

# Scientific machine learning: ways to incorporate scientific principles into ML

Ben Moseley

Postdoctoral fellow, ETH AI Center /  
ETH Computational and Applied Mathematics Laboratory




**ETH** zürich



# Talk overview

- What is scientific machine learning (SciML)?
- Ways to incorporate scientific principles into ML
- Our research: scaling SciML techniques to complex, real-world problems



Later in the talk: peering into shadows on the Moon using SciML

# Talk overview

- What is scientific machine learning (SciML)?
- Ways to incorporate scientific principles into ML
- Our research: scaling SciML techniques to complex, real-world problems

# What is deep learning?


Trained using:

- (Stochastic) gradient descent
- An appropriate loss function
- Many (millions of) training examples



Input

$\mathbf{x}$



Model

$\text{NN}(\mathbf{x}, \theta)$

Probability = Dog

Output

$y = \text{NN}(\mathbf{x}, \theta)$

For example:

$$y = W_2 \sigma(W_1 x + b_1) + b_2$$



# The challenge of generalisation

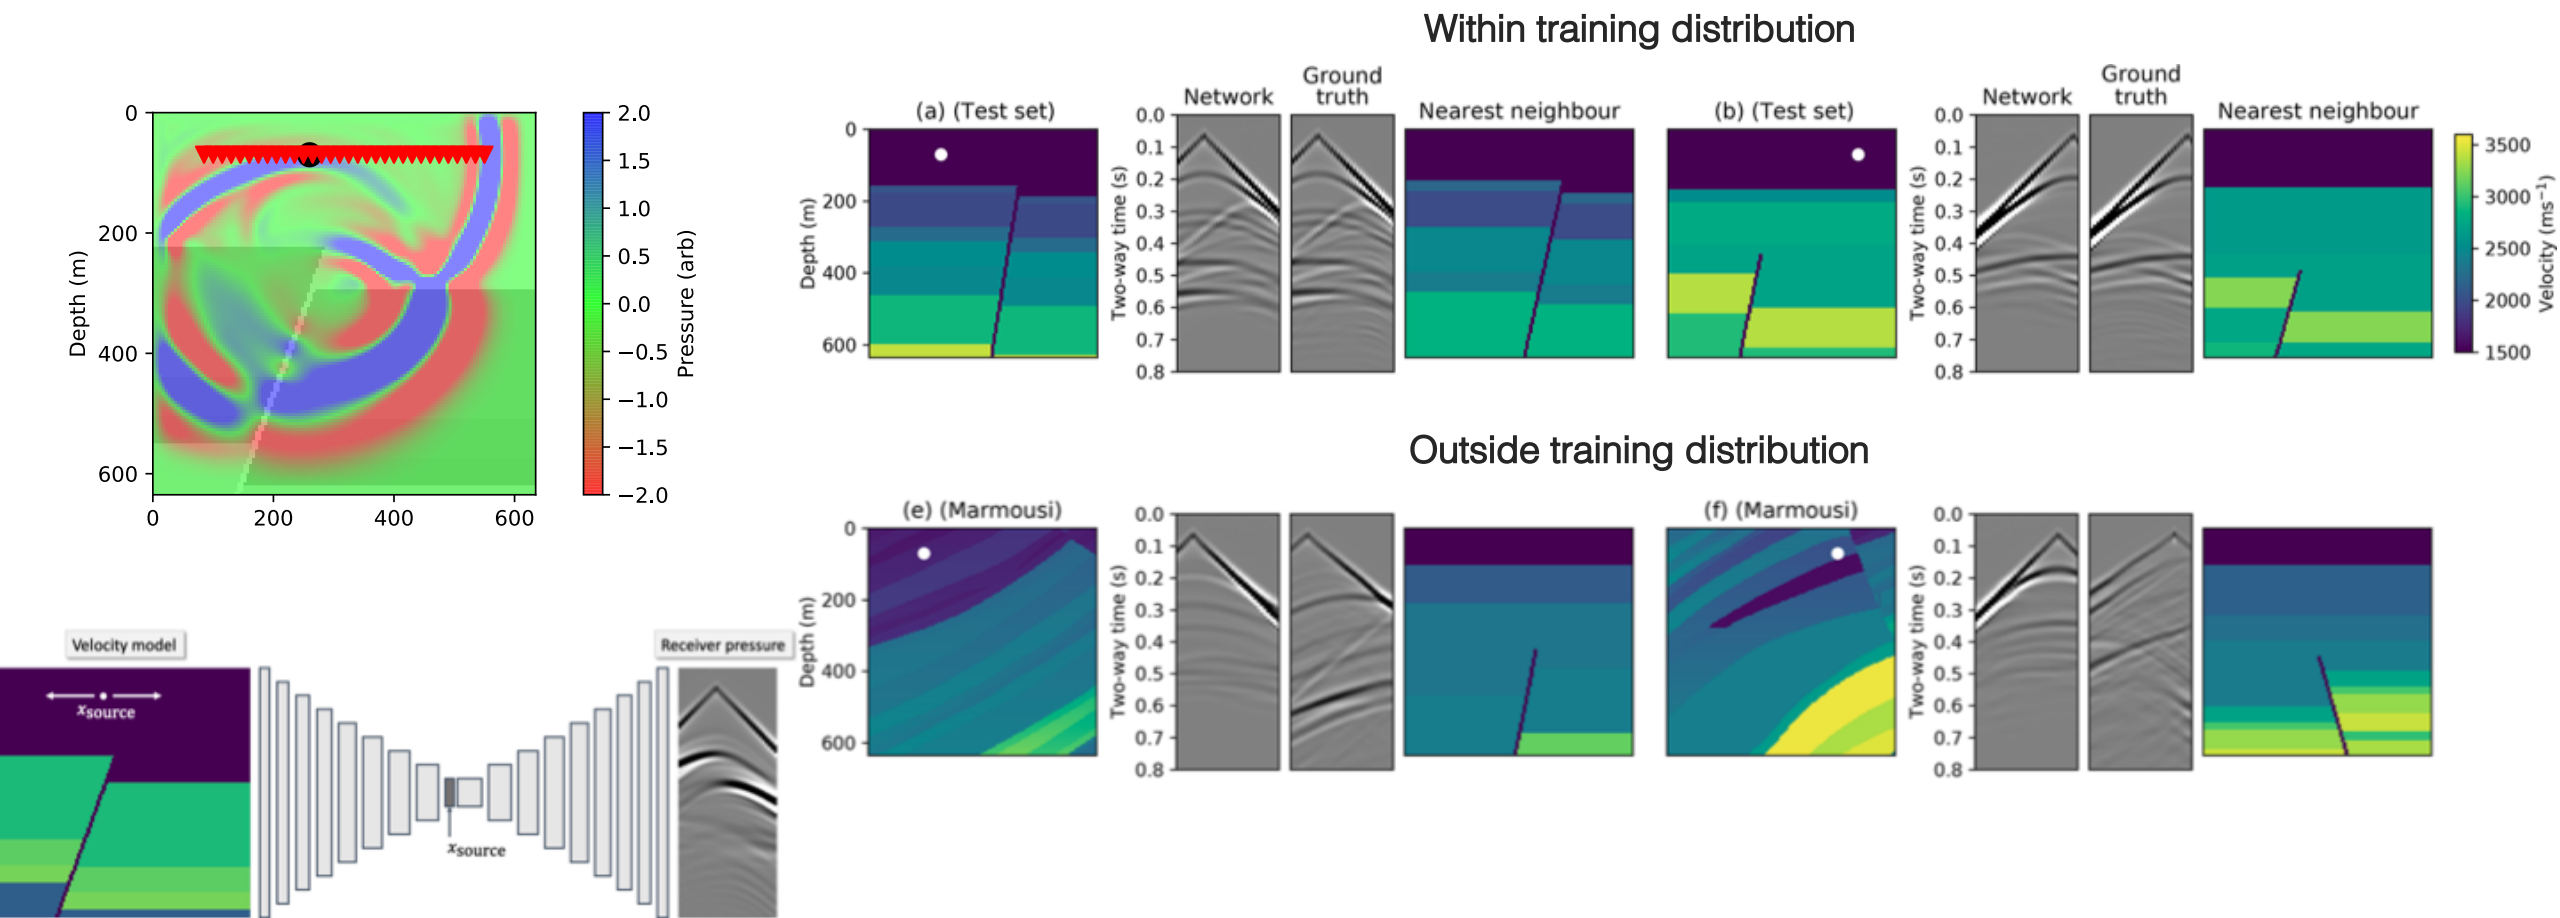


Labrador retriever 52%  
Chesapeake Bay retriever 7%  
golden retriever 5%  
Canis dingo 4%  
bloodhound, sleuthound 3%



laboratory coat 40%  
jeweler's loupe 8%  
English foxhound 6%  
soccer ball 4%  
neck brace 3%

# “Naive” application of ML



# Scientific machine learning (SciML)

## Major problem

Despite big breakthroughs in science + AI


**Naively** using deep learning for scientific tasks usually leads to:

- Lack of interpretability
- Poor generalization
- Lots of training data required

Do neural networks really “**understand**” the scientific tasks they are being applied to?

A good scientific theory = makes novel predictions

## Solution



# Scientific machine learning (SciML)



Physics-informed neural networks: # citations  
(reproduced from Cuomo et al, 2022)

## National Science Foundation announces MIT-led Institute for Artificial Intelligence and Fundamental Interactions

IAIFI will advance physics knowledge — from the smallest building blocks of nature to the largest structures in the universe — and galvanize AI research innovation.

Laboratory for Nuclear Science  
August 26, 2020



The U.S. National Science Foundation (NSF) announced today an investment of more than \$100 million to establish five artificial intelligence (AI) institutes, each receiving roughly \$20 million over five years. One of these, the NSF AI Institute for Artificial Intelligence and Fundamental Interactions ([IAIFI](#)), will be led by MIT's Laboratory for Nuclear Science (LNS) and become the intellectual home of more than 25 physics and AI senior researchers at MIT and Harvard, Northeastern, and Tufts universities.


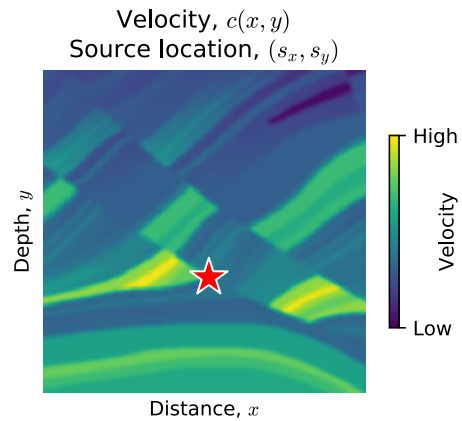
Do neural networks really “understand” the phenomena they are being applied to?

A good scientific theory = makes novel predictions

# Talk overview

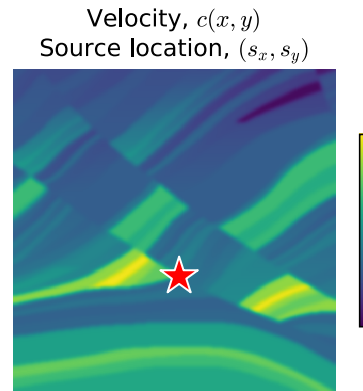
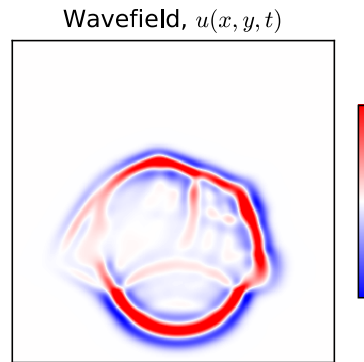
- What is scientific machine learning (SciML)?
- Ways to incorporate scientific principles into ML
- Our research: scaling SciML techniques to complex, real-world problems

# Typical SciML tasks




## Forward simulation

Estimate wavefield  $u(x, y, t)$   
Given velocity  $c(x, y)$   
and source location  $(s_x, s_y)$



## Inversion

Estimate velocity  $c(x, y)$   
and source location  $(s_x, s_y)$   
Given wavefield  $u(x, y, t)$



Wave equation


$$\nabla^2 u - \frac{1}{c^2} \frac{\partial^2 u}{\partial t^2} = f$$

## Equation discovery


Estimate governing equation  
Given wavefield  $u(x, y, t)$ ,  
velocity  $c(x, y)$ ,  
and source location  $(s_x, s_y)$

# Ways to incorporate scientific principles into machine learning


Loss function

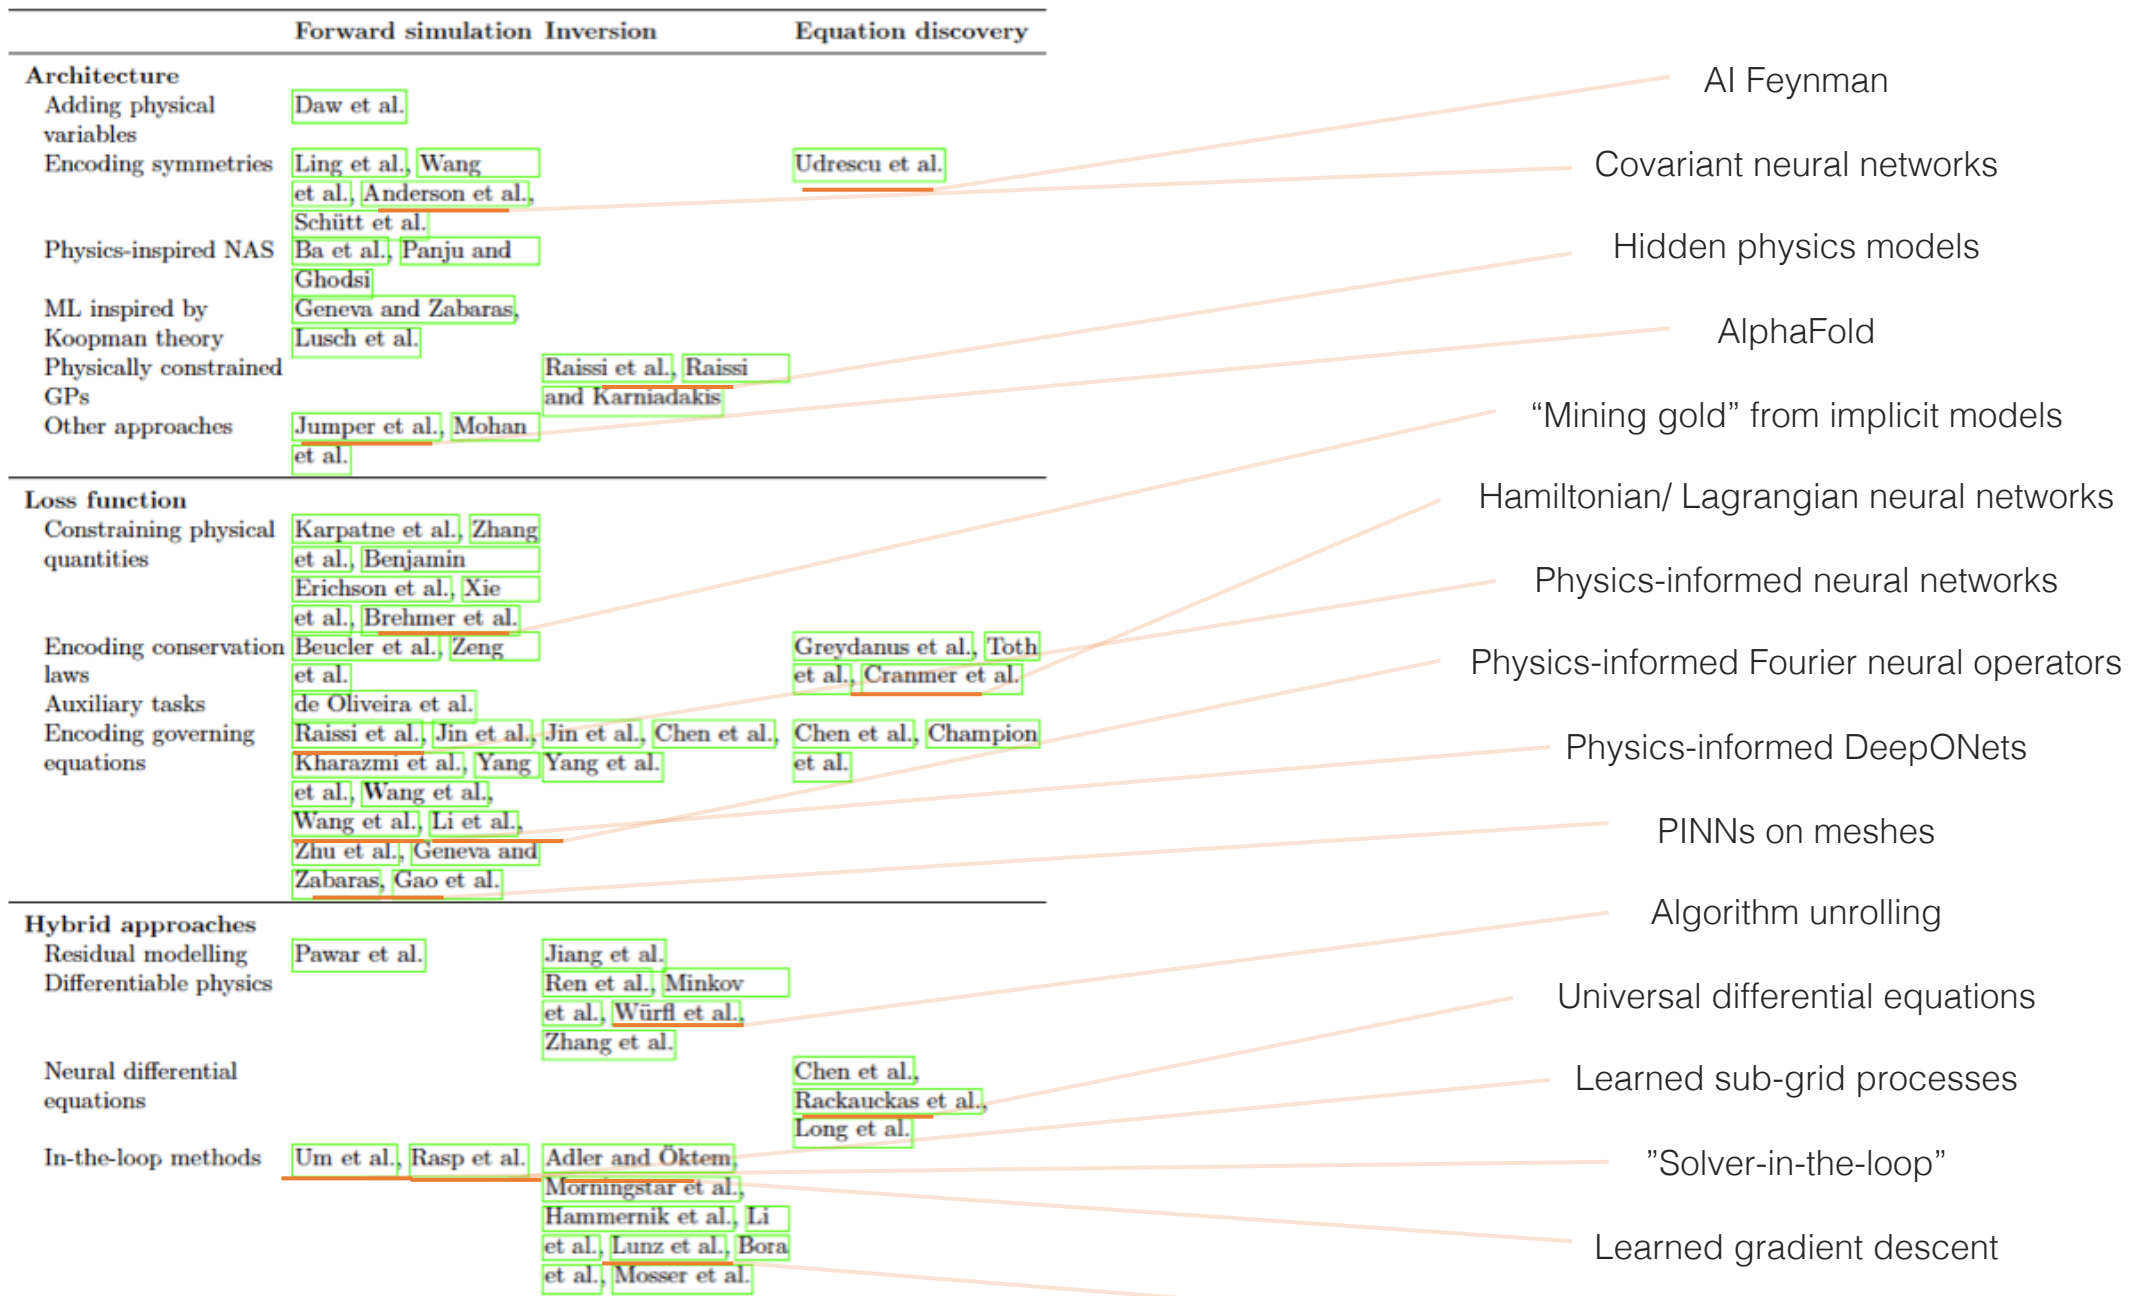


Architecture



Hybrid approaches

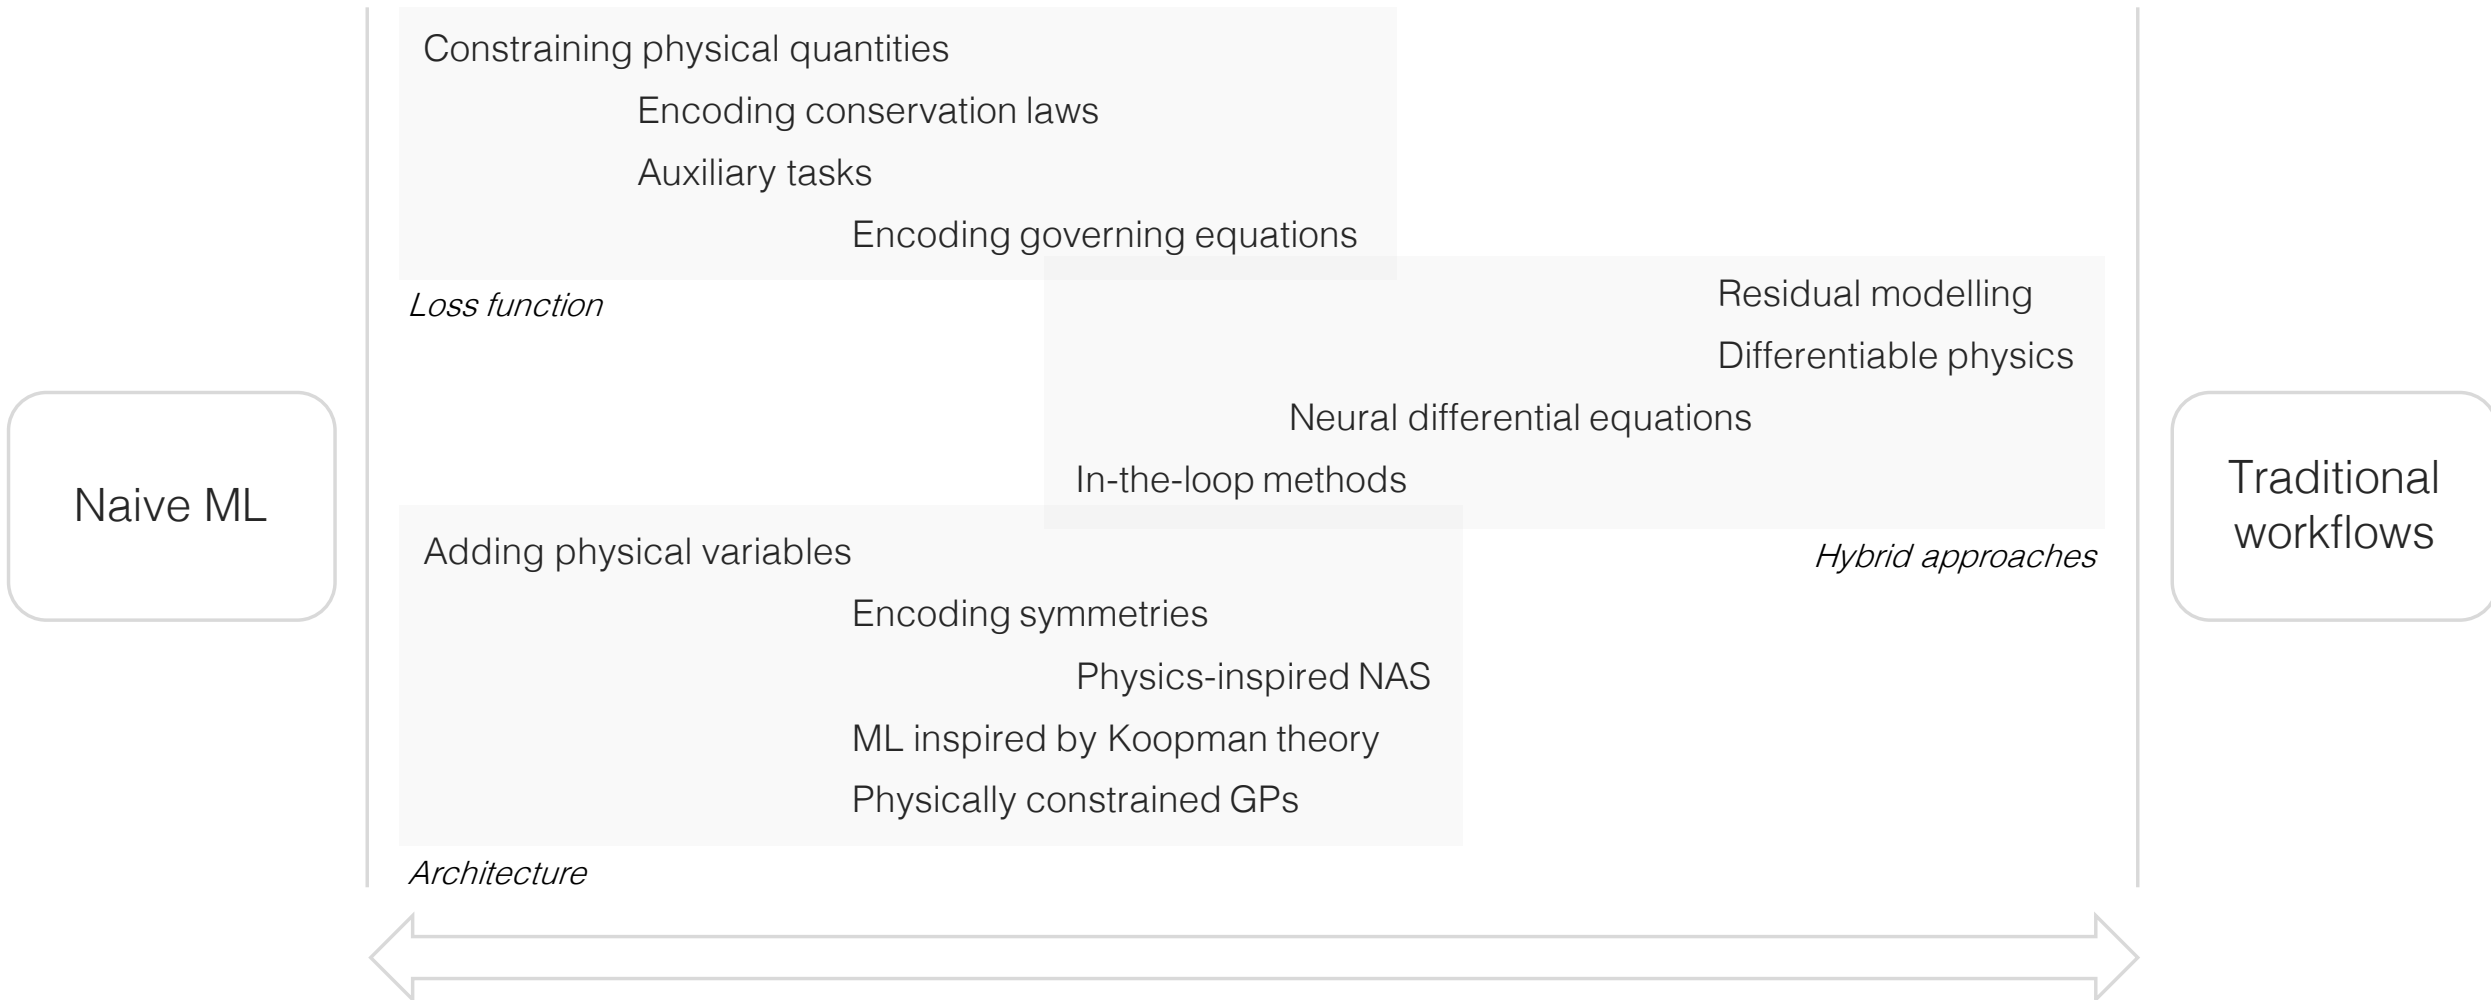




Source: My PhD thesis: Physics-informed machine learning: from concepts to real-world applications  
 read it here: [tinyurl.com/mw39wdps](https://tinyurl.com/mw39wdps)



# A wide plethora of techniques



Source: My PhD thesis: Physics-informed machine learning: from concepts to real-world applications

read it here: [tinyurl.com/mw39wdps](https://tinyurl.com/mw39wdps)

© Ben Moseley 2022


	Forward simulation	Inversion	Equation discovery
<b>Architecture</b>			
Adding physical variables	Daw et al.		
Encoding symmetries	Ling et al., Wang et al., Anderson et al., Schütt et al.		Udrescu et al.
Physics-inspired NAS	Ba et al., Panju and Ghodsi		
ML inspired by Koopman theory	Geneva and Zabaras, Lusch et al.		
Physically constrained GPs		Raissi et al., Raissi and Karniadakis	
Other approaches	Jumper et al., Mohan et al.		
<b>Loss function</b>			
Constraining physical quantities	Karpatne et al., Zhang et al., Benjamin Erichson et al., Xie et al., Brehmer et al.		
Encoding conservation laws	Beucler et al., Zeng et al.		Greydanus et al., Toth et al., Cranmer et al.
Auxiliary tasks	de Oliveira et al.		
Encoding governing equations	Raissi et al., Jin et al., Kharazmi et al., Yang et al., Wang et al., Wang et al., Li et al., Zhu et al., Geneva and Zabaras, Gao et al.	Jin et al., Chen et al., Yang et al.	Chen et al., Champion et al.
<b>Hybrid approaches</b>			
Residual modelling	Pawar et al.	Jiang et al.	
Differentiable physics		Ren et al., Minkov et al., Würfl et al., Zhang et al.	
Neural differential equations			Chen et al., Rackauckas et al., Long et al.
In-the-loop methods	Um et al., Rasp et al.	Adler and Öktem, Morningstar et al., Hammernik et al., Li et al., Lunz et al., Bora et al., Mosser et al.	

Source: My PhD thesis: Physics-informed machine learning: from concepts to real-world applications

read it here: [tinyurl.com/mw39wdps](https://tinyurl.com/mw39wdps)

# Loss function:

## Physics-informed neural networks




## “Naive” neural network




$$\min_{\theta} \frac{1}{N} \sum_i^N (u_{\text{NN}}(x_i; \theta) - \underline{u_{\text{true}}(x_i)})^2$$


## “Naive” neural network




$$\min_{\theta} \frac{1}{N} \sum_i^N (u_{\text{NN}}(x_i; \theta) - \underline{u_{\text{true}}(x_i)})^2$$




## “Naive” neural network



$$\min_{\theta} \frac{1}{N} \sum_i (u_{\text{NN}}(x_i; \theta) - \underline{u_{\text{true}}(x_i)})^2$$



## Physics-informed neural network




$$\min_{\theta} \frac{1}{N} \sum_i (u_{\text{NN}}(x_i; \theta) - \underline{u_{\text{true}}(x_i)})^2 + \frac{1}{M} \sum_j \left( \left[ m \frac{d^2}{dx^2} + \mu \frac{d}{dx} + k \right] u_{\text{NN}}(\underline{x_j}; \theta) \right)^2$$

Raissi et al (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations.



Lagaris et al (1998). Artificial neural networks for solving ordinary and partial differential equations.

# PINNs for simulation

Ground truth FD



Velocity model



Boundary loss

$$\frac{1}{N_b} \sum_{i=1}^{N_b} \|u_{\text{FD}}(x_i, t_i) - NN(x_i, t_i; \theta)\|^2$$

Physics loss

$$\frac{\lambda}{N_p} \sum_{j=1}^{N_p} \left\| \left[ \nabla^2 - \frac{1}{c(x_j)^2} \frac{\partial^2}{\partial t^2} \right] NN(x_j, t_j; \theta) \right\|$$

$$\mathcal{L}(\theta) = \frac{1}{N_b} \sum_{i=1}^{N_b} \|u_{\text{FD}}(x_i, t_i) - NN(x_i, t_i; \theta)\|^2$$

$$+ \frac{\lambda}{N_p} \sum_{j=1}^{N_p} \left\| \left[ \nabla^2 - \frac{1}{c(x_j)^2} \frac{\partial^2}{\partial t^2} \right] NN(x_j, t_j; \theta) \right\|$$

# PINNs for simulation


Ground truth FD

“Naïve” NN


PINN

Difference (NN)

Difference (PINN)



Velocity model




NN architecture:  
Fully-connected, 10 layers, 512  
hidden channels

Training time: ~3 hours




# PINNs for inversion

Shukla K et al, Physics-Informed Neural Network for Ultrasound Nondestructive Quantification of Surface Breaking Cracks, Journal of Nondestructive Evaluation (2020)




(a) Actual data at  $t = 12.38 \mu s$ .




(b) Data recovered from PINN simulation at  $t = 12.38 \mu s$ .

$$\mathcal{L}(\theta) = \frac{1}{N_b} \sum_{i=1}^{N_b} \|u_{FD}(x_i, t_i) - NN(x_i, t_i; \theta)\|^2 + \frac{\lambda}{N_p} \sum_{j=1}^{N_p} \left\| \left[ \nabla^2 - \frac{1}{c(x_j)^2} \frac{\partial^2}{\partial t^2} \right] NN(x_j, t_j; \theta) \right\|$$



- Simultaneously learn velocity model too!



(d) Speed  $v(x, y)$  recovered from PINN simulation.

# Try PINNs for yourself!

benmoseley / harmonic-oscillator-pinn Public

Notifications Fork 47 Star 132

<> Code Issues 1 Pull requests Actions Projects Security Insights

main 2 branches 0 tags Go to file Code

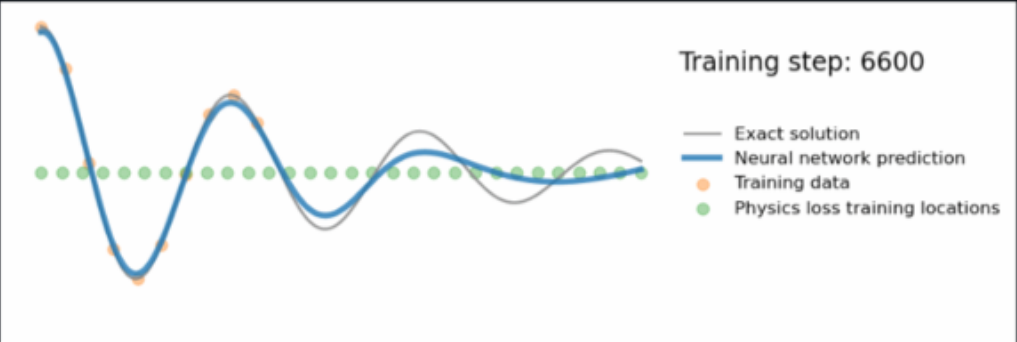
benmoseley Update README.md 1 f00c5d9 on 28 Aug 2021 3 commits

figures	update	13 months ago
Harmonic oscillator PINN.ipynb	update	13 months ago
LICENSE	Initial commit	13 months ago
README.md	Update README.md	13 months ago

README.md

## harmonic-oscillator-pinn

Code accompanying my blog post: [So, what is a physics-informed neural network?](#)



Training step: 6600

- Exact solution
- Neural network prediction
- Training data
- Physics loss training locations

About

Code accompanying my blog post: [So, what is a physics-informed neural network?](#)

Readme MIT license 132 stars 3 watching 47 forks

Releases

No releases published

Packages

No packages published

Languages


Jupyter Notebook 100.0%

[github.com/benmoseley](https://github.com/benmoseley)

© Ben Moseley 2022


# Architecture:

ML inspired by Koopman theory

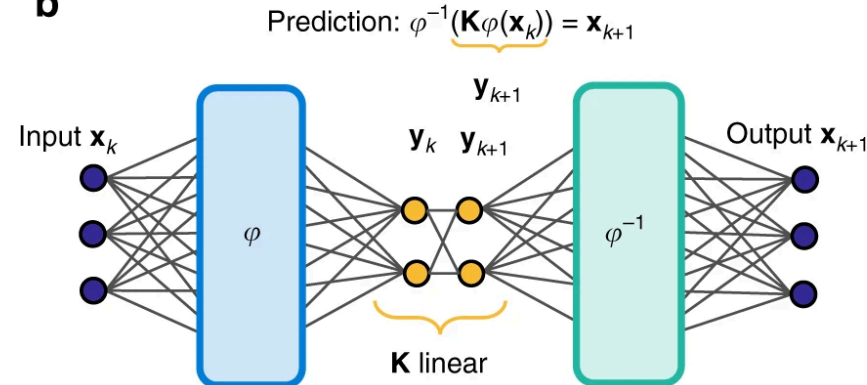


Roughly: Koopman theory states that any (potentially nonlinear) dynamical system can be represented in terms of an infinite-dimensional linear operator

**a**

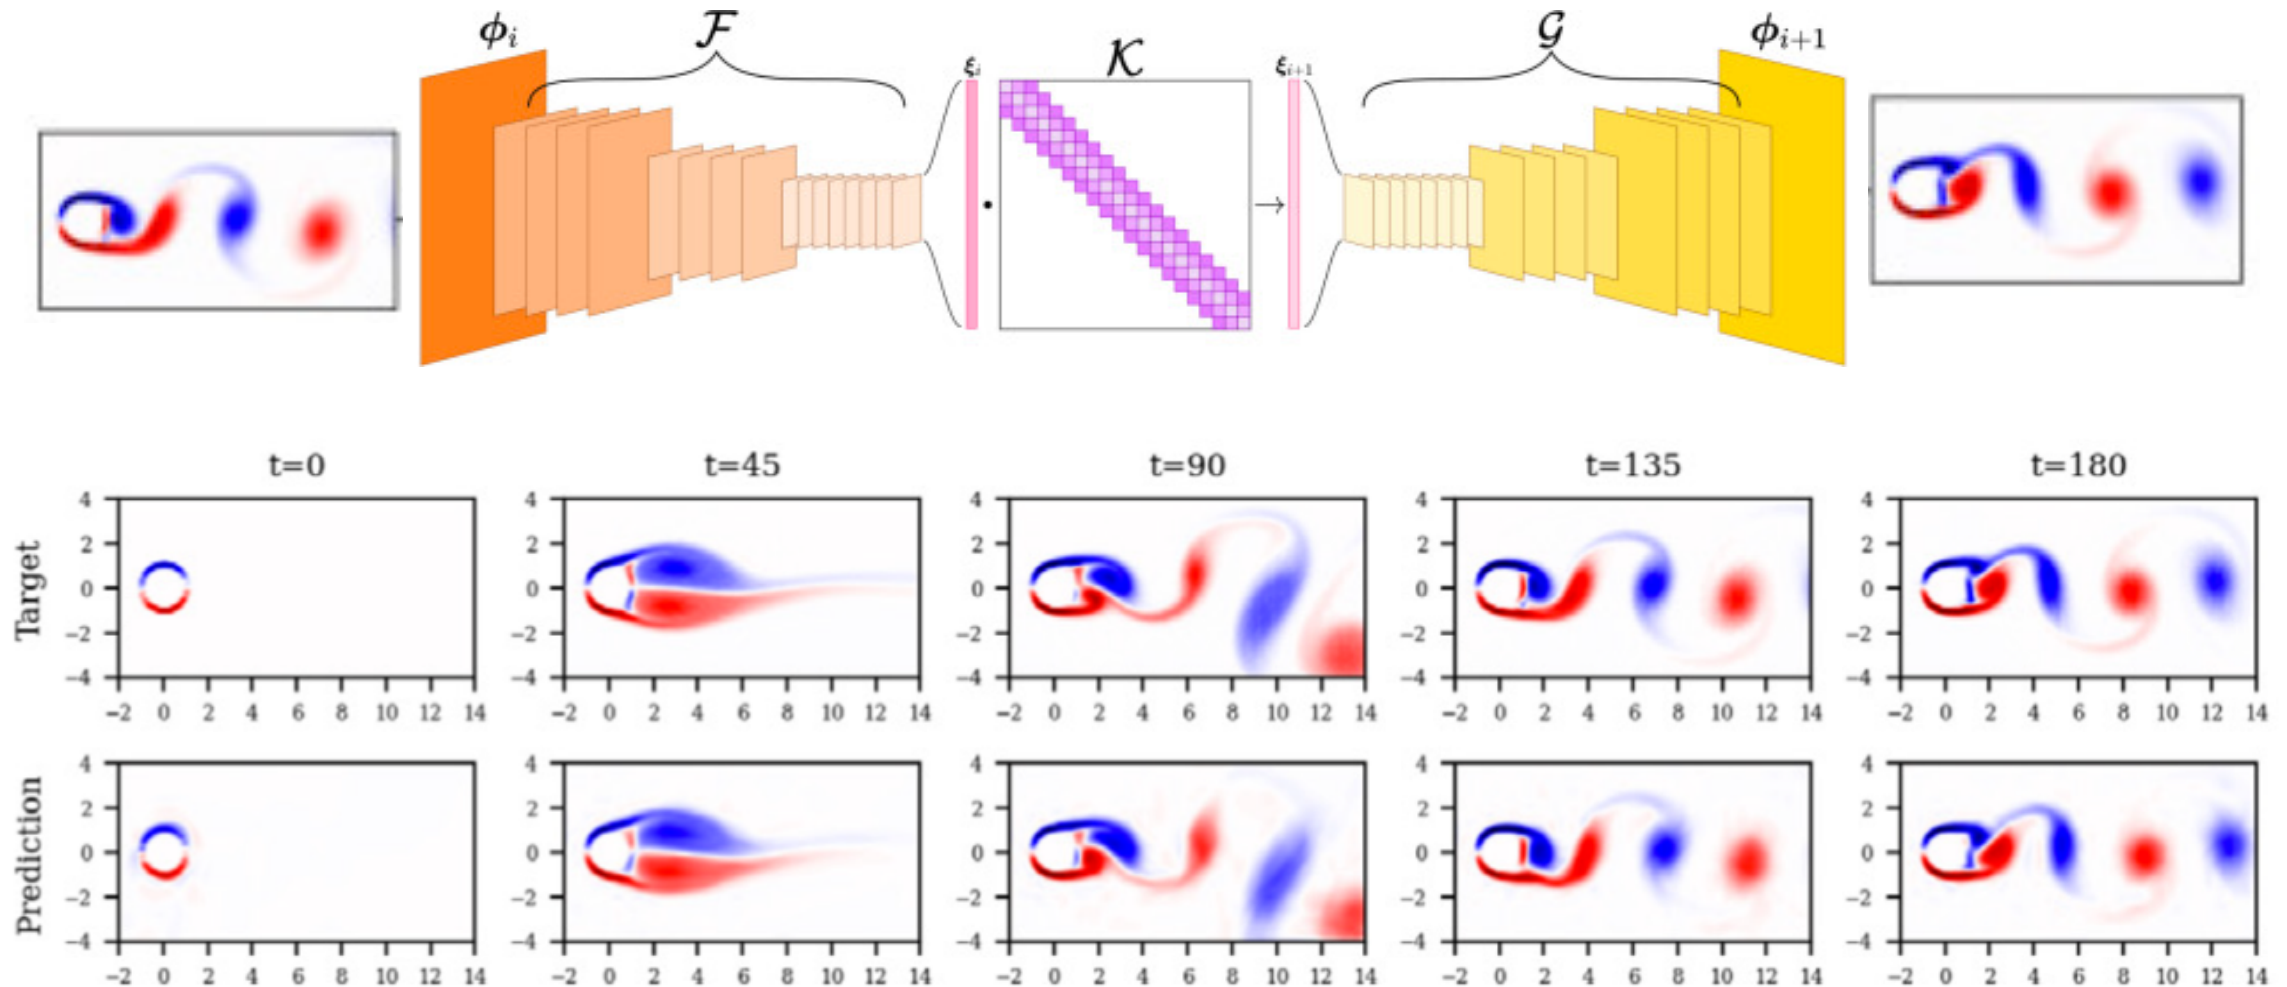


**b**




Loss = reconstruction loss + prediction loss

Roughly: Koopman theory states that any (potentially nonlinear) dynamical system can be represented in terms of an infinite-dimensional linear operator



# Hybrid approaches: Differentiable physics




# Differentiable physics

```
In [1]: import torch

x = torch.arange(0, 2.1, 0.1, requires_grad=True)
y = x**2

y.backward(torch.ones_like(x))


import matplotlib.pyplot as plt
plt.plot(x.detach(), y.detach(), label="$y=x^2$")
plt.plot(x.detach(), x.grad.detach(), label="$\frac{dy}{dx}$")
plt.legend()
plt.show()
```




Key ideas:

1. Auto-differentiation is very powerful!!
2. We can write (nearly) any traditional scientific algorithm in a differentiable programming language
3. And insert and train ML components / flexible parameters anywhere inside them

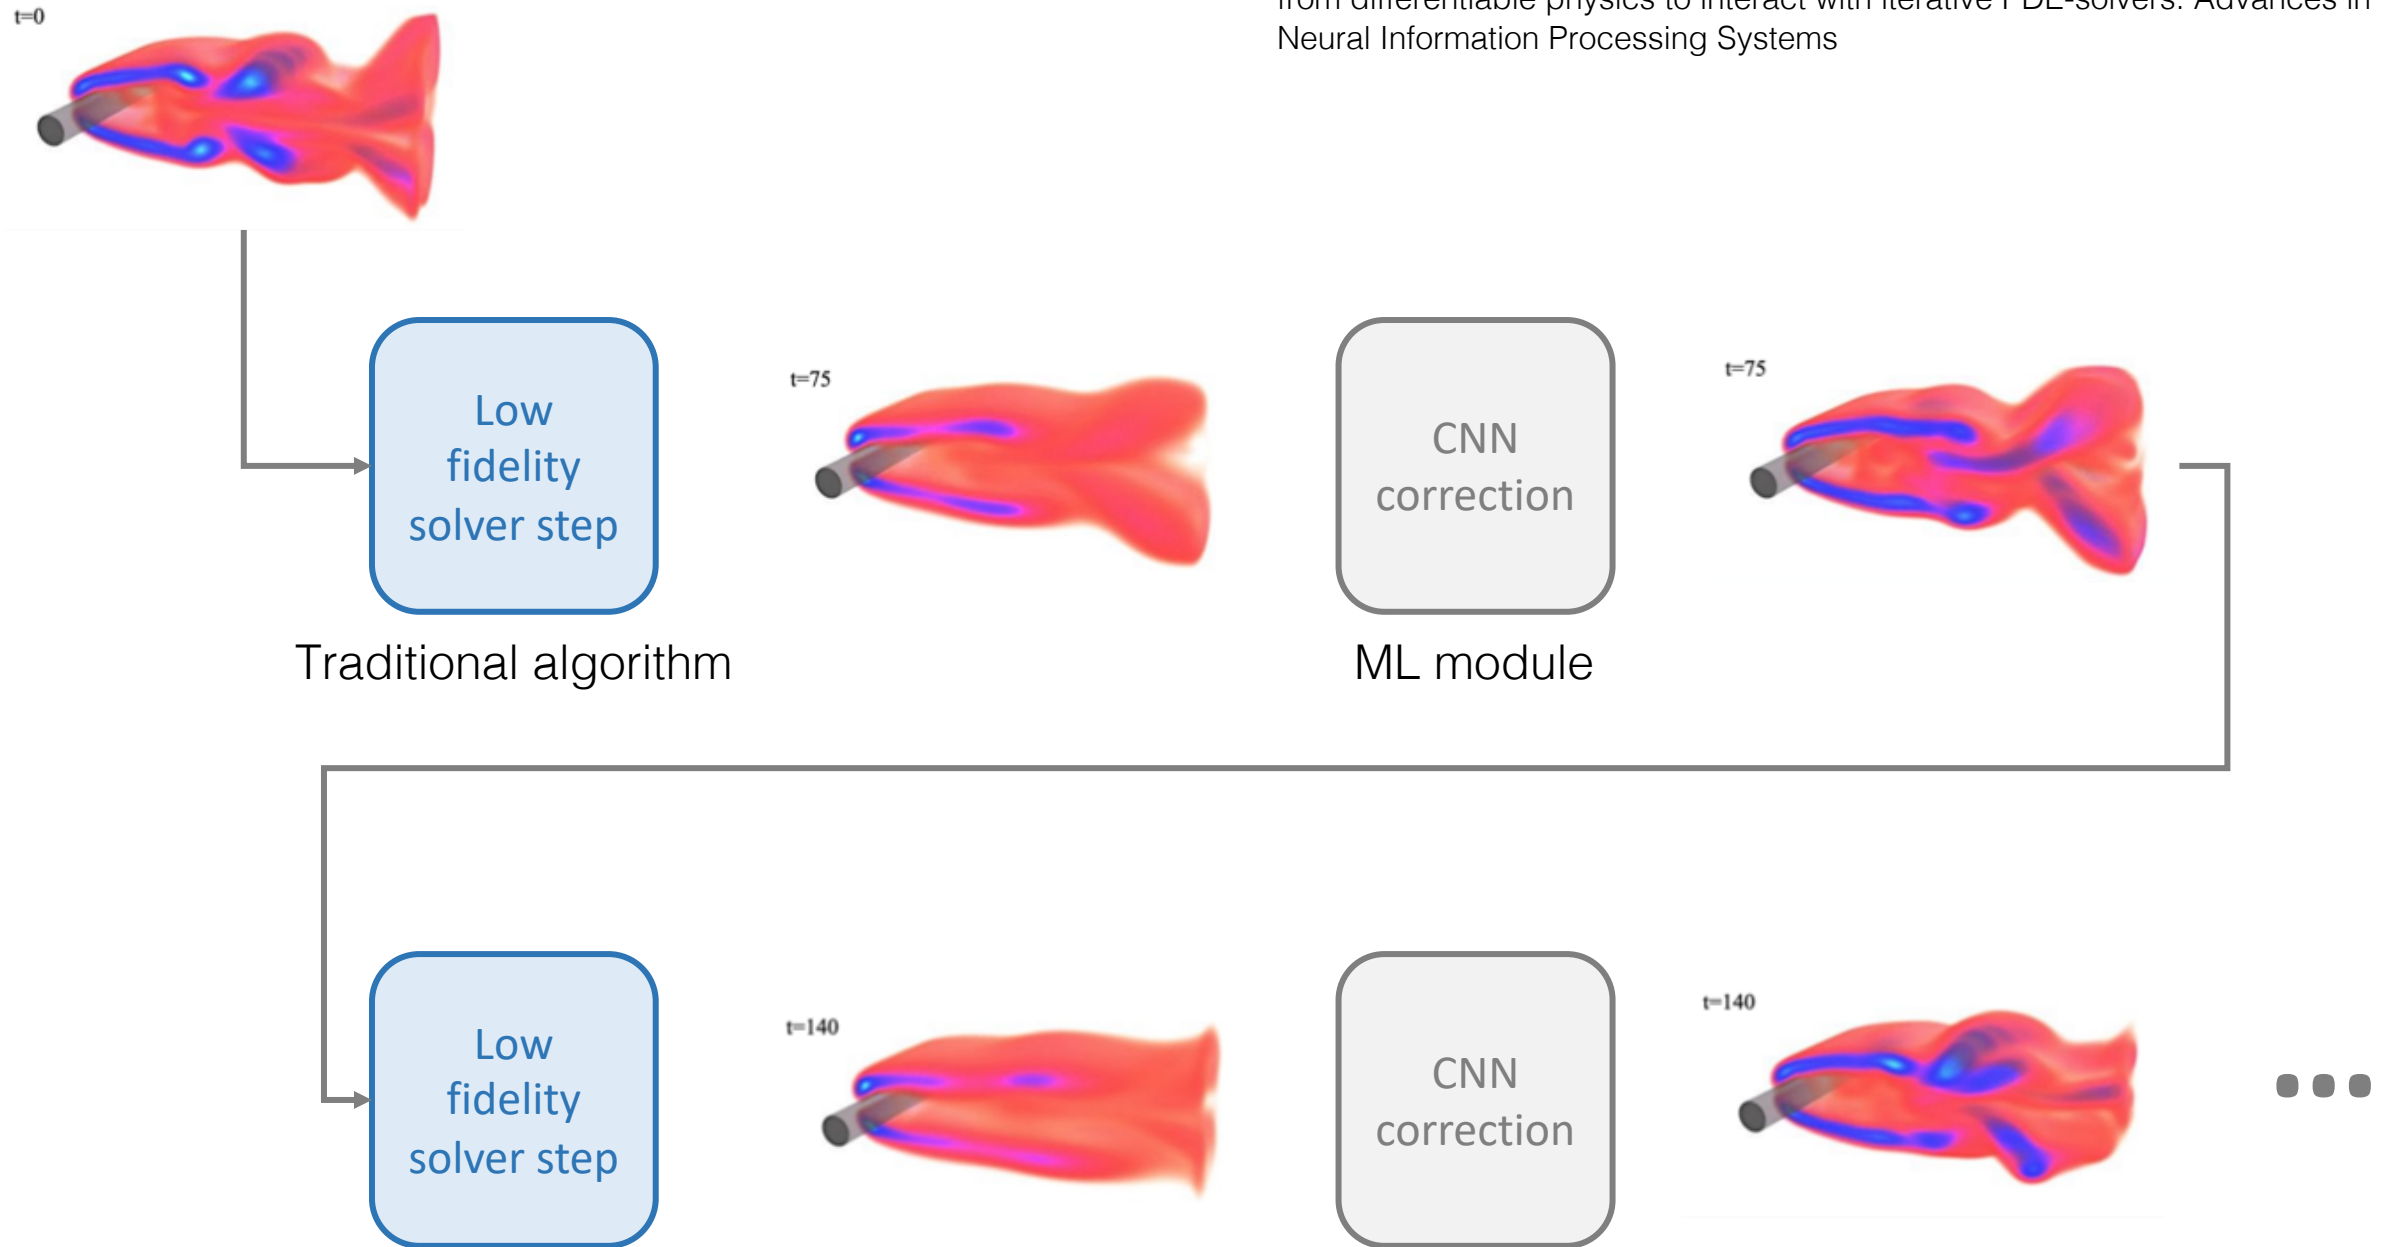
Low fidelity solver

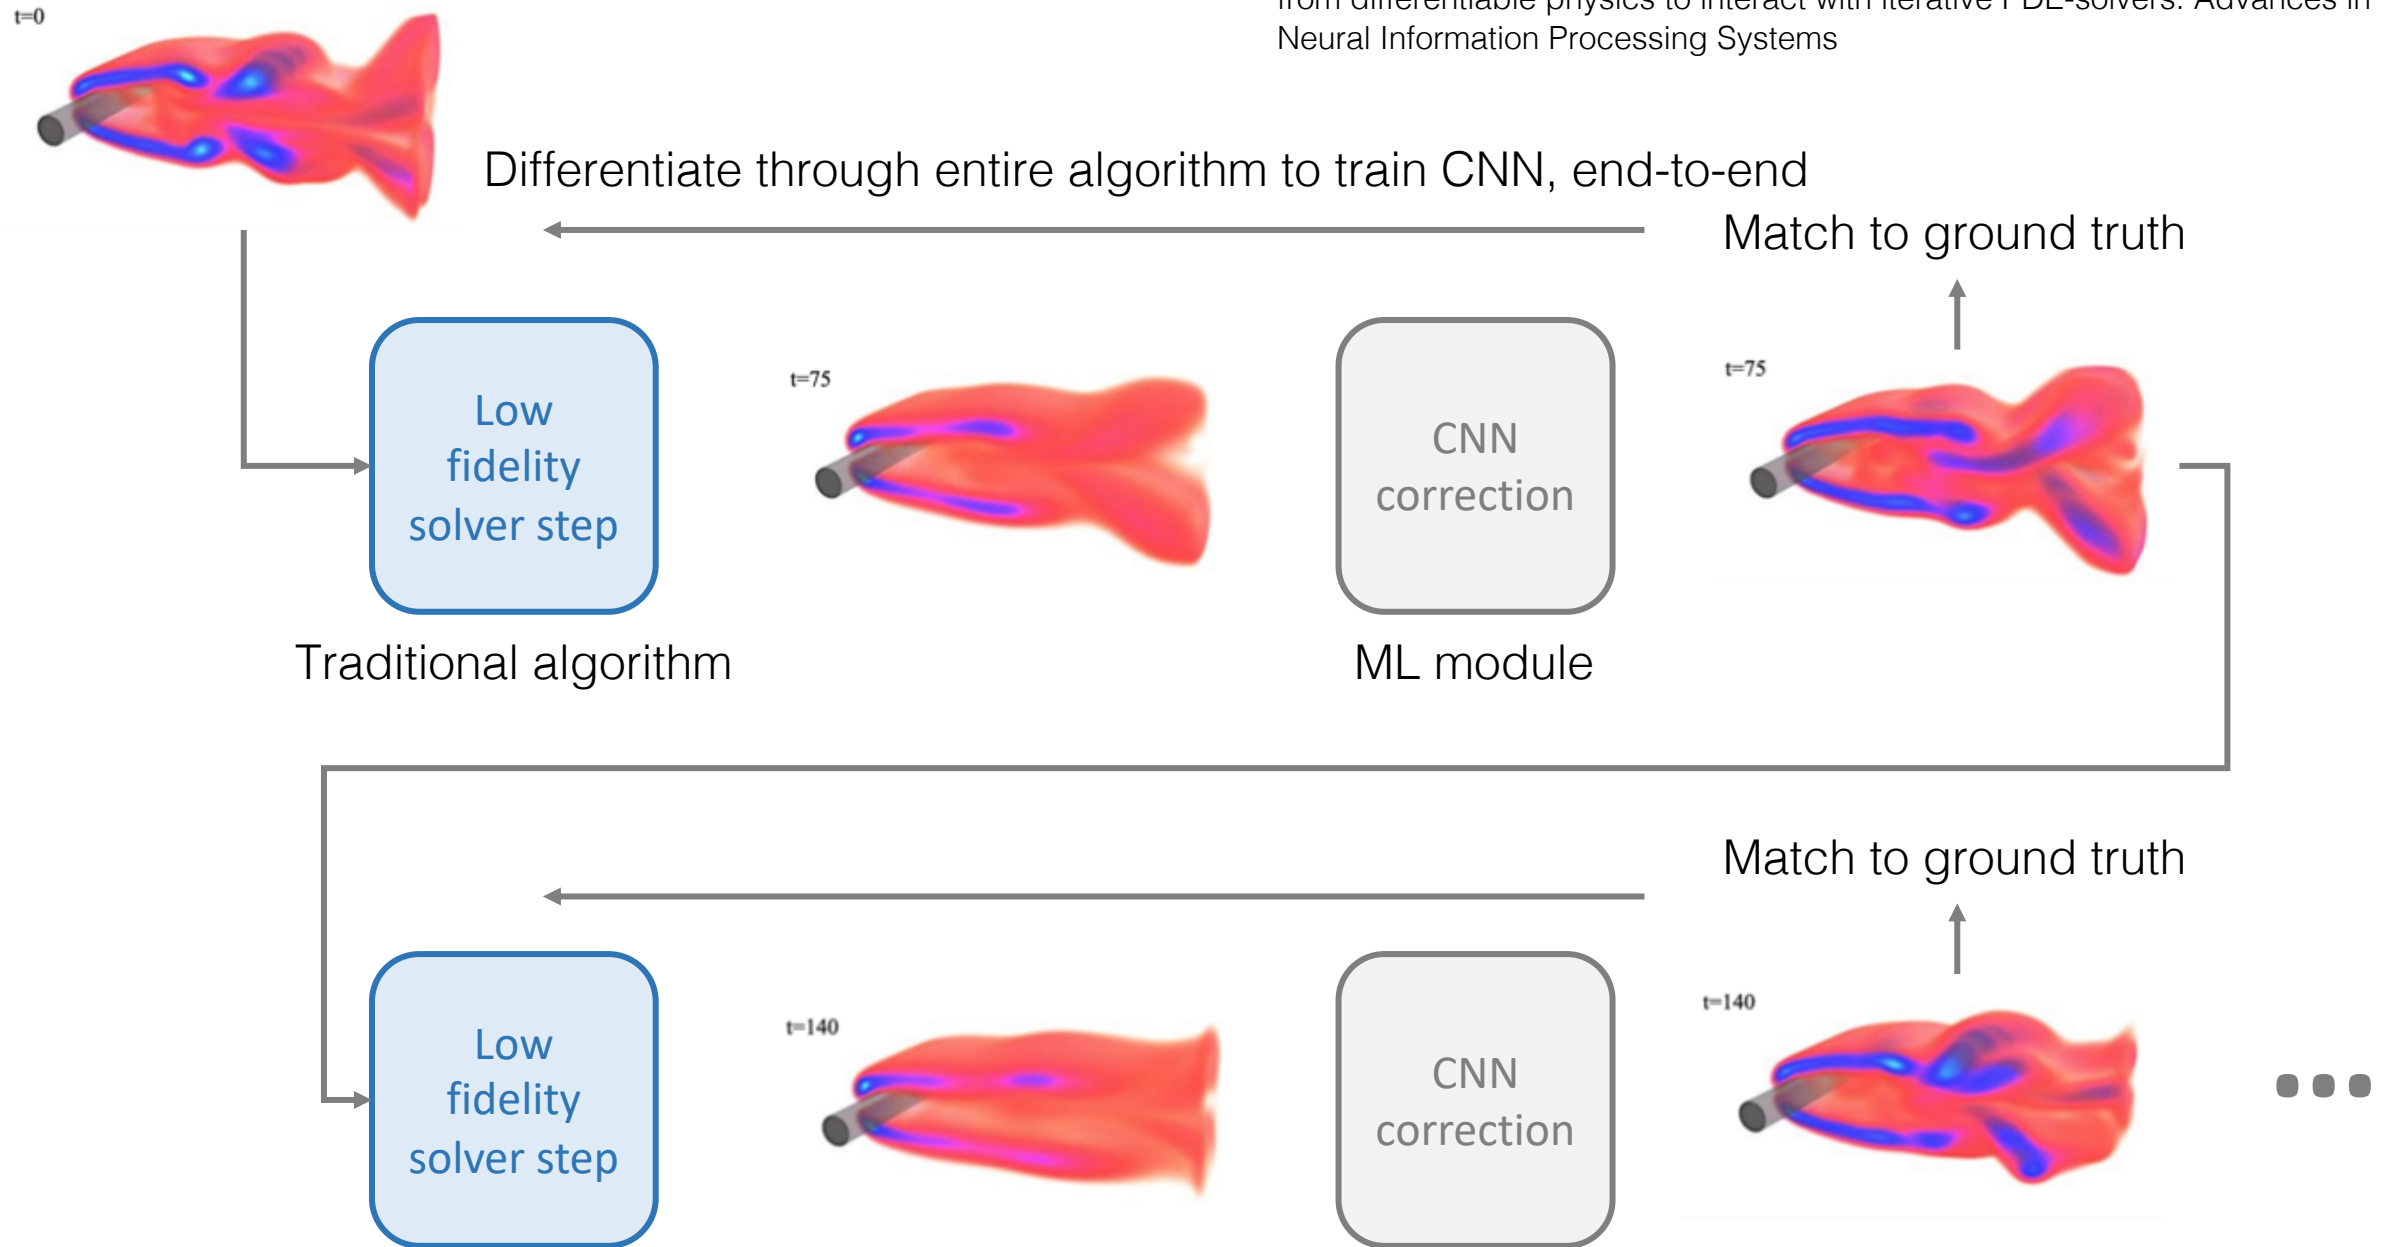


High fidelity solver  
“Ground truth”












Low fidelity solver



Solver-in-the-loop



High fidelity solver  
"Ground truth"







# Talk overview

- What is scientific machine learning (SciML)?
- Ways to incorporate scientific principles into ML
- Our research: scaling SciML techniques to complex, real-world problems

# Our research:

Many SciML works focus on solving **toy/simplified problems** to validate their techniques

How well do SciML techniques scale to complex, real-world problems?




**Traditional scientific methods**  
struggle to scale when:

Adding more complex phenomena  
(multi-scale, multi-physics)


Increasing the domain size /  
adding higher frequencies (high  
computational cost)

Incorporating real, noisy and  
sparse data


# Scaling to large problems: The spectral bias issue




(a) Iteration 100



(b) Iteration 1000




(c) Iteration 10000



(d) Iteration 80000

- NNs prioritise learning lower frequency functions first
- Under certain assumptions can be proved via neural tangent kernel theory

Rahaman, N., et al, On the spectral bias of neural networks. 36th International Conference on Machine Learning, ICML (2019)







# A motivating problem

321 free parameters


(a) PINN ( $\omega = 1$ , 2 layers, 16 hidden units)




(b) PINN ( $\omega = 15$ , 2 layers, 16 hidden units)



(c) PINN ( $\omega = 15$ , 4 layers, 64 hidden units)




(d) PINN ( $\omega = 15$ , 5 layers, 128 hidden units)



$$\frac{du}{dx} = \cos(\omega x) ,$$
$$u(0) = 0 ,$$

66,433 free parameters


(e) Test loss



# Key scaling issues


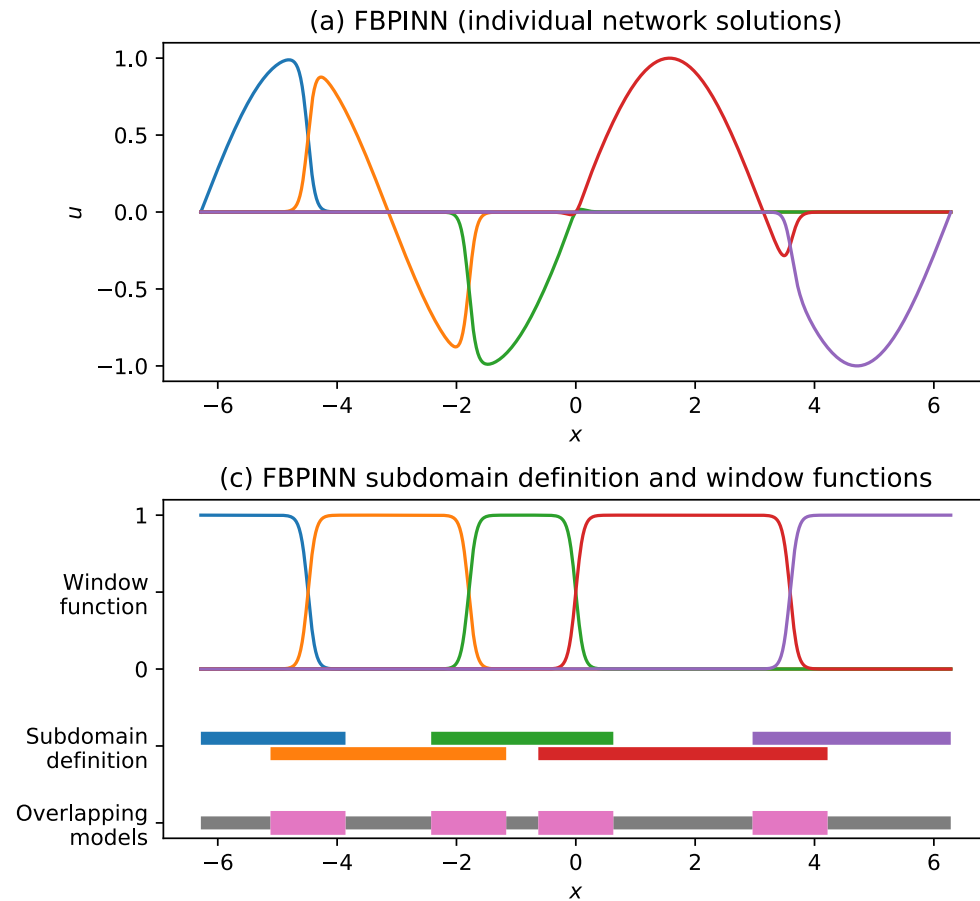
Multiple related issues when scaling PINNs to large domains:

- As the domain increases:
  - The complexity of the solution increases;
  - Requiring a larger neural network (more free parameters);
  - And more training points to sample the domain;
  - -> Leading to a harder optimisation problem.
- As the frequency increases;
  - The neural network takes longer to converge (spectral bias).
- As the size of the network, number of training points, and convergence time grows, the computational resources grows significantly





# Finite basis physics-informed neural networks (FBPINNs)



- Use **domain decomposition**, **individual subdomain normalisation** and **flexible training schedules** to allow PINNs to scale to large domains

$$\hat{u}(x; \theta) = \mathcal{C} \left[ \sum_i^n \underbrace{w_i(x)}_{\text{Window function}} \cdot \underbrace{\text{unnorm} \circ NN_i \circ \text{norm}_i(x)}_{\text{Subdomain network}} \right]$$

Window function

Subdomain network

Separate subdomain normalisation


$$\frac{du}{dx} = \cos(\omega x), \quad u(x) = \frac{1}{\omega} \sin(\omega x)$$

$$u(0) = 0,$$


$$\hat{u}(x; \theta) = \tanh(\omega x) \overline{NN}(x; \theta)$$

$$\omega = 15$$


(a) FBPINN (individual network solutions)




(b) FBPINN (full solution)




(c) FBPINN subdomain definition and window functions



(d) Test loss



(e) Test loss



- Subdomain network size: 2 layers, 16 hidden units



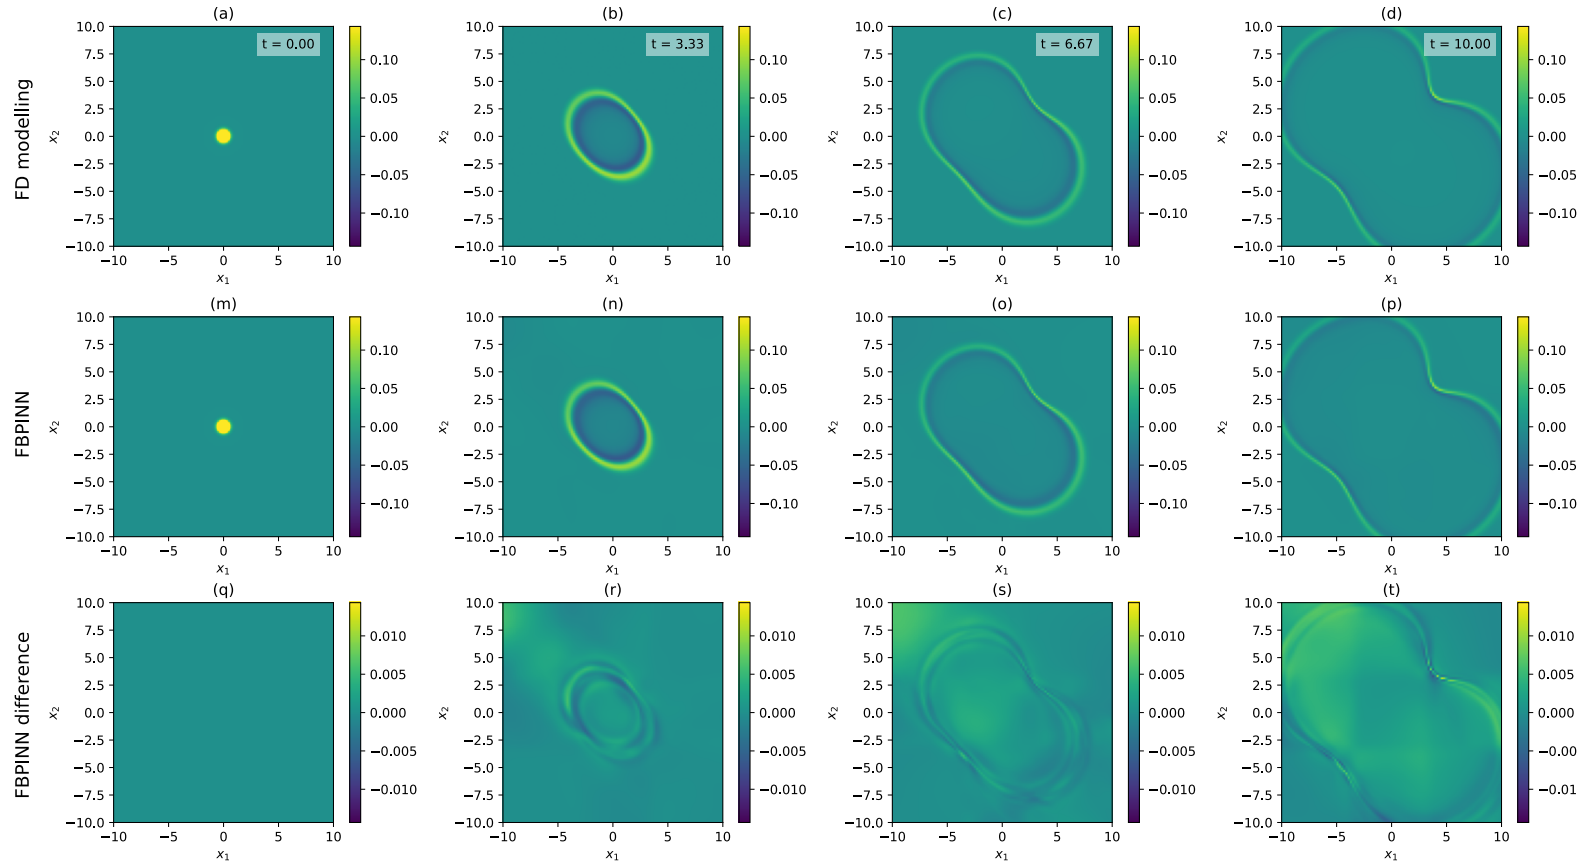
- (2+1)D time-dependent wave equation

$$\left[ \nabla^2 - \frac{1}{c(x)^2} \frac{\partial^2}{\partial t^2} \right] u(x, t) = 0 ,$$

$$u(x, 0) = e^{-\frac{1}{2}(\|x-\mu\|/\sigma)^2} ,$$

$$\frac{\partial u}{\partial t}(x, 0) = 0 ,$$


$$\hat{u}(x, t; \theta) = \phi(5(2 - t/t_1)) e^{-\frac{1}{2}(\|x-\mu\|/\sigma)^2} + \tanh^2(t/t_1) \overline{NN}(x, t; \theta)$$



[github.com/benmoseley/FBPINNs](https://github.com/benmoseley/FBPINNs)


# Peering into shadows on the Moon with SciML

Shackleton crater, lunar south pole




21 km

Lunar south pole




200 km



Brightness


1 km

Wapowski crater




Credits: LROC / GSFC / ASU





Credits: LROC / GSFC / ASU

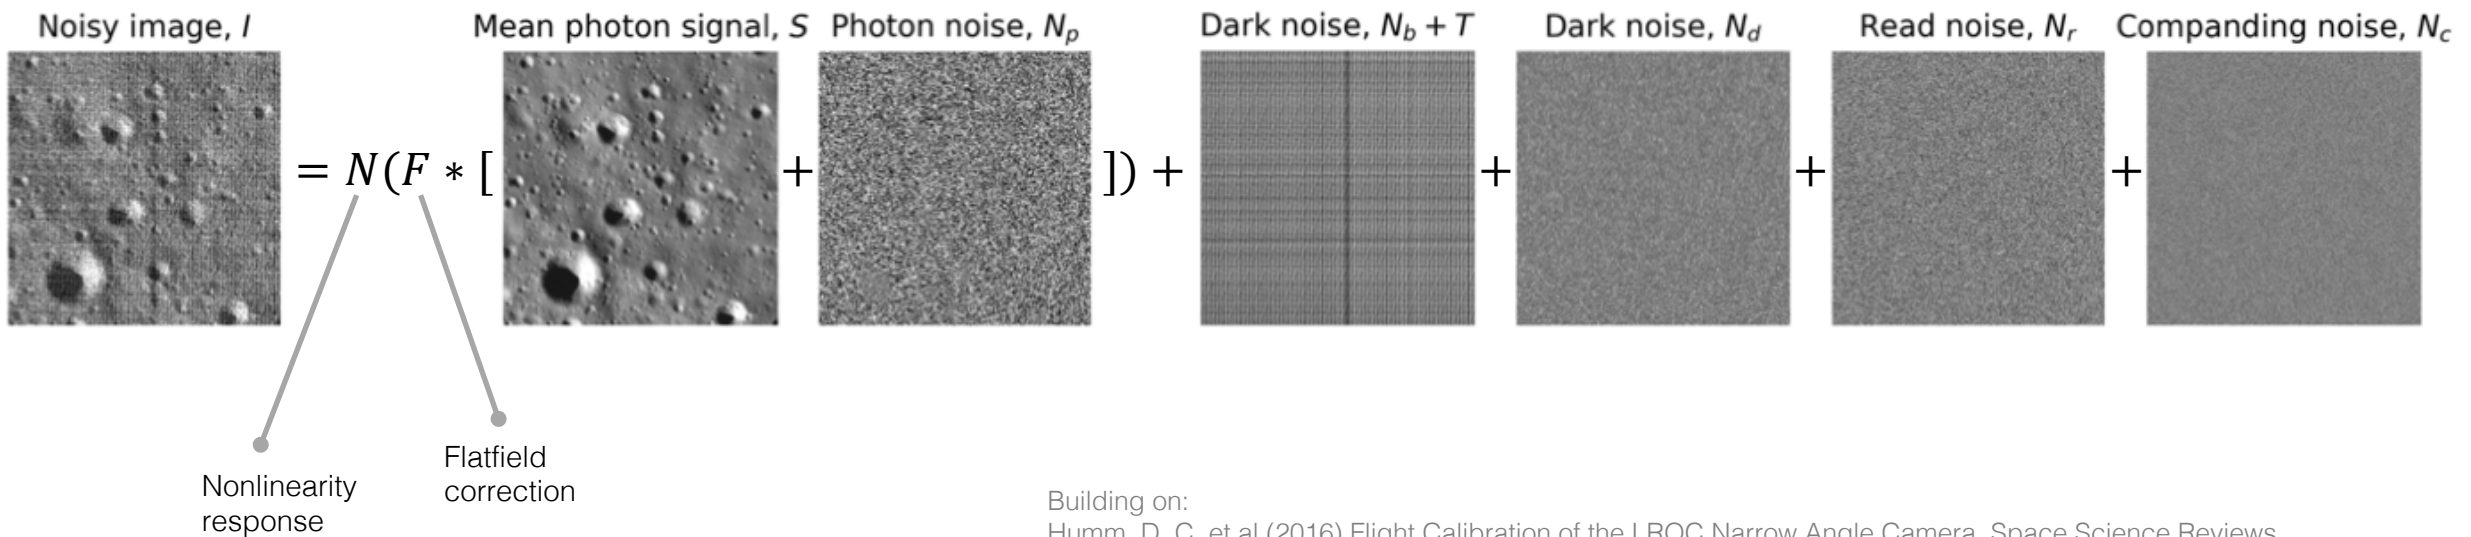
Lunar Reconnaissance Orbiter Camera



Smartphone



## Physical noise model



Building on:

Humm, D. C. et al (2016) Flight Calibration of the LROC Narrow Angle Camera. Space Science Reviews

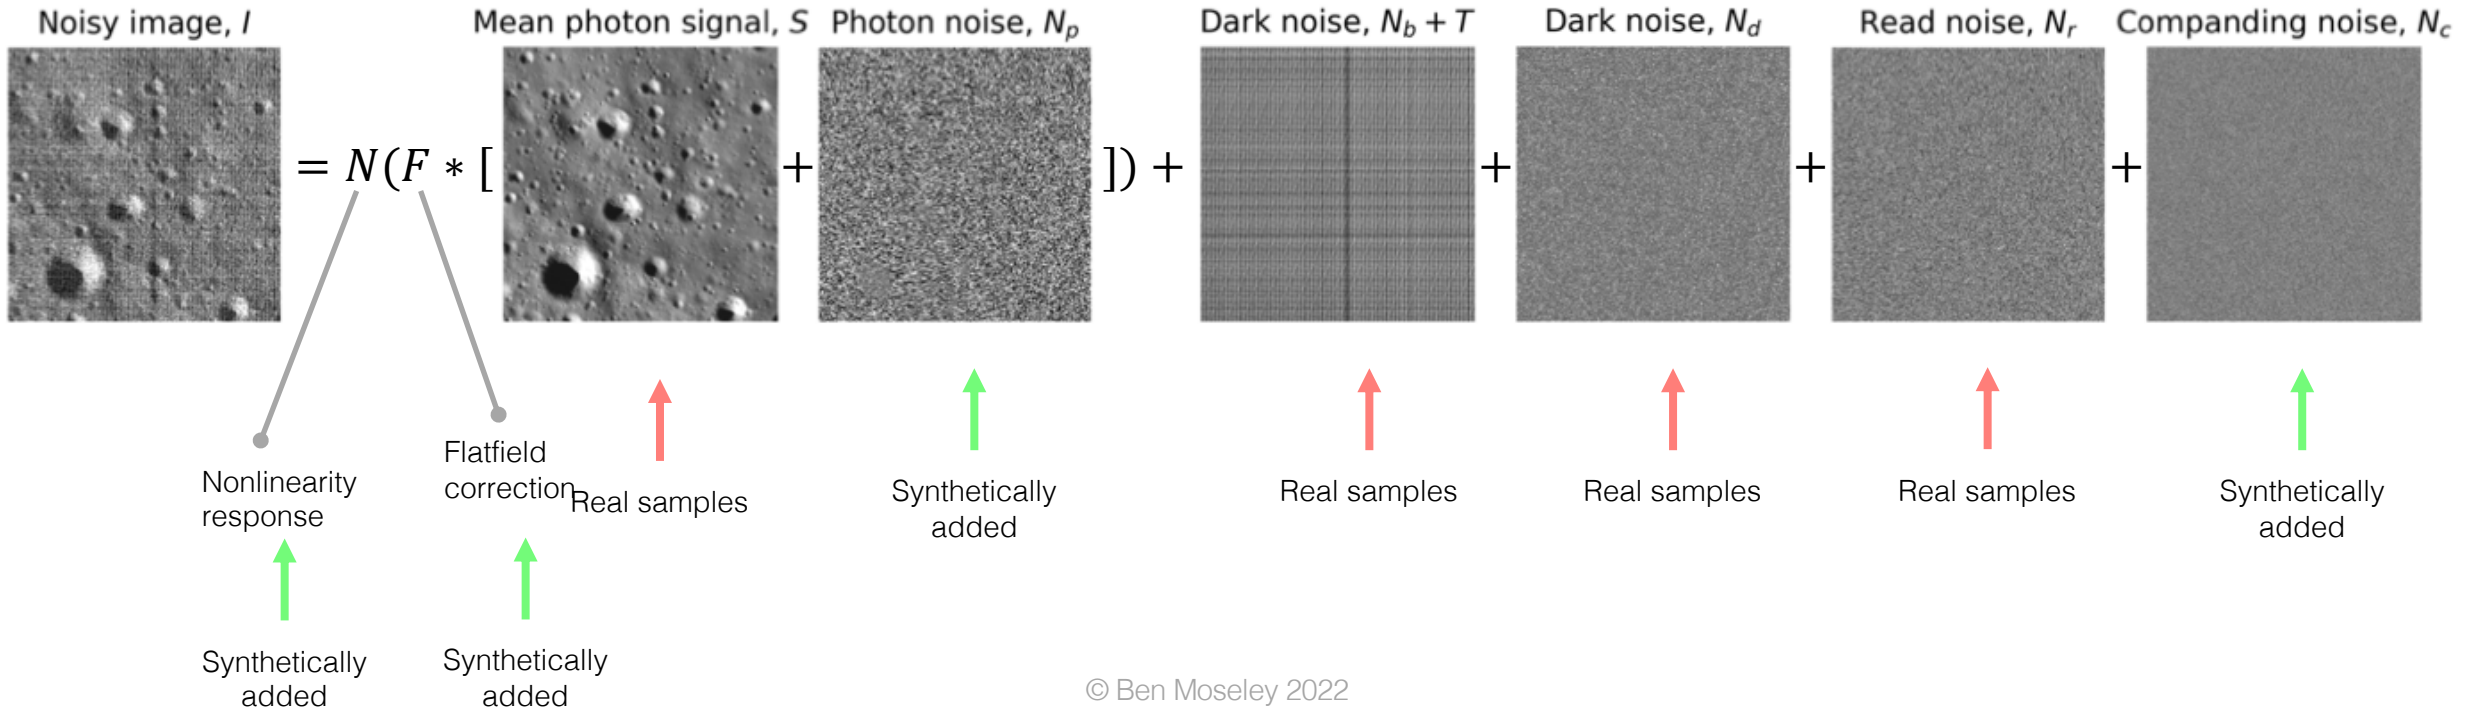

Robinson, M. S. et al (2010) Lunar reconnaissance orbiter camera (LROC) instrument overview. Space Science Reviews

© Ben Moseley 2022



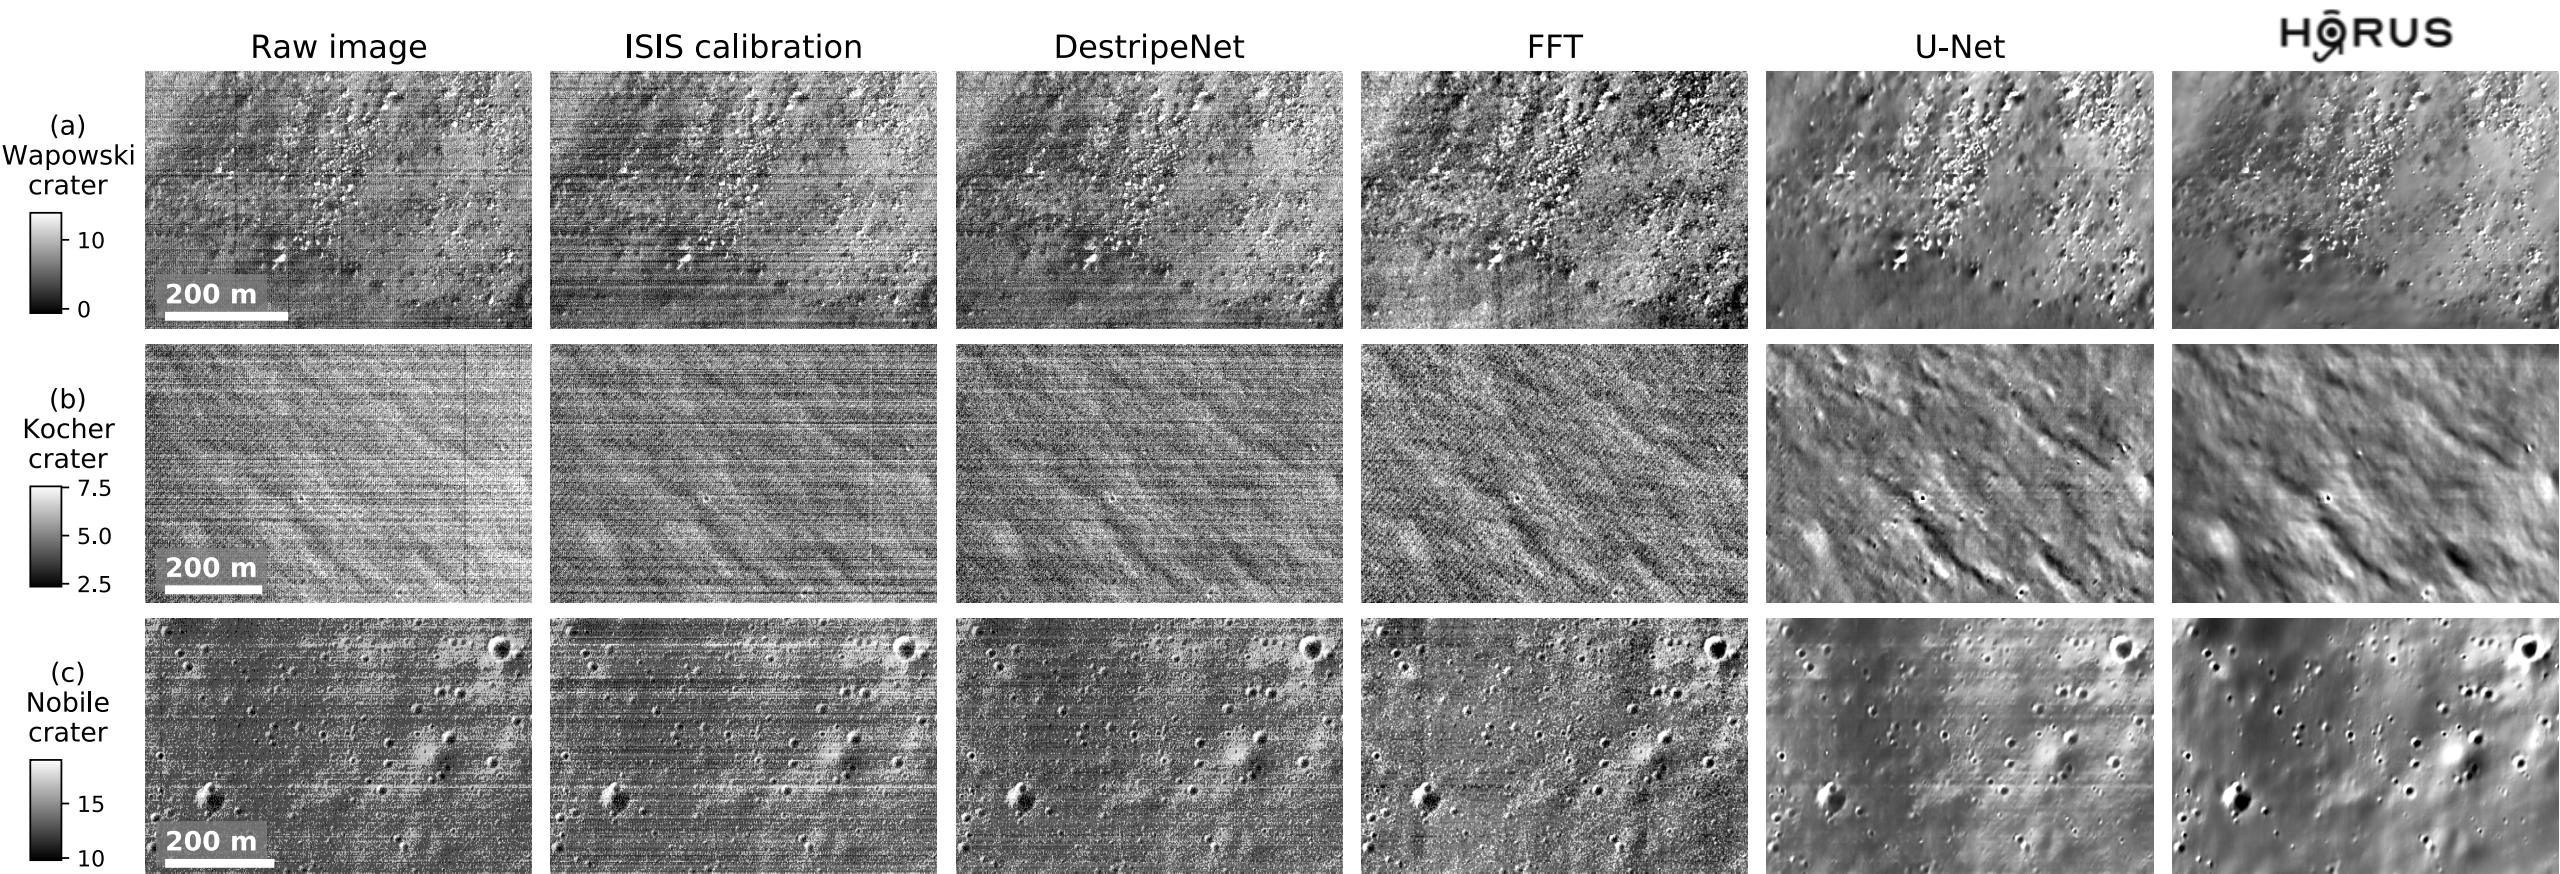
# HORUS

## HYPER-EFFECTIVE NOISE REMOVAL UNET SOFTWARE







# Comparison to other methods





Peering into lunar permanently shadowed regions with deep learning



VIPER Lunar rover  
Credits: NASA Ames/Daniel Rutter

# Takeaways

- SciML is a blossoming field of research
- There are a plethora of different SciML techniques...
- ...which range in the way scientific constraints are added, and their intended scientific task
- Scaling SciML techniques to more complex, multi-scale, multi-physics problems remains an exciting field of research!
- Check out my blog/ GitHub for more!

[benmoseley.blog](https://benmoseley.blog) [github/benmoseley](https://github.com/benmoseley)