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)

International School for Advanced Studies (SISSA), Trieste, December 2024

Towards a fundamental understanding of our Universe

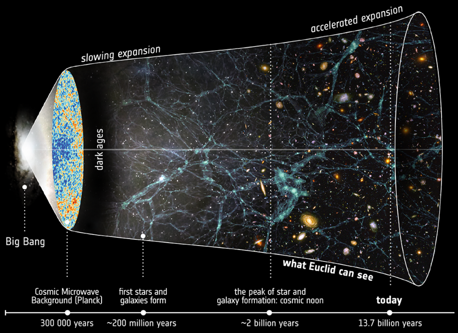

- ▷ Unanswered fundamental questions
- ▷ Imminent new data
- ▷ How can we bring AI to bear?

Outline

1. Towards scientific AI

2. Statistical characterisation and generative modelling of cosmological fields

Towards scientific AI

The AI hammer

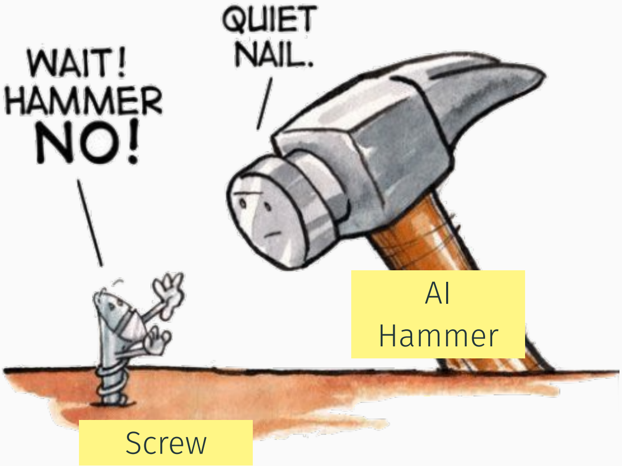

The AI cog

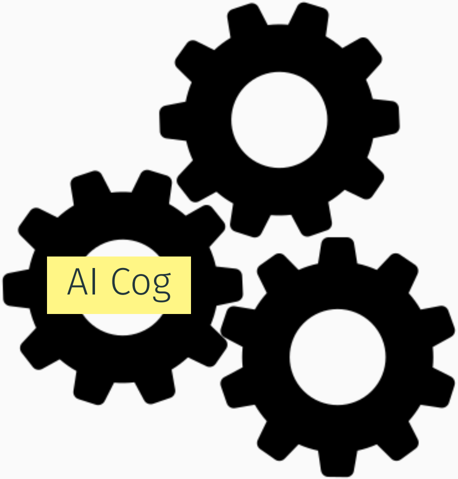

Merging paradigms

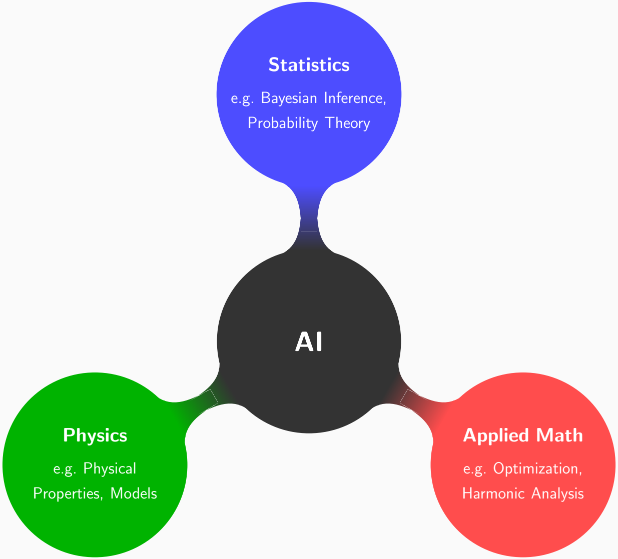

Scientific AI for the physical sciences

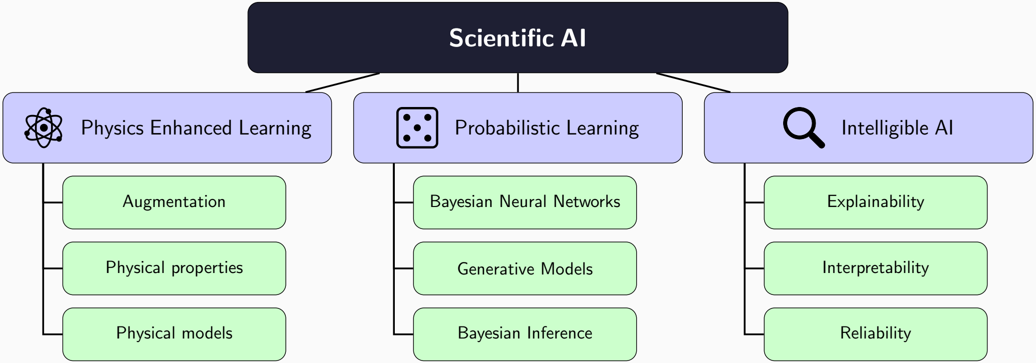

Physics Enhanced Learning

Embed physical understanding of the world into machine learning models.

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

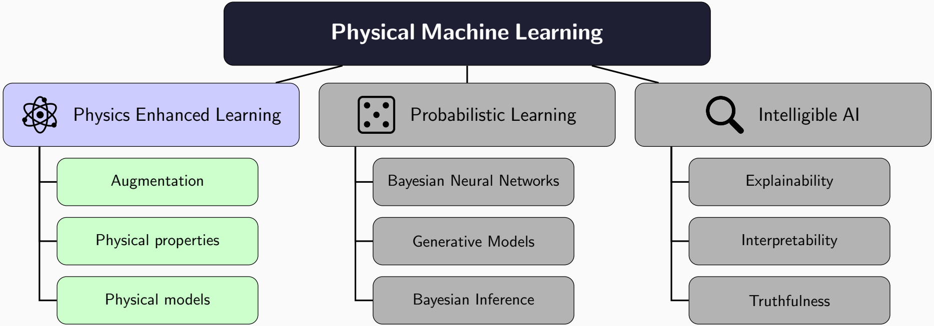

Augmentation

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

Augmentation

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

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

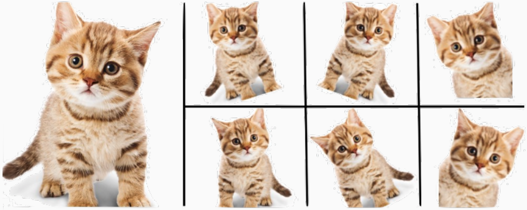

Image augmentation

Apply **physical transformations** that data known to satisfy to augment training data ⇝ ML model **learns physics through training**.

- Redshift augmentation of supernovae observations (Boone 2019; Alves, *et al.* 2022, 2023)

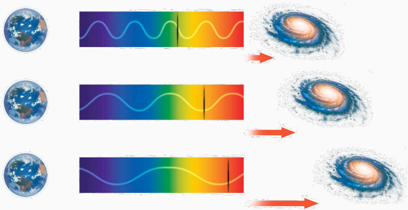

Redshift augmentation

Augmentation

Apply **physical transformations** that data known to satisfy to augment training data $\rightsquigarrow$ ML 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 ML models (e.g. geometry, symmetries, conservation laws) $\rightsquigarrow$ **Physics embedded in architecture** of ML model.

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

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

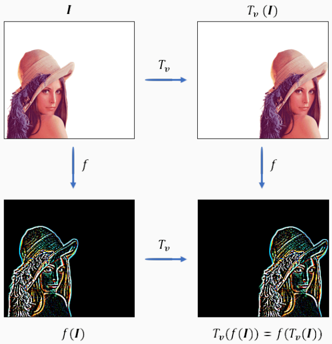

Translational equivariance

Physical properties: geometries, symmetries, conservation laws

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

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

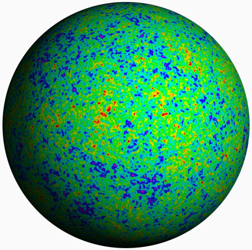

CMB observed on the celestial sphere

Physical properties: geometries, symmetries, conservation laws

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

- Equivariant machine learning, structured like classical physics (Villar *et al.* 2021)

Orthogonal	$\mathrm{O}(d) = \{Q \in \mathbb{R}^{d\times d} : Q^\top Q = Q\,Q^\top = I_d\}$,
Rotation	$\mathrm{SO}(d) = \{Q \in \mathbb{R}^{d\times d} : Q^\top Q = Q\,Q^\top = I_d,\ \det(Q) = 1\}$
Translation	$\mathrm{T}(d) = \{w \in \mathbb{R}^d\}$
Euclidean	$\mathrm{E}(d) = \mathrm{T}(d) \rtimes \mathrm{O}(d)$
Lorentz	$\mathrm{O}(1,d) = \{Q \in \mathbb{R}^{(d+1)\times(d+1)} : Q^\top \Lambda\, Q = \Lambda,\ \Lambda = \mathrm{diag}([1,-1,..,-1])\}$
Poincaré	$\mathrm{IO}(1,d) = \mathrm{T}(d+1) \rtimes \mathrm{O}(1,d)$
Permutation	$\mathrm{S}_n = \{\sigma : [n] \to [n] \text{ bijective function}\}$

Groups considered

SciAI UCL
Jason McEwen

Physical properties: geometries, symmetries, conservation laws

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

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

$\rightsquigarrow$ **Physics learned in training and embedded in model.**

Physical models: PINNS and differentiable physics

Encode physical models of world into ML models:

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

$\rightsquigarrow$ **Physics learned in training and embedded in model.**

- Physics informed neural networks (PINNs) encode differentiable equations (e.g. boundary conditions) in loss.

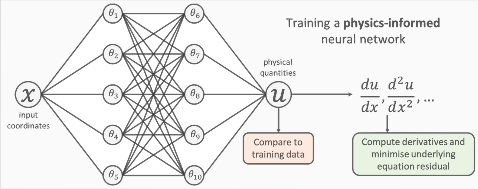

PINNs

SciAI UCL
Jason McEwen

Physical models: PINNS and differentiable physics

Encode physical models of world into ML models:

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

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

- Differentiable physical models
  - Radio interferometric telescope (Mars *et al.* 2023, 2024, Liaudat *et al.* McEwen 2024)
  - Optical PSF (Liaudat *et al.* 2023)
  - JAX-Cosmo (Campagne *et al.* 2023)

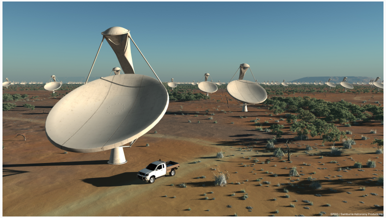

SKA (artist impression)

Physical models: PINNS and differentiable physics

Encode physical models of world into ML models:

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

$\rightsquigarrow$ **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)

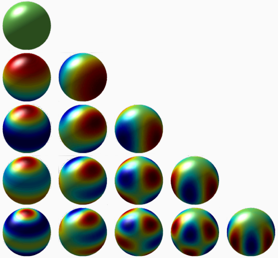

Spherical harmonics

Physical models: PINNS and differentiable physics

Encode physical models of world into ML models:

1. Encode dynamics (differential equations) via loss functions (PINNs).
2. Embed full (differentiable) physical models inside ML 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 (e.g. `CosmoPower`; Spurio Mancini et al. 2021).
- Write new differentiable codes (e.g. `s2fft`; Price & McEwen 2023).

Probabilistic Learning

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

(See Murray 2022.)

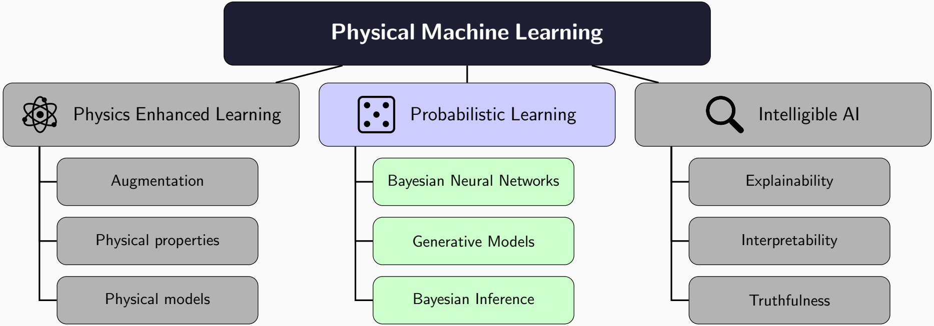

Bayesian neural networks for uncertainty quantification

Bayesian neural networks incorporate **probabilistic representation** to quantify **uncertainty of outputs** (idea pioneered by MacKay 1992).

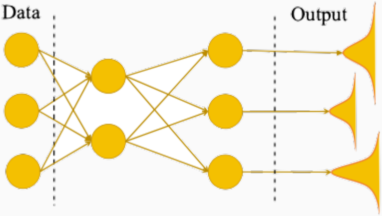

Bayesian neural networks for uncertainty quantification

Bayesian neural networks incorporate **probabilistic representation** to quantify **uncertainty of outputs** (idea pioneered by MacKay 1992).

- MC Dropout (Gal & Ghahramani 2016): drop nodes probabilistically to sample an ensemble of networks.

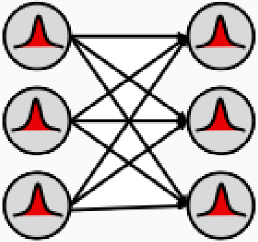

Bayesian neural networks for uncertainty quantification

Bayesian neural networks incorporate **probabilistic representation** to quantify **uncertainty of outputs** (idea pioneered by MacKay 1992).

- Bayes by Backprop (Blundel *et al.* 2015): model distribution of weights (by variational inference).

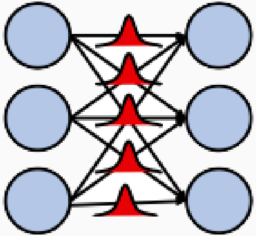

Bayesian neural networks for uncertainty quantification

Bayesian neural networks incorporate **probabilistic representation** to quantify **uncertainty of outputs** (idea pioneered by MacKay 1992).

- Probabilistic ML frameworks (*e.g.* TensorFlow Probability).

Bayesian neural networks for uncertainty quantification

Bayesian neural networks incorporate **probabilistic representation** to quantify **uncertainty of outputs** (idea pioneered by MacKay 1992).

- Encode epistemic uncertainty of model.
- But what does the output distribution represent?
- Requires careful consideration of training data.

- Statistical validation (hold that thought... see upcoming Truthfulness section).

Generative models

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

Generative models

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

- Emulation: sample from learned prior
  (Perraudin *et al.* 2020, Allys *et al.* 2020, Price *et al.* 2023, Price *et al.* in prep., Mousset *et al.* McEwen 2024)

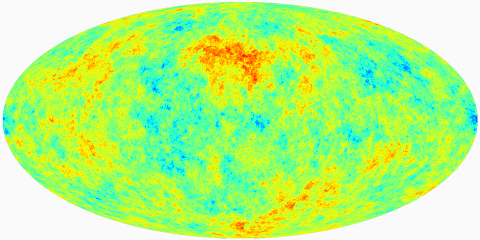

Emulated cosmic string maps
(`stringgen`, Price *et al.* 2023, Price *et al.* in prep.)

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

- Integrate learned priors into analysis
  (Remy *et al.* 2022, McEwen *et al.* 2023, Liaudat *et al.* McEwen 2024)

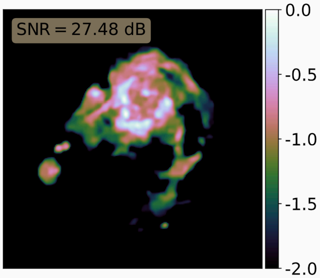

Learn radio galaxy prior
(Liaudat *et al.* McEwen 2024)

Generative models

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

- Availability and representativeness of training data.
- Truthfulness, *e.g.* diversity of ML model 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).
- Truthfulness (hold that thought... see upcoming Truthfulness section).

Bayesian inference

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

Bayesian inference

ML 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).

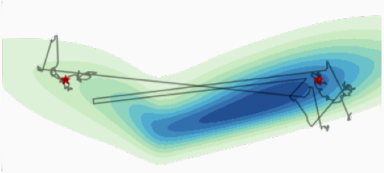

Learned proposal distributions

Bayesian inference

ML 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.* 2024, Piras *et al.* McEwen 2024, Spurio Mancini *et al.* McEwen 2023, 2024).
  - Only requires posterior samples.
  - Agnostic to sampling technique.
  - Scale to high dimensions.

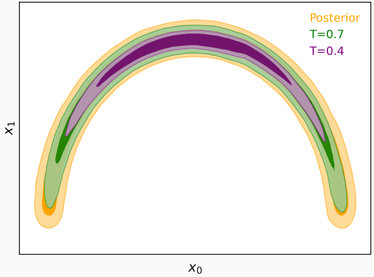

Learned harmonic mean estimator (**`harmonic`**)

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

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

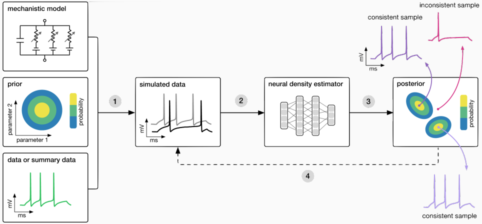

sbi

Bayesian inference

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

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

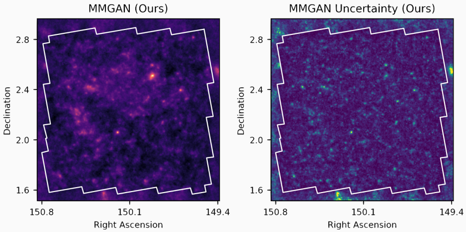

Mass mapping with uncertainties by variational inference (Whitney *et al.* McEwen 2024)

SciAI UCL
Jason McEwen

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

- Availability and representativeness of training data.
- Cost of training.
- Truthfulness?

- 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 statistical guarantees.
- Statistical validation (hold that thought... see upcoming Truthfulness section).

Intelligible AI

Machine learning methods that are able to be understood by humans.

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

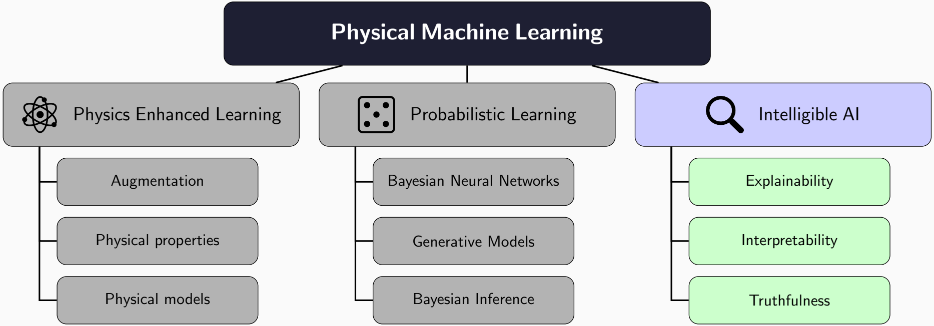

Explainability

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

Explainability

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

- Feature importances
  (Lochner *et al.* 2016)

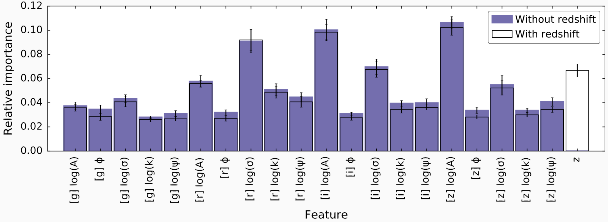

Supernova feature importances

Jason McEwen

Explainability

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

- Saliency maps (Bhambra *et al.* 2022)

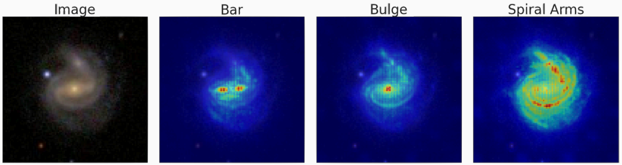

Galaxy saliency mapping

Explainability

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

Poking the black box: may provide some explanation of outputs but humans still not able to comprehend underlying process.

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

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

- Designed models such as wavelet scattering networks
  (Allys *et al.* 2020, Cheng *et al.* 2020, McEwen *et al.* 2022, Mousset *et al.* McEwen 2024)

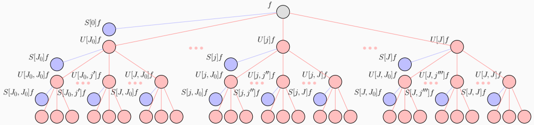

Scattering network (McEwen *et al.* 2022)

Interpretability

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

- Designed models such as wavelet scattering networks
  (Allys *et al.* 2020, Cheng *et al.* 2020, McEwen *et al.* 2022, Mousset *et al.* McEwen 2024)

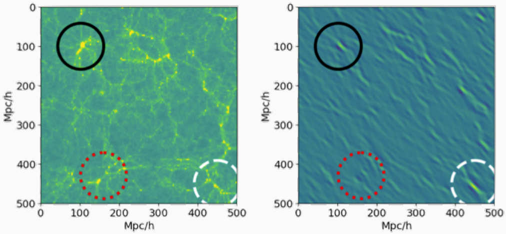

LSS features captured by wavelets
(Allys *et al.* 2020)

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

- Interpretable constraints on ML models, *e.g.* convexity
  (Liaudat *et al.* McEwen 2024)

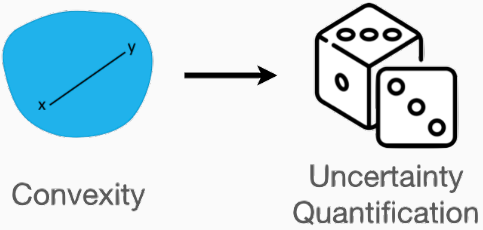

Impose convexity on learned model

Interpretability

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

- Deep priors learned from training data (hybrid model-based and data-driven)
  (Remy *et al.* 2022, McEwen *et al.* 2023, Liaudat *et al.* McEwen 2024)

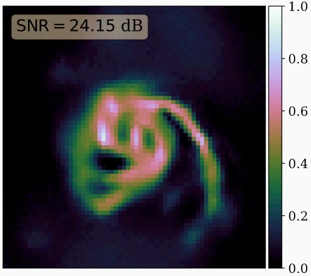

Compute Bayesian evidence for model selection
(**proxnest**, McEwen *et al.* 2023)

Interpretability

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

- ▷ Designed models limit flexibility.
- ▷ Availability and representativeness of training data.

- ▷ Benefits of designed models often outweigh (minimal) reduced flexibility.
- ▷ Public datasets/benchmarks (*e.g.* IllustrisTNG, CAMELS, Quijote, CosmoGrid, Gower St).

Truthfulness

Truthfulness **critical for science** in order for humans to have confidence in results of ML models. Closely coupled with a **meaningful statistical distribution** of outputs.

Truthfulness

Truthfulness **critical for science** in order for humans to have confidence in results of ML 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)
  - Design to ensure conservative and avoid mode collapse (Delaunoy *et al.* 2022, Price *et al.* 2023, Whitney *et al.* McEwen 2024)
  - Coverage testing (Lemos *et al.* 2023)
  - Simulation-based calibration checks (Talts *et al.* 2020)

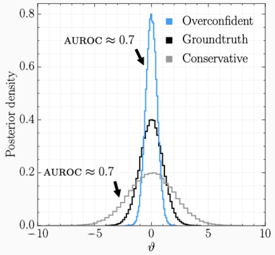

Validity of distribution (Hermans *et al.* 2022)

Truthfulness

Truthfulness **critical for science** in order for humans to have confidence in results of ML models. Closely coupled with a **meaningful statistical distribution** of outputs.

- Uncertainties not aways meaningful.
- Diversity of ML model often lacking.

- Integrate in statistical framework to inherit theoretical guarantees.
- Design to be conservative and avoid mode collapse.
- Extensive validation tests.
- Well-posed frameworks (*e.g.* physics enhanced, probabilistic).

Statistical characterisation and generative modelling of cosmological fields

Wavelet scattering networks and representations inspired by CNNs but designed rather than learned filters (Mallat 2012).

⇝ **Scattering networks on the sphere**
(McEwen et al. 2022, ICLR, arXiv:2102.02828)

⇝ **Generative models of astrophysical fields with scattering transforms on the sphere**
(Mousset *et al.* McEwen 2024, A&A, arXiv:2407.07007)

Wavelets on the sphere

Adopt scale-discretized wavelets on the sphere (e.g. McEwen et al. 2018, McEwen et al. 2015).

Wavelets $\psi_j \in \mathrm{L}^2(\mathbb{S}^2)$ capture spatially-localised, high-frequency signal content at scale *j*.

Scaling function $\phi \in \mathrm{L}^2(\mathbb{S}^2)$ captures spatially-localised, low-frequency content.

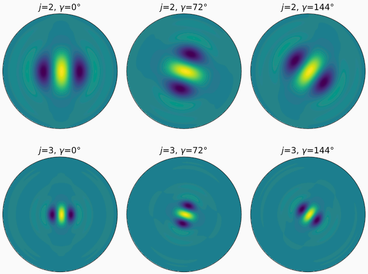

Orthographic plot of spherical wavelets.

Wavelets on the sphere

Adopt scale-discretized wavelets on the sphere (e.g. McEwen et al. 2018, McEwen et al. 2015).

Wavelets $\psi_j \in \mathrm{L}^2(\mathbb{S}^2)$ capture spatially-localised, high-frequency signal content at scale *j*.

Scaling function $\phi \in \mathrm{L}^2(\mathbb{S}^2)$ captures spatially-localised, low-frequency content.

Spherical wavelet transform given by

$$W_j(\rho) = (f \star \psi_j)(\rho) = \int_{\mathbb{S}^2} \mathrm{d}\mu(\omega') f(\omega') (R_\rho \psi_j)^*(\omega').$$

Spherical convolution

Rotated wavelet

Fast algorithms available

(e.g. McEwen et al. 2007, 2013, 2015).

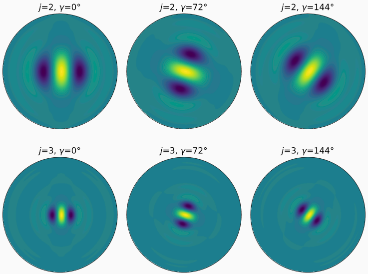

Orthographic plot of spherical wavelets.

Scattering transform on the sphere

Spherical scattering propagator for scale j:

$$U[j]f = |f \star \psi_j|.$$

Modulus function is adopted for the activation function (since non-expansive and preserves stability of wavelet representation).

Scattering transform on the sphere

Spherical scattering propagator for scale j:

$$U[j]f = |f \star \psi_j|.$$

Modulus function is adopted for the activation function (since non-expansive and preserves stability of wavelet representation).

Spherical cascade of propagators:

$$U[p]f = |||f \star \psi_{j_1}| \star \psi_{j_2}| \dots \star \psi_{j_d}|,$$

for the path $p = (j_1, j_2, \dots, j_d)$ with depth d.

Spherical scattering propagator for scale j:

$$U[j]f = |f \star \psi_j|.$$

Modulus function is adopted for the activation function (since non-expansive and preserves stability of wavelet representation).

Spherical cascade of propagators:

$$U[p]f = |||f \star \psi_{j_1}| \star \psi_{j_2}| \dots \star \psi_{j_d}|,$$

for the path $p = (j_1, j_2, \dots, j_d)$ with depth d.

Scattering coefficients:

$$S[p]f = |||f \star \psi_{j_1}| \star \psi_{j_2}| \dots \star \psi_{j_d}| \star \phi.$$

Spherical scattering network is collection of scattering transforms for a number of paths: $\mathcal{S}_{\mathbb{P}} f = \{S[p]f : p \in \mathbb{P}\}$, where the general path set $\mathbb{P}$ denotes the infinite set of all possible paths $\mathbb{P} = \{p = (j_1, j_2, \ldots, j_d) : J_0 \leq j_i \leq J, 1 \leq i \leq d,\ d \in \mathbb{N}_0\}$.

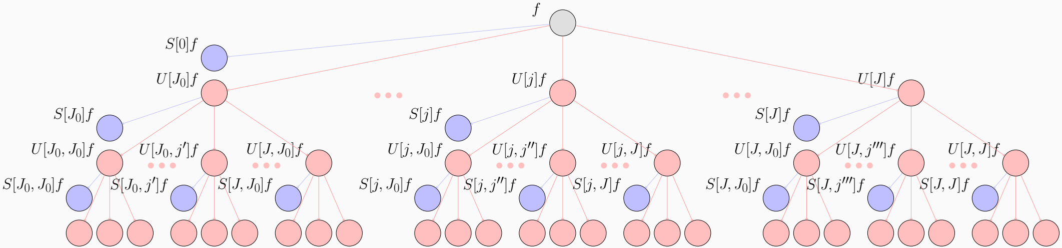

Capture all information content at infinite depth and typically $> 99\%$ for depth $d = 3$.

Properties

Latent representation is very well-behaved and satisfies a number of important properties:

1. **Rotational equivariance**
2. **Isometric invariance**
3. **Stability to diffeomorphisms**

Rotationally equivariance

Rotational Equivariance

$$((\mathcal{R}_\rho f) \star \psi)(\rho') = (\mathcal{R}_\rho (f \star \psi))(\rho').$$

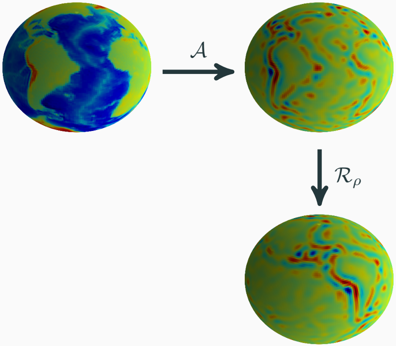

Rotational Equivariance

$$((\mathcal{R}_\rho f) \star \psi)(\rho') = (\mathcal{R}_\rho (f \star \psi))(\rho').$$

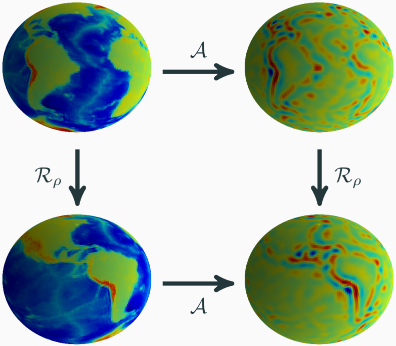

Isometric invariance

Isometric Invariance

Let $\zeta \in \mathrm{Isom}(\mathbb{S}^2)$, then there exists a constant C such that for all $f \in \mathrm{L}^2(\mathbb{S}^2)$,

$$\|\mathcal{S}_{\mathbb{P}_D} f - \mathcal{S}_{\mathbb{P}_D} V_\zeta f\|_2 \leq C L^{5/2} (D+1)^{1/2} \lambda^{J_0} \|\zeta\|_\infty \|f\|_2 .$$

Difference in representation.

Scattering network representation is invariant to isometries up to a scale.

Isometric invariance

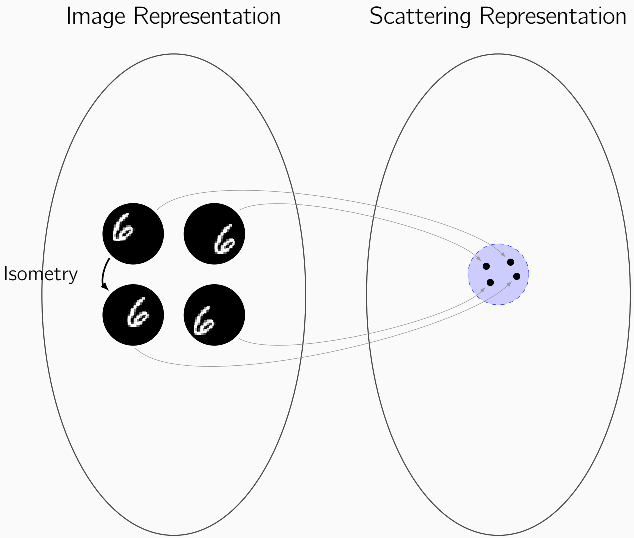

Stability to Diffeomorphisms

Let $\zeta \in \mathrm{Diff}(\mathbb{S}^2)$. If $\zeta = \zeta_1 \circ \zeta_2$ for some isometry $\zeta_1 \in \mathrm{Isom}(\mathbb{S}^2)$ and diffeomorphism $\zeta_2 \in \mathrm{Diff}(\mathbb{S}^2)$, then there exists a constant C such that for all $f \in \mathrm{L}^2(\mathbb{S}^2)$,

$$\|\mathcal{S}_{\mathbb{P}_D} f - \mathcal{S}_{\mathbb{P}_D} V_\zeta f\|_2 \leq C L^2 \big[L^2 \|\zeta_2\|_\infty + L^{1/2}(D+1)^{1/2} \lambda^{J_0} \|\zeta_1\|_\infty \big] \|f\|_2 .$$

Difference in representation.

Scattering network representation is stable to small diffeomorphisms about isometry.

Jason McEwen

Stability to diffeomorphisms

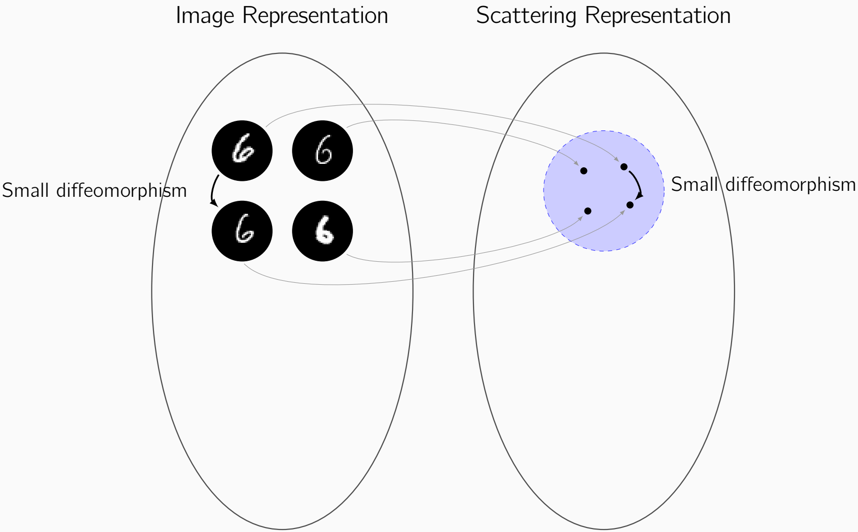

Stability to diffeomorphisms

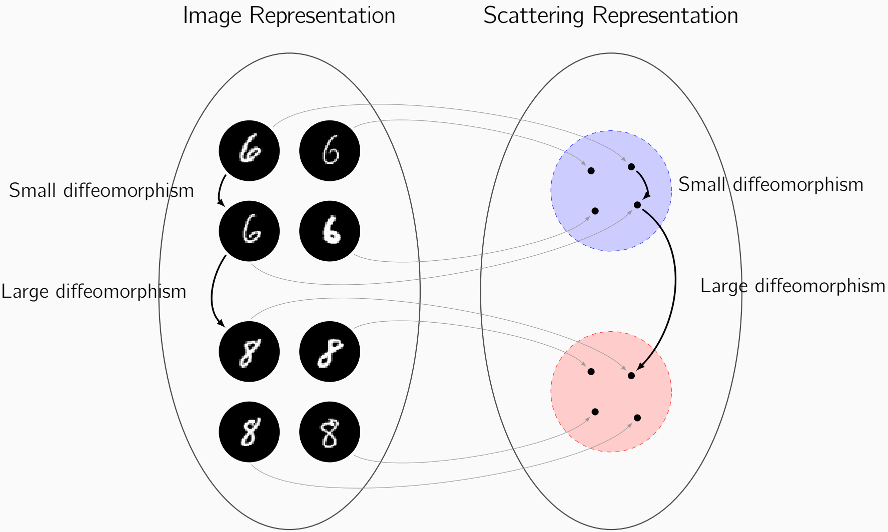

Toy problem: Gaussianity of the cosmic microwave background (CMB)

Wavelet scattering as a representation space for classification.

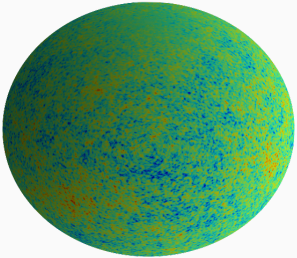

Gaussian

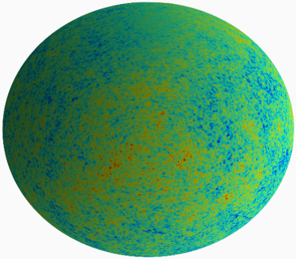

Non-Gaussian

⇝ 53% classification accuracy without scattering versus 95% with scattering.

Scattering for simulation-based inference (SBI)

Wavelet scattering as a representation space for SBI (Lin, Joachimi & McEwen 2024).

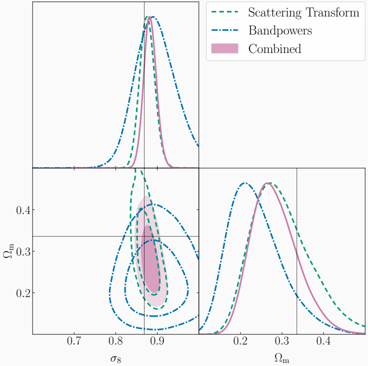

Posterior

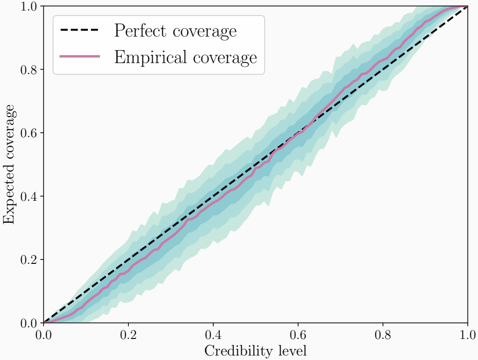

Coverage

SciAI UCL

Jason McEwen

Spherical scattering covariance for generative modelling

Generative models of astrophysical fields with scattering transforms on the sphere

(Mousset *et al.* McEwen 2024; `s2scat` code)

Scattering covariance statistics:

1. $S_1[\lambda]\, f = \mathbb{E}\big[\, |f \star \psi_\lambda| \,\big]$.
2. $S_2[\lambda]\, f = \mathbb{E}\big[\, |f \star \psi_\lambda|^2 \,\big]$.
3. $S_3[\lambda_1, \lambda_2]\, f = \mathrm{Cov}\big[\, f \star \psi_{\lambda_2}, |f \star \psi_{\lambda_1}| \star \psi_{\lambda_2} \,\big]$.
4. $S_4[\lambda_1, \lambda_2, \lambda_3]\, f = \mathrm{Cov}\big[\, |f \star \psi_{\lambda_1}| \star \psi_{\lambda_3}, |f \star \psi_{\lambda_2}| \star \psi_{\lambda_3} \,\big]$.

Generative models of astrophysical fields with scattering transforms on the sphere

(Mousset *et al.* McEwen 2024; **`s2scat`** code)

Scattering covariance statistics:

1. $S_1[\lambda]\, f = \mathbb{E}\big[\, |f \star \psi_\lambda| \,\big]$.
2. $S_2[\lambda]\, f = \mathbb{E}\big[\, |f \star \psi_\lambda|^2 \,\big]$.
3. $S_3[\lambda_1, \lambda_2]\, f = \mathrm{Cov}\big[\, f \star \psi_{\lambda_2}, |f \star \psi_{\lambda_1}| \star \psi_{\lambda_2} \,\big]$.
4. $S_4[\lambda_1, \lambda_2, \lambda_3]\, f = \mathrm{Cov}\big[\, |f \star \psi_{\lambda_1}| \star \psi_{\lambda_3}, |f \star \psi_{\lambda_2}| \star \psi_{\lambda_3} \,\big]$.

Generative modelling by **matching set of scattering covariance statistics** $\mathcal{S}(f)$ with a (single) target simulation:

$$\min_f \|\mathcal{S}(f) - \mathcal{S}(f_{\text{target}})\|^2 .$$

SciAI UCL

Jason McEwen

Differentiable and GPU-accelerated spherical transform codes (in JAX)

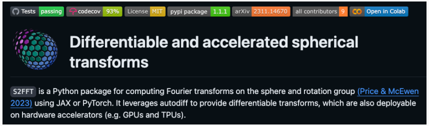
Tests passing | codecov 93% | License MIT | pypi package 1.1.1 | arXiv 2311.14670 | all contributors 9 | Open in Colab

Differentiable and accelerated spherical transforms

S2FFT is a Python package for computing Fourier transforms on the sphere and rotation group (Price & McEwen 2023) using JAX or PyTorch. It leverages autodiff to provide differentiable transforms, which are also deployable on hardware accelerators (e.g. GPUs and TPUs).

`s2fft`: Spherical harmonic transforms

https://github.com/astro-informatics/s2fft

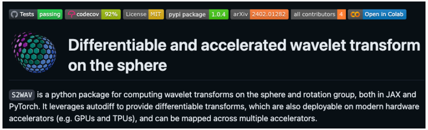
Tests passing | codecov 92% | License MIT | pypi package 1.0.4 | arXiv 2402.01282 | all contributors 4 | Open in Colab

Differentiable and accelerated wavelet transform on the sphere

S2WAV is a python package for computing wavelet transforms on the sphere and rotation group, both in JAX and PyTorch. It leverages autodiff to provide differentiable transforms, which are also deployable on modern hardware accelerators (e.g. GPUs and TPUs), and can be mapped across multiple accelerators.

`s2wav`: Spherical wavelet transforms

https://github.com/astro-informatics/s2wav

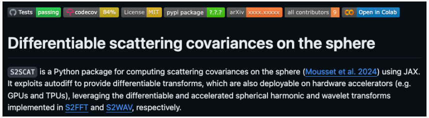
Tests passing | codecov 84% | License MIT | pypi package ?.?.? | arXiv xxxx.xxxxx | all contributors 9 | Open in Colab

Differentiable scattering covariances on the sphere

S2SCAT is a Python package for computing scattering covariances on the sphere (Mousset et al. 2024) using JAX. It exploits autodiff to provide differentiable transforms, which are also deployable on hardware accelerators (e.g. GPUs and TPUs), leveraging the differentiable and accelerated spherical harmonic and wavelet transforms implemented in S2FFT and S2WAV, respectively.

`s2scat`: Spherical scattering transforms

https://github.com/astro-informatics/s2scat

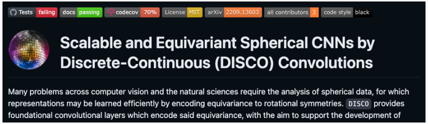
Tests failing | docs passing | codecov 70% | License MIT | arXiv 2209.13603 | all contributors 3 | code style black

Scalable and Equivariant Spherical CNNs by Discrete-Continuous (DISCO) Convolutions

Many problems across computer vision and the natural sciences require the analysis of spherical data, for which representations may be learned efficiently by encoding equivariance to rotational symmetries. DISCO provides foundational convolutional layers which encode said equivariance, with the aim to support the development of

`s2ai`: Spherical AI

Coming very soon! Contact us for early access.

Generative modelling of large scale structure (LSS)

Which field is emulated and which simulated?

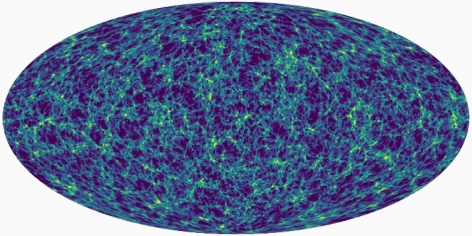

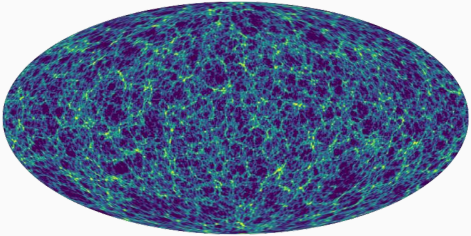

Logarithm (for visualization) of weak lensing field.

Generative modelling of large scale structure (LSS)

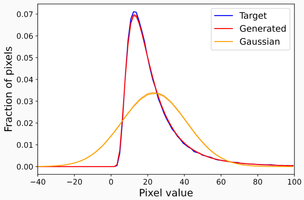

Pixel distribution

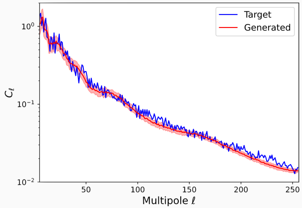

Power spectrum

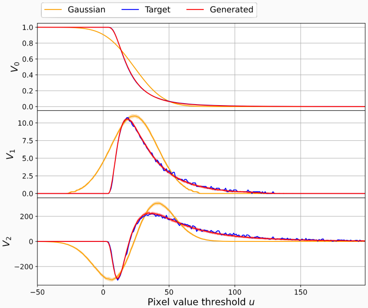

Minkowski functionals

Statistical validation.

Jason McEwen

Generative modelling of cosmic strings in the CMB

Need to **simulate full physics**, evolving a network of strings through cosmic time, and then ray-trace CMB photons through the string network (Ringeval et al. 2012).

A single simulation requires 800,000 CPU hours on a supercomputer.

There are **only three full-sky string maps in existence**.

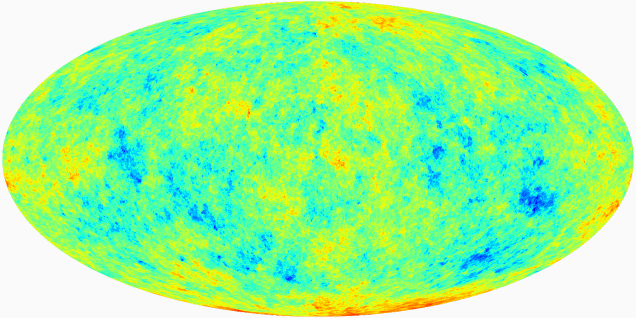

Generative modelling of cosmic strings in the CMB

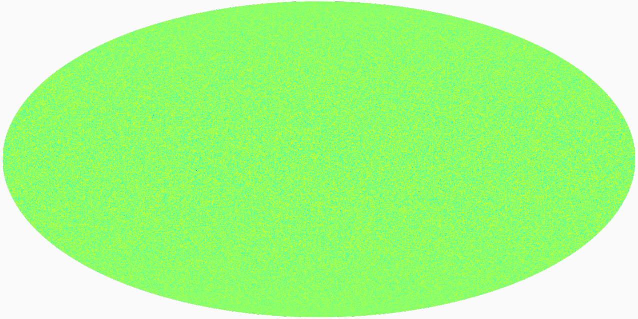

Computation time: 800,000 CPU hours on supercomputer $\rightarrow \mathcal{O}(1)$ hours on A100 GPU.

Still work in progress (statistical validation in progress).

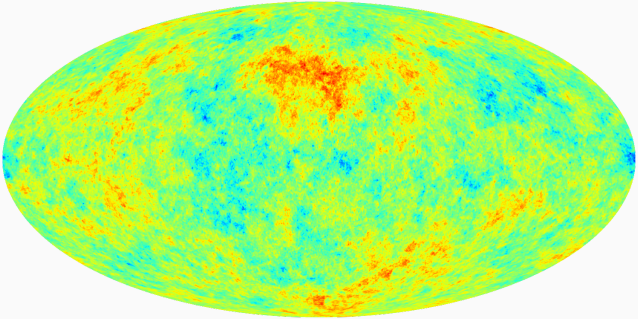

Computation time: 800,000 CPU hours on supercomputer $\rightarrow \mathcal{O}(1)$ hours on A100 GPU.

Still work in progress (statistical validation in progress).

Summary

Scientific AI

- **Physics Enhanced Learning**
  - Augmentation
  - Physical properties
  - Physical models
- **Probabilistic Learning**
  - Bayesian Neural Networks
  - Generative Models
  - Bayesian Inference
- **Intelligible AI**
  - Explainability
  - Interpretability
  - Reliability

With great power comes great responsibility!