: how more generally

More examples:

- Contrastive learning [1] (previous example)
- Learning disentangled representations [2]

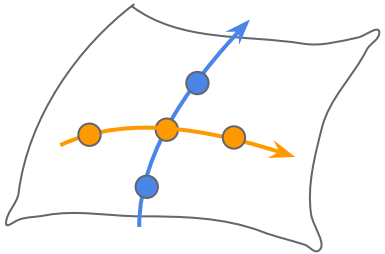

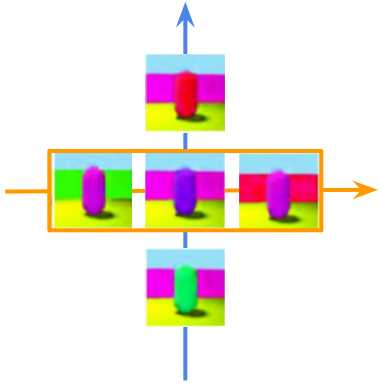

[1] Le-Khac et. al., *Contrastive Representation Learning: A Framework and Review* in IEEE Access, vol. 8, pp. 193907-193934, 2020.
[2] Locatello et al., *A Sober Look at the Unsupervised Learning of Disentangled Representations and their Evaluation*, Journal of Machine Learning Research (JMLR) 2020.

Representation Learning: how more generally

More examples:

- Contrastive learning [1] (previous example)
- Learning disentangled representations [2]
- Manifold learning
  - Isomap [3]
  - t-distributed Stochastic Neighbor Embedding (t-SNE) [4]
  - UMAP [5]
  - ….

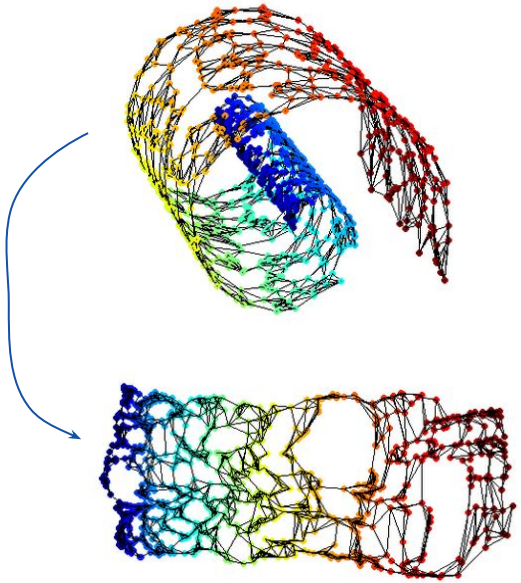

[1] Le-Khac et. al., *Contrastive Representation Learning: A Framework and Review* in IEEE Access, vol. 8, pp. 193907-193934, 2020
[2] Locatello et al., *A Sober Look at the Unsupervised Learning of Disentangled Representations and their Evaluation*, Journal of Machine Learning Research (JMLR) 2020
[3] Tenenbaum et. al. *A Global Geometric Framework for Nonlinear Dimensionality Reduction [http://wearables.cc.gatech.edu/paper_of_week/isomap.pdf]*
[4] McInnes et. al. *Umap: Uniform manifold approximation and projection for dimension reduction*, *arXiv preprint arXiv:1802.03426* (2018).
[5] Van der Maaten and Hinton, *Visualizing data using t-SNE*, *Journal of Machine Learning Research* (JMLR) 2008.

Representation Learning: how more generally

More examples:

- Contrastive learning [1] (previous example)
- Learning disentangled representations [2]
- Manifold learning [3, 4, 5]
- Self-supervised representation learning
  - Solving jigsaw puzzles [6]
  - Predicting rotations [7]
  - …

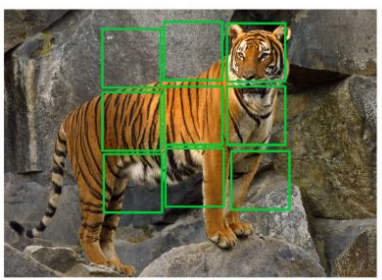

[1] Le-Khac et. al., *Contrastive Representation Learning: A Framework and Review* in IEEE Access, vol. 8, pp. 193907-193934, 2020
[2] Locatello et al., *A Sober Look at the Unsupervised Learning of Disentangled Representations and their Evaluation*, Journal of Machine Learning Research (JMLR) 2020
[3] Tenenbaum et. al. *A Global Geometric Framework for Nonlinear Dimensionality Reduction [http://wearables.cc.gatech.edu/paper_of_week/isomap.pdf]*
[4] McInnes et. al. *Umap: Uniform manifold approximation and projection for dimension reduction*, *arXiv preprint arXiv:1802.03426* (2018).
[5] Van der Maaten and Hinton, *Visualizing data using t-SNE*, *Journal of Machine Learning Research* (JMLR) 2008.
[6] Noroozi et. at. *Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles,* European Conference on Computer Vision (ECCV) 2016.
[7] Gidaris et. al. *Unsupervised Representation Learning by Predicting Image Rotations*, International Conference on Learning Representations (ICLR) 2018.

Representation Learning: how more generally

More examples:

- Contrastive learning [1] (previous example)
- Learning disentangled representations [2]
- Manifold learning [3, 4, 5]
- Self-supervised representation learning [6, 7]
- etc ...

How can we use data representations?

[1] Le-Khac et. al., *Contrastive Representation Learning: A Framework and Review* in IEEE Access, vol. 8, pp. 193907-193934, 2020
[2] Locatello et al., *A Sober Look at the Unsupervised Learning of Disentangled Representations and their Evaluation*, Journal of Machine Learning Research (JMLR) 2020
[3] Tenenbaum et. al. *A Global Geometric Framework for Nonlinear Dimensionality Reduction [http://wearables.cc.gatech.edu/paper_of_week/isomap.pdf]*
[4] McInnes et. al. *Umap: Uniform manifold approximation and projection for dimension reduction*, *arXiv preprint arXiv:1802.03426* (2018).
[5] Van der Maaten and Hinton, *Visualizing data using t-SNE*, *Journal of Machine Learning Research* (JMLR) 2008.
[6] Noroozi et. at. *Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles,* European Conference on Computer Vision (ECCV) 2016.
[7] Gidaris et. al. *Unsupervised Representation Learning by Predicting Image Rotations*, International Conference on Learning Representations (ICLR) 2018.

Using learned data representations

- Semi-supervised learning [8]
  - Train a small classifier using few labeled data on extracted representations
  - Evaluate on a test set
- Transfer learning [8]
  - Extract representations of a new dataset and train a classifier
  - Evaluate on a test set
- Density estimation [9]
  - Extract representations of a dataset and train a density estimation model
- Specific applications: robotics, biology, etc

[8] Chen et.al., *A Simple Framework for Contrastive Learning of Visual Representations*, International Conference on Machine Learning (ICML) 2020
[9] Kirichenko et. al., *Why Normalizing Flows Fail to Detect Out-of-Distribution Data,* Neural Information Processing Systems (NeurIPS) 2020

How do we know what is a in fact good representation?

- Low dimensional
- Captures similarities
- View invariant
- Disentangled
- Reflects input manifold

"The one that makes the subsequent learning tasks easier."

What if we cannot define an evaluation task? For example:

- No labeled data
- Too many potential evaluation tasks

What if we cannot define an evaluation task?

Example: generative models

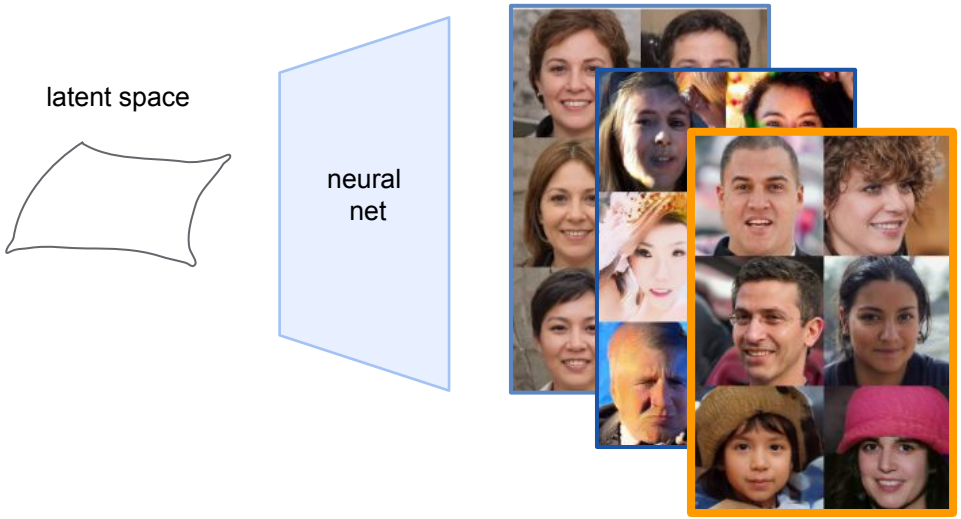

How to evaluate the quality of generated images?

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

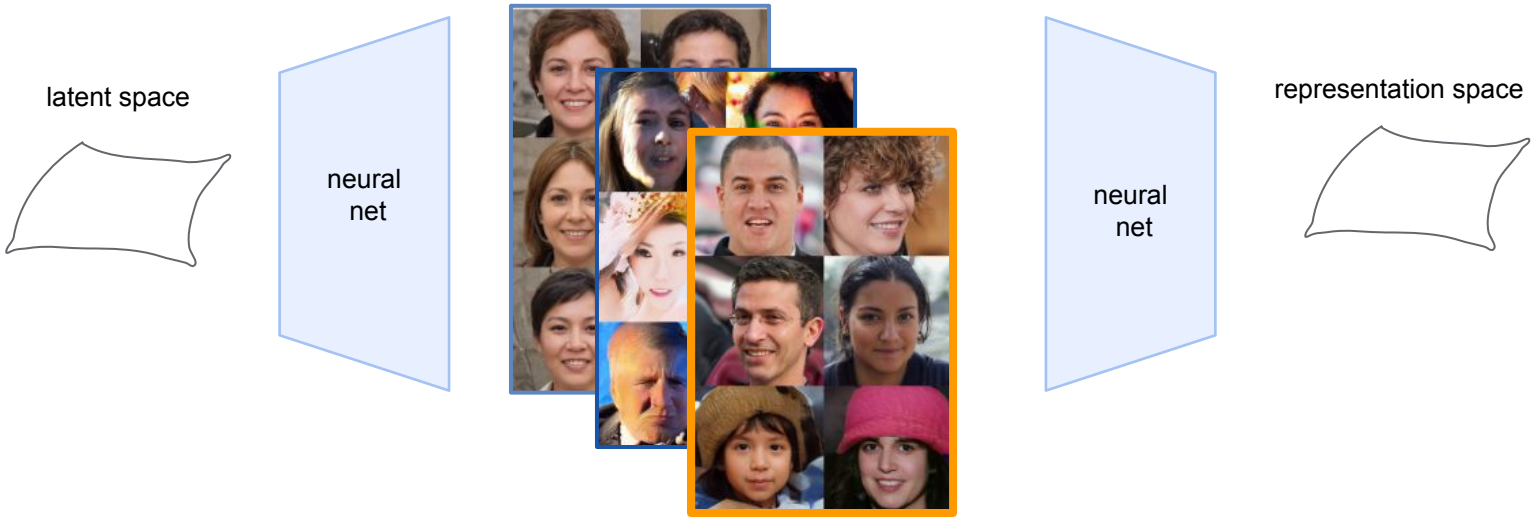

Image source: Kynkäänniemi et. al., *Improved Precision and Recall Metric for Assessing Generative Models* in Advances in Neural Information Processing Systems (NeurIPS) 2019

Evaluation of data representations

GeomCA: Geometric Evaluation of Data Representations

Petra Poklukar, Anastasiia Varava and Danica Kragic

International Conference on Machine Learning (ICML) 2021

Example: representations of **generated** images and **training** images

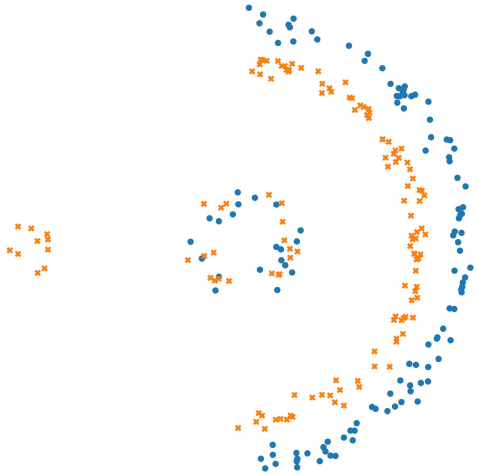

Geometric Component Analysis (GeomCA)

Example: Representations of **generated** images and **training** images

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

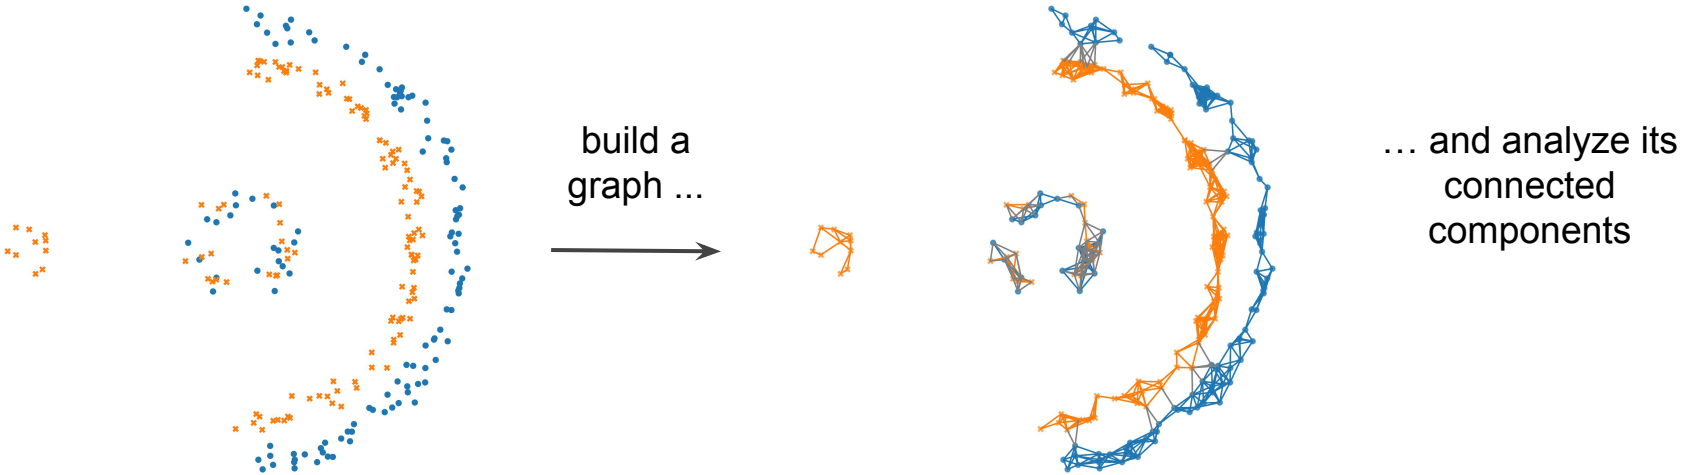

What if we cannot define an evaluation task?

Example: generative models

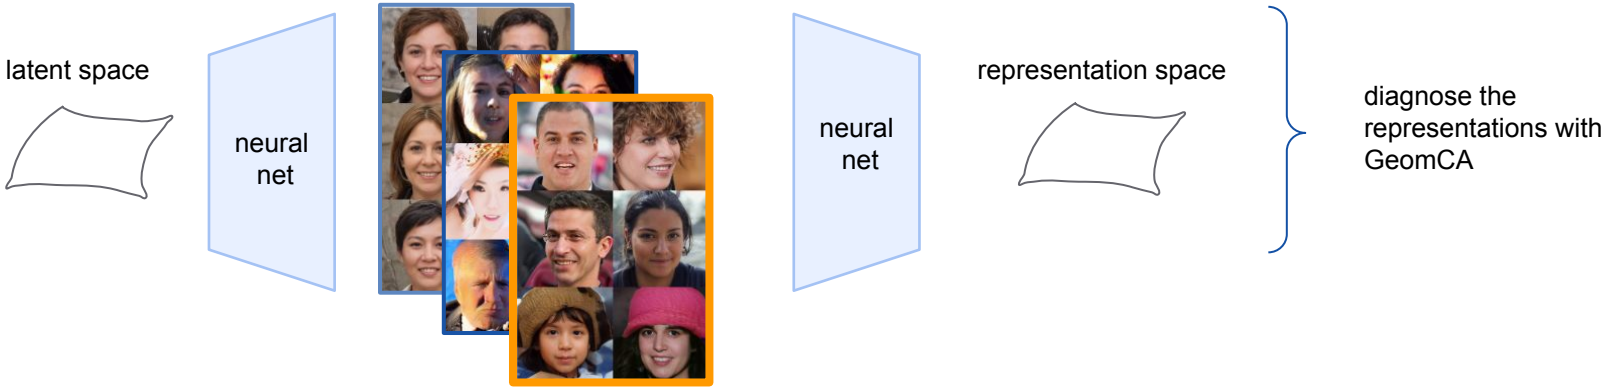

Image source: Kynkäänniemi et. al., *Improved Precision and Recall Metric for Assessing Generative Models* in Advances in Neural Information Processing Systems (NeurIPS) 2019

Conclusion

- Useful data representations are perhaps a step towards general AI
- There are many ways to learn data representations
  - example from robotics: contrastive learning
  - many more: disentangled representations, manifold learning, self-supervised representation learning, …
- There are many ways to use learned data representations
  - robotics, semi-supervised learning, transfer learning, density estimation, …
- It is not always trivial to evaluate the quality of learned data representations
  - GeomCA

Learning and Evaluating Data Representations

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