Fermilab

EuCAIFCon, June 2025

AI in Astrophysics:

Tackling Domain Shift, Model Robustness and Uncertainty

Aleksandra Ciprijanvoic

Associate Scientist, Fermilab
aleksand@fnal.gov

U.S. DEPARTMENT of ENERGY

Fermi National Accelerator Laboratory is managed by
FermiForward for the U.S. Department of Energy Office of Science

FERMI NATIONAL ACCELERATOR LABORATORY

AI in Astrophysics is booming

- Finding and characterizing astrophysical objects
- Anomaly searches and alert systems
- Telescope scheduling and design

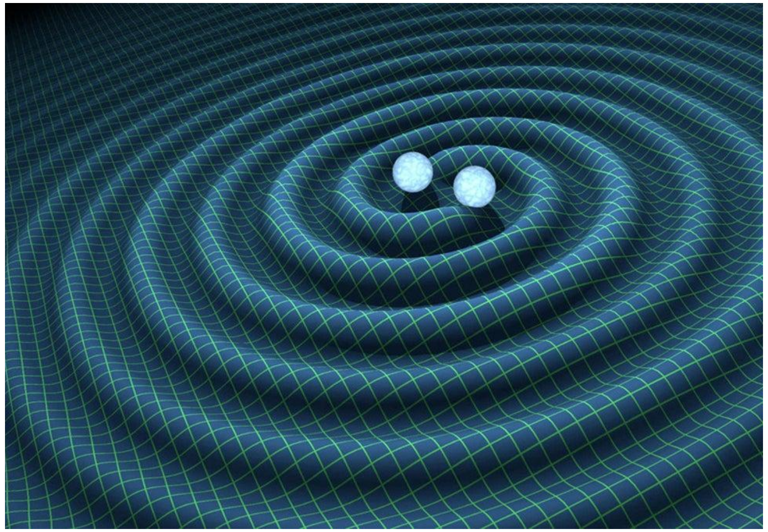

AI in Astrophysics is booming

- Finding and characterizing astrophysical objects
- Anomaly searches and alert systems
- Telescope scheduling and design

- **Foundation models, multimodal learning**
- **Large Language Models (LLMs), better search systems, hypothesis generation**

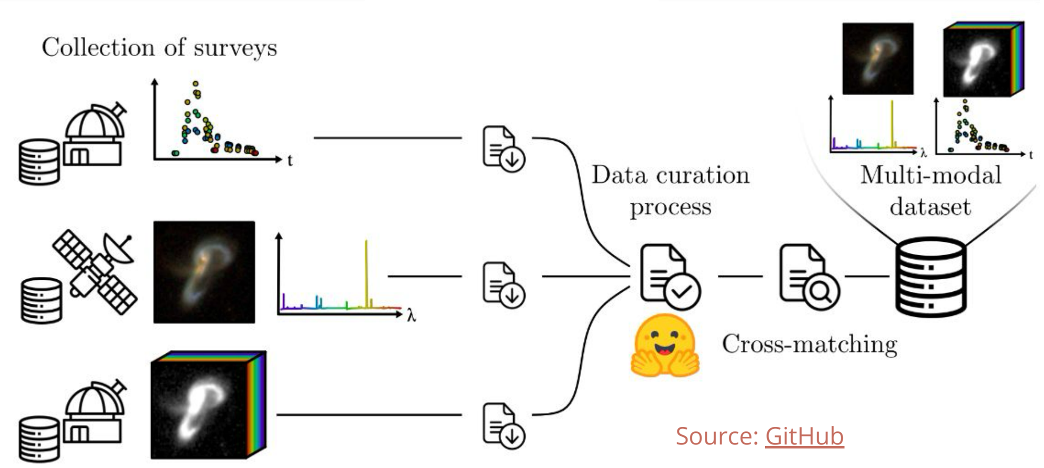

Source: GitHub

MULTIMODAL UNIVERSE

Pathfinder
LLM enabled literature search

Iyer et al. 2024.

Major challenges for AI applications in Science

As listed in The Dawes Review 10: The impact of deep learning for the analysis of galaxy surveys
[Huertas-Company & Lanusse 2022]

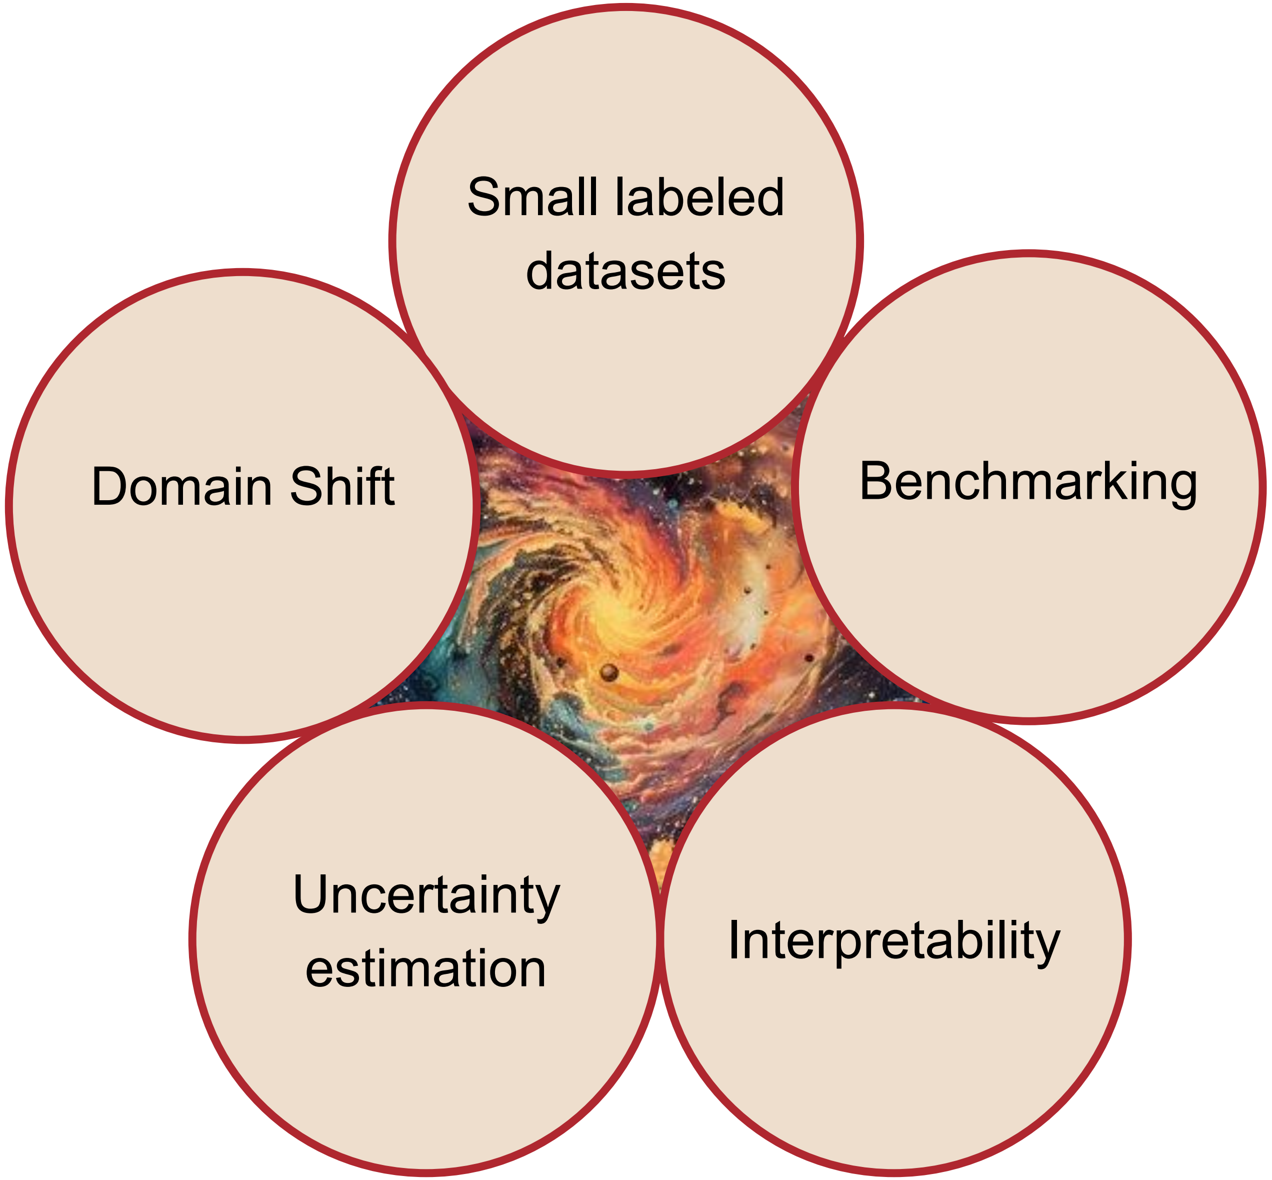

Major challenges for AI applications in Science

As listed in The Dawes Review 10: The impact of deep learning for the analysis of galaxy surveys
[Huertas-Company & Lanusse 2022]

I want to train my AI model by I don’t have enough trustworthy labels?

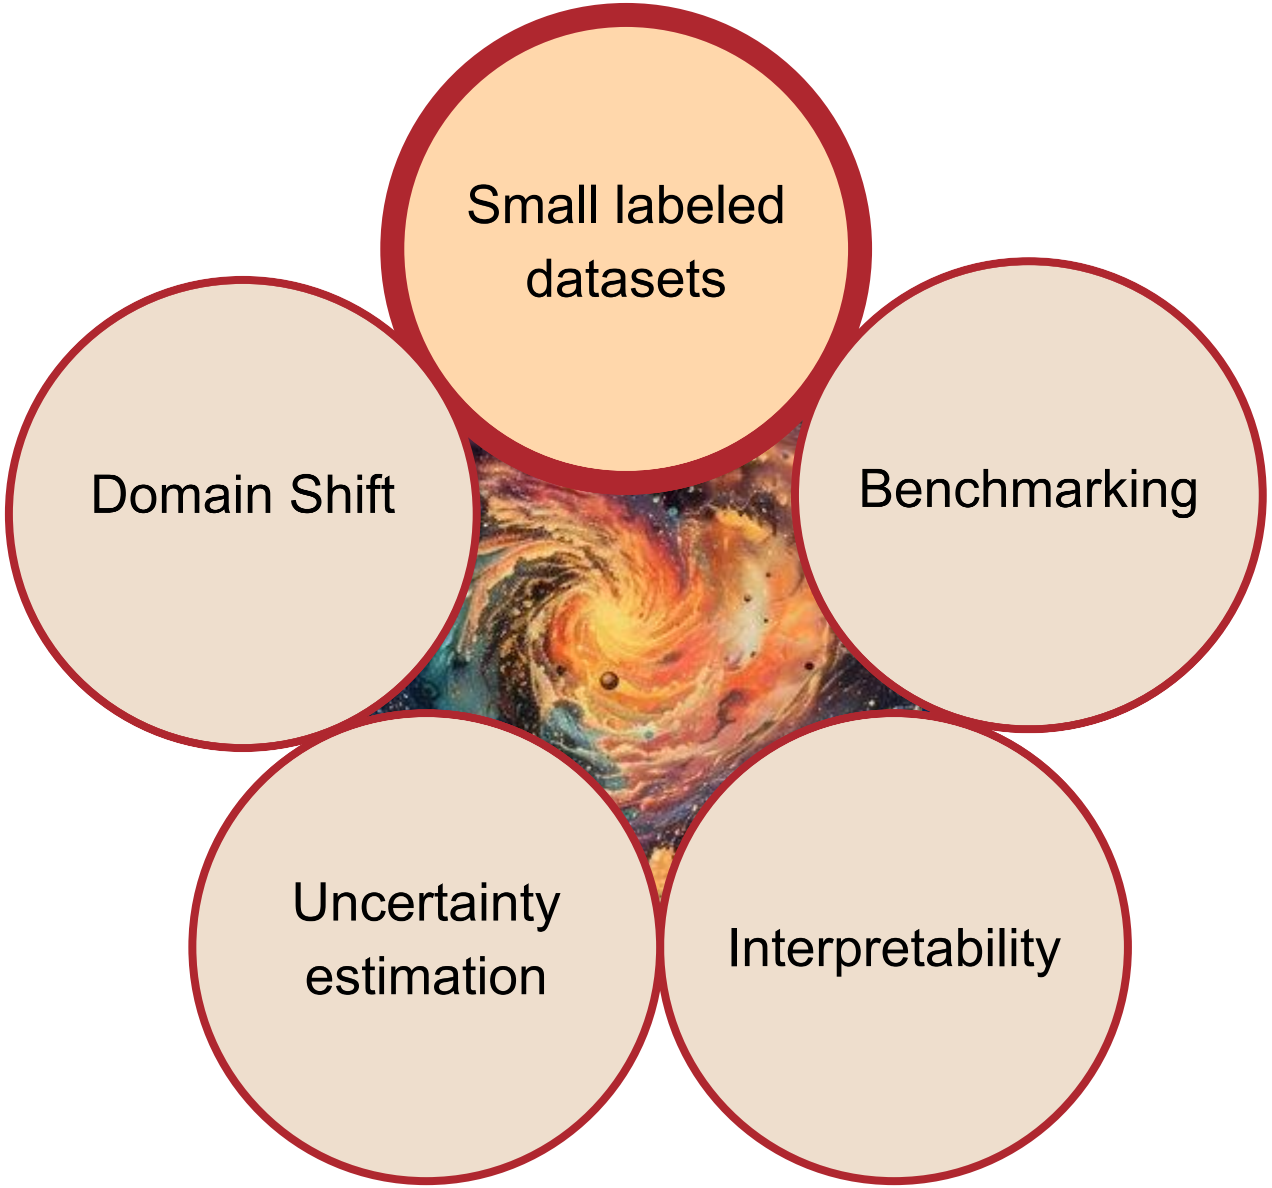

Major challenges for AI applications in Science

As listed in The Dawes Review 10: The impact of deep learning for the analysis of galaxy surveys
[Huertas-Company & Lanusse 2022]

I want to train my AI model by I don’t have enough trustworthy labels?

Can I train on simulations and apply the model to my real data?

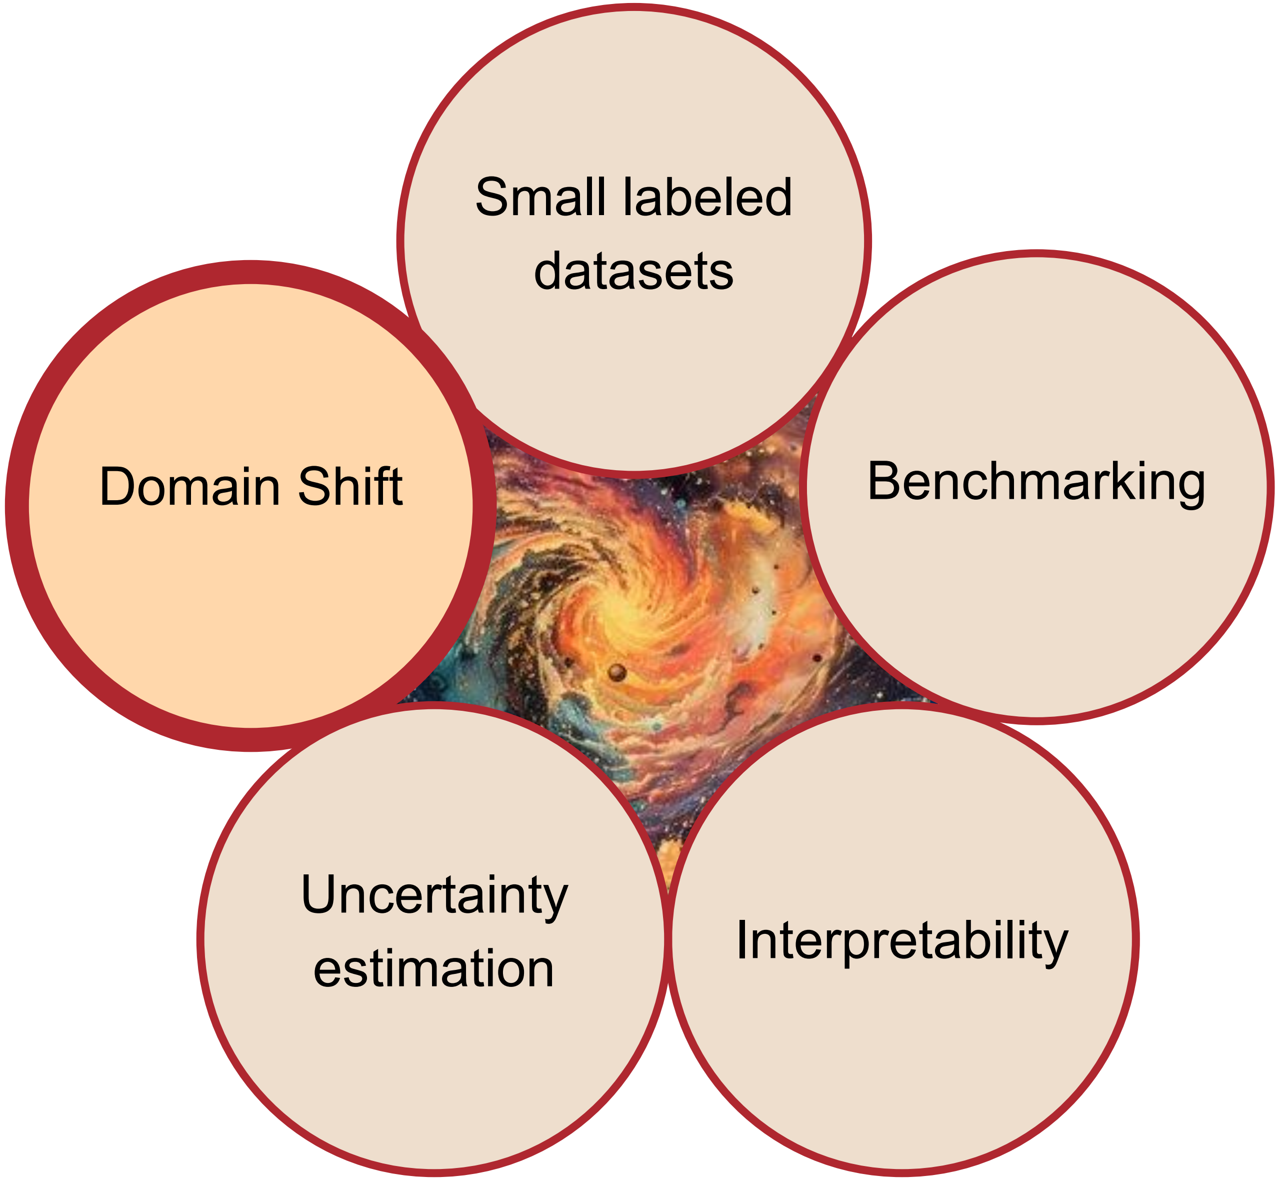

Major challenges for AI applications in Science

As listed in The Dawes Review 10: The impact of deep learning for the analysis of galaxy surveys
[Huertas-Company & Lanusse 2022]

I want to train my AI model by I don't have enough trustworthy labels?

Can I train on simulations and apply the model to my real data?

How well is my model going to perform?

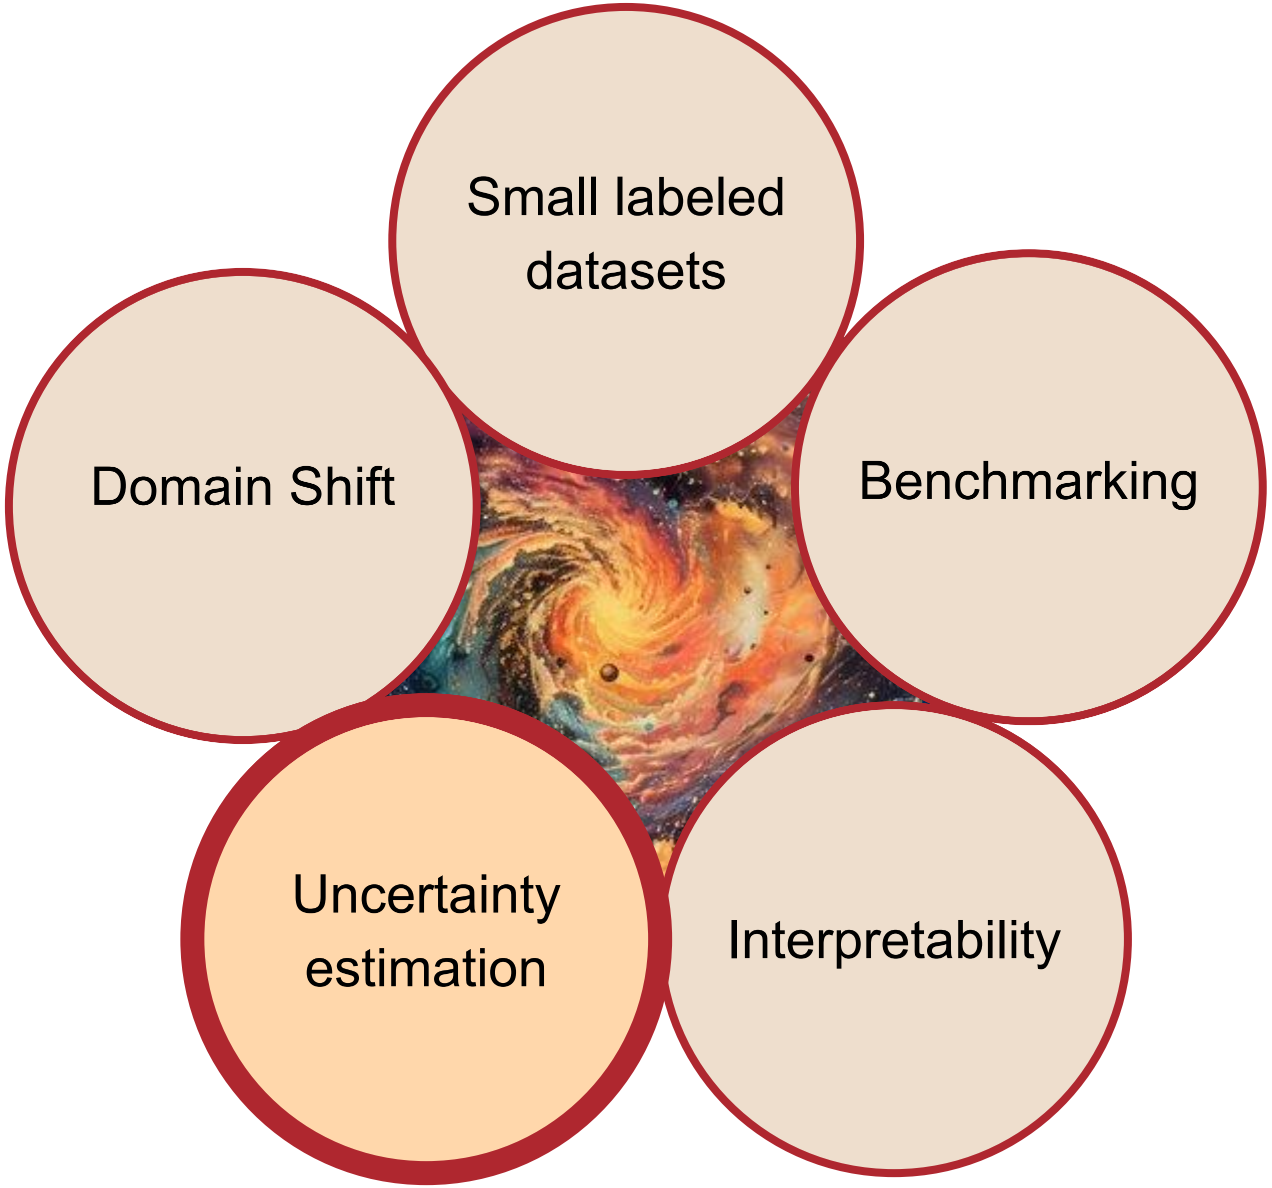

Major challenges for AI applications in Science

As listed in The Dawes Review 10: The impact of deep learning for the analysis of galaxy surveys
[Huertas-Company & Lanusse 2022]

I want to train my AI model by I don’t have enough trustworthy labels?

Can I train on simulations and apply the model to my real data?

How well is my model going to perform?

Can I understand its performance?

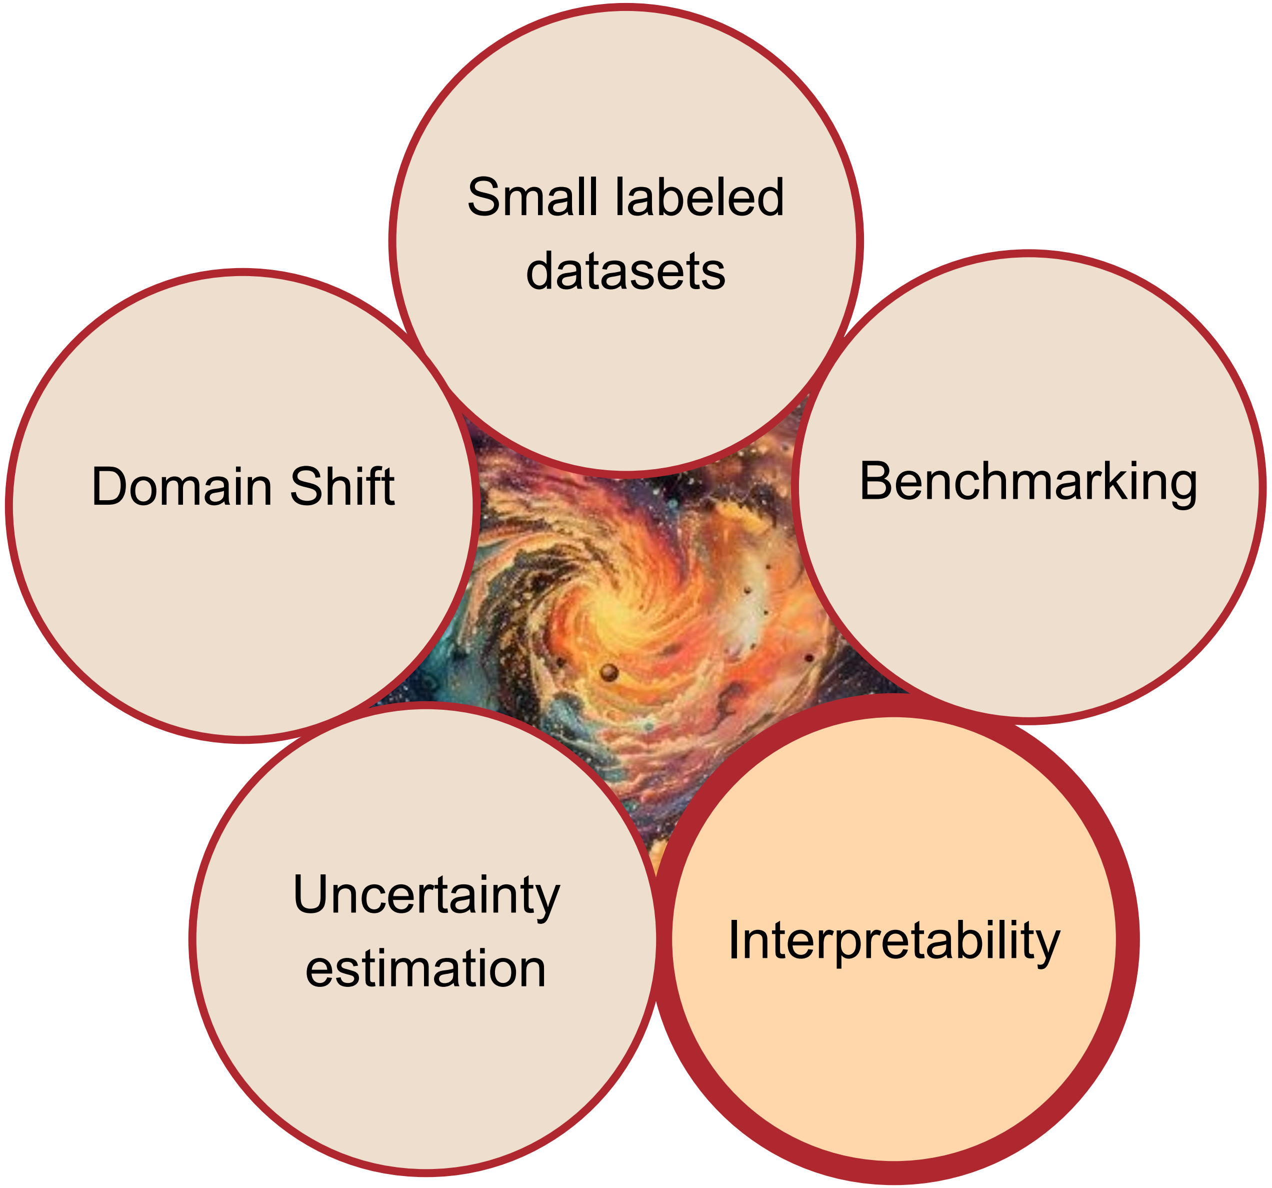

Major challenges for AI applications in Science

As listed in The Dawes Review 10: The impact of deep learning for the analysis of galaxy surveys
[Huertas-Company & Lanusse 2022]

I want to train my AI model by I don’t have enough trustworthy labels?

Can I train on simulations and apply the model to my real data?

How well is my model going to perform?

Can I understand its performance?

Can I compare my model with other researchers?

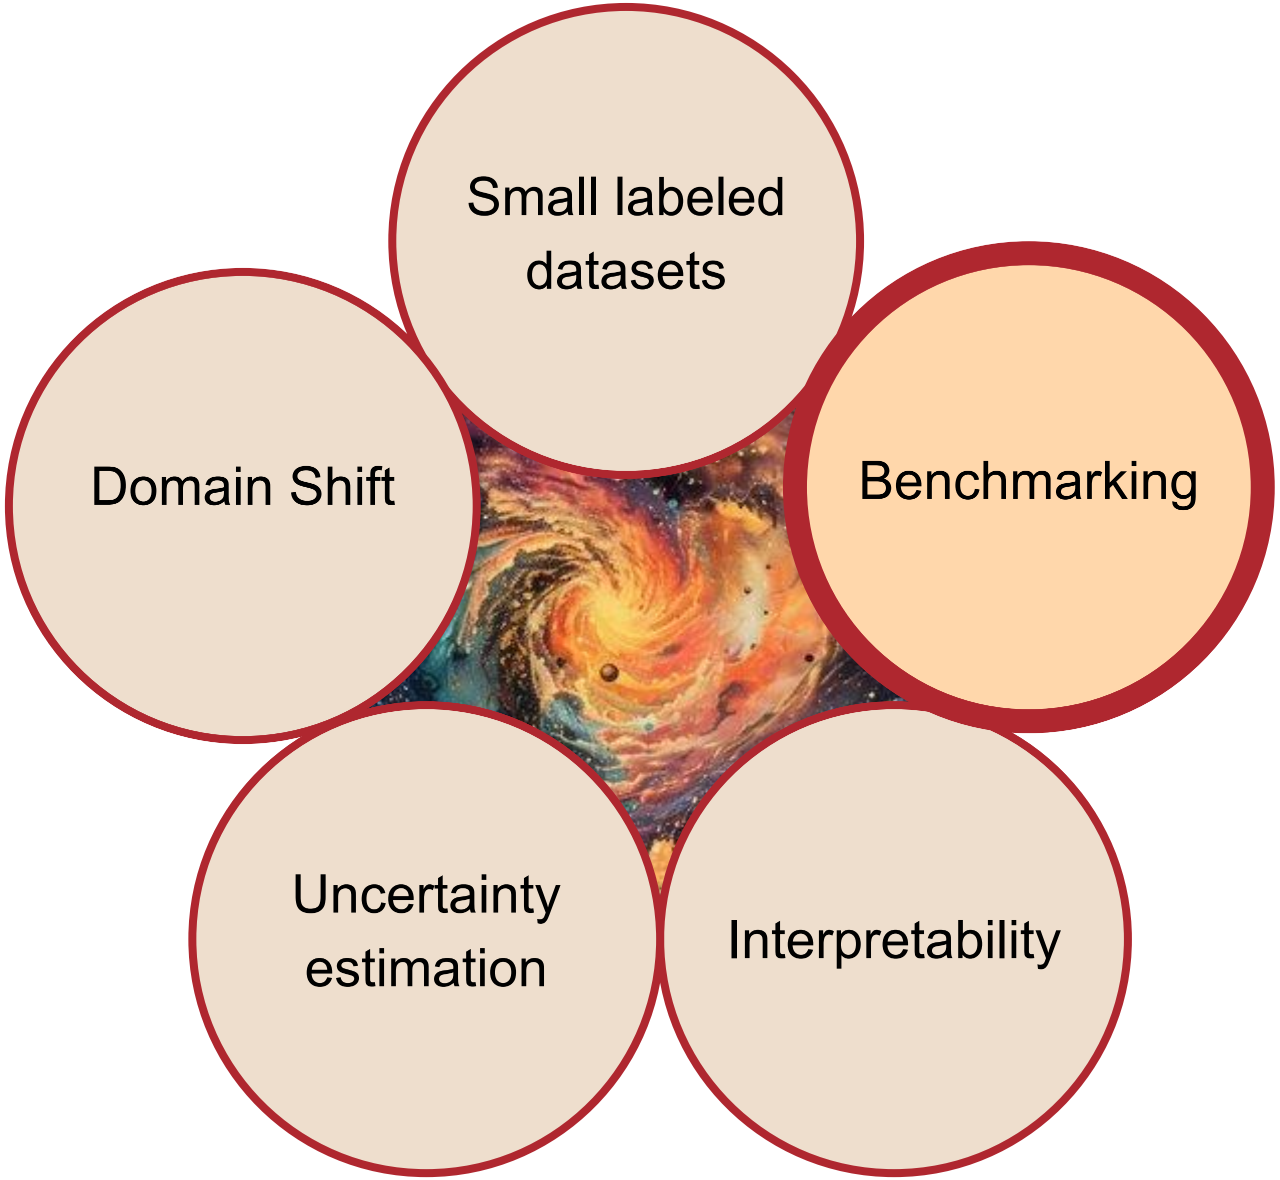

Major challenges for AI applications in Science

As listed in The Dawes Review 10: The impact of deep learning for the analysis of galaxy surveys
[Huertas-Company & Lanusse 2022]

I want to train my AI model by I don’t have enough trustworthy labels?

Can I train on simulations and apply the model to my real data?

How well is my model going to perform?

Can I understand its performance?

Can I compare my model with other researchers?

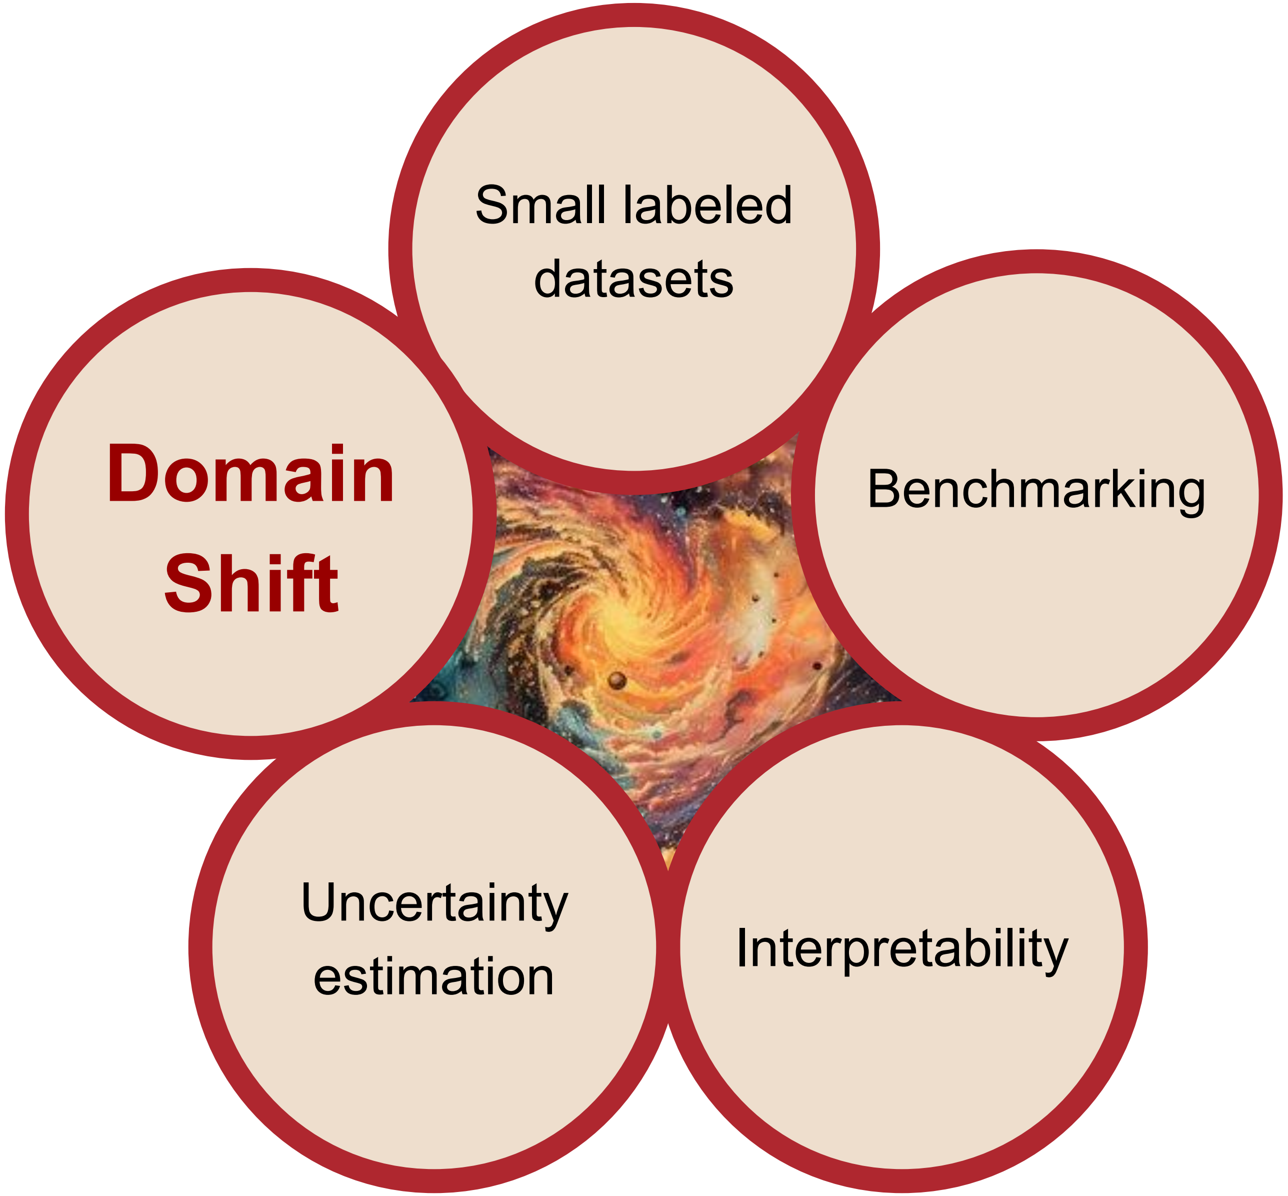

Talk Outline

1. Domain Shift Problem

2. What can we do about it?

3. Some Interesting Applications

4. Uncertainty Quantification (UQ)

01

Domain Shift Problem

More Data - More problems

All areas of astro(physics) often need to create **model trained on simulated data, that also work on real data**!

DOMAIN SHIFT

- Missing and unknown physics, wrong geometry, background levels
- Computational constraints for simulations
- Detector problems, transients, errors, data compression
- Imperfect addition of observational effects
- Different detectors or telescopes

	SIMULATED	REAL
MicroBooNE (neutrinos)	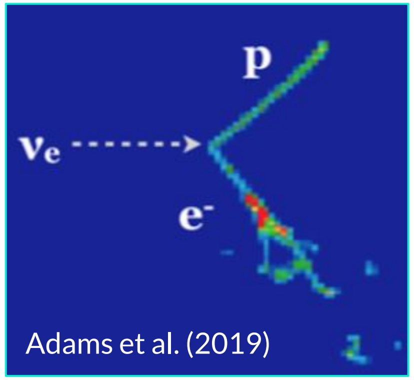	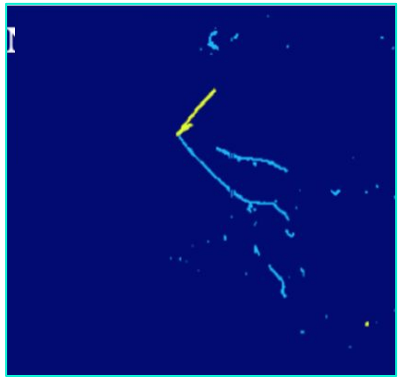
Illustris / Hubble (merging galaxies)	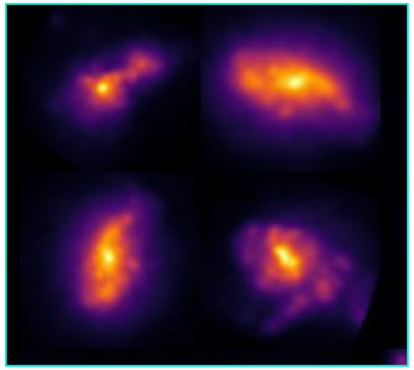	

FERMI NATIONAL ACCELERATOR LABORATORY

Traditional ML

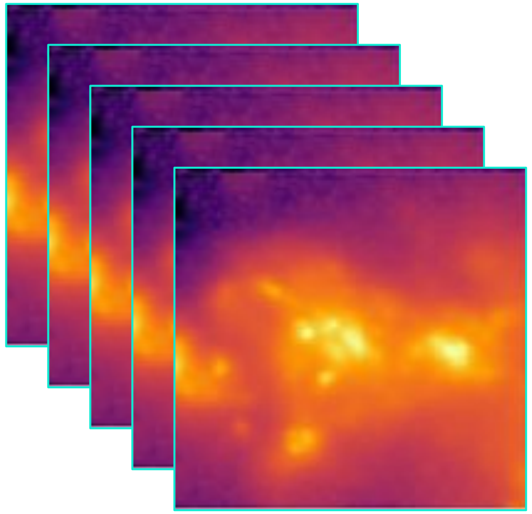

Regular Training

FIND AND REFINE FEATURES

SIMULATED IMAGES + LABELS

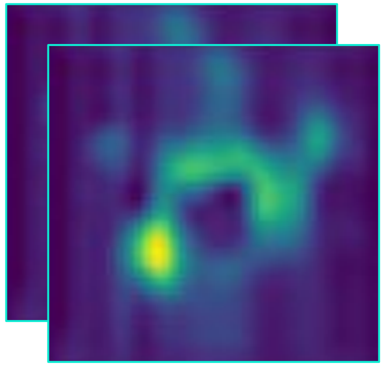

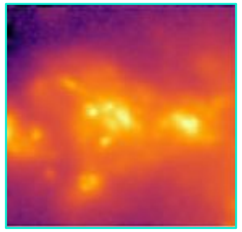

(NON)MERGER

Traditional ML

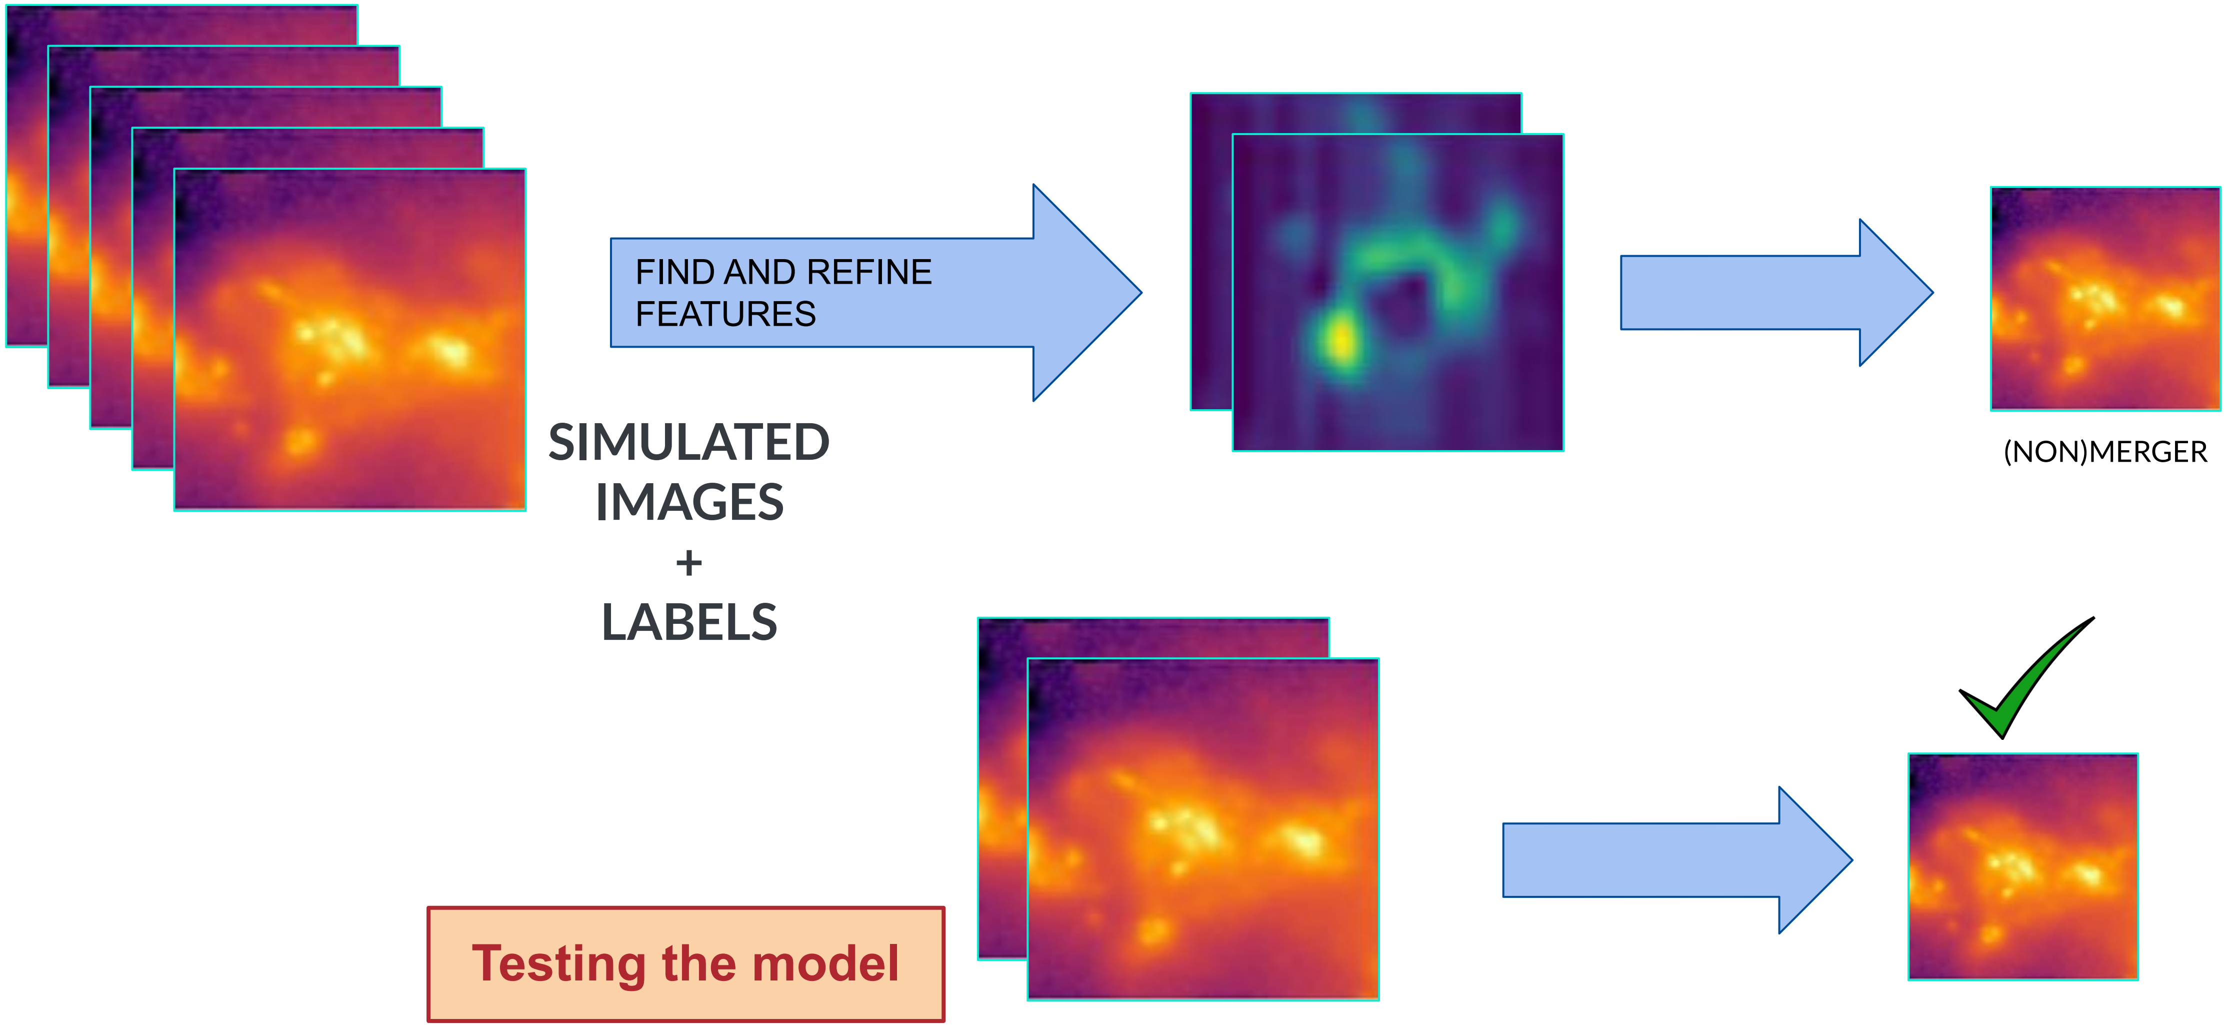

Traditional ML

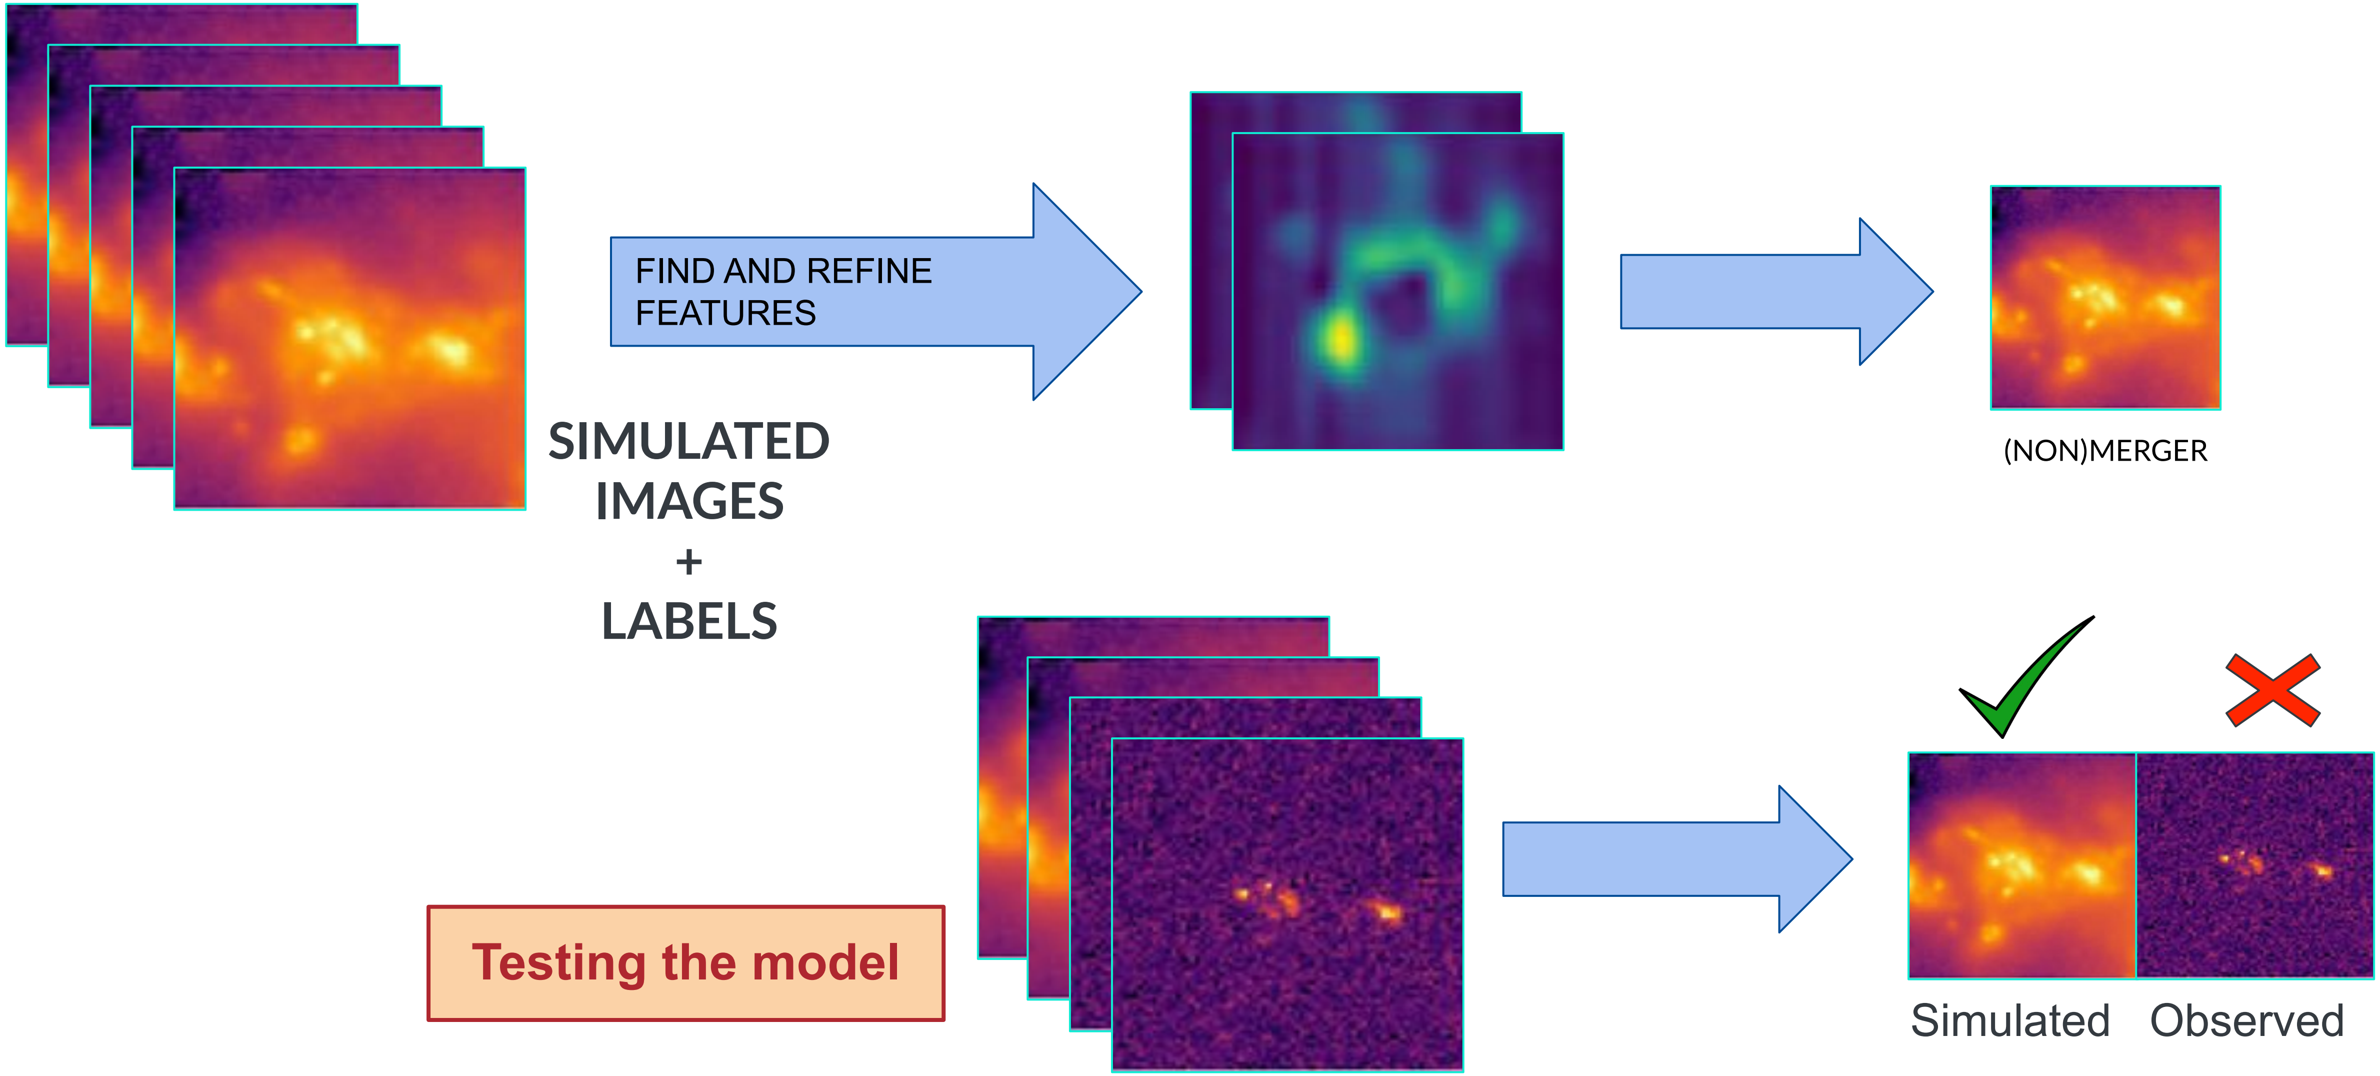

Why do models fail?

Train the model on source dataset and find the decision boundary.

Source Domain

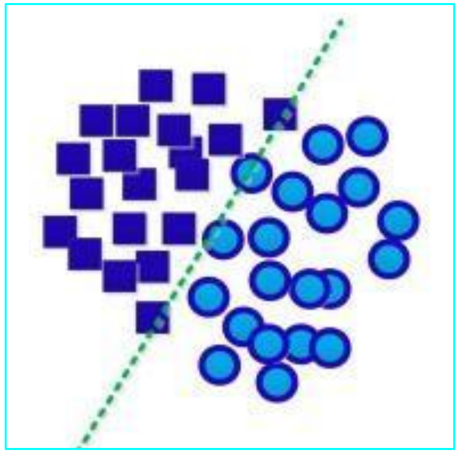

Why do models fail?

New domain is shifted,
learned decision boundary doesn't work.

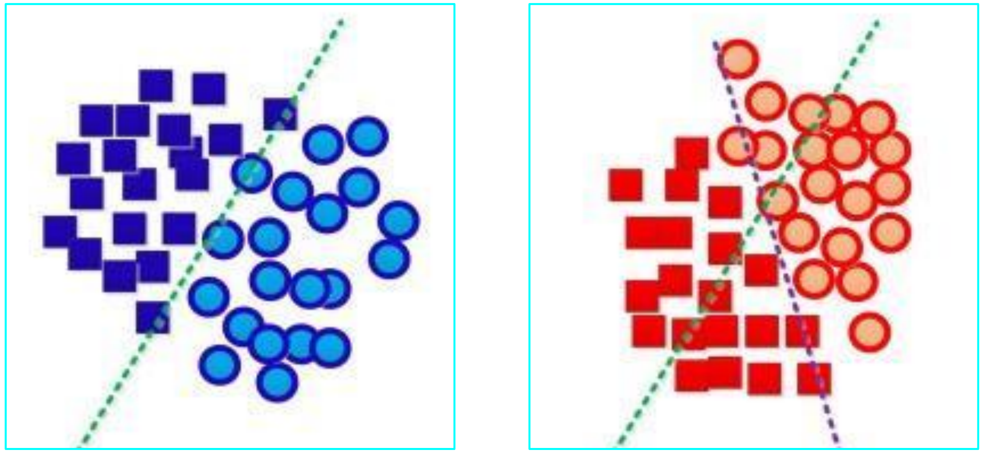

Why do models fail?

We need to align the data during training!

Source Domain

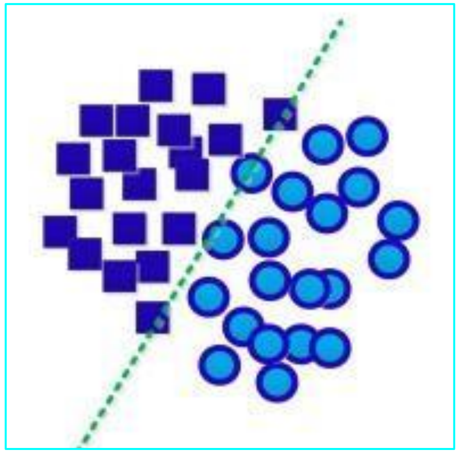

Target Domain

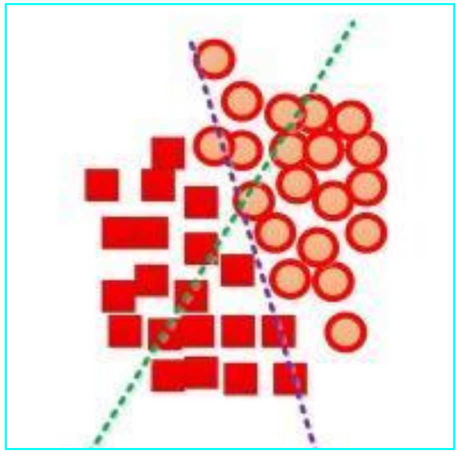

Domain Alignment

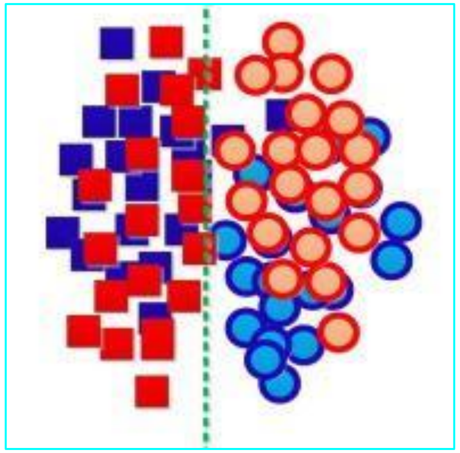

02

Domain Shift Problem
What can we do about it?

Combining Datasets

DOMAIN ADAPTATION

Align data distributions in the latent space of the network by forcing the network to **find more robust domain-invariant features.**

Combining Datasets

DOMAIN ADAPTATION

Align data distributions in the latent space of the network by forcing the network to **find more robust domain-invariant features.**

Distance-based methods

Adversarial methods

Reconstruction-based

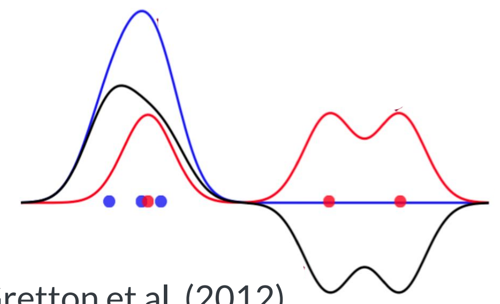

Gretton et al. (2012)

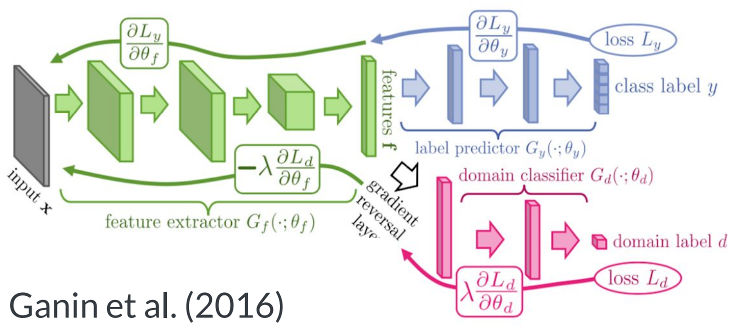

Ganin et al. (2016)

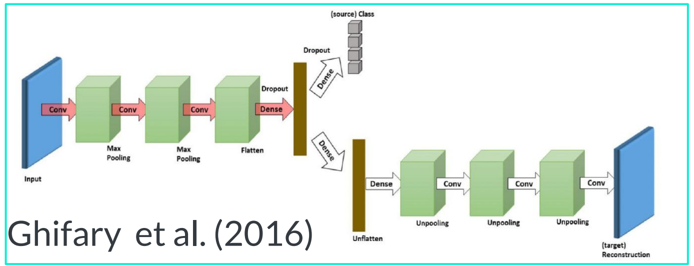

Ghifary et al. (2016)

Minimize the distance metric between two latent distributions.

Using domain discriminators to encourage domain confusion through an adversarial objective.

Data reconstruction as an auxiliary task to ensure feature invariance.

Combining Datasets

DOMAIN ADAPTATION

Align data distributions in the latent space of the network by forcing the network to **find more robust domain-invariant features.**

- Distance-based methods
- Adversarial methods
- Reconstruction-based

Works on **unlabeled target domain**!
Can be applied to **new data**, no need for scientists to label anything.

Training with Domain Adaptation

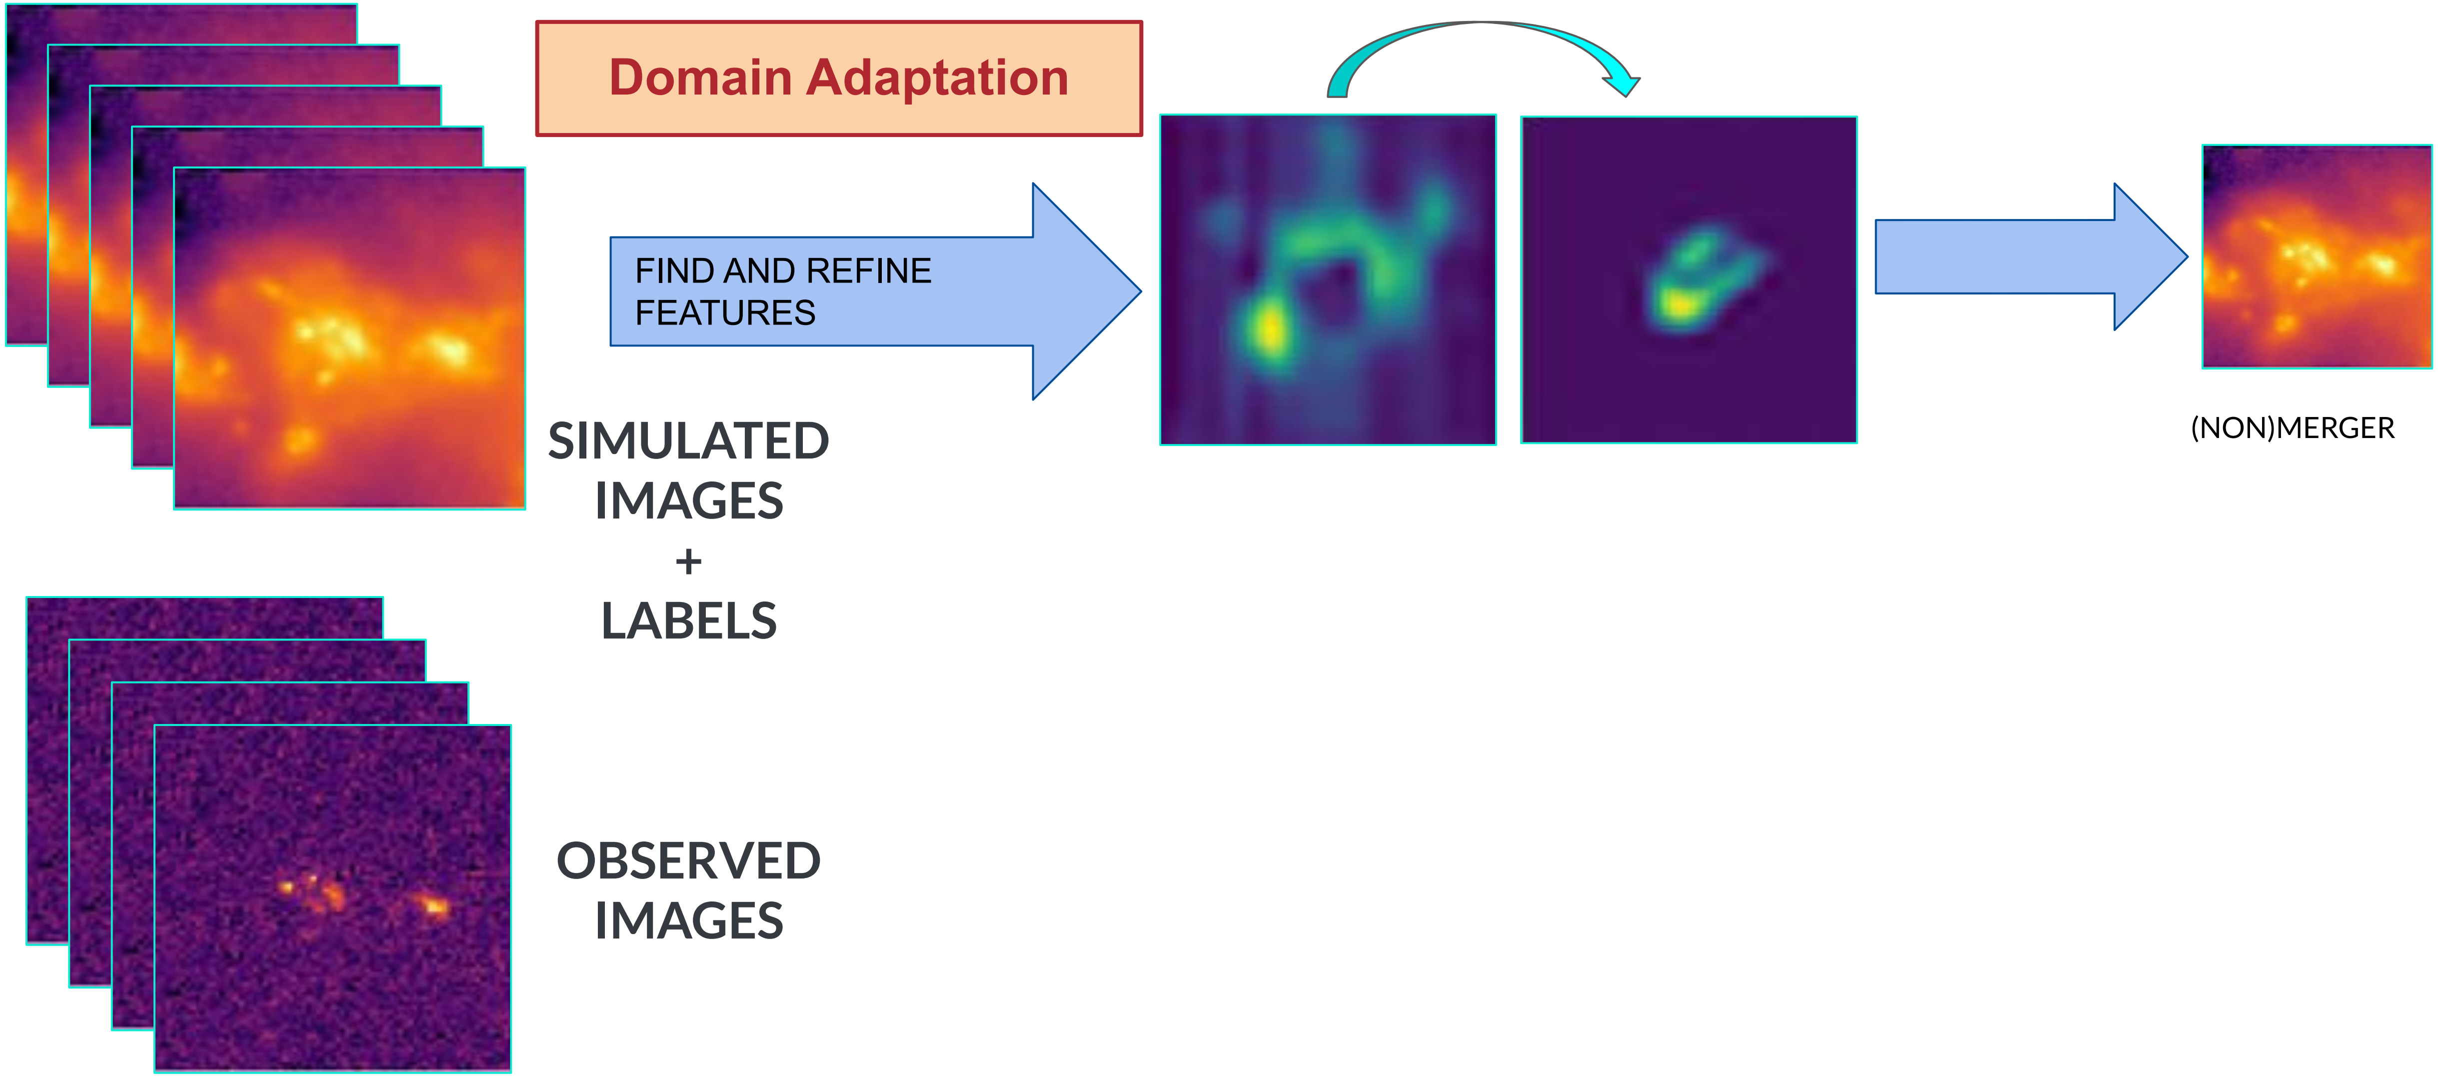

Training with Domain Adaptation

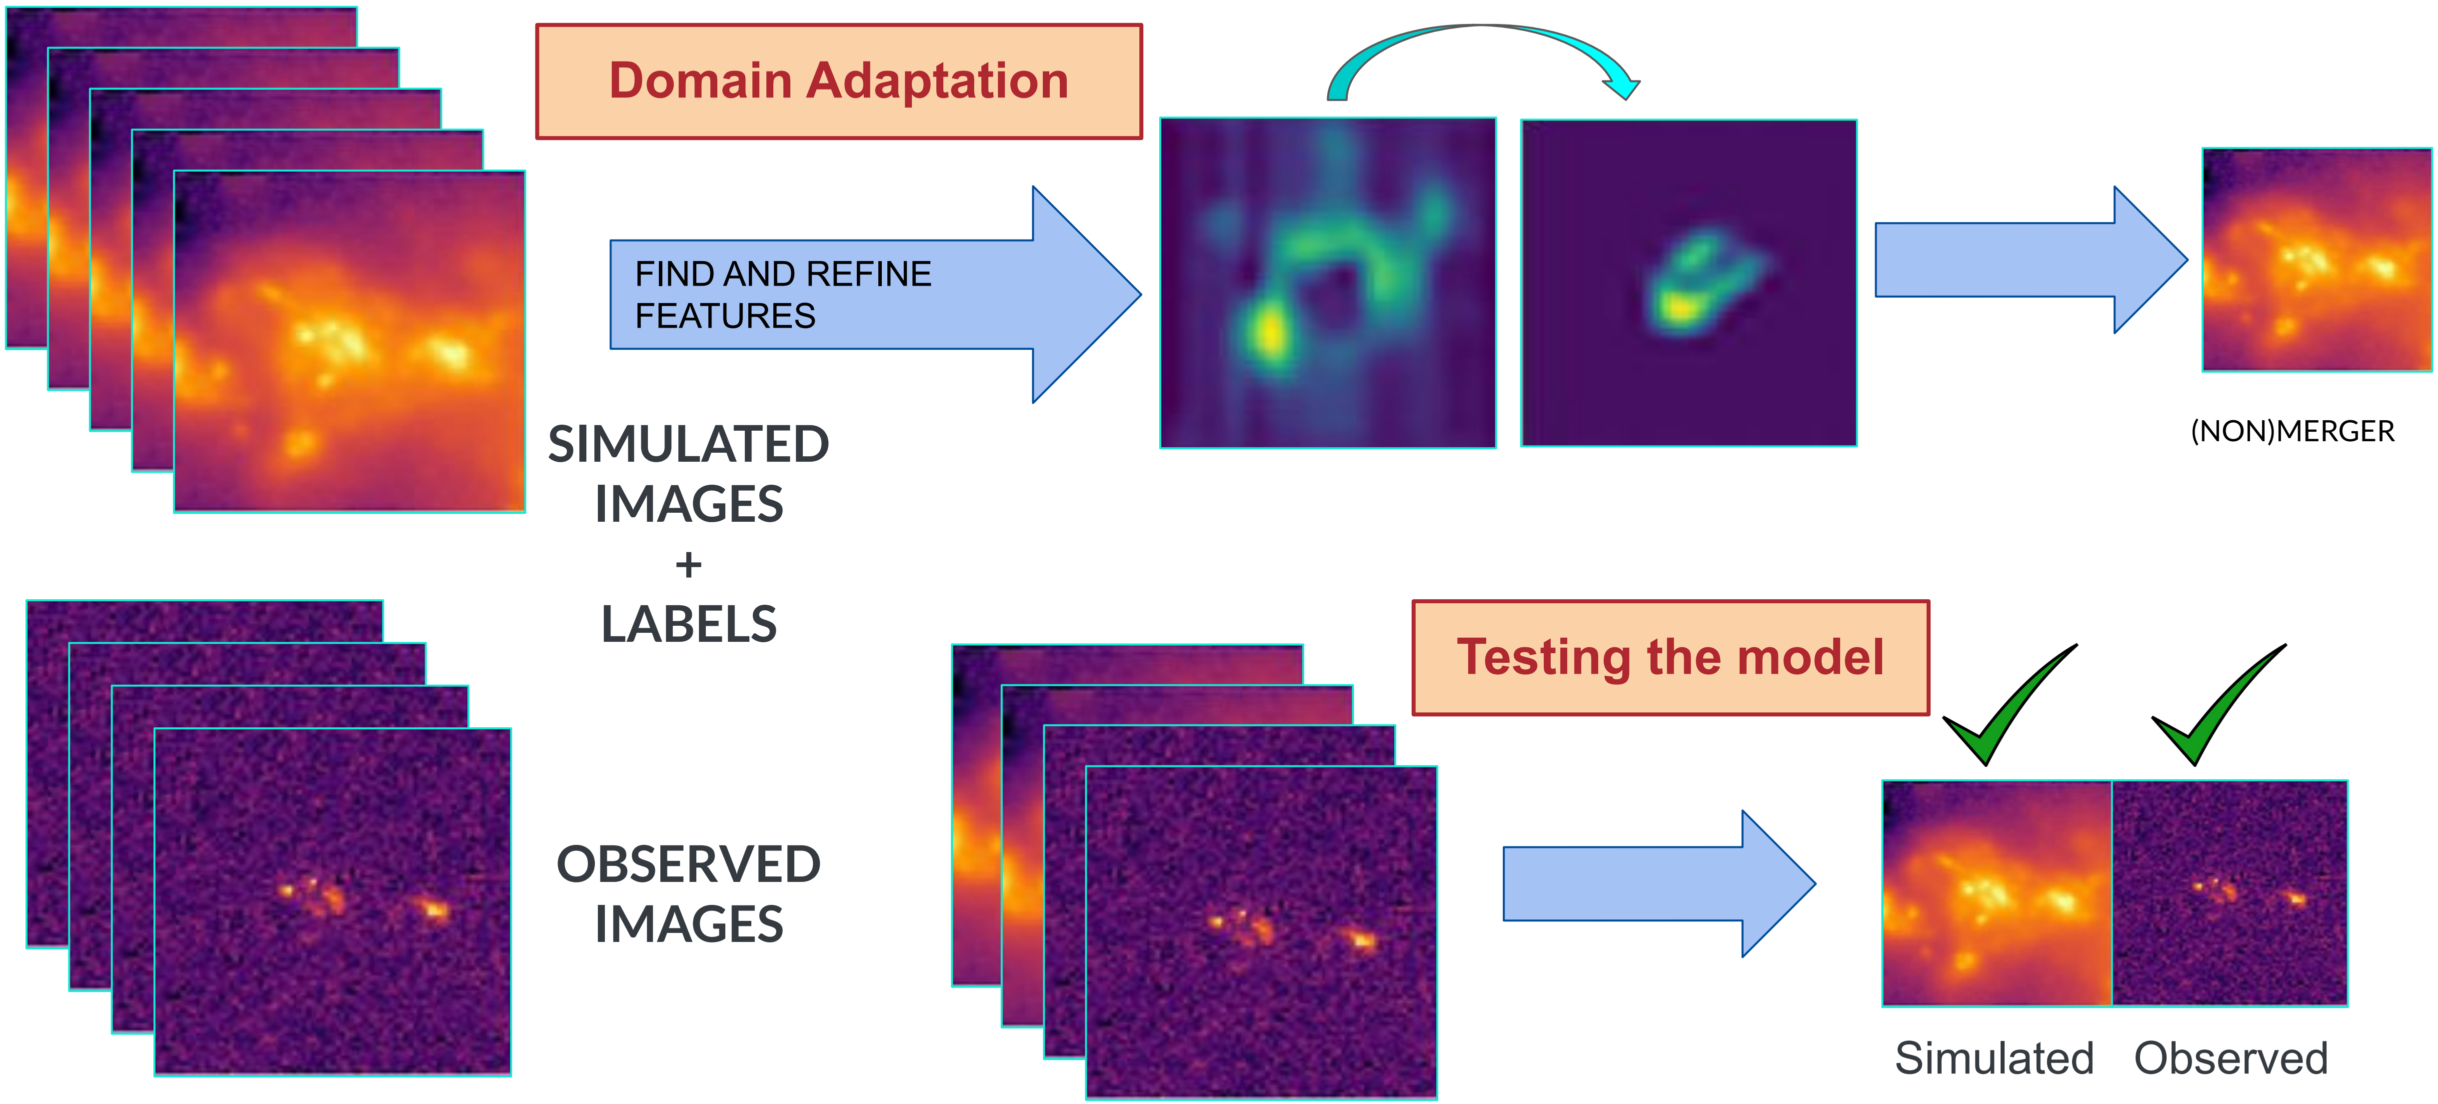

Latent Space Alignment

Source - Illustris

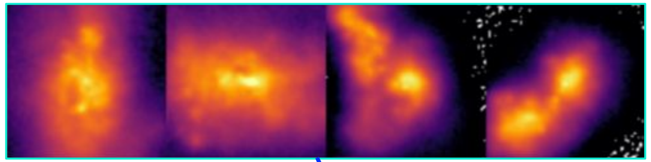

Target - SDSS observations

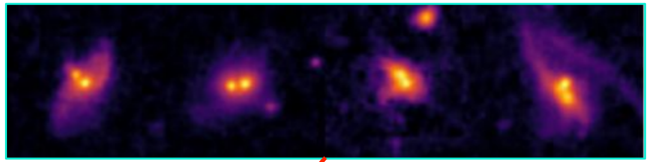

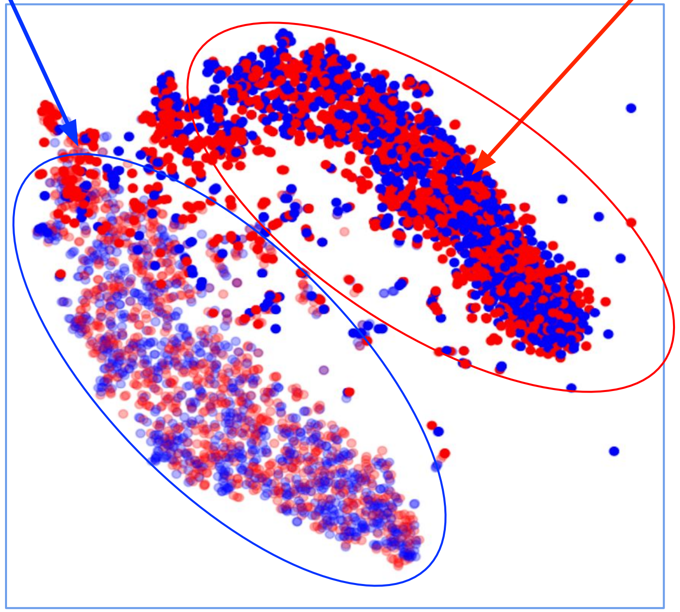

This is how the network sees the data.
2D representation of network’s latent space.

Ćiprijanović et al. 2020.
Ćiprijanović et al. 2021.

Latent Space Alignment

Source - Illustris

Target - SDSS observations

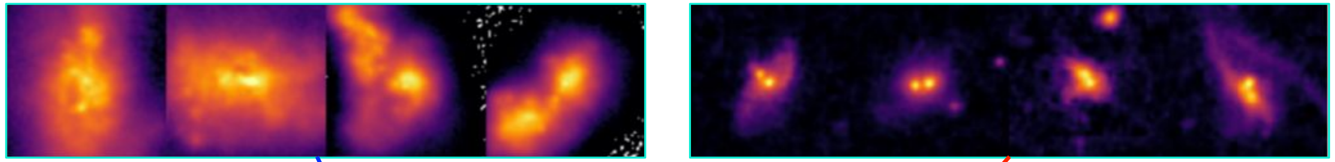

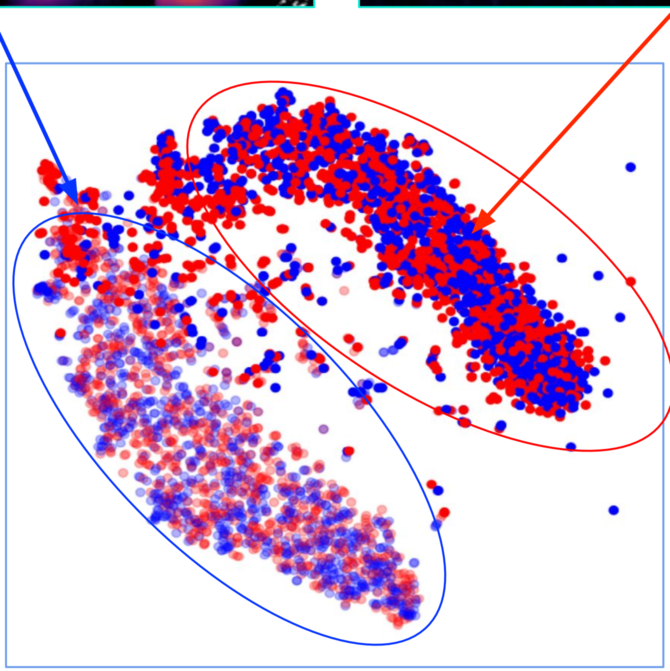

Regular Training

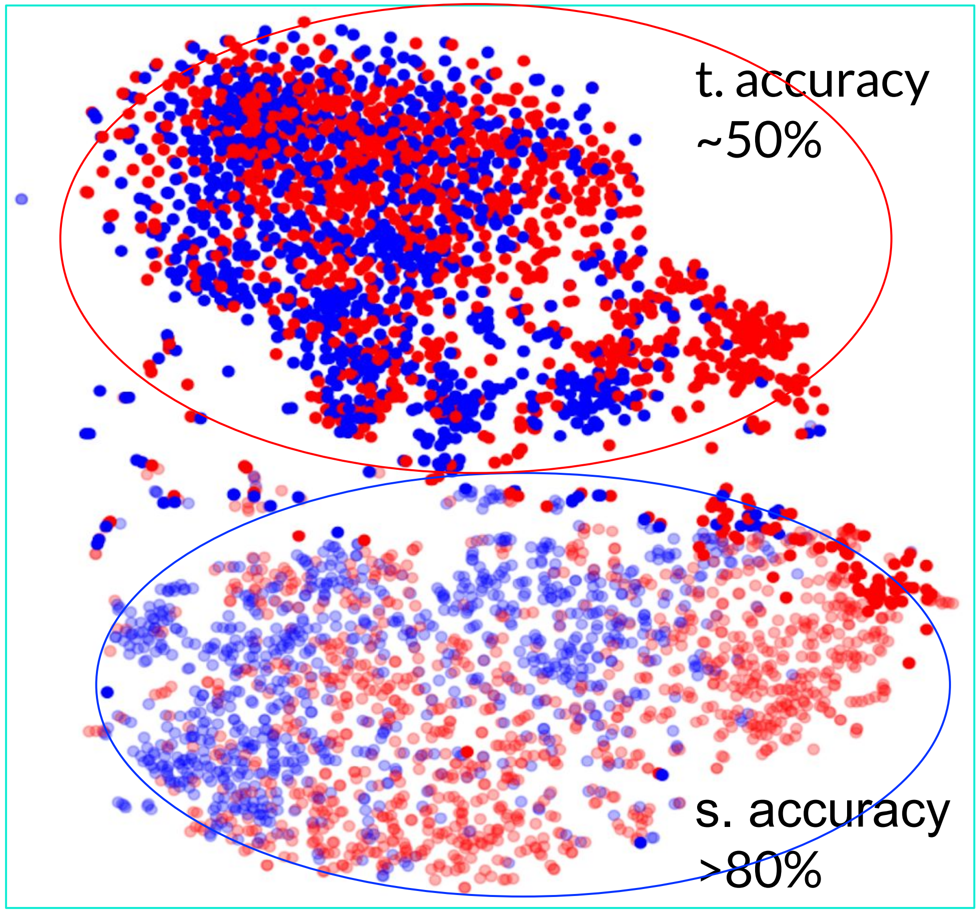

Ćiprijanović et al. 2020.
Ćiprijanović et al. 2021.

Latent Space Alignment

Source - Illustris

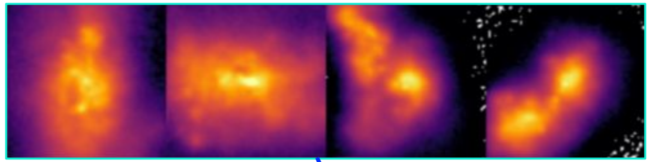

Target - SDSS observations

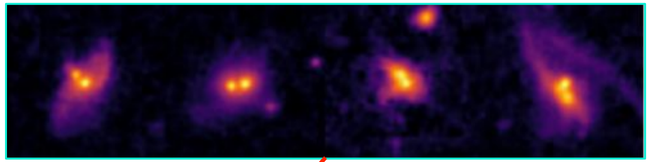

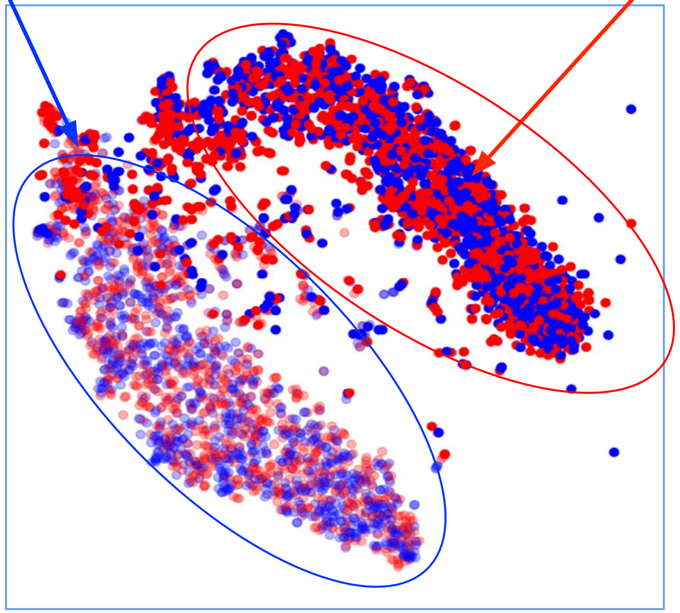

Ćiprijanović et al. 2020.
Ćiprijanović et al. 2021.

Domain Adaptation

Up to 30% increase!

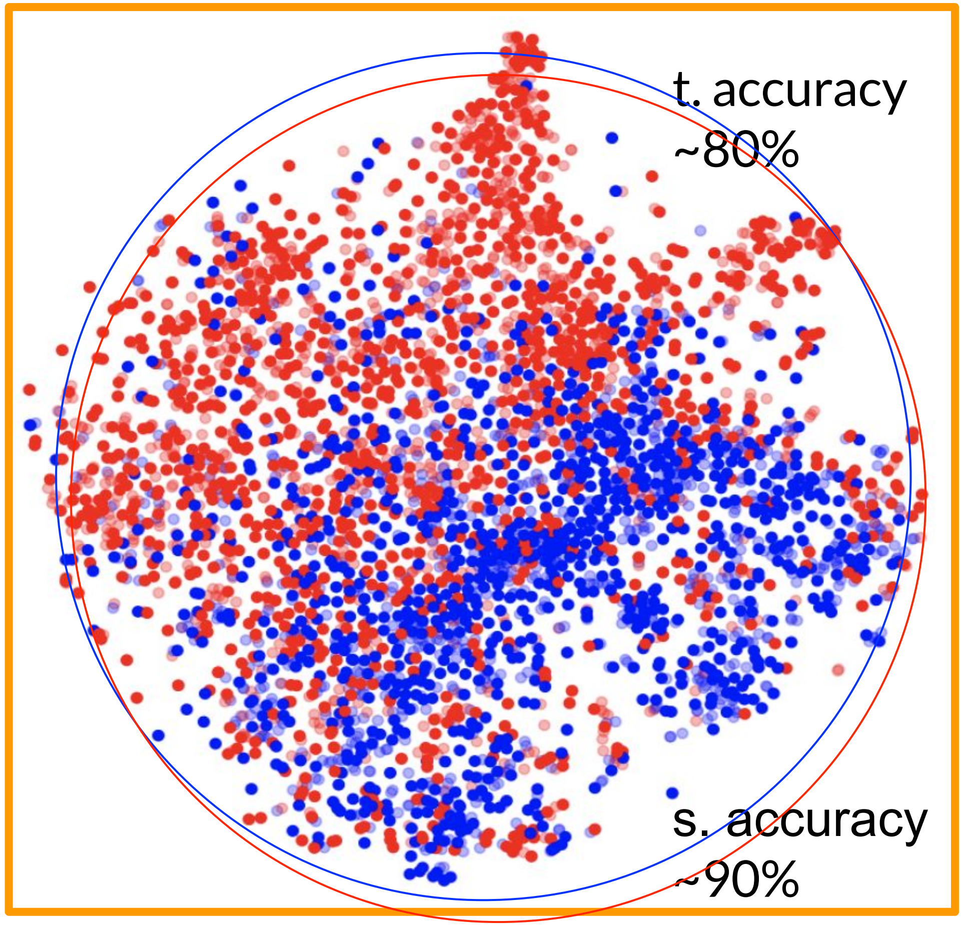

Where can we use it?

Any data, any task, any problem (within reason!)

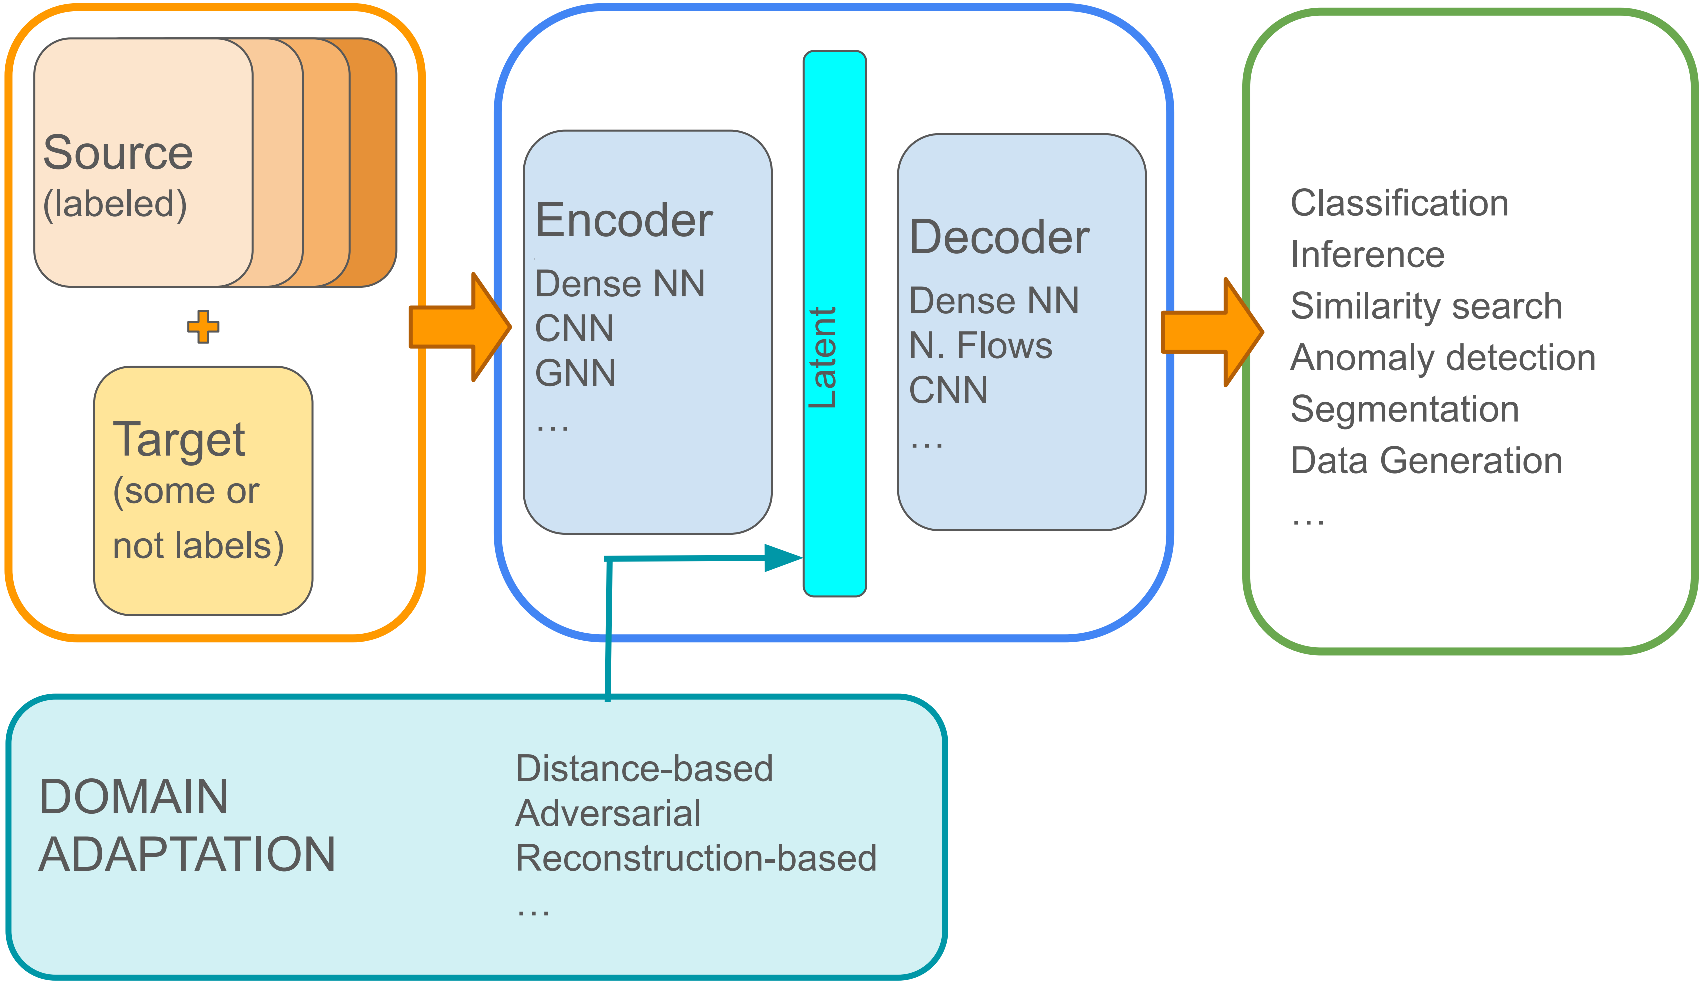

Where can we use it?

Any data, any task, any problem (within reason!)

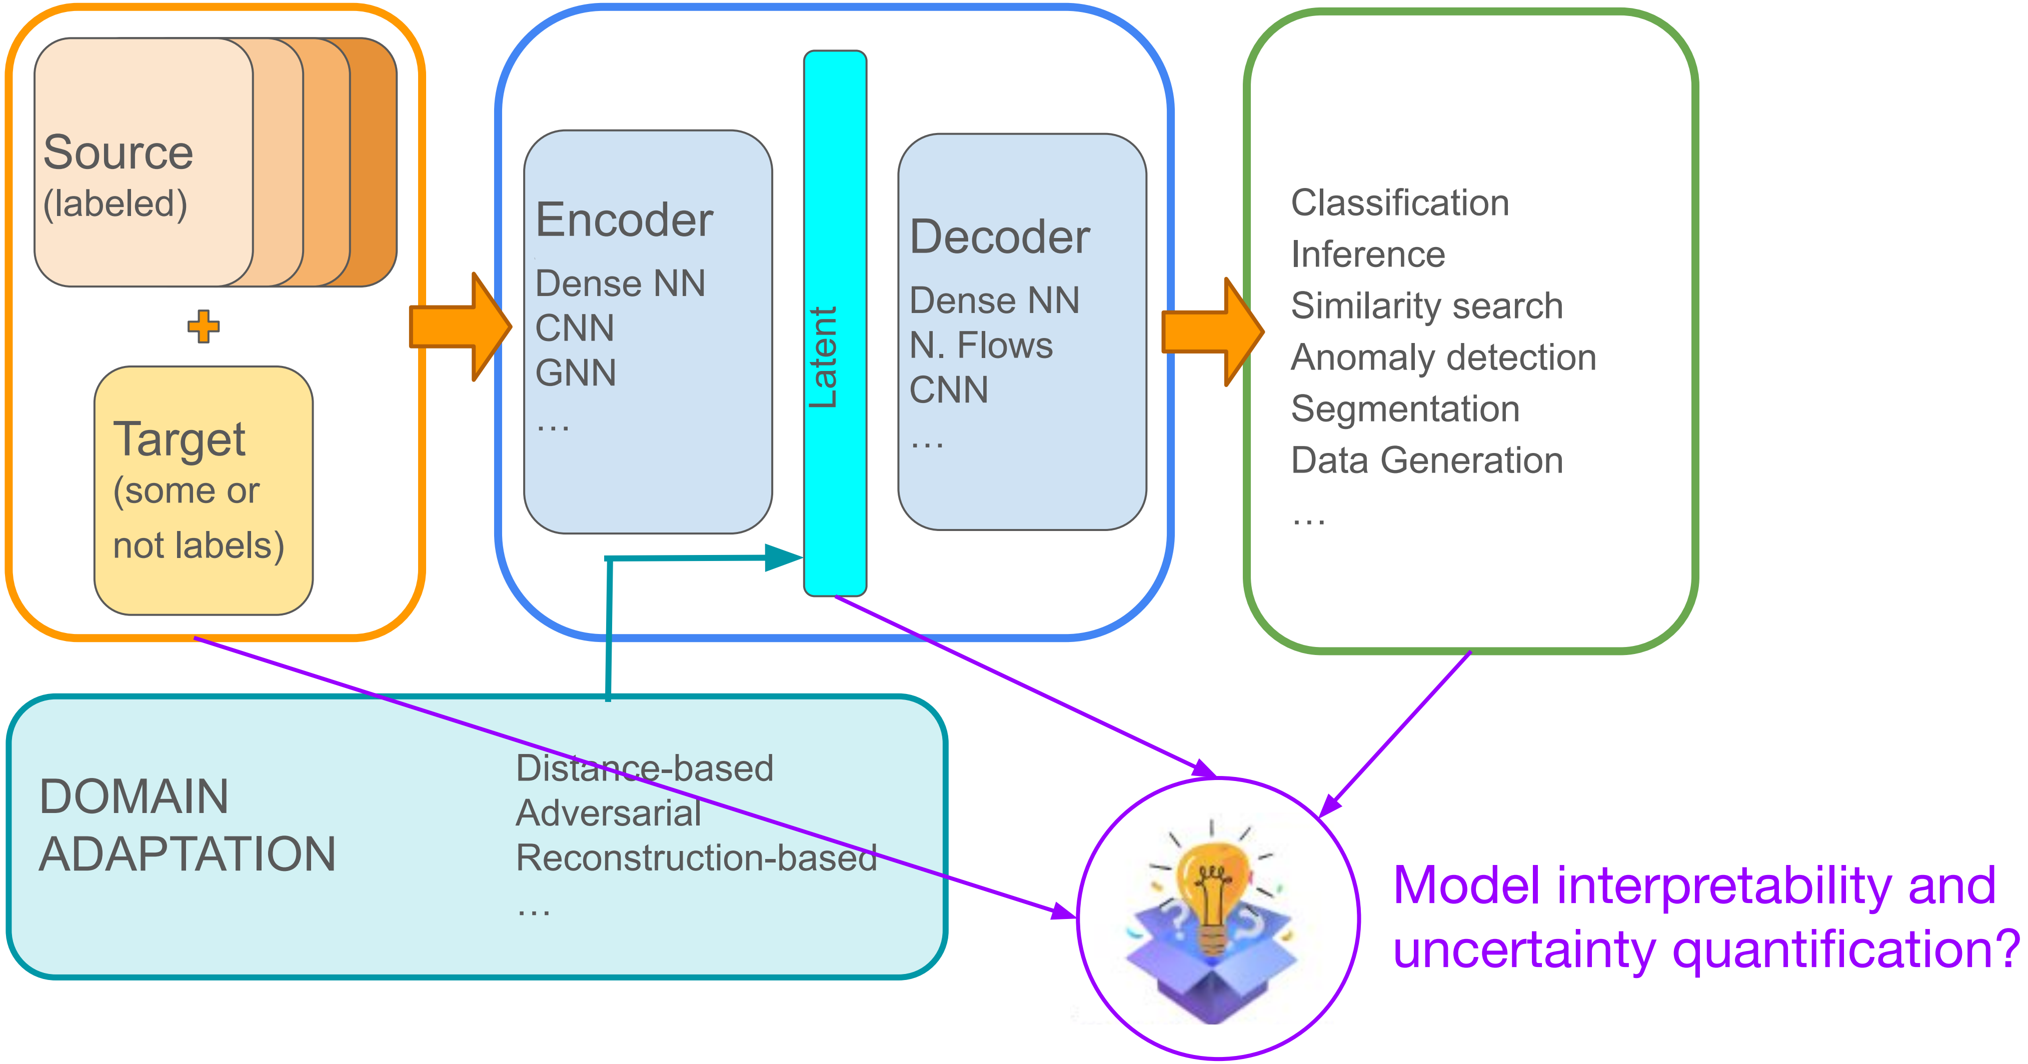

03

Some interesting applications

SIDDA: SInkhorn Dynamic Domain Adaptation

Can we perform automated domain alignment and avoid time-consuming hyperparameter tuning?

SIDDA: SInkhorn Dynamic Domain Adaptation

Can we perform automated domain alignment?

- **Trainable scaling** of the main task loos and DA loss.
- **Trainable Sinkhorn plan** i.e., how detailed should the distance measure be.
- **No tuning needed**!

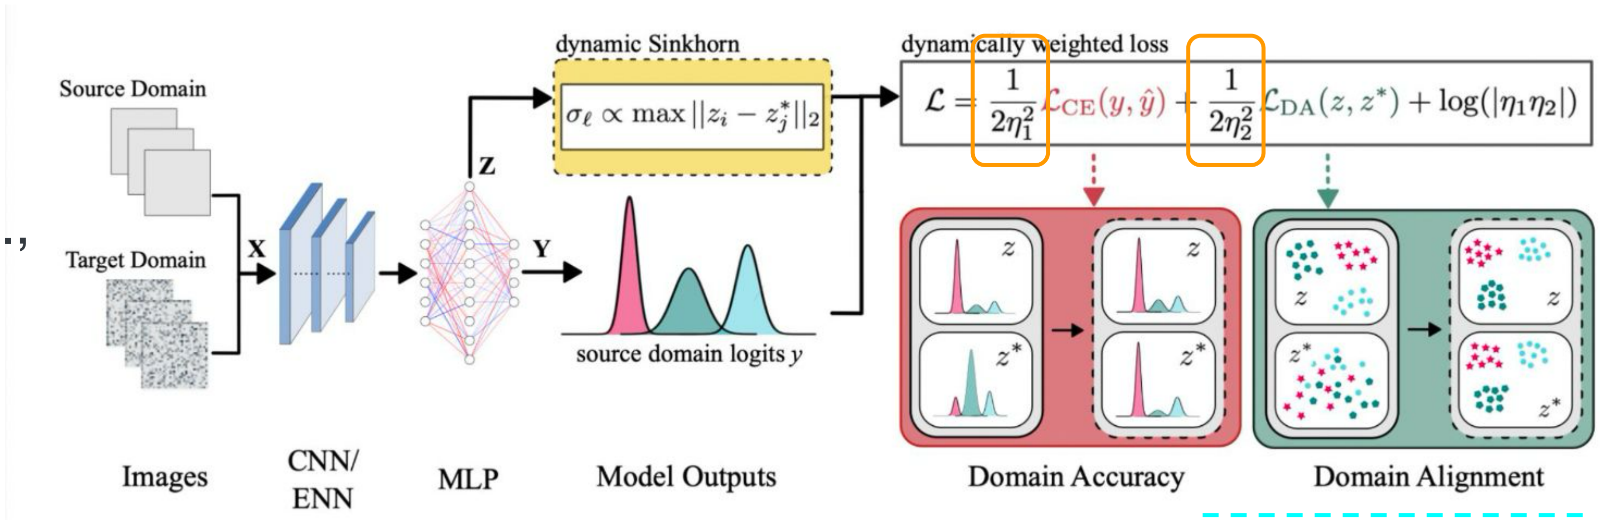

Pandya et al. 2025.

Source: Charpentier 2024.

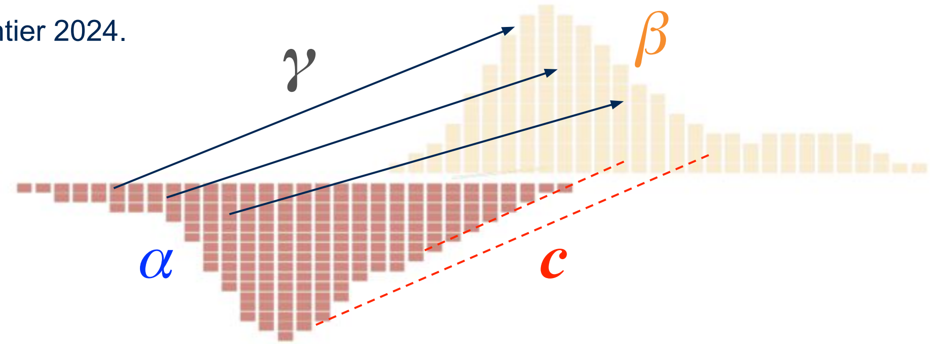

OPTIMAL TRANSPORT:
Finding the "most efficient" transportation plan γ that minimizes the total transportation cost c of moving the entire mass from probability α distribution to β.

SIDDA: SInkhorn Dynamic Domain Adaptation

Can we perform automated domain alignment?

- **Trainable scaling** of the main task loos and DA loss.
- **Trainable Sinkhorn plan** i.e., how detailed should the distance measure be.
- **No tuning needed**!

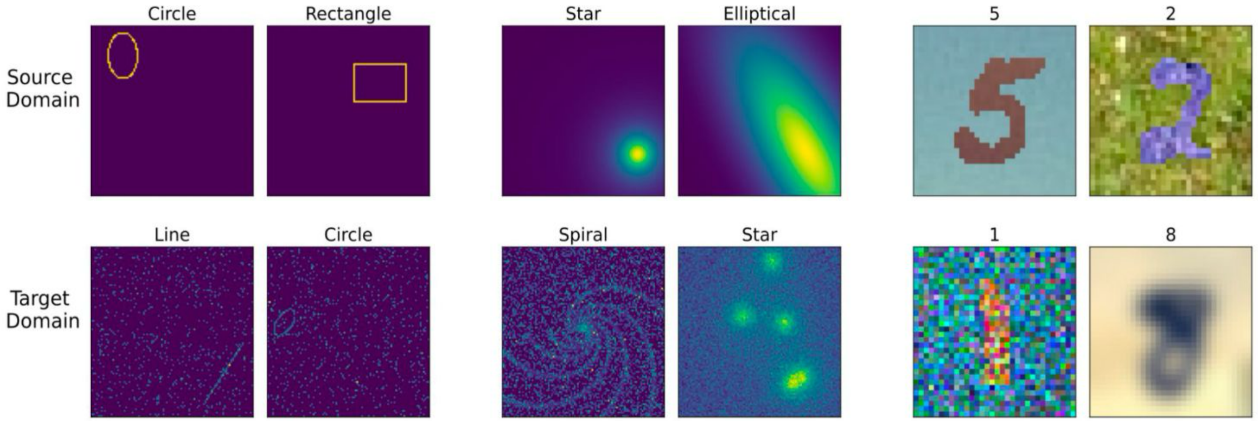

Voetberg et al. 2025

- Up to **40% better accuracy** on unlabeled data.

Pandya et al. 2025.

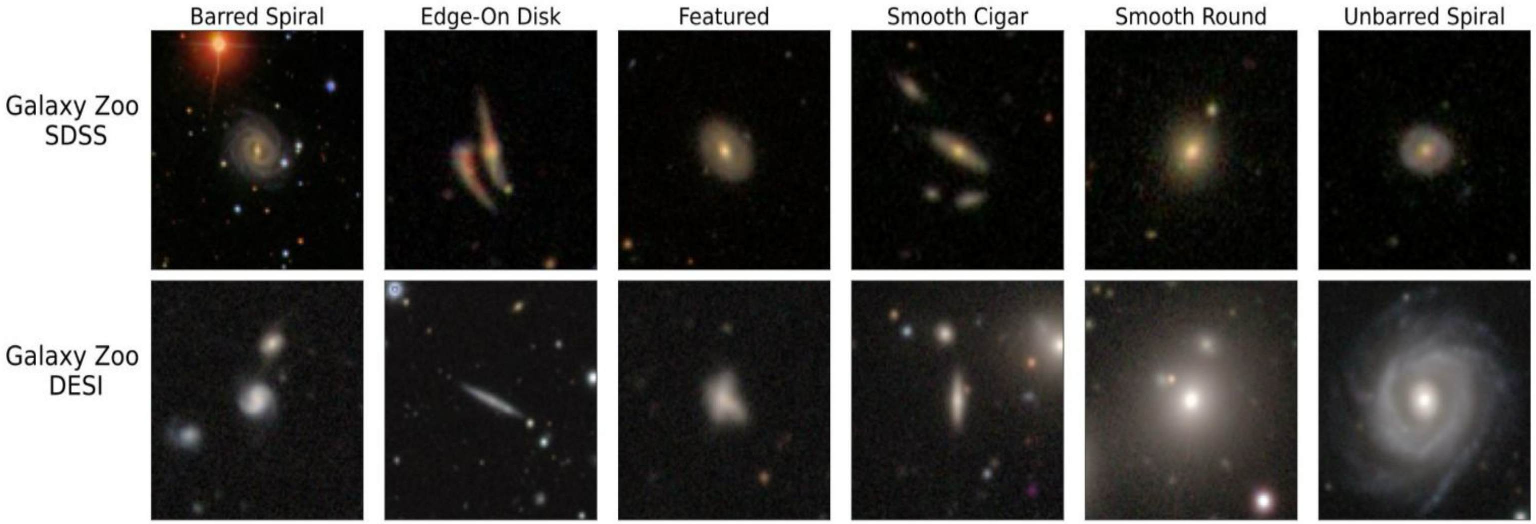

Cosmology with Graphs

NeurIPS 2023.
Roncoli et al. 2023.

Can we correctly predict cosmology across different cosmological simulations?

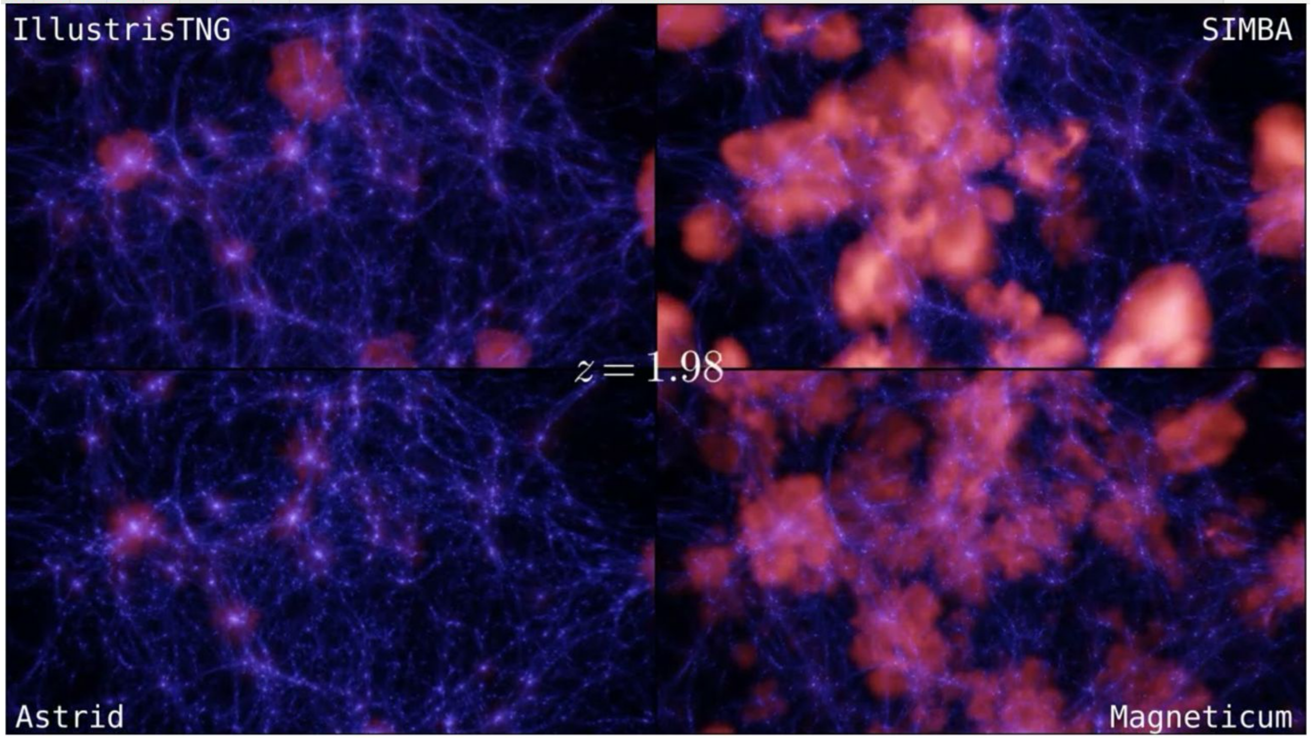

Cosmology with Graphs

Can we correctly infer cosmology across different simulations?

NeurIPS 2023.
Roncoli et al. 2023.

Graph Neural Networks:
ideal for sparse
galaxy catalogs!

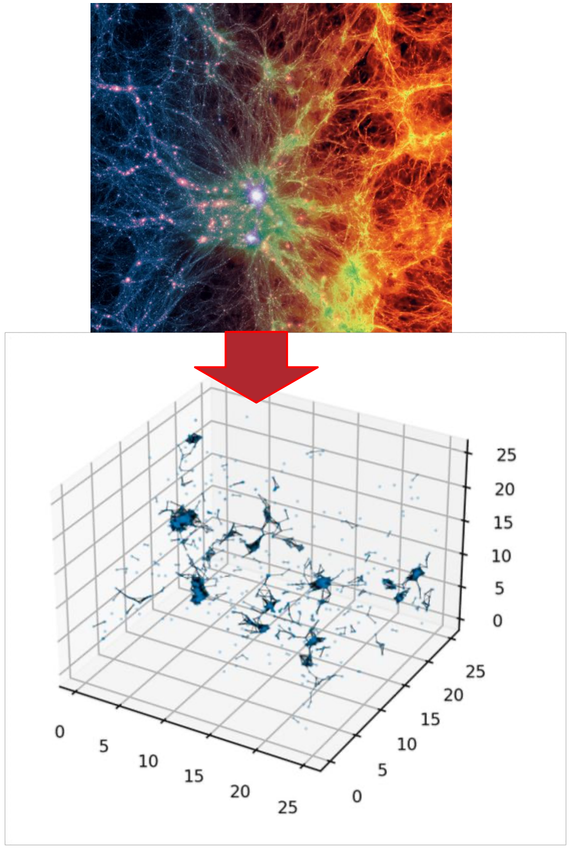

Cosmology with Graphs

Can we correctly infer cosmology across different simulations?

NeurIPS 2023.
Roncoli et al. 2023.

Graph Neural Networks:
ideal for sparse
galaxy catalogs!

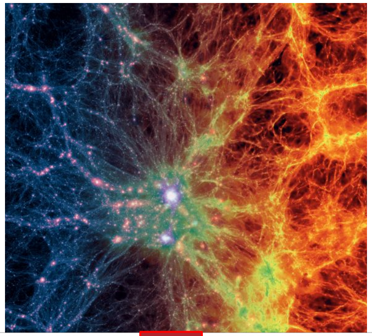

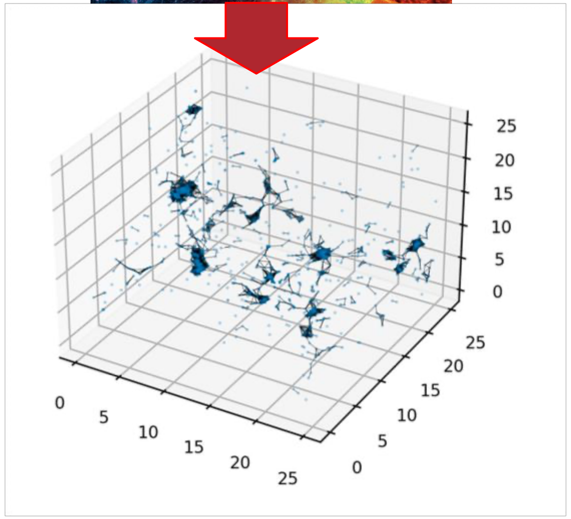

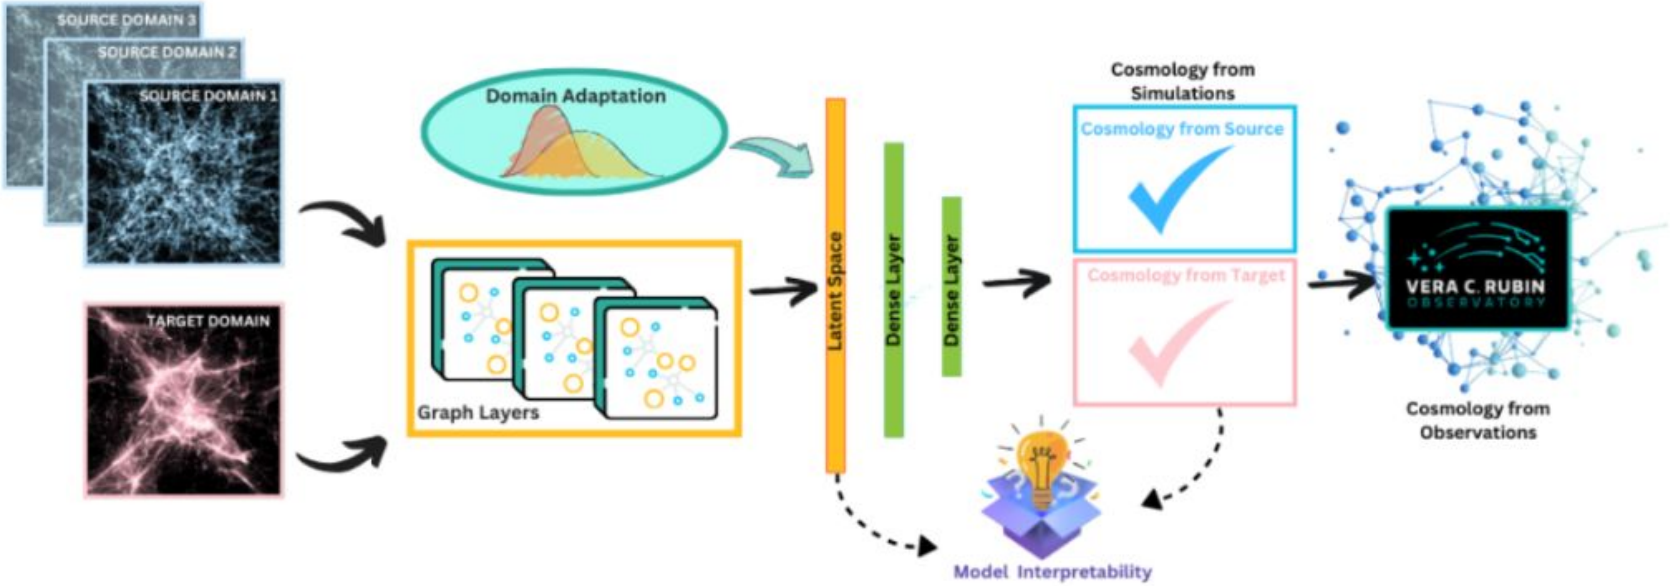

Cosmology with Graphs

Can we correctly infer cosmology across different simulations?

NeurIPS 2023.
Roncoli et al. 2023.

Graph Neural Networks: ideal for sparse galaxy catalogs!

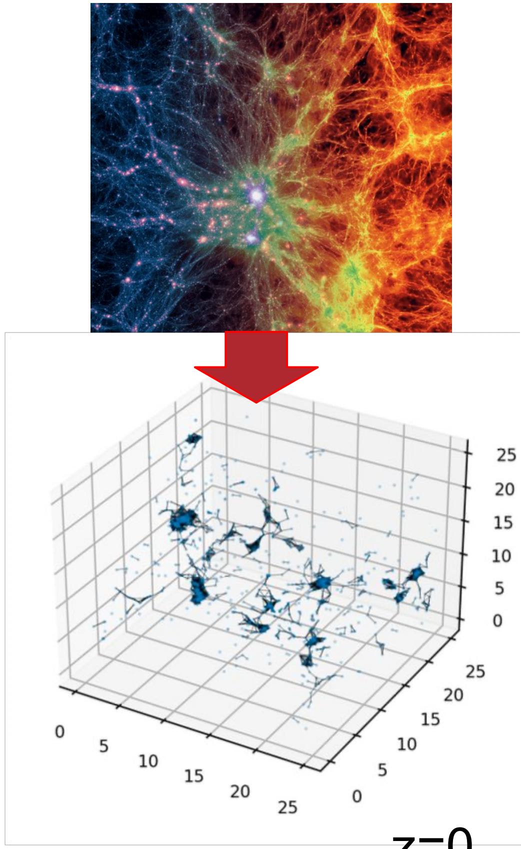

z=0 1000 simulations each

Source | Target | Latents

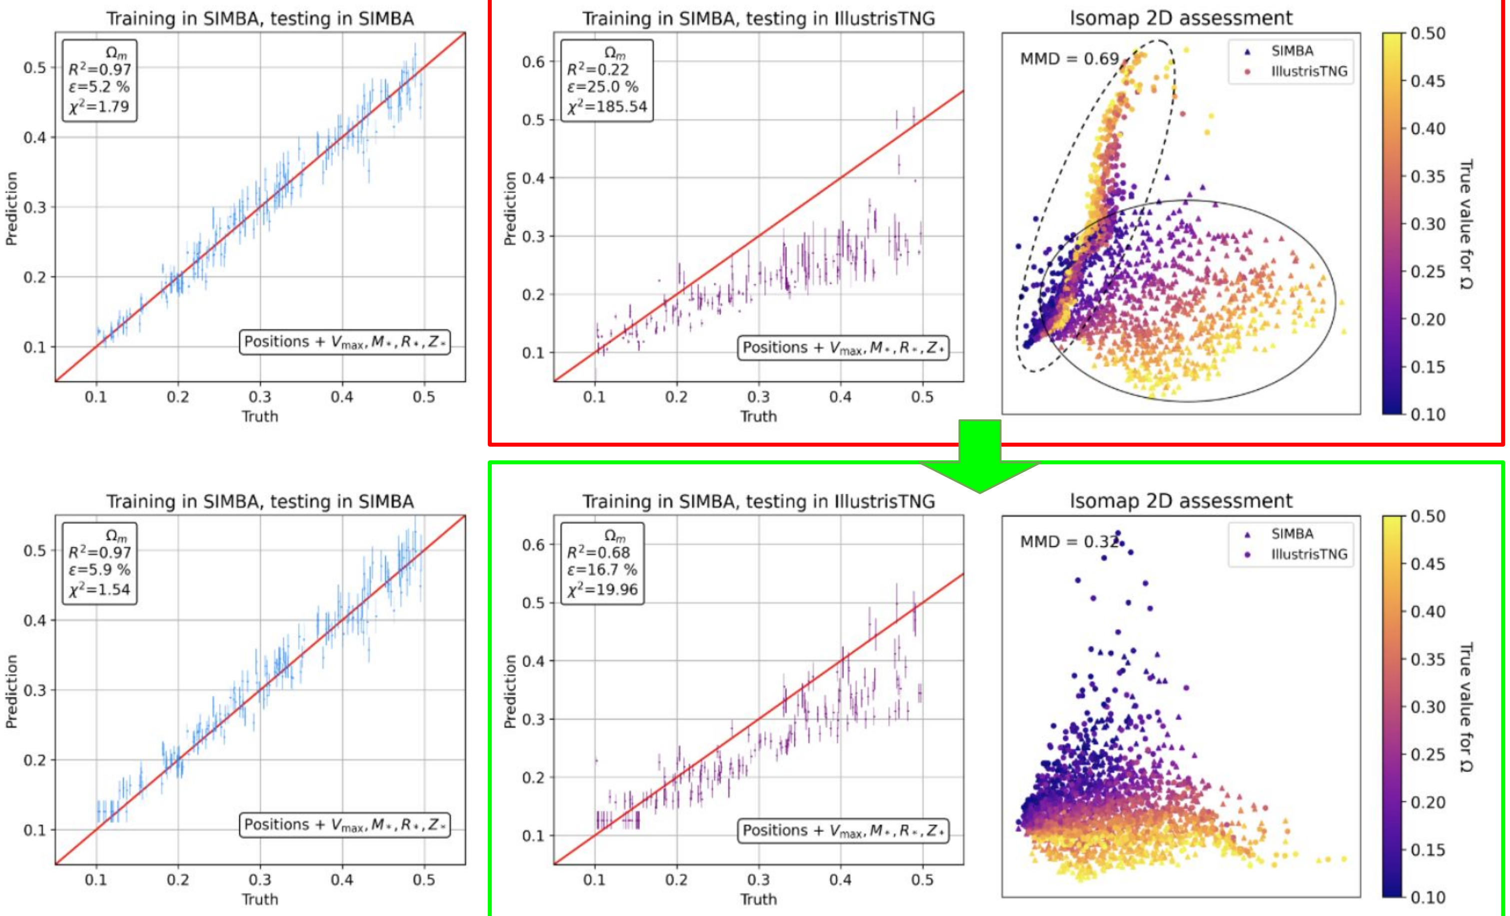

28% better relative error

Order of mag. better χ^2

FERMI NATIONAL ACCELERATOR LABORATORY

Strong gravitational lensing

Can we infer strong lensing parameters (Einstein radius, ellipticity, position/offset) robustly in both simulated and real data?

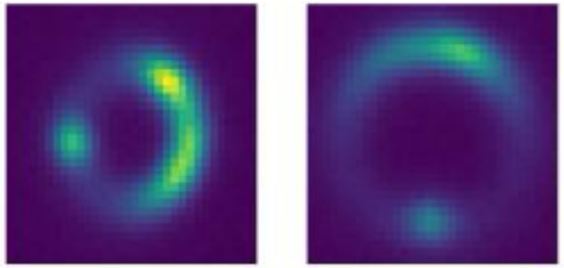

Source - Clean

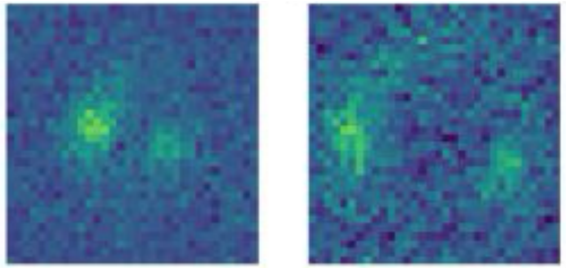

Target - DES

Strong gravitational lensing

Infer Einstein radius, ellipticity, position/offset

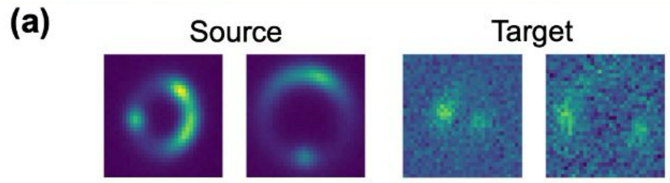

a. 100k images (1-filter) with and without noise
b. **Latent distributions are correctly aligned**.
c. Differences between true values and predicted posteriors are small. **Accuracy improves up to 2 orders of magnitude**.

Swierc et al. 2024.

Strong gravitational lensing

Infer Einstein radius, ellipticity, position/offset

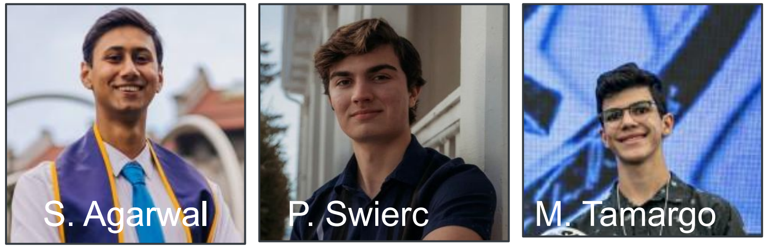

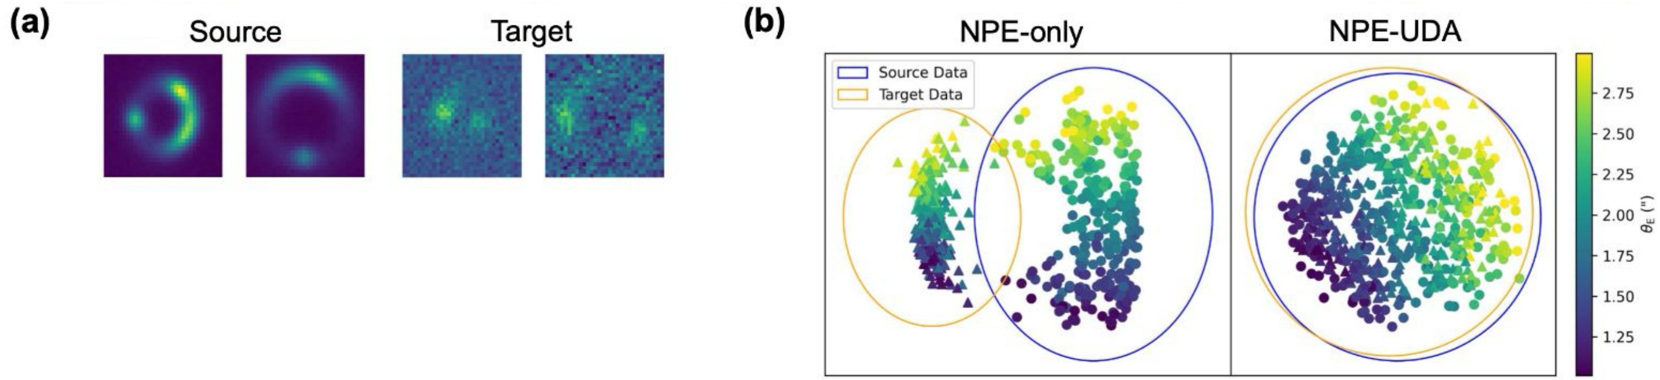

a. 100k images (1-filter) with and without noise
b. **Latent distributions are correctly aligned**.
c. Differences between true values and predicted posteriors are small. **Accuracy improves up to 2 orders of magnitude**.

Swierc et al. 2024.

Strong gravitational lensing

SBI - infer Einstein radius, ellipticity, position/offset

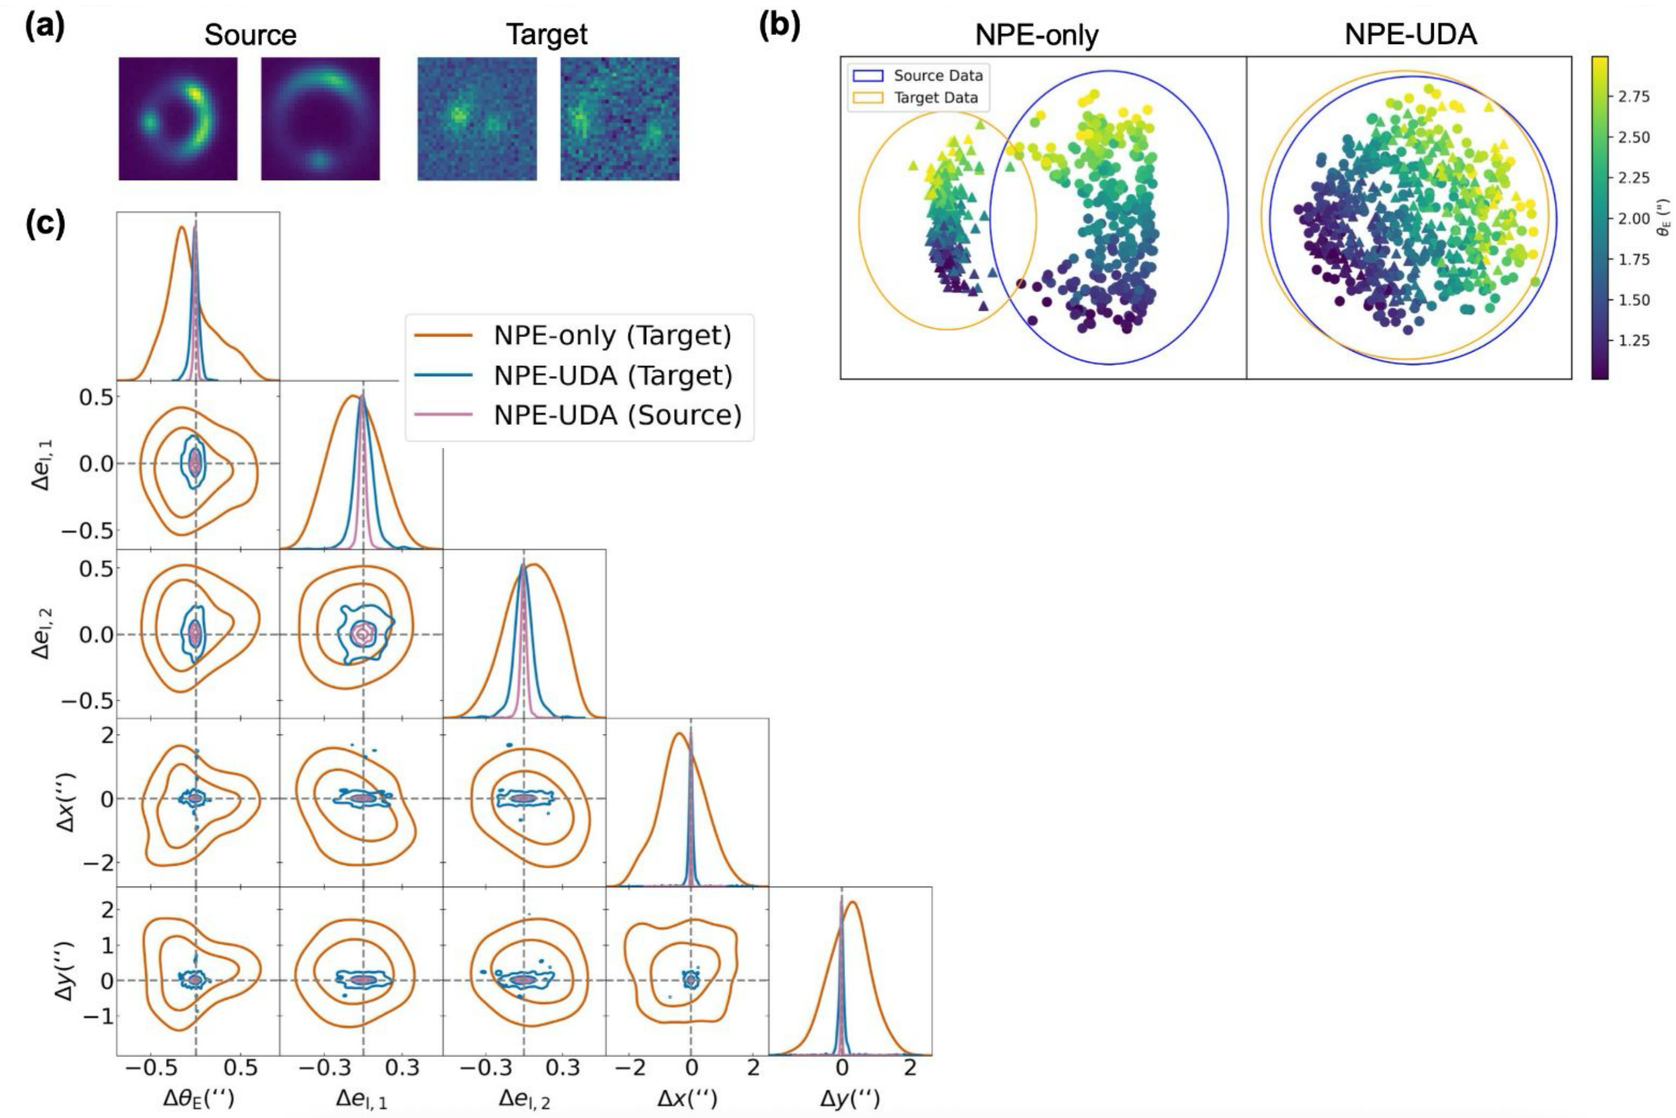

a. 100k images (1-filter) with and without noise
b. **Latent distributions are correctly aligned**.
c. Differences between true values and predicted posteriors are small. **Accuracy improves up to 2 orders of magnitude**.

Swierc et al. 2024.

04

Uncertainty Quantification (UQ)

Sources of uncertainties

All of this will influence model parameters and model outputs!

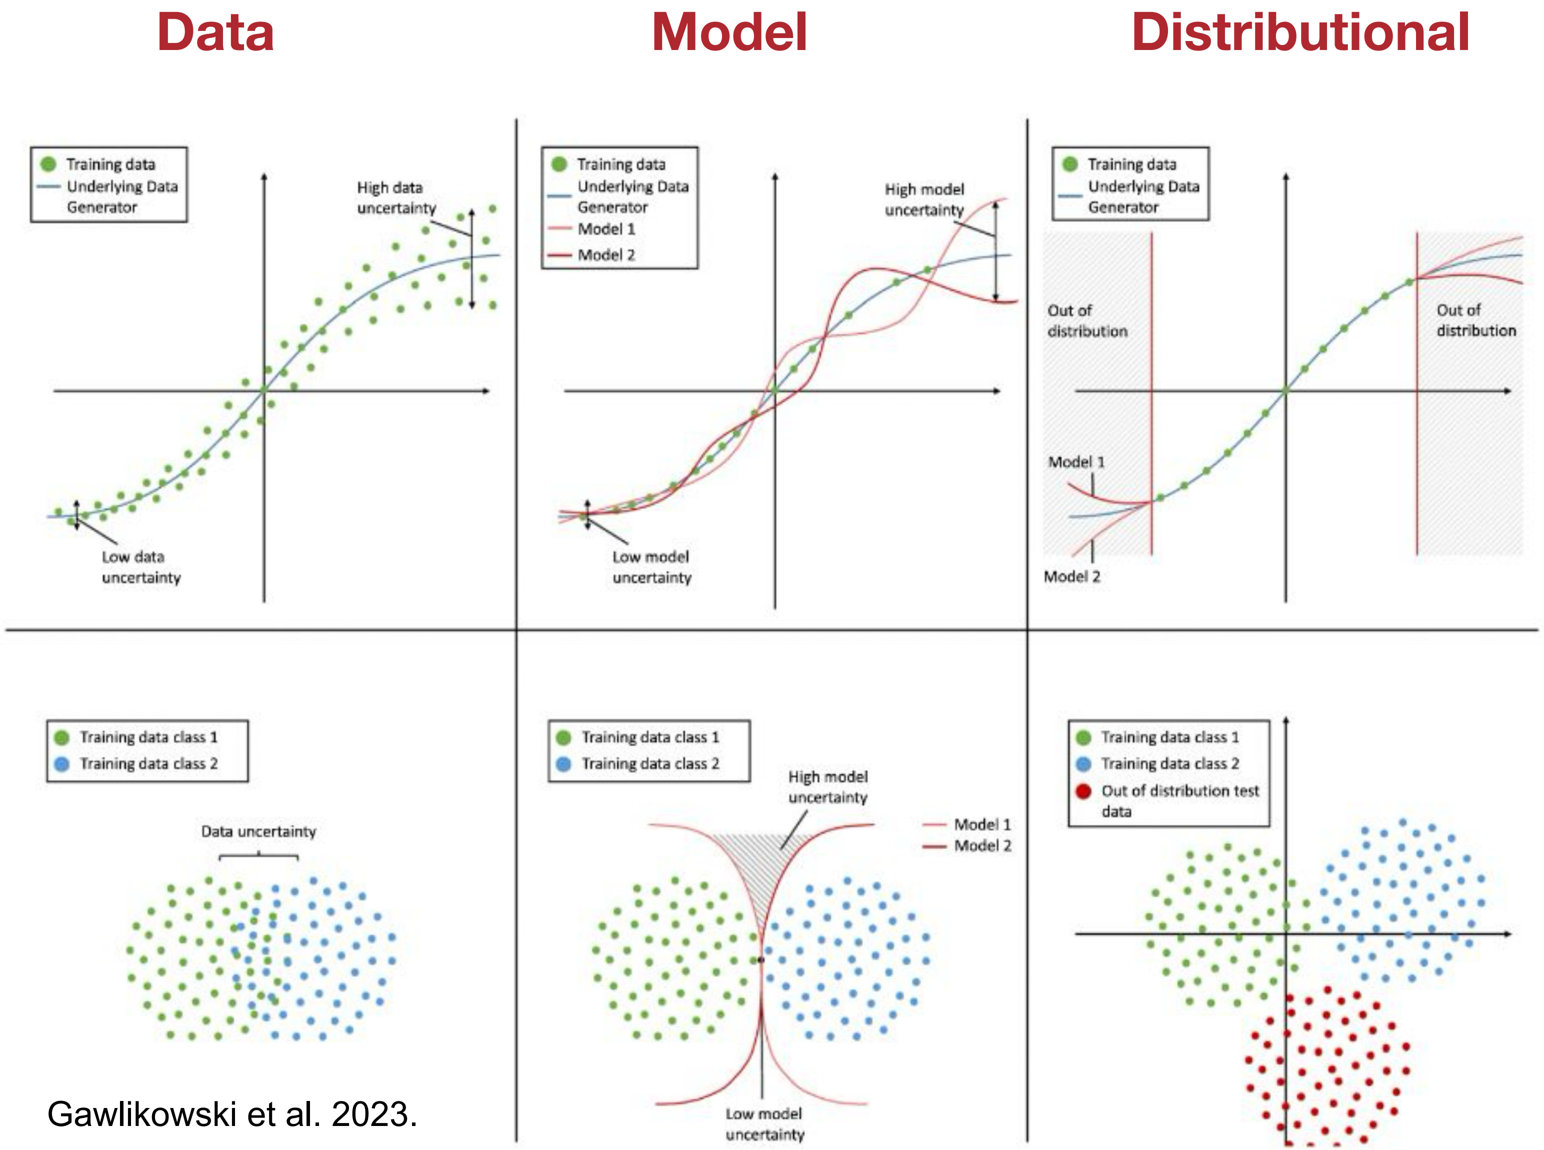

Gawlikowski et al. 2023.

Sources of uncertainties

All of this will influence model parameters and model outputs!

Systematic
persistent bias from instruments, calibration, or imperfect theory
Statistical
variation due to limited data

Epistemic
model or data-driven uncertainty
Aleatoric
irreducible noise in the data

These vocabularies are not interchangeable!

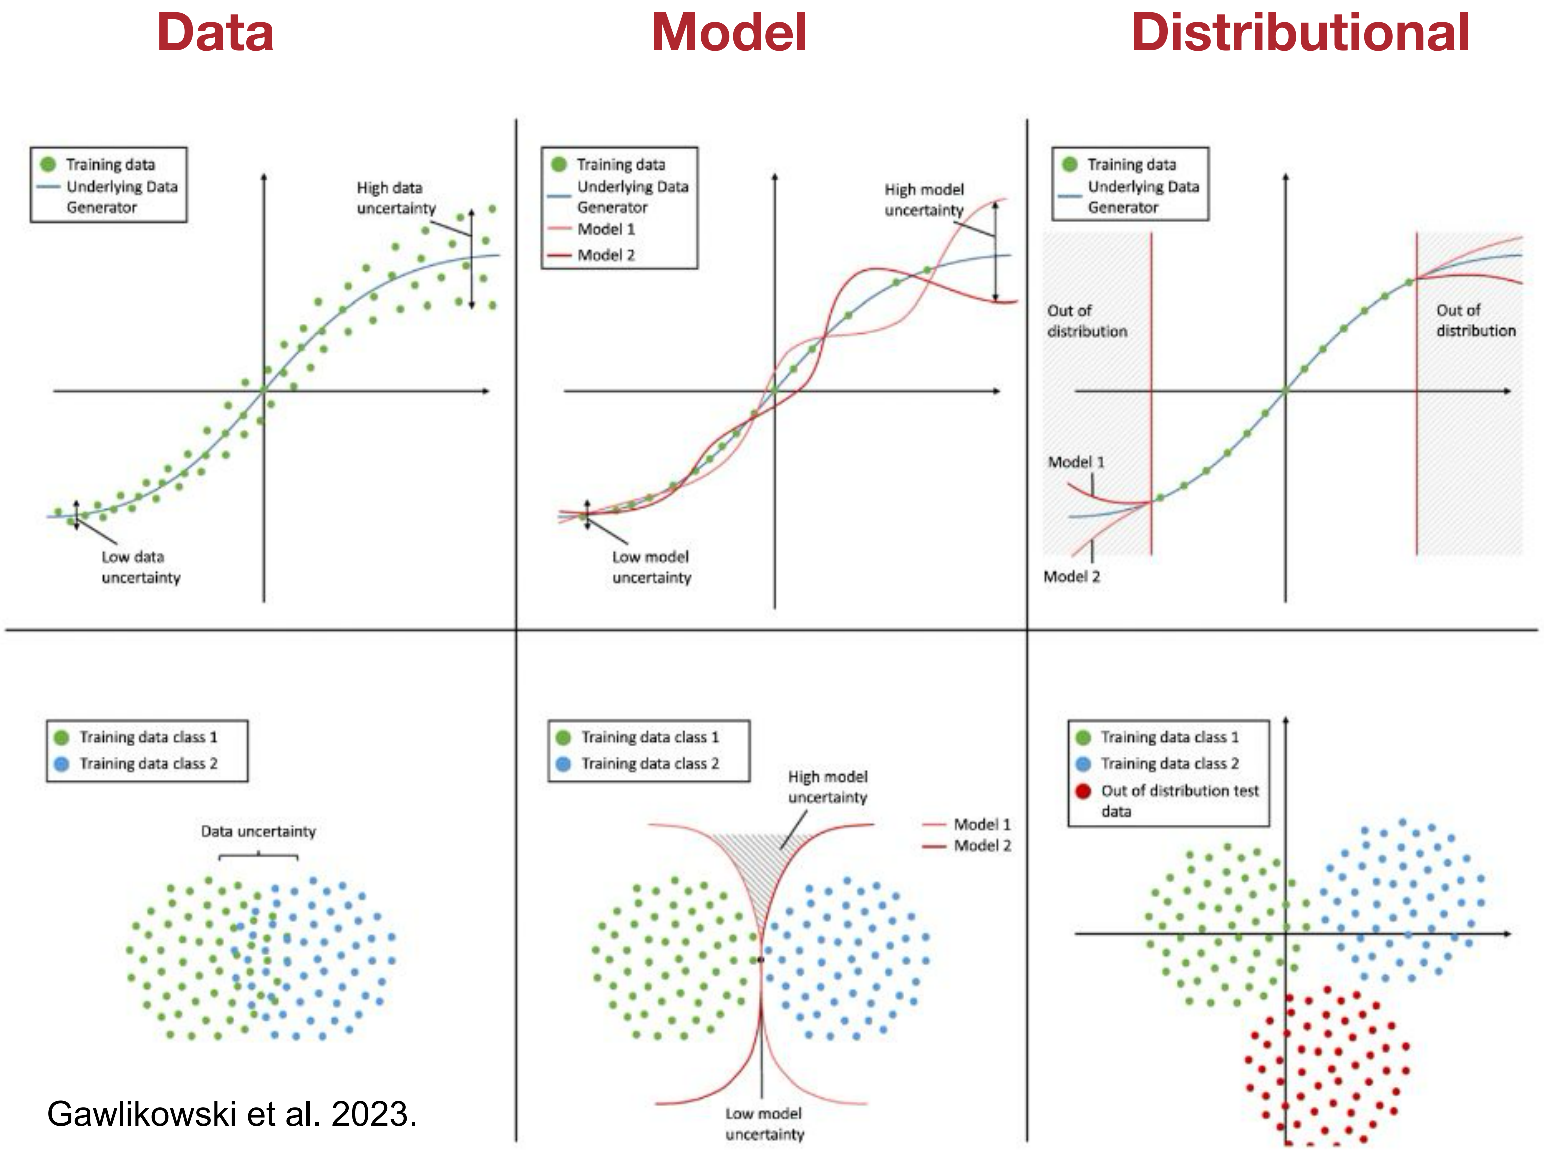

Gawlikowski et al. 2023.

Prediction / Inference with UQ

The way we think about AI models
and UQ is and should be evolving.

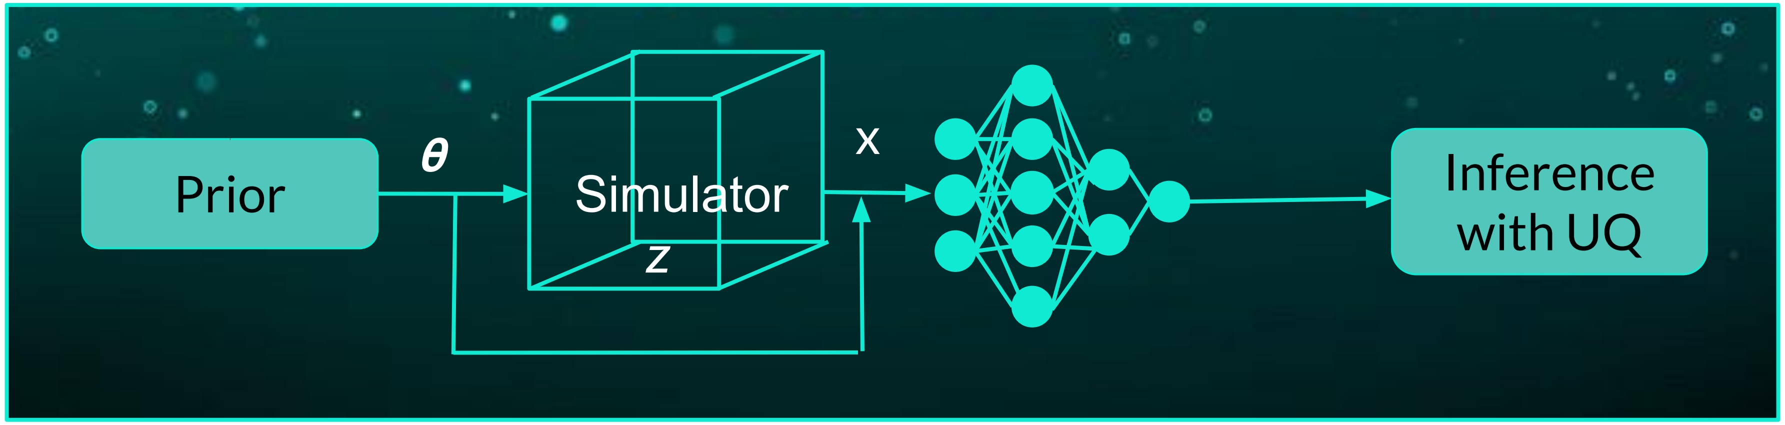

Prediction / Inference with UQ

The way we think about AI models and UQ is and should be evolving.

Deterministic Model

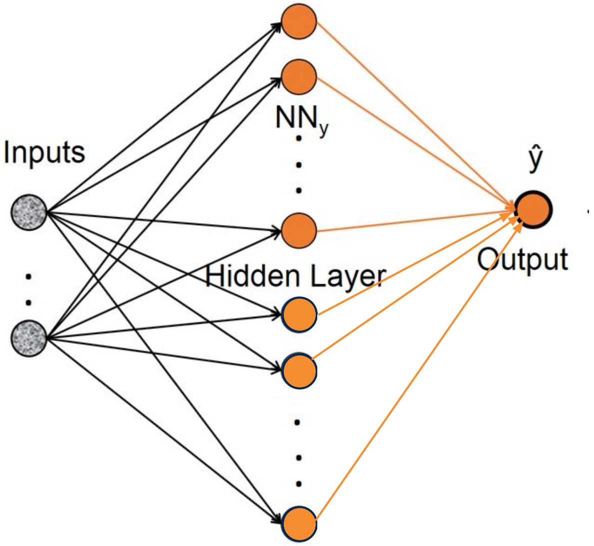

Point prediction

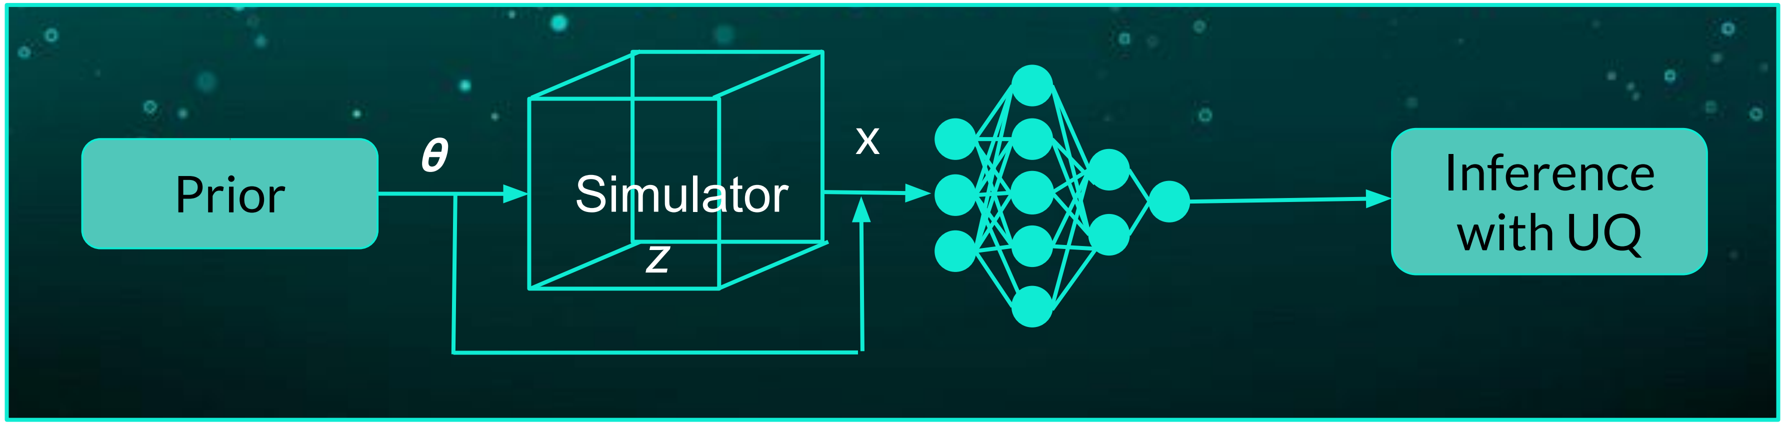

Prediction / Inference with UQ

The way we think about AI models and UQ is and should be evolving.

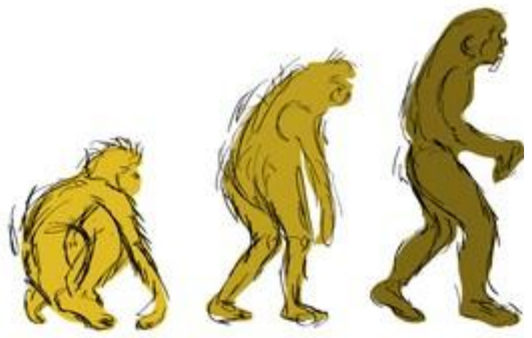

MVE / BNN

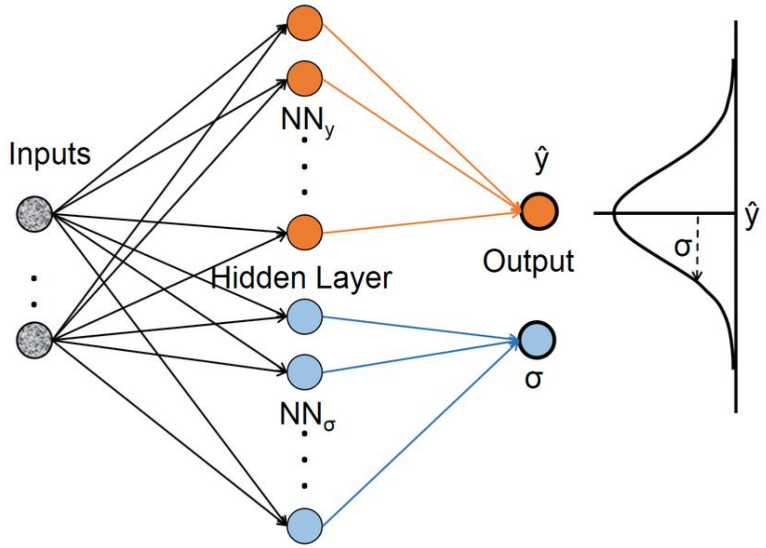

Point prediction

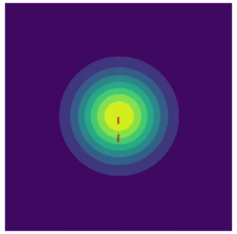

Some UQ

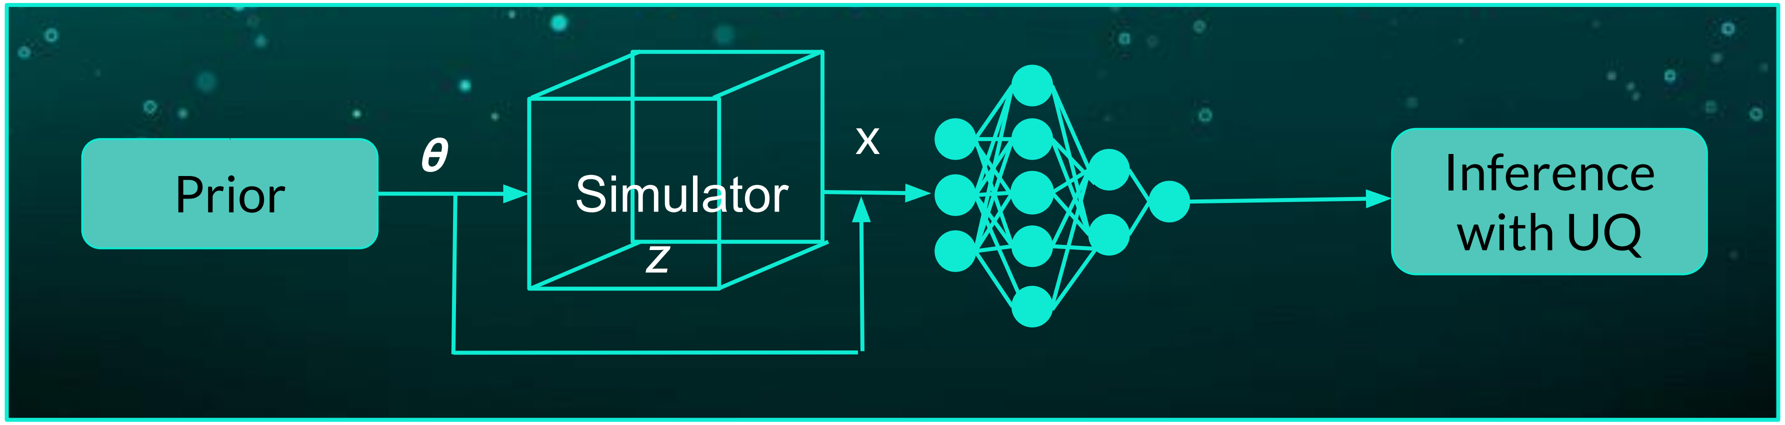

Prediction / Inference with UQ

The way we think about AI models and UQ is and should be evolving.

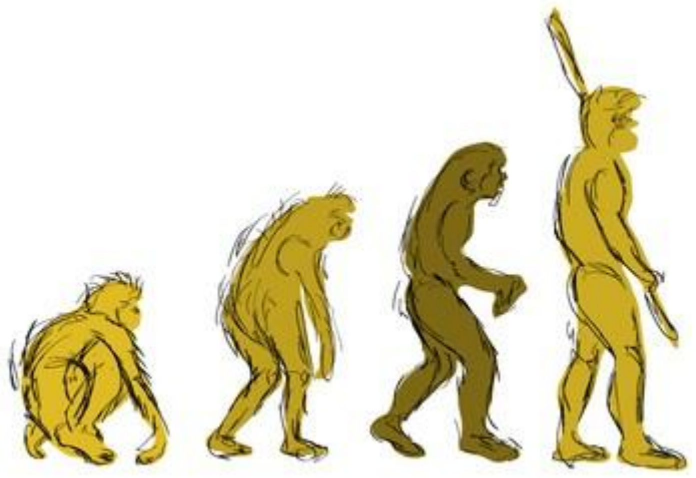

Neural Density Estimator

Normalizing Flows

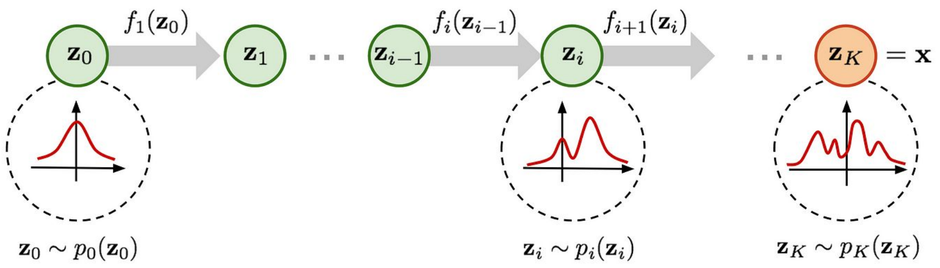

Simulation-Based Inference (SBI)

Point prediction

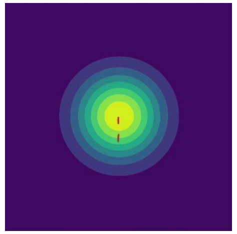

Some UQ

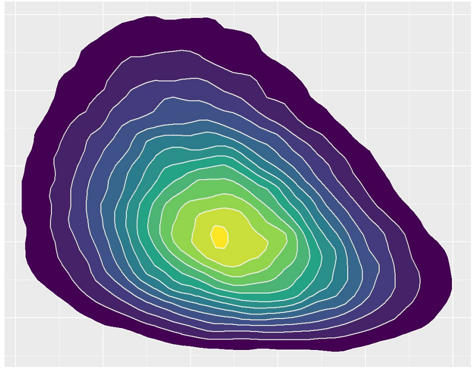

Improved UQ

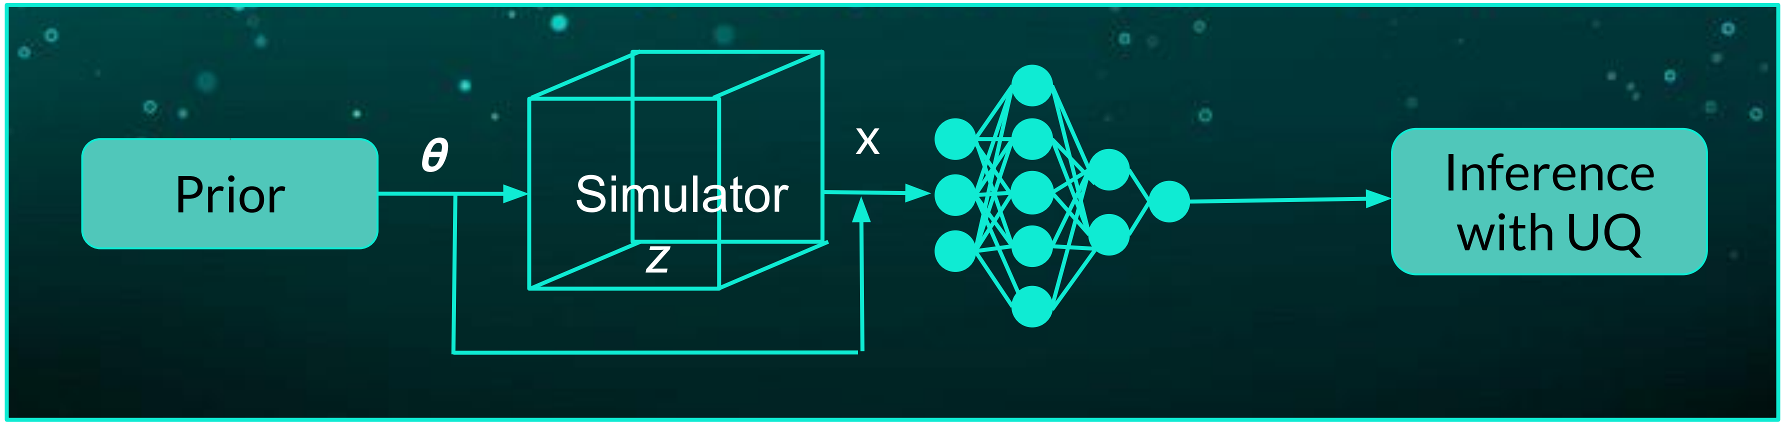

Taking the next step in UQ

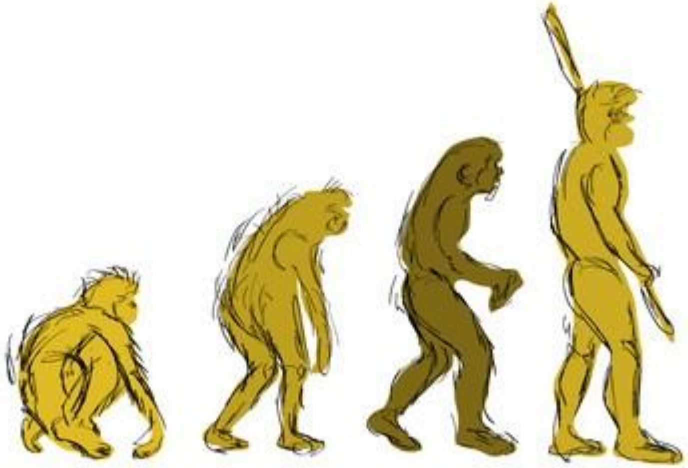

Taking the next step in UQ

Aligning Intent and Implementation in UQ for ML

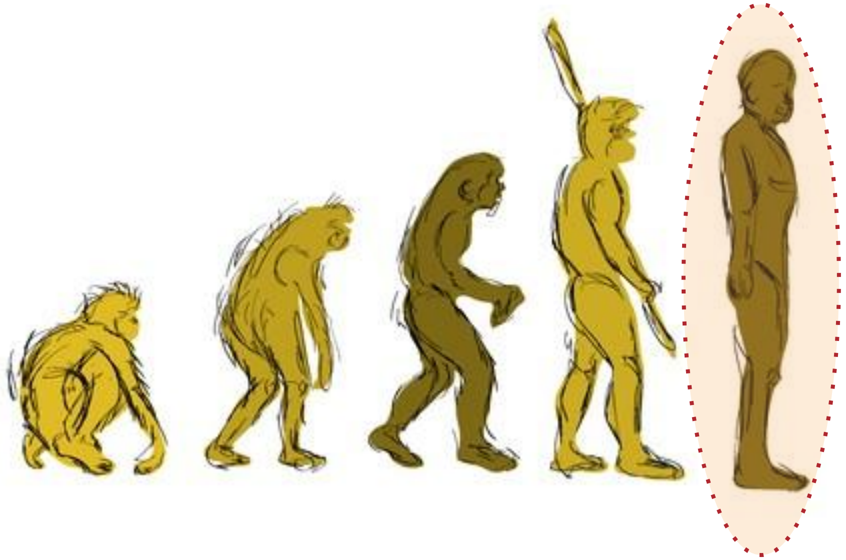

Taking the next step in UQ

Aligning Intent and Implementation in UQ for ML

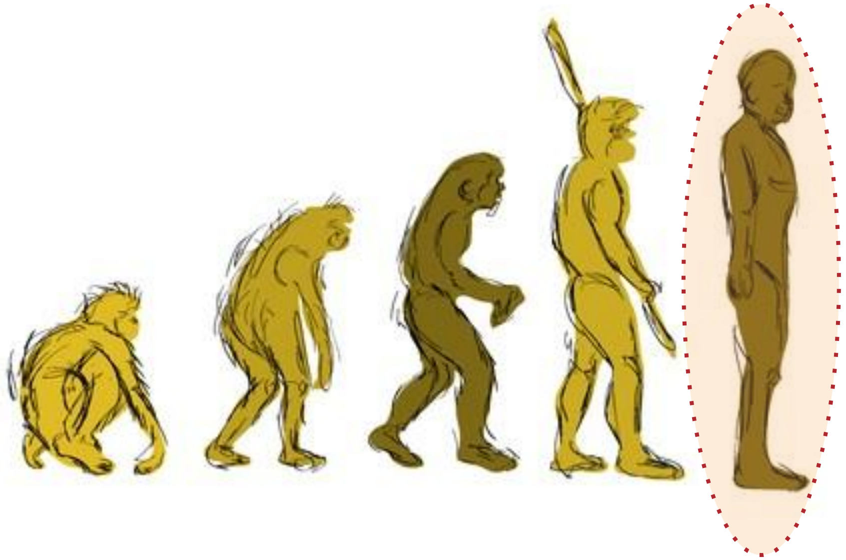

Improving UQ in scientific ML does not require solving foundational questions in the philosophy of science. But it does require practical discipline: defining what is being estimated, justifying the uncertainty attached to it, and validating whether that uncertainty supports the decisions or claims it is meant to inform. We advocate for this kind of *epistemic hygiene*—not as a constraint, but as an enabling structure. The payoff is not just cleaner semantics, but more decisive modeling. By aligning estimation targets with uncertainty constructs and validation tools, we enable models to play a trustworthy role in scientific inference—making uncertainty a vehicle for insight, not confusion, and supporting the kind of transparent, cumulative progress that scientific ML now urgently demands.

Trivedi & Nord 2025. arXiv:2506.03037

Taking the next step in UQ

Aligning Intent and Implementation in UQ for ML

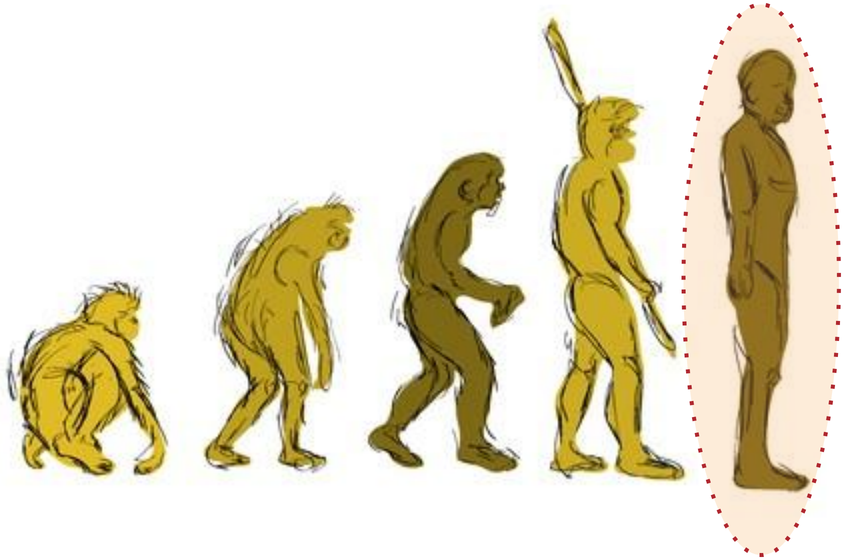

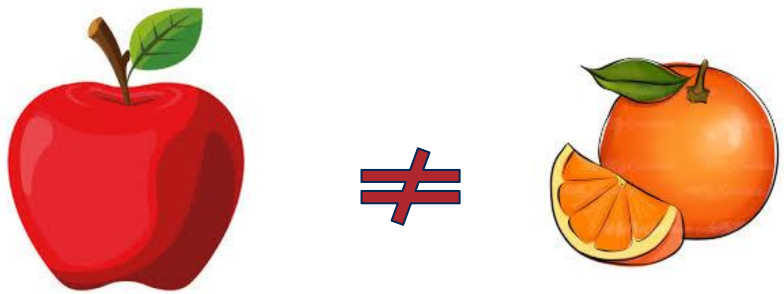

Improving UQ in scientific ML does not require solving foundational questions in the philosophy of science. But it does require practical discipline: defining what is being estimated, justifying the uncertainty attached to it, and validating whether that uncertainty supports the decisions or claims it is meant to inform. We advocate for this kind of *epistemic hygiene*—not as a constraint, but as an enabling structure. The payoff is not just cleaner semantics, but more decisive modeling. By aligning estimation targets with uncertainty constructs and validation tools, we enable models to play a trustworthy role in scientific inference—making uncertainty a vehicle for insight, not confusion, and supporting the kind of transparent, cumulative progress that scientific ML now urgently demands.

Trivedi & Nord 2025. arXiv:2506.03037

Taking the next step in UQ

Aligning Intent and Implementation in UQ for ML

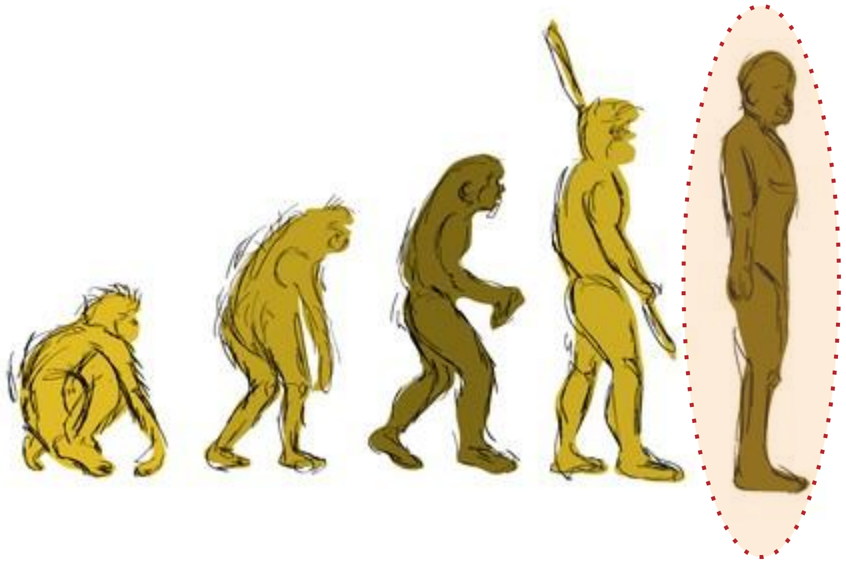

- **Declare the inference chain**
  the estimation target, the loss or decision goal, the uncertainty construct and what it means.
- **Match UQ method to context**
  what is the uncertainty is meant to support (forcasting, control, hypothesis testing)
- **Check both forward and inverse validity**
  in data and parameter space
- **Use the simulator as an instrument**
  not just to train, but to test: perturb parameters and measure UQ behavior
- **Benchmark beyond IID**
  domain shifts
- **Study model stability**
  sensitivity to dataset realizations, initialization seeds - stability is a proxy for epistemic trustworthiness

Trivedi & Nord 2025. arXiv:2506.03037

- Era of big astro surveys - **large amount of data but not all of it labeled.**
- Simulations and old data are different - **domain shift problem**!
- **Domain Adaptation** can help but:
  - we need more work related to **model interpretability**;
  - better **understanding of data and model errors;**
  - **correctly choose UQ metric** for a given type of task.

- Era of big astro surveys - **large amount of data but not all of it labeled.**
- Simulations and old data are different - **domain shift problem**!
- **Domain Adaptation** can help but:
  - we need more work related to **model interpretability**;
  - better **understanding of data and model errors;**
  - **correctly choose UQ metric** for a given type of task.

THANK YOU!

aleksand@fnal.gov

Fermilab

FermiFORWARD

U.S. DEPARTMENT of ENERGY

FERMI NATIONAL ACCELERATOR LABORATORY