When a computer scientist meets a physicist

Gilles Louppe
g.louppe@uliege.be
November 29, 2024

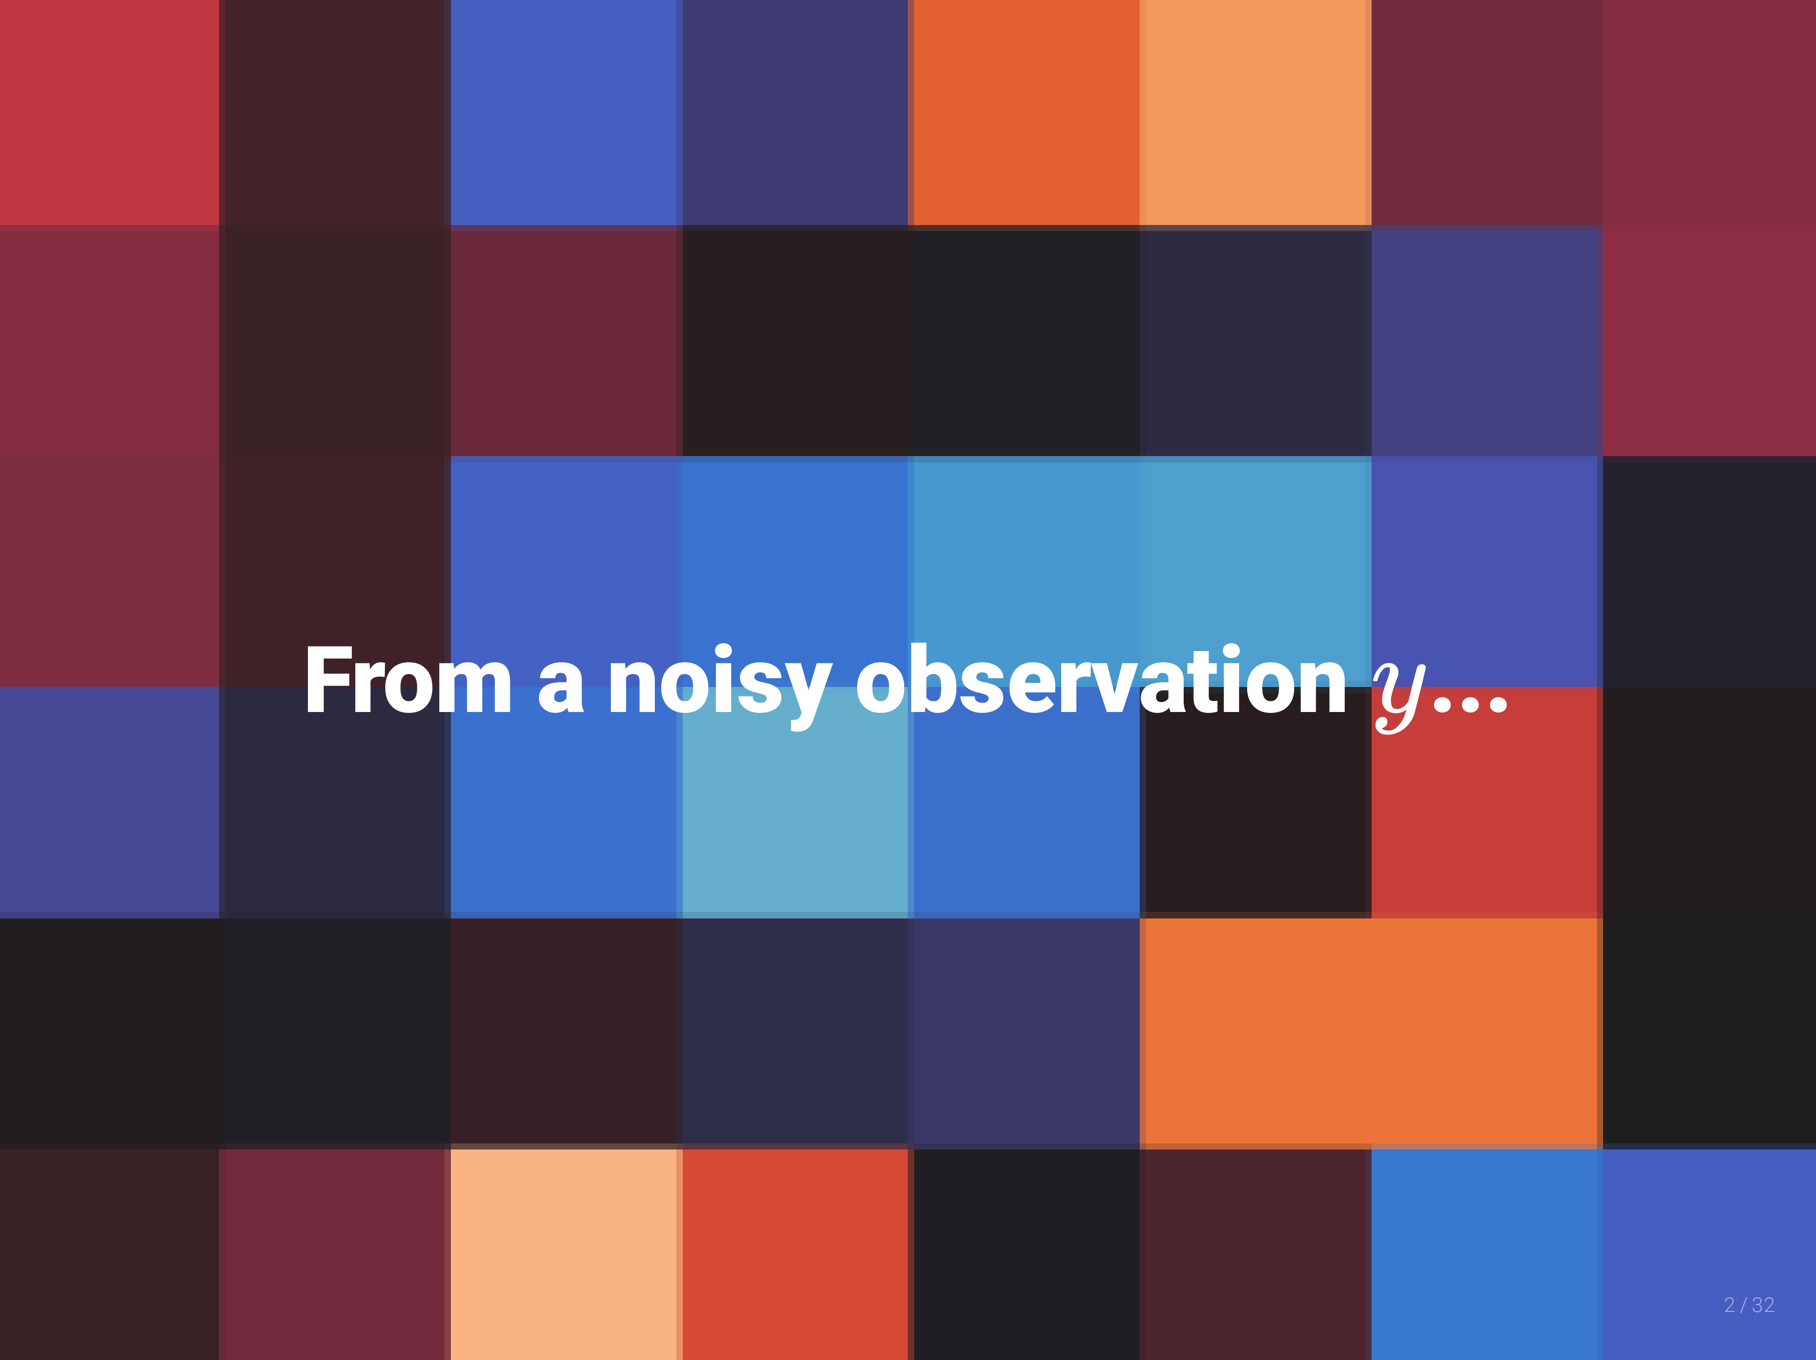

From a noisy observation y...

... can we recover images x?

$$\dot{u} = -u\nabla u + \frac{1}{Re}\nabla^2 u - \frac{1}{\rho}\nabla p + f$$

$$0 = \nabla \cdot u$$

... or parameters $\theta = \{Re, \rho, f\}$?

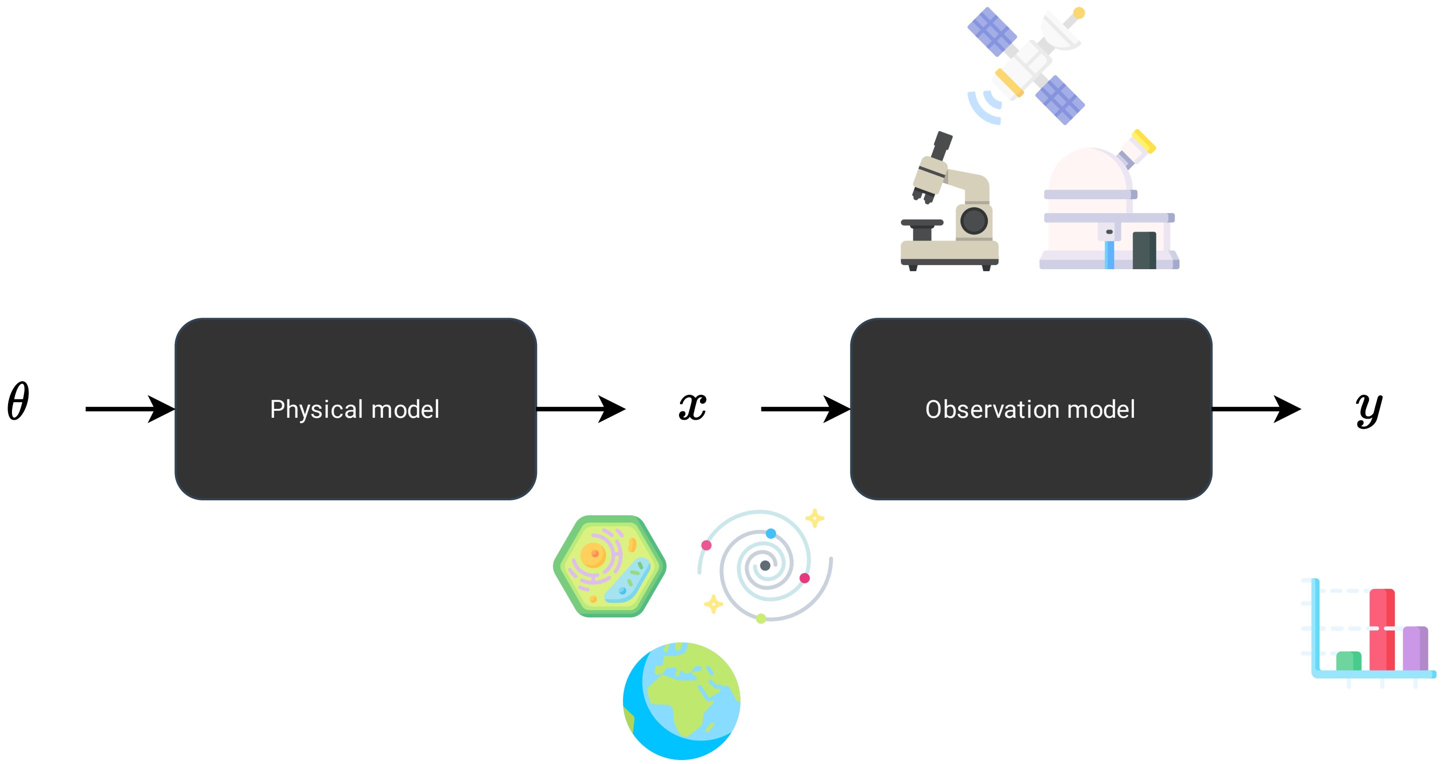

Goal: Estimate parameters $\boldsymbol{\theta}$ or latents $\mathbf{x}$ from noisy or incomplete observations $\mathbf{y}$.

AI for particle physics at the LHC

Kyle Cranmer
(New York University)

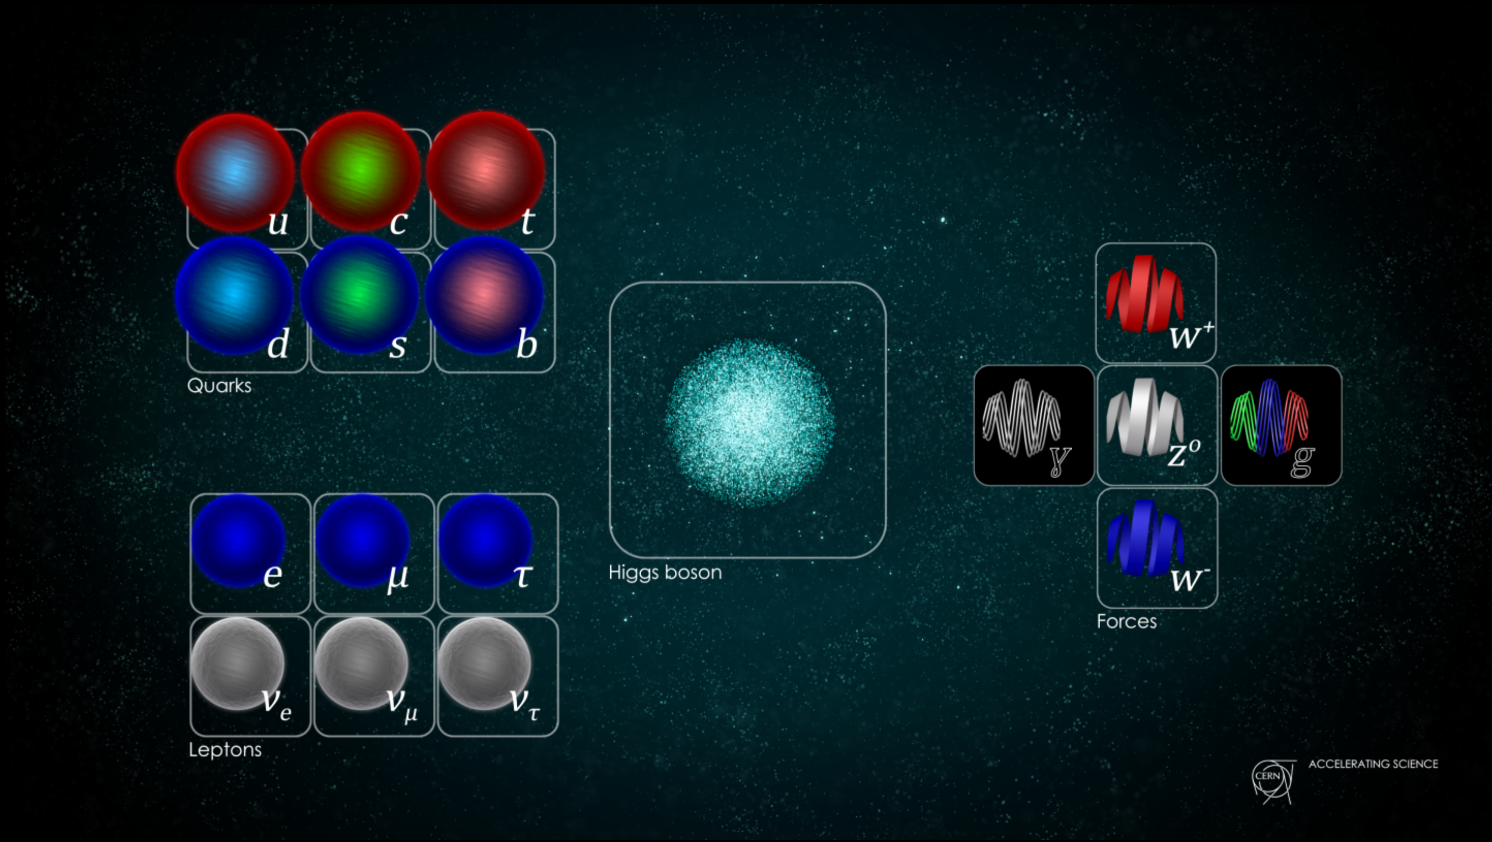

$$\theta = \{m_e, m_\mu, m_\tau, \ldots\}$$

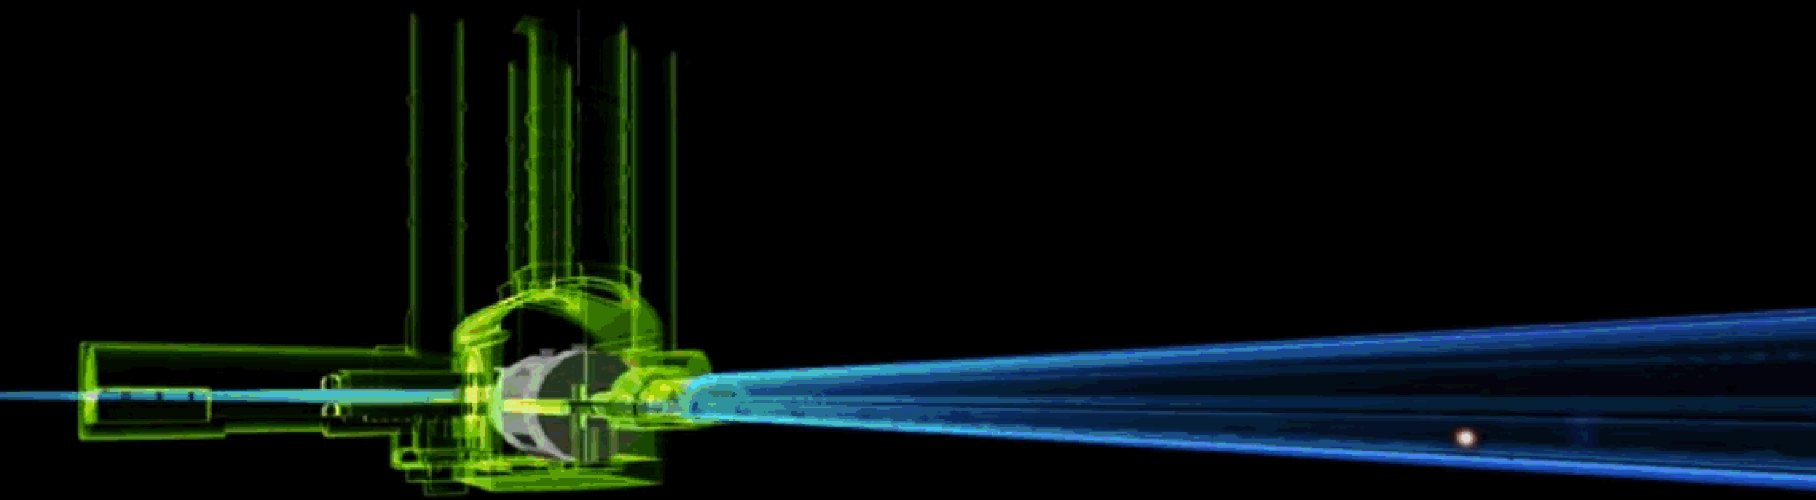

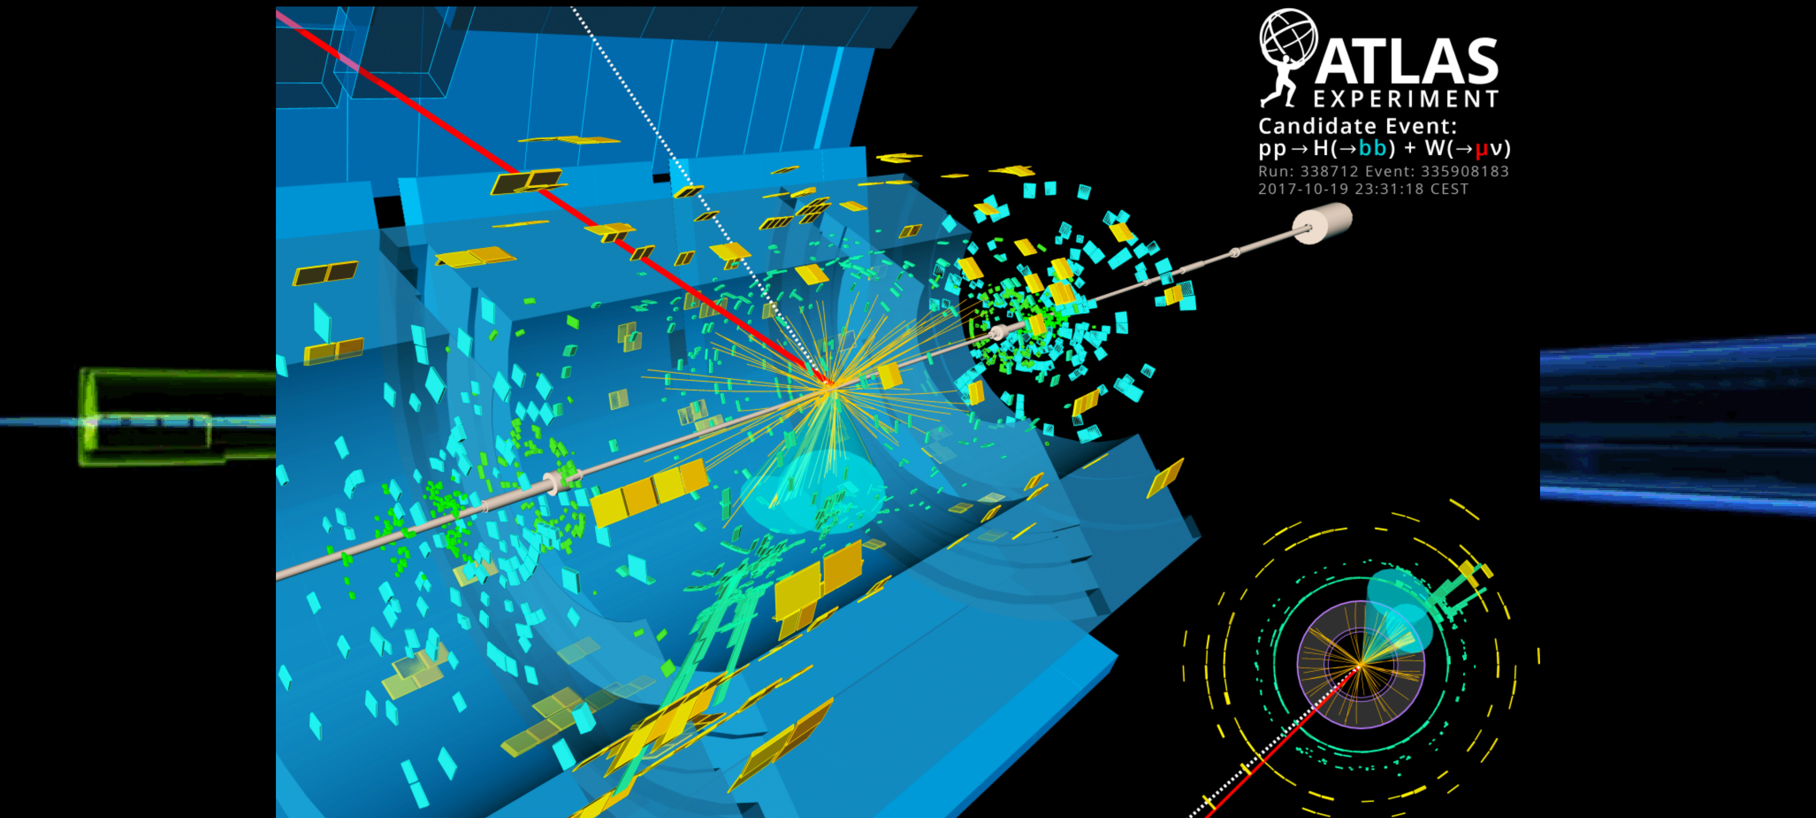

$$\arg\max_{\theta} p(\mathbf{y}|\theta)?$$

pre-2019

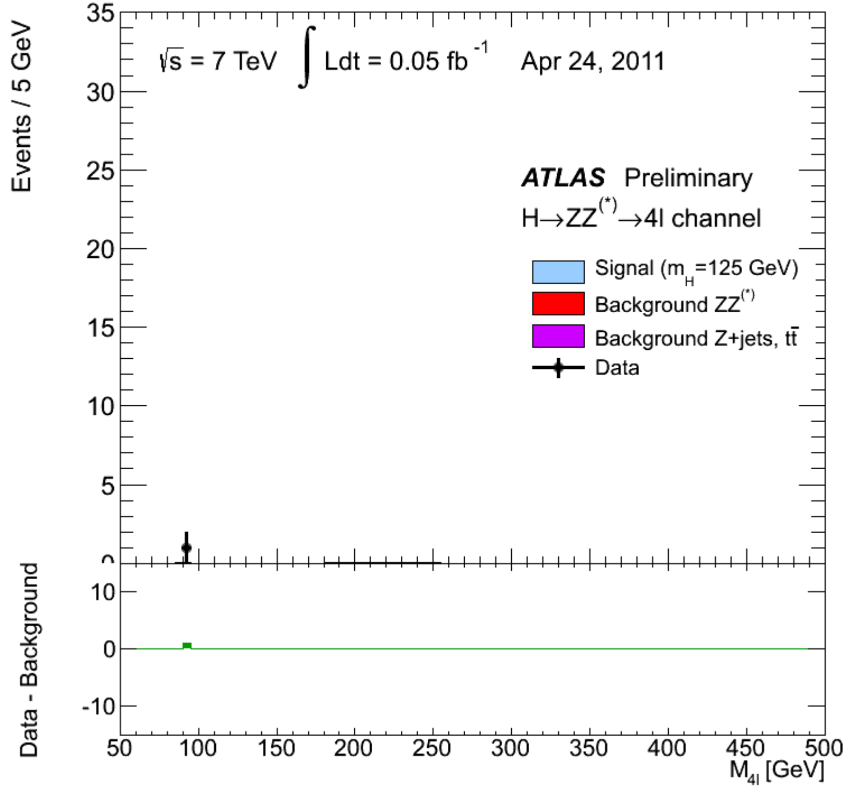

Reject the null hypothesis that the Higgs boson does not exist by a likelihood-ratio test $\lambda(\theta) = -2\log p(\mathbf{x}|\theta)/p(\mathbf{x}|\hat{\theta})$, where the likelihood $p(\mathbf{x}|\theta)$ is approximated as $p(s(\mathbf{x})|\theta)$, for some summary statistic $s(\cdot)$.

Wait a minute Kyle... how do we pick the right summary statistic s?

Neural simulation-based inference

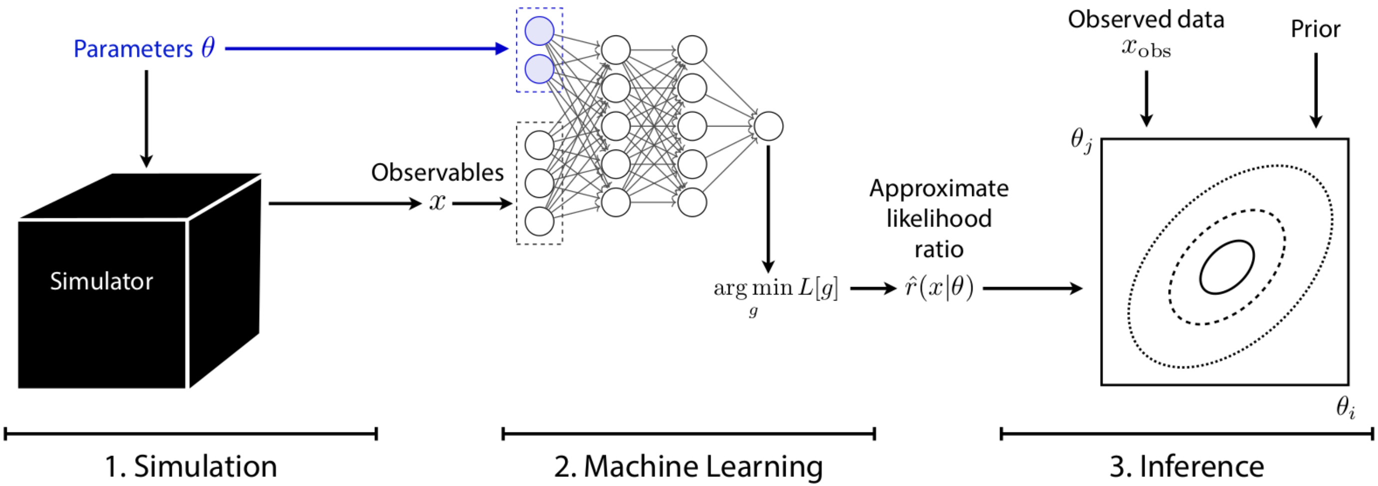

Run simulator and save data

Train NN classifier, interpret as likelihood ratio estimator

Amortized: cheap to repeat for new data

Learn the statistic $s(\cdot)$ with a neural network approximating the likelihood ratio $r(\mathbf{x}|\theta) = p(\mathbf{x}|\theta)/p_{\text{ref}}(\mathbf{x})$.

—

Credits: Cranmer et al, 2015; Brehmer et al, 2018.

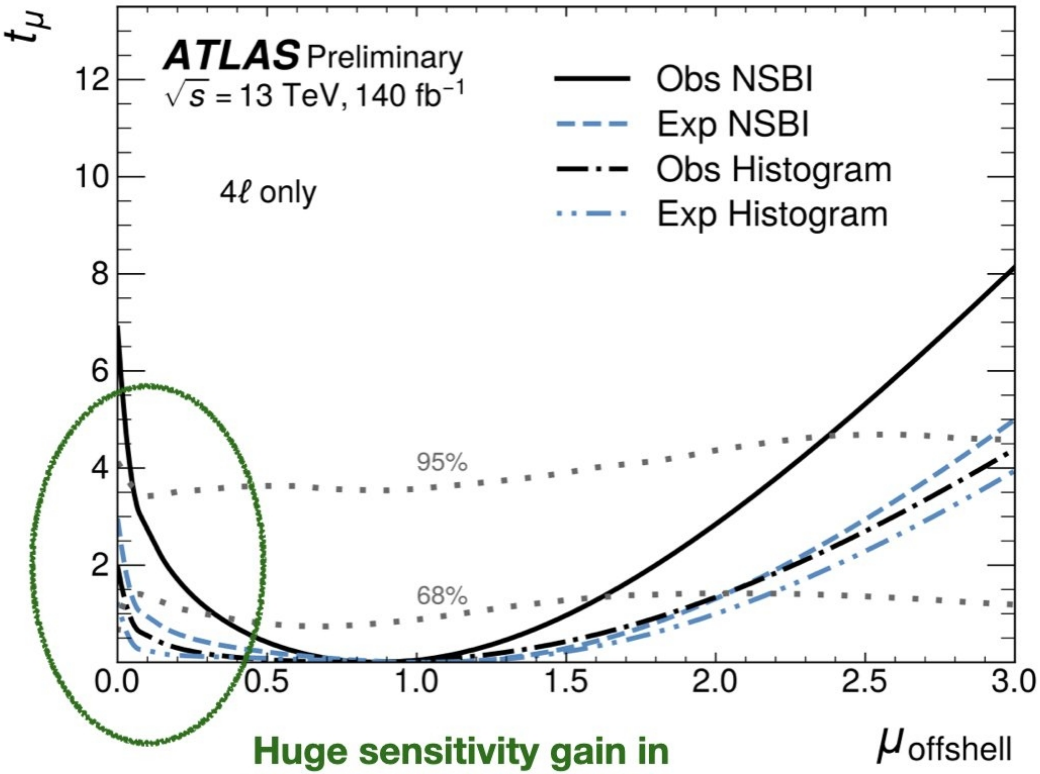

Huge sensitivity gain in interference rich regions

$$\sqrt{\mu} \cdot p_I(x) \gg \mu \cdot p_S(x)$$

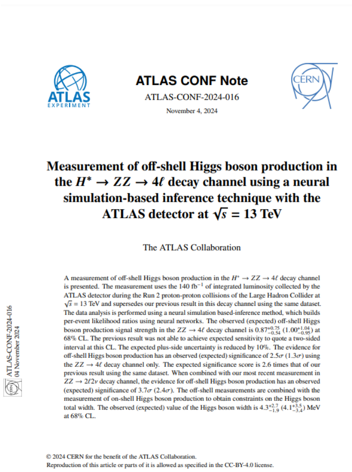

ATLAS CONF Note

ATLAS-CONF-2024-016

November 4, 2024

Measurement of off-shell Higgs boson production in the $H^* \to ZZ \to 4\ell$ decay channel using a neural simulation-based inference technique with the ATLAS detector at $\sqrt{s} = 13$ TeV

The ATLAS Collaboration

A measurement of off-shell Higgs boson production in the $H^* \to ZZ \to 4\ell$ decay channel is presented. The measurement uses the 140 fb^{-1} of integrated luminosity collected by the ATLAS detector during the Run 2 proton-proton collisions of the Large Hadron Collider at $\sqrt{s} = 13$ TeV and supersedes our previous result in this decay channel using the same dataset. The data analysis is performed using a neural simulation based-inference method, which builds per-event likelihood ratios using neural networks. The observed (expected) off-shell Higgs boson production signal strength in the $ZZ \to 4\ell$ decay channel is $0.87^{+0.75}_{-0.54}$ ($1.00^{+1.04}_{-0.95}$) at 68% CL. The previous result was not able to achieve expected sensitivity to quote a two-sided interval at this CL. The expected plus-side uncertainty is reduced by 10%. The evidence for off-shell Higgs boson production has an observed (expected) significance of 2.5σ (1.3σ) using the $ZZ \to 4\ell$ decay channel only. The expected significance score is 2.6 times that of our previous result using the same dataset. When combined with our most recent measurement in $ZZ \to 2\ell 2\nu$ decay channel, the evidence for off-shell Higgs boson production has an observed (expected) significance of 3.7σ (2.4σ). The off-shell measurements are combined with the measurement of on-shell Higgs boson production to obtain constraints on the Higgs boson total width. The observed (expected) value of the Higgs boson width is $4.3^{+2.7}_{-1.9}$ ($4.1^{+3.5}_{-3.4}$) MeV at 68% CL.

ATLAS-CONF-2024-016
04 November 2024

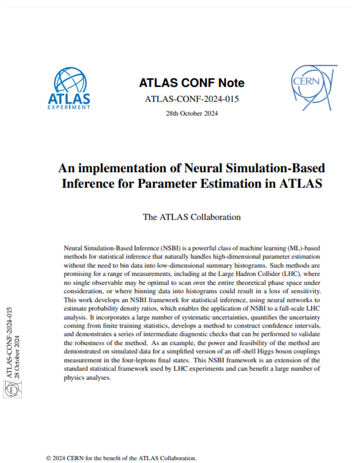

ATLAS CONF Note

ATLAS-CONF-2024-015

28th October 2024

An implementation of Neural Simulation-Based Inference for Parameter Estimation in ATLAS

The ATLAS Collaboration

Neural Simulation-Based Inference (NSBI) is a powerful class of machine learning (ML)-based methods for statistical inference that naturally handles high-dimensional parameter estimation without the need to bin data into low-dimensional summary histograms. Such methods are promising for a range of measurements, including at the Large Hadron Collider (LHC), where no single observable may be optimal to scan over the entire theoretical phase space under consideration, or where binning data into histograms could result in a loss of sensitivity. This work develops an NSBI framework for statistical inference, using neural networks to estimate probability density ratios, which enables the application of NSBI to a full-scale LHC analysis. It incorporates a large number of systematic uncertainties, quantifies the uncertainty coming from finite training statistics, develops a method to construct confidence intervals, and demonstrates a series of intermediate diagnostic checks that can be performed to validate the robustness of the method. As an example, the power and feasibility of the method are demonstrated on simulated data for a simplified version of an off-shell Higgs boson couplings measurement in the four-leptons final states. This NSBI framework is an extension of the standard statistical framework used by LHC experiments and can benefit a large number of physics analyses.

ATLAS-CONF-2024-015
28 October 2024

—

Credits: ATLAS Collaboration, 2024a; ATLAS Collaboration, 2024b.

Cosmological inference from stellar streams

Christoph Weniger
(University of Amsterdam)

What is the nature of dark matter?

$$\theta = m_{\mathrm{WDM}}$$

Constraining dark matter with stellar streams

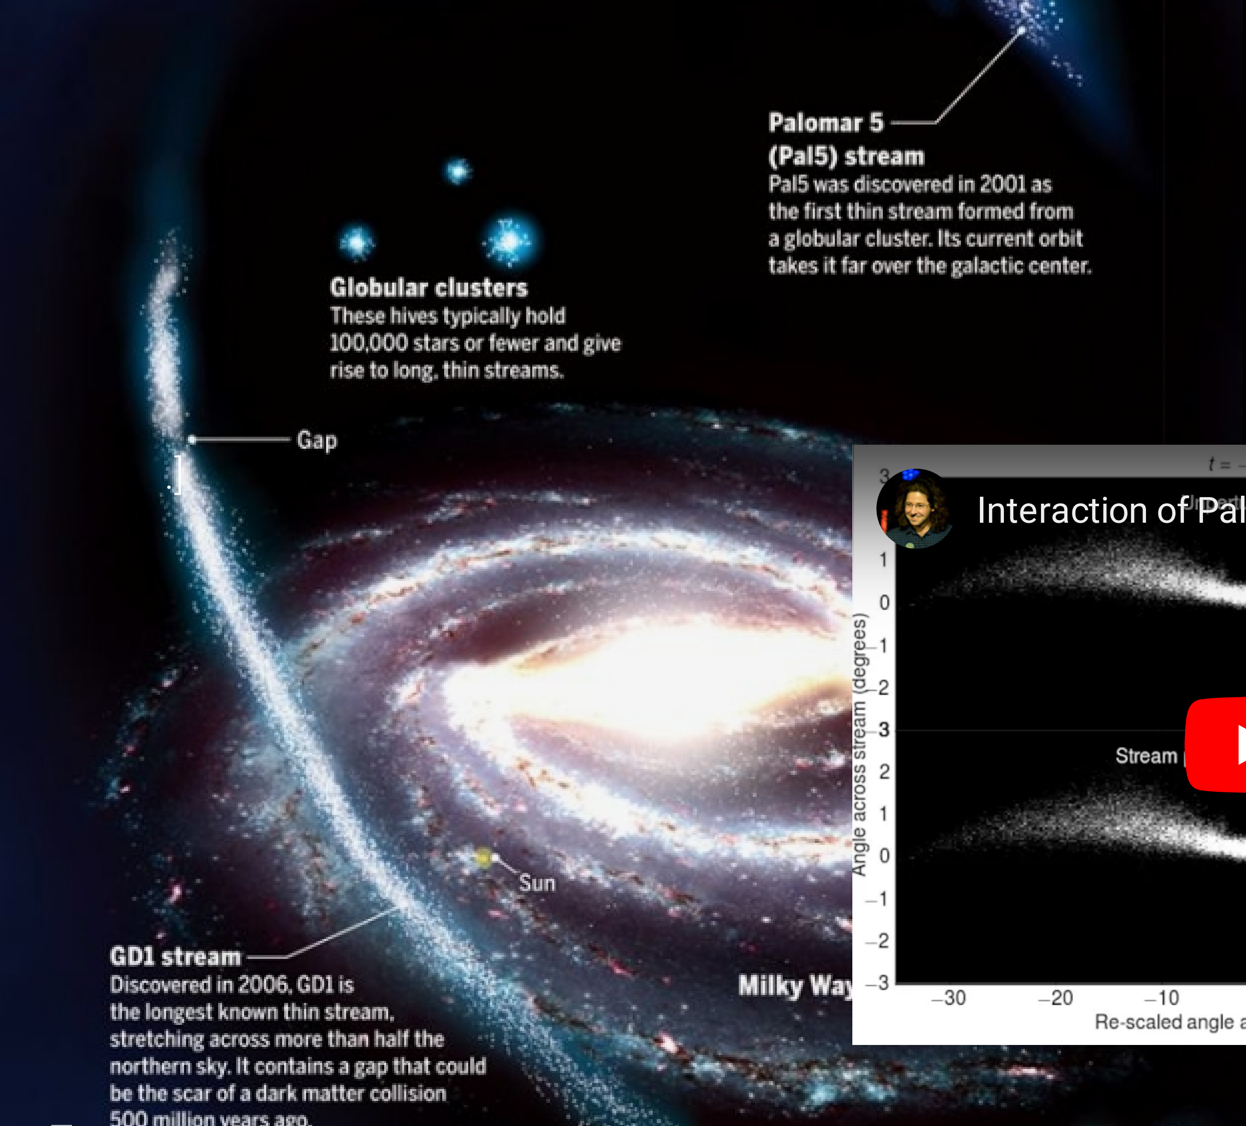

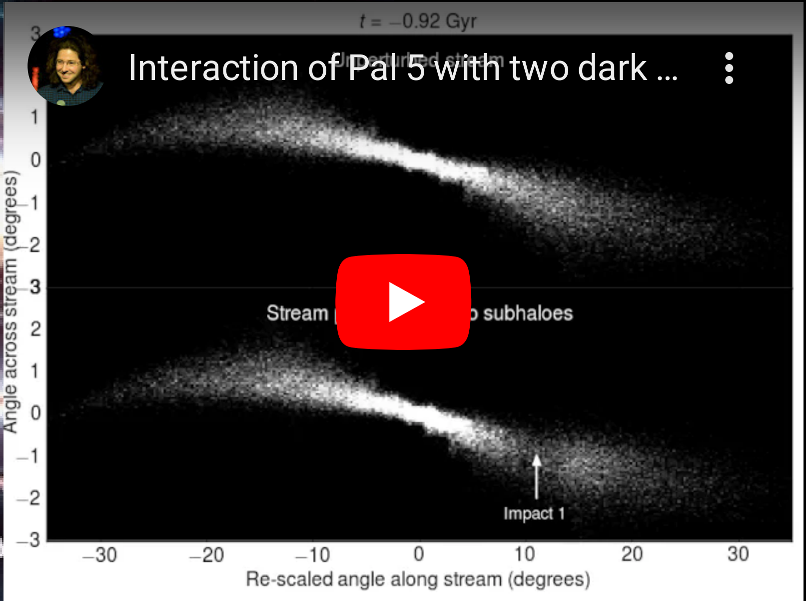

Image credits: C. Bickel/Science; D. Erkal.

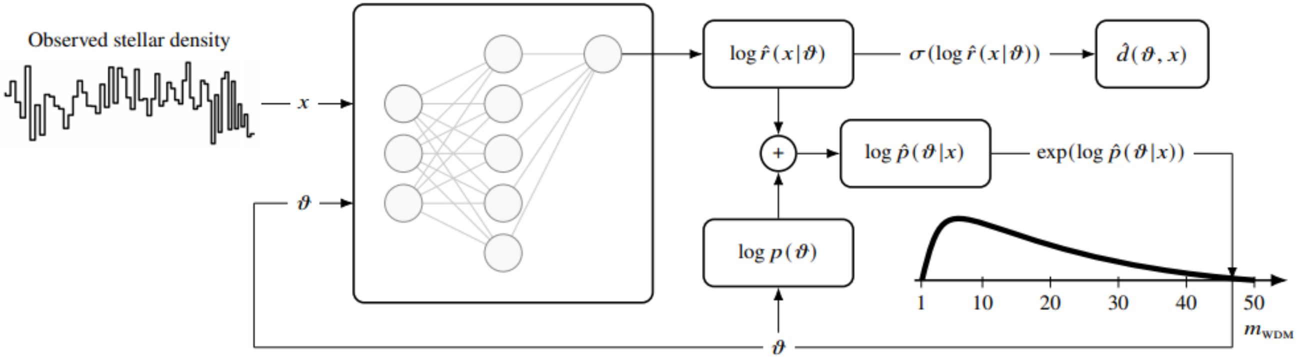

—

Credits: Hermans et al, 2021.

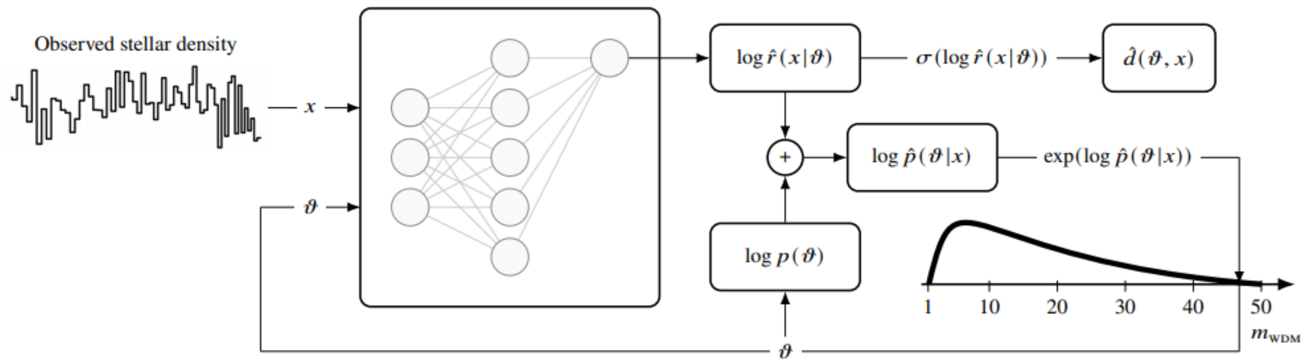

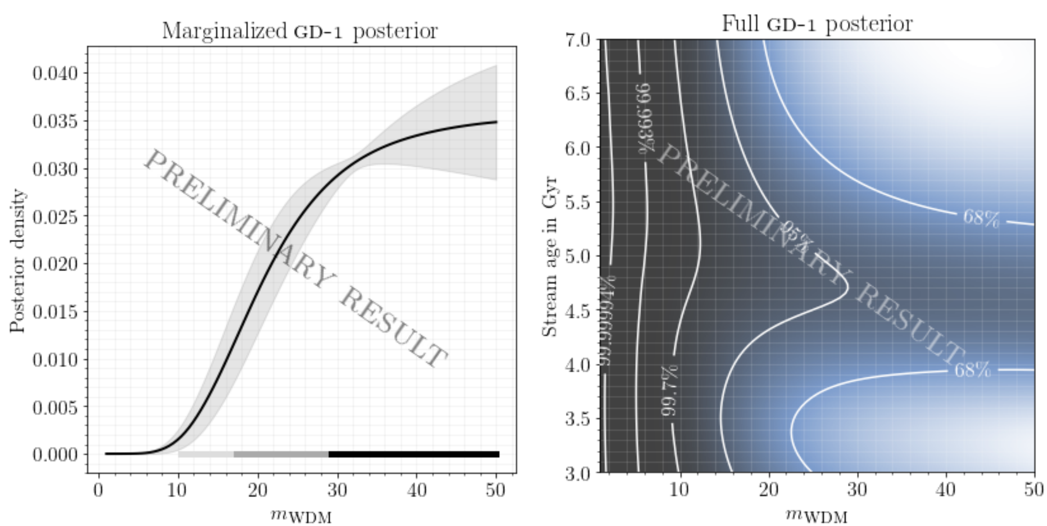

Preliminary results for GD-1 suggest a preference for CDM over WDM.

—

Credits: Hermans et al, 2021.

Wait a minute Gilles... I can't claim that in a paper!
Your neural network must be wrong!

Enforcing conservative posteriors

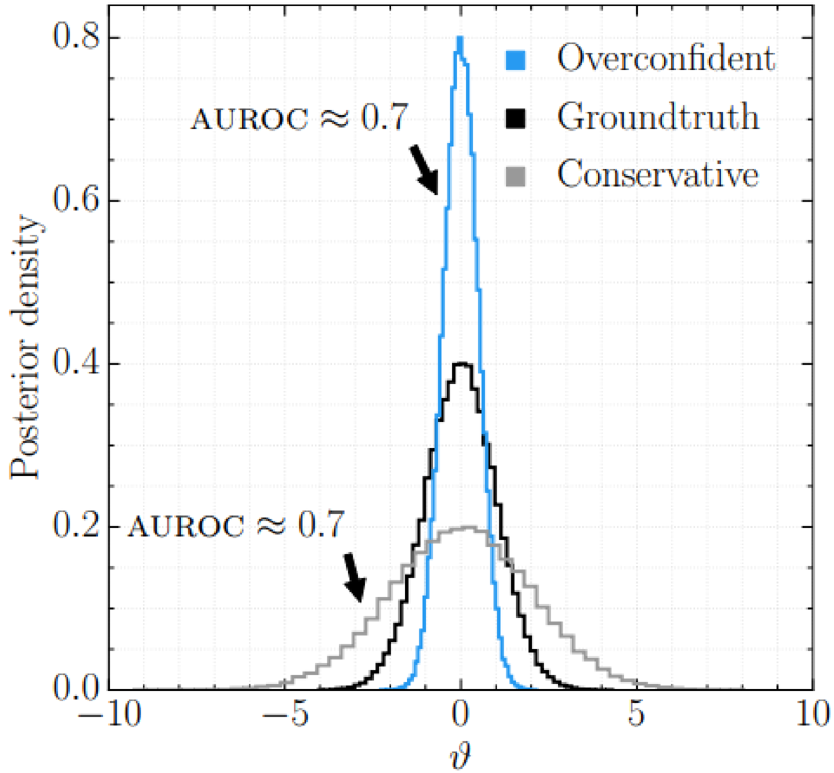

Posterior approximations can be either overconfident (dangerous and wrong) or underconfident (safe but inefficient).

—

Credits: Hermans et al, 2021.

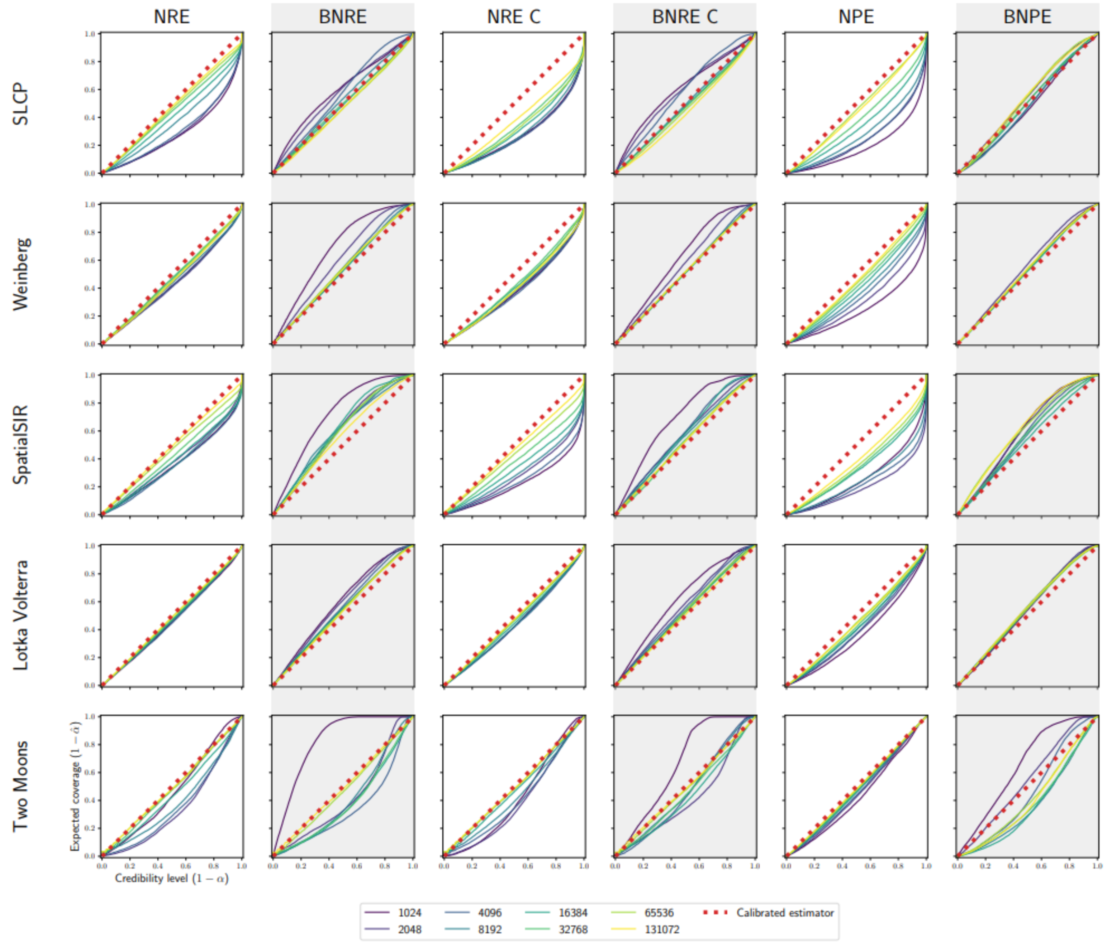

Conservative posteriors can be enforced algorithmically!

Credits: Delaunoy et al, 2022; Delaunoy et al, 2023.

Data assimilation in weather and climate models

Marilaure Grégoire, Xavier Fettweis
(University of Liège)

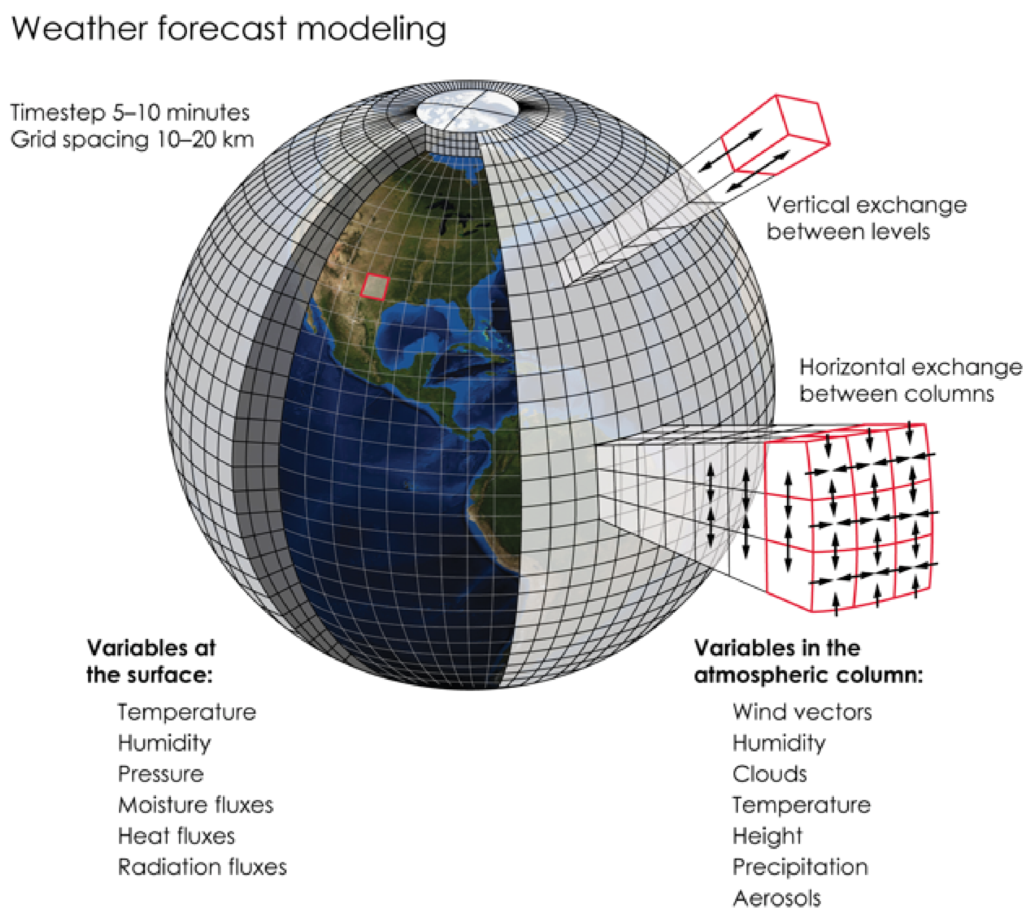

Assuming an initial state x_0, weather forecasts are obtained by propagating the state forward in time using a dynamical model $p(x_{i+1}|x_i)$ based on (costly) numerical simulations.

GraphCast (Google Deepmind) demonstrated that graph neural networks can be used for skilful weather forecasts, at a fraction of the computational cost.

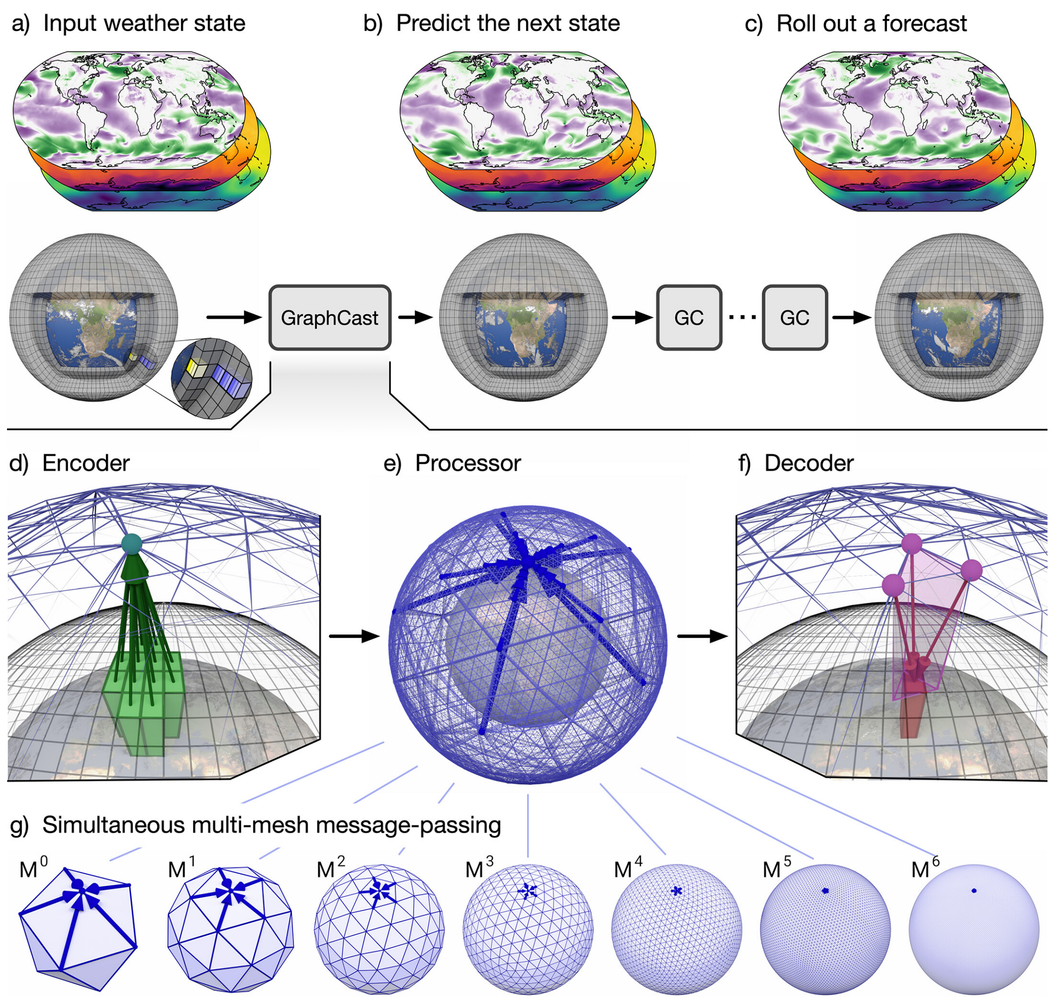

Wait a minute... how do we know x_0 in the first place?

We only have noisy observations y!

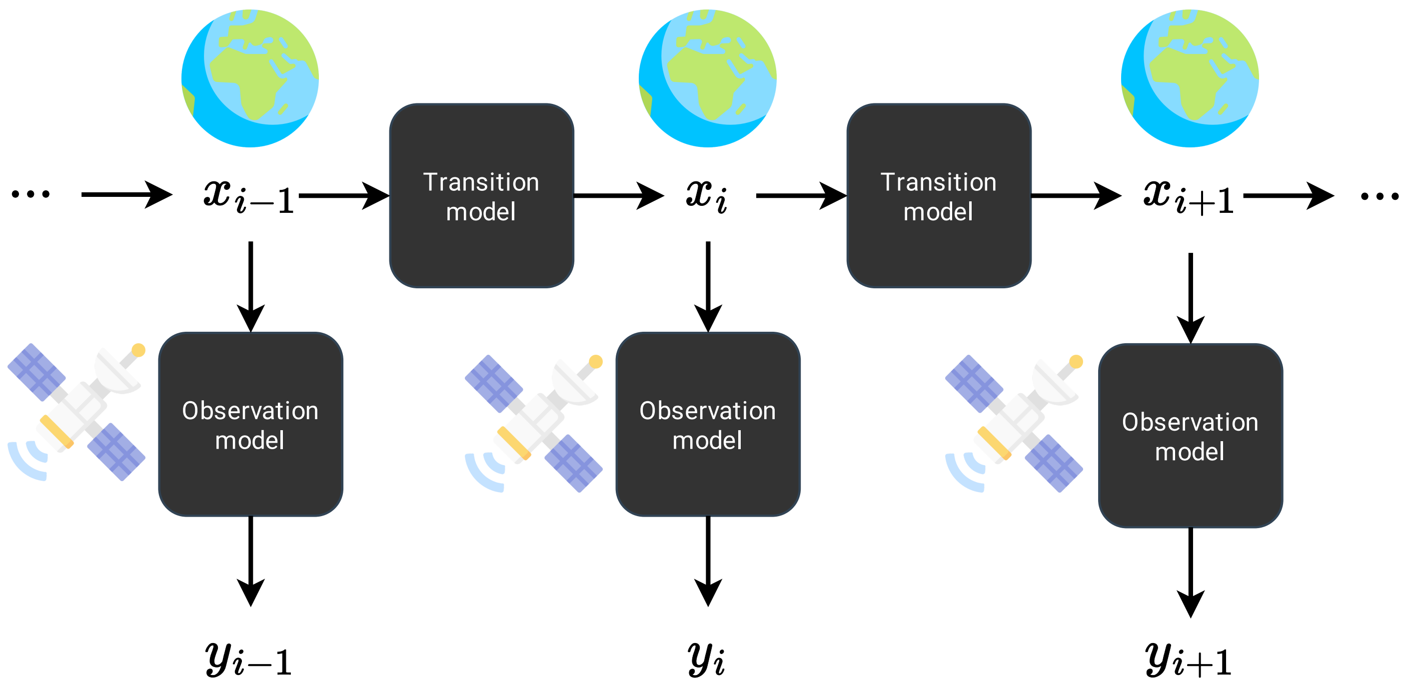

Data assimilation: Estimate plausible trajectories $x_{1:L}$ given one or more noisy observations y (or $y_{1:L}$) as the posterior

$$p(x_{1:L}|y) = \frac{p(y|x_{1:L})}{p(y)} p(x_0) \prod_{i=1}^{L-1} p(x_{i+1}|x_i).$$

Score-based data assimilation

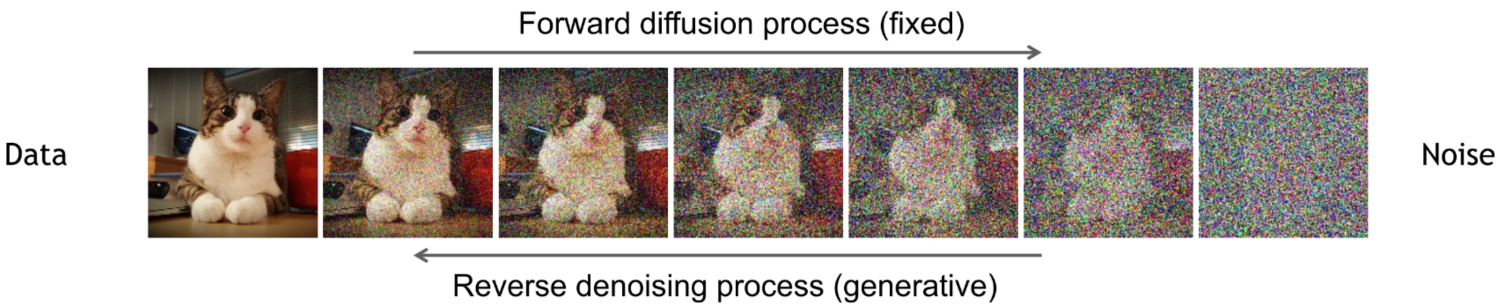

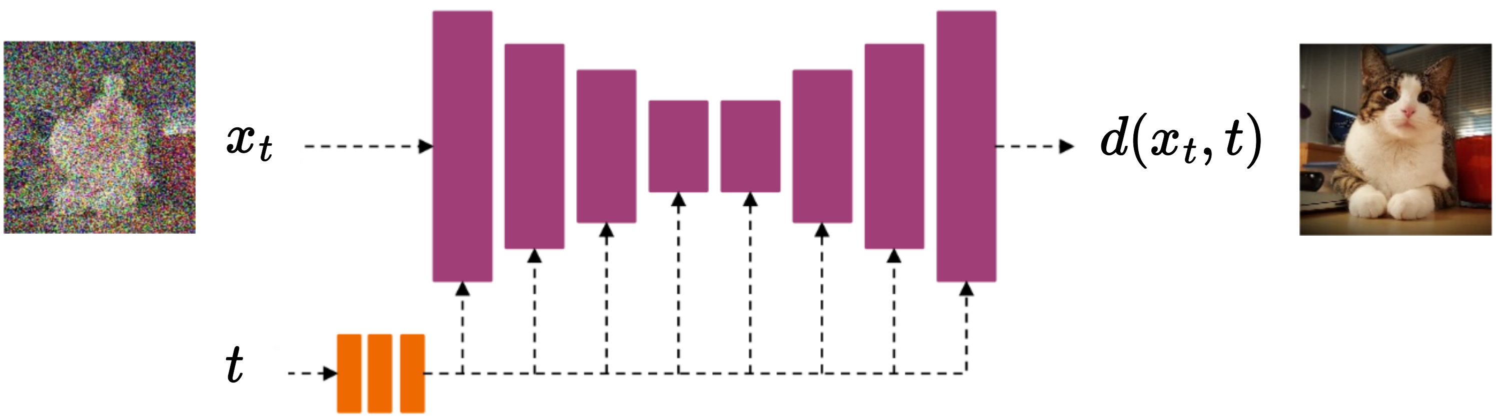

Diffusion models are deep generative models capable of producing data from pure noise.

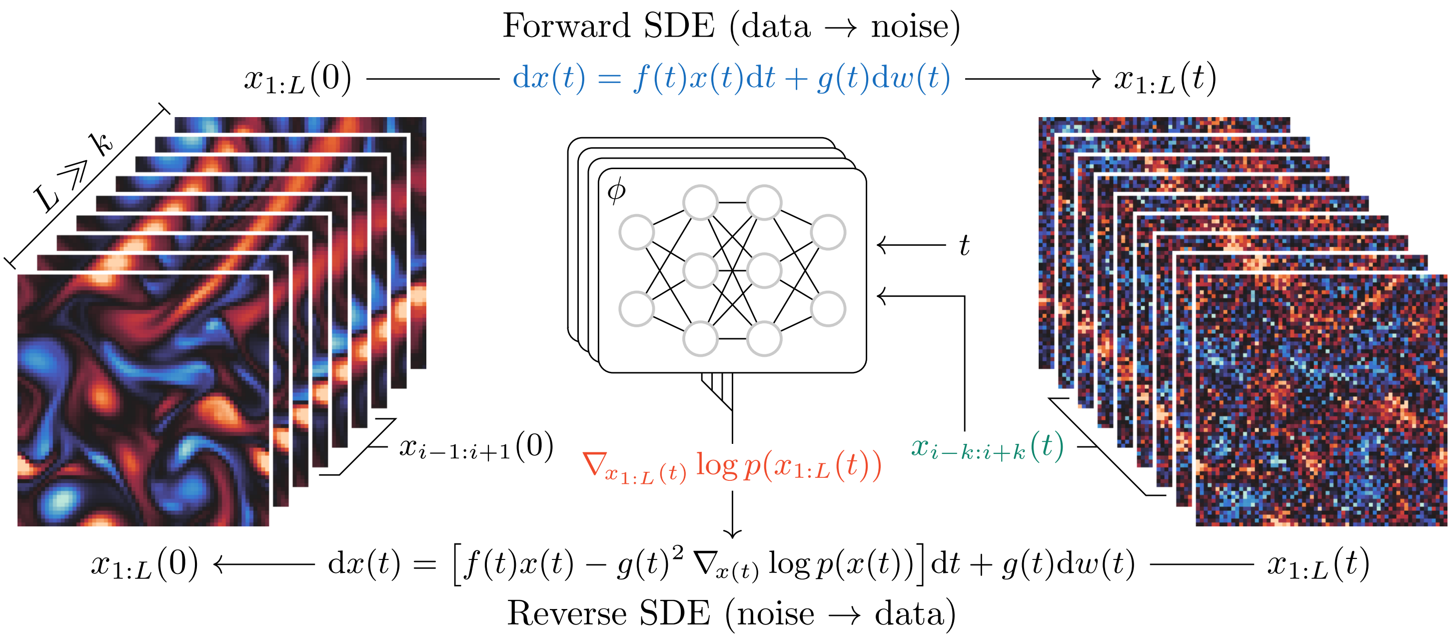

Our approach:

- Build a diffusion model $p(x_{1:L})$ of arbitrary-length trajectories.
- Use zero-shot posterior sampling to generate plausible trajectories from noisy observations y.

—

Credits: Rozet and Louppe, 2023.

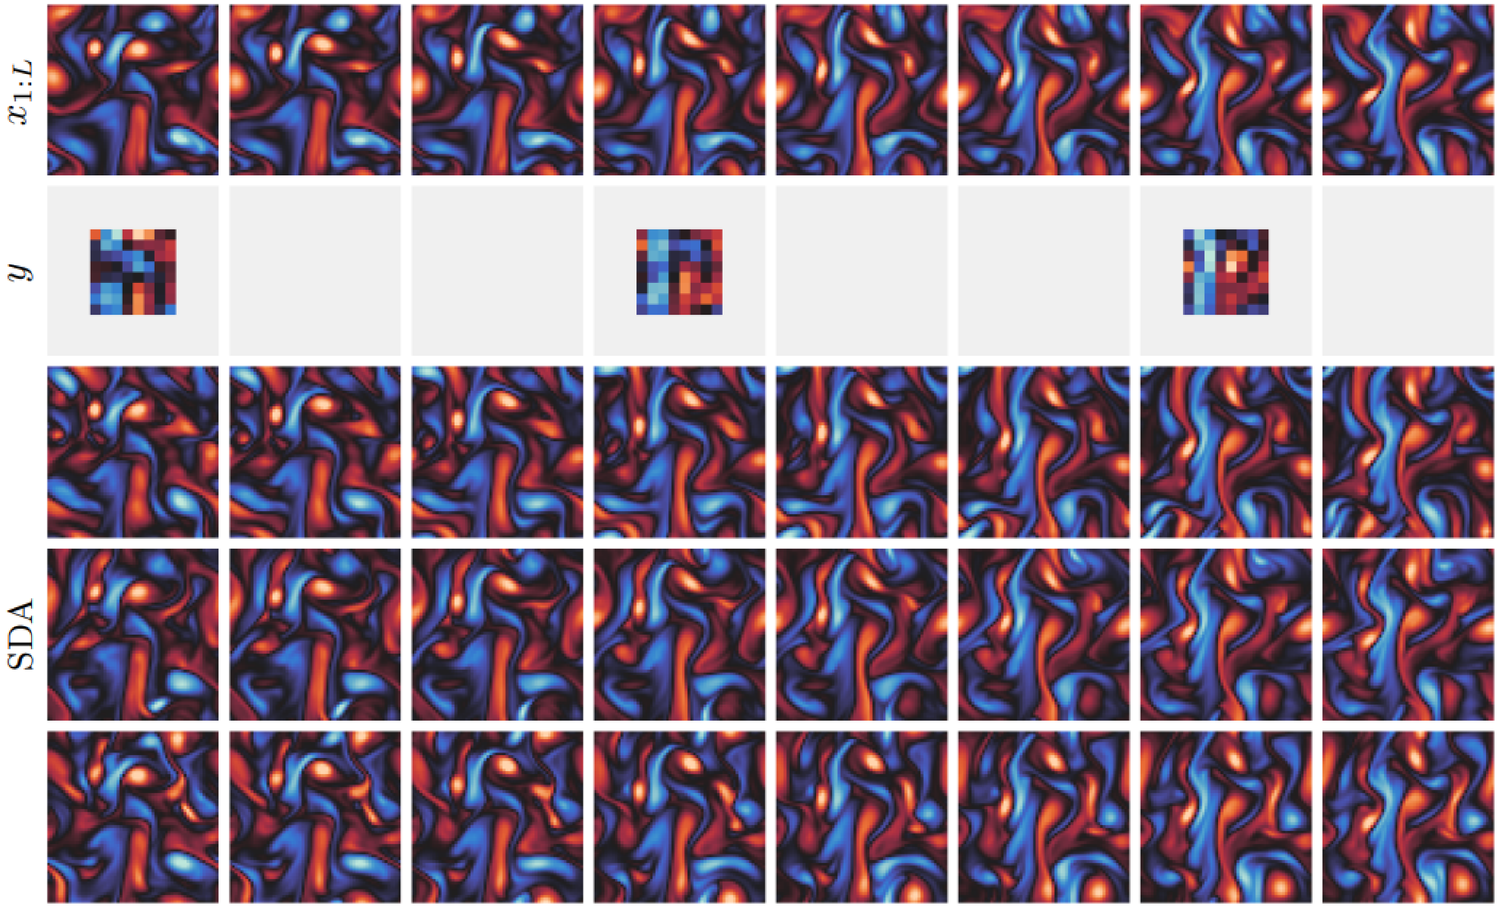

Sampling trajectories $x_{1:L}$ from
noisy, incomplete and coarse-grained observations y.

—
Credits: Rozet and Louppe, 2023.

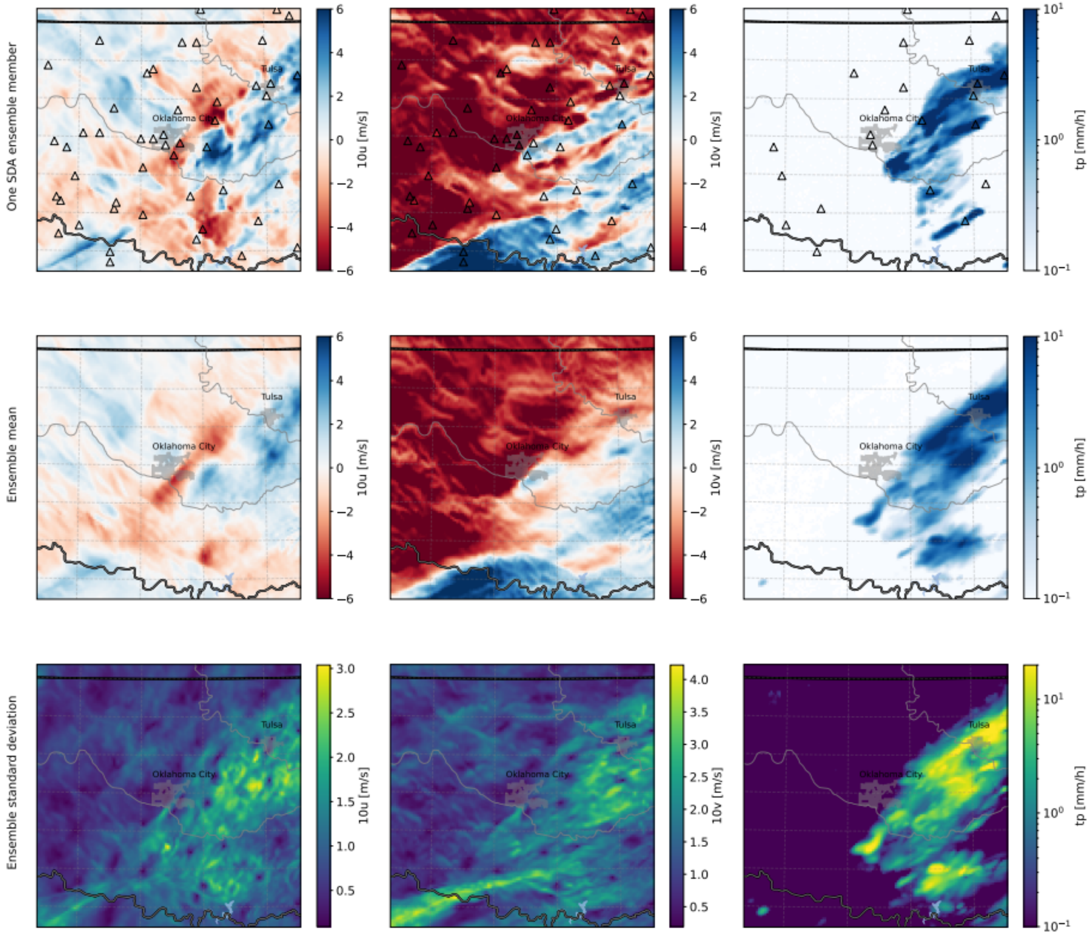

SDA can assimilate noisy weather observations to produce stochastic ensembles.

—

Credits: Manshausen et al, 2024.

Score-based data assimilation for regional models (ongoing)

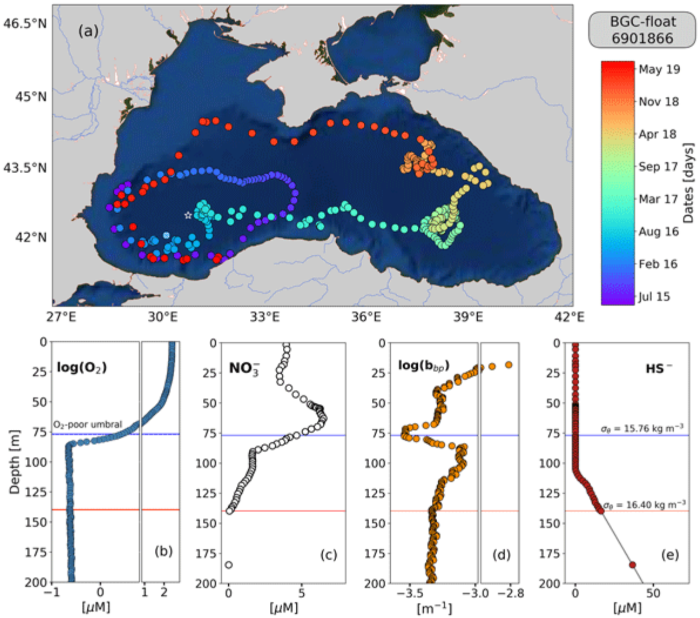

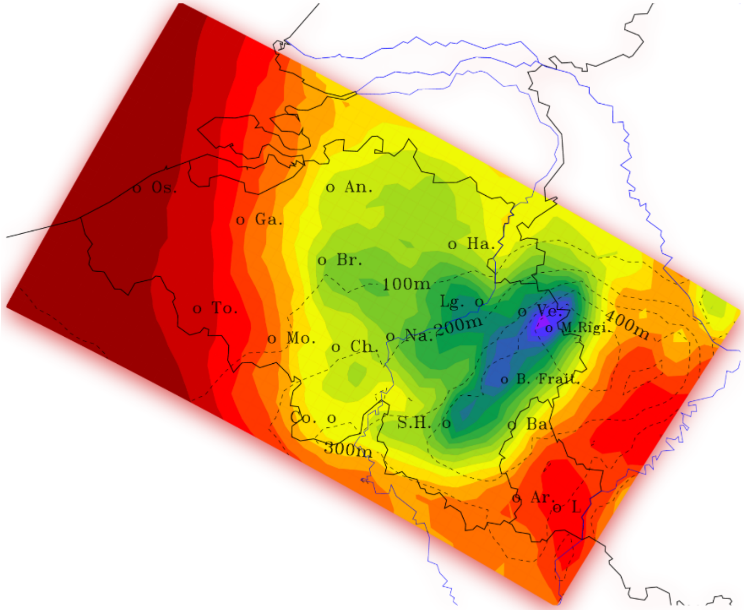

Assimilating satellite and float observations in ocean models (with Marilaure Grégoire).

Assimilating local observations in regional climate models (with Xavier Fettweis).

Earth-scale data assimilation at 0.25° resolution
with latent diffusion models (ongoing).

Conclusions

- AI can help us make sense of complex systems.
- AI unlocks problems we couldn't solve before!
- AI should be used in a principled way if we aim for scientific progress.
- Collaborations between computer scientists and scientists are key.

The end.