

Scientific AI for Cosmology and Beyond



Jason D. McEwen

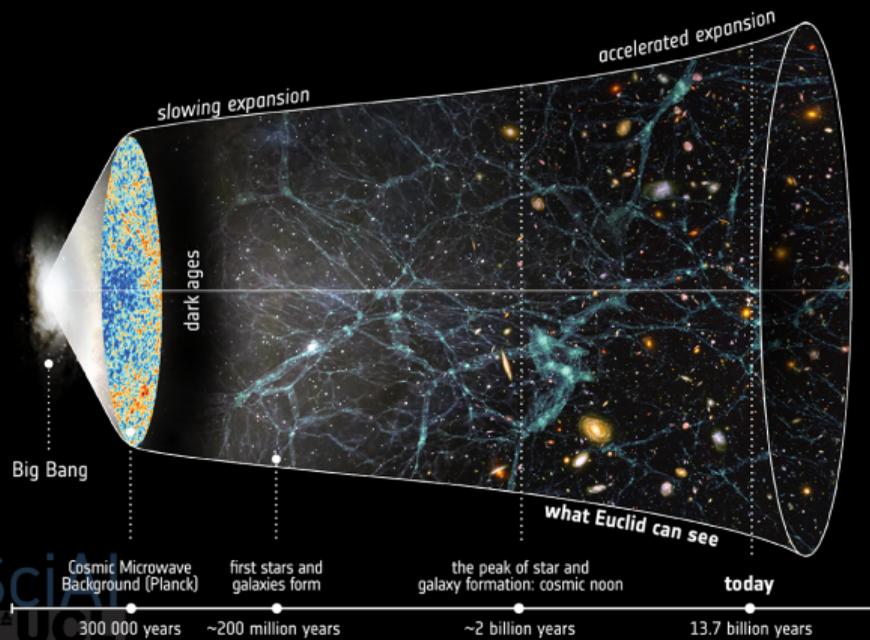
www.jasonmcewen.org

Scientific AI (SciAI) Group, Mullard Space Science Laboratory (MSSL)
University College London (UCL)

Alan Turing Institute

Accelerating Science with AI, Royal Holloway, February 2026

Towards a fundamental understanding of our Universe

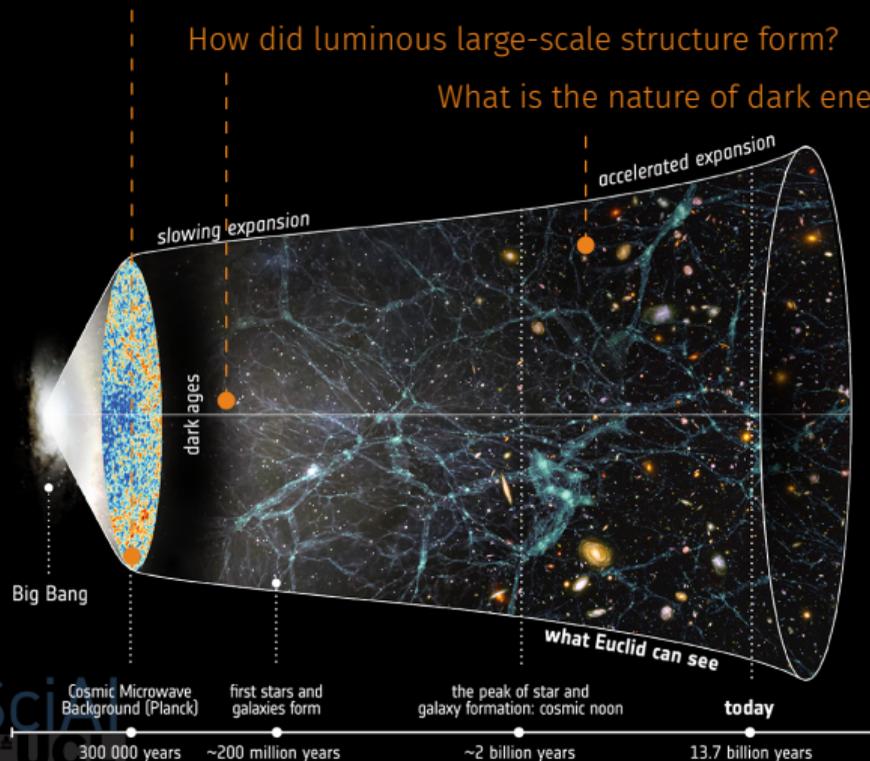


Towards a fundamental understanding of our Universe

What is the origin of structure?

How did luminous large-scale structure form?

What is the nature of dark energy and dark matter?

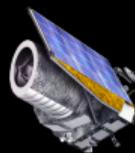
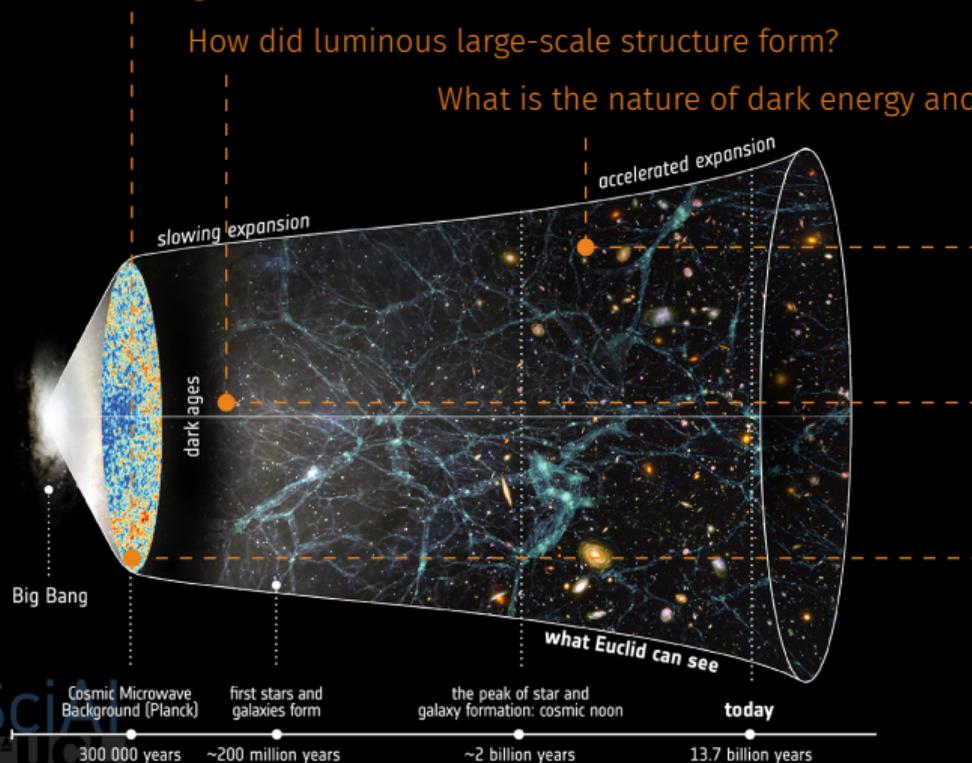


Towards a fundamental understanding of our Universe

What is the origin of structure?

How did luminous large-scale structure form?

What is the nature of dark energy and dark matter?



Euclid



Roman



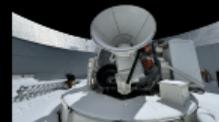
Rubin-LSST



SKA

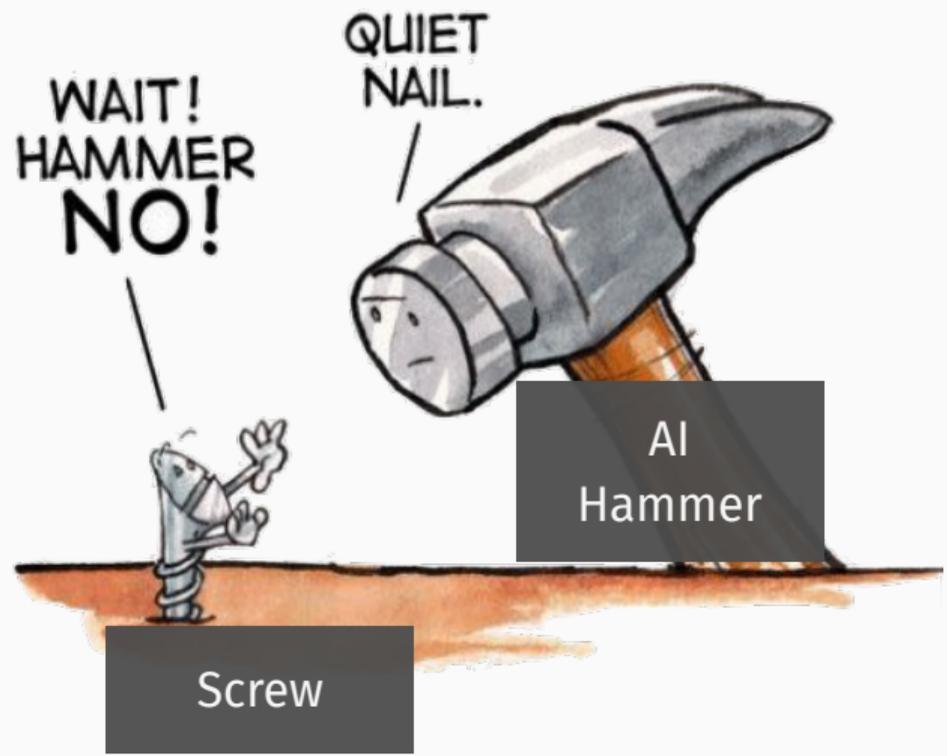


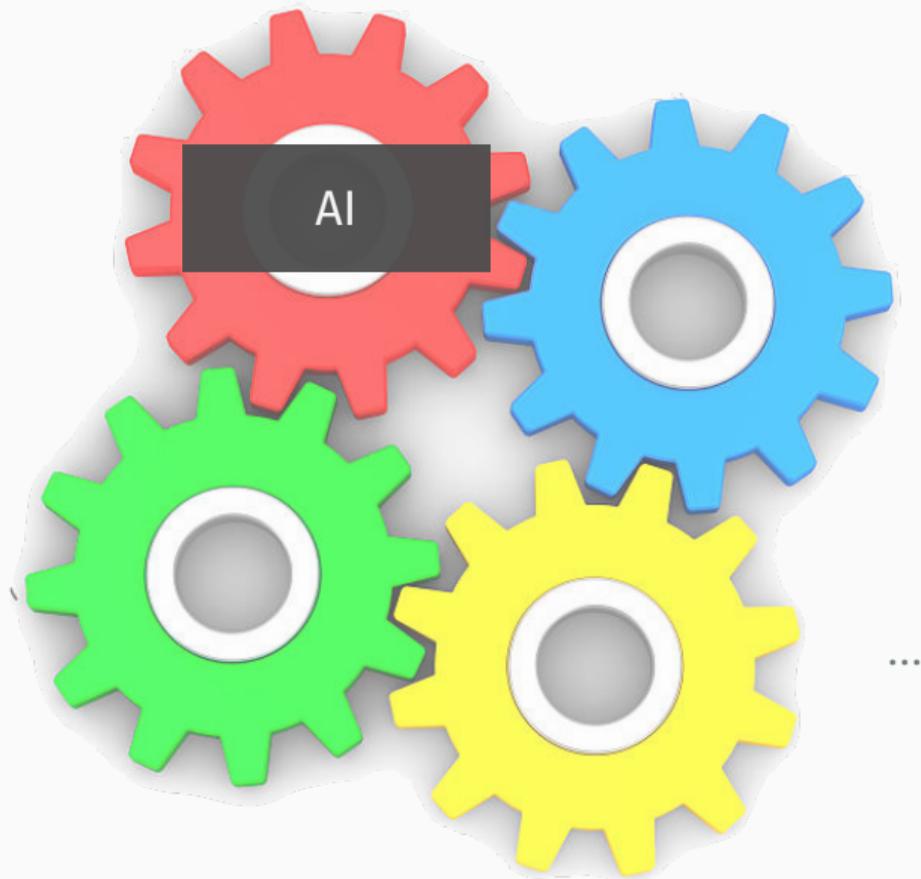
LiteBIRD

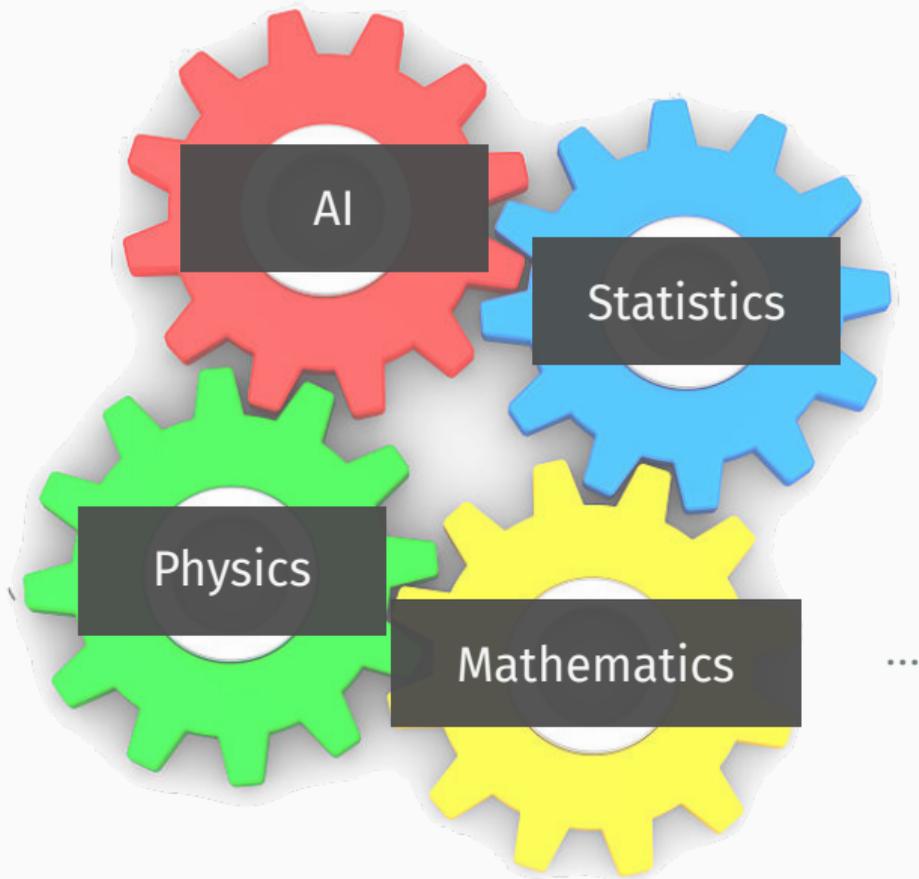


Simons

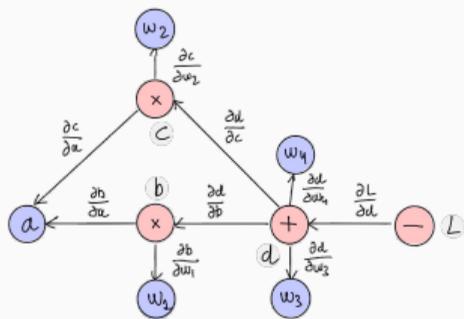
The AI hammer



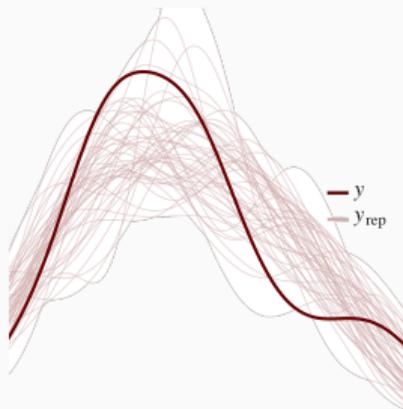




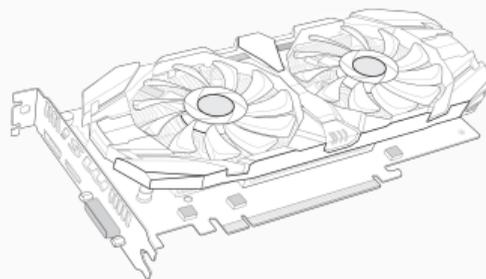
Harnessing modern computing paradigms



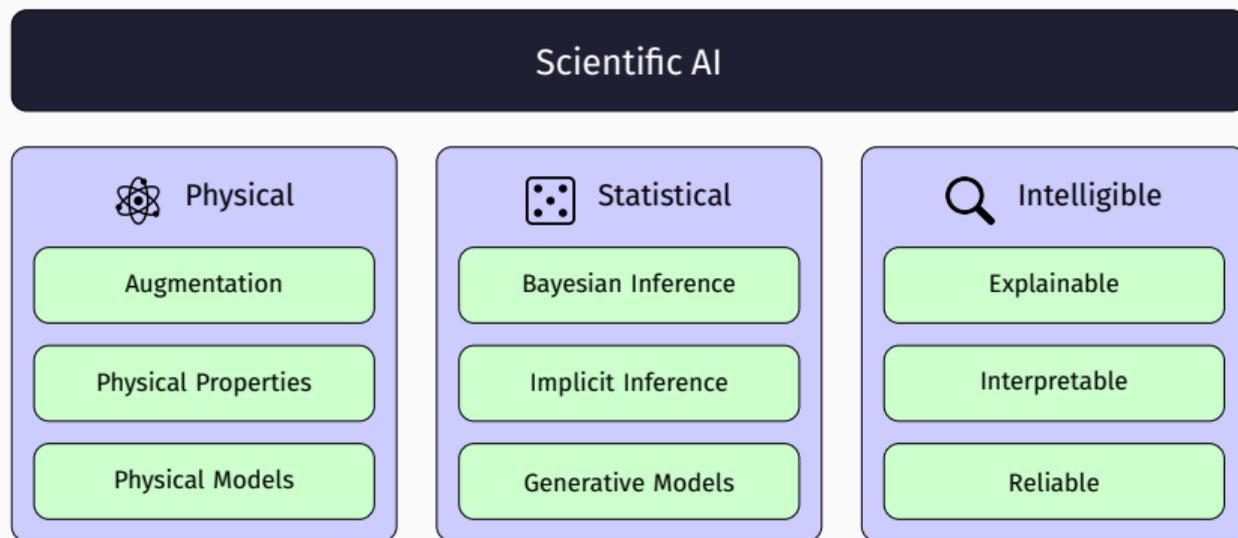
Automatic differentiation



Probabilistic programming



GPU acceleration

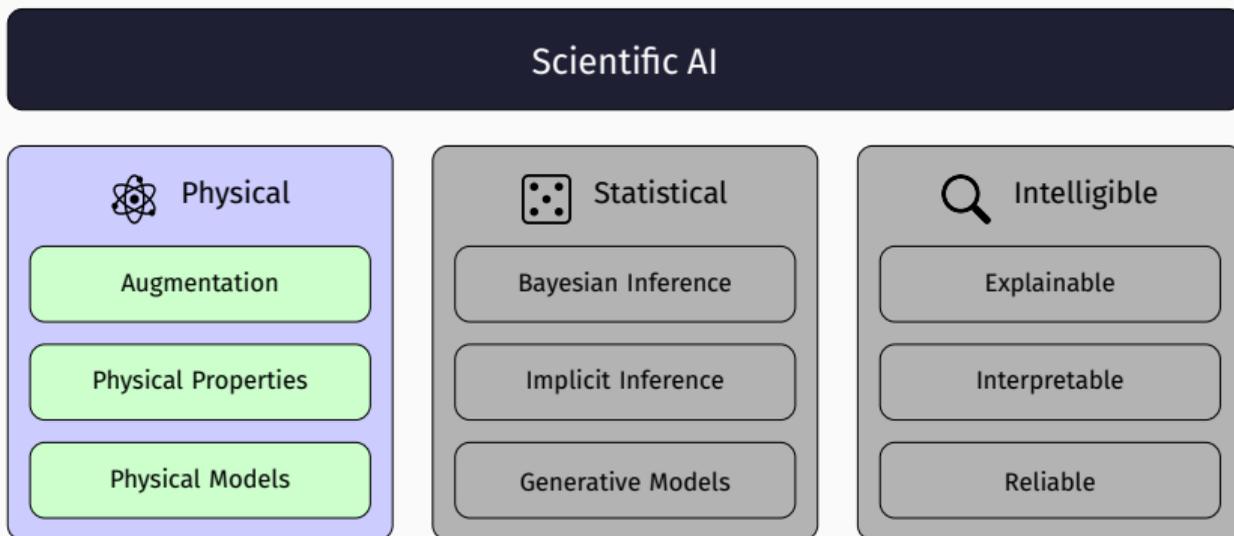


Physics enhanced AI

Physics Enhanced AI

Embed physical understanding of the world into AI models.

(See review by Karniadakis *et al.* 2021.)





Apply **physical transformations** that data known to satisfy to augment training data \rightsquigarrow AI model **learns physics through training**.



Apply **physical transformations** that data known to satisfy to augment training data \rightsquigarrow AI model **learns physics through training**.

- ▶ Common to augment image data-set with rotations, flips, shifts, scales, contrast, ...

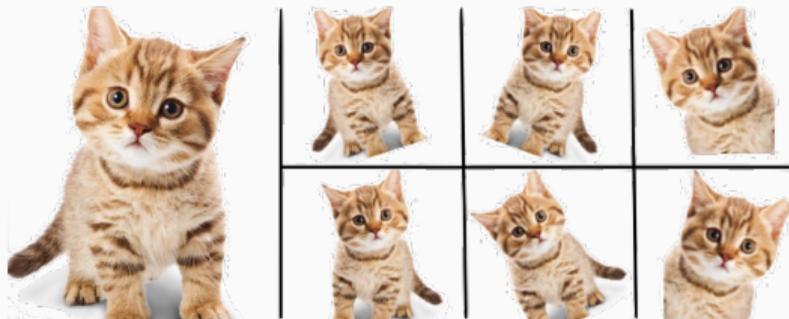
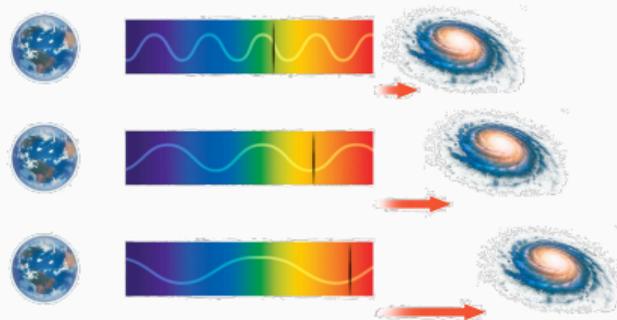


Image augmentation



Apply **physical transformations** that data known to satisfy to augment training data \rightsquigarrow AI model **learns physics through training**.

- ▷ Redshift augmentation of supernovae
(Boone 2019, Alves *et al.* (inc. McEwen) 2022, 2023)



Redshift augmentation

Physical properties: geometries, symmetries, conservation laws



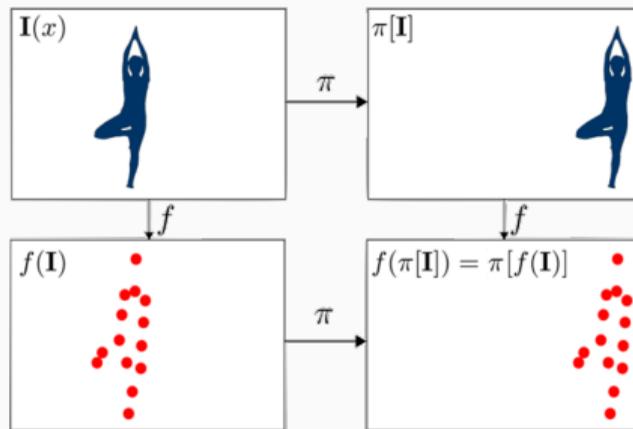
Encode **physical properties** of the world into AI models (e.g. geometry, symmetries, conservation laws) \rightsquigarrow **Physics embedded in architecture** of AI model.

Physical properties: geometries, symmetries, conservation laws



Encode physical properties of the world into AI models (e.g. geometry, symmetries, conservation laws) \rightsquigarrow Physics embedded in architecture of AI model.

- ▷ Key factor CNNs so successful is due to encoding translational equivariance.



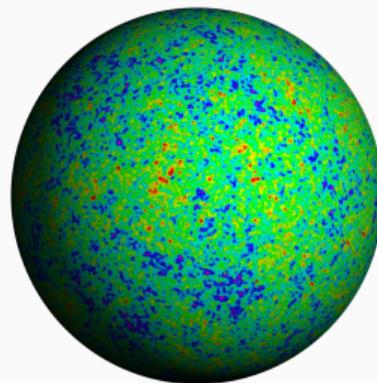
Translational equivariance

Physical properties: geometries, symmetries, conservation laws



Encode **physical properties** of the world into AI models (e.g. geometry, symmetries, conservation laws) \rightsquigarrow **Physics embedded in architecture** of AI model.

- ▷ Geometric deep learning on the sphere
(Cobb et al. 2021, McEwen et al. 2022,
Ocampo, Price & McEwen 2023)



CMB observed on the
celestial sphere

Physical models: PINNs and differentiable physics

Encode physical models of world into AI models:



1. Encode dynamics (differential equations) via loss functions (PINNs).
2. Embed full (differentiable) physical models inside AI model.

↪ Physics learned in training and embedded in model.

Physical models: PINNs and differentiable physics

Encode physical models of world into AI models:

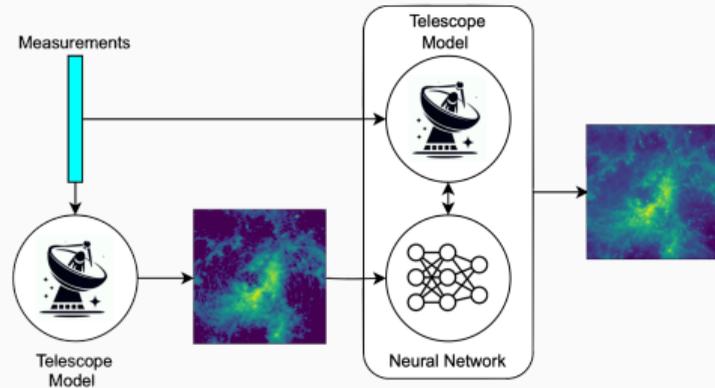


1. Encode dynamics (differential equations) via loss functions (PINNs).
2. Embed full (differentiable) physical models inside AI model.

↪ **Physics learned in training and embedded in model.**

▷ Differentiable physical models

- Instrument models
(Mars *et al.* McEwen 2023, 2024, Liaudat *et al.* McEwen 2024)
- Physical models
(Piras *et al.* McEwen 2024, Spurio Mancini *et al.* McEwen 2024, Whitney *et al.* McEwen in prep.)



Hybrid physics-enhanced AI model
(Mars *et al.* McEwen 2023, 2024)

Physical models: PINNs and differentiable physics

Encode physical models of world into AI models:

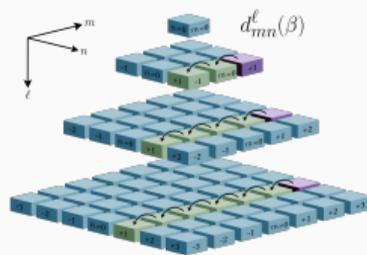


1. Encode dynamics (differential equations) via loss functions (PINNs).
2. Embed full (differentiable) physical models inside AI model.

↪ **Physics learned in training and embedded in model.**

▷ Differentiable mathematical methods

- Fourier transforms
- Spherical harmonic transforms (s2fft; Price & McEwen 2023)
- Spherical wavelet transforms (s2wav; Price *et al.* McEwen 2024)
- Spherical scattering transforms (s2scat; Mousset *et al.* McEwen 2024)



Initialise Recursion

$$d_{mn}^l(\beta) = \sqrt{\frac{(2l)!}{(l+n)!(l-n)!}} \left(-\sin \frac{\beta}{2}\right)^{l-n} \left(\cos \frac{\beta}{2}\right)^{l+n}$$

Execute Recursion

$$d_{\ell m-1, n}^{\ell}(\beta) = \lambda_m a_{m-1} d_{mn}^{\ell}(\beta) - \frac{a_{m-1}}{a_m} d_{m+1, n}^{\ell}(\beta)$$

where $\lambda_m = \frac{n-m \cos \beta}{\sin \beta}$ and $a_m = \frac{2}{\sqrt{(l-m)(l+m+1)}}$

Avoid Singularities

$$d_{m0}^{\ell}(0) = \delta_{m0} \text{ and } d_{m0}^{\ell}(\pi) = (-1)^{\ell+m} \delta_{m, -n}$$

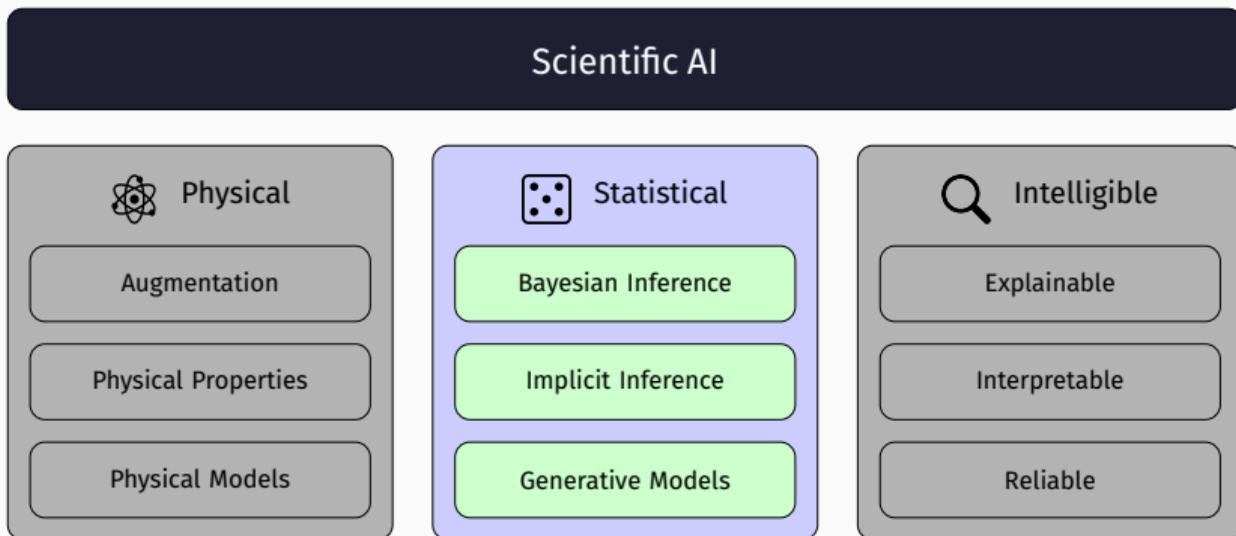
Differentiable and GPU-friendly recursions
(Price & McEwen 2023)

Statistical AI

Statistical AI

Embed a statistical representation of data, models and/or outputs.

(See Murray 2022.)



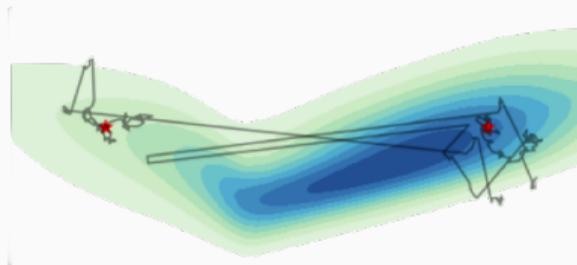


AI techniques can be integrated into Bayesian frameworks to **enhance accuracy and computational efficiency**, making some approaches accessible that were previously intractable.



AI techniques can be integrated into Bayesian frameworks to **enhance accuracy and computational efficiency**, making some approaches accessible that were previously intractable.

- ▷ Enhanced MCMC for parameter estimation (Grabrie *et al.* 2022, Karamanis *et al.* 2022).

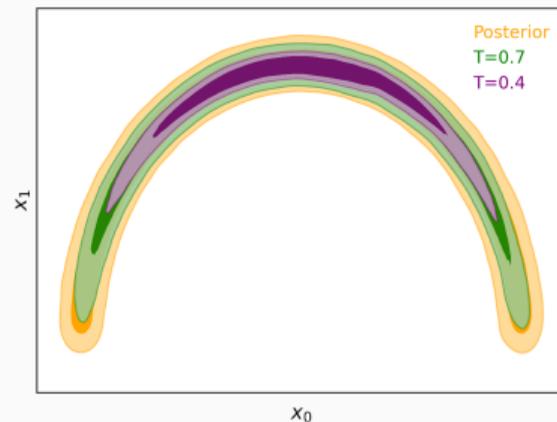


Learned proposal distributions



AI techniques can be integrated into Bayesian frameworks to **enhance accuracy and computational efficiency**, making some approaches accessible that were previously intractable.

- ▷ Enhanced Bayesian model selection (**harmonic**; McEwen *et al.* 2021, Polanska *et al.* McEwen 2023, 2024, Piras *et al.* McEwen 2024, Spurio Mancini *et al.* McEwen 2023, 2024).



Learned harmonic mean estimator
(harmonic)



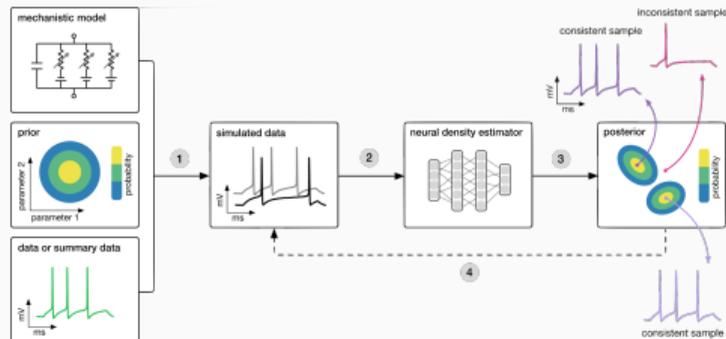
AI techniques can be used to **learn surrogates of implicit distributions** when they are not tractable or are computationally infeasible.

Implicit inference



AI techniques can be used to **learn surrogates of implicit distributions** when they are not tractable or are computationally infeasible.

- ▶ Simulation-based inference (Cranmer *et al.* 2021).



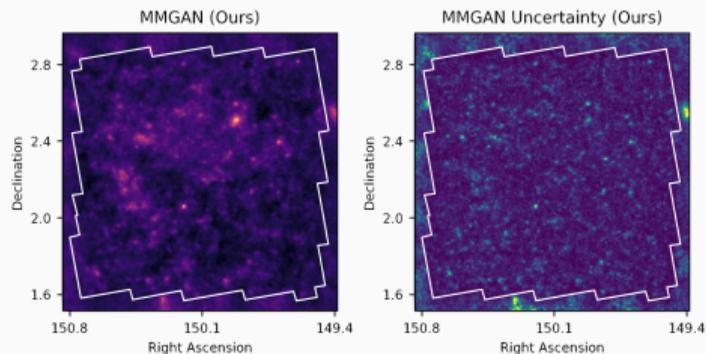
sbi

Implicit inference



AI techniques can be used to **learn surrogates of implicit distributions** when they are not tractable or are computationally infeasible.

- ▷ Variational inference
(Whitney *et al.* McEwen 2024).



Mass mapping with uncertainties
by variational inference
(Whitney *et al.* McEwen 2024)



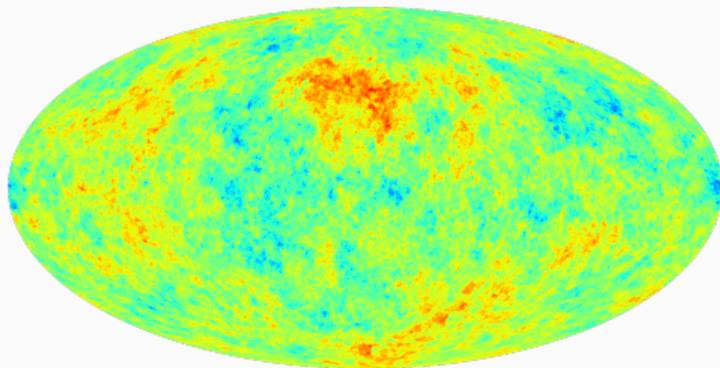
Generative models **learn a prior distribution** from data for sampling and/or evaluating probability densities.

Generative models



Generative models **learn a prior distribution** from data for sampling and/or evaluating probability densities.

- ▷ Emulation: sample from learned prior
(Price *et al.* McEwen 2023, Price *et al.* McEwen in prep., Mousset *et al.* McEwen 2024)



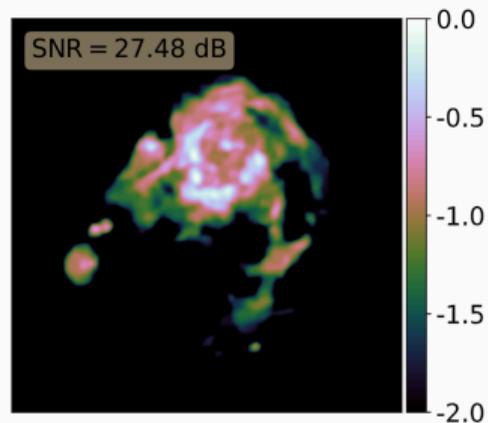
Emulated cosmic string maps
(**stringgen**, Price *et al.* McEwen 2023,
Price *et al.* McEwen in prep.)

Generative models



Generative models **learn a prior distribution** from data for sampling and/or evaluating probability densities.

- ▷ Integrate learned priors into analysis
(McEwen *et al.* 2023, Liaudat *et al.* McEwen 2024)



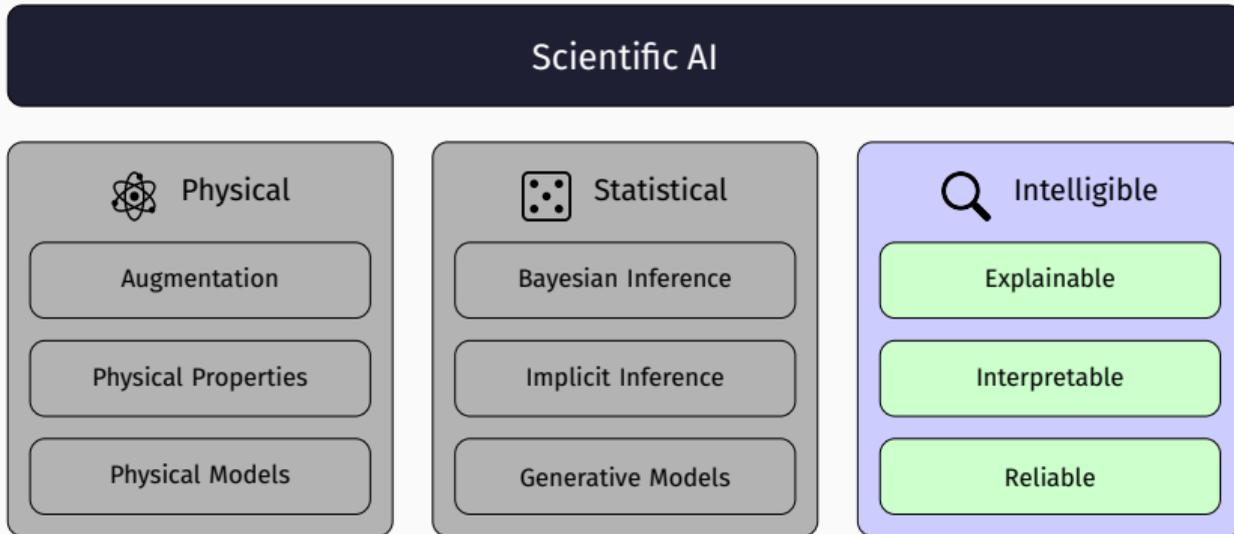
Learn radio galaxy prior
(Liaudat *et al.* McEwen 2024)

Intelligible AI

Intelligible AI

AI methods that are able to be understood by humans and are reliable.

(See Weld & Bansal 2018, Ras *et al.* 2020.)



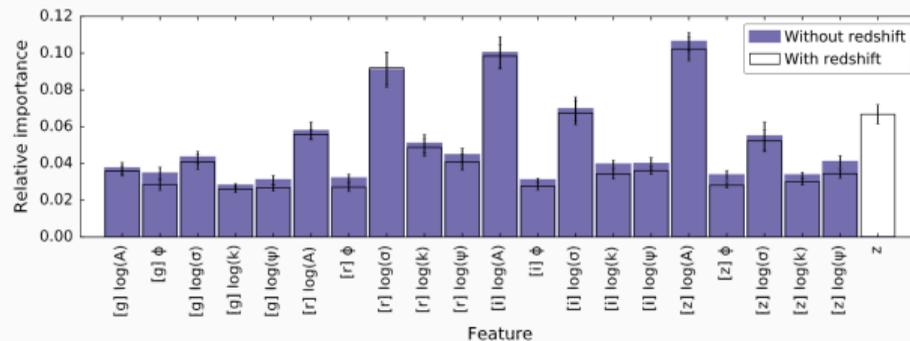


Explainable AI techniques may or may not be interpretable themselves but their outputs can be explained to humans.



Explainable AI techniques may or may not be interpretable themselves but their outputs can be explained to humans.

- ▷ Feature importances
(Lochner, McEwen *et al.* 2016)

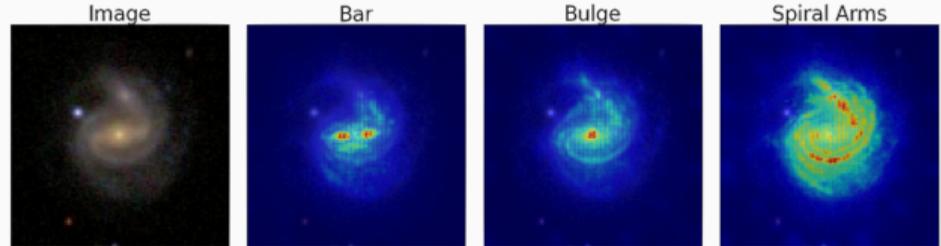


Supernova feature importances



Explainable AI techniques may or may not be interpretable themselves but their outputs can be explained to humans.

- ▷ Saliency maps
(Bhambra, Joachimi & Lahav 2022)



Galaxy saliency mapping

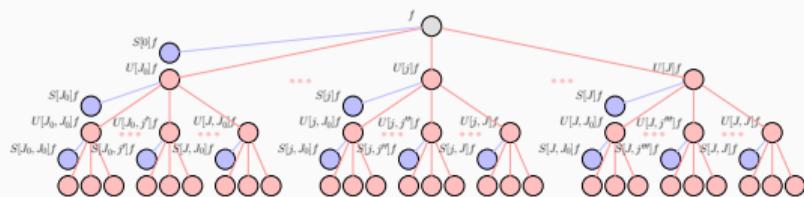


Interpretable AI models are **white boxes** that can be understood by humans.



Interpretable AI models are **white boxes** that can be understood by humans.

- ▷ Designed models such as wavelet scattering networks
(McEwen *et al.* 2022, Mousset *et al.* McEwen 2024)

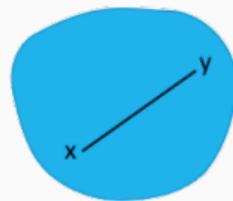


Scattering network (McEwen *et al.* 2022)



Interpretable AI models are **white boxes** that can be understood by humans.

- ▷ Interpretable constraints on AI models
(Liaudat *et al.* McEwen 2024)



Convexity



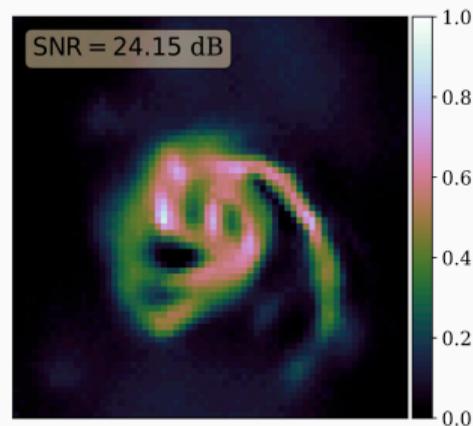
Uncertainty
Quantification

Impose convexity on learned model



Interpretable AI models are **white boxes** that can be understood by humans.

- ▷ Deep priors learned from training data (McEwen *et al.* 2023, Liaudat *et al.* McEwen 2024)



Compute Bayesian evidence for model selection (McEwen *et al.* 2023)

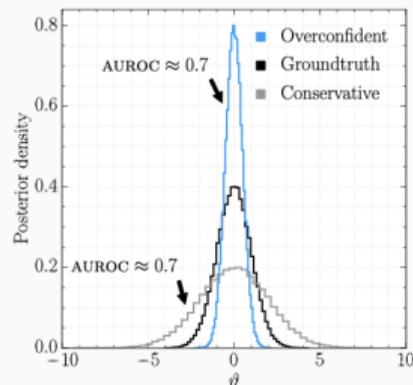


Reliability and validity **critical for science** to have confidence in results of AI models. Closely coupled with a **meaningful statistical distribution** of outputs.



Reliability and validity **critical for science** to have confidence in results of AI models. Closely coupled with a **meaningful statistical distribution** of outputs.

- ▷ Validity of statistical distributions
(Lueckmann *et al.* 2021, Hermans *et al.* 2022, Cannon *et al.* 2023)

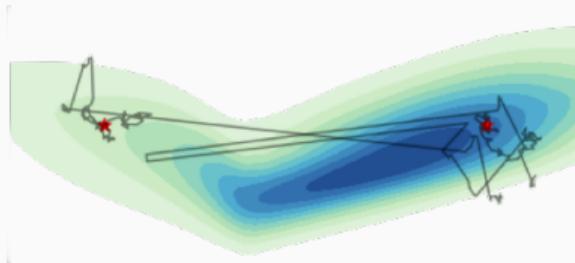


Validity of distribution
(Hermans *et al.* 2022)



Reliability and validity **critical for science** to have confidence in results of AI models. Closely coupled with a **meaningful statistical distribution** of outputs.

- ▷ Integrate into statistical frameworks to inherit theoretical guarantees
(McEwen *et al.* 2023, Liaudat *et al.* McEwen 2024, McEwen *et al.* 2021, Polanska *et al.* McEwen 2023, 2024, Piras *et al.* McEwen 2024)

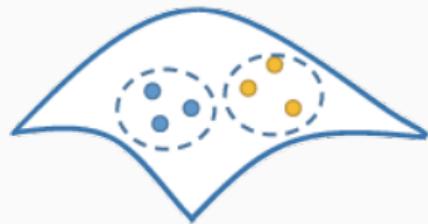


Inherit guarantees from overarching statistical frameworks



Reliability and validity **critical for science** to have confidence in results of AI models. Closely coupled with a **meaningful statistical distribution** of outputs.

- ▷ Design to ensure conservative and avoid mode collapse (Delaunoy *et al.* 2022, Price *et al.* McEwen 2023, Whitney *et al.* McEwen 2024)



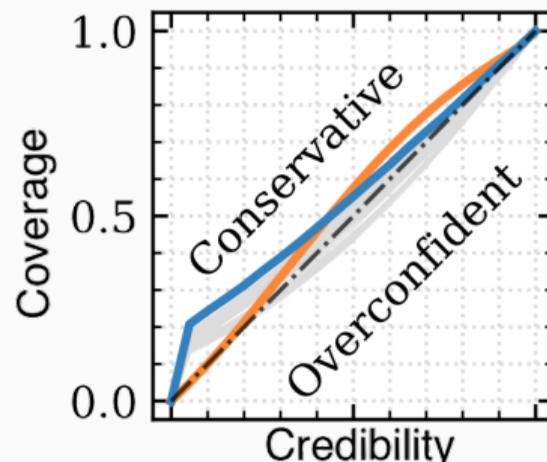
Recover probability distribution over full underlying data manifold (Price *et al.* McEwen 2023)



Reliability and validity **critical for science** to have confidence in results of AI models. Closely coupled with a **meaningful statistical distribution** of outputs.

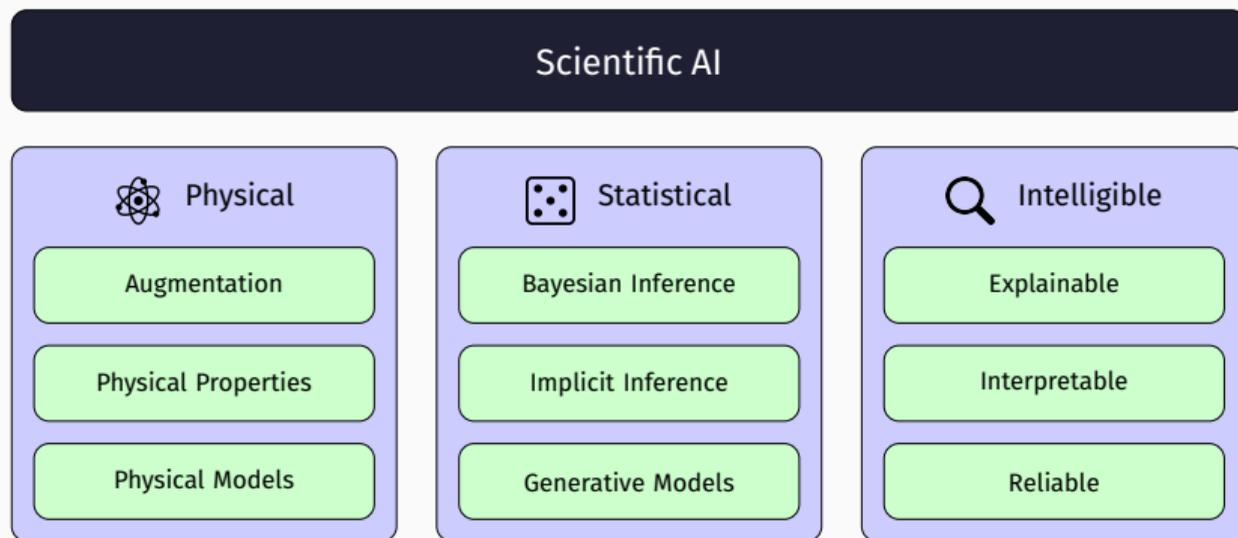
▷ Extensive validation checks:

- Coverage testing (Lemos *et al.* 2023)
- Simulation-based calibration checks (Talts *et al.* 2020)
- Classifier two-sample tests (C2ST) (Lopez-Paz & Oquab 2017)
- ...



Coverage analysis
(Cannon *et al.* 2023)

Summary

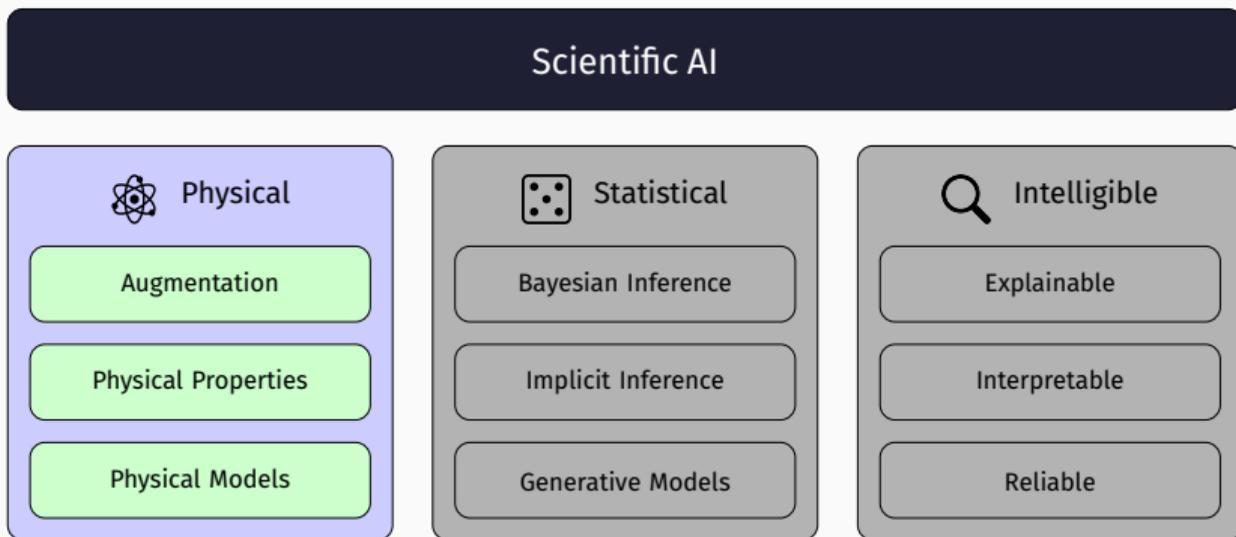


Extra Slides

Physics Enhanced AI

Embed physical understanding of the world into AI models.

(See review by Karniadakis *et al.* 2021.)



Augmentation



Apply **physical transformations** that data known to satisfy to augment training data \rightsquigarrow AI model **learns physics through training**.



▷ Data efficiency suffers: data “used” to learn physics, rather than problem.



▷ Simple and easy to implement.

Physical properties: geometries, symmetries, conservation laws



Encode **physical properties** of the world into AI models (e.g. geometry, symmetries, conservation laws) \rightsquigarrow **Physics embedded in architecture** of AI model.



- ▷ Inductive biases required? Should we just learn from data?
- ▷ Highly computationally demanding.



- ▷ Improved data-efficiency.
- ▷ Inductive biases not necessarily strictly enforced.
- ▷ Develop efficient algorithms (e.g. Ocampo, Price & McEwen 2023).

Physical models: PINNs and differentiable physics

Encode physical models of world into AI models:



1. Encode dynamics (differential equations) via loss functions (PINNs).
2. Embed full (differentiable) physical models inside AI model.

↪ **Physics learned in training and embedded in model.**



- ▷ PINNs only capture limited dynamics via loss.
- ▷ Full physical models requires differentiable programming frameworks.

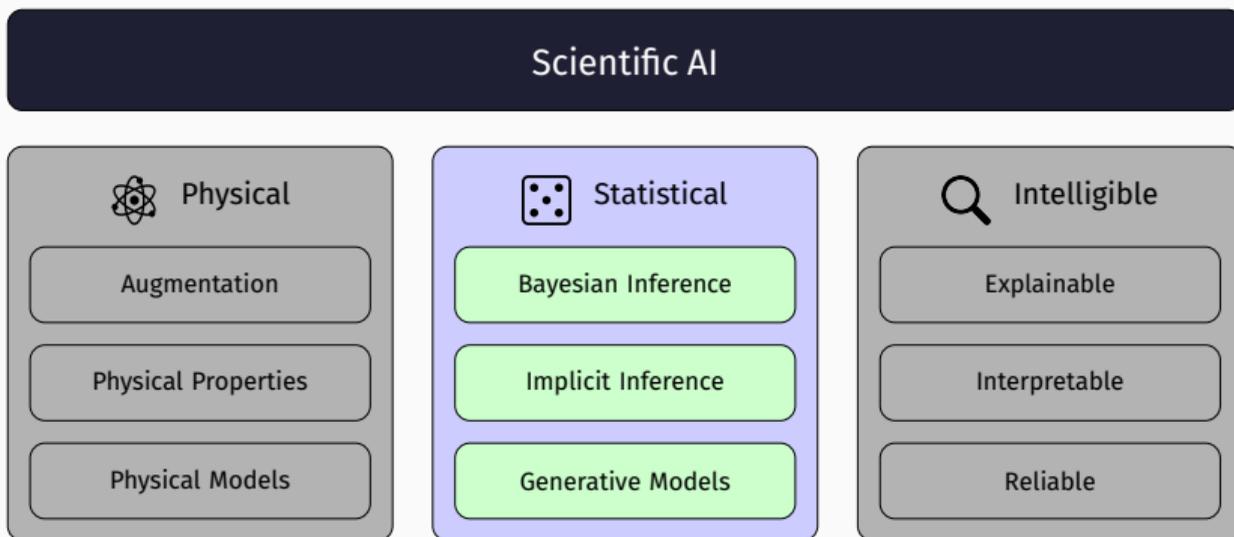
- ▷ Capture full physics with differentiable models!
- ▷ Emulators also provide differentiability.
- ▷ Write new differentiable codes (e.g. `s2fft`; Price & McEwen 2023).



Statistical AI

Embed a statistical representation of data, models and/or outputs.

(See Murray 2022.)



Bayesian inference



AI techniques can be integrated into Bayesian frameworks to **enhance accuracy and computational efficiency**, making some approaches accessible that were previously intractable.



▷ Learned components (e.g. proposals) must satisfy appropriate properties.



▷ Enforce required properties on learned components.

▷ Inherit statistical guarantees of overarching Bayesian framework.

Implicit inference



AI techniques can be used to **learn surrogates of implicit distributions** when they are not tractable or are computationally infeasible.



- ▷ Availability and representativeness of training data.
- ▷ Cost of training.
- ▷ Reliability?



- ▷ Public datasets/benchmarks (e.g. IllustrisTNG, CAMELS, Quijote, CosmoGrid, Gower St).
- ▷ Amortized inference (training **not** repeated for new observations).
- ▷ Integrate in Bayesian framework to provide some statistical guarantees.
- ▷ Statistical validation (hold that thought... see Reliability section).

Generative models



Generative models **learn a prior distribution** from data for sampling and/or evaluating probability densities.



- ▷ Availability and representativeness of training data.
- ▷ Reliability, *e.g.* diversity of AI models often lacking.

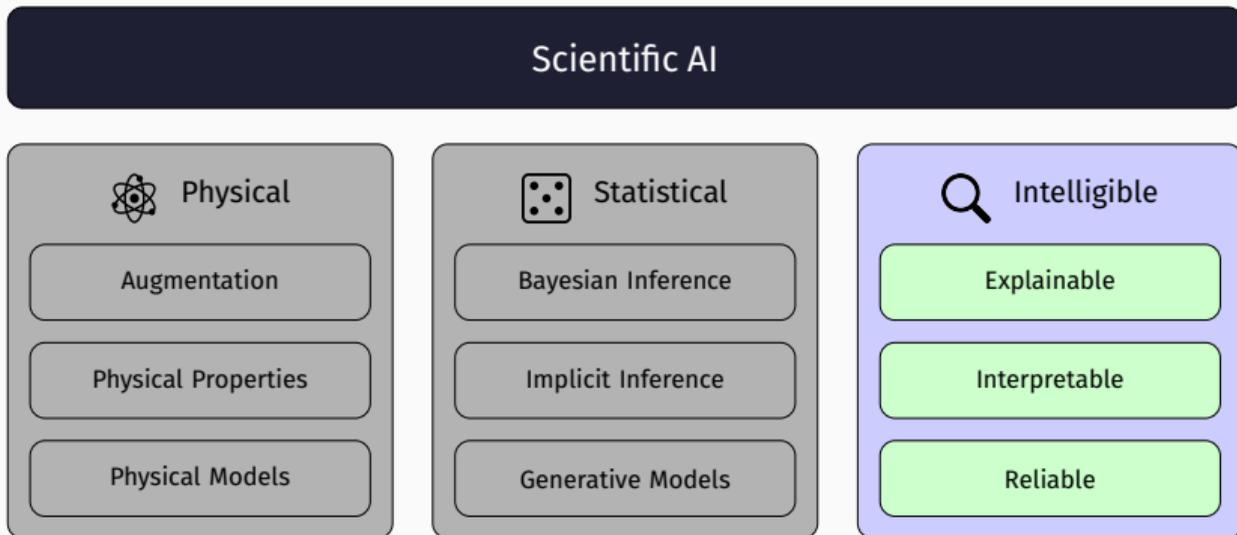


- ▷ Public datasets/benchmarks (*e.g.* IllustrisTNG, CAMELS, Quijote, CosmoGrid, Gower St).
- ▷ Meta sampling to recover distribution over manifold (*e.g.* Price *et al.* 2023).
- ▷ Reliability (hold that thought... see Reliability section).

Intelligible AI

AI methods that are able to be understood by humans and are reliable.

(See Weld & Bansal 2018, Ras *et al.* 2020.)





Explainable AI techniques may or may not be interpretable themselves but their outputs can be explained to humans.



- ▷ Poking the black box.
- ▷ Humans still not able to comprehend underlying process.

Interpretability



Interpretable AI models are **white boxes** that can be understood by humans.



▷ Designed models limit flexibility.



▷ Benefits of designed models often outweigh (minimal) reduced flexibility.



Reliability and validity **critical for science** to have confidence in results of AI models. Closely coupled with a **meaningful statistical distribution** of outputs.



- ▷ Uncertainties not always meaningful.
- ▷ Diversity of AI model often lacking.



- ▷ Integrate in statistical framework to inherit theoretical guarantees.
- ▷ Design to be conservative and avoid mode collapse
(*e.g.* Price *et al.* McEwen 2023, Whitney *et al.* McEwen 2024).
- ▷ Extensive validation tests.
- ▷ Well-posed frameworks (*e.g.* physics enhanced, statistical).