Machine Learning for Physics; Physics for Machine Learning

Jason D. McEwen

www.jasonmcewen.org

Mullard Space Science Laboratory (MSSL), University College London (UCL)

European Astronomical Society (EAS) Annual Meeting 2023

The machine learning hammer

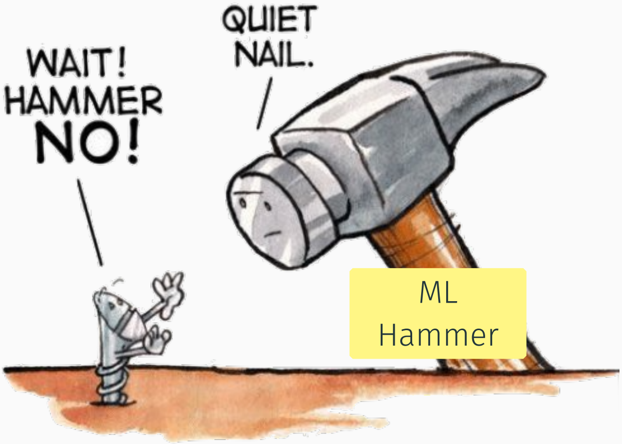

The machine learning cog

Merging paradigms

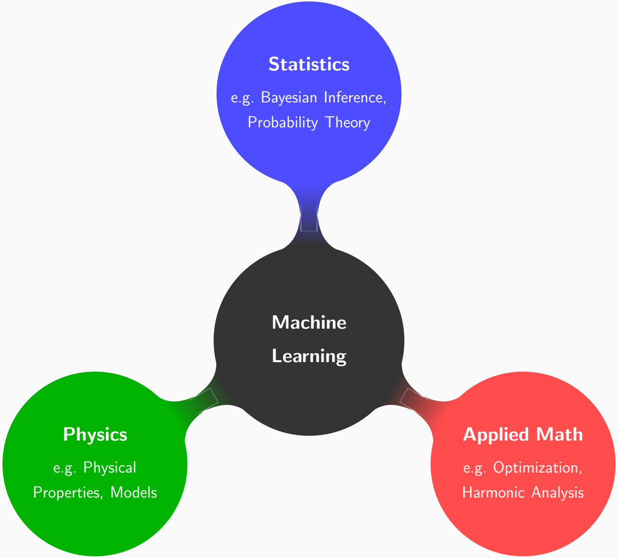

Outline

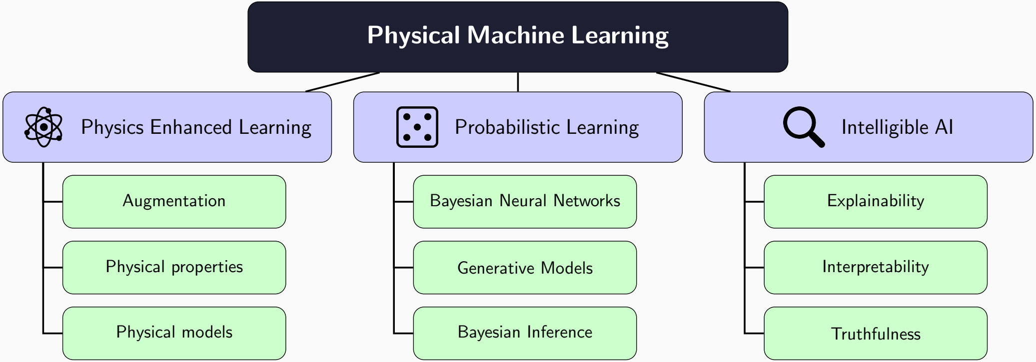

Physical Machine Learning
Physics Enhanced Learning
Augmentation
Physical properties
Physical models
Probabilistic Learning
Bayesian Neural Networks
Generative Models
Bayesian Inference
Intelligible AI
Explainability
Interpretability
Truthfulness

Physics Enhanced Learning

Physics Enhanced Learning

Embed physical understanding of the world into machine learning models.

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

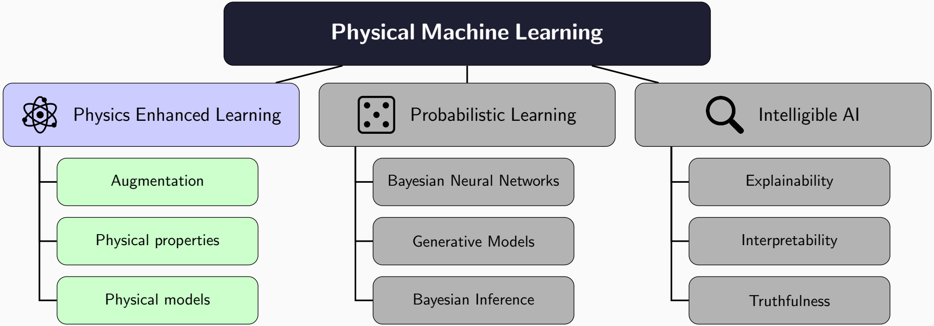

Augmentation

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, ...

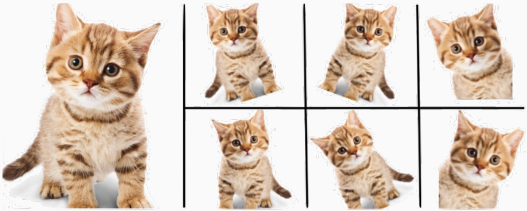

Image augmentation

Augmentation

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

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

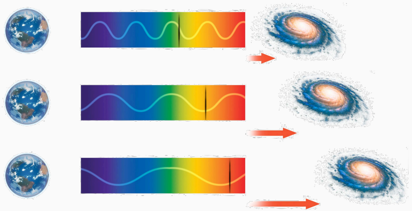

Redshift augmentation

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.

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.

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.

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

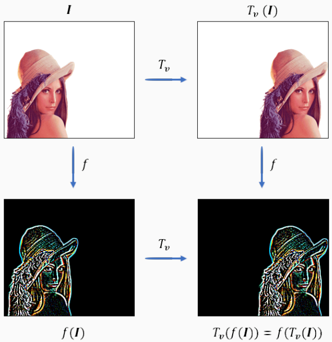

Physical properties: geometries, symmetries, conservation laws

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)

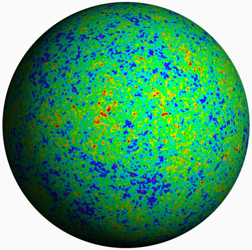

CMB observed on the celestial sphere

Physical properties: geometries, symmetries, conservation laws

Encode physical properties of the world into ML models (e.g. geometry, symmetries, conservation laws) $\rightsquigarrow$ **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	$\mathrm{O}(1,d) = \{Q \in \mathbb{R}^{(d+1)\times(d+1)} : Q^\top \Lambda\, Q = \Lambda,\ \Lambda = \mathrm{diag}([1,-1,\ldots,-1])\}$
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

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.

- Highly computationally demanding.
- Always required?

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.

- Highly computationally demanding.
- Always required?

- Develop efficient algorithms (e.g. Ocampo, Price & McEwen 2023).
- Inductive biases not enforced.

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.

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

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

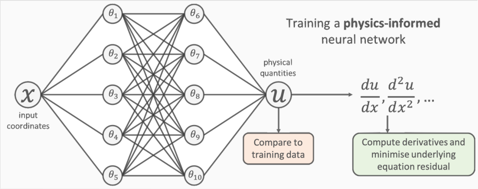

PINNs

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

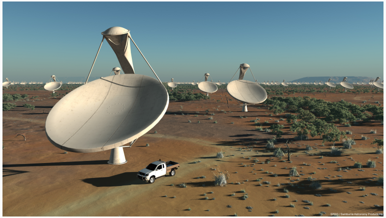

SKA (artist impression)

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.

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

- Differentiable mathematical methods
  - Fourier transforms
  - Spherical harmonic transforms (**`s2fft`**; Price & McEwen, in prep.)
  - Spherical wavelet transforms (**`s2wav`**; Price *et al.* in prep.)
  - Spherical scattering transforms (Mousset, Price, Allys, McEwen, in prep.)

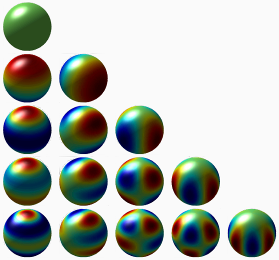

Spherical harmonics

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.

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

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

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, in prep.).

Probabilistic Learning

Probabilistic Learning

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

(See Murray 2022.)

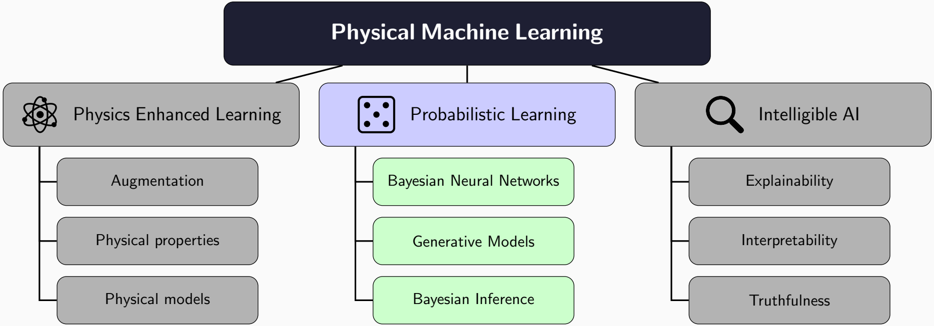

Bayesian neural networks for uncertainty quantification

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

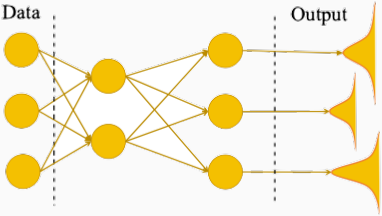

Bayesian neural networks for uncertainty quantification

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.

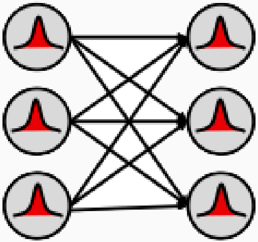

Bayesian neural networks for uncertainty quantification

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).

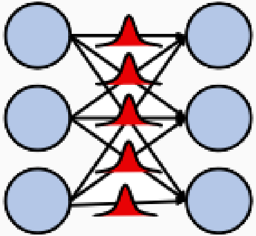

Bayesian neural networks for uncertainty quantification

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

Bayesian neural networks for uncertainty quantification

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

- Encode epsitemic 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 **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.)

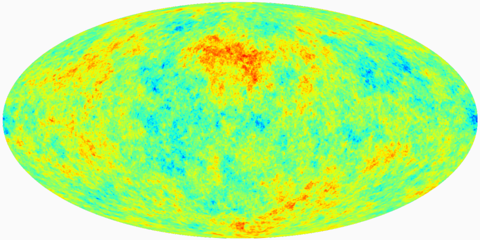

Emulated cosmic string maps
(**stringgen**, Price *et al.* 2023, Price *et al.* in prep.)

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)

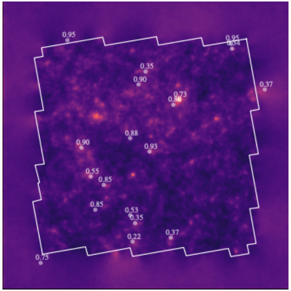

Learn convergence field prior (Remy *et al.* 2022)

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.

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).
- ▷ Meta sampling to recover distribution over manifold (*e.g.* Price *et al.* 2023).
- ▷ Truthfulness (hold that thought... see upcoming Truthfulness section).

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).

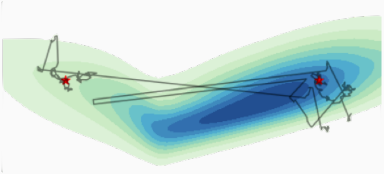

Learned proposal distributions

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
(McEwen *et al.* 2021, Polanska *et al.* 2023).

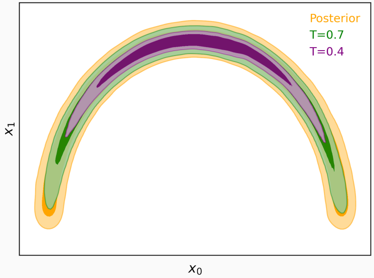

Learned harmonic mean estimator (**harmonic**)

Bayesian inference

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 (Alsing *et al.* 2018, Cranmer *et al.* 2021).
- Model selection for simulation-based inference (`harmonic`; Spurio Mancini *et al.* 2022)

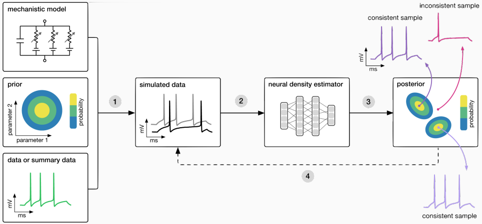

`sbi`

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.* in prep.)

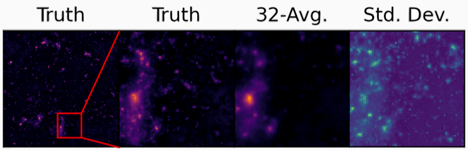

Mass mapping with uncertatinties
by variational inference

Bayesian inference

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

- Cost of training.
- Truthfulness?

Bayesian inference

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

- Cost of training.
- Truthfulness?

- 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

Intelligible AI

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

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

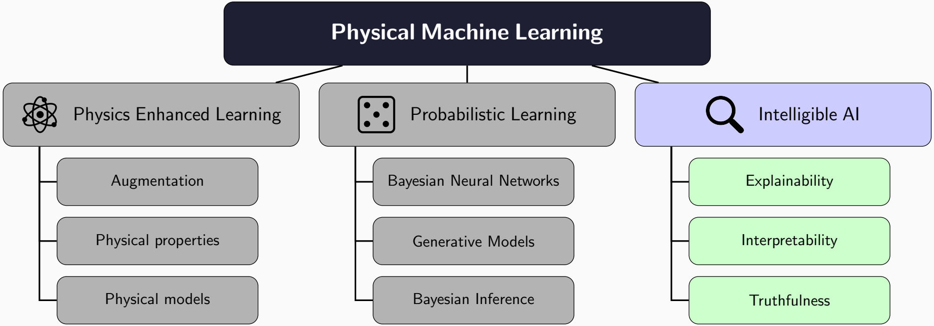

Explainability

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

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

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

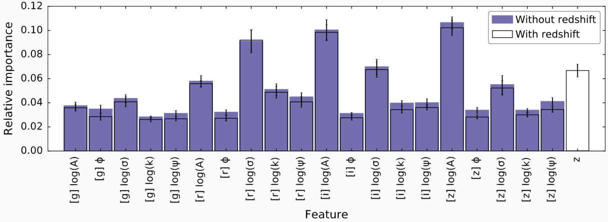

Supernova feature importances

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)

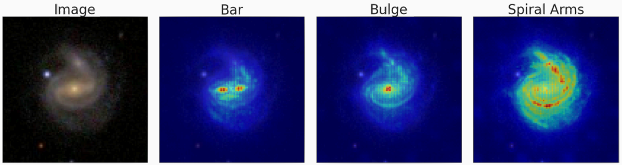

Galaxy saliency mapping

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.

Interpretability

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

Interpretability

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

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

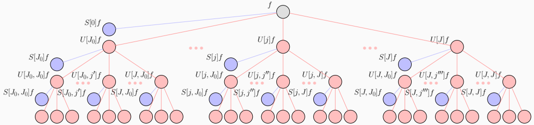

Scattering network (McEwen *et al.* 2022)

Interpretability

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

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

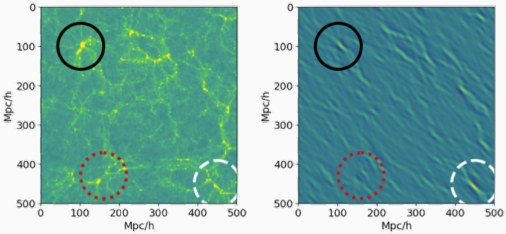

LSS features captured by wavelets (Allys *et al.* 2020)

Interpretability

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

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

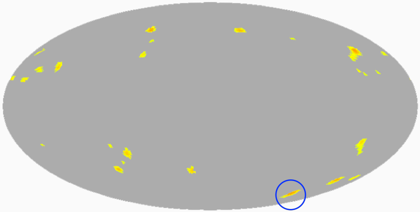

First evidence that CMB cold spot due to supervoid (McEwen *et al.* 2007)

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

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

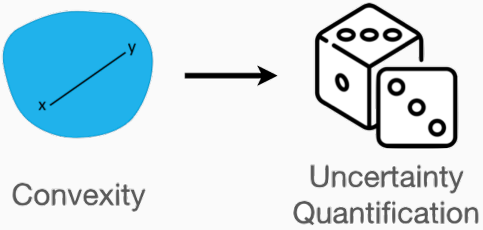

Impose convexity on learned model

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)

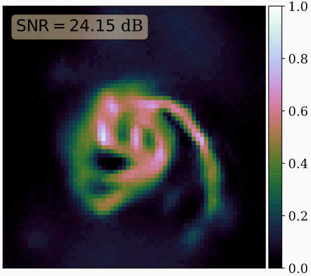

Compute Bayesian evidence for model selection (**`proxnest`**, McEwen *et al.* 2023)

Interpretability

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

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

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 outway (minimal) reduced flexibility.
- Public datasets/benchmarks (*e.g.* IllustrisTNG, CAMELS, Quijote, CosmoGrid).
- Transfer learning, self-supervised learning.

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
  (Hermans *et al.* 2022, Lemos *et al.* 2023)

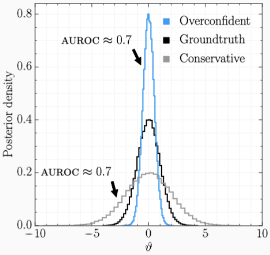

Validity of distribution
(Hermans *et al.* 2022)

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.

- Validity of statistical distributions
  (Hermans *et al.* 2022, Lemos *et al.* 2023)

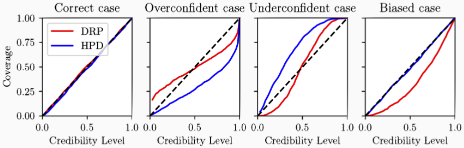

Coverage analysis (Lemos *et al.* 2023)

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.

- Diversity (avoiding mode-collapse)
  (Price *et al.* 2023, Whitney *et al.* in prep.)

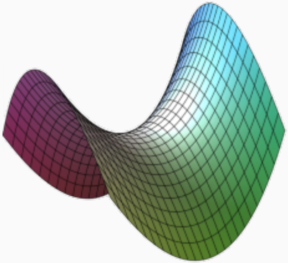

Recover probabily ditribution
over full underlying manifold

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.

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.
- Extensive validation tests (*e.g.* Hermans *et al.* 2022, Lemos *et al.* 2023).
- Meta sampling to recover distribution over manifold (*e.g.* Price *et al.* 2023).
- Well-posed frameworks (*e.g.* physics enhanced, probabilistic).

Summary

Summary

Physical Machine Learning

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

With great power comes great responsibility!