Improving Representations of Ice Crystals in Atmospheric Models: Insights from *in situ* Imagery and Machine Learning

Joseph Ko, Columbia University
ELLIIT Focus Period Linköping 2024
Machine Learning for Climate Science
October 14, 2024

Outline

- Motivation
- Part 1: Predicting 3-D properties from images (supervised)
- Part 2: Latent representations of ice crystals (unsupervised)
- Part 3: Latent metrics for LES cirrus simulations (unsupervised)
- Conclusion

Outline

- Motivation
- Part 1: Predicting 3-D properties from images (supervised)

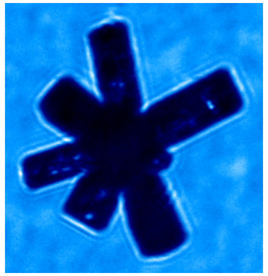

- Part 2: Latent representations of ice crystals (unsupervised)
- Part 3: Latent metrics for LES cirrus simulations (unsupervised)
- Conclusion

Outline

- Motivation
- Part 1: Predicting 3-D properties from images (supervised)
- Part 2: Latent representations of ice crystals (unsupervised)

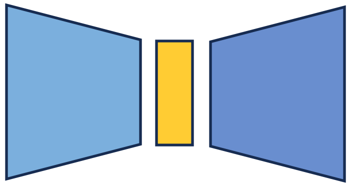

- Part 3: Latent metrics for LES cirrus simulations (unsupervised)
- Conclusion

Outline

- Motivation
- Part 1: Predicting 3-D properties from images (supervised)
- Part 2: Latent representations of ice crystals (unsupervised)
- Part 3: Latent metrics for LES cirrus simulations (unsupervised)

- Conclusion

Outline

- Motivation
- Part 1: Predicting 3-D properties from images (supervised)
- Part 2: Latent representations of ice crystals (unsupervised)
- Part 3: Latent metrics for LES cirrus simulations (unsupervised)
- Conclusion

Clouds strongly impact the climate

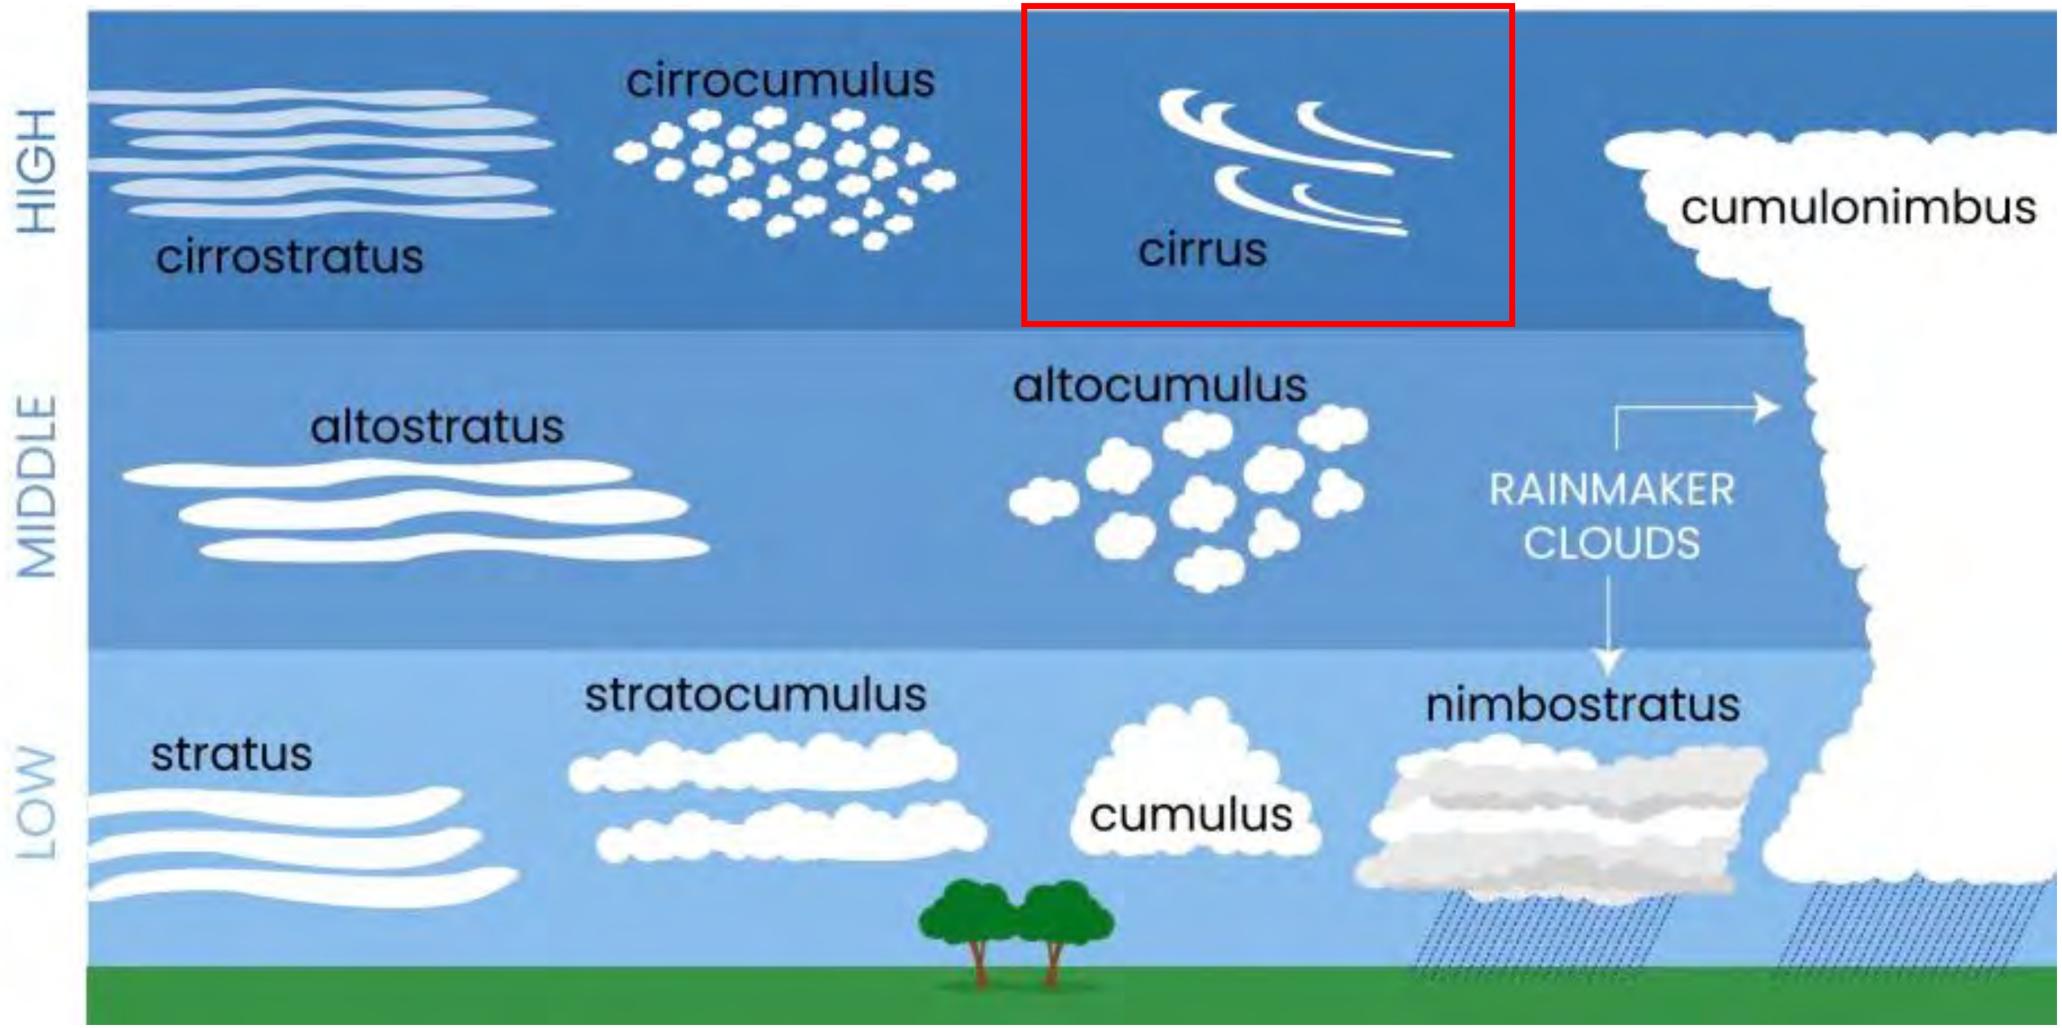

Source: UCAR

Clouds impact Earth's energy balance and hydrologic cycle

Ice clouds are poorly understood

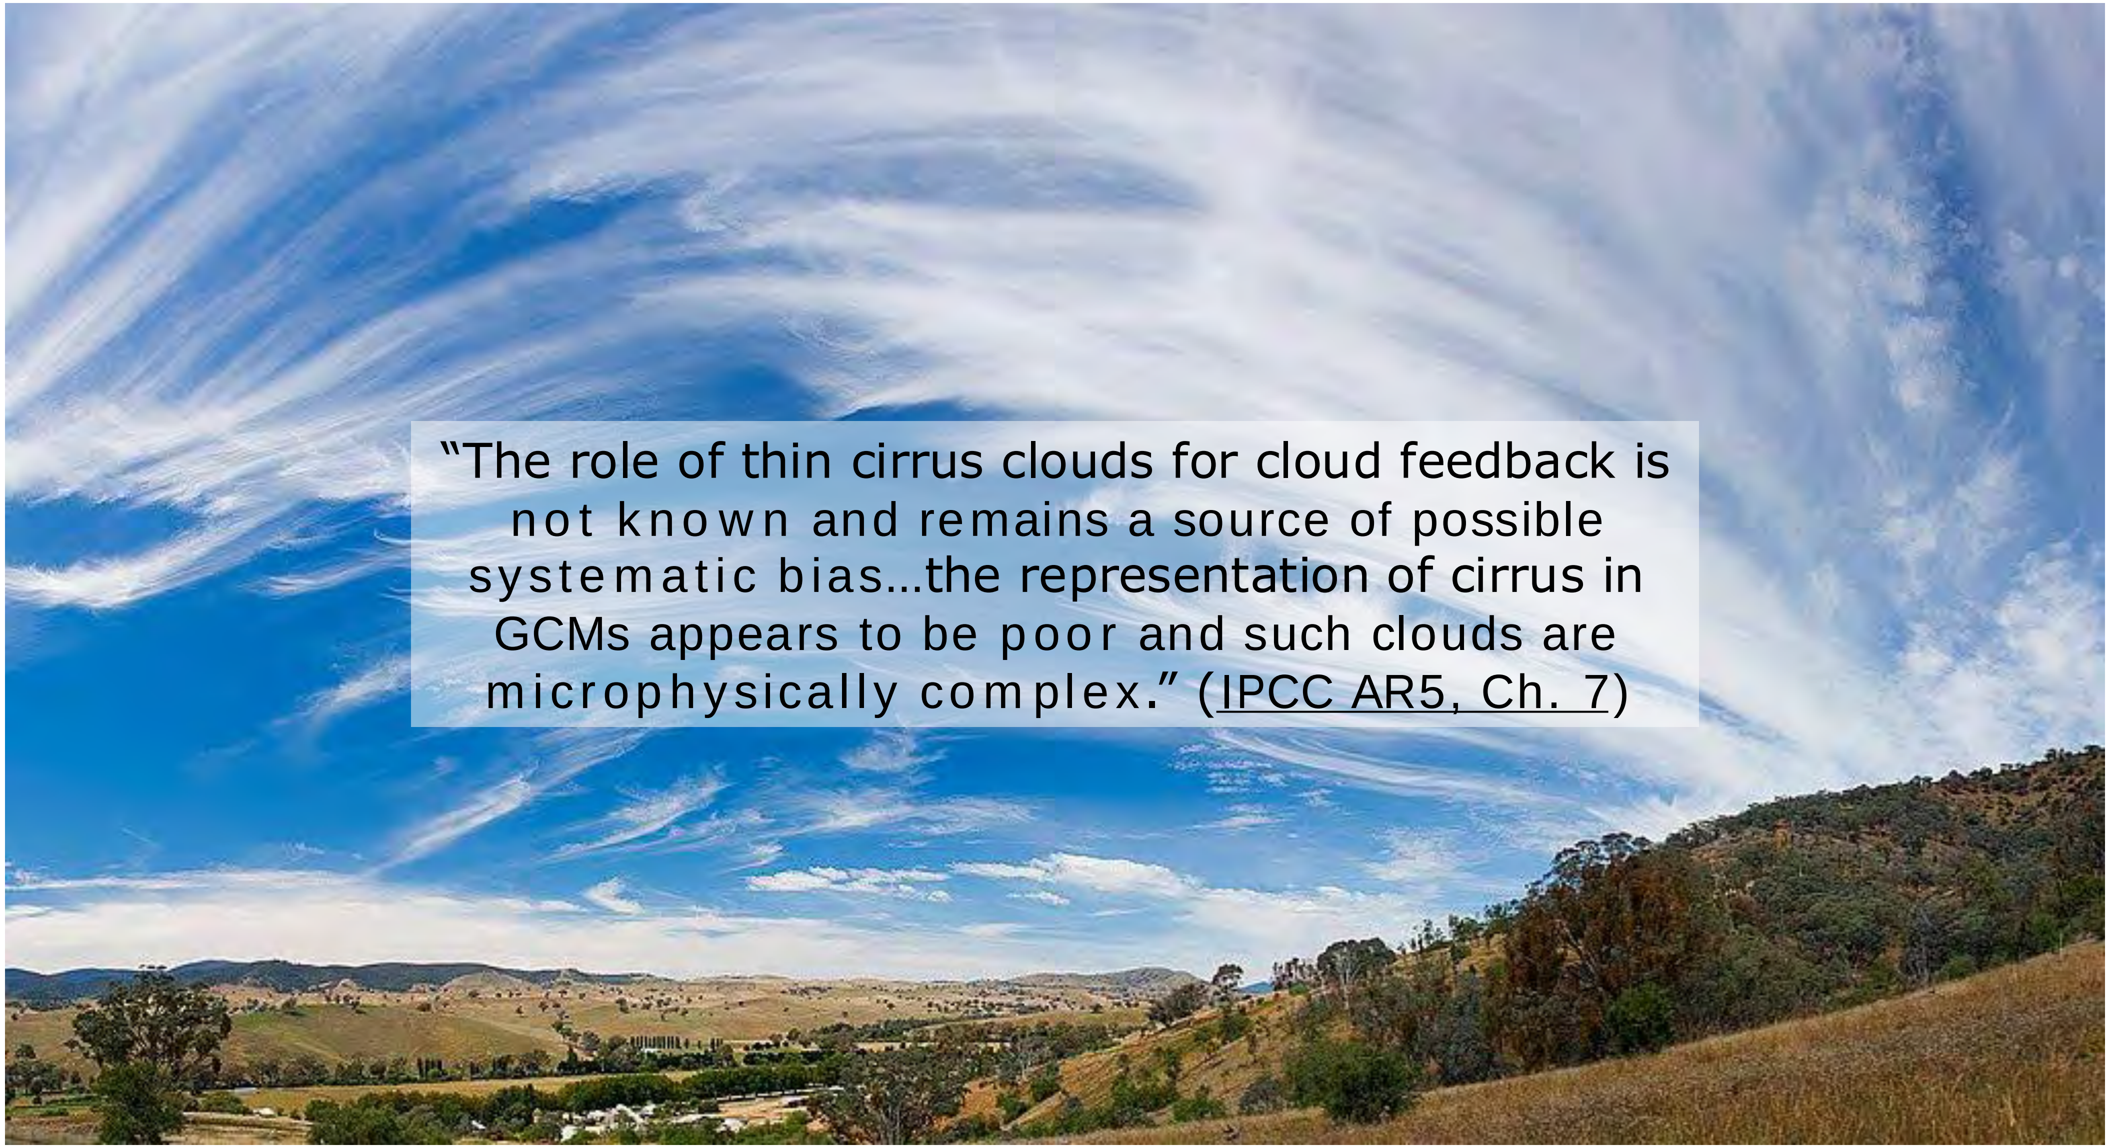

"The role of thin cirrus clouds for cloud feedback is not known and remains a source of possible systematic bias…the representation of cirrus in GCMs appears to be poor and such clouds are microphysically complex." (IPCC AR5, Ch. 7)

Source: Fir0002/Flagstaffotos

Ice habit (i.e., shape) matters

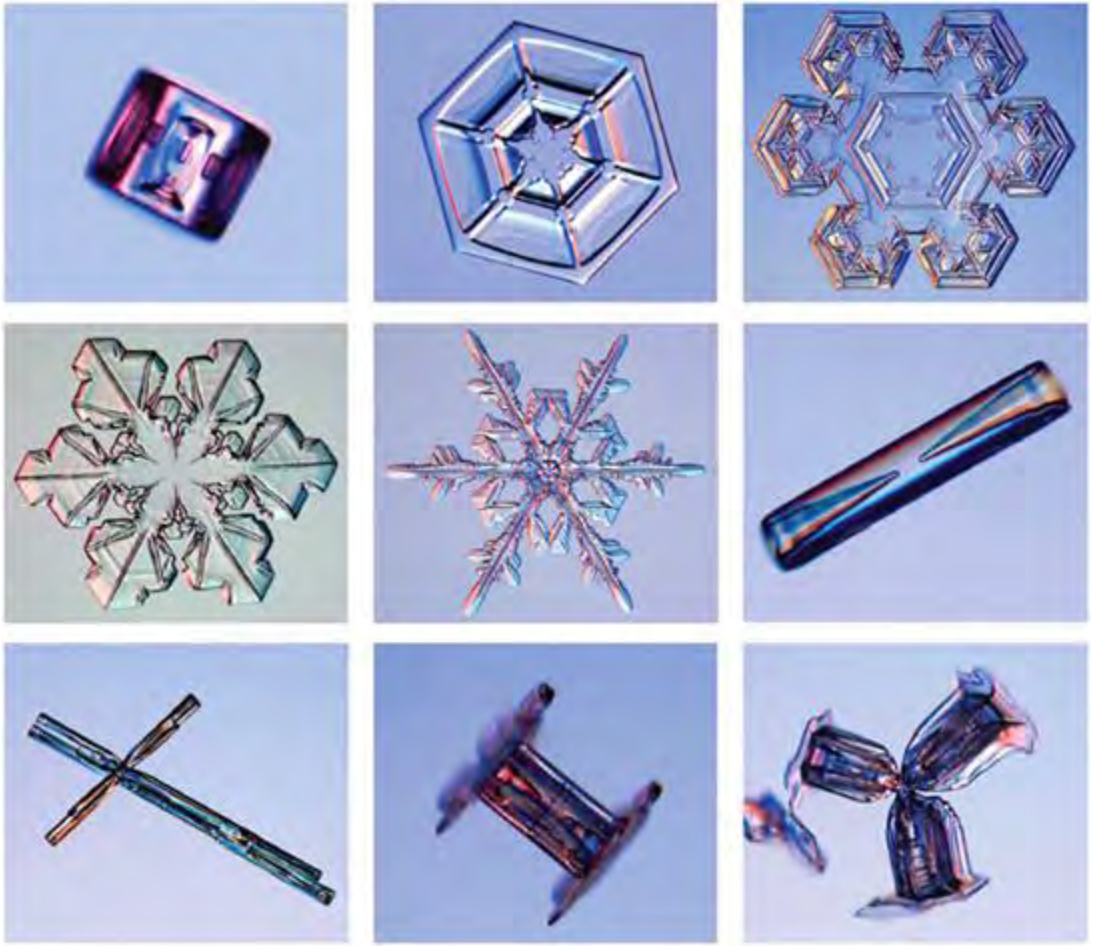

- Habit = Shape
- Habit ~ function of *temperature* and *supersaturation* (i.e., humidity)
- Habit influences:
  - microphysical process rates
  - fall speeds
  - optical properties
- E.g. Ice complexity may induce additional cooling effect of -1.1 W m^{-2} (Jarvinen et al. 2018)
  - For reference: CO_2 forcing is ~2 W m^{-2}

Source: Kenneth Libbrecht, snowcrystals.com

Millions of in situ CPI images available

*CPI = Cloud Particle Imager

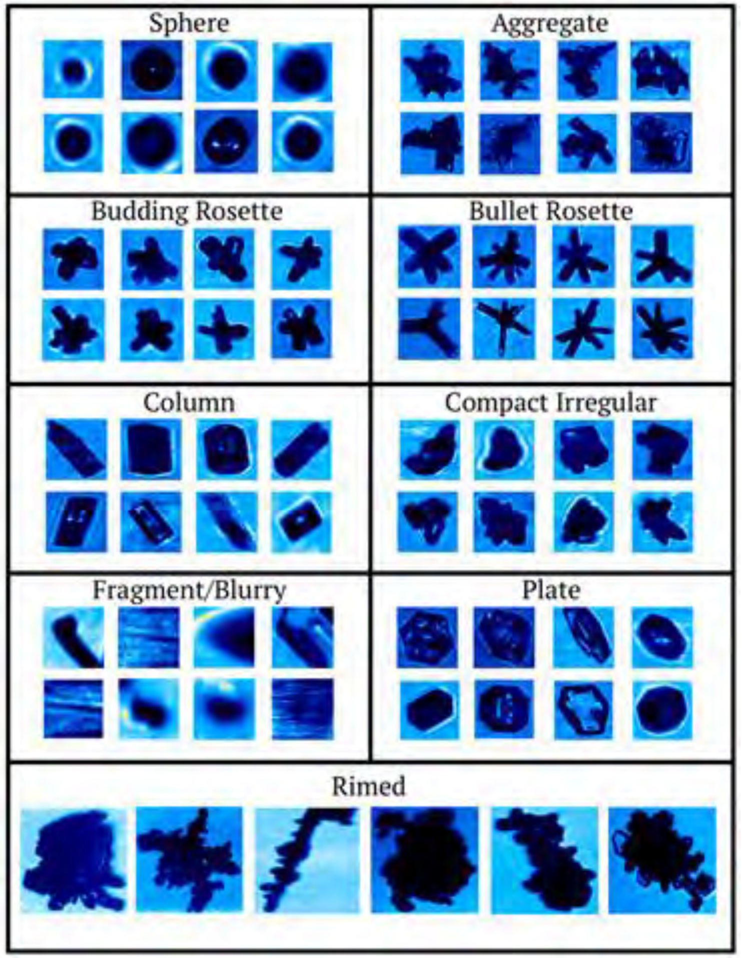

Source: [Przybylo et al. (2022)]

Source: [Xiao et al. 2019]

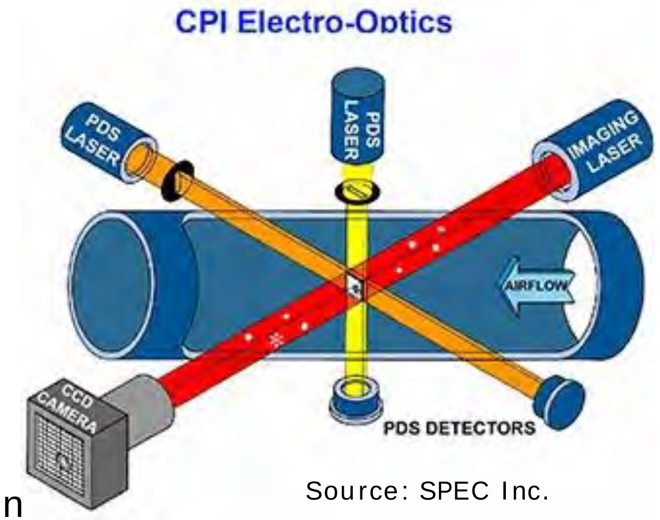

2.3 µm pixel resolution

Source: SPEC Inc.

Millions of in situ CPI images available

*CPI = Cloud Particle Imager

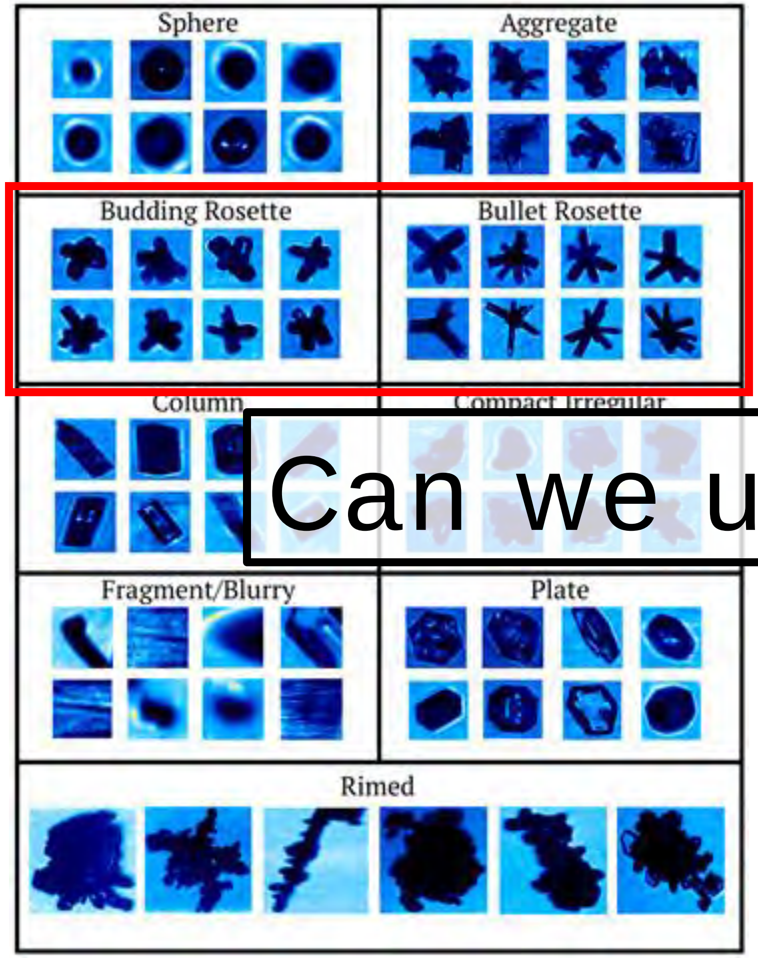

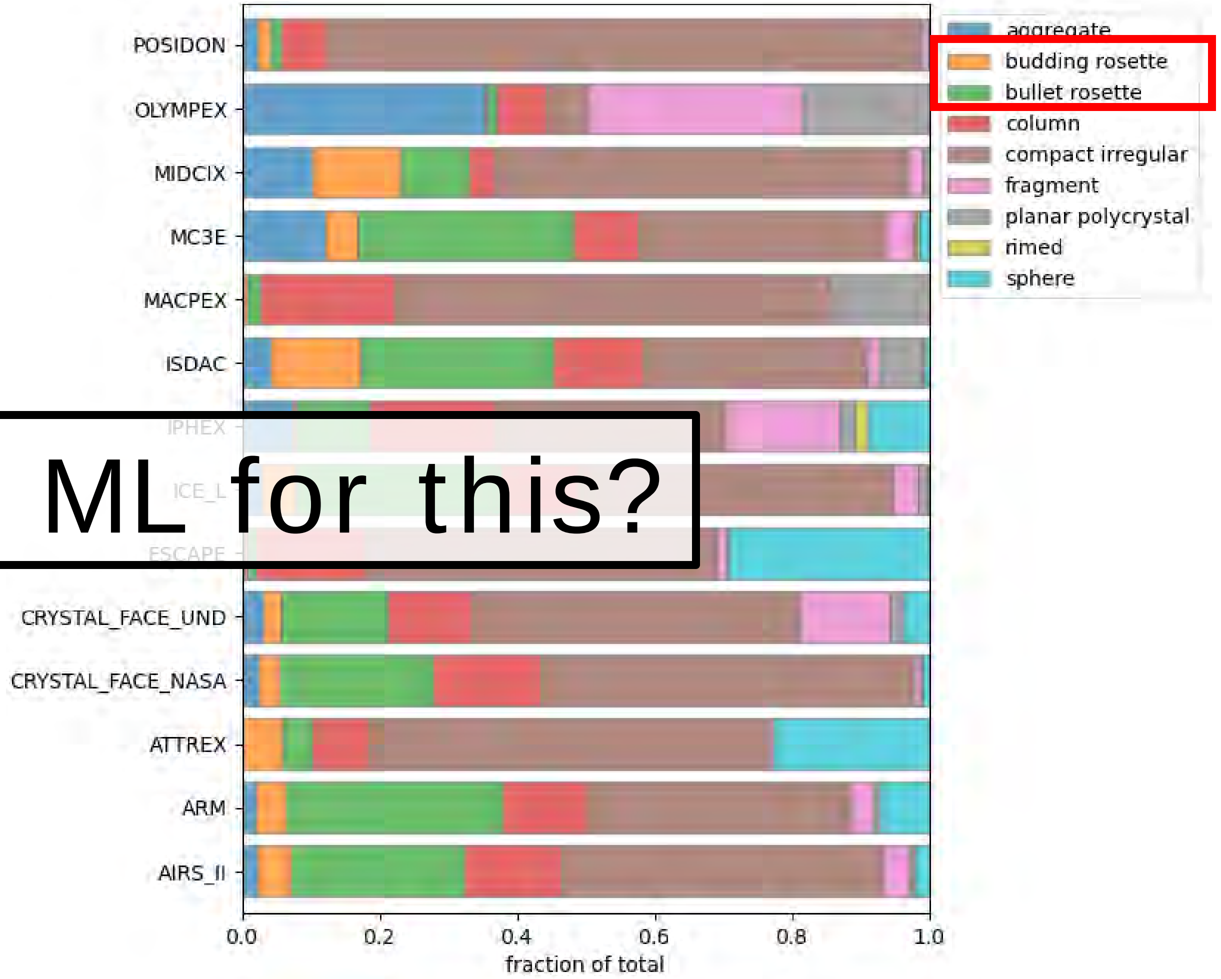

Can we use ML for this?

Source: Przybylo et al. (2022)

Challenge: We have 2-D images but we want 3-D relevant features

Part 1: Predicting 3-D properties from images (supervised)

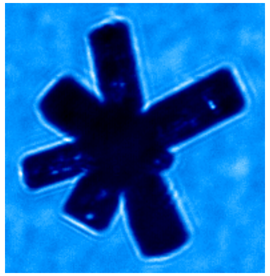

No ground truth? → Use synthetic data

A priori geometric model of bullet rosette

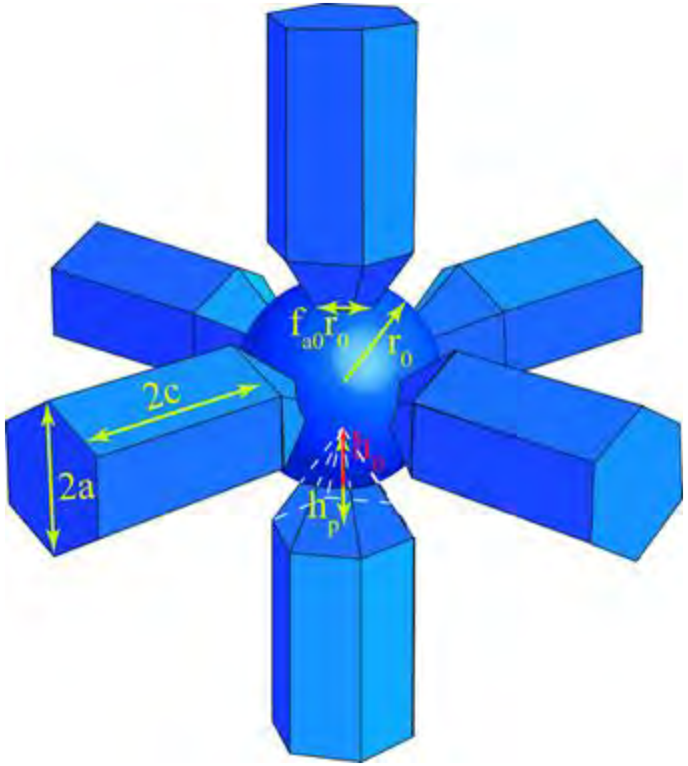

Source: [Pokrifka et al., 2023]

Computationally generate random variations

Synthetic 3-D models

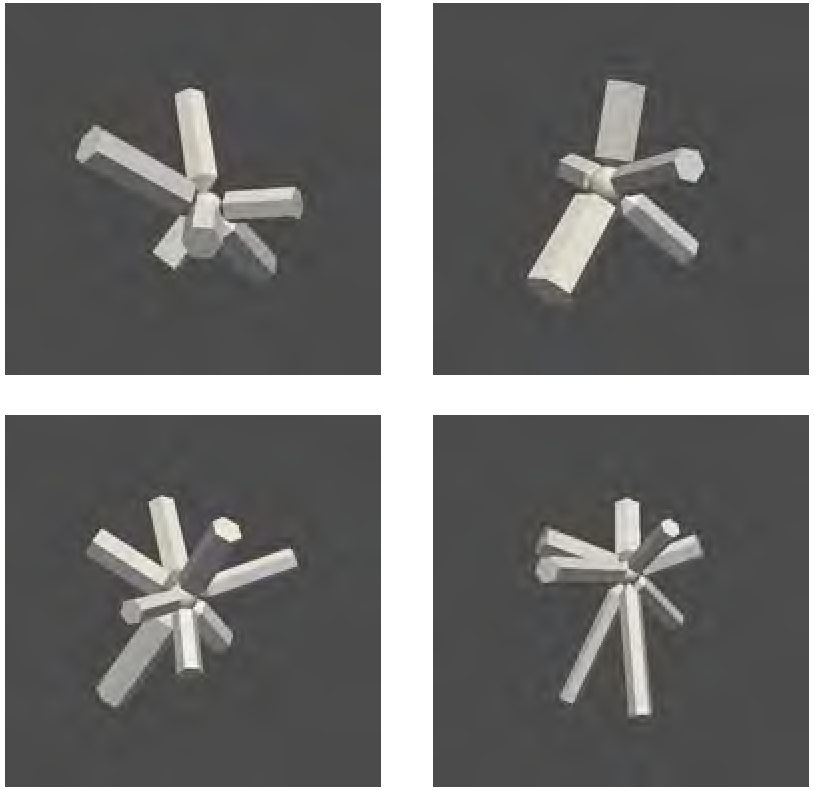

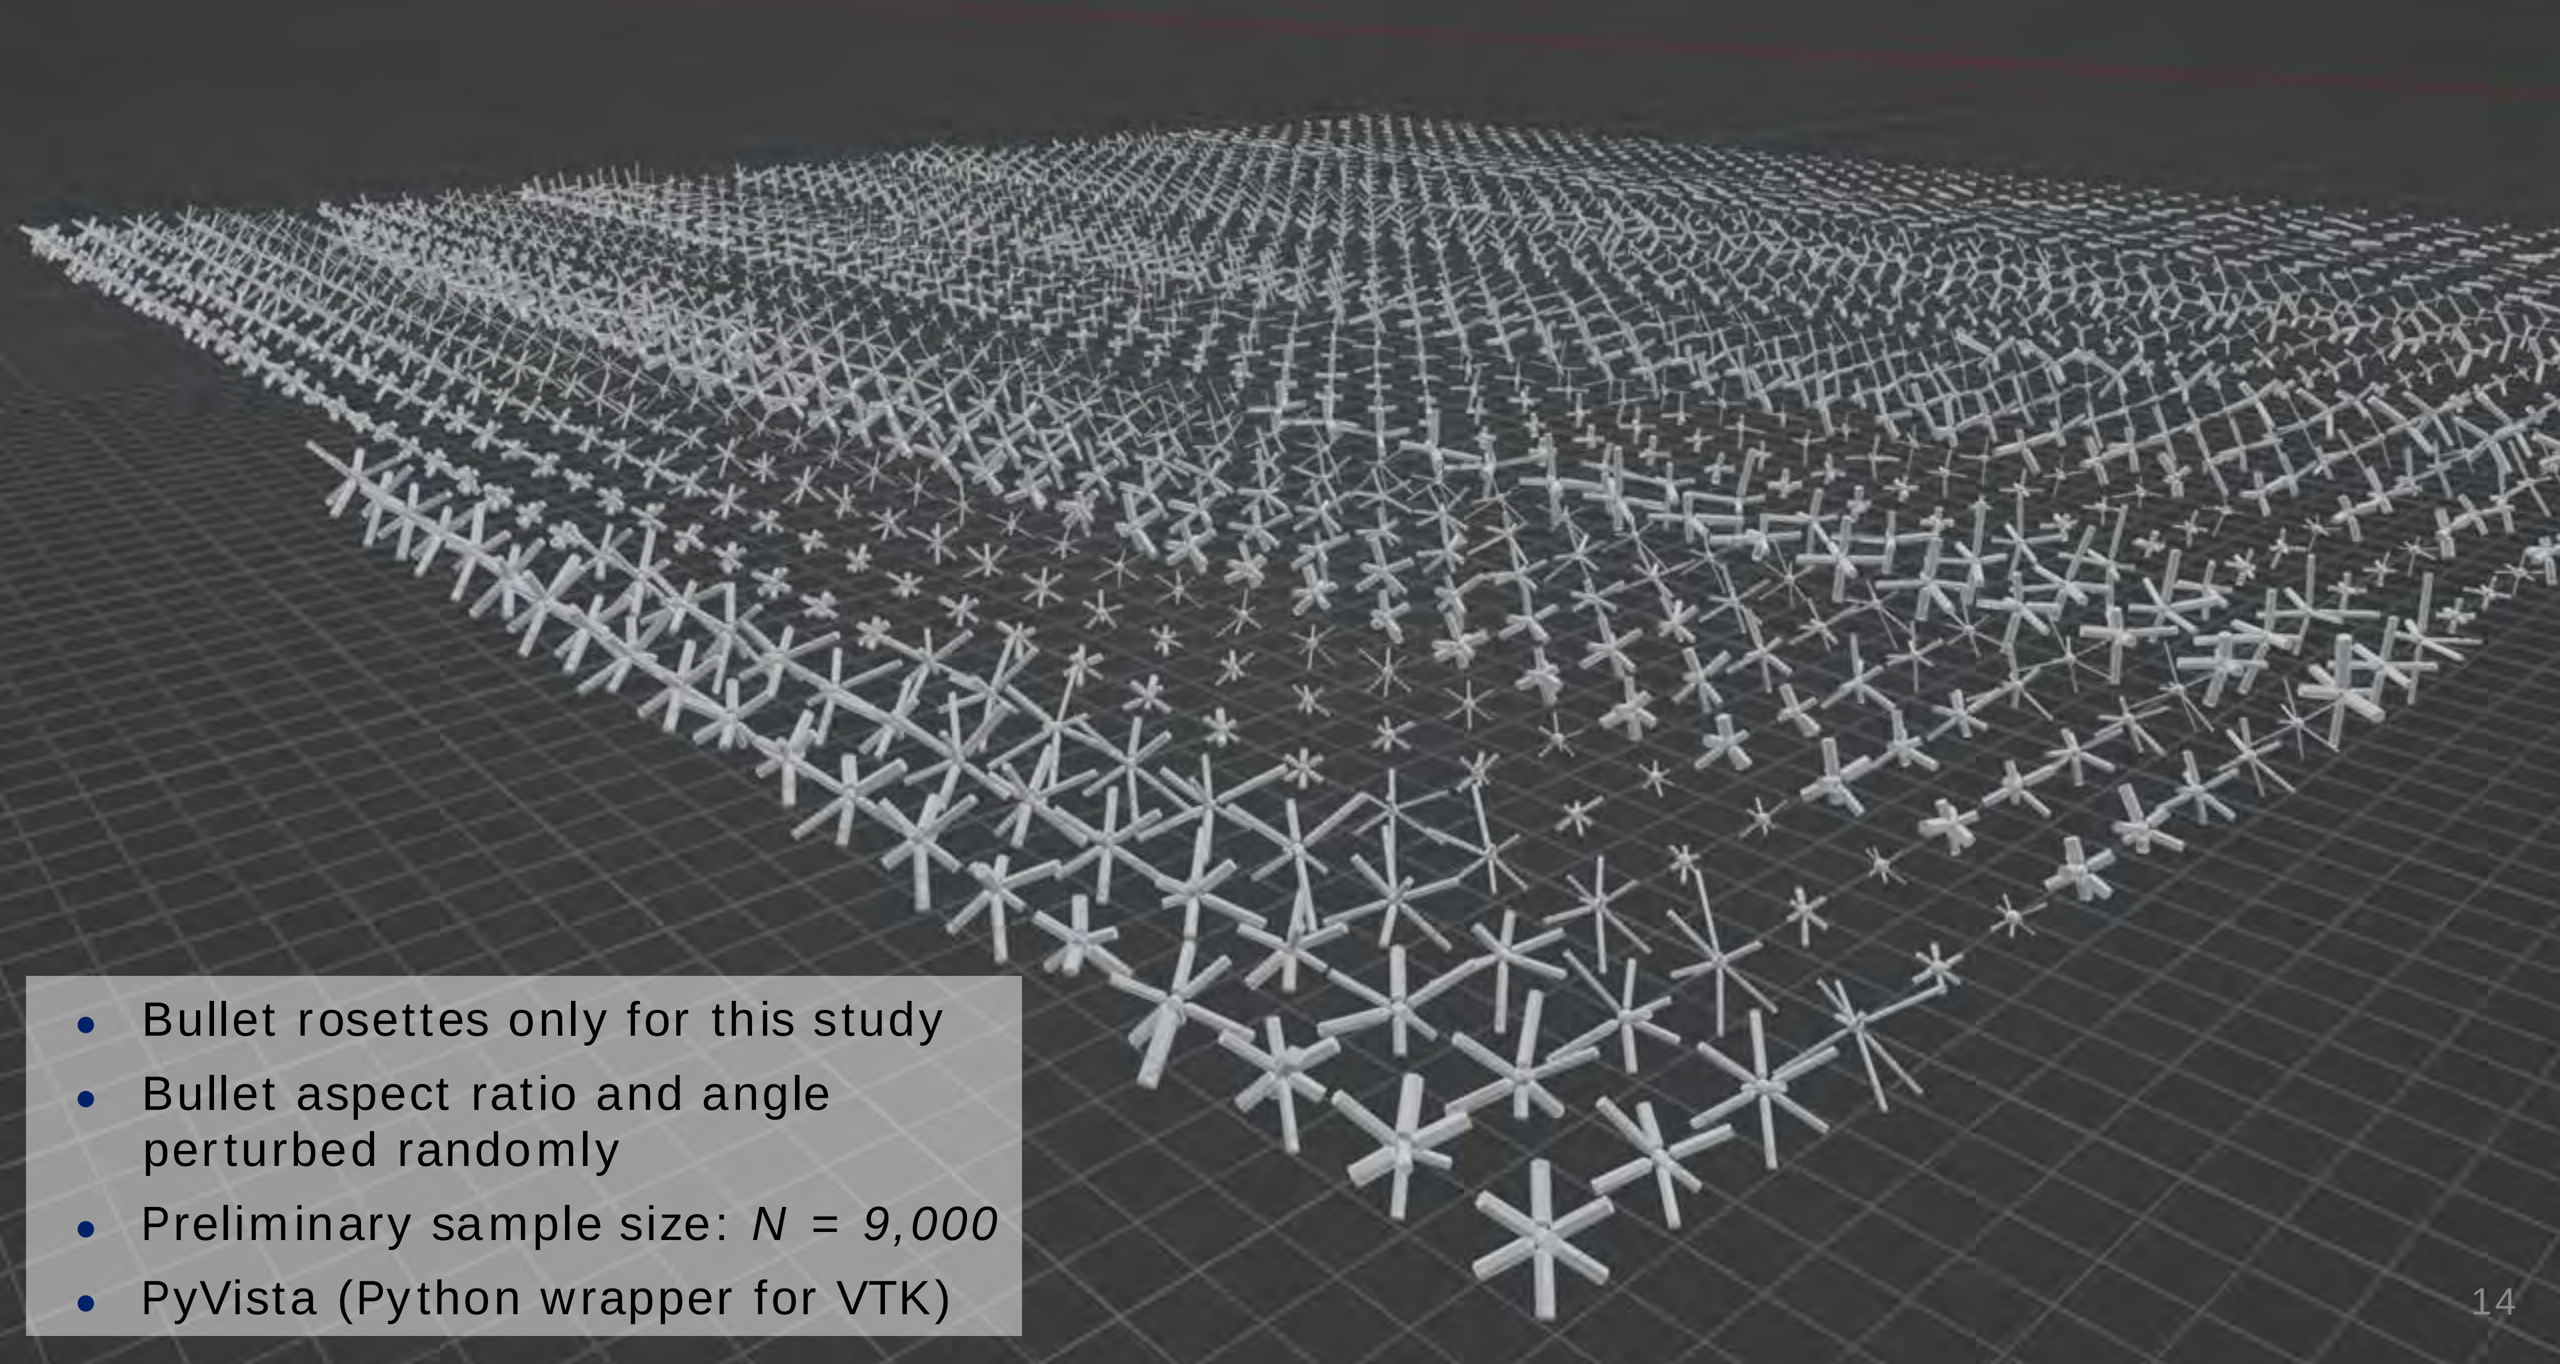

- Bullet rosettes only for this study
- Bullet aspect ratio and angle perturbed randomly
- Preliminary sample size: $N = 9{,}000$
- PyVista (Python wrapper for VTK)

Initial route: single-view 3-D reconstruction

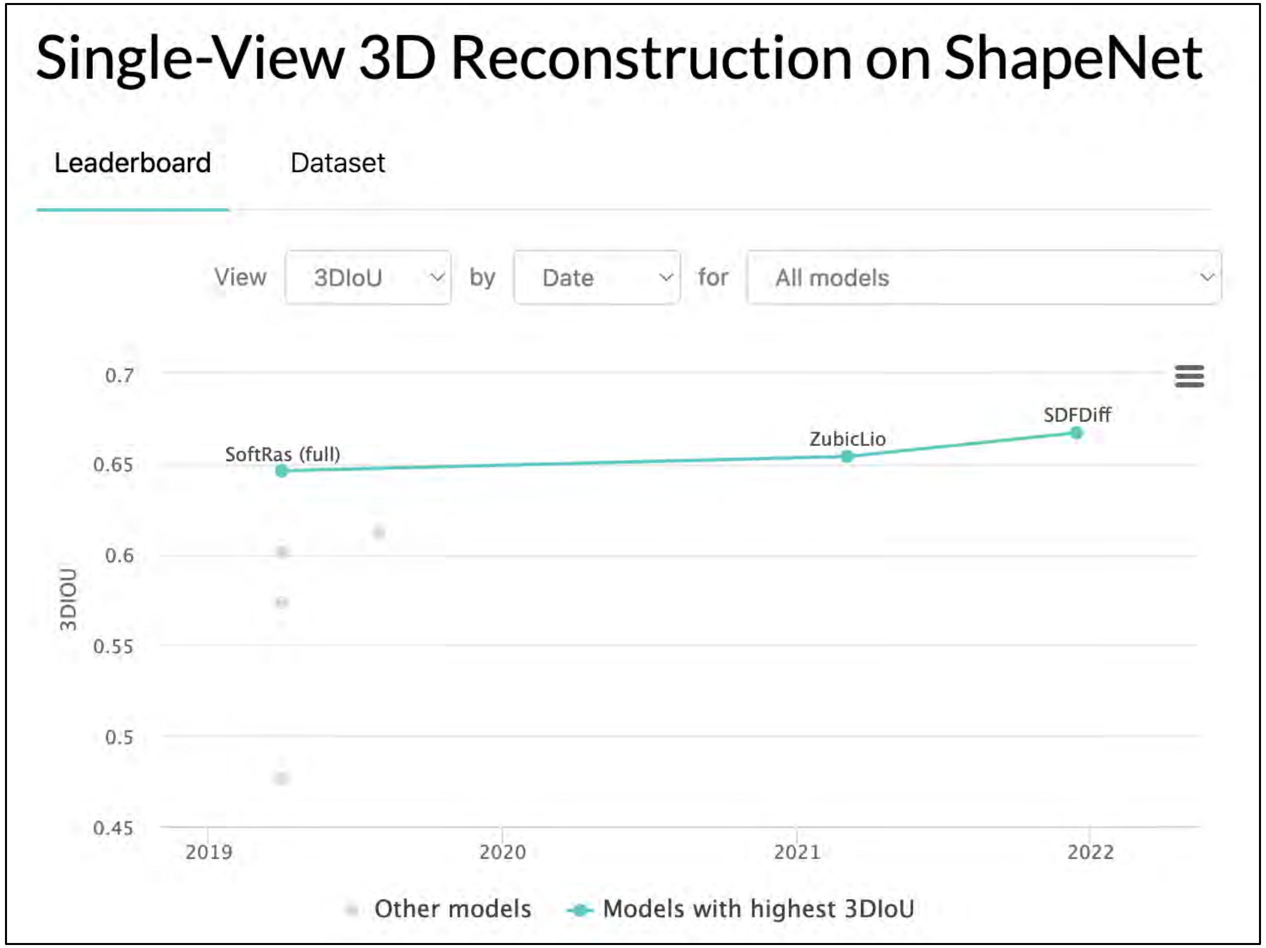

https://paperswithcode.com/sota/single-view-3d-reconstruction-on-shapenet

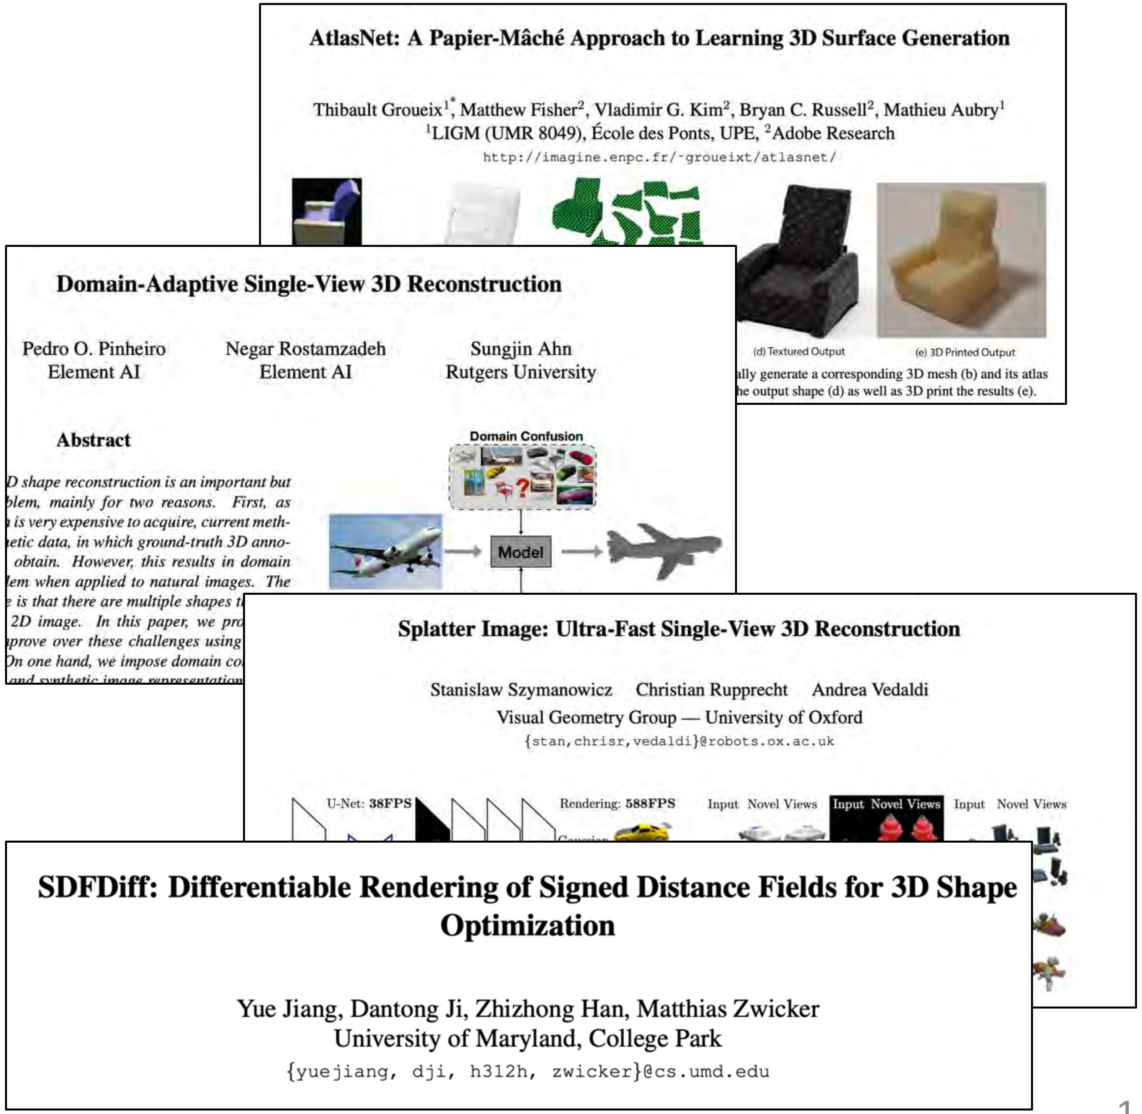

AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation

Thibault Groueix[1,*], Matthew Fisher[2], Vladimir G. Kim[2], Bryan C. Russell[2], Mathieu Aubry[1]
[1]LIGM (UMR 8049), École des Ponts, UPE, [2]Adobe Research
http://imagine.enpc.fr/~groueixt/atlasnet/

Domain-Adaptive Single-View 3D Reconstruction

Pedro O. Pinheiro, Element AI
Negar Rostamzadeh, Element AI
Sungjin Ahn, Rutgers University

Abstract

Splatter Image: Ultra-Fast Single-View 3D Reconstruction

Stanislaw Szymanowicz Christian Rupprecht Andrea Vedaldi
Visual Geometry Group — University of Oxford
{stan,chrisr,vedaldi}@robots.ox.ac.uk

SDFDiff: Differentiable Rendering of Signed Distance Fields for 3D Shape Optimization

Yue Jiang, Dantong Ji, Zhizhong Han, Matthias Zwicker
University of Maryland, College Park
{yuejiang, dji, h312h, zwicker}@cs.umd.edu

3D-R2N2: out of the box implementation

3D Recurrent Reconstruction Neural Network (3D-R2N2): Choy et al., 2016

Training:

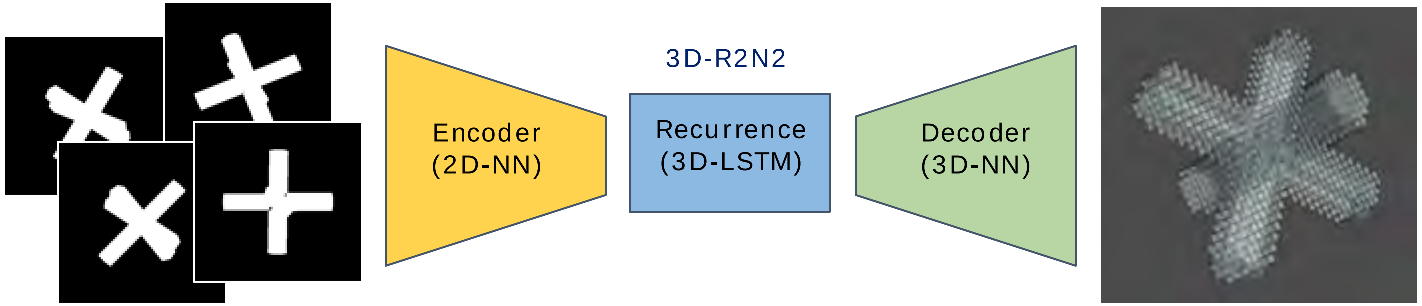

2-D projections from training set

corresponding voxelized model

Testing/Inference:

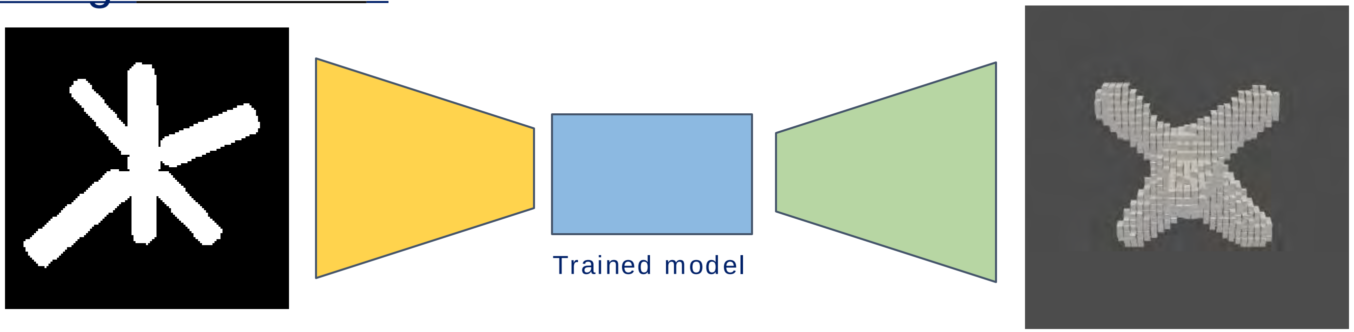

Testing: unseen projections
Inference: CPI images

inferred voxelized model

Maybe explicit reconstruction isn't the way to go…

What Do Single-view 3D Reconstruction Networks Learn?

Maxim Tatarchenko*[1], Stephan R. Richter*[2], René Ranftl[2], Zhuwen Li[2], Vladlen Koltun[2], and Thomas Brox[1]

[1]University of Freiburg [2]Intel Labs

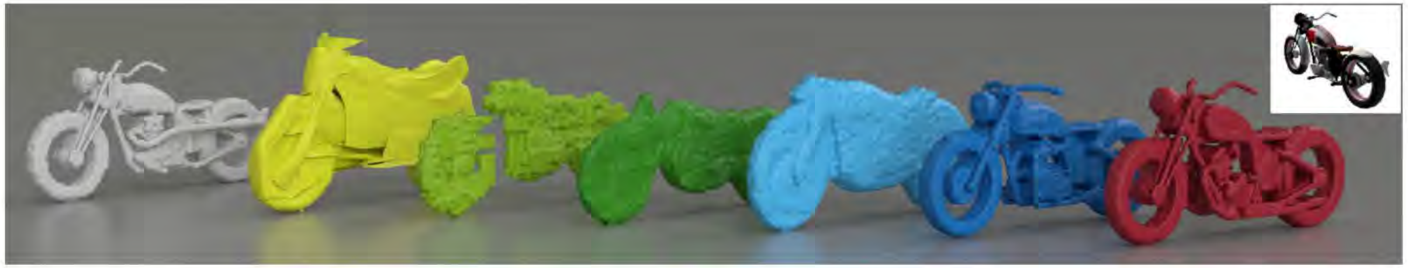

Figure 1. We provide evidence that state-of-the-art single-view 3D reconstruction methods (AtlasNet (light green, 0.38 IoU) [12], OGN (green, 0.46 IoU) [46], Matryoshka Networks (dark green, 0.47 IoU) [37]) do not actually perform reconstruction but image classification. We explicitly design pure recognition baselines (Clustering (light blue, 0.46 IoU) and Retrieval (dark blue, 0.57 IoU)) and show that they produce similar or better results both qualitatively and quantitatively. For reference, we show the ground truth (white) and a nearest neighbor from the training set (red, 0.76 IoU). The inset shows the input image.

Tatarchenko et al. 2019 (CVPR)

"In this work, we set up two alternative approaches that perform image classification and retrieval respectively. These simple baselines yield better results than state-of-the-art methods, both qualitatively and quantitatively. We show that encoder-decoder methods are statistically indistinguishable from these baselines, thus indicating that the current state of the art in single-view object reconstruction does not actually perform reconstruction but image classification."

Back to basics: simple supervised learning

Training:

Extract 2d features from projections

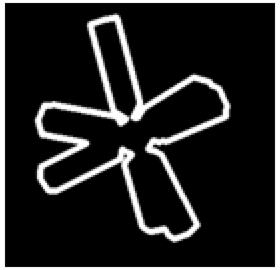

Data in tabular form

X (input): the 2d features

Y (output): e.g., effective density, surface area, # arms

Train models

Testing:

unseen data

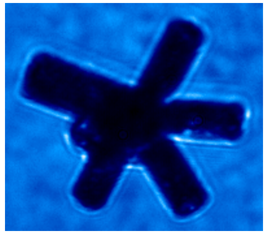

Trained ML model

Predictions:
- mass
- # arms
- surface area
- etc.

Pipeline to predict 3-D targets

9,000
synthetic models

Calculate & save
target outputs:

(1) # arms,
(2) effective density
(3) effective surface area

9,000 x 24 views
= 216,000 images (samples)

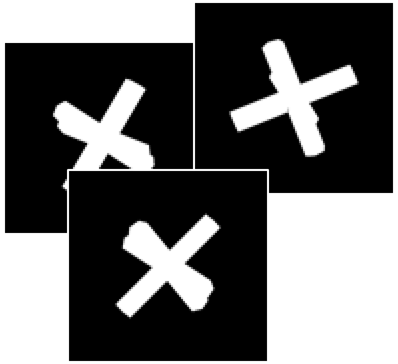

Take 24 random projections & calculate 2-D features for each projection:

(1) Aspect ratio
(2) Elliptical aspect ratio
(3) # extreme points
(4) Contour area
(5) Area ratio
(6) Complexity
(7) Circularity

Create merged
tabular dataset

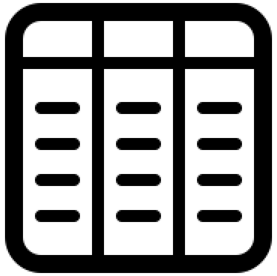

- 216,000 rows
- 7 feature columns (inputs)
- 3 target columns (outputs)

80%

Train model

20%

Test model

Random forest regression: predicting surface area and mass

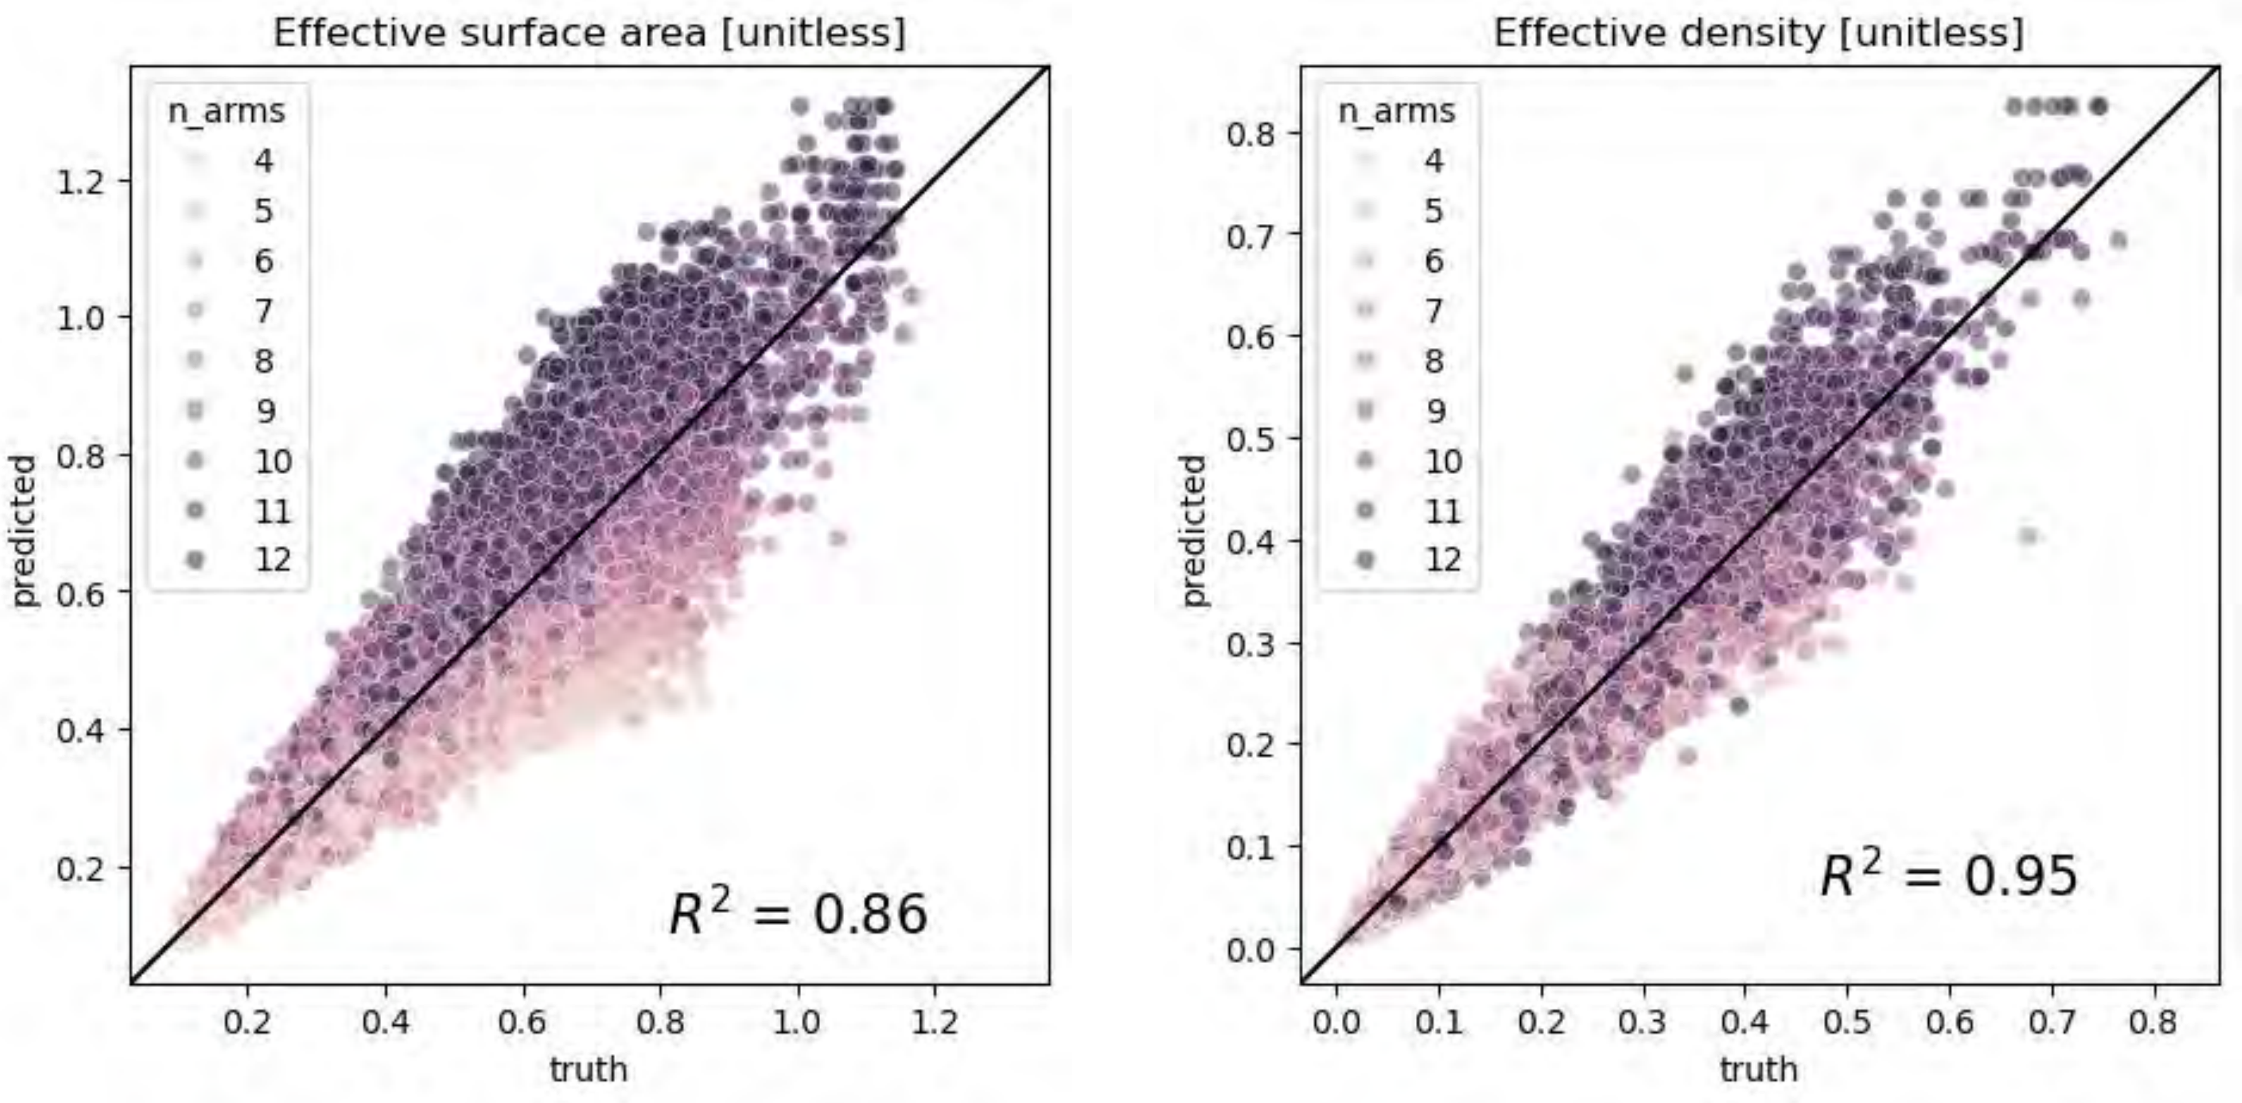

Random forest > linear baseline

Multiple Linear Regression

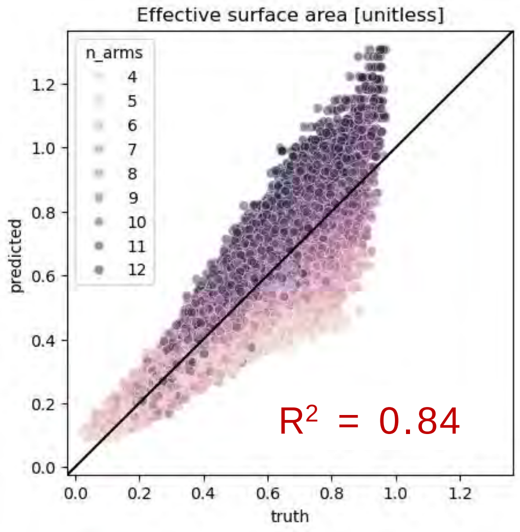

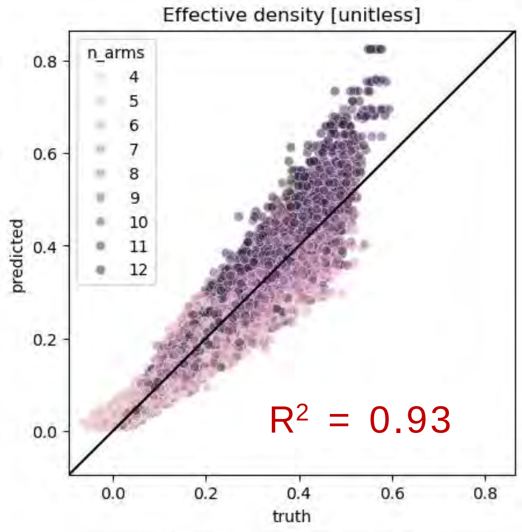

Random Forest

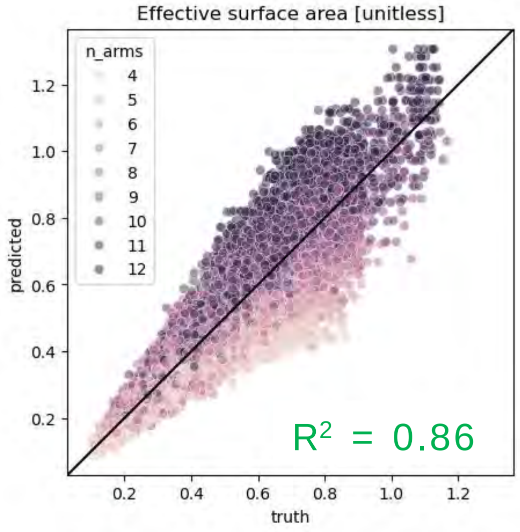

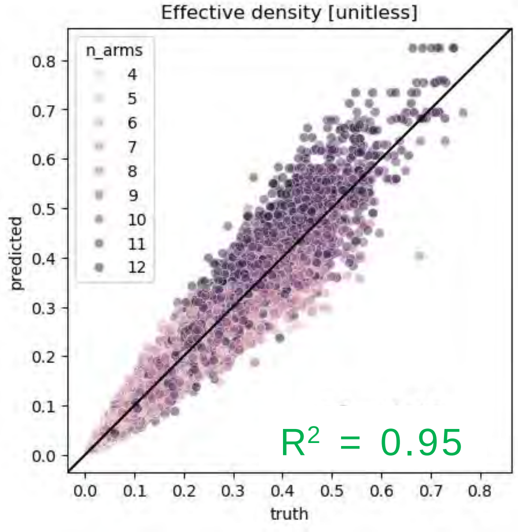

This indicates some non-linearities being captured by the random forest

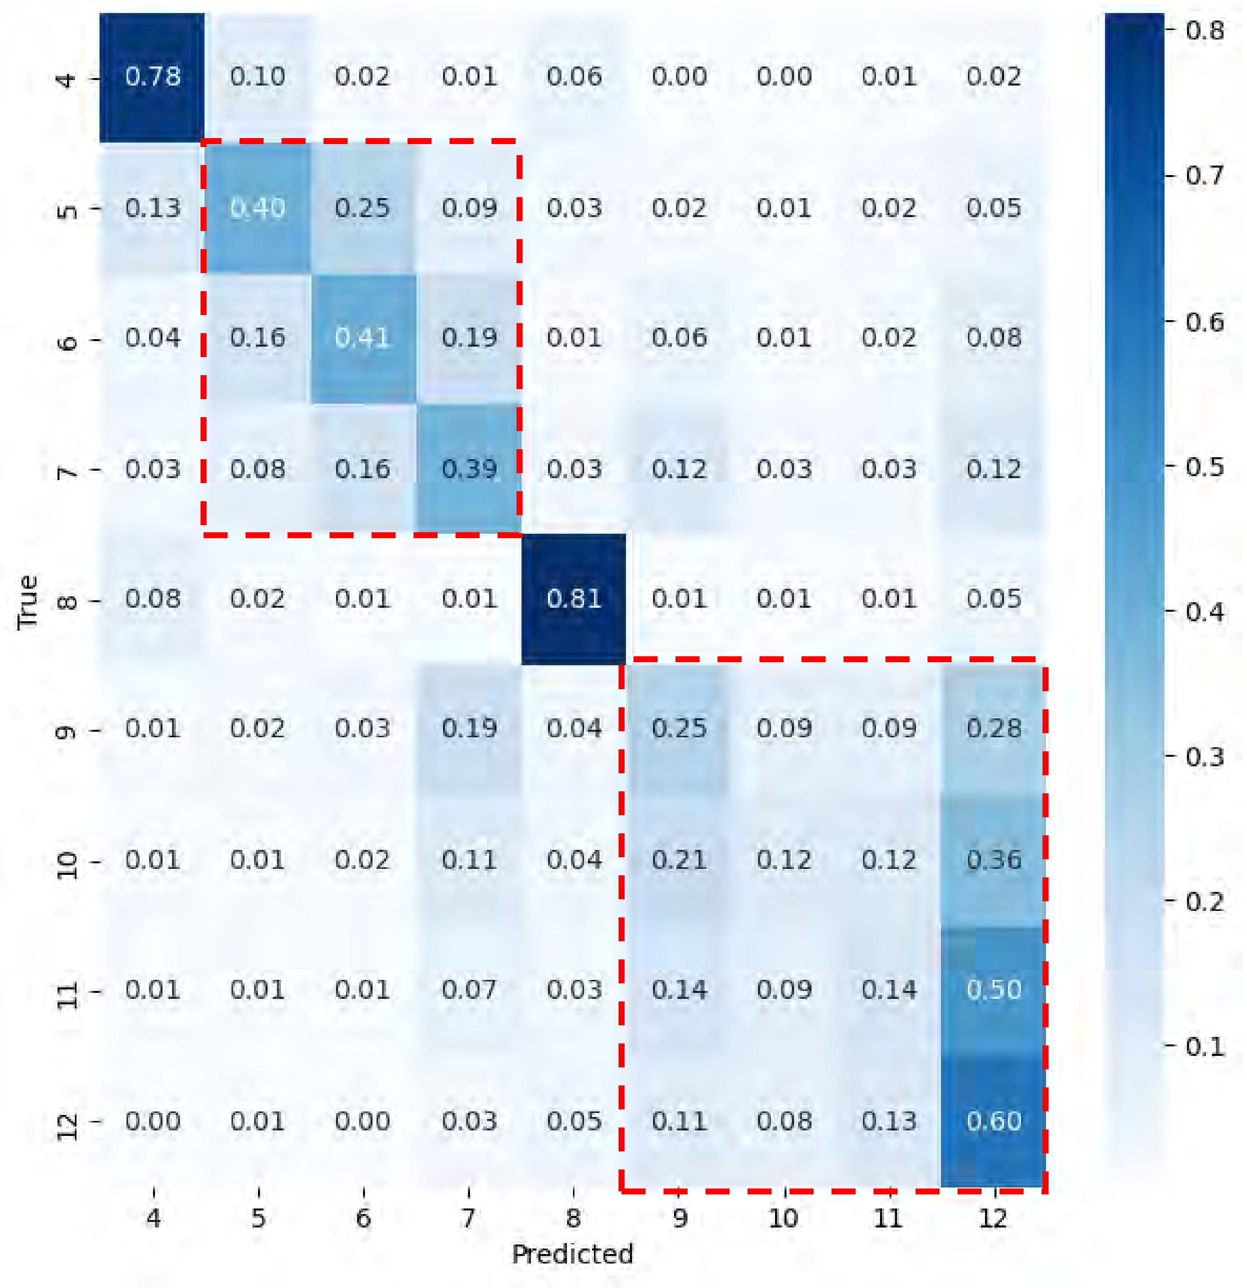

Confusion matrix for # of bullets predicted by random forest

Normalized by row (i.e., each row sums up to 1.0)

A perfect predictor would show 1.0 values in the diagonal

We can also use deep learning (if we want)

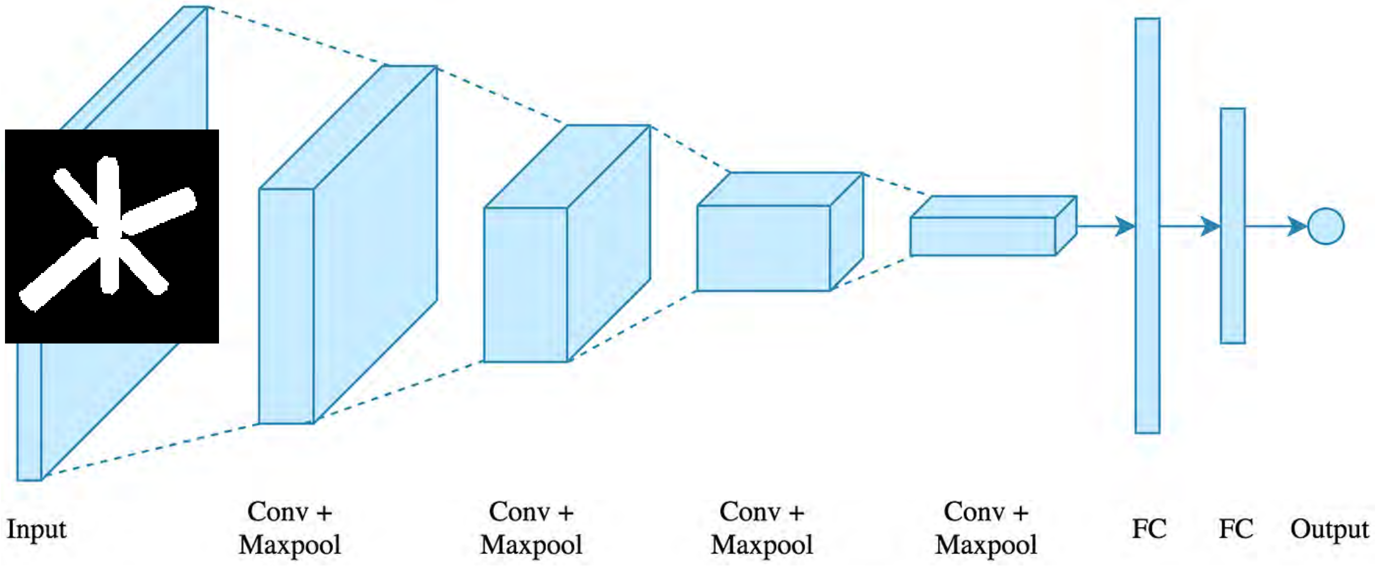

Credit: [Arden Dertat]

Predict:
- Effective density
- Effective surface area
- # arms
- Etc.

Preliminary results:

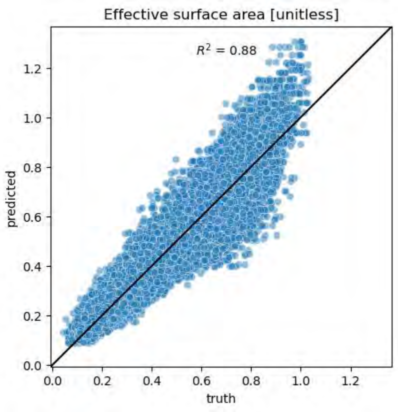

0.88 (CNN) > 0.86 (RF)

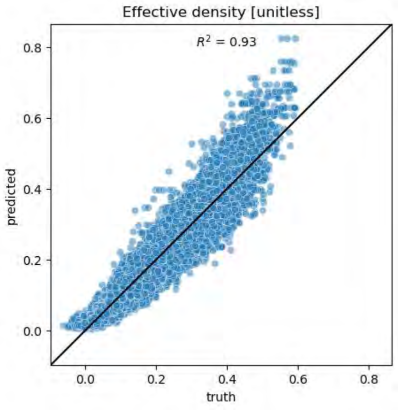

0.93 (CNN) < 0.95 (RF)

What if we have additional views?

2D-S probe

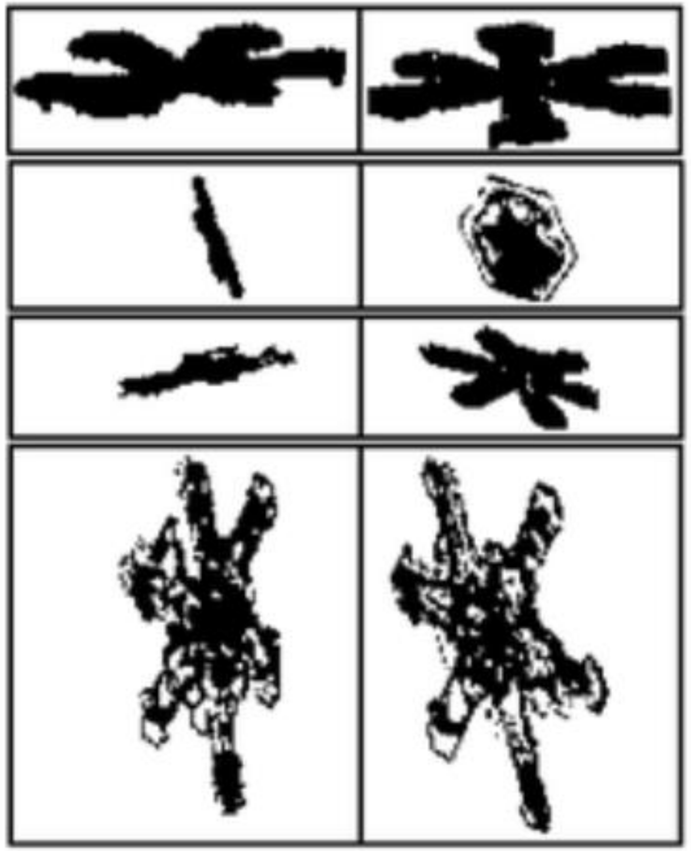

Lawson et al. 2006

PHIPS-HALO

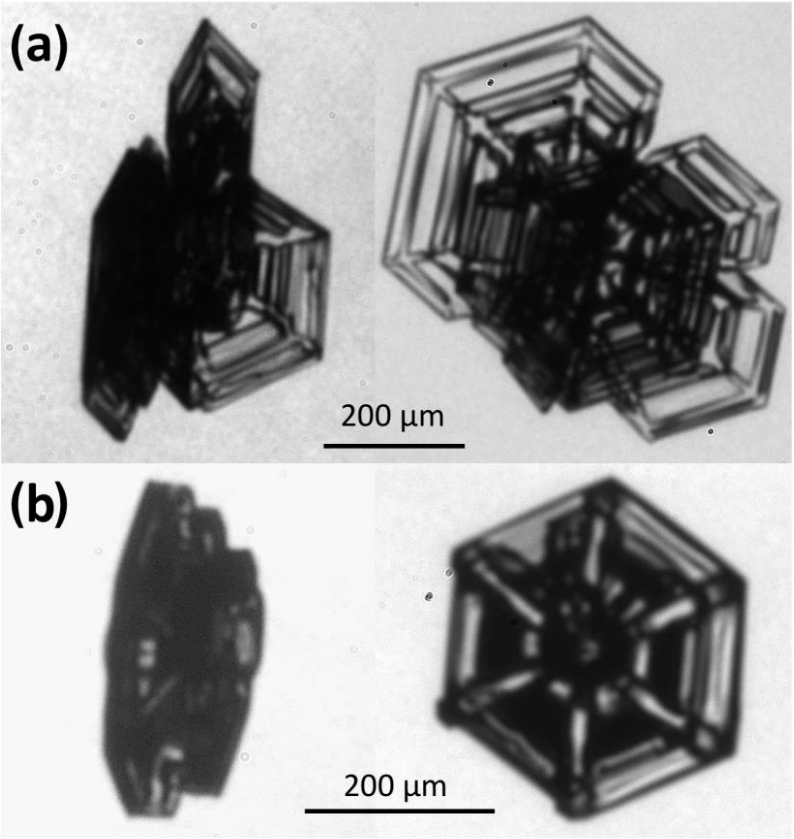

Schnaiter et al. 2018

Pipeline w/ two views

9,000
synthetic
models

Calculate & save target outputs:

(1) # arms,
(2) effective density
(3) effective surface area

9,000 x *24 image pairs*
= 216,000 image pairs

View 1 View 2

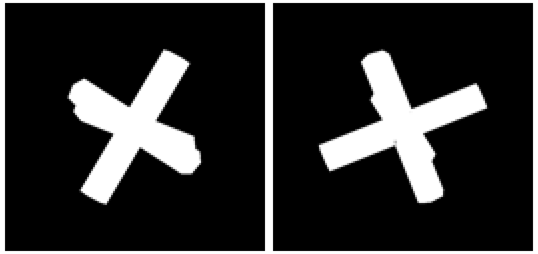

Take 24 random *projection pairs* & calculate 2-D features for each projection image of each pair:

(1) Aspect ratio
(2) Elliptical aspect ratio
(3) # extreme points
(4) Contour area
(5) Area ratio
(6) Complexity
(7) Circularity

Create merged tabular dataset

- 216,000 rows
- *7x2 = 14 feature columns (inputs)*
- 3 target columns (outputs)

80% → Train model

20% → Test model

Train model → Test model

Two view are better than one

Random Forest
w/ single view

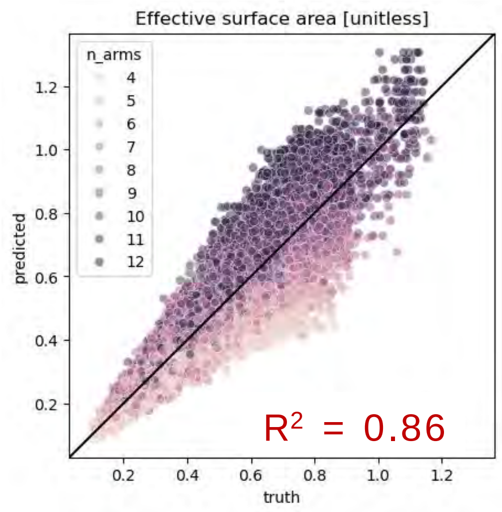

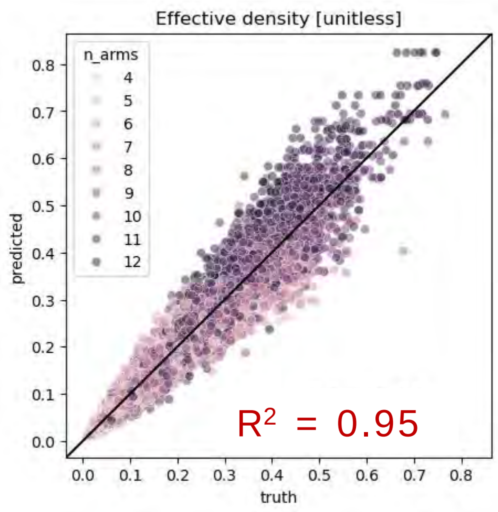

Random Forest
w/ two views

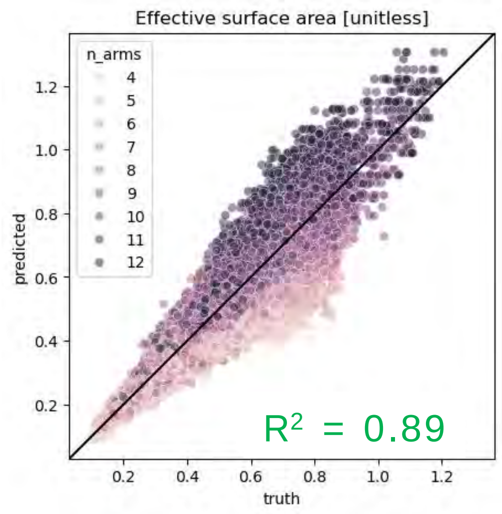

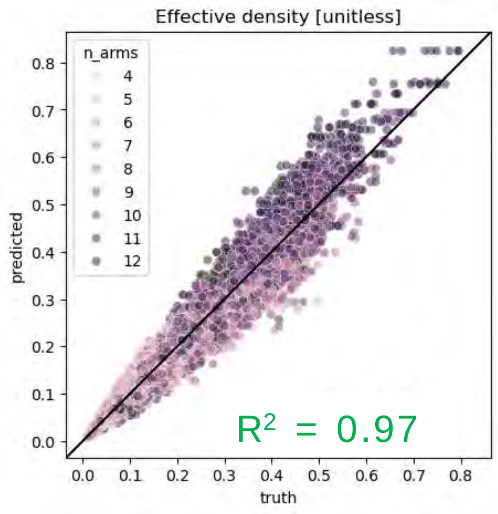

Two view are better than one

Single view

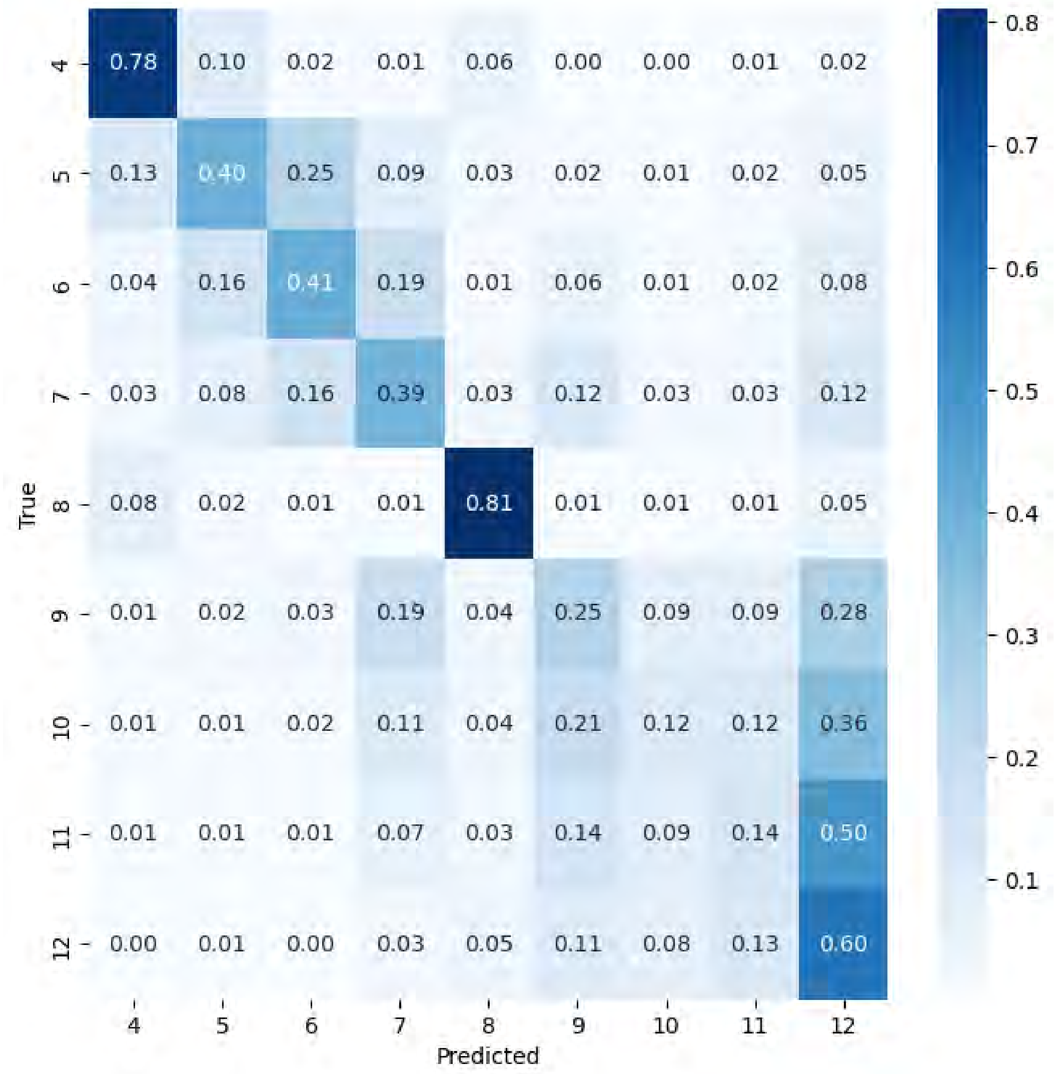

Two views

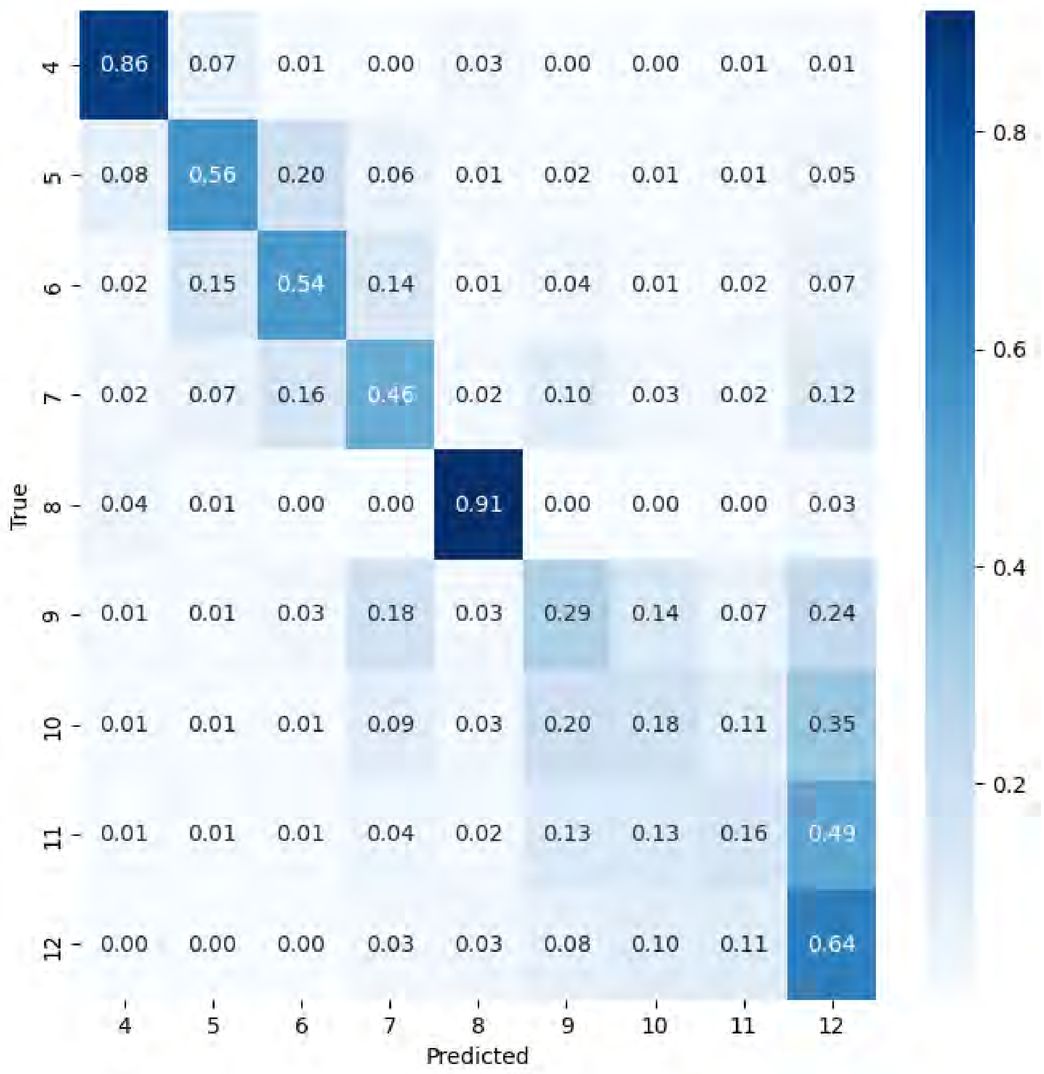

Part 1: Key Points

- A dataset of synthetic bullet rosettes was created
- ML was able to predict effective density and surface area with encouraging skill (to be improved)
- The classification of # arms was more challenging
- Inferring 3-D properties from CPI images will allow us to improve parameterizations moving forward

Part 2: Latent representations of ice crystals (unsupervised)

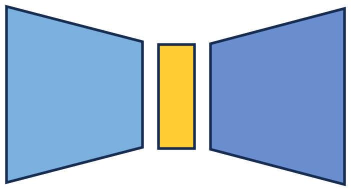

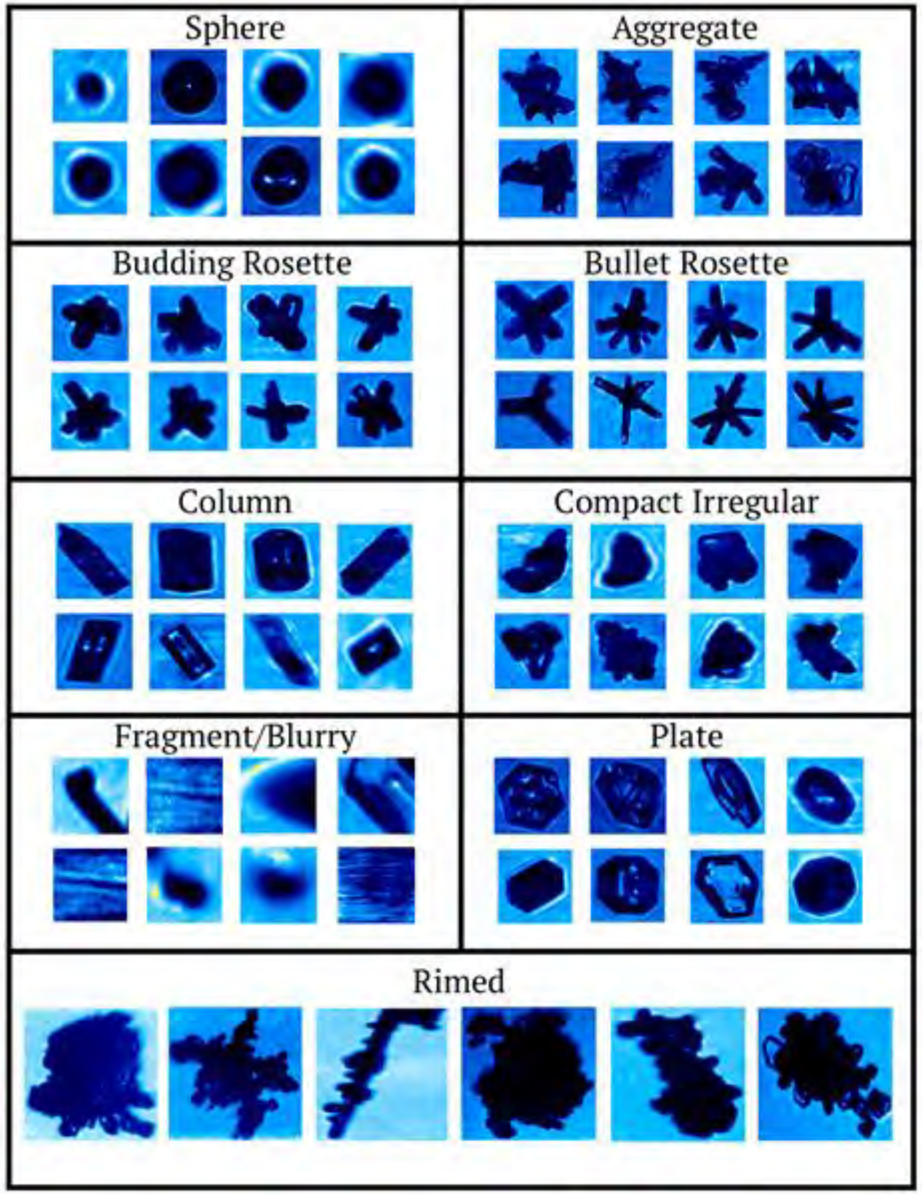

Source: [Przybylo et al. (2022)]

Initial question: Can we classify crystals in an unsupervised manner?

Unsupervised clustering pipeline

Step 1:
Train VAE with CPI images

Step 2:
Unsupervised clustering in latent space

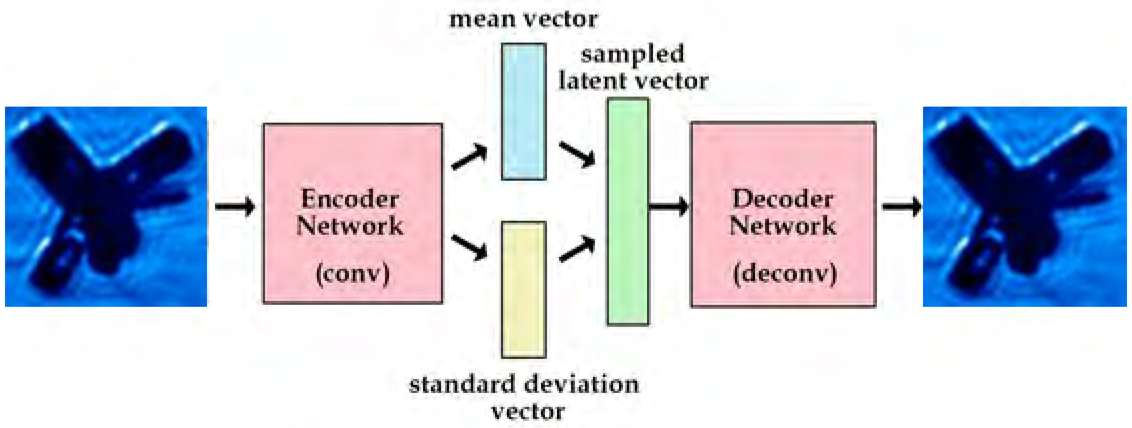

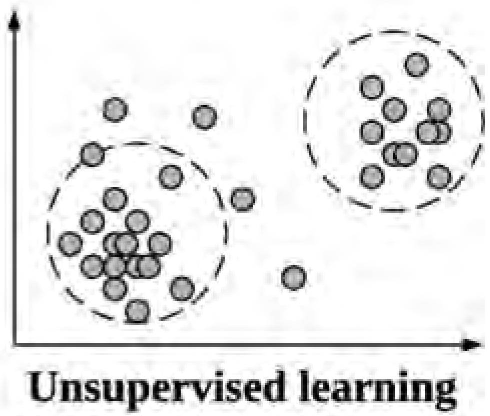

Simple case study: spheres vs. rosettes

?

Simple case study: Spheres vs. Rosettes

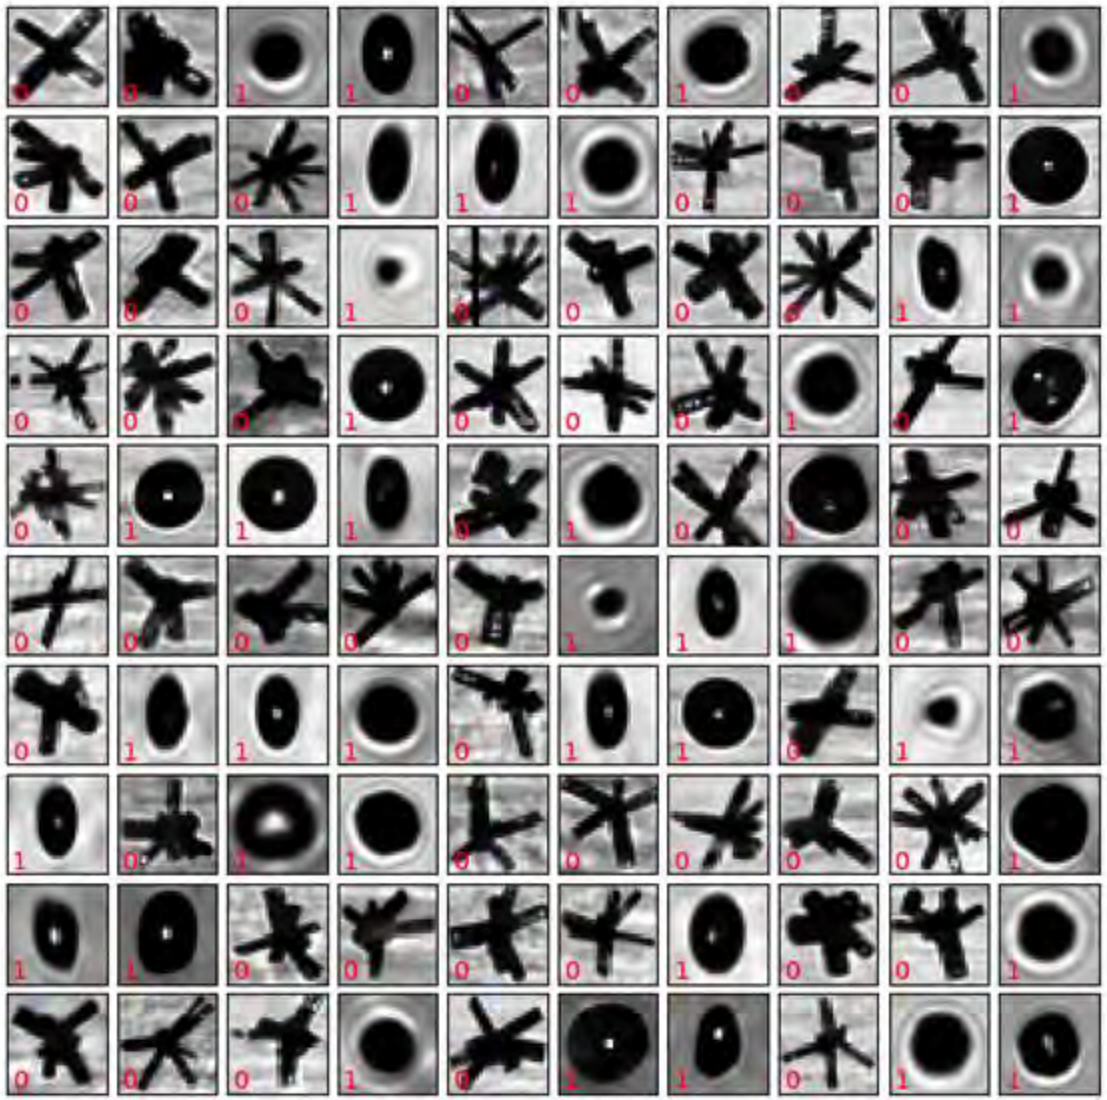

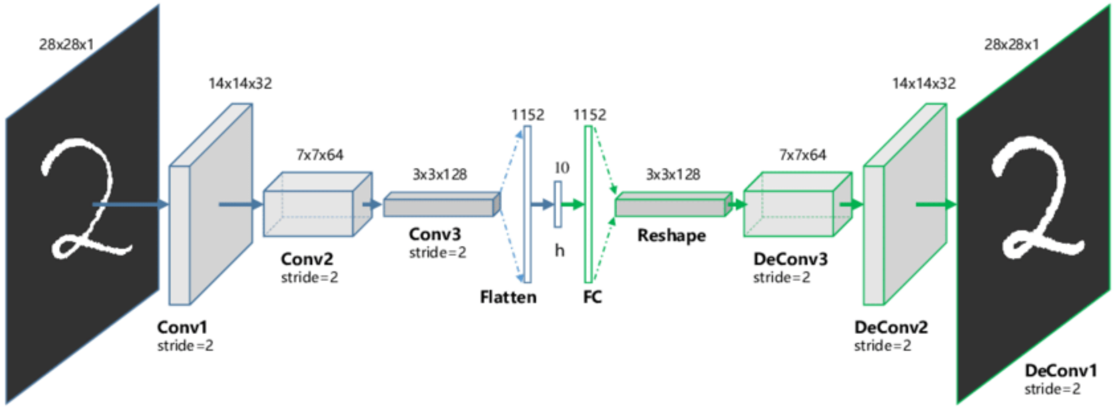

- Using PyroVED Python package
- Training: 100 images from each class
- 2 latent variables
- 28 x 28 resolution

Qualitatively inspecting the latent manifold

Visualizing the latent manifold
(z = latent variable)

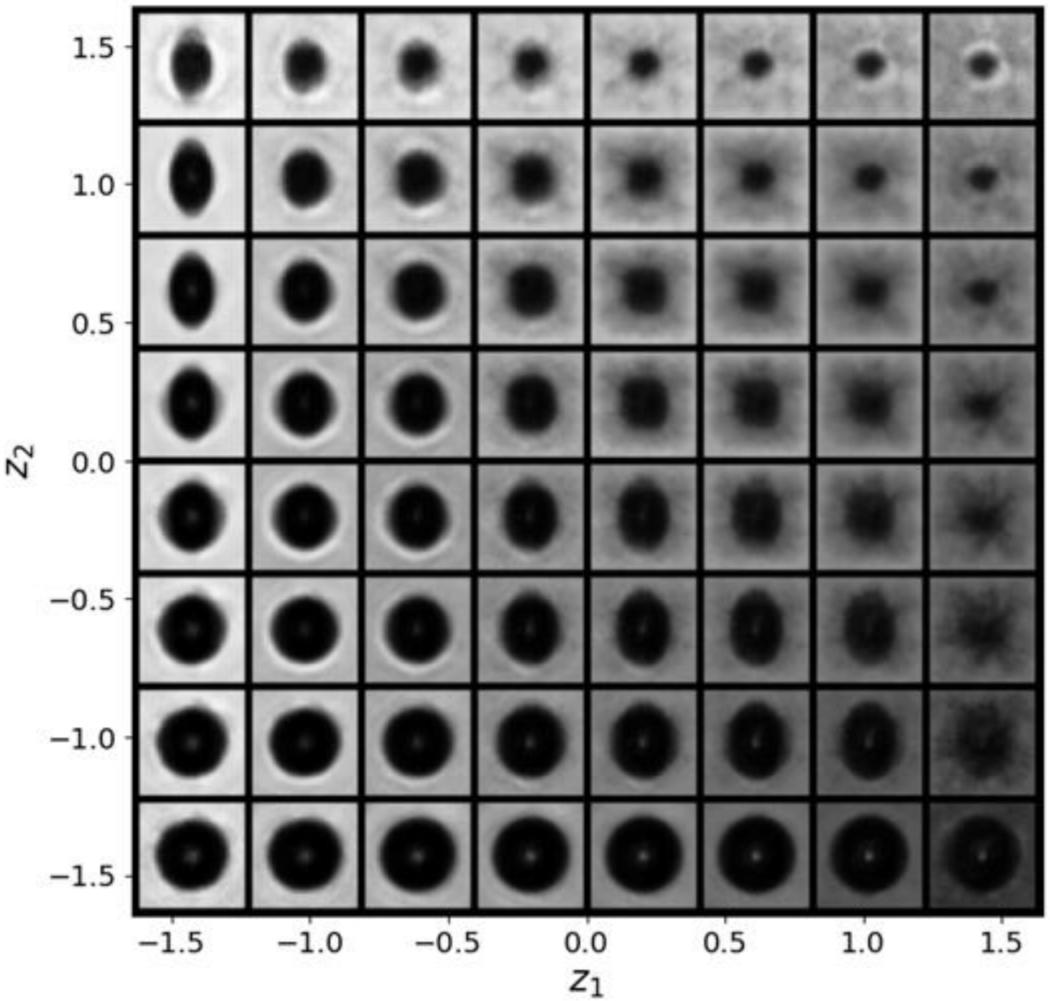

Scatter plot in latent space

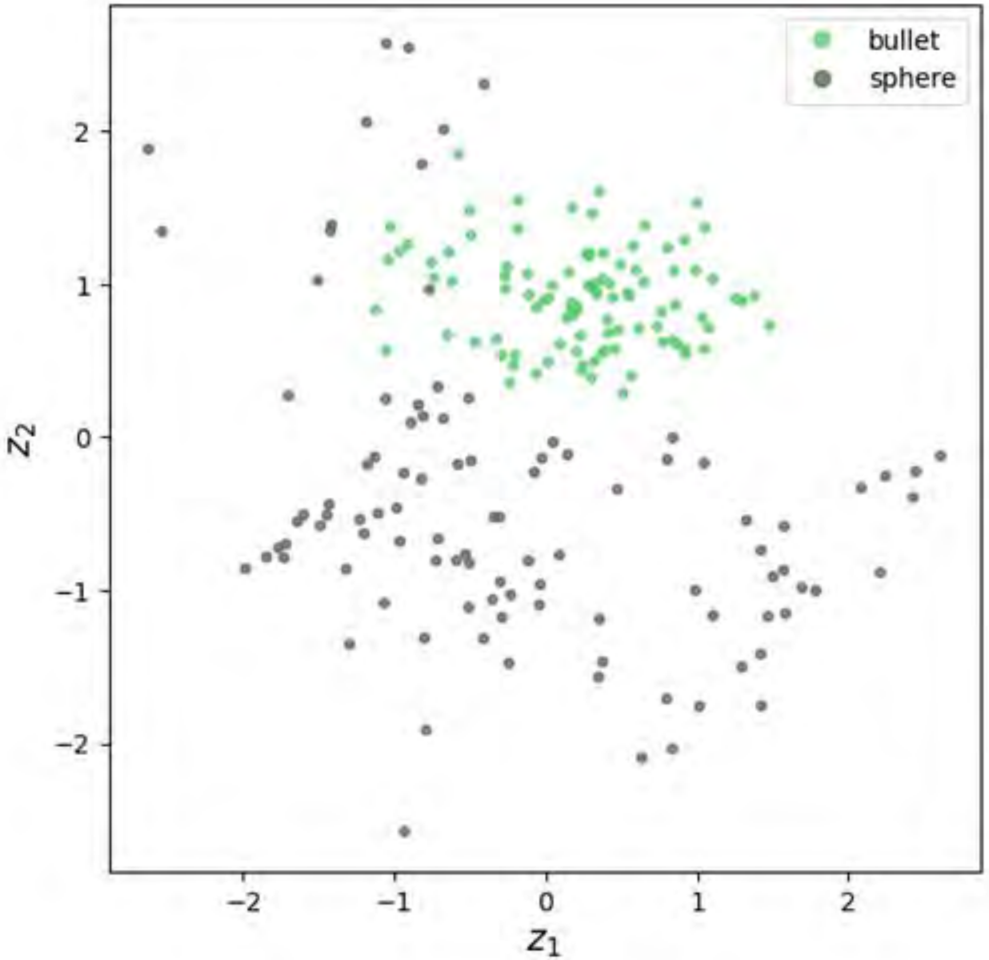

K-means used to cluster data

scatter plot in latent space w/ true labels

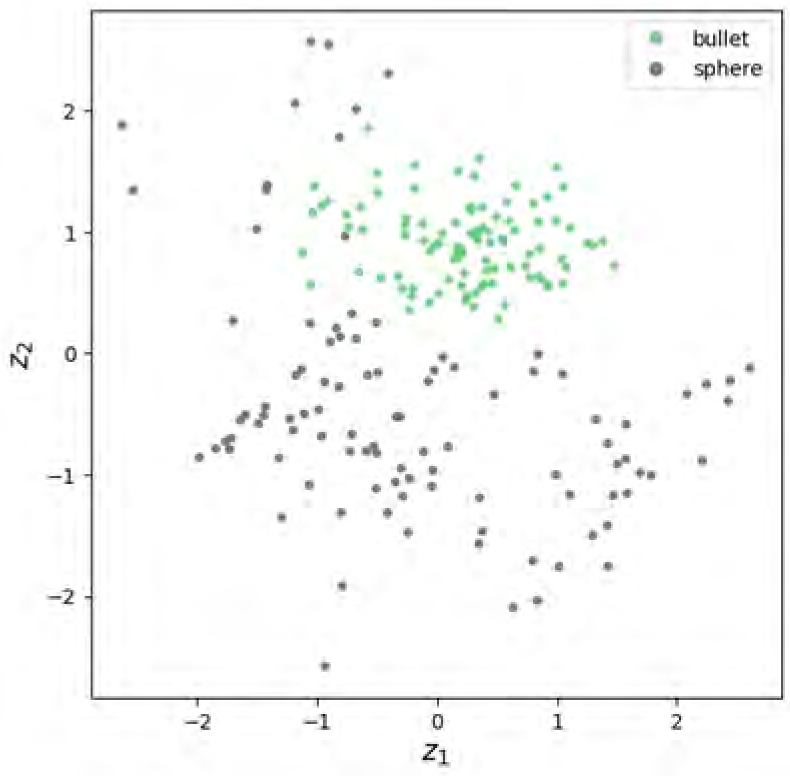

scatter plot in latent space w/ predicted labels

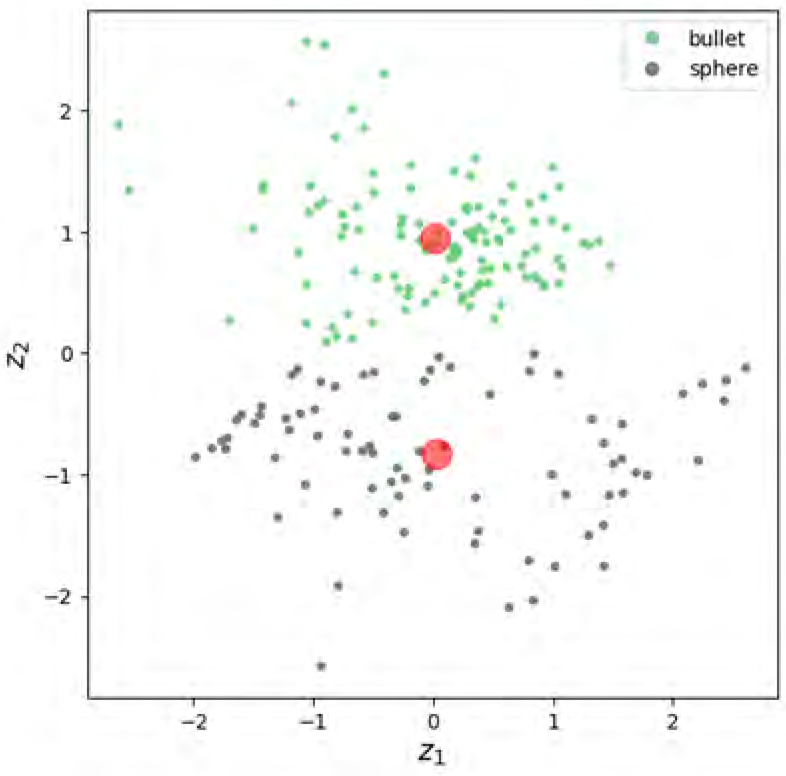

confusion matrix for k-means clustering

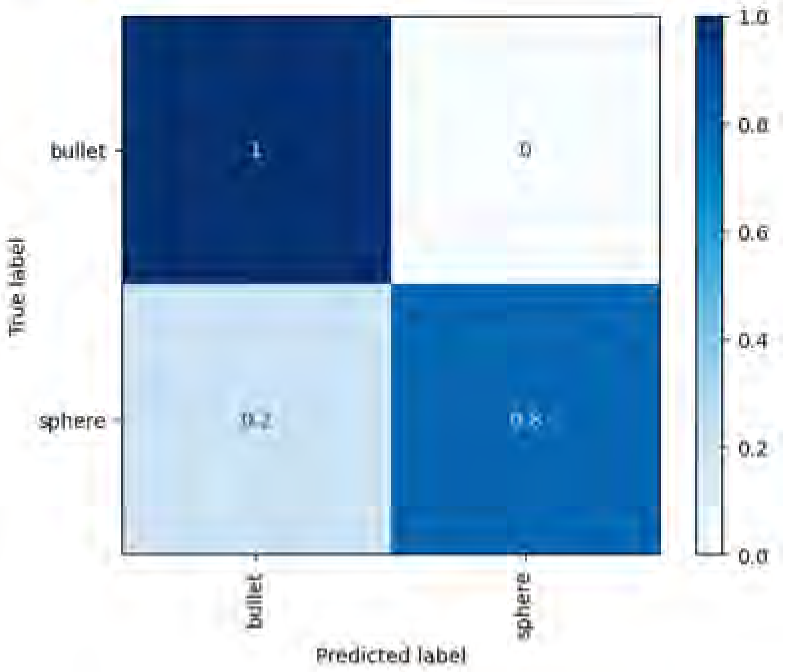

What happens with more classes?

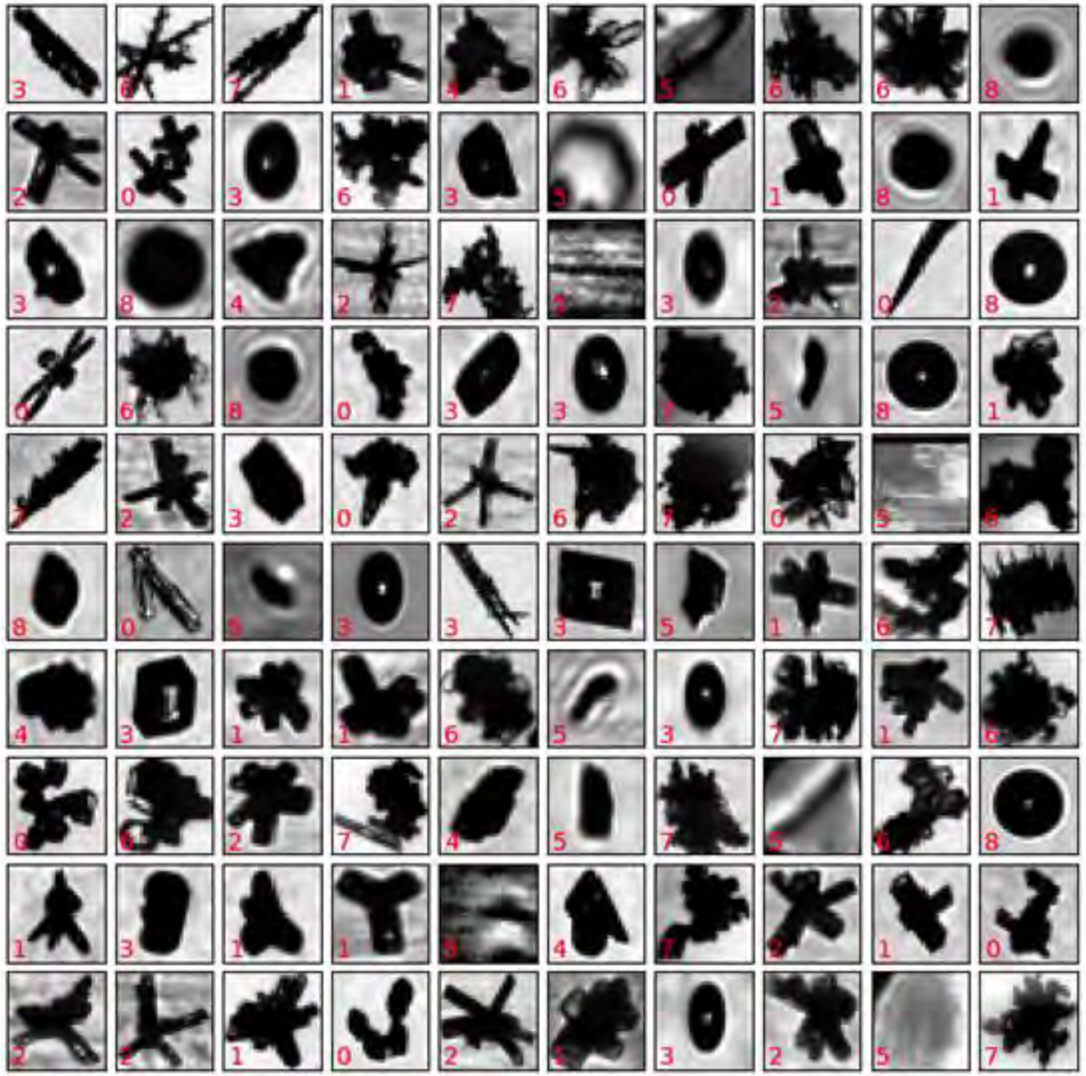

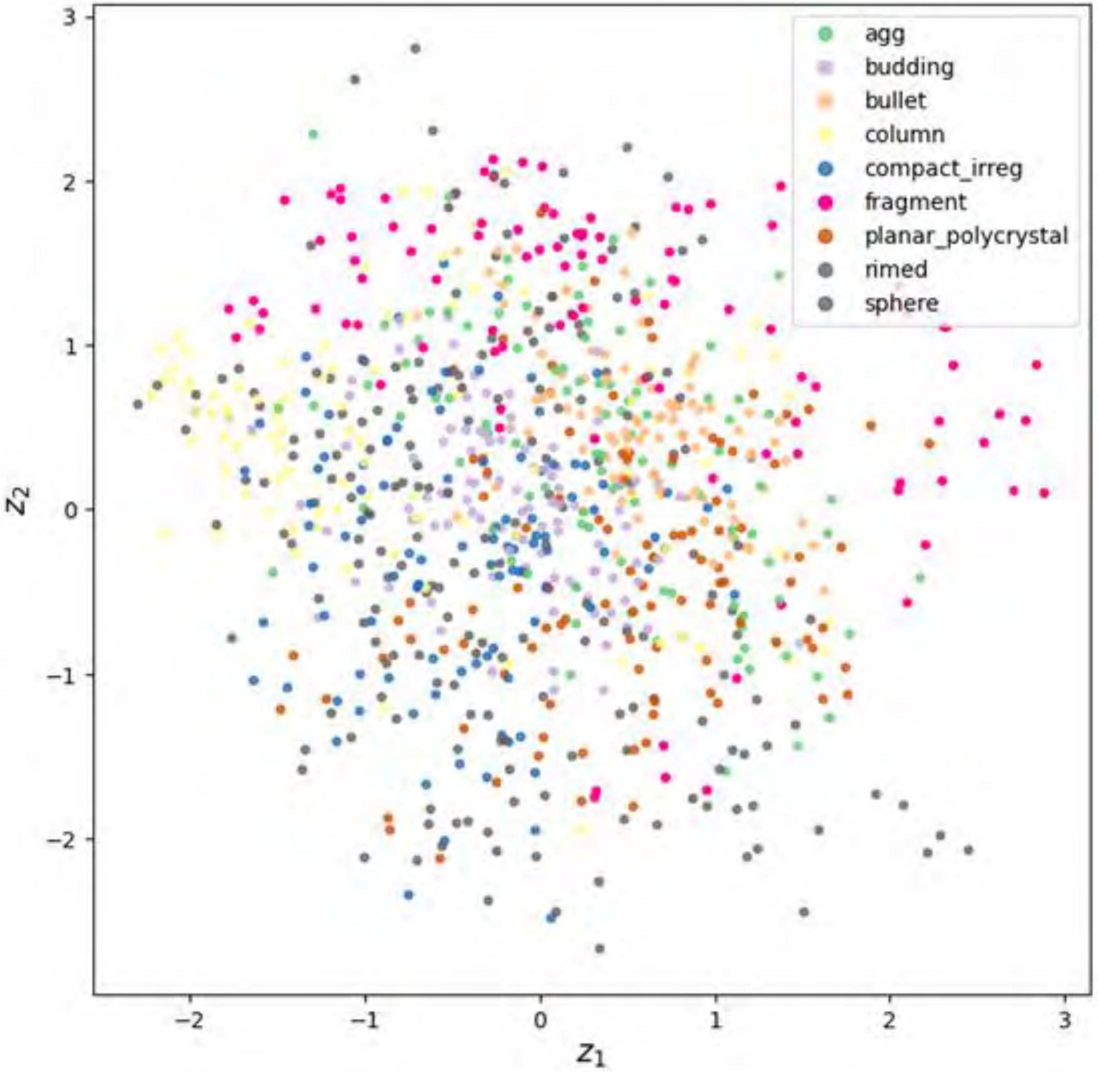

Increasing latent dimension: 2-d → 128-d

Step 1: Train VAE

latent_dim = 128

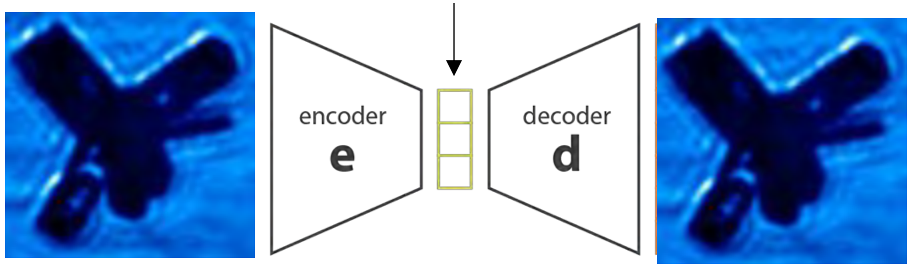

Input resolution also increased to 224x224

Step 2: Reduce Dimensionality

UMAP used to reduce 128-d latent code to 2-d

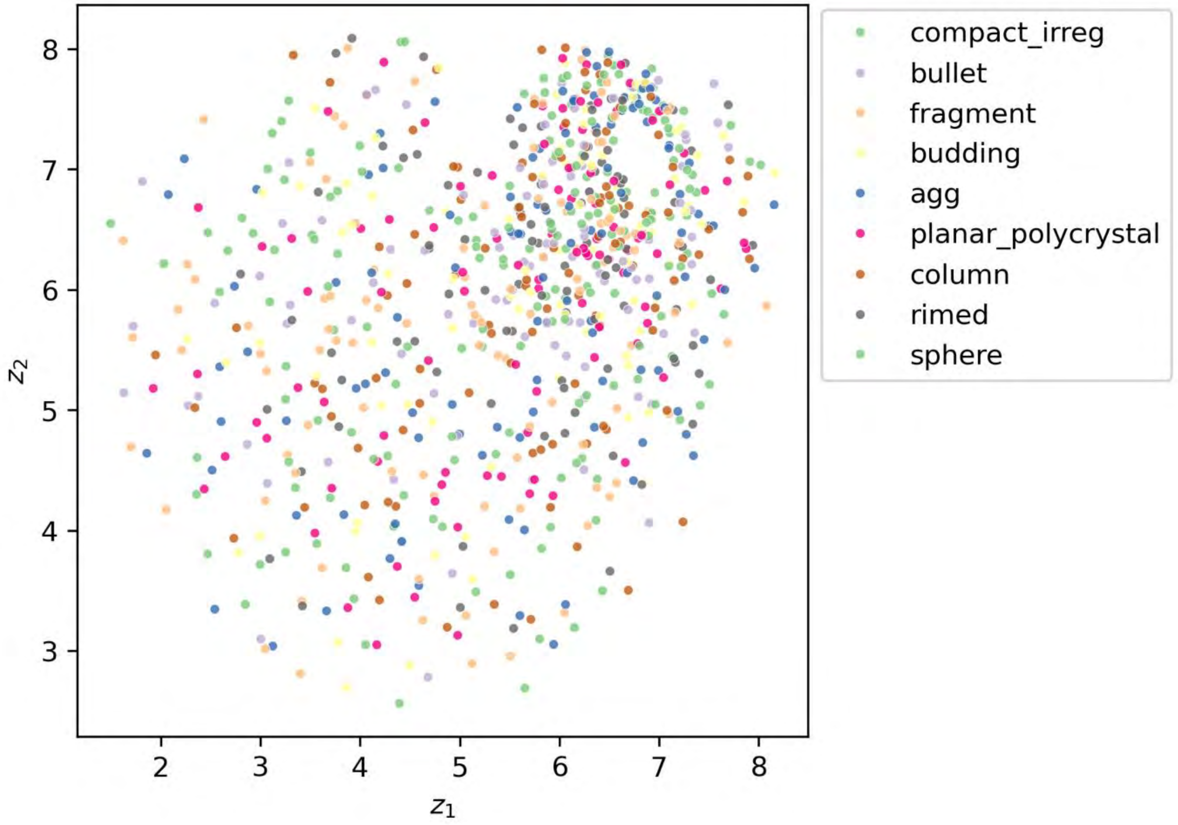

Not so great ☹

Increasing the latent dimensionality alone did not improve clustering

Using latent_dim = 2

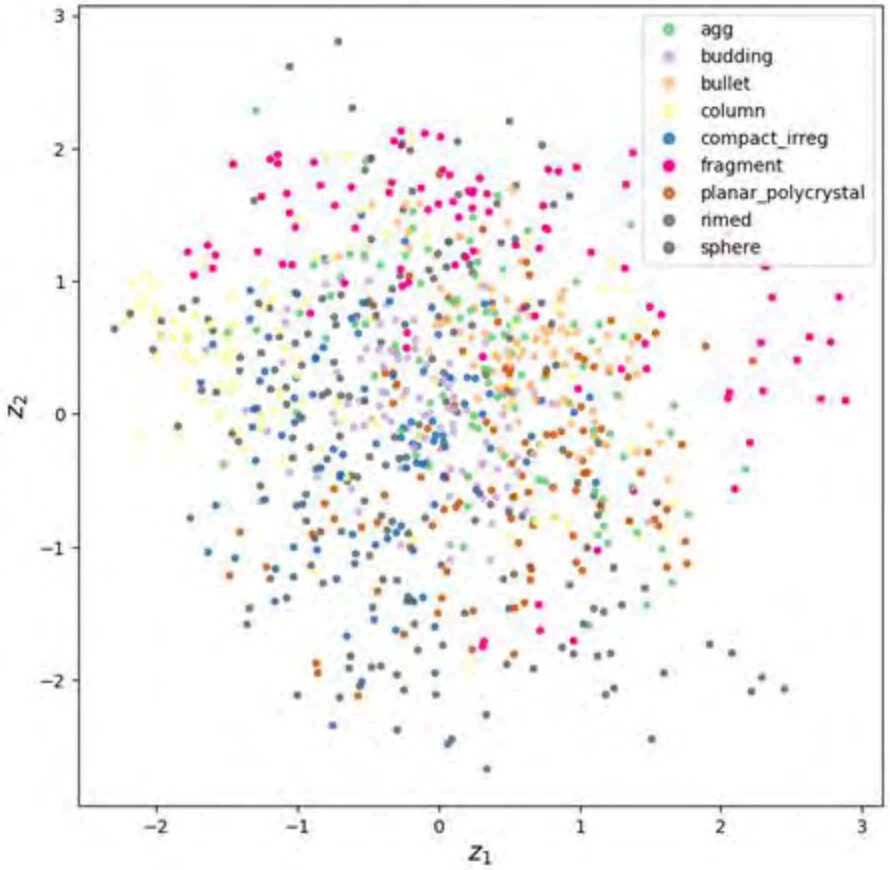

Using latent_dim = 128

+ dimensionality reduction w/ UMAP

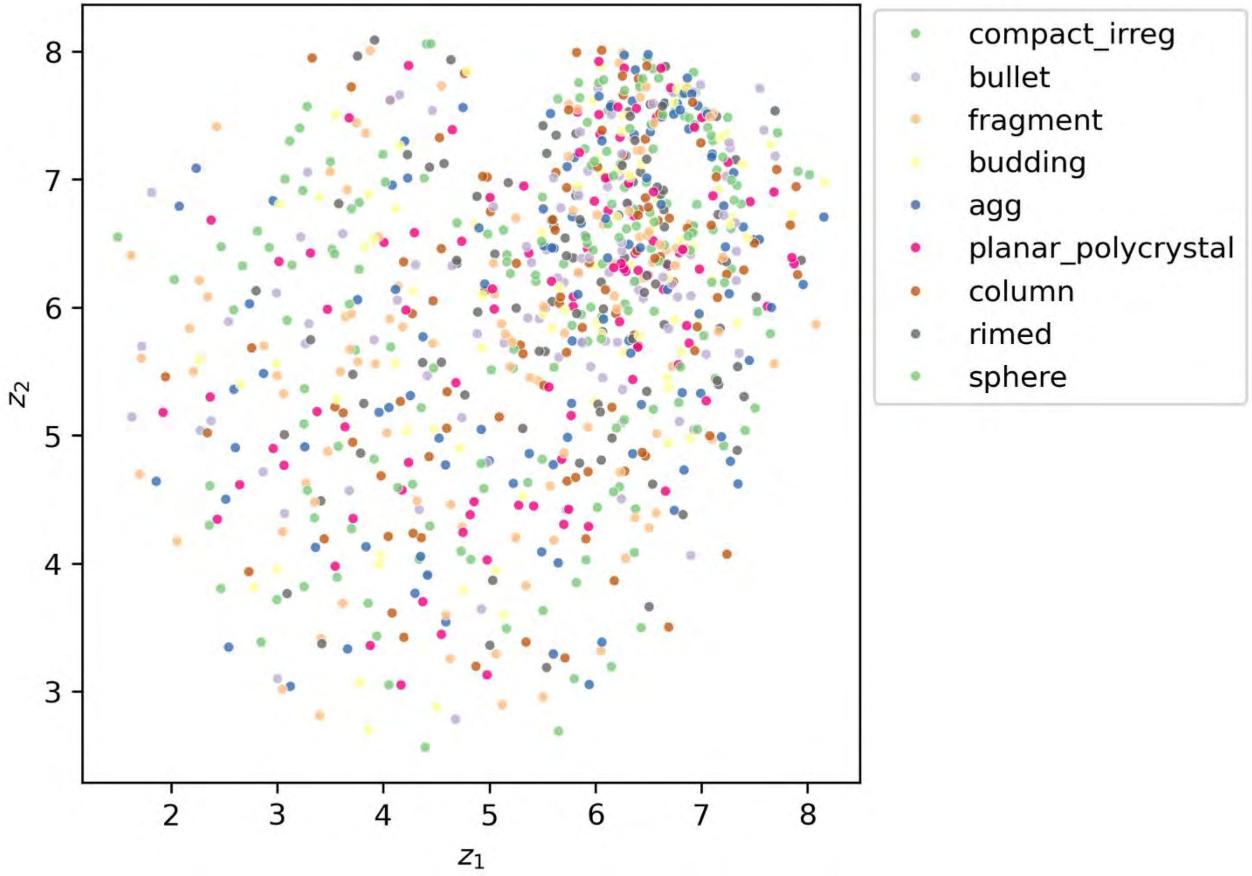

Experimenting with architecture - ResNetVAE

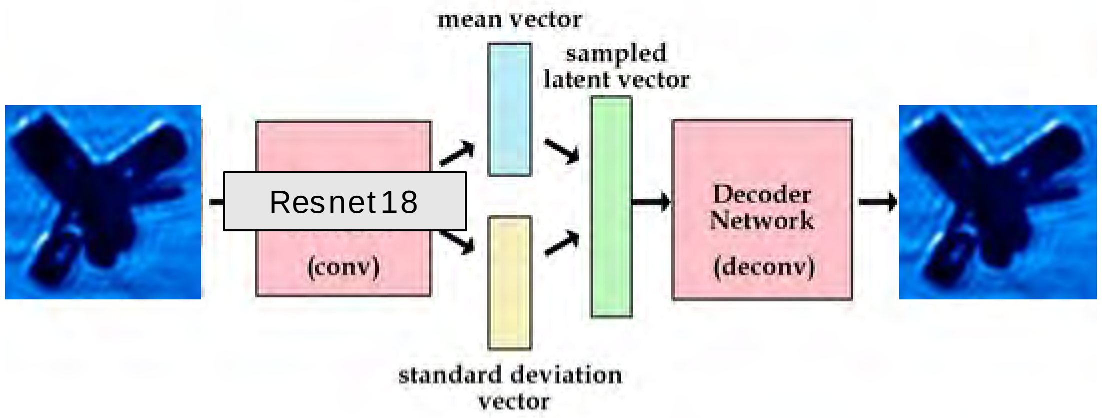

- N = 100,000 samples (9 classes)
- Input resolution: 224 x 224
- Mask applied for better geometric isolation

Experimenting with architecture - ResNetVAE

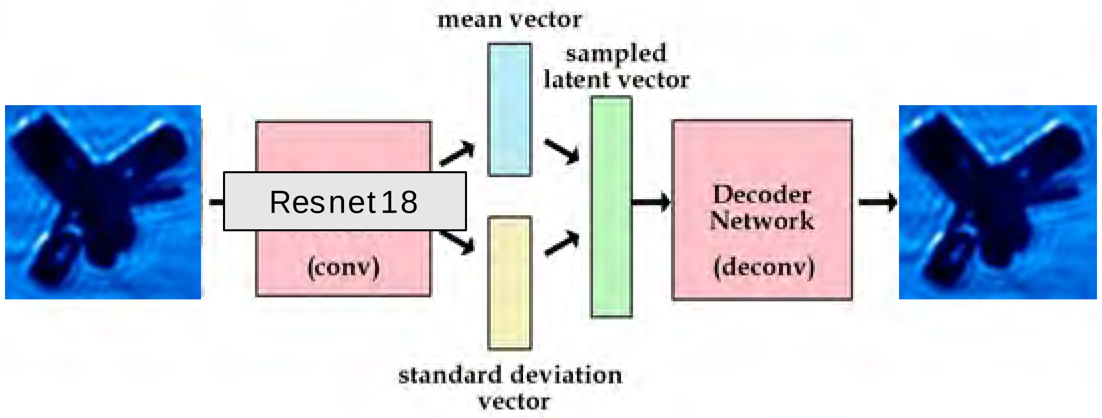

- N = 100,000 samples (9 classes)
- Input resolution: 224 x 224
- Mask applied for better shape isolation

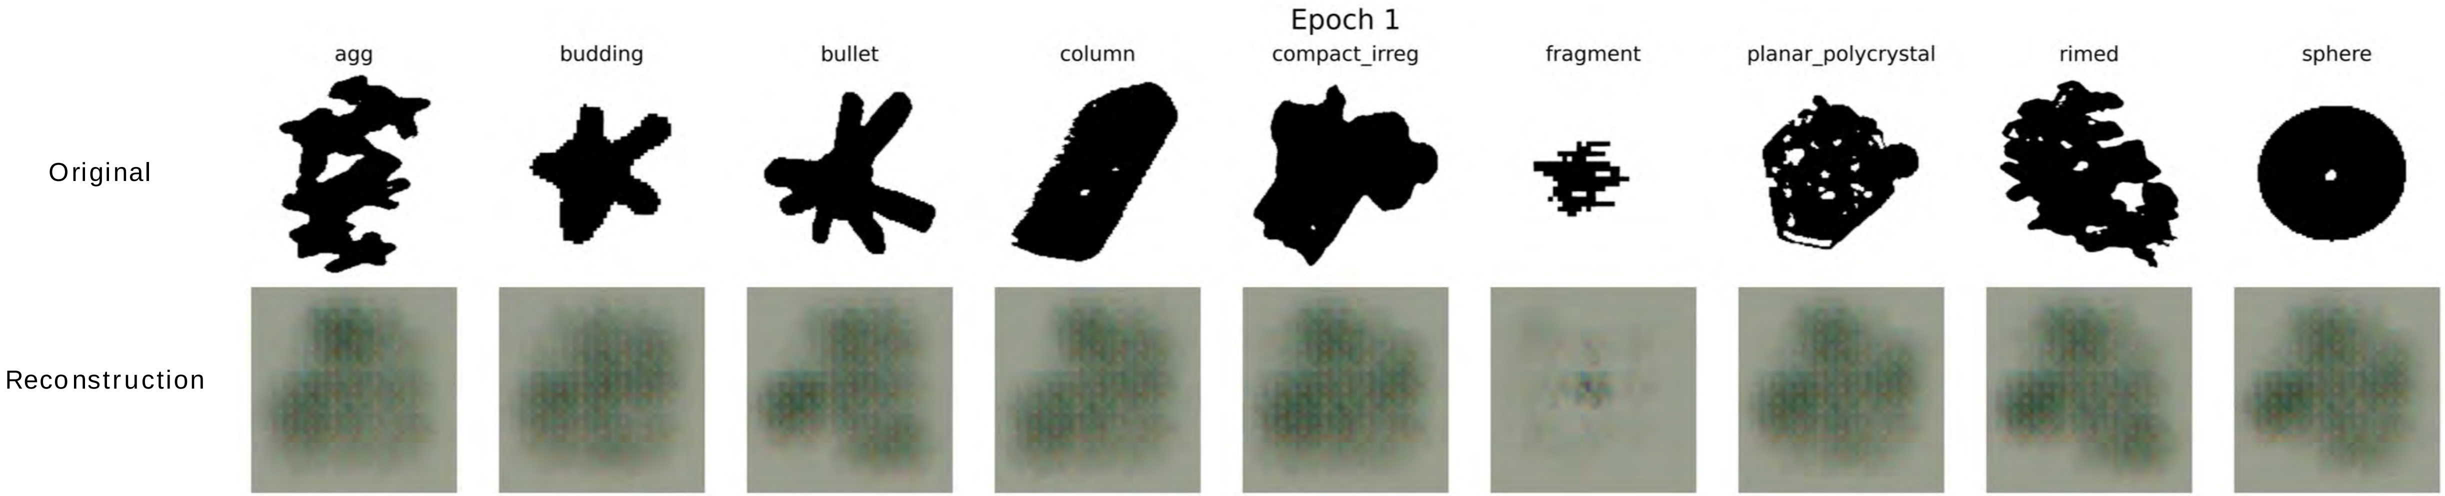

Results using a ResNetVAE

1) Train VAE with laten_dim = 128

[N ~ 100,000 samples]

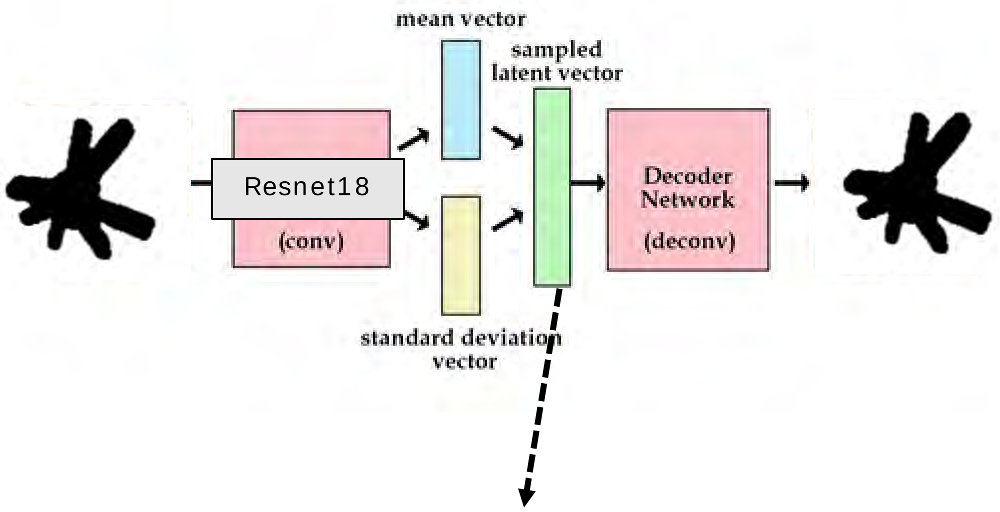

2) Reduce 128-d to 2-d with UMAP

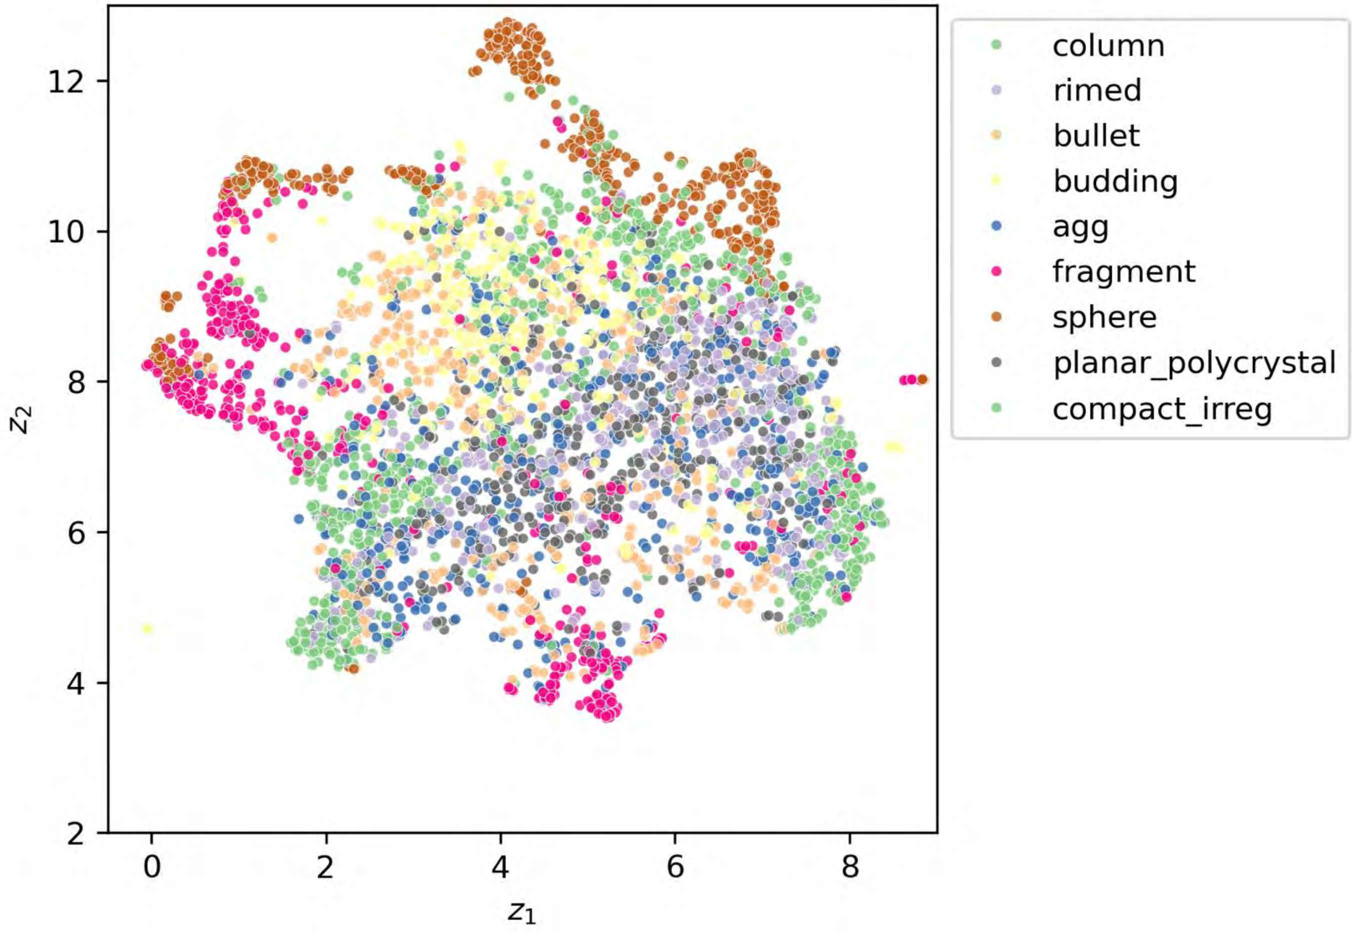

We start to see structure and disentanglement in the latent space

Default Conv-VAE
w/ latent_dims = 2

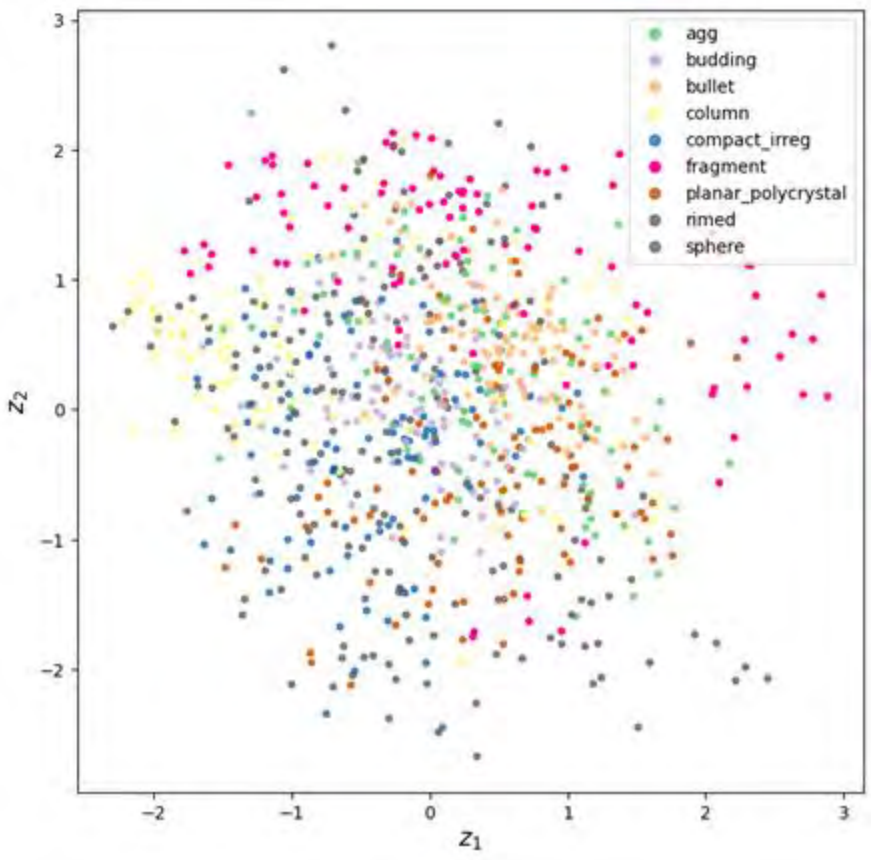

ResNetVAE w/ latent_dims = 128
(+ masked images + UMAP)

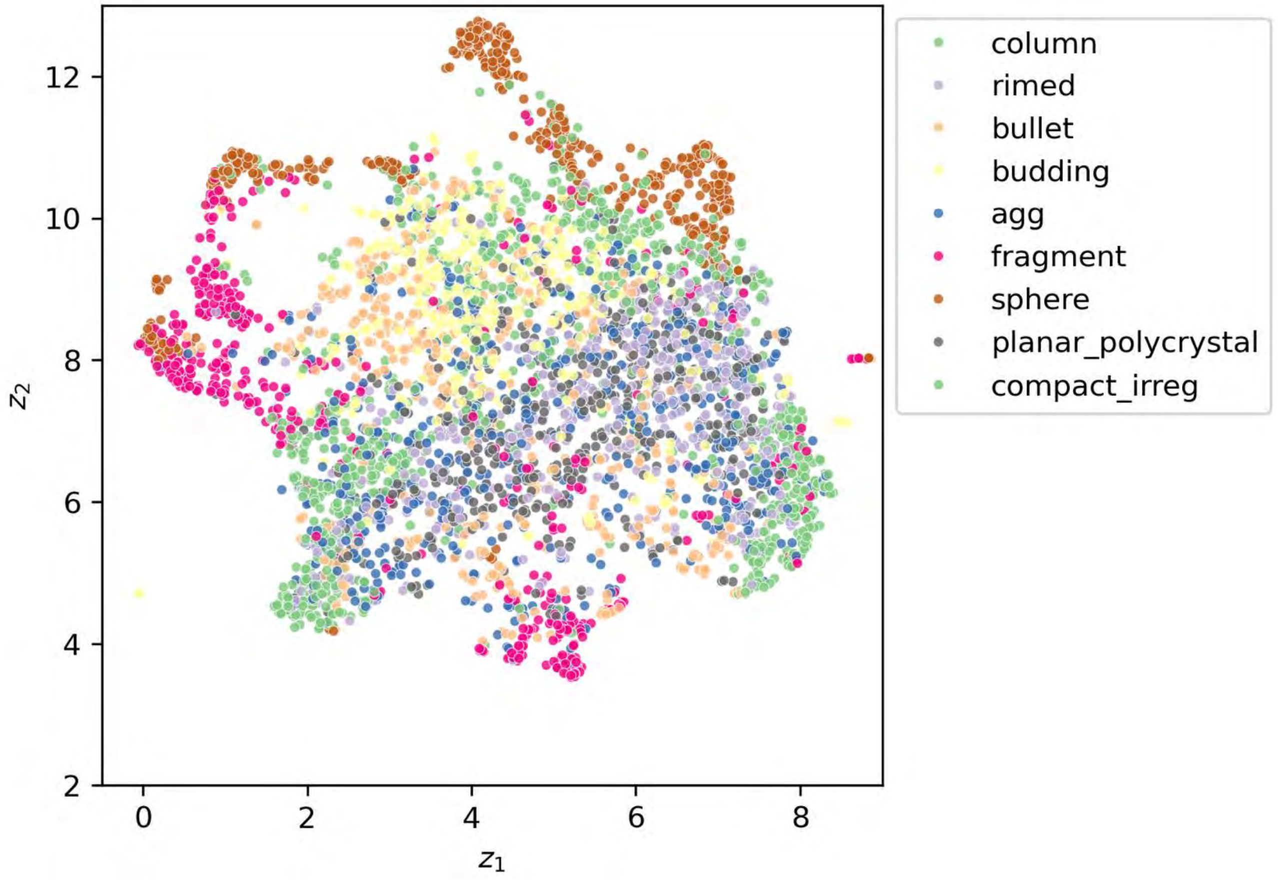

Can we use representation learning to improve/validate our synthetic data?

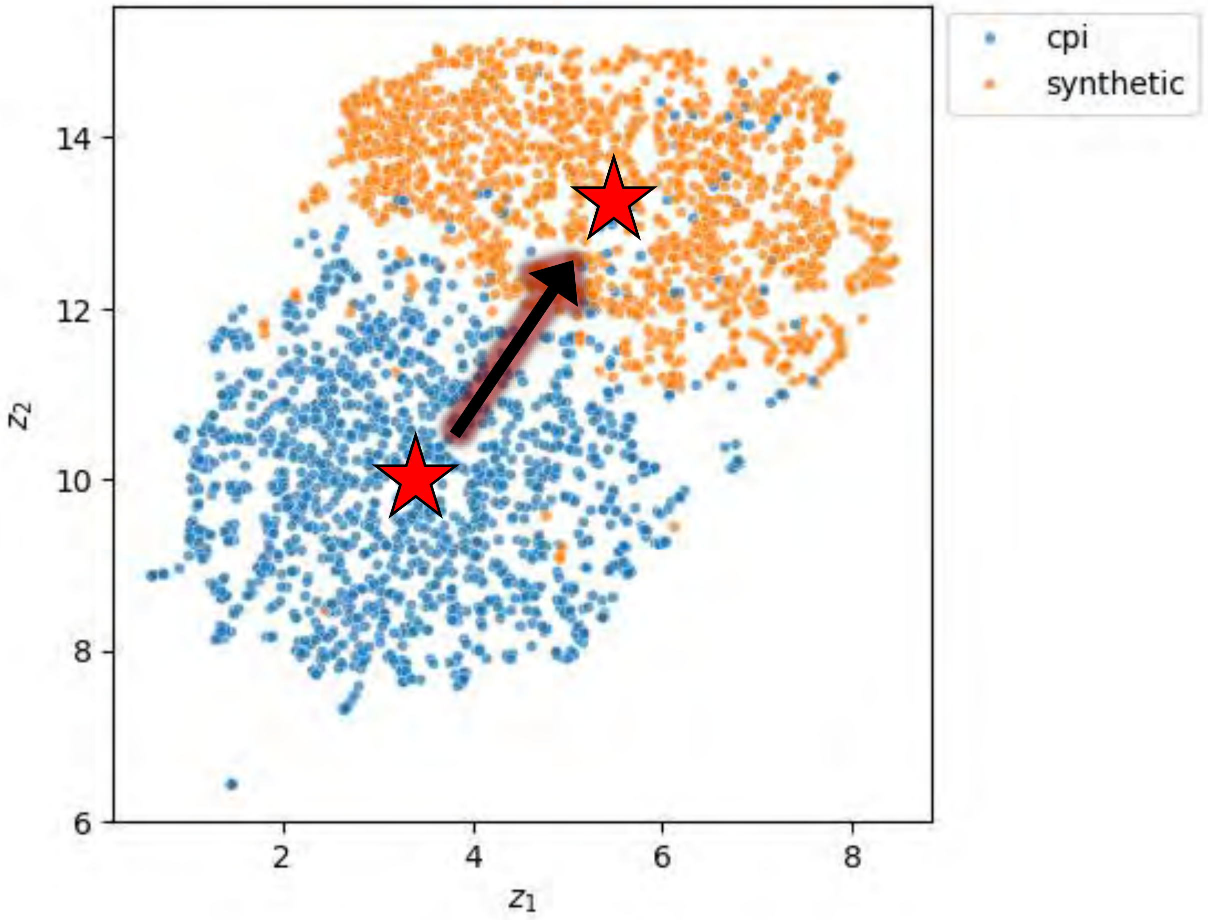

1) Train ResNetVAE with latent_dim = 128

- N_{CPI} = 1,700
- N_{synth} = 1,700

2) UMAP: 128-d → 2-d

Part 2: Key Points

- Viability of the VAE + k-means pipeline demonstrated with a 2-class proof-of-concept
- Unsupervised learning may help us understand the distribution of particle shapes at scale, without labels
- Latent representations may be useful in making synthetic crystals more realistic

Part 3: Latent metrics for LES cirrus simulations (unsupervised)

Microphysics represented in 3 different ways

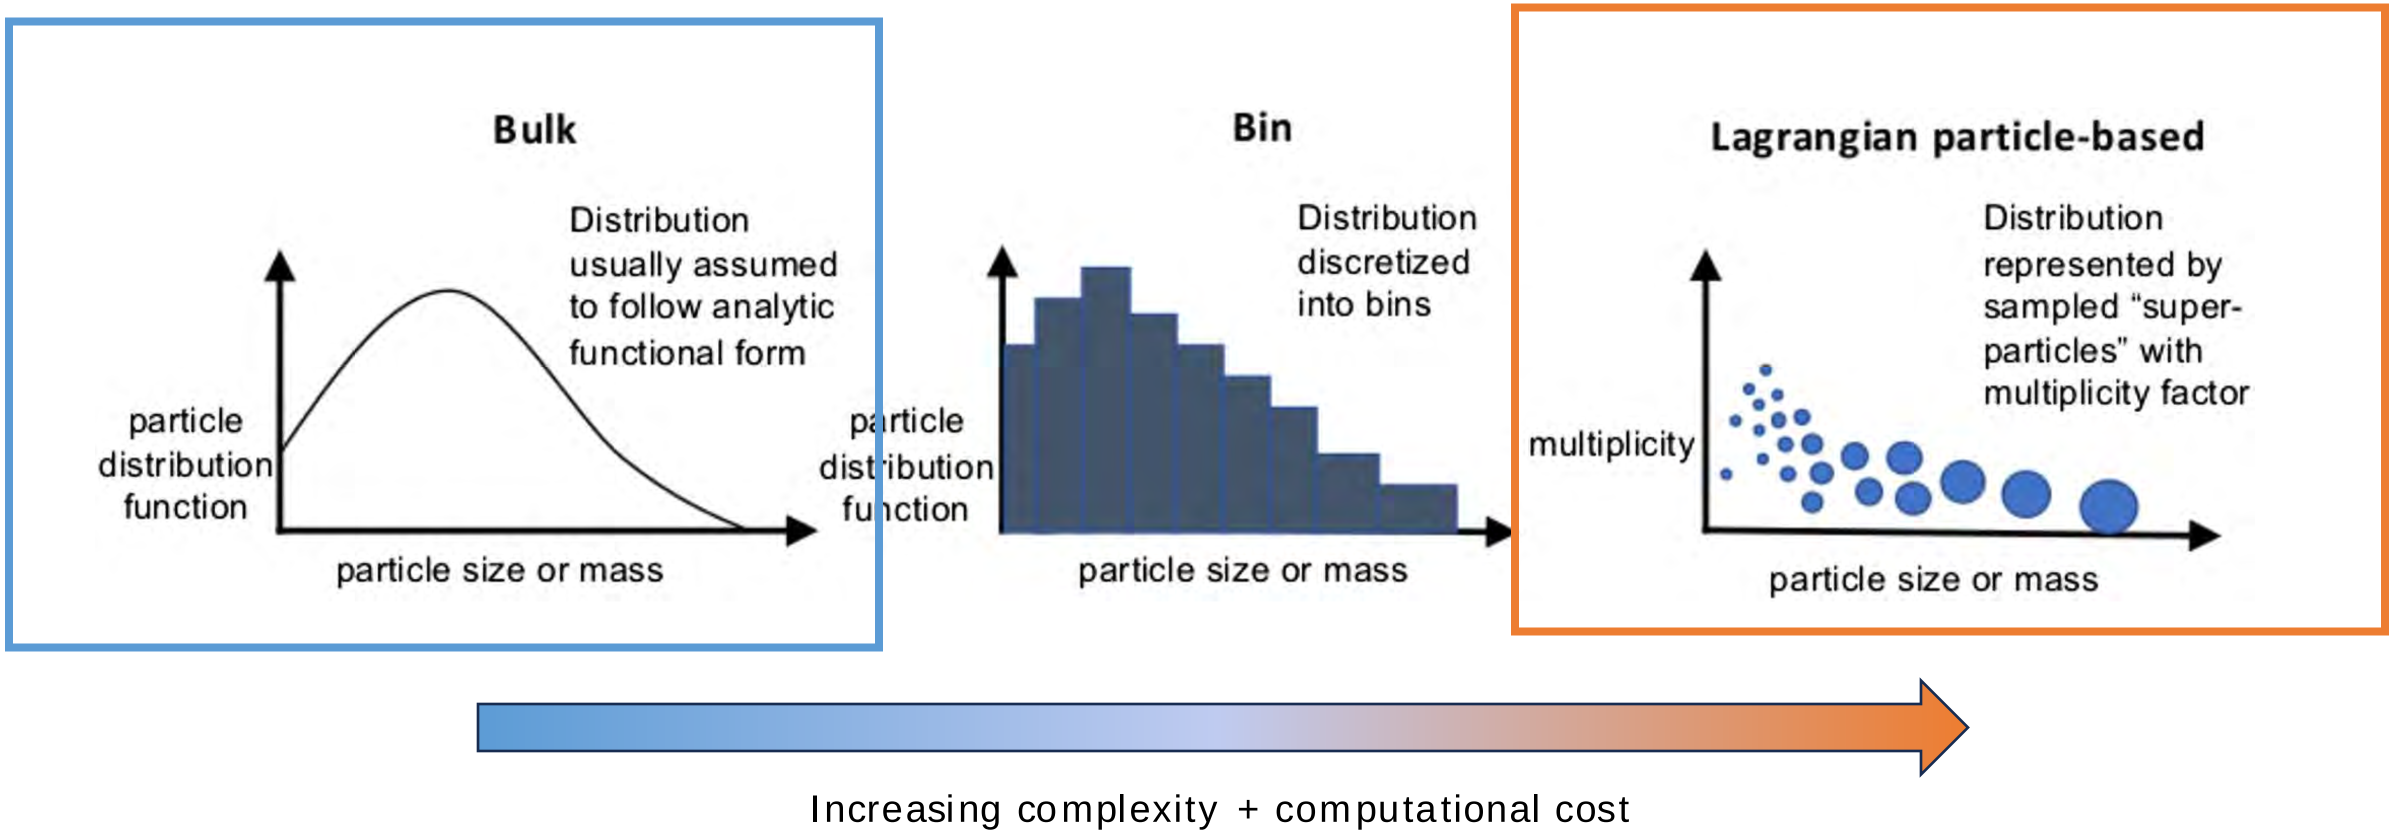

Microphysics represented in 3 different ways

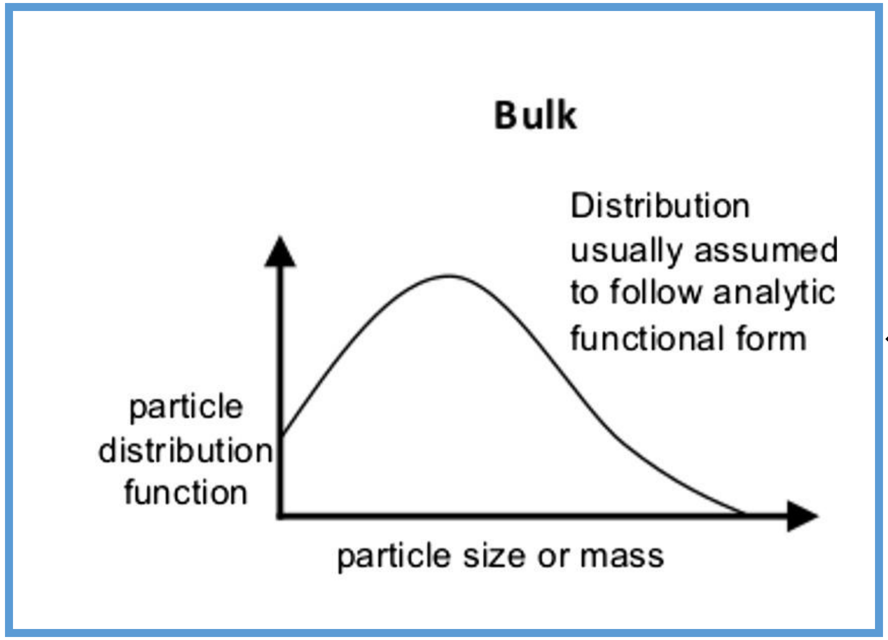

How do we compare the outputs meaningfully?

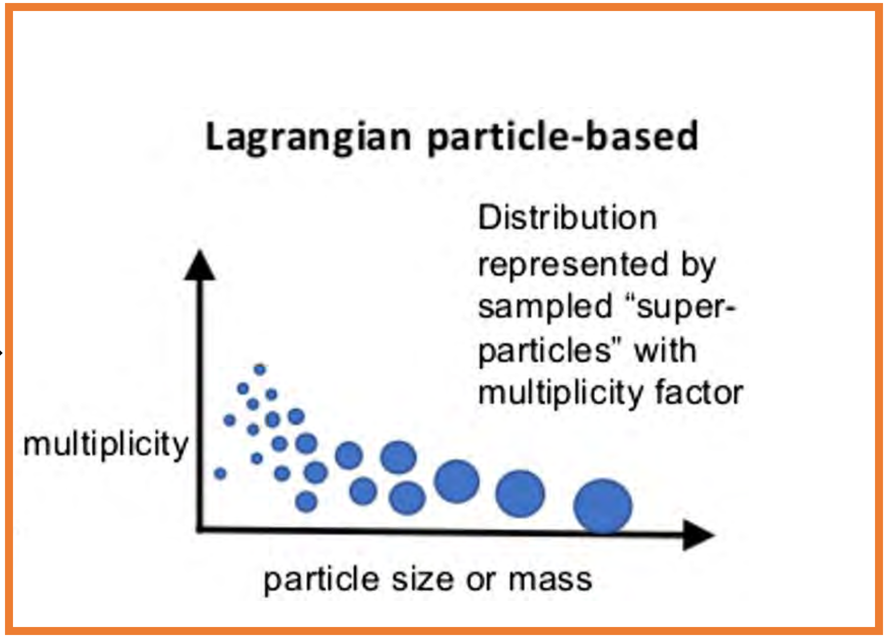

Microphysics represented in 3 different ways

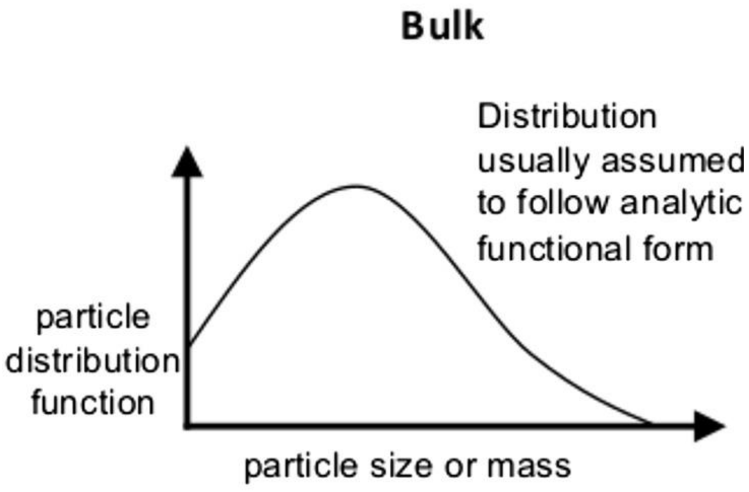

How do we compare the outputs meaningfully?

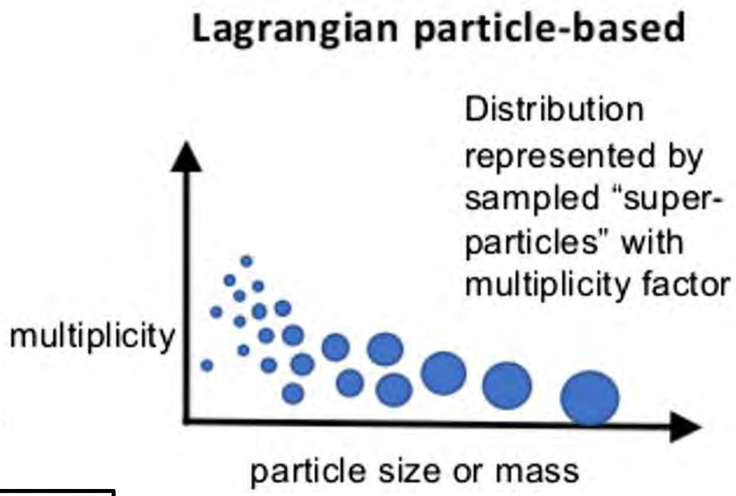

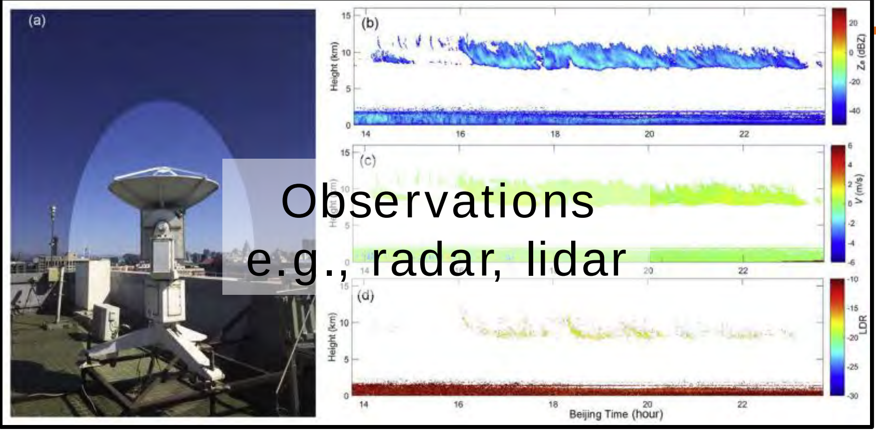

Source: Huo et al. 2020

Ensemble of bulk LES cirrus simulations

- Numerical model: NCAR's CM1
- Case study: ICEBall field campaign
  - Location: Billings, Oklahoma
- Process of interest: depositional growth
- 2 perturbed parameters (to start)
- 100 member ensemble
- 50 m resolution
- 12.8 km (W) x 12.8 km (L) x 14 km (H)

Lagrangian LES cirrus simulation

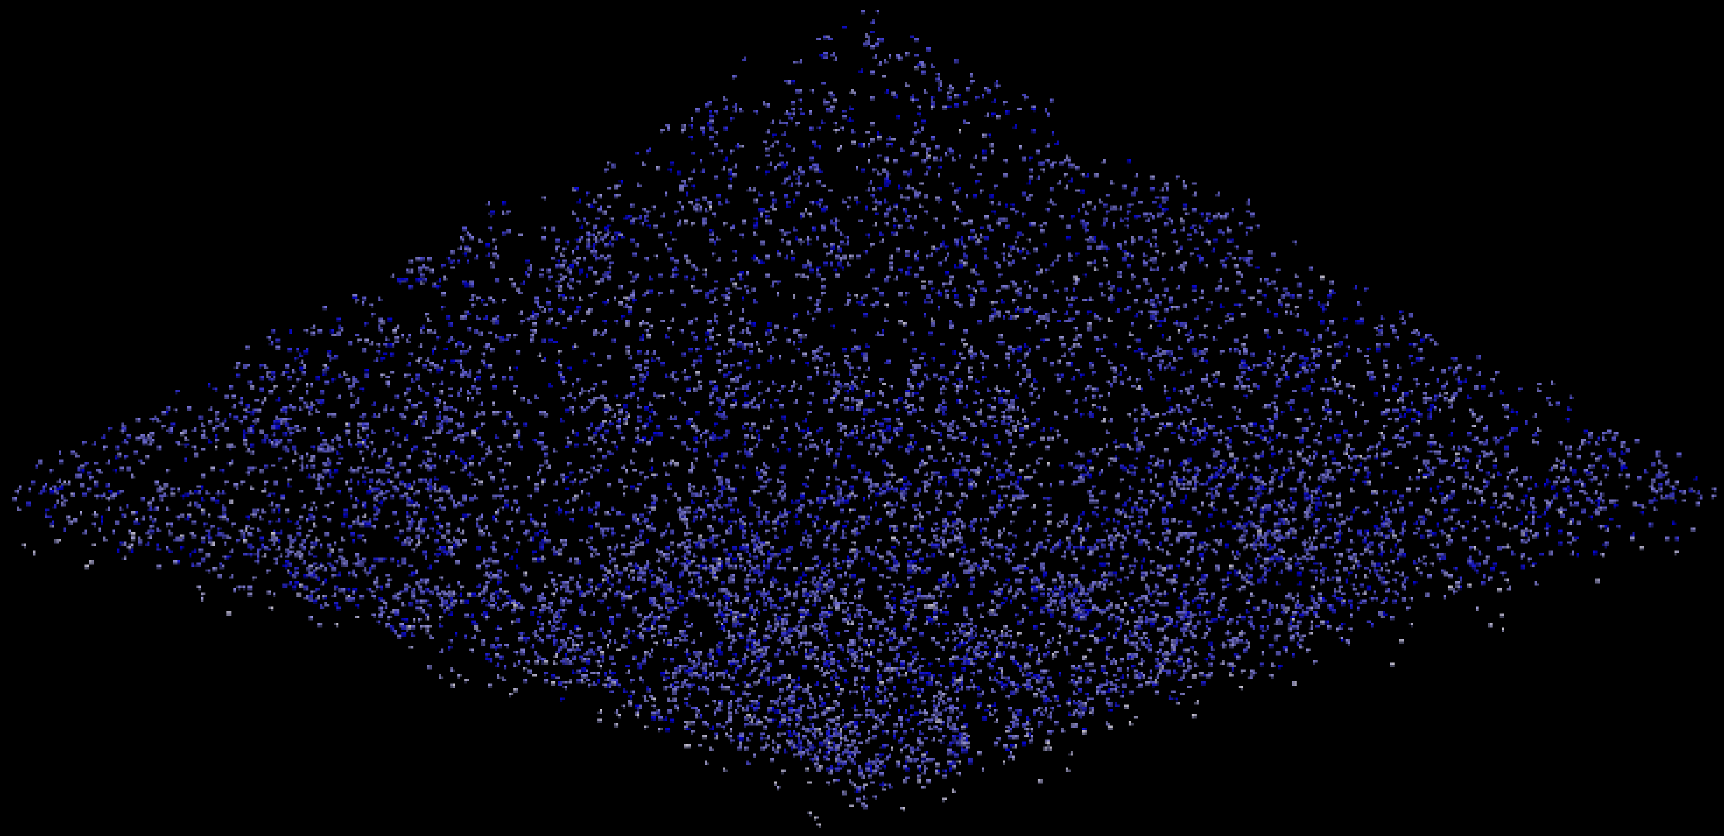

Credit: Kamal Chandrakar (sim), Obin Sturm (viz)

Which variable(s) to use for comparison?

Mass mixing ratio

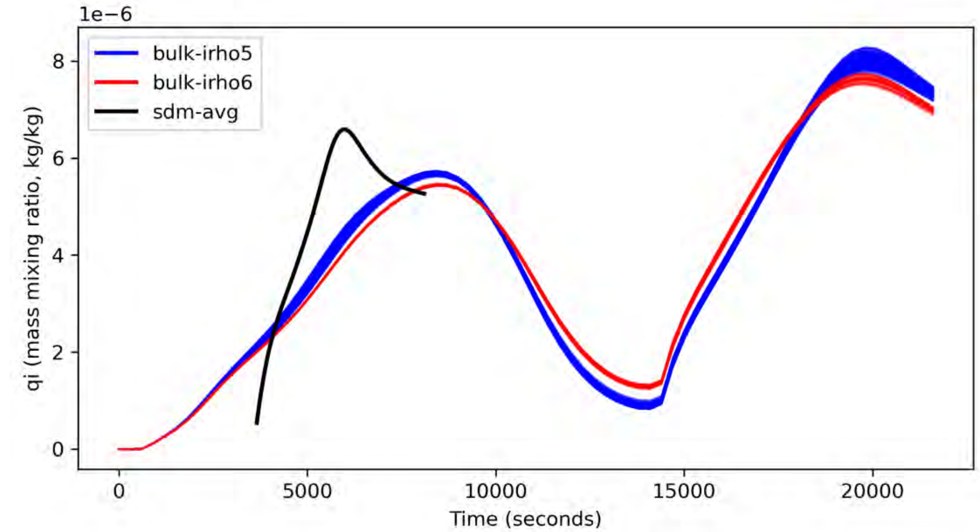

Number mixing ratio

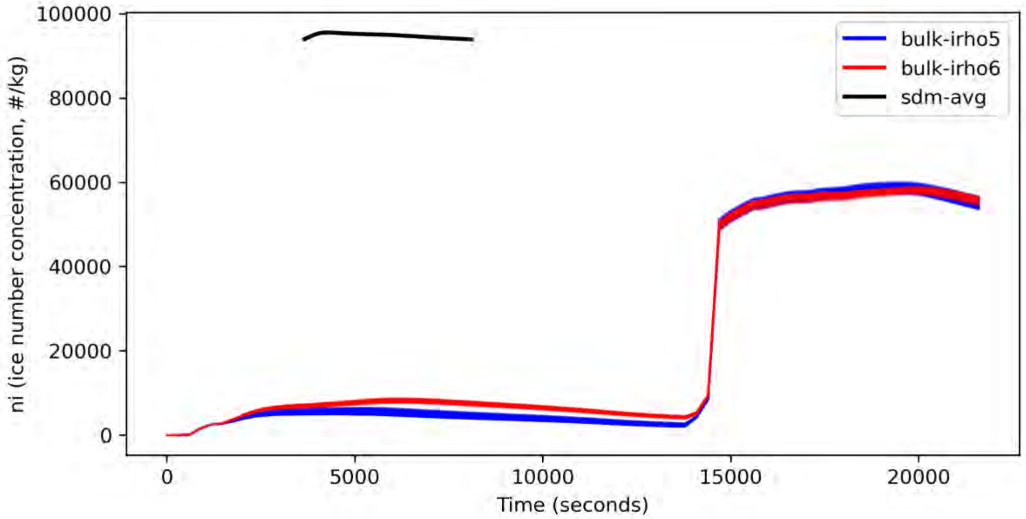

Mass-weighted fall speed

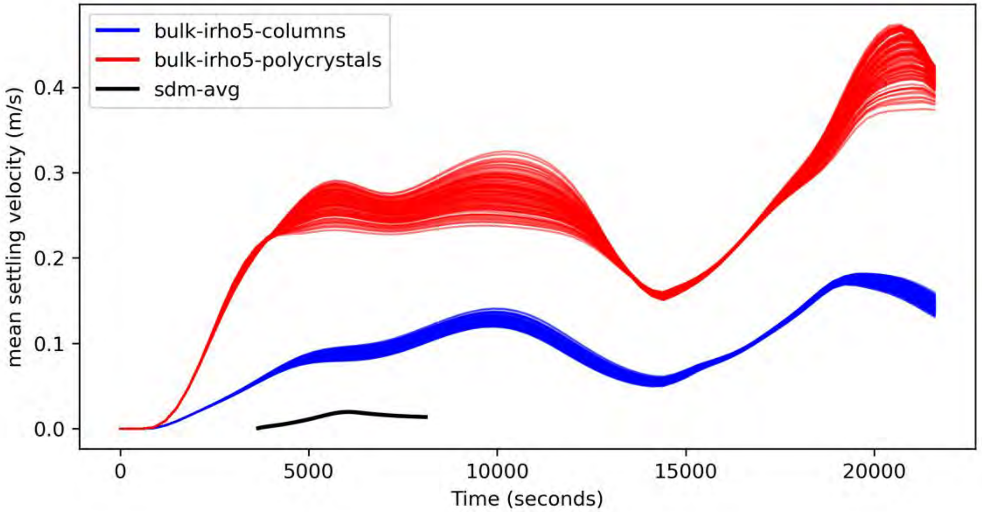

Mass-weighted particle density

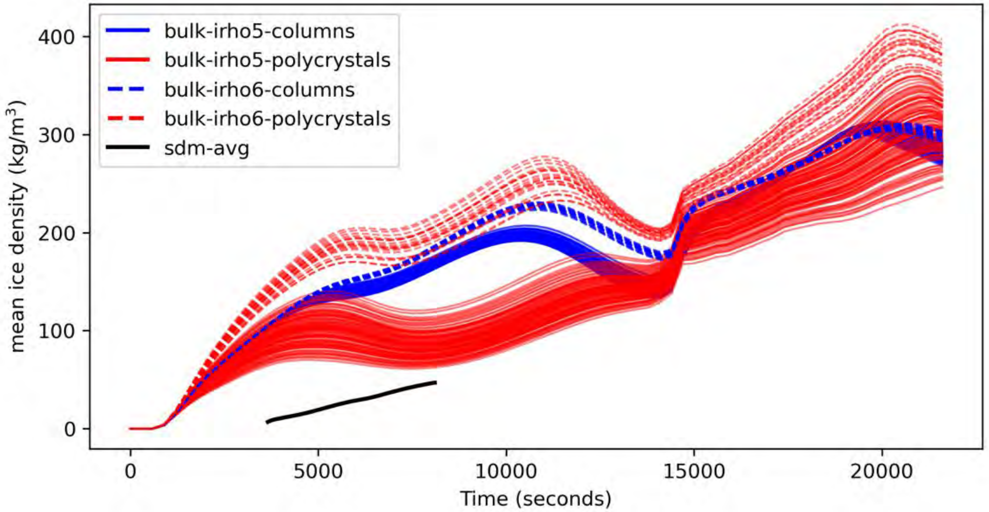

There are large structural differences between bulk and Lagrangian microphysics!

Mass mixing ratio

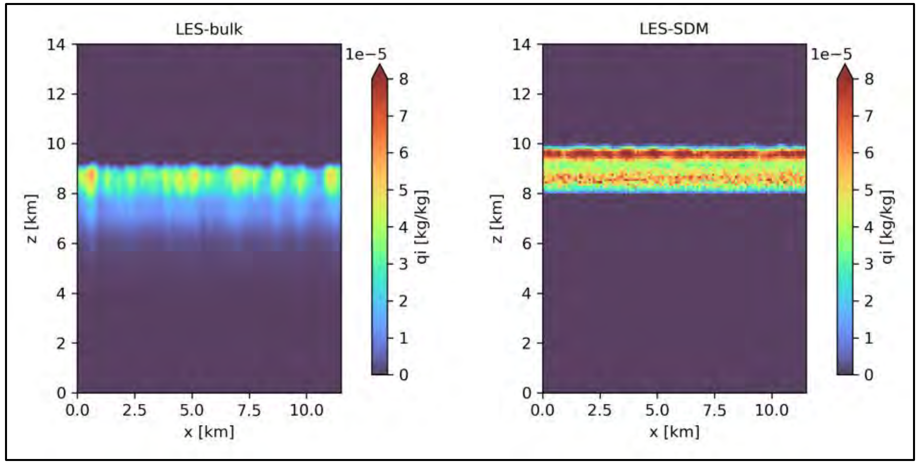

Number mixing ratio

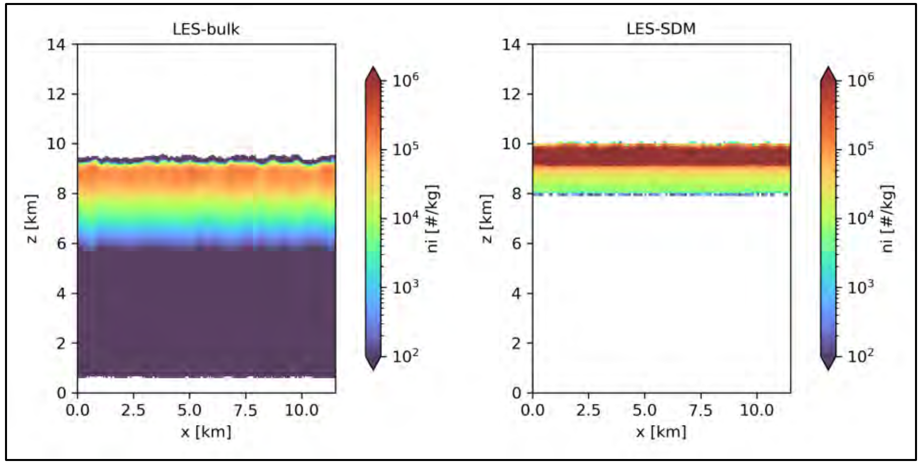

Mass-weighted fall speed

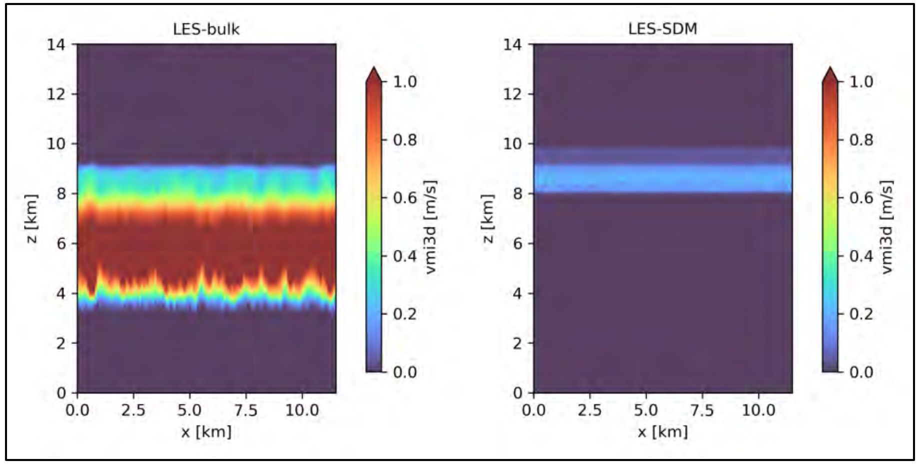

Mass-weighted particle density

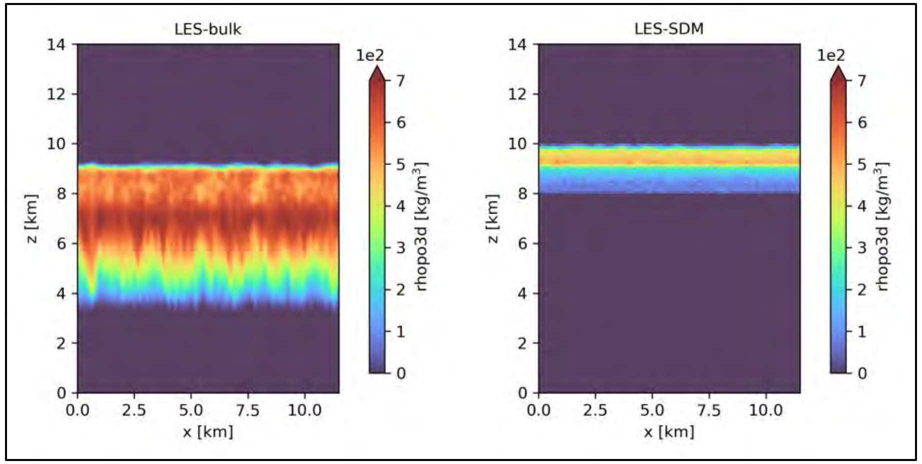

Is there a way to compare simulations in a compressed latent space?

input dim: (3x64x64)

output dim: (3x64x64)

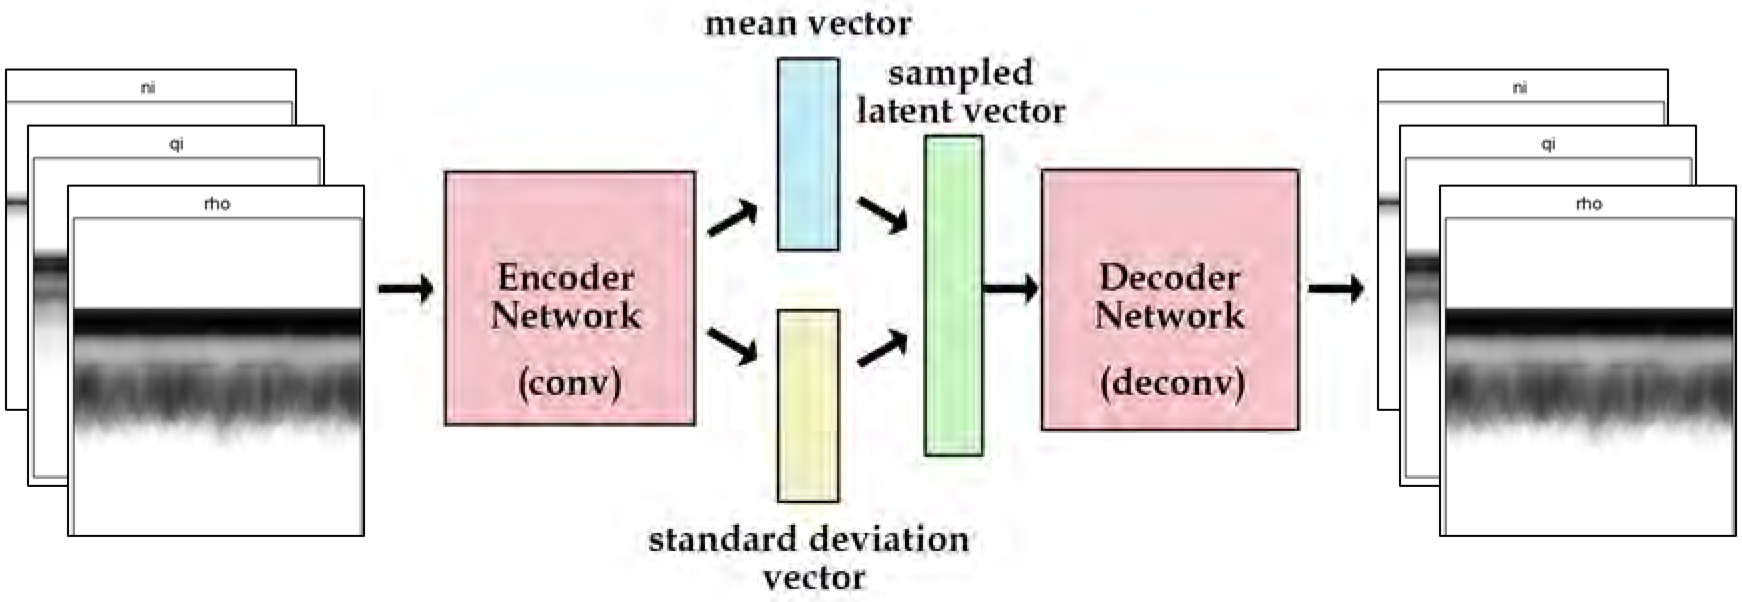

3-channel, convolutional VAE

- 3 hidden layers
- 16 latent dimensions

Four examples: reconstructions of three 2-D variable fields

Top:
Actual

Bottom:
Reconstruction

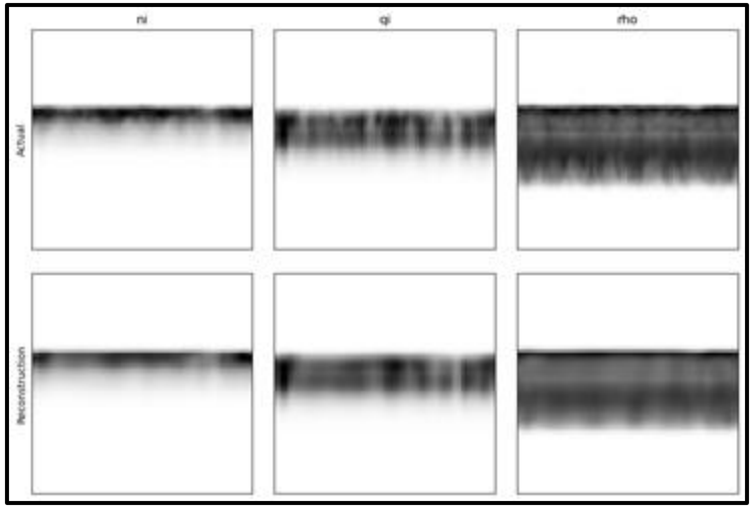

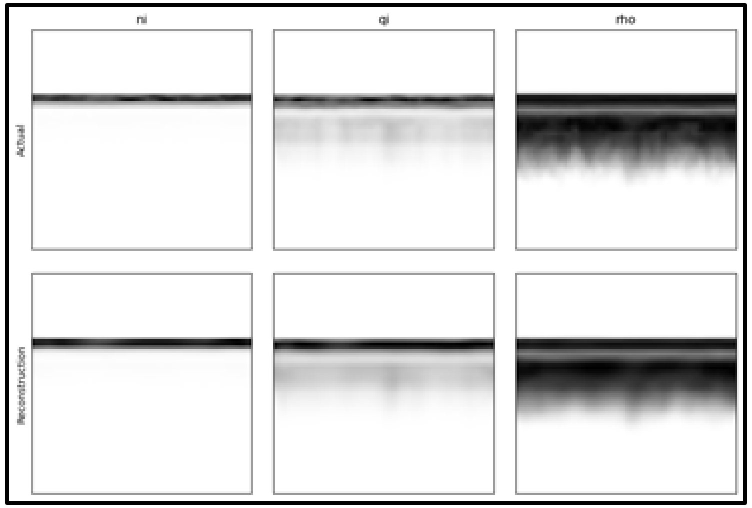

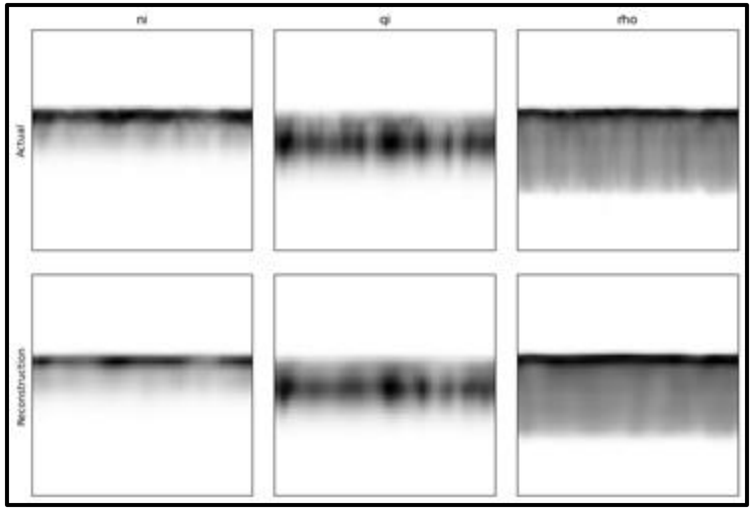

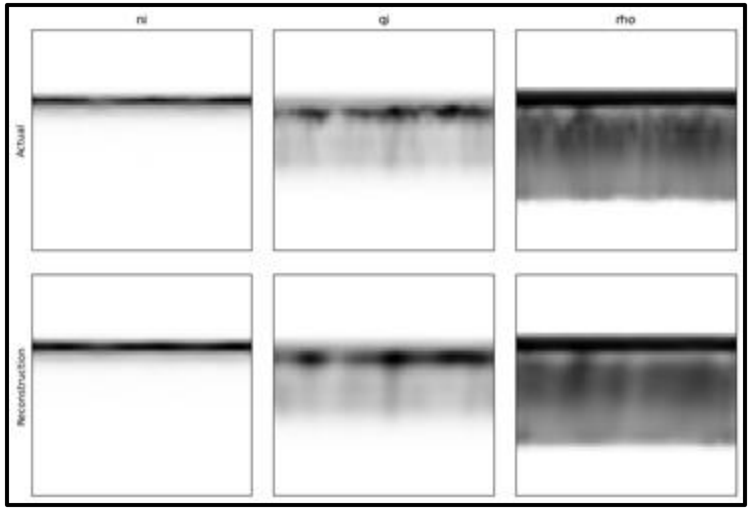

2-D projections of the 16-D latent manifold

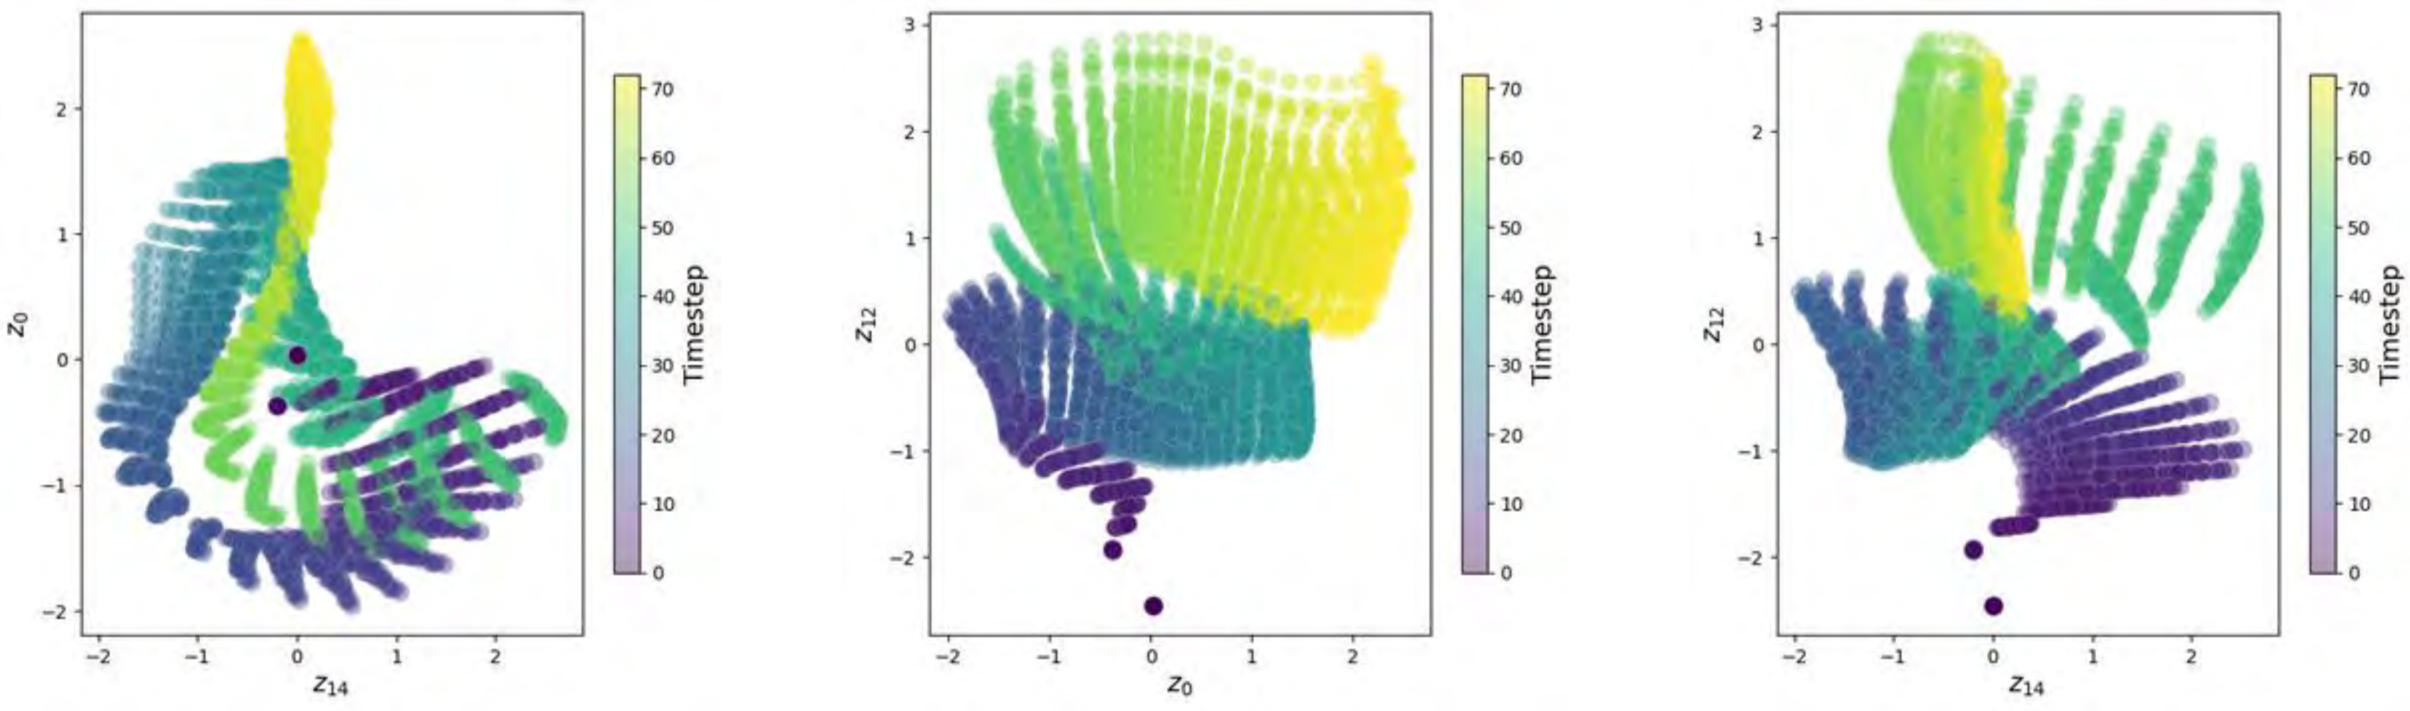

Note: z0, z12, z14 are the 3 latent dimensions with highest variance

The time evolution of 3 variables represented as a 'smooth' spiral manifold

Pushing it to the extreme: reduce latent space to 2-D

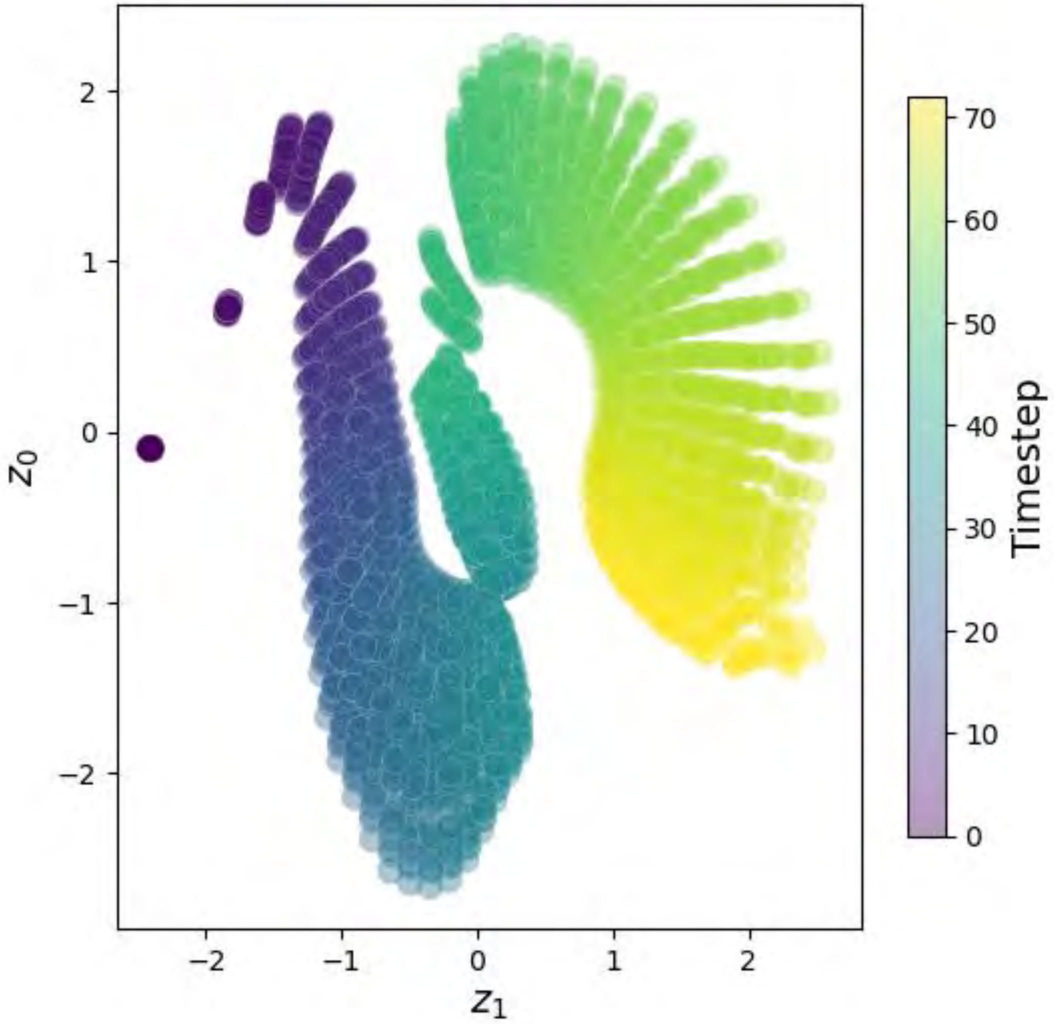

- Smoothly varying manifold in 2 latent dimensions
- Captures most of variability of three microphysical variables (qi, mi, rho) with only 2 variables!
- Variance along manifold trajectory shows when parameters affect variables of interest the most

Conclusion

Part 1: Predicting 3-D properties from images (supervised)

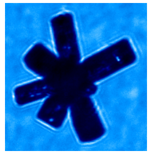

- An ML framework was developed to predict 3-D crystal properties of interest from 2-D imagery, using synthetic crystals

Part 2: Latent representations of ice crystals (unsupervised)

- Using a VAE, a ‘structured’ and ‘disentangled’ latent representation of crystal shapes was found

Part 3: Latent metrics for LES cirrus simulations (unsupervised)

- Latent variables may be a promising avenue to compare disparate simulation outputs in a more objective manner

Conclusion

Part 1: Predicting 3-D properties from images (supervised)

- An ML framework was developed to predict 3-D crystal properties of interest from 2-D imagery, using synthetic crystals

Part 2: Latent representations of ice crystals (unsupervised)

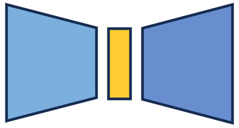

- Using a VAE, a 'structured' and 'disentangled' latent representation of crystal shapes was found

Part 3: Latent metrics for LES cirrus simulations (unsupervised)

- Latent variables may be a promising avenue to compare disparate simulation outputs in a more objective manner

Conclusion

Part 1: Predicting 3-D properties from images (supervised)

- An ML framework was developed to predict 3-D crystal properties of interest from 2-D imagery, using synthetic crystals

Part 2: Latent representations of ice crystals (unsupervised)

- Using a VAE, a ‘structured’ and ‘disentangled’ latent representation of crystal shapes was found

Part 3: Latent metrics for LES cirrus simulations (unsupervised) 

- Latent variables may be a promising avenue to compare disparate simulation outputs in a more objective manner

Conclusion

Part 1: Predicting 3-D properties from images (supervised)

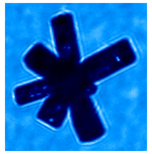

- An ML framework was developed to predict 3-D crystal properties of interest from 2-D imagery, using synthetic crystals

Part 2: Latent representations of ice crystals (unsupervised)

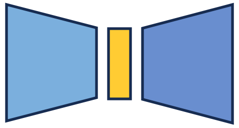

- Using a VAE, a ‘structured’ and ‘disentangled’ latent representation of crystal shapes was found

Part 3: Latent metrics for LES cirrus simulations (unsupervised)

- Latent variables may be a promising avenue to compare disparate simulation outputs in a more objective manner

Thanks!

Contact: jk4730@columbia.edu

Advisors, collaborators & help:

Kara Lamb (Columbia)
Marcus van Lier-Walqui (Columbia, NASA GISS)
Jerry Harrington (Penn State)
Kamal Chandrakar (NCAR)
Hugh Morrison (NCAR)
Kaitlyn Loftus (Columbia)
Kara Sulia (U. Albany)
Vanessa Przybylo (formerly U. Albany)
Obin Sturm (University of Southern California)
...et al.

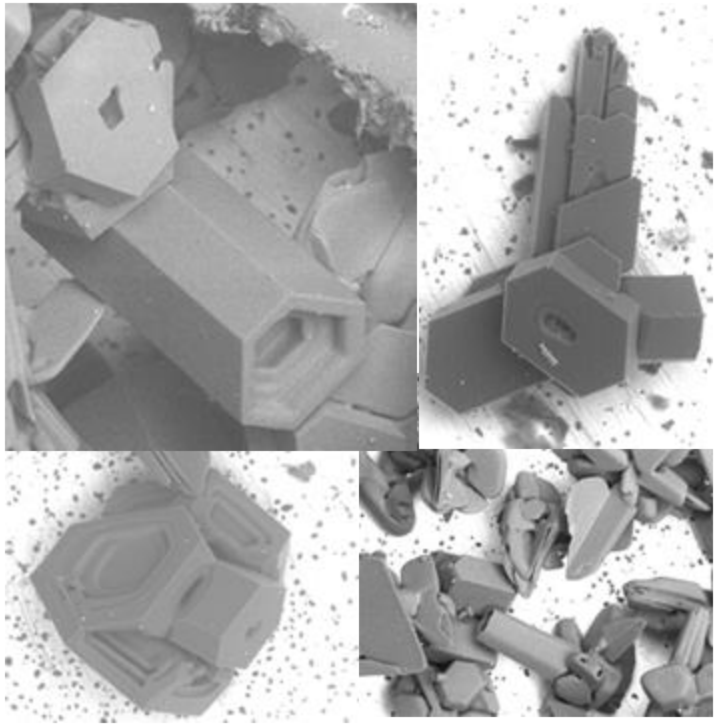

Images from Nathan Magee

Funding: LEAP NSF U.S. DEPARTMENT OF ENERGY

Conclusion

Part 1: Predicting 3-D properties from images (supervised)

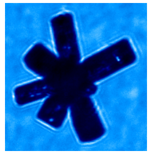

- An ML framework was developed to predict 3-D crystal properties of interest from 2-D imagery, using synthetic crystals

Part 2: Latent representations of ice crystals (unsupervised)

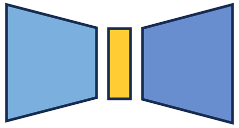

- Using a VAE, a ‘structured’ and ‘disentangled’ latent representation of crystal shapes was found

Part 3: Latent metrics for LES cirrus simulations (unsupervised)

- Latent variables may be a promising avenue to compare disparate simulation outputs in a more objective manner