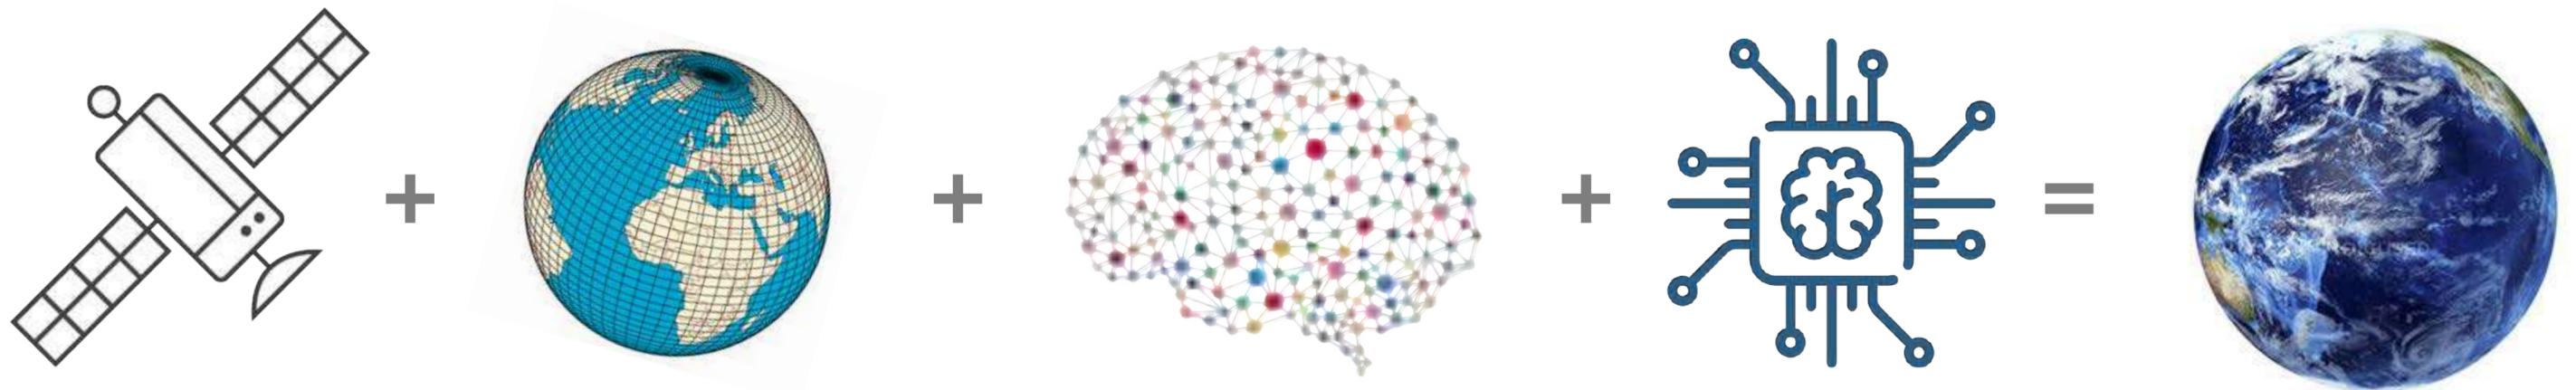


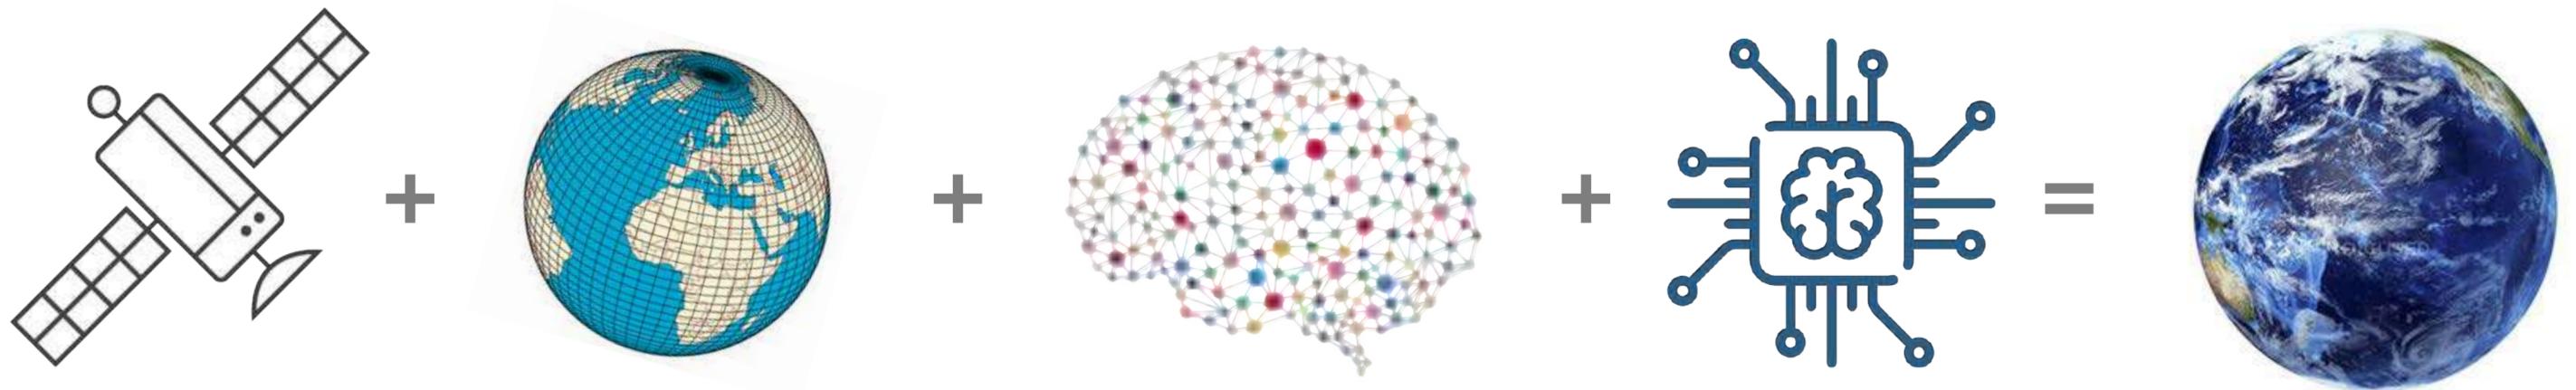
4.

# The plumbing challenges of **hybrid** modelling





# The plumbing challenges of **hybrid** modelling

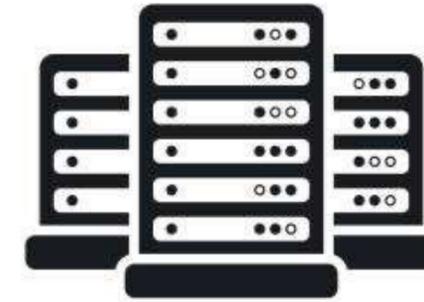


# Interfacing ocean models with DL frameworks (1/3)

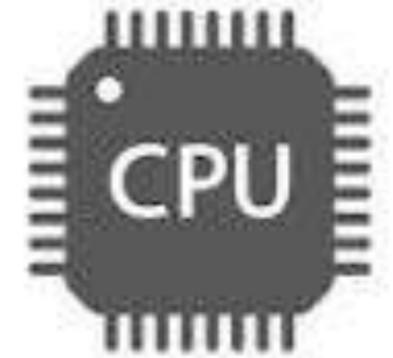


stable, robust, low abstraction languages

+



supercomputers



runs only on CPUs

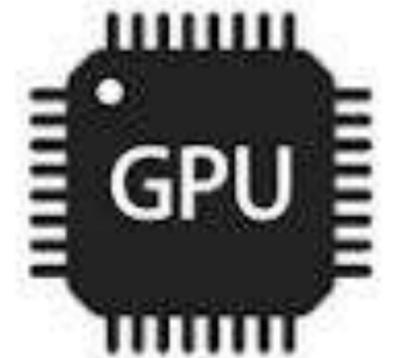


high abstraction, fast evolving languages

+

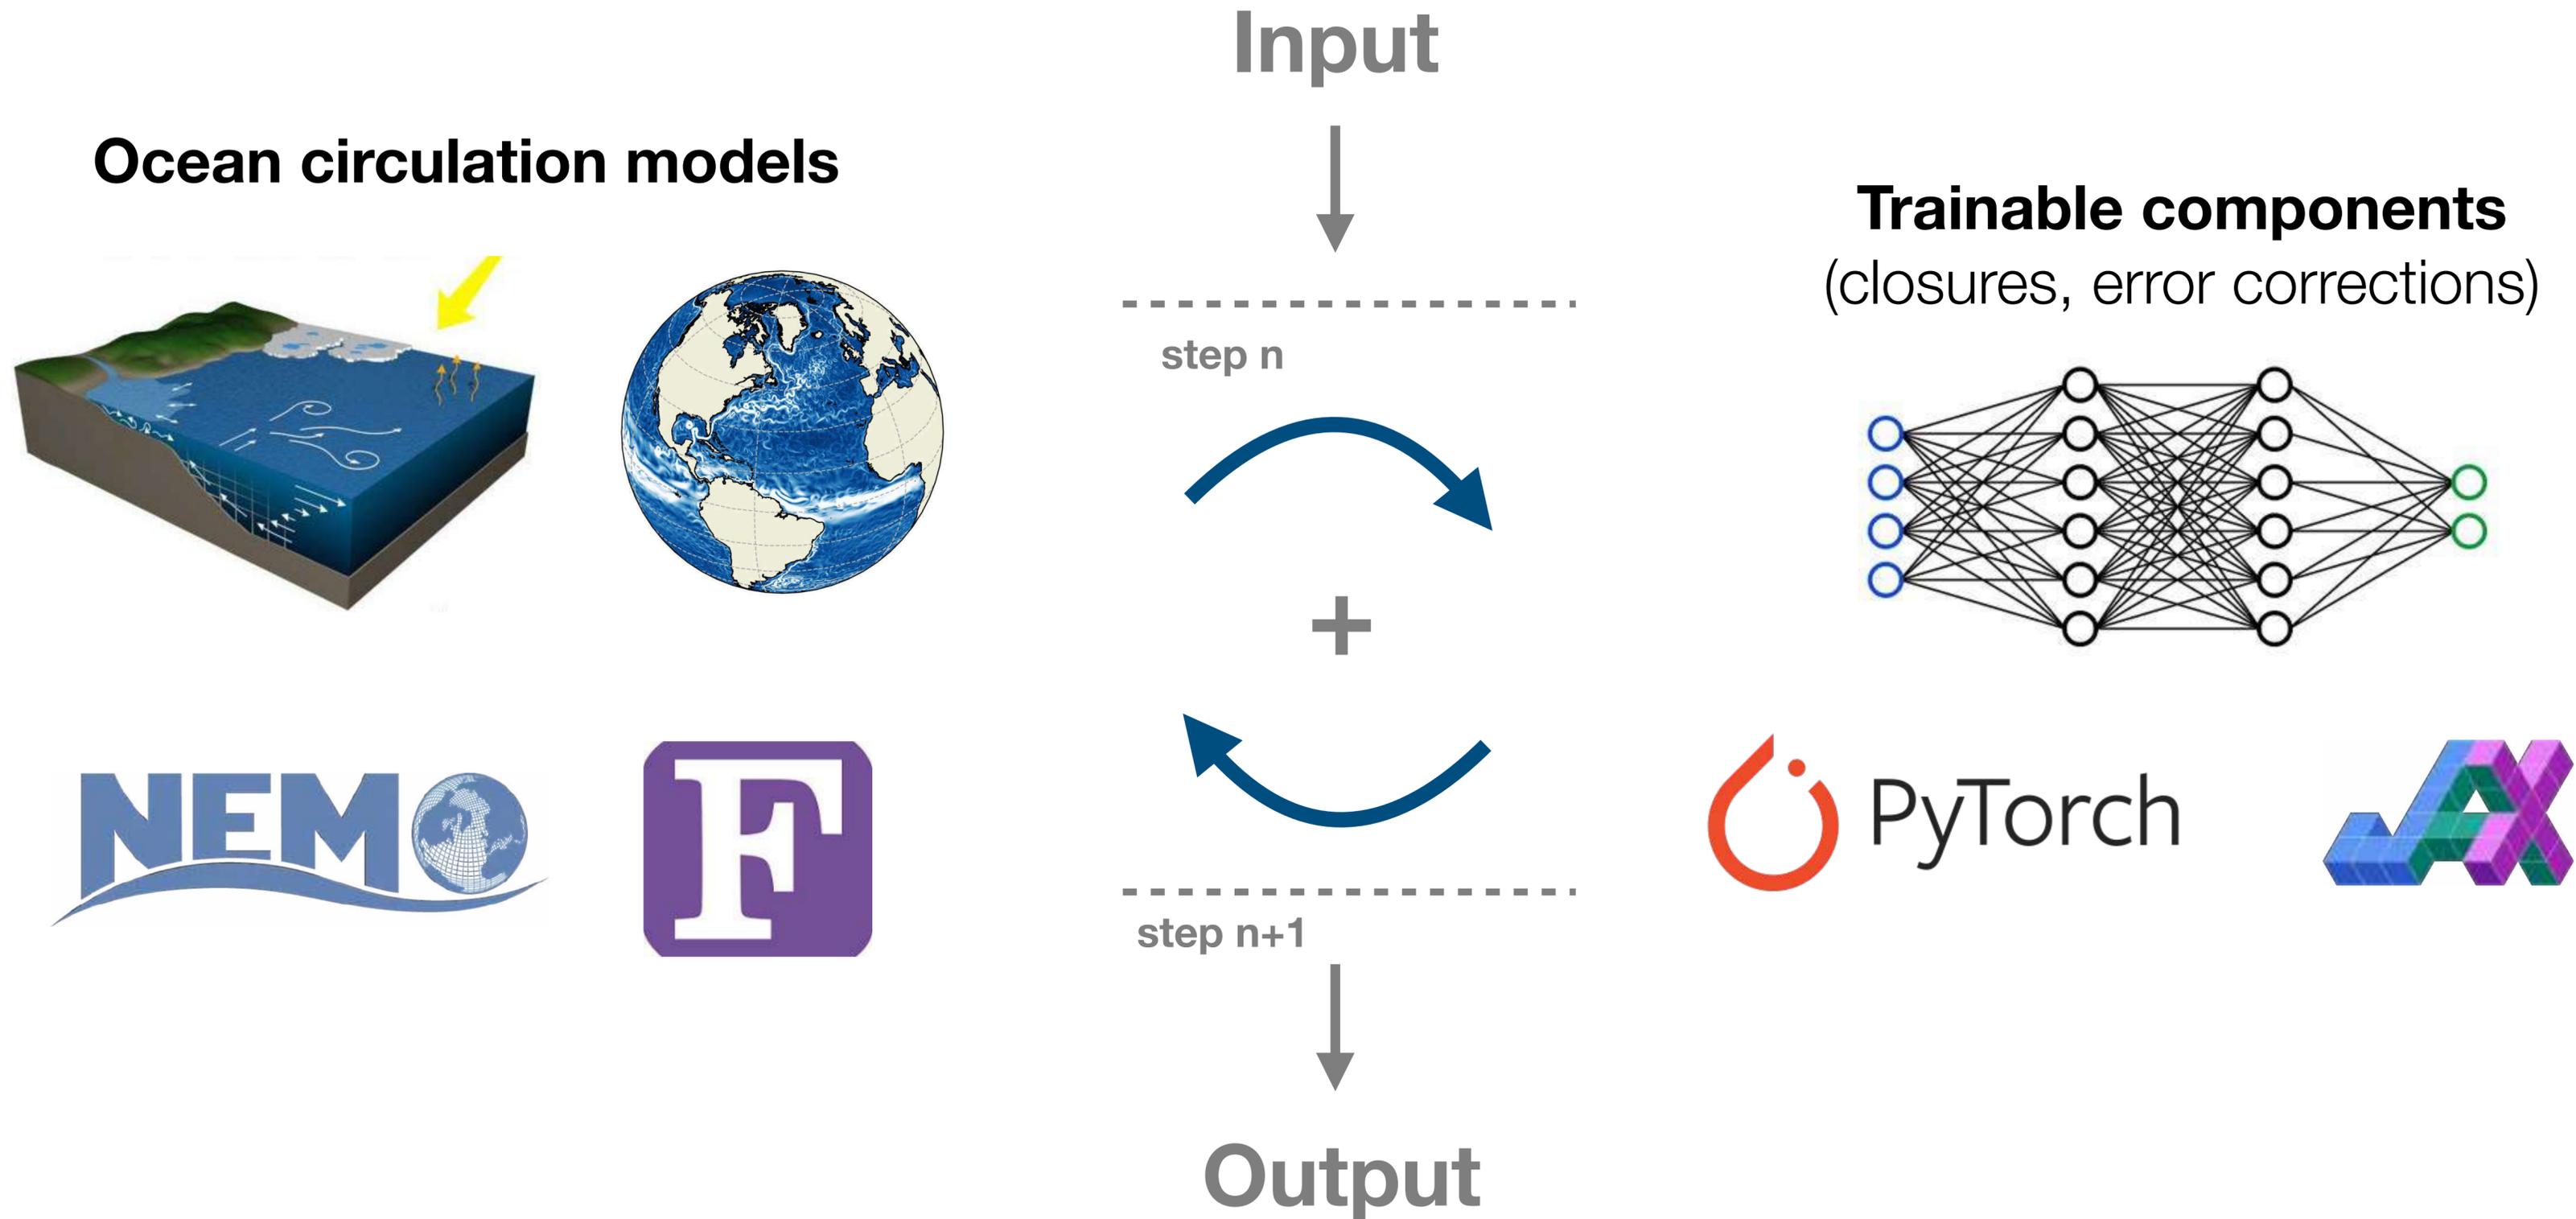


cloud ready

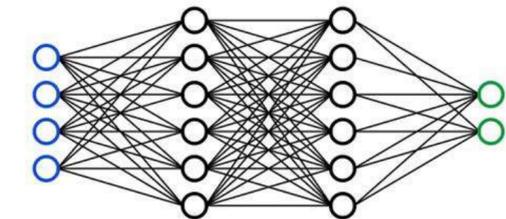
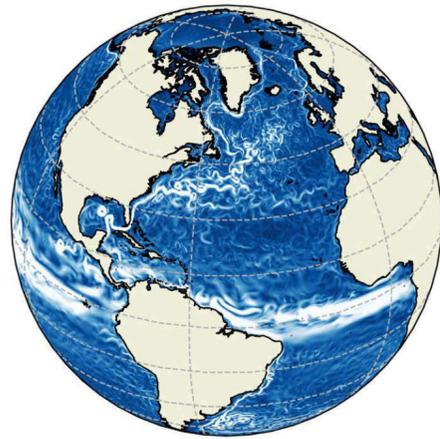


natively runs on GPUs

# Interfacing ocean models with DL frameworks (2/3)

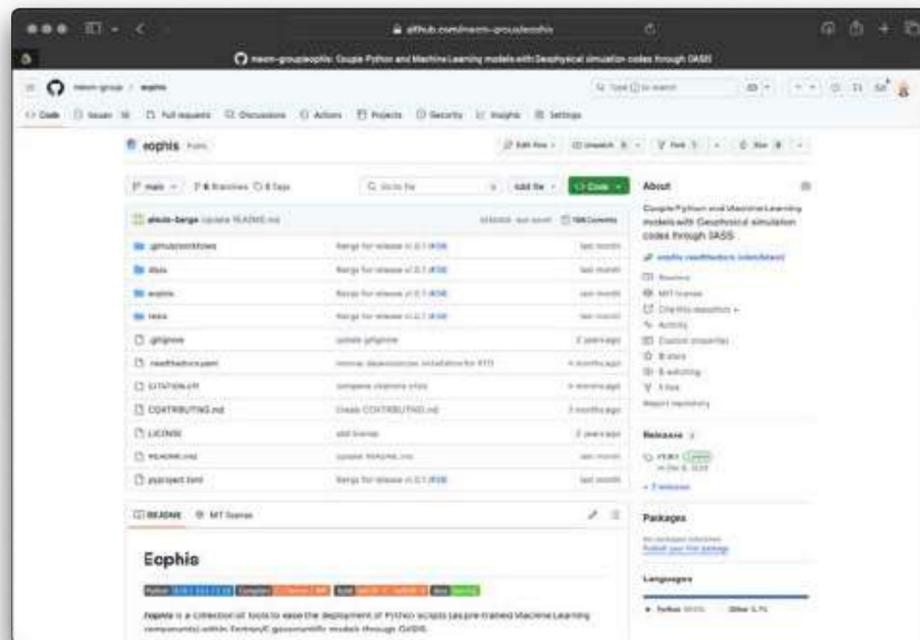


# Interfacing ocean models with DL frameworks (3/3)



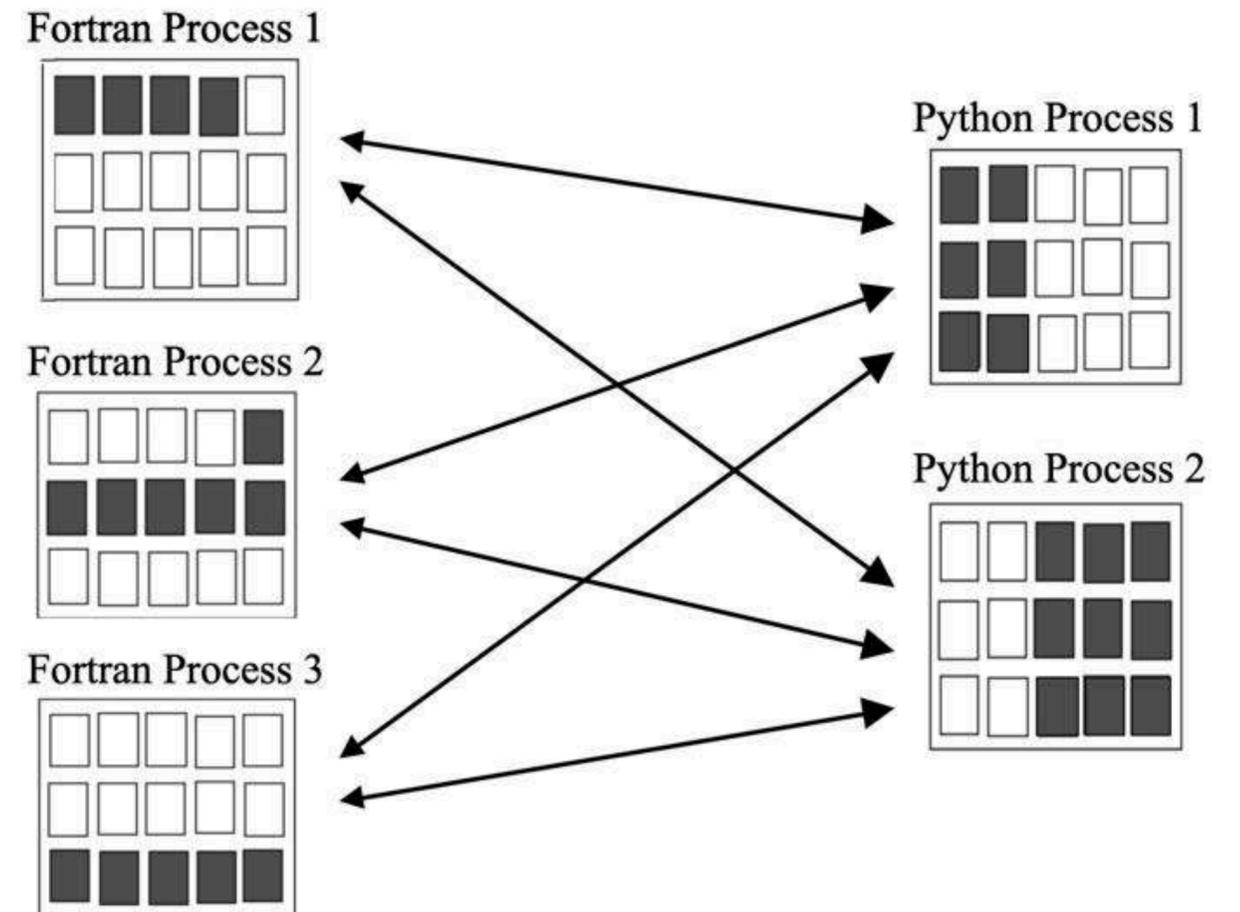
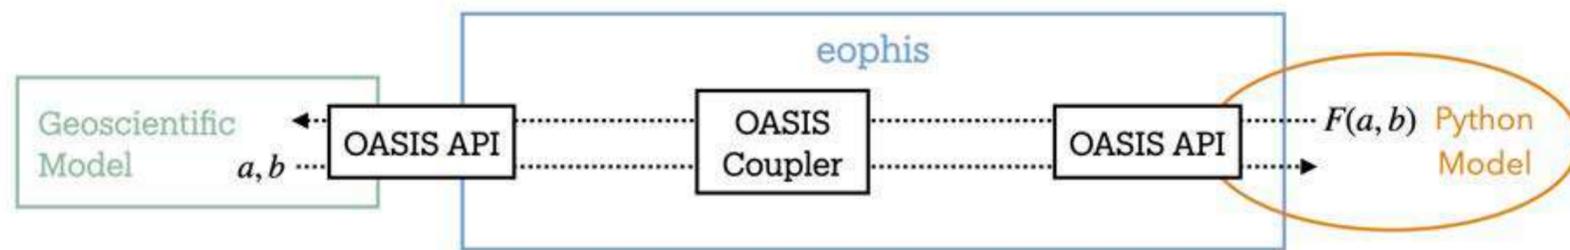
A. Barge

Work by Alexis Barge at IGE

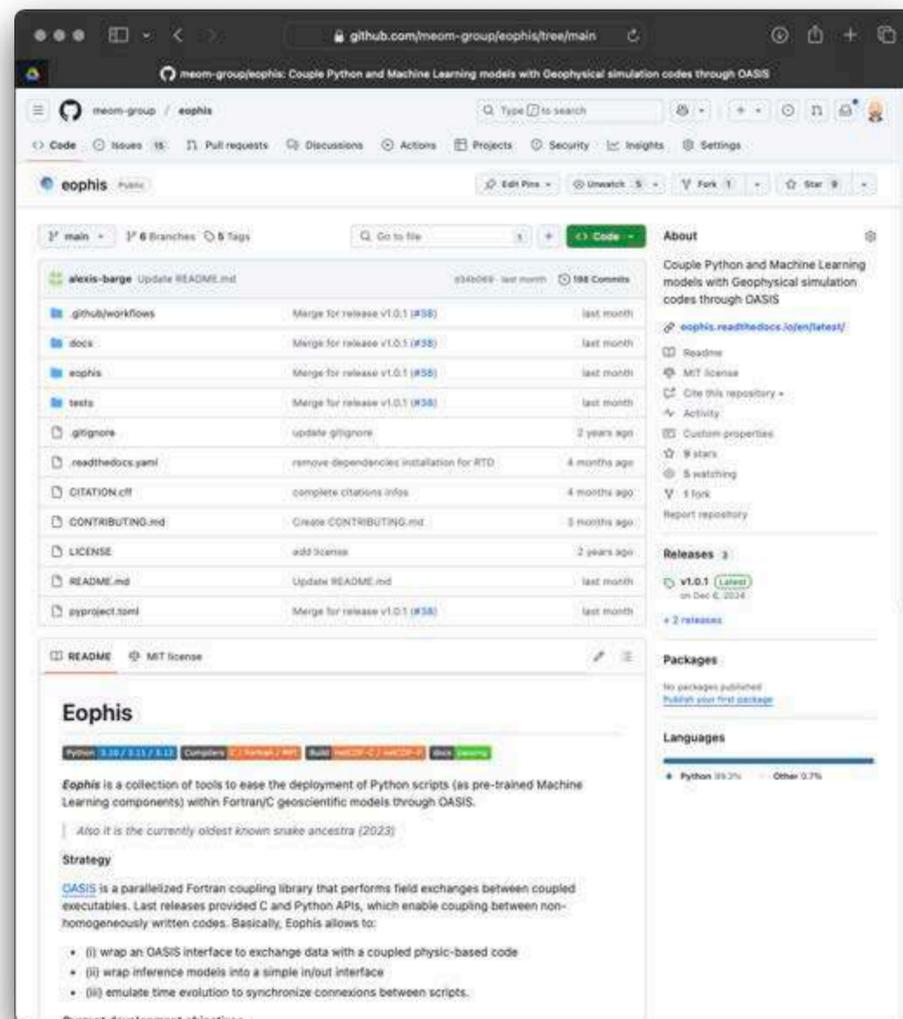


- OASIS : exchange of 3D data between different codes
- Eophis : simplified deployment of ML models w/ OASIS
- Requires some change to the NEMO code
- Key : portability, domain decomposition

# Interfacing ocean models with DL frameworks (3/3)



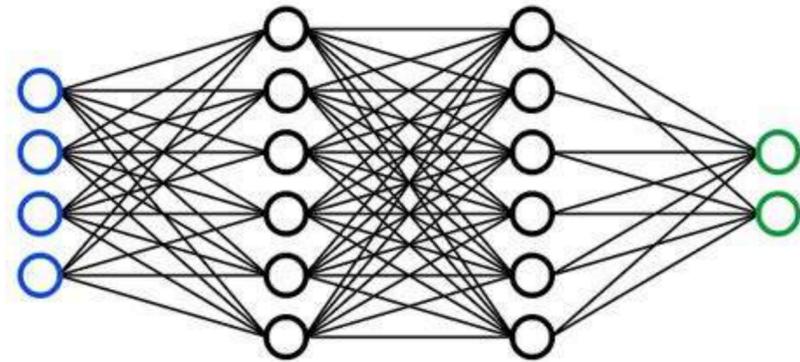
See eophys tutorial next week



- OASIS : exchange of 3D data between different codes
- Eophys : simplified deployment of ML models w/ OASIS
- Requires some change to the NEMO code
- Key : portability, domain decomposition

# The (real) challenge of online training (1/2)

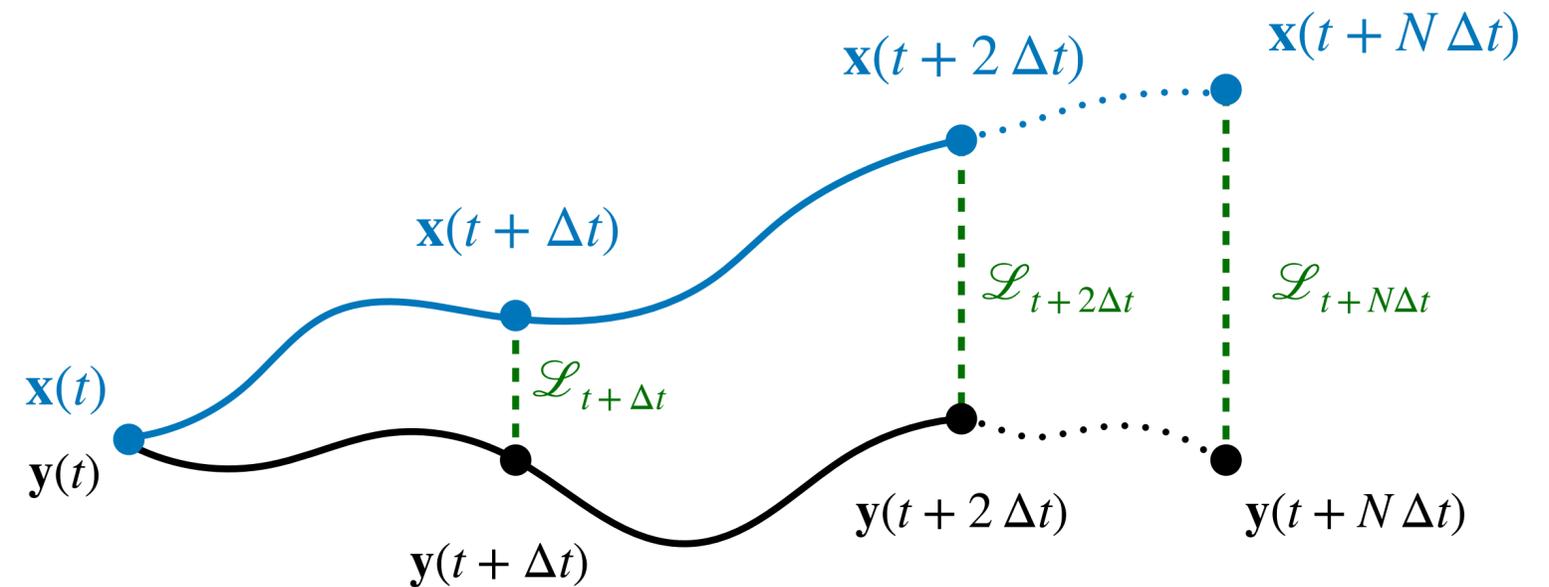
## offline learning



mapping  $\bar{\mathbf{x}} \rightarrow \overline{\mathcal{N}(\mathbf{x})}$

from pre-existing data

## online learning



$$\partial_t \mathbf{y} + G(\mathbf{y}) + \mathcal{M}_{NN}(\mathbf{y}) = f$$

along a trajectory

*(a.k.a : a posteriori, solver-in-the-loop, end-to-end, auto-regressive roll-outs)*

**Online training improves performance, stability, generalisation**

*Frezat et al. 2022; List et al. 2024*

# The (real) challenge of online training (2/2)

$$\arg \min_{\theta} \mathcal{L}(\mathbf{z}, \mathcal{M}(\mathbf{y} | \theta))$$

target                      prediction

$$\frac{\partial \mathcal{L}}{\partial \theta}(\mathbf{z}, \mathcal{M}(\mathbf{y} | \theta)) = \frac{\partial \mathcal{M}}{\partial \theta}(\mathbf{y} | \theta) \frac{\partial \mathcal{L}}{\partial \mathcal{M}}$$

gradient of the loss

For time evolving problems, with

$$\mathbf{y}(t + \Delta t) = E_m \circ \dots \circ E_1(\mathbf{y}(t))$$

$\mathcal{M} \equiv E$     Auto-regressive operator (time)

The gradient of the loss involves

tricky without Automatic  
Differentiation (AD) !

$$\frac{\partial \mathcal{M}}{\partial \theta} \equiv \frac{\partial E}{\partial \theta} = \frac{\partial(E_m \circ \dots \circ E_1)}{\partial \theta} = \frac{\partial E_m}{\partial E_{m-1}} \dots \frac{\partial E_2}{\partial E_1} \frac{\partial E_1}{\partial \theta}$$

# The challenge of online training strategies (2/2)

$$\arg \min_{\theta} \mathcal{L}(\mathbf{z}, \mathcal{M}(\mathbf{y} | \theta))$$

target                      prediction

AD is readily available in some language



For time evolving problems, with

$$\mathbf{y}(t + \Delta t) = E_m \circ \dots \circ E_1(\mathbf{y}(t))$$

$\mathcal{M} \equiv E$  temporal evolution operator

The gradient of the loss involves

$$\frac{\partial \mathcal{M}}{\partial \theta} \equiv \frac{\partial E}{\partial \theta} = \frac{\partial (E_m \circ \dots \circ E_1)}{\partial \theta} = \frac{\partial E_m}{\partial E_{m-1}} \dots \frac{\partial E_2}{\partial E_1} \frac{\partial E_1}{\partial \theta}$$

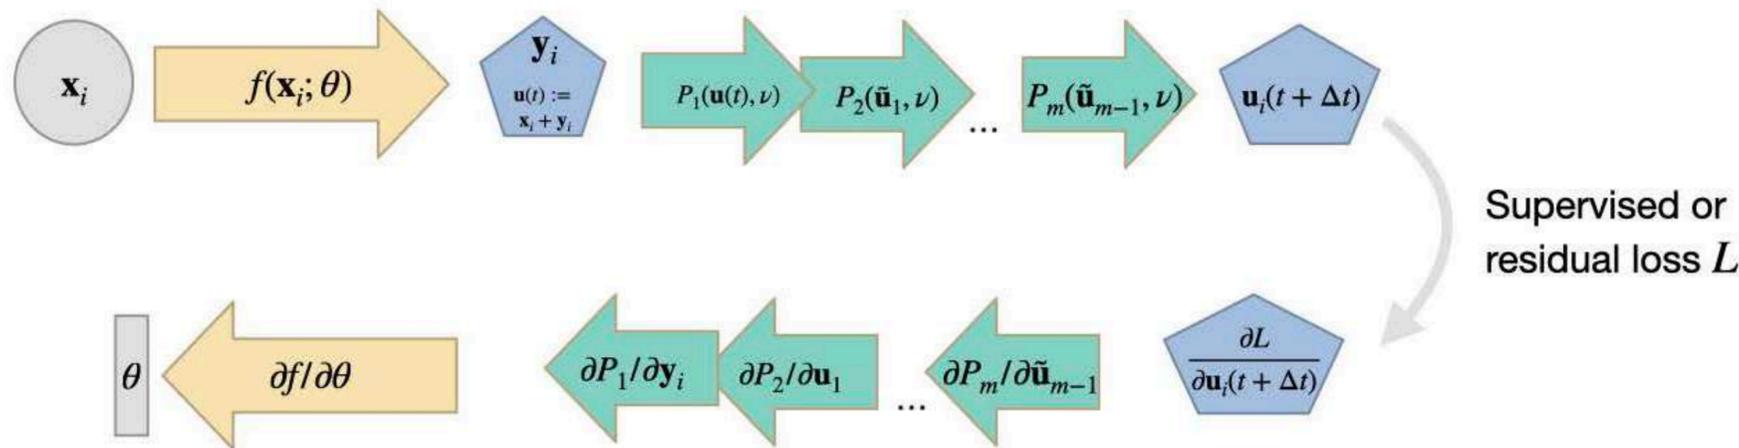
tricky without Automatic Differentiation (AD) !

# The challenge of online training strategies (2/2)

$$\arg \min_{\theta} \mathcal{L}(\mathbf{z}, \mathcal{M}(\mathbf{y} | \theta))$$

target prediction

AD is readily available in some language



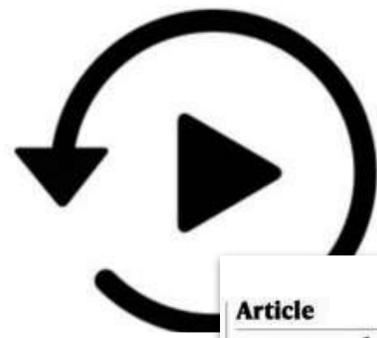
But AD used yet in climate models...

## Differentiable programming

- programs composed of **differentiable** building blocks
- building blocks : trainable and procedural **code components**
- **trainable end-to-end** with gradient based optimisation

See eg Sapienza et al. 2024

<https://arxiv.org/abs/2406.09699>



# AI-native hybrid geoscientific models

**Article**  
**Neural general circulation models for weather and climate**

<https://doi.org/10.1038/s41586-024-07744-y>  
 Received: 13 November 2023  
 Accepted: 15 June 2024  
 Published online: 22 July 2024

**Open access**  
 Check for updates

Dmitrii Kochkov<sup>1,2,3,4</sup>, Janni Yuval<sup>1,2,3,4</sup>, Ian Langmore<sup>1,2</sup>, Peter Norgaard<sup>1,2</sup>, Jamie Smith<sup>1,2</sup>, Griffin Mooers<sup>1,2</sup>, Milan Klöwer<sup>1,2</sup>, James Lottes<sup>1,2</sup>, Stephan Rasp<sup>1,2</sup>, Peter Düben<sup>1,2</sup>, Sam Hatfield<sup>1,2</sup>, Peter Battaglia<sup>1,2</sup>, Alvaro Sanchez-Gonzalez<sup>1,2</sup>, Matthew Willson<sup>1,2</sup>, Michael P. Brenner<sup>1,2</sup> & Stephan Hoyer<sup>1,2,3,4</sup>

General circulation models (GCMs) are the foundation of weather and climate prediction<sup>1,2</sup>. GCMs are physics-based simulators that combine a numerical solver for large-scale dynamics with tuned representations for small-scale processes such as cloud formation. Recently, machine-learning models trained on reanalysis data have achieved comparable or better skill than GCMs for deterministic weather forecasting<sup>3,4</sup>. However, these models have not demonstrated improved ensemble forecasts, or shown sufficient stability for long-term weather and climate simulations. Here we present a GCM that combines a differentiable solver for atmospheric dynamics with machine-learning components and show that it can generate forecasts of deterministic weather, ensemble weather and climate on par with the best machine-learning and physics-based methods. NeuralGCM is competitive with machine-learning models for one- to ten-day forecasts, and with the European Centre for Medium-Range Weather Forecasts ensemble prediction for one- to fifteen-day forecasts. With prescribed sea surface temperature, NeuralGCM can accurately track climate metrics for multiple decades, and climate forecasts with 140-kilometre resolution show emergent phenomena such as realistic frequency and trajectories of tropical cyclones. For both weather and climate, our approach offers orders of magnitude computational savings over conventional GCMs, although our model does not extrapolate to substantially different future climates. Our results show that end-to-end deep learning is compatible with tasks performed by conventional GCMs and can enhance the large-scale physical simulations that are essential for understanding and predicting the Earth system.

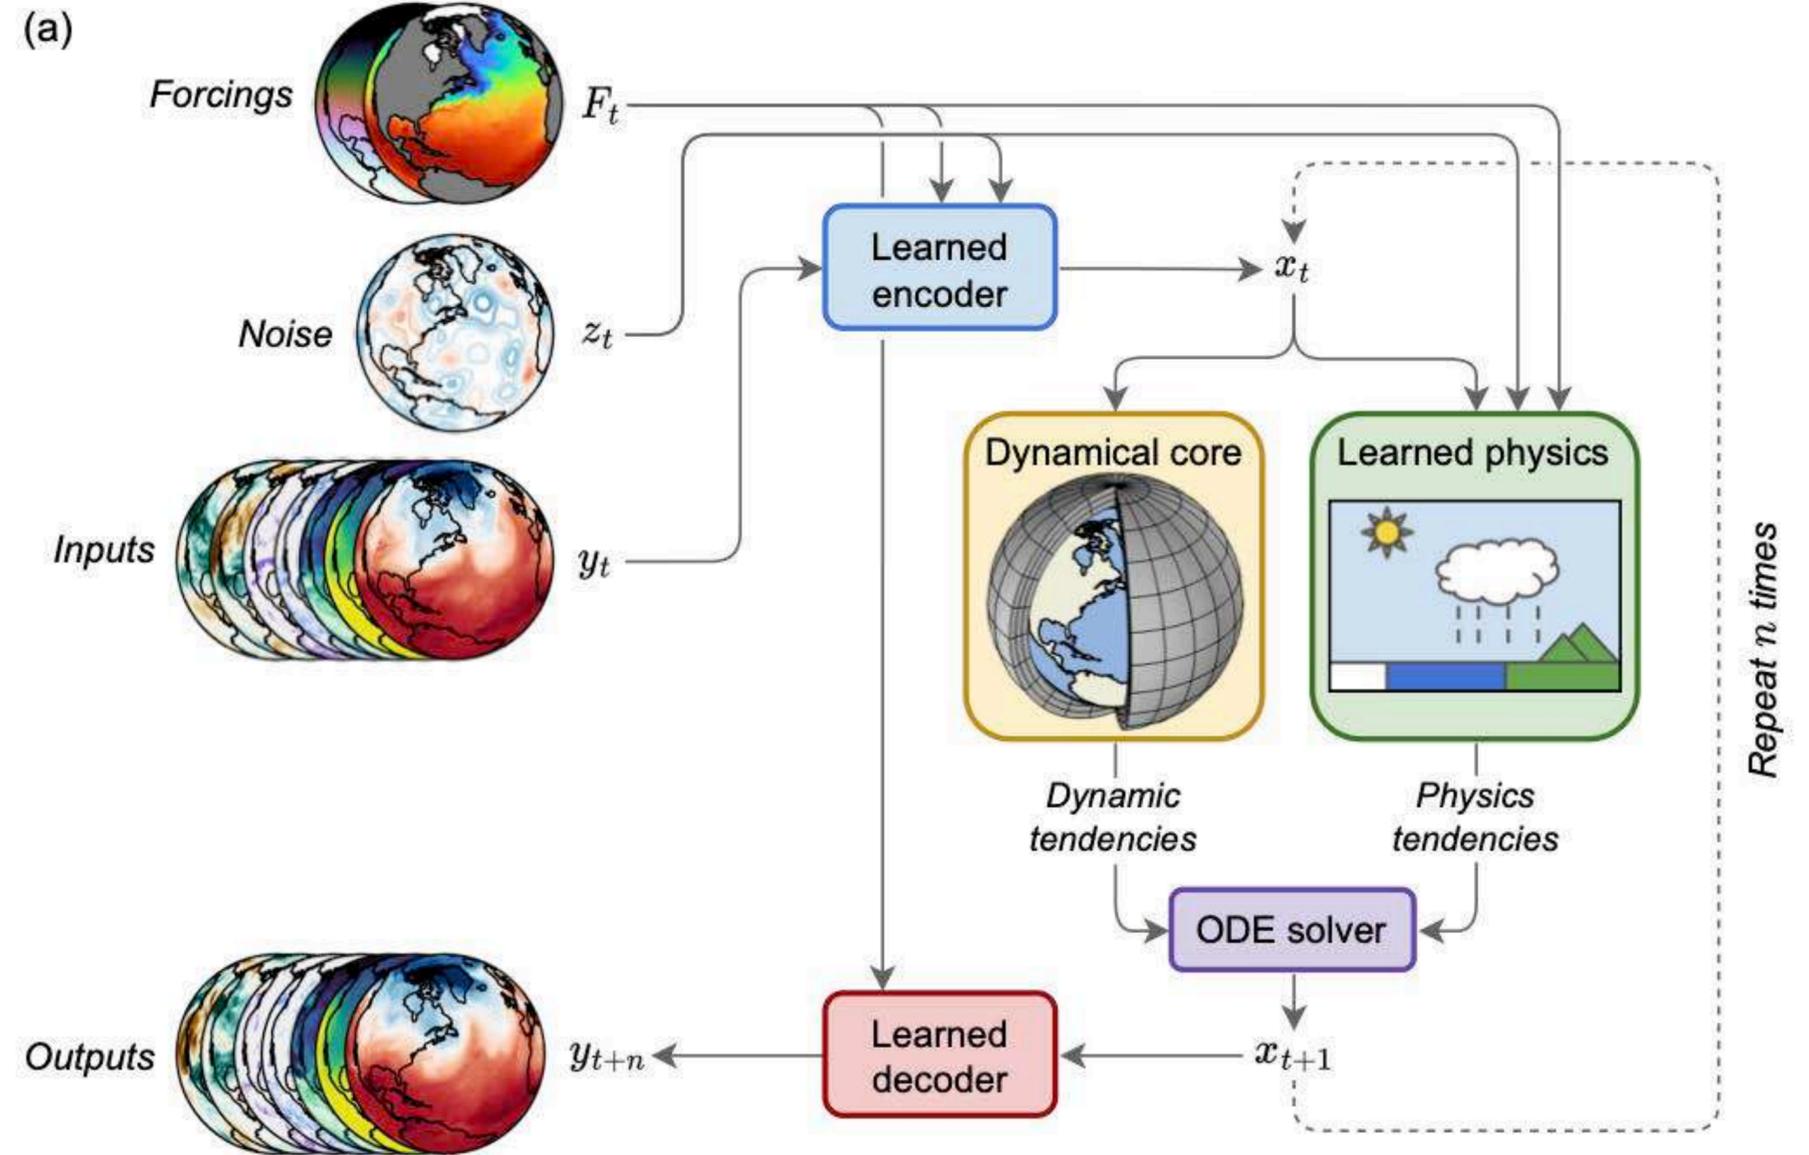
Solving the equations for Earth's atmosphere with general circulation models (GCMs) is the basis of weather and climate prediction<sup>1,2</sup>. Over the past 70 years, GCMs have been steadily improved with better numerical methods and more detailed physical models, while exploiting faster computers to run at higher resolution. Inside GCMs, the unresolved physical processes such as clouds, radiation and precipitation are represented by semi-empirical parameterizations. Tuning GCMs to match historical data remains a manual process<sup>3</sup>, and GCMs retain many persistent errors and biases<sup>4</sup>. The difficulty of reducing uncertainty in long-term climate projections<sup>5</sup> and estimating distributions of extreme weather events<sup>6</sup> presents major challenges for climate mitigation and adaptation<sup>7</sup>.

Recent advances in machine learning have presented an alternative for weather forecasting<sup>8,9,10,11</sup>. These models rely solely on machine-learning techniques, using roughly 40 years of historical data from the European Centre for Medium-Range Weather Forecasts (ECMWF) reanalysis v5 (ERAS)<sup>12</sup> for model training and forecast initialization. Machine-learning methods have been remarkably successful, demonstrating state-of-the-art deterministic forecasts for 1- to 10-day weather prediction at a fraction of the computational cost of traditional models<sup>8,9</sup>. Machine-learning atmospheric models also require considerably less code, for example GraphCast<sup>10</sup> has 5,417 lines versus 376,578 lines for the National Oceanic and Atmospheric Administration's FV3 atmospheric model<sup>13</sup> (see Supplementary Information section A for details).

Nevertheless, machine-learning approaches have noteworthy limitations compared with GCMs. Existing machine-learning models have focused on deterministic prediction, and surpass deterministic numerical weather prediction in terms of the aggregate metrics for which they are trained<sup>8,9</sup>. However, they do not produce calibrated uncertainty estimates<sup>4</sup>, which is essential for useful weather forecasts<sup>1</sup>. Deterministic machine-learning models using a mean-squared-error loss are rewarded for averaging over uncertainty, producing unrealistically blurry predictions when optimized for multi-day forecasts<sup>8,9</sup>. Unlike physical models, machine-learning models misrepresent derived (diagnostic) variables such as geostrophic wind<sup>14</sup>. Furthermore,

<sup>1</sup>Google Research, Mountain View, CA, USA. <sup>2</sup>Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA. <sup>3</sup>European Centre for Medium-Range Weather Forecasts, Reading, UK. <sup>4</sup>Google DeepMind, London, UK. <sup>5</sup>School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA. \*These authors contributed equally: Dmitrii Kochkov, Janni Yuval, Ian Langmore, Peter Norgaard, Jamie Smith, Stephan Hoyer. \*e-mail: dkochkov@google.com; janniyuval@google.com; shoyer@google.com

1060 | Nature | Vol 632 | 29 August 2024



<https://doi.org/10.1038/s41586-024-07744-y>

Kochkov et al. (2024)



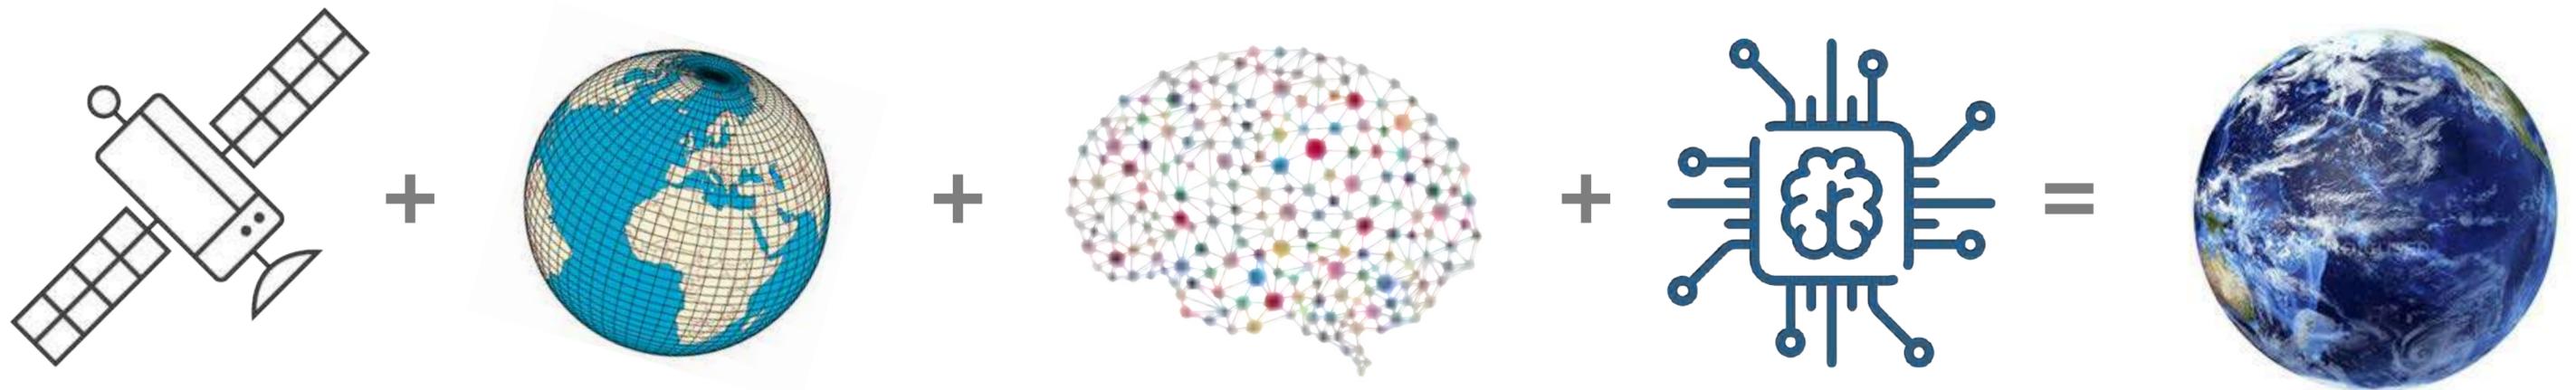
<https://github.com/google-research/dinosaur>

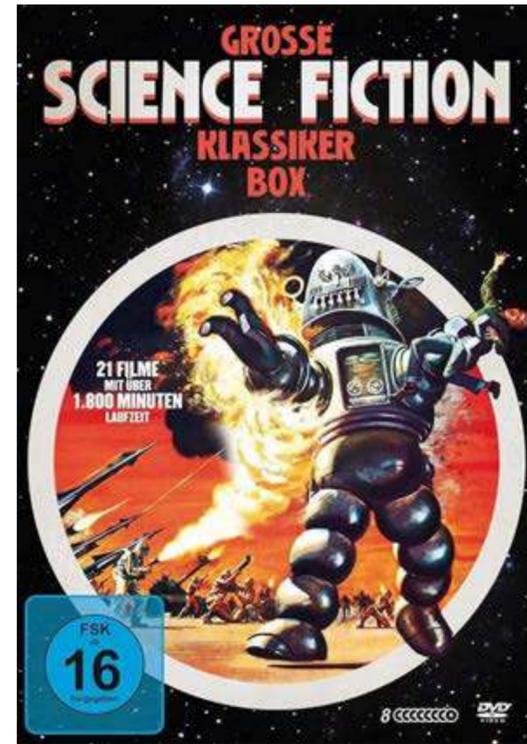
<https://github.com/google-research/neuralgcm>



5.

# **Towards** AI-native hybrid climate models ?



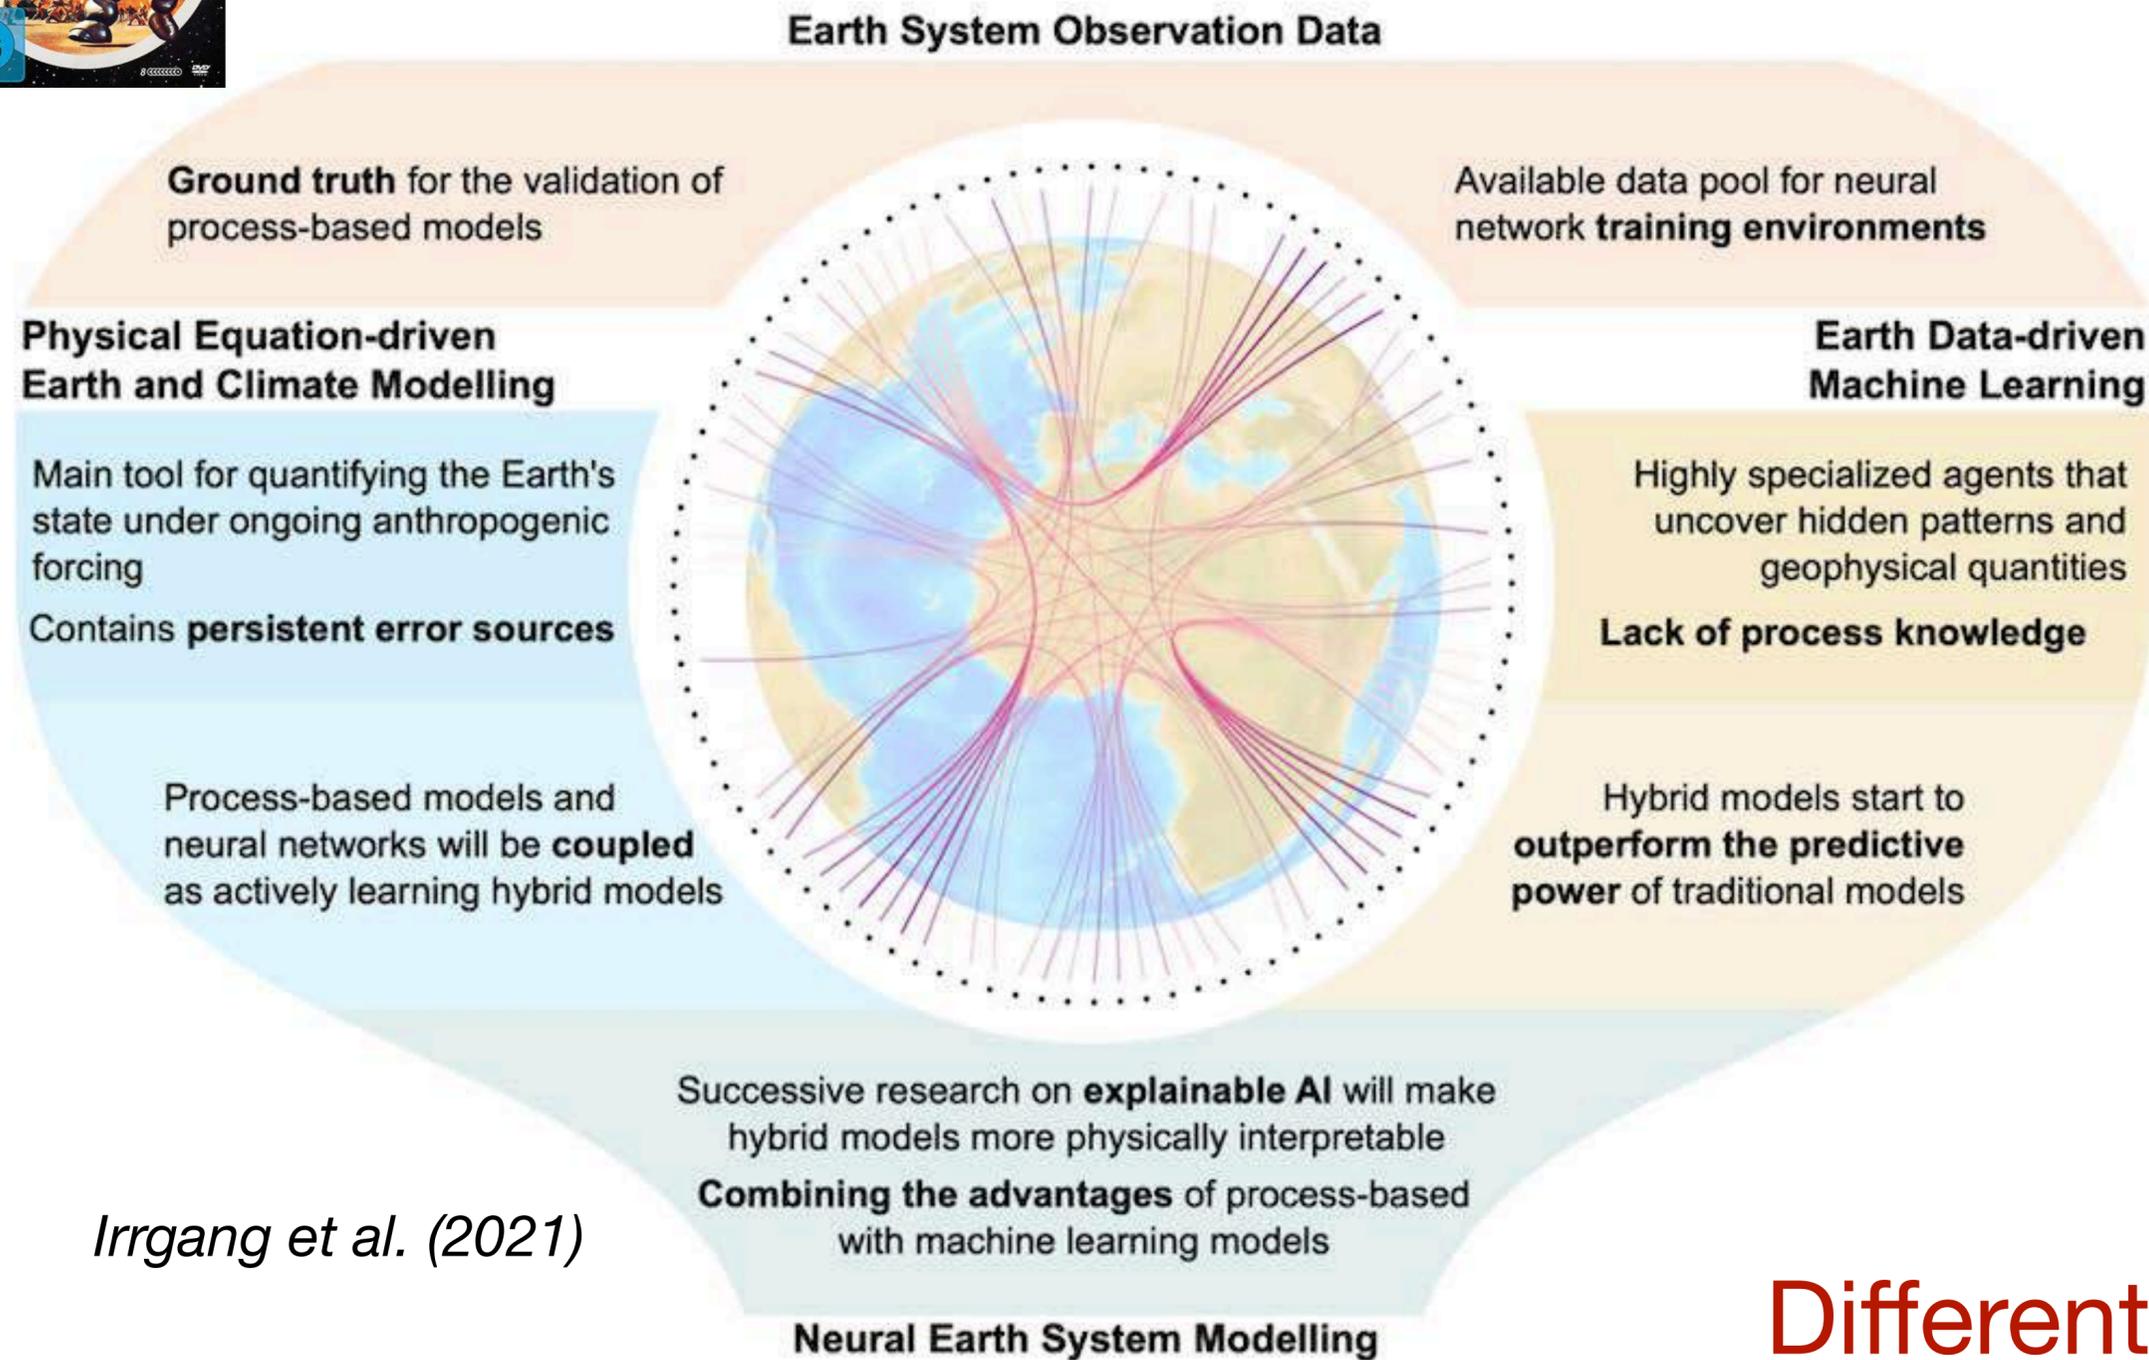


# Towards AI-native hybrid climate models ?





# AI-native hybrid climate models ?



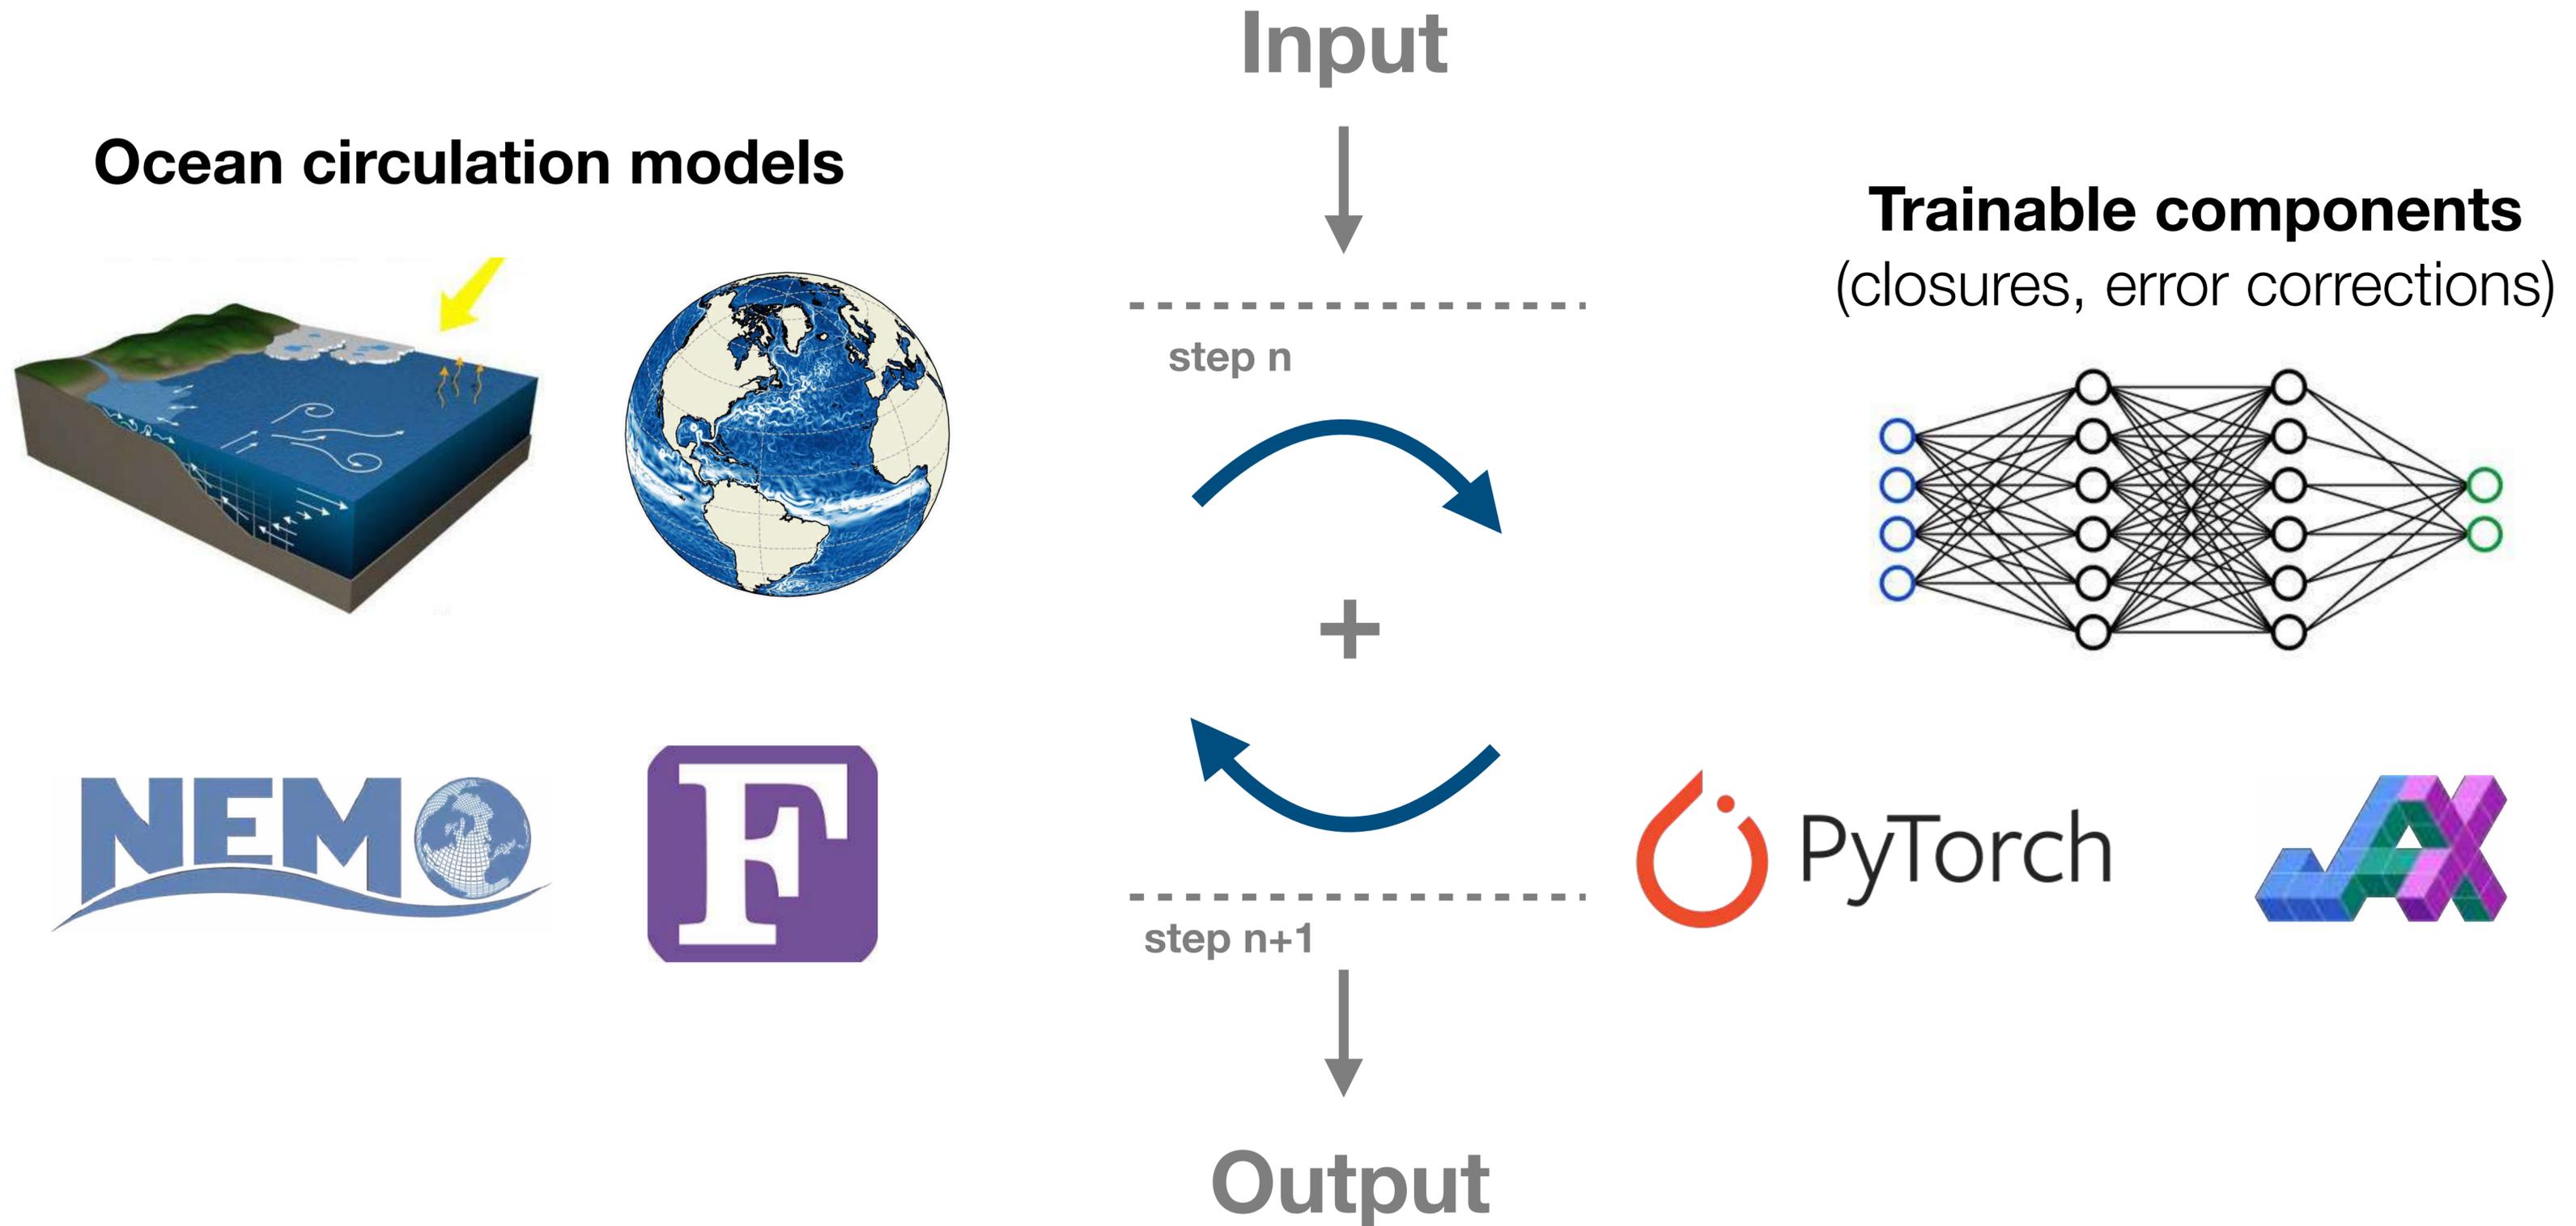
Irrgang et al. (2021)

Betting harnessing **observations & hi-fidelity simulations**

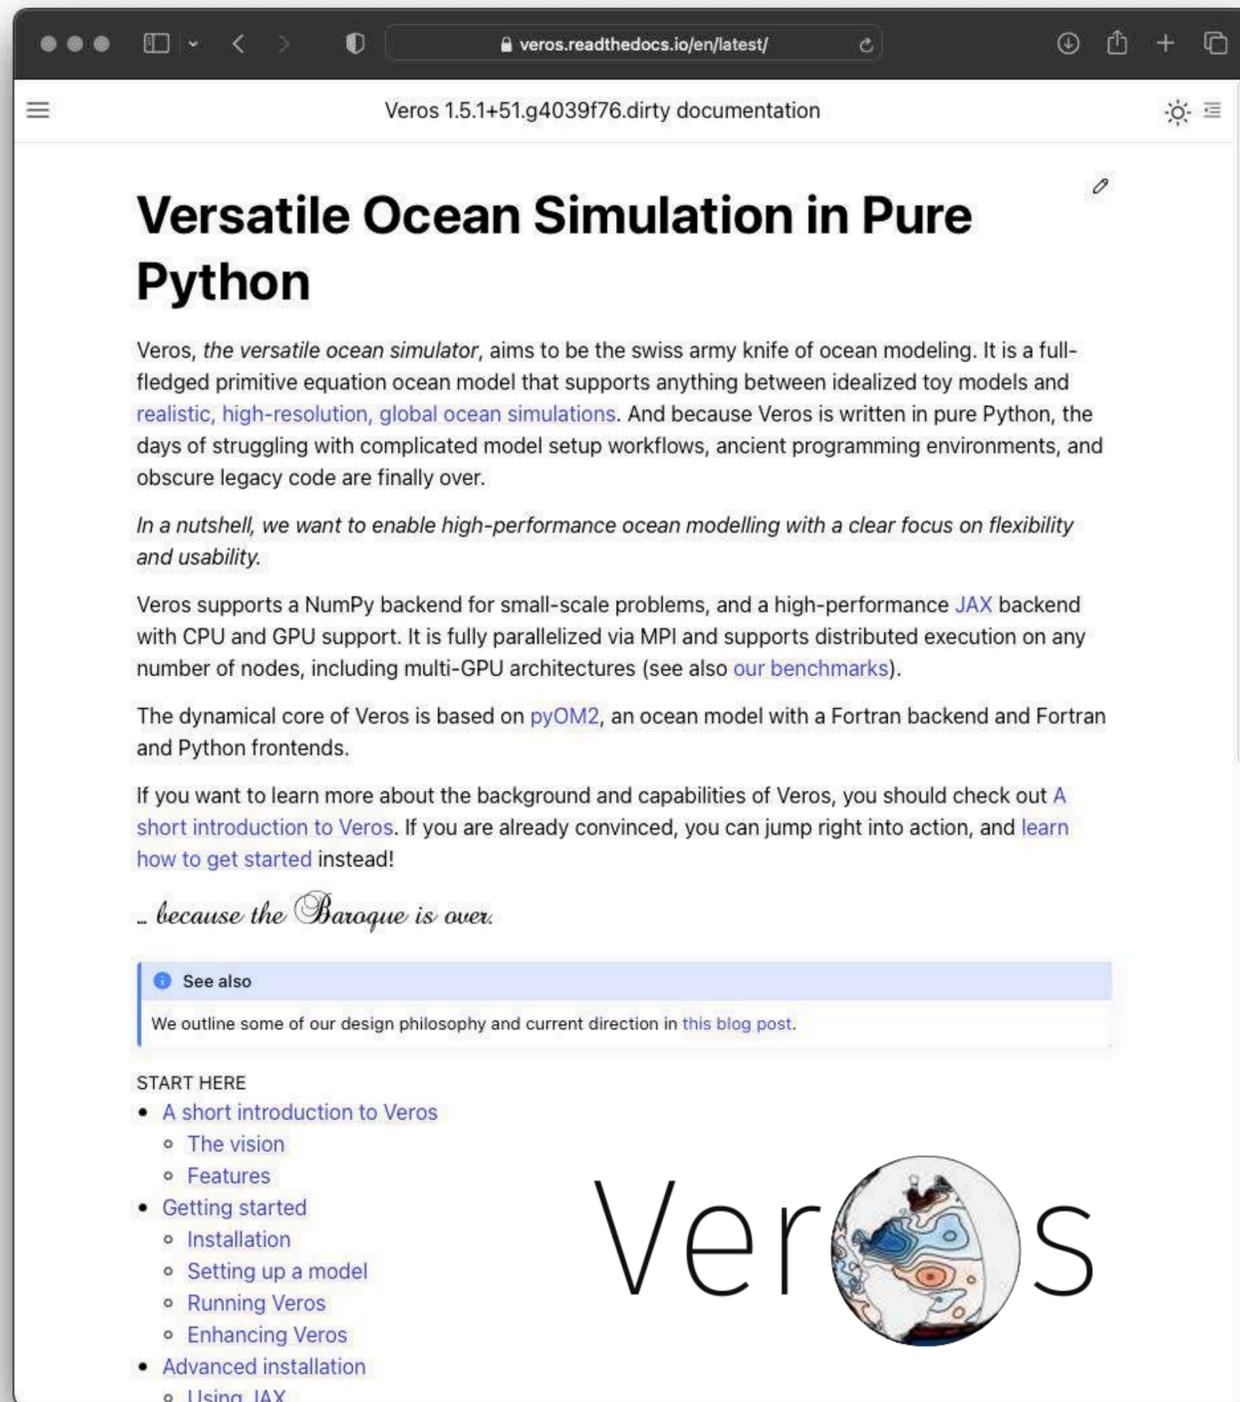
- ... for optimising
- model **parameters**
- **numerical** schemes
- subgrid **closures**
- ...

**Differentiable programming**  
in earth system models ?

# Current generation hybrid models



# A new generation of geoscientific models



veros.readthedocs.io/en/latest/

## Versatile Ocean Simulation in Pure Python

Veros, the versatile ocean simulator, aims to be the swiss army knife of ocean modeling. It is a full-fledged primitive equation ocean model that supports anything between idealized toy models and realistic, high-resolution, global ocean simulations. And because Veros is written in pure Python, the days of struggling with complicated model setup workflows, ancient programming environments, and obscure legacy code are finally over.

*In a nutshell, we want to enable high-performance ocean modelling with a clear focus on flexibility and usability.*

Veros supports a NumPy backend for small-scale problems, and a high-performance JAX backend with CPU and GPU support. It is fully parallelized via MPI and supports distributed execution on any number of nodes, including multi-GPU architectures (see also our benchmarks).

The dynamical core of Veros is based on pyOM2, an ocean model with a Fortran backend and Fortran and Python frontends.

If you want to learn more about the background and capabilities of Veros, you should check out [A short introduction to Veros](#). If you are already convinced, you can jump right into action, and [learn how to get started](#) instead!

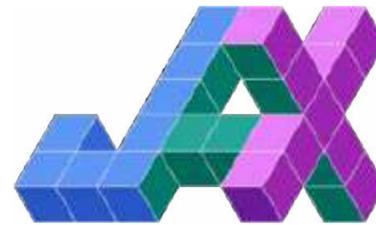
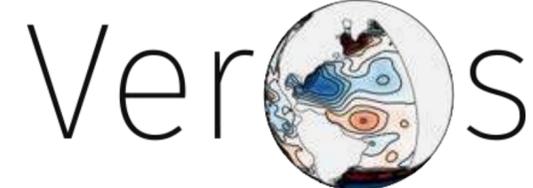
*... because the Baroque is over.*

**See also**

We outline some of our design philosophy and current direction in [this blog post](#).

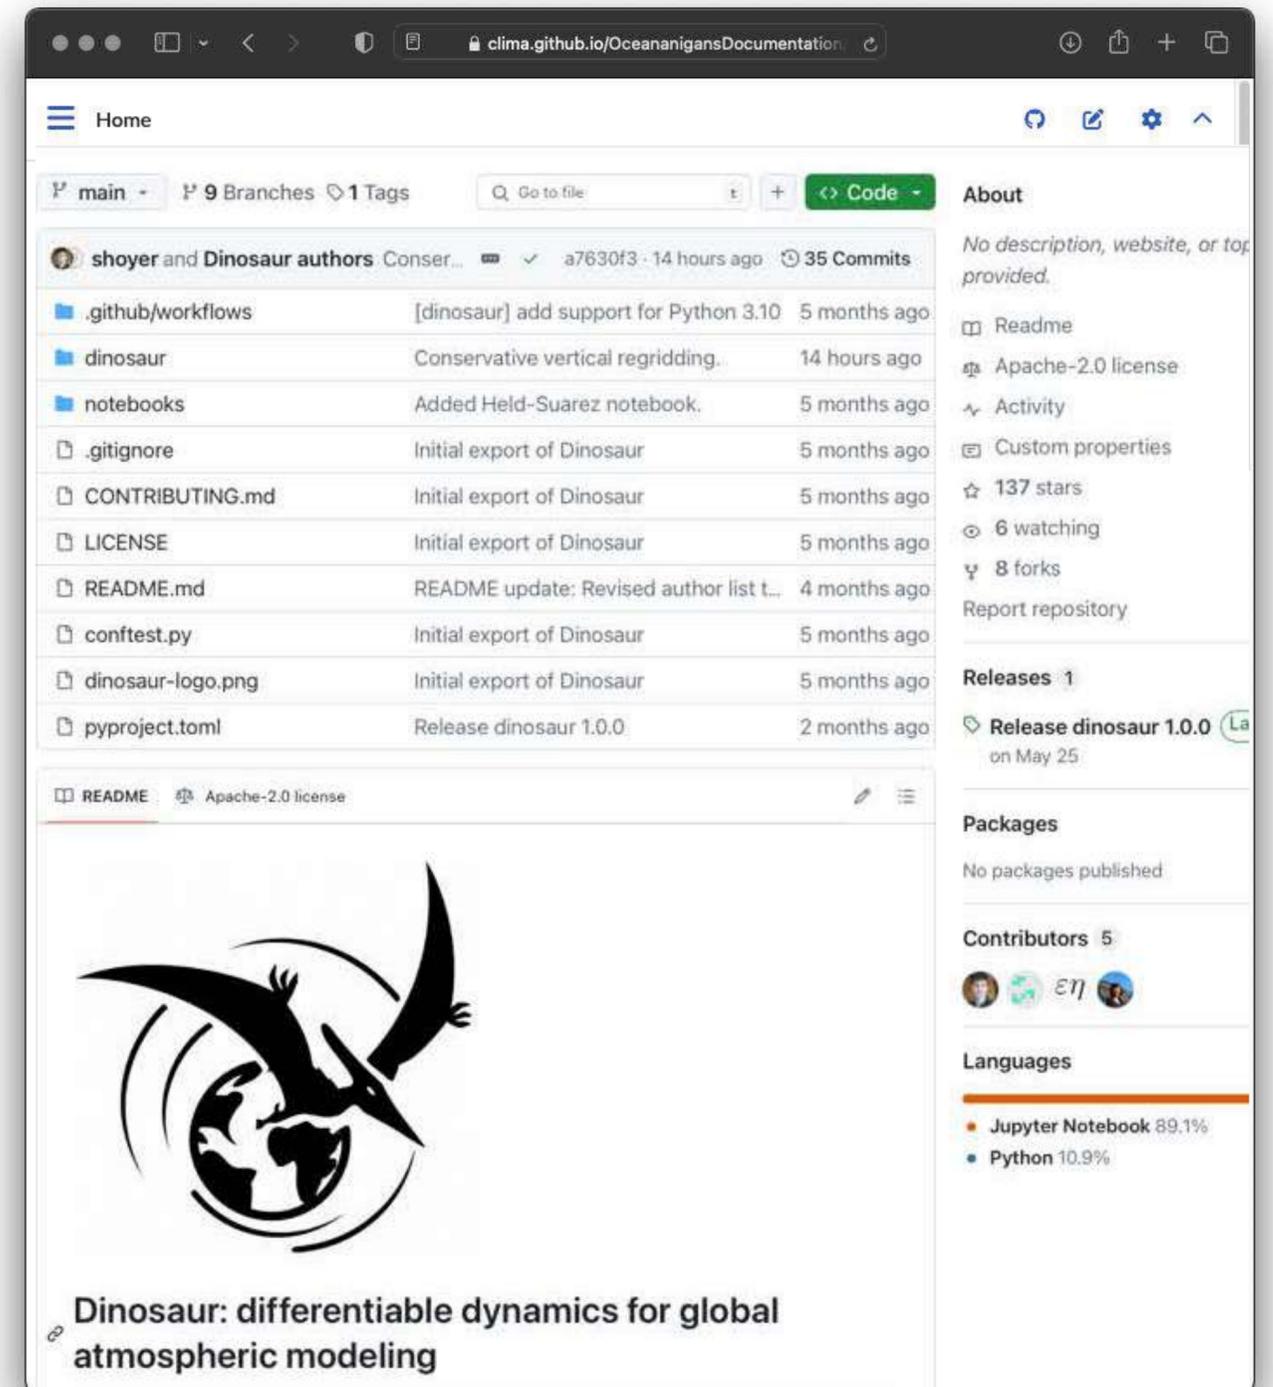
**START HERE**

- [A short introduction to Veros](#)
  - The vision
  - Features
- [Getting started](#)
  - Installation
  - Setting up a model
  - Running Veros
  - Enhancing Veros
- [Advanced installation](#)
  - Using JAX



Atmos

Ocean



clima.github.io/OceananigansDocumentation

## Home

main 9 Branches 1 Tags

shoyer and Dinosaur authors

| File              | Commit                                  | Time         |
|-------------------|---|--------------|
| .github/workflows | [dinosaur] add support for Python 3.10  | 5 months ago |
| dinosaur          | Conservative vertical regridding.       | 14 hours ago |
| notebooks         | Added Held-Suarez notebook.             | 5 months ago |
| .gitignore        | Initial export of Dinosaur              | 5 months ago |
| CONTRIBUTING.md   | Initial export of Dinosaur              | 5 months ago |
| LICENSE           | Initial export of Dinosaur              | 5 months ago |
| README.md         | README update: Revised author list t... | 4 months ago |
| confstest.py      | Initial export of Dinosaur              | 5 months ago |
| dinosaur-logo.png | Initial export of Dinosaur              | 5 months ago |
| pyproject.toml    | Release dinosaur 1.0.0                  | 2 months ago |

**README** Apache-2.0 license



**Dinosaur: differentiable dynamics for global atmospheric modeling**

**Releases** 1

Release dinosaur 1.0.0 on May 25

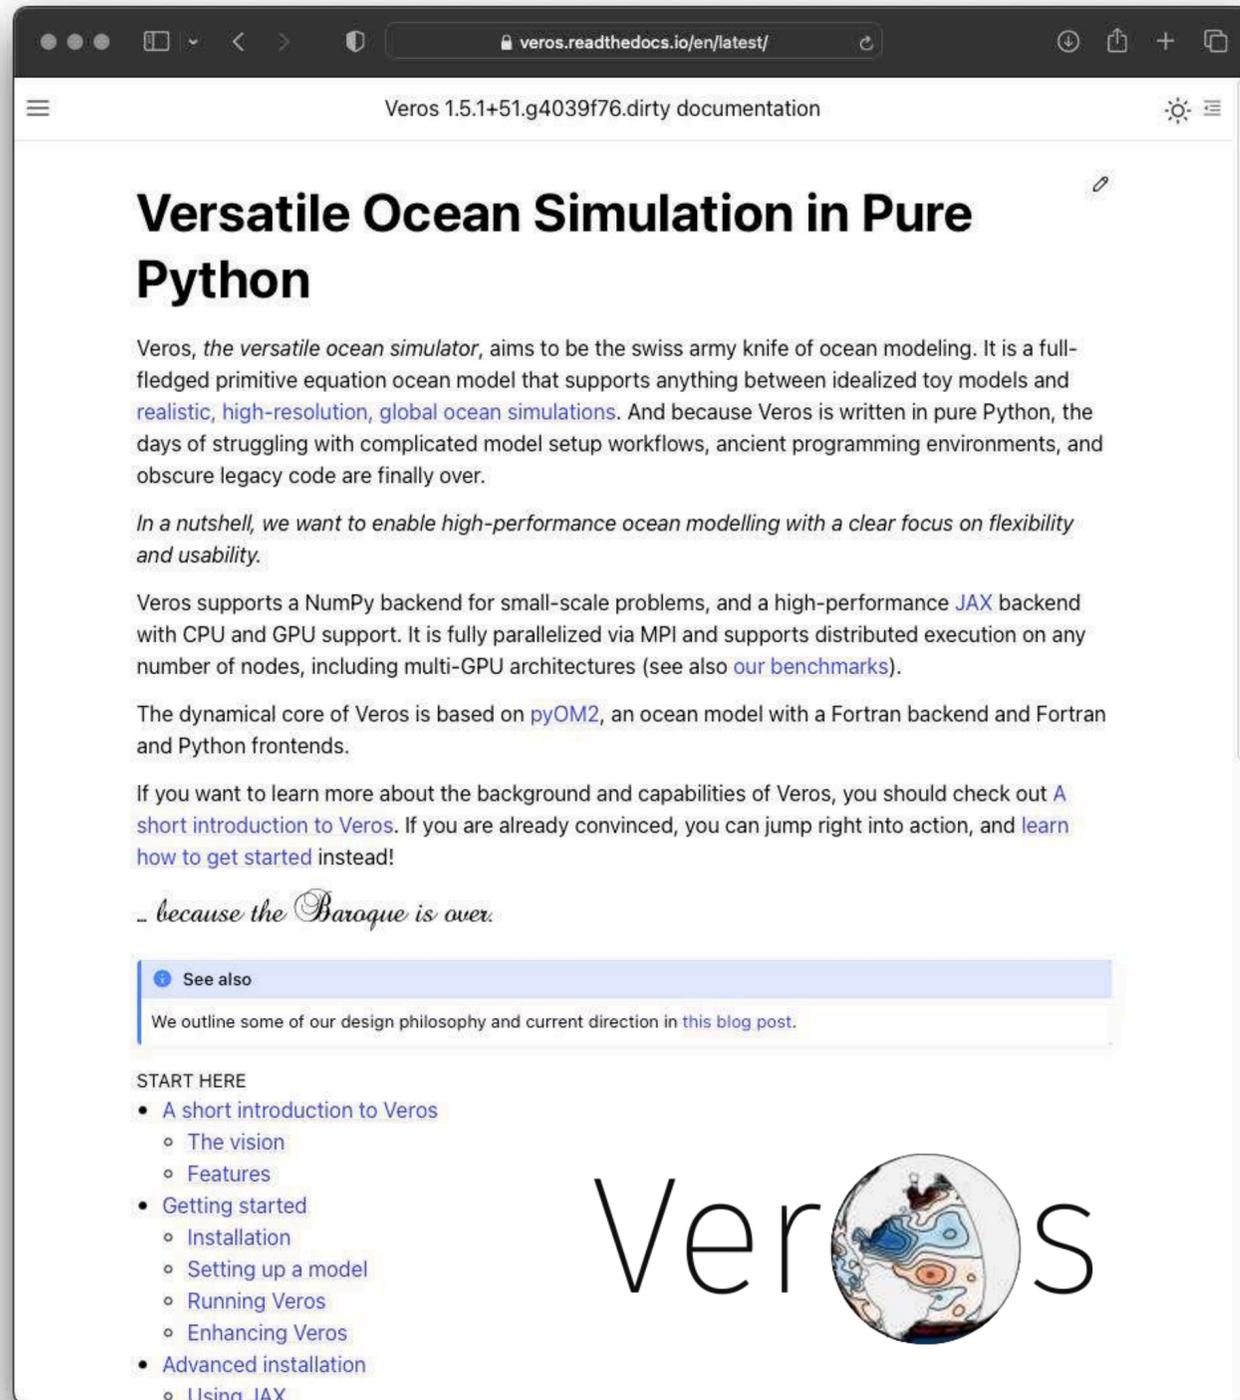
**Contributors** 5

**Languages**

- Jupyter Notebook 89.1%
- Python 10.9%

Modern code and compute : simple to write, scales, runs on any hardware

# A new generation of geoscientific models



Veros 1.5.1+51.g4039f76.dirty documentation

## Versatile Ocean Simulation in Pure Python

Veros, the versatile ocean simulator, aims to be the swiss army knife of ocean modeling. It is a full-fledged primitive equation ocean model that supports anything between idealized toy models and realistic, high-resolution, global ocean simulations. And because Veros is written in pure Python, the days of struggling with complicated model setup workflows, ancient programming environments, and obscure legacy code are finally over.

In a nutshell, we want to enable high-performance ocean modelling with a clear focus on flexibility and usability.

Veros supports a NumPy backend for small-scale problems, and a high-performance JAX backend with CPU and GPU support. It is fully parallelized via MPI and supports distributed execution on any number of nodes, including multi-GPU architectures (see also our benchmarks).

The dynamical core of Veros is based on pyOM2, an ocean model with a Fortran backend and Fortran and Python frontends.

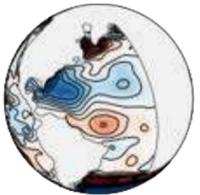
If you want to learn more about the background and capabilities of Veros, you should check out [A short introduction to Veros](#). If you are already convinced, you can jump right into action, and [learn how to get started](#) instead!

*... because the Baroque is over.*

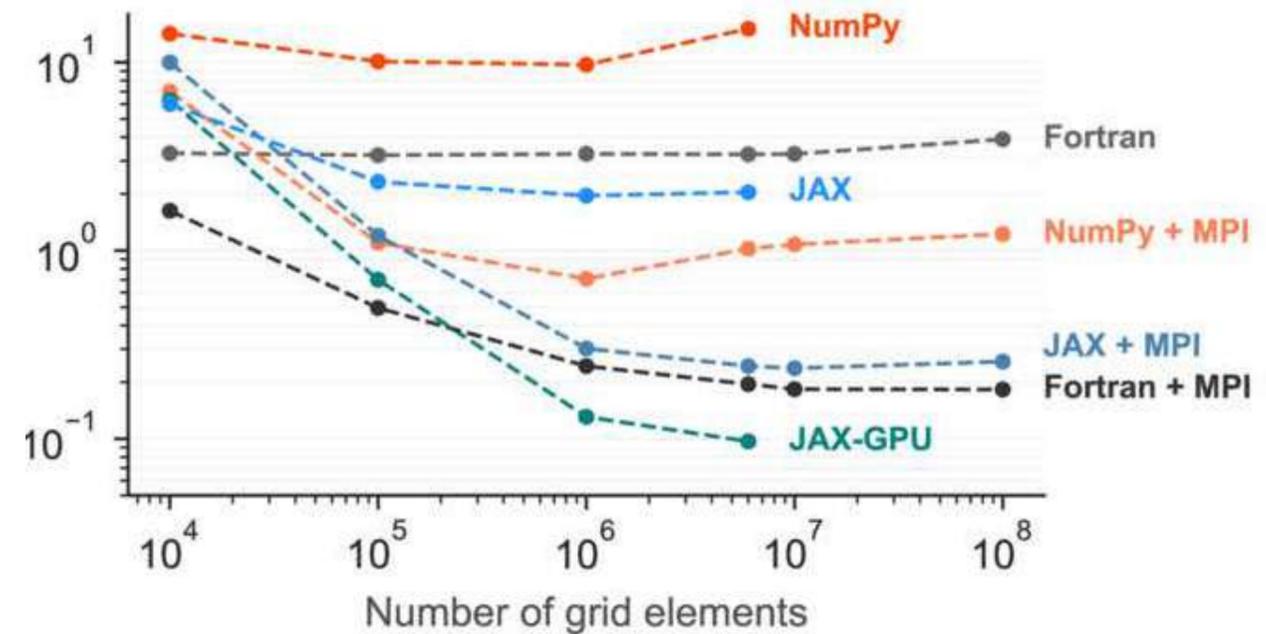
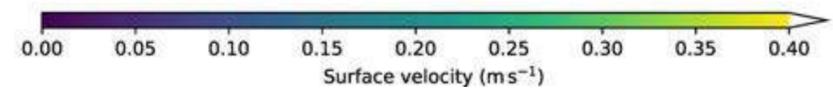
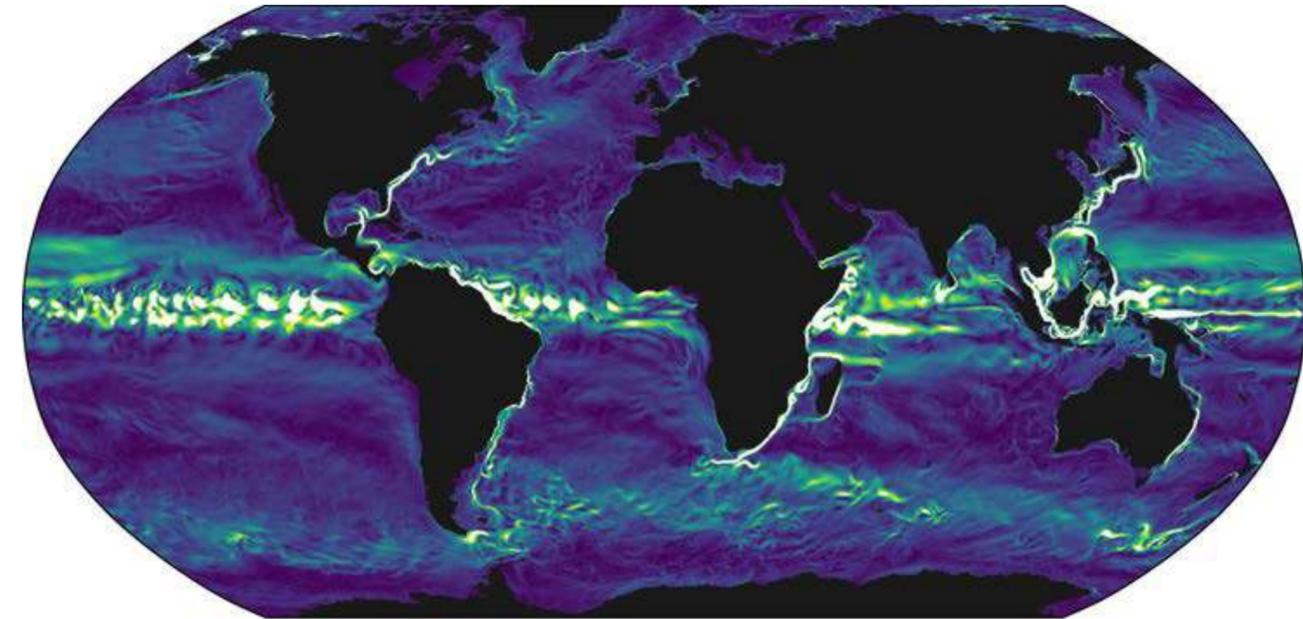
**See also**  
We outline some of our design philosophy and current direction in [this blog post](#).

**START HERE**

- [A short introduction to Veros](#)
  - [The vision](#)
  - [Features](#)
- [Getting started](#)
  - [Installation](#)
  - [Setting up a model](#)
  - [Running Veros](#)
  - [Enhancing Veros](#)
- [Advanced installation](#)
  - [Using JAX](#)

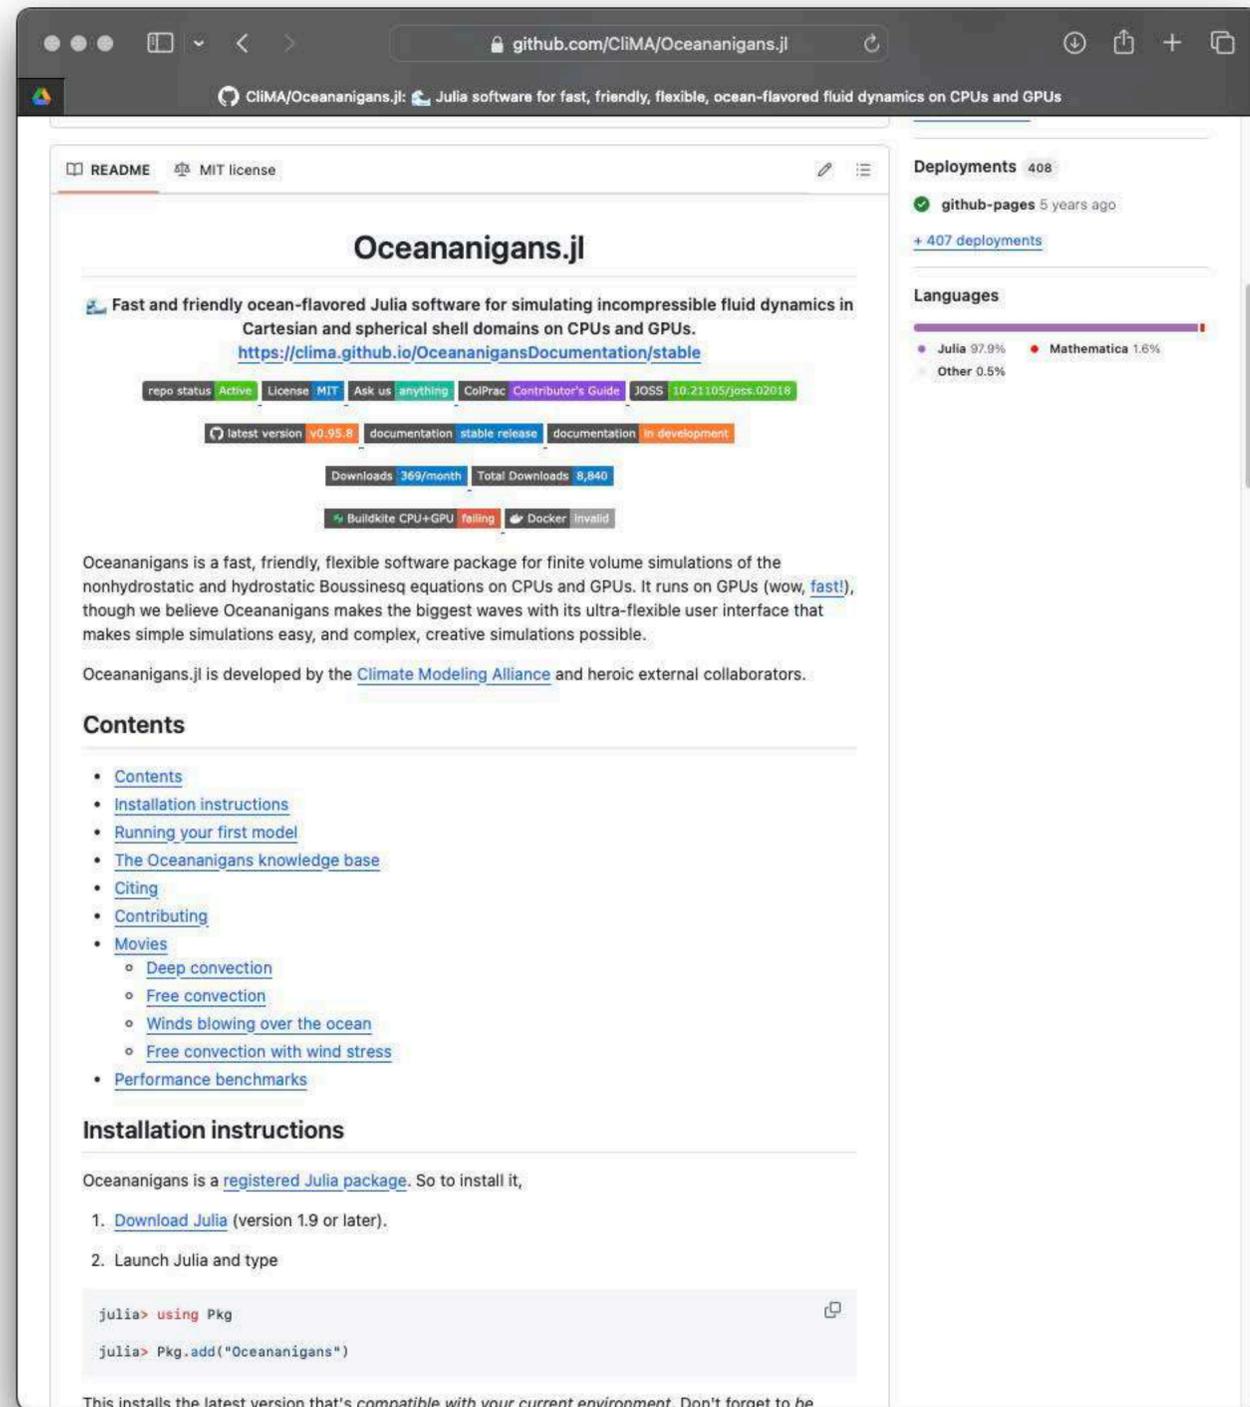


# Veros



Modern code and compute : simple to write, scales, runs on any hardware

# A new generation of geoscientific models



github.com/CliMA/Oceananigans.jl

CliMA/Oceananigans.jl: Julia software for fast, friendly, flexible, ocean-flavored fluid dynamics on CPUs and GPUs

## Oceananigans.jl

Fast and friendly ocean-flavored Julia software for simulating incompressible fluid dynamics in Cartesian and spherical shell domains on CPUs and GPUs.  
<https://clima.github.io/OceananigansDocumentation/stable>

repo status: Active License: MIT Ask us anything ColPrac Contributor's Guide JOSS 10.21105/joss.02018

latest version: v0.95.8 documentation: stable release documentation: in development

Downloads: 369/month Total Downloads: 8,840

Buildkite CPU+GPU: failing Docker: invalid

Oceananigans is a fast, friendly, flexible software package for finite volume simulations of the nonhydrostatic and hydrostatic Boussinesq equations on CPUs and GPUs. It runs on GPUs (wow, *fast!*), though we believe Oceananigans makes the biggest waves with its ultra-flexible user interface that makes simple simulations easy, and complex, creative simulations possible.

Oceananigans.jl is developed by the [Climate Modeling Alliance](#) and heroic external collaborators.

### Contents

- [Contents](#)
- [Installation instructions](#)
- [Running your first model](#)
- [The Oceananigans knowledge base](#)
- [Citing](#)
- [Contributing](#)
- [Movies](#)
  - [Deep convection](#)
  - [Free convection](#)
  - [Winds blowing over the ocean](#)
  - [Free convection with wind stress](#)
- [Performance benchmarks](#)

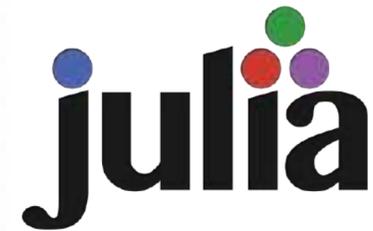
### Installation instructions

Oceananigans is a [registered Julia package](#). So to install it,

1. [Download Julia](#) (version 1.9 or later).
2. Launch Julia and type

```
julia> using Pkg
julia> Pkg.add("Oceananigans")
```

This installs the latest version that's *compatible with your current environment*. Don't forget to be



Ocean

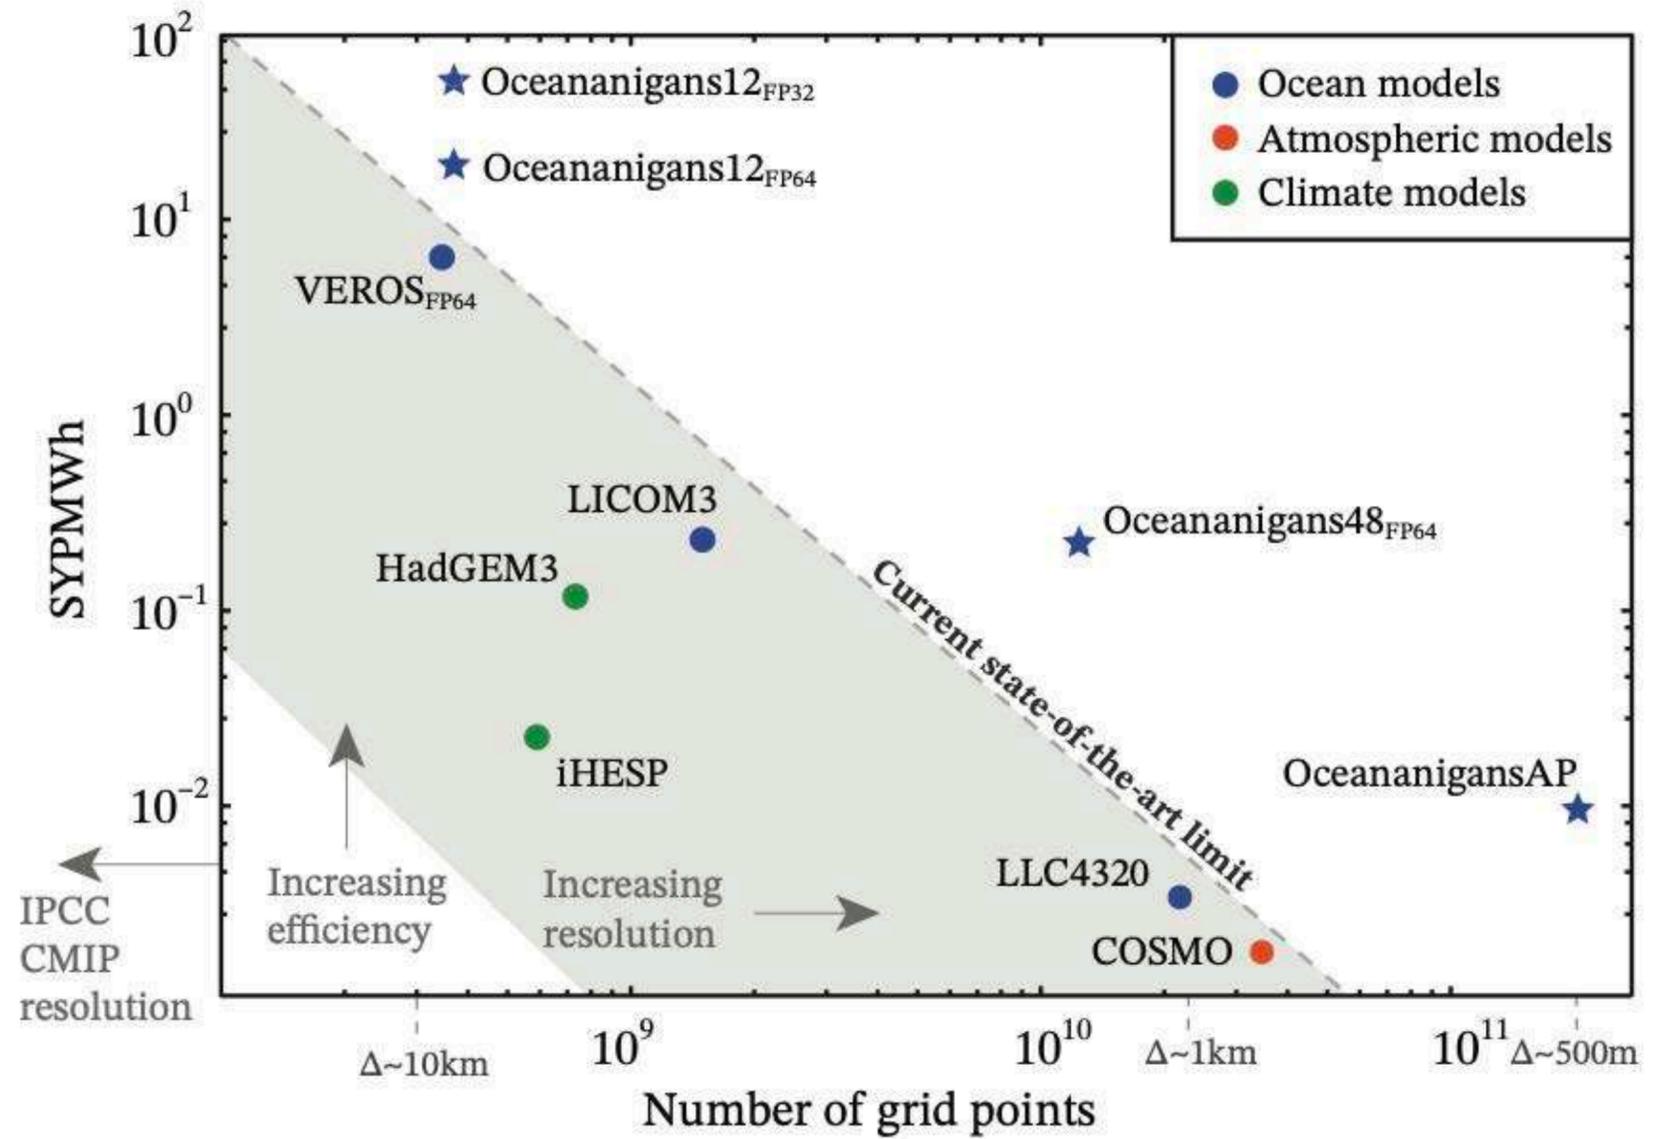


Modern code and compute : simple to write, scales, runs on any hardware

# A new generation of geoscientific models

The screenshot shows the GitHub repository for Oceananigans.jl. The main content area includes a README with a description: "Fast and friendly ocean-flavored Julia software for simulating incompressible fluid dynamics in Cartesian and spherical shell domains on CPUs and GPUs." It also features a list of contents, installation instructions, and a code block showing how to install the package using Julia's package manager.

```
julia> using Pkg
julia> Pkg.add("Oceananigans")
```



Ocean

Modern code and compute : simple to write, scales, runs on any hardware

# A new generation of geoscientific models

The screenshot shows the GitHub repository for Oceananigans.jl. The main heading is "Oceananigans.jl" with a sub-heading: "Fast and friendly ocean-flavored Julia software for simulating incompressible fluid dynamics in Cartesian and spherical shell domains on CPUs and GPUs." It includes a link to the documentation, a badge for "latest version v0.95.8", and download statistics (369/month, 8,840 total). The "Contents" section lists links for installation, running a model, and performance benchmarks. The "Installation instructions" section provides steps to install the package via Julia.



Atmos



Ocean



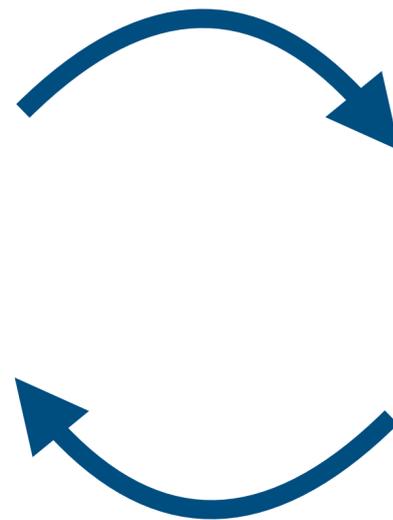
The screenshot shows the GitHub repository for SpeedyWeather.jl. The main heading is "SpeedyWeather.jl" with a sub-heading: "Play atmospheric modelling like it's LEGO." It includes a link to the documentation, a badge for "CI passing", and deployment statistics (500+). The "Dynamics and physics" section lists features like different physical equations and tracer advection. The "Numerics and computing" section lists features like different spatial grids and resolutions. The "User interface" section lists features like data visualization and extensibility. The "Vision and roadmap" section discusses the model's goals and why it was chosen.

Modern code and compute : simple to write, scales, runs on any hardware

# Allowing a **seamless** integration with AI



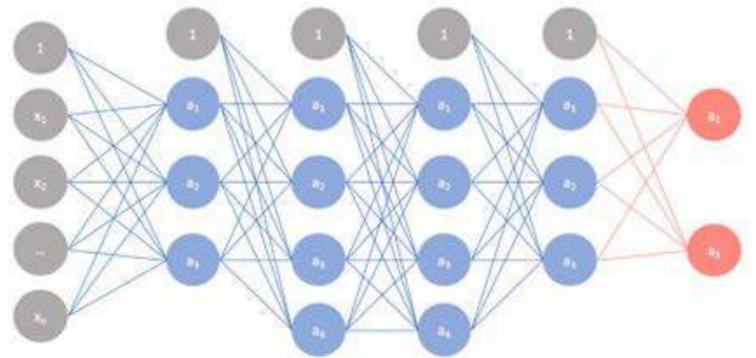
stable, robust, low abstraction languages



high abstraction, fast evolving languages

# Allowing a **seamless** integration with AI

## Data-driven modelling



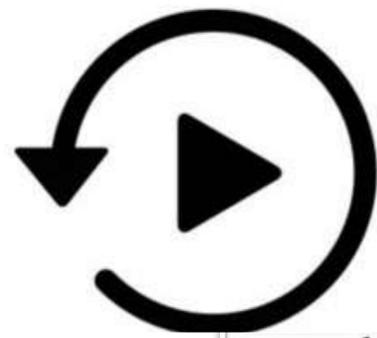
## Physics-based modelling

$$\frac{Du}{Dt} = \frac{uv \tan \phi}{r} - \frac{uw}{r} + fv - f'w - \frac{c_p \theta}{r \cos \phi} \frac{\partial \Pi}{\partial \lambda} + D(u),$$
$$\frac{Dv}{Dt} = -\frac{u^2 \tan \phi}{r} - \frac{vw}{r} - uf - \frac{c_p \theta}{r} \frac{\partial \Pi}{\partial \phi} + D(v),$$
$$\delta \frac{Dw}{Dt} = \frac{u^2 + v^2}{r} + uf' - g(r) - c_p \theta \frac{\partial \Pi}{\partial r},$$



high abstraction, fast evolving languages

high abstraction, fast evolving languages



# AI-native hybrid climate models ?



## Neural general circulation models for weather and climate

https://doi.org/10.1038/s41586-024-07744-y  
Received: 13 November 2023  
Accepted: 15 June 2024  
Published online: 22 July 2024  
Open access  
Check for updates

Dmitrii Kochkov<sup>1,2,3,4</sup>, Janni Yuval<sup>1,4,5,6</sup>, Ian Langmore<sup>1,4</sup>, Peter Norgaard<sup>1,4</sup>, Jamie Smith<sup>1,4</sup>, Griffin Mooers<sup>1</sup>, Milan Klöwer<sup>7</sup>, James Lottes<sup>1</sup>, Stephan Rasp<sup>1</sup>, Peter Düben<sup>1</sup>, Sam Hatfield<sup>1</sup>, Peter Battaglia<sup>1</sup>, Alvaro Sanchez-Gonzalez<sup>1</sup>, Matthew Willson<sup>1</sup>, Michael P. Brenner<sup>1,4</sup> & Stephan Hoyer<sup>1,6,8,9</sup>

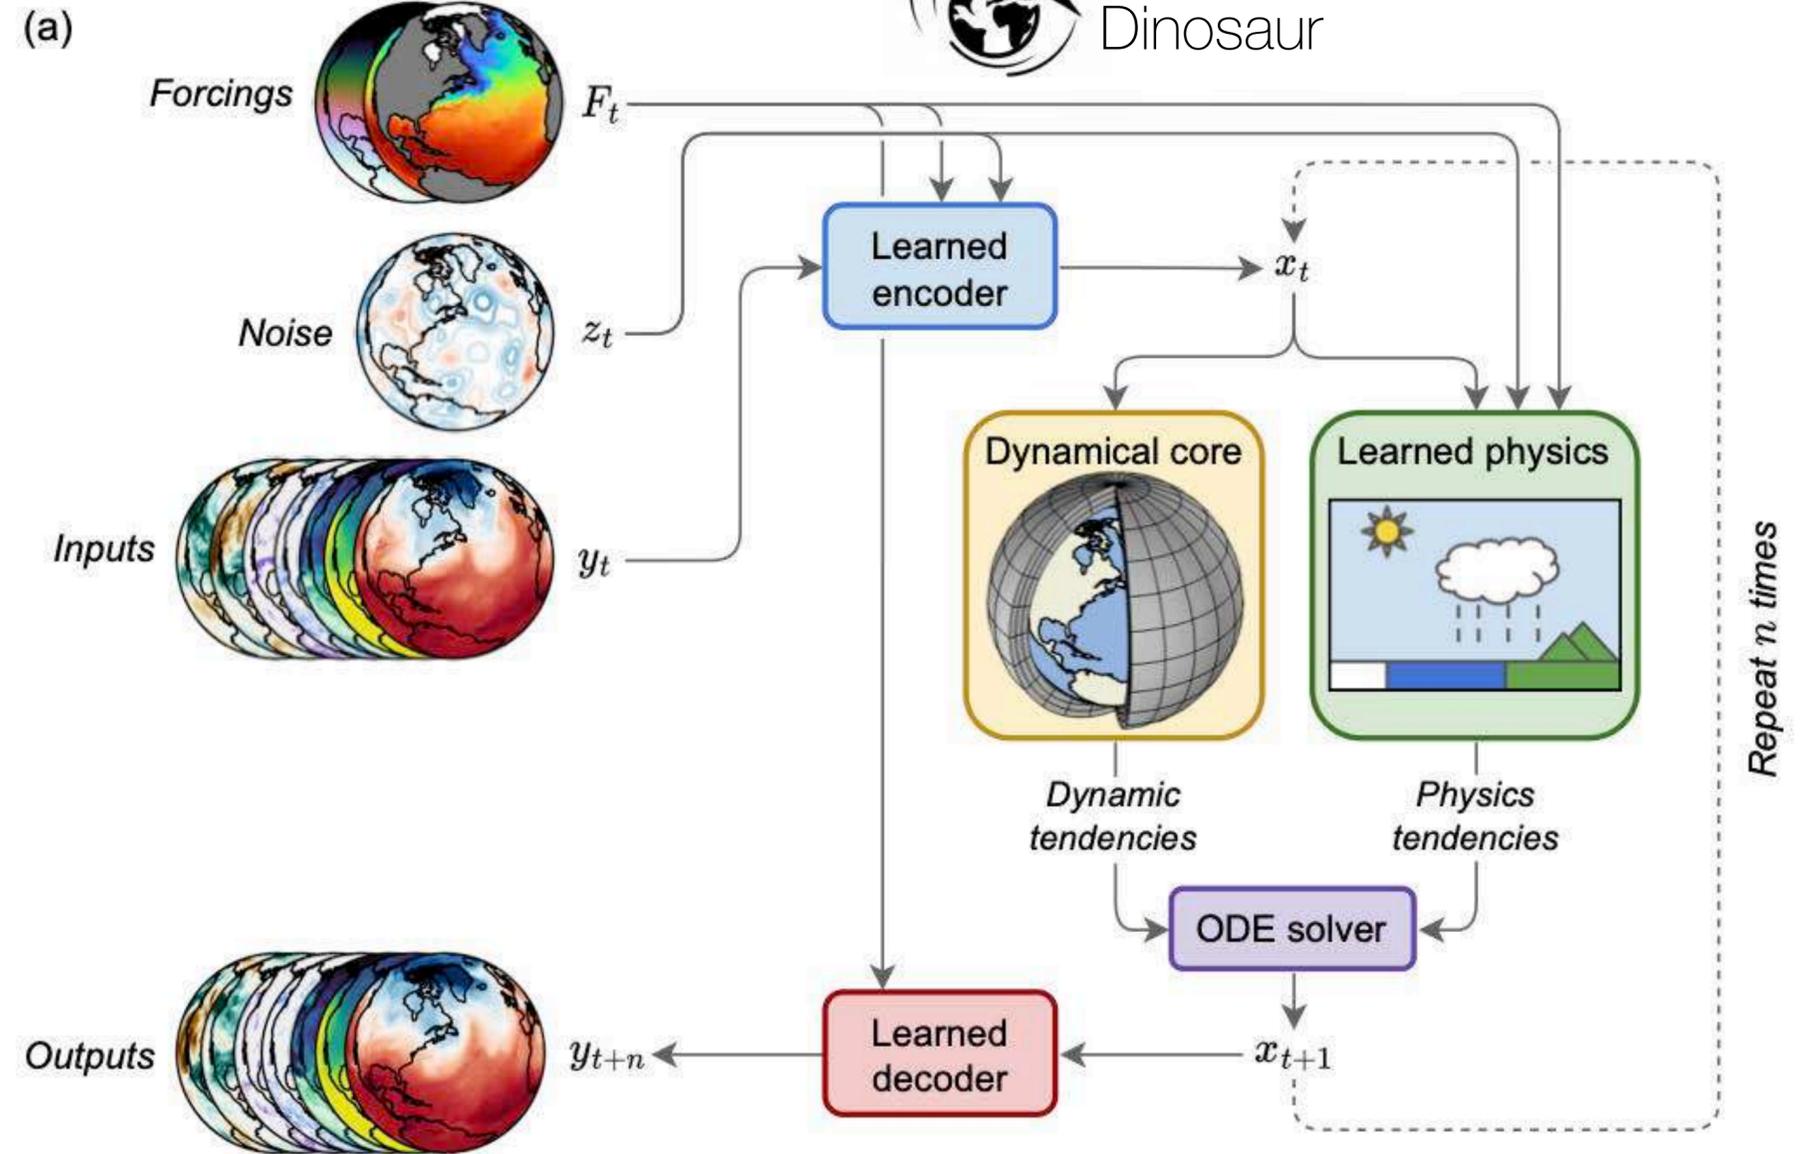
General circulation models (GCMs) are the foundation of weather and climate prediction<sup>1,2</sup>. GCMs are physics-based simulators that combine a numerical solver for large-scale dynamics with tuned representations for small-scale processes such as cloud formation. Recently, machine-learning models trained on reanalysis data have achieved comparable or better skill than GCMs for deterministic weather forecasting<sup>3,4</sup>. However, these models have not demonstrated improved ensemble forecasts, or shown sufficient stability for long-term weather and climate simulations. Here we present a GCM that combines a differentiable solver for atmospheric dynamics with machine-learning components and show that it can generate forecasts of deterministic weather, ensemble weather and climate on par with the best machine-learning and physics-based methods. NeuralGCM is competitive with machine-learning models for one- to ten-day forecasts, and with the European Centre for Medium-Range Weather Forecasts ensemble prediction for one- to fifteen-day forecasts. With prescribed sea surface temperature, NeuralGCM can accurately track climate metrics for multiple decades, and climate forecasts with 140-kilometre resolution show emergent phenomena such as realistic frequency and trajectories of tropical cyclones. For both weather and climate, our approach offers orders of magnitude computational savings over conventional GCMs, although our model does not extrapolate to substantially different future climates. Our results show that end-to-end deep learning is compatible with tasks performed by conventional GCMs and can enhance the large-scale physical simulations that are essential for understanding and predicting the Earth system.

Solving the equations for Earth's atmosphere with general circulation models (GCMs) is the basis of weather and climate prediction<sup>1,2</sup>. Over the past 70 years, GCMs have been steadily improved with better numerical methods and more detailed physical models, while exploiting faster computers to run at higher resolution. Inside GCMs, the unresolved physical processes such as clouds, radiation and precipitation are represented by semi-empirical parameterizations. Tuning GCMs to match historical data remains a manual process<sup>3</sup>, and GCMs retain many persistent errors and biases<sup>4</sup>. The difficulty of reducing uncertainty in long-term climate projections<sup>5</sup> and estimating distributions of extreme weather events<sup>6</sup> presents major challenges for climate mitigation and adaptation<sup>7</sup>.

Recent advances in machine learning have presented an alternative for weather forecasting<sup>3,4,8,9</sup>. These models rely solely on machine-learning techniques, using roughly 40 years of historical data from the European Centre for Medium-Range Weather Forecasts (ECMWF) reanalysis v5 (ERAS)<sup>10</sup> for model training and forecast initialization. Machine-learning methods have been remarkably successful,

demonstrating state-of-the-art deterministic forecasts for 1- to 10-day weather prediction at a fraction of the computational cost of traditional models<sup>3,4</sup>. Machine-learning atmospheric models also require considerably less code, for example GraphCast<sup>3</sup> has 5,417 lines versus 376,578 lines for the National Oceanic and Atmospheric Administration's FV3 atmospheric model<sup>11</sup> (see Supplementary Information section A for details).

Nevertheless, machine-learning approaches have noteworthy limitations compared with GCMs. Existing machine-learning models have focused on deterministic prediction, and surpass deterministic numerical weather prediction in terms of the aggregate metrics for which they are trained<sup>3,4</sup>. However, they do not produce calibrated uncertainty estimates<sup>1</sup>, which is essential for useful weather forecasts<sup>1</sup>. Deterministic machine-learning models using a mean-squared-error loss are rewarded for averaging over uncertainty, producing unrealistically blurry predictions when optimized for multi-day forecasts<sup>3,4</sup>. Unlike physical models, machine-learning models misrepresent derived (diagnostic) variables such as geostrophic wind<sup>12</sup>. Furthermore,



<https://doi.org/10.1038/s41586-024-07744-y>



<https://github.com/google-research/dinosaur>

<https://github.com/google-research/neuralgcm>

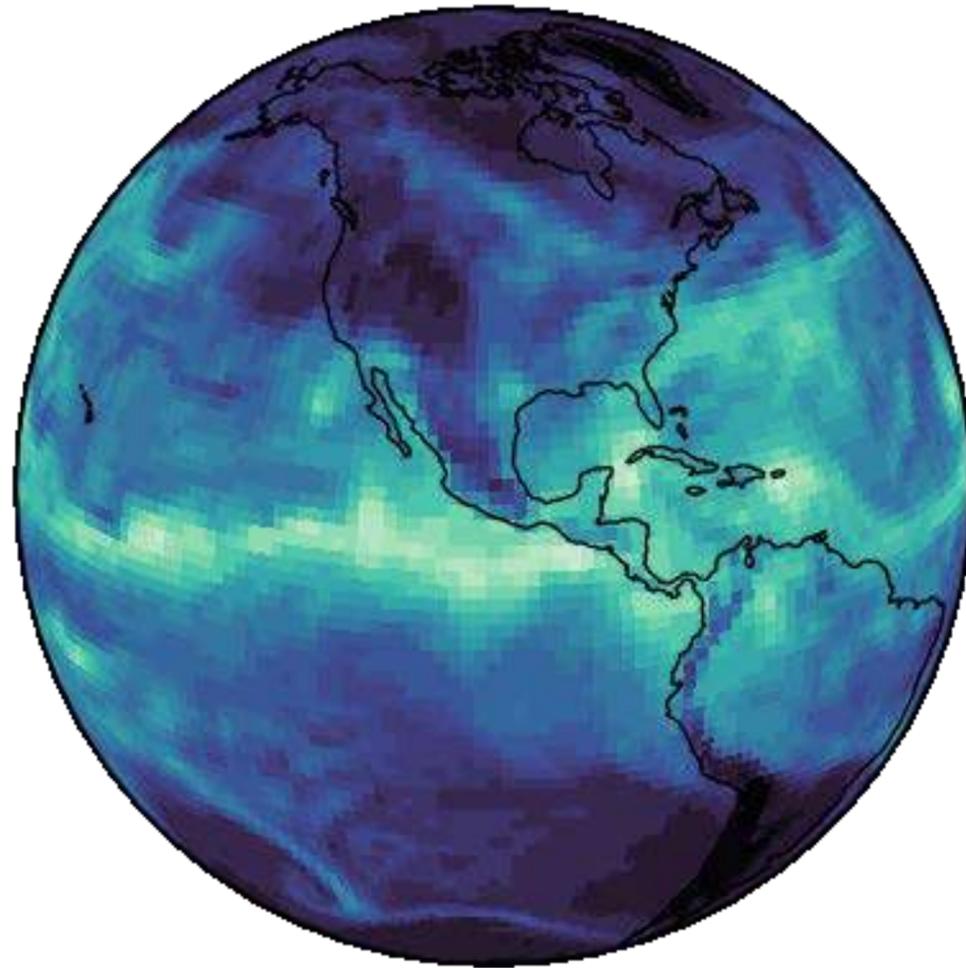
Kochkov et al. (2024)

# AI-native hybrid climate models ?

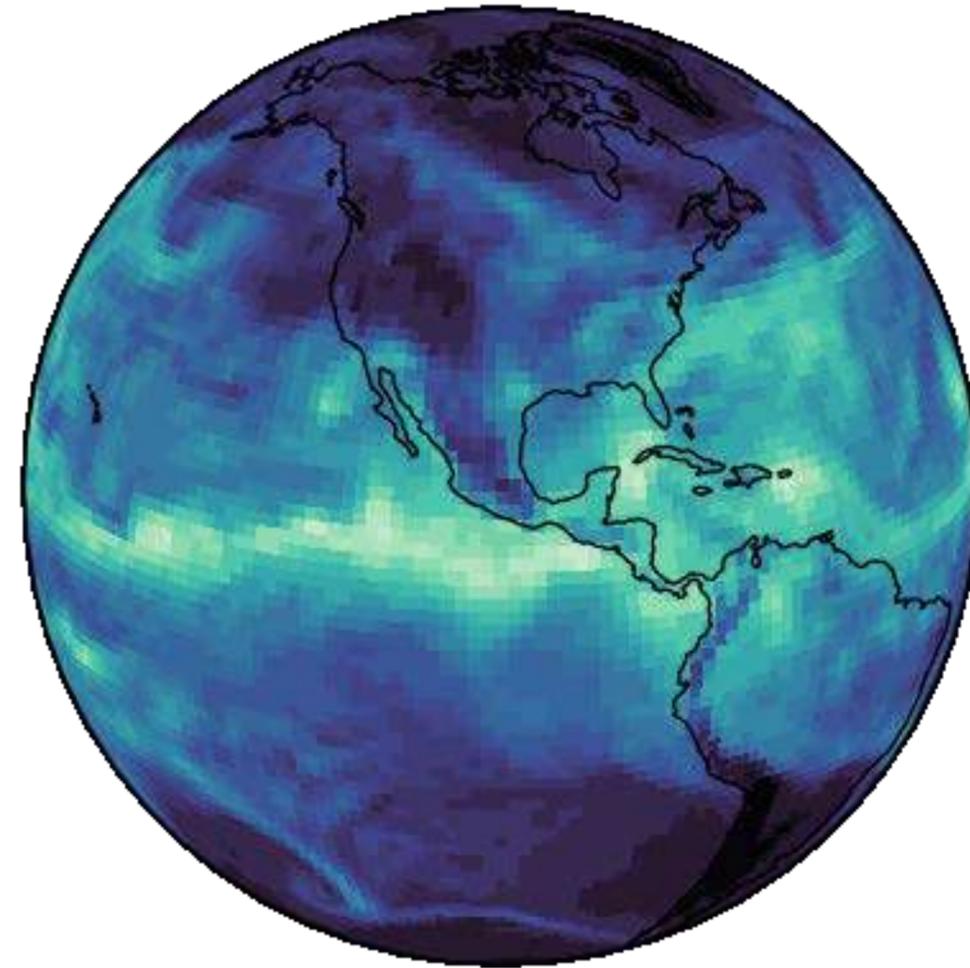
*Kochkov et al. (2024)*

<https://arxiv.org/abs/2311.07222>

Total column water, 0-15 days



ERA5



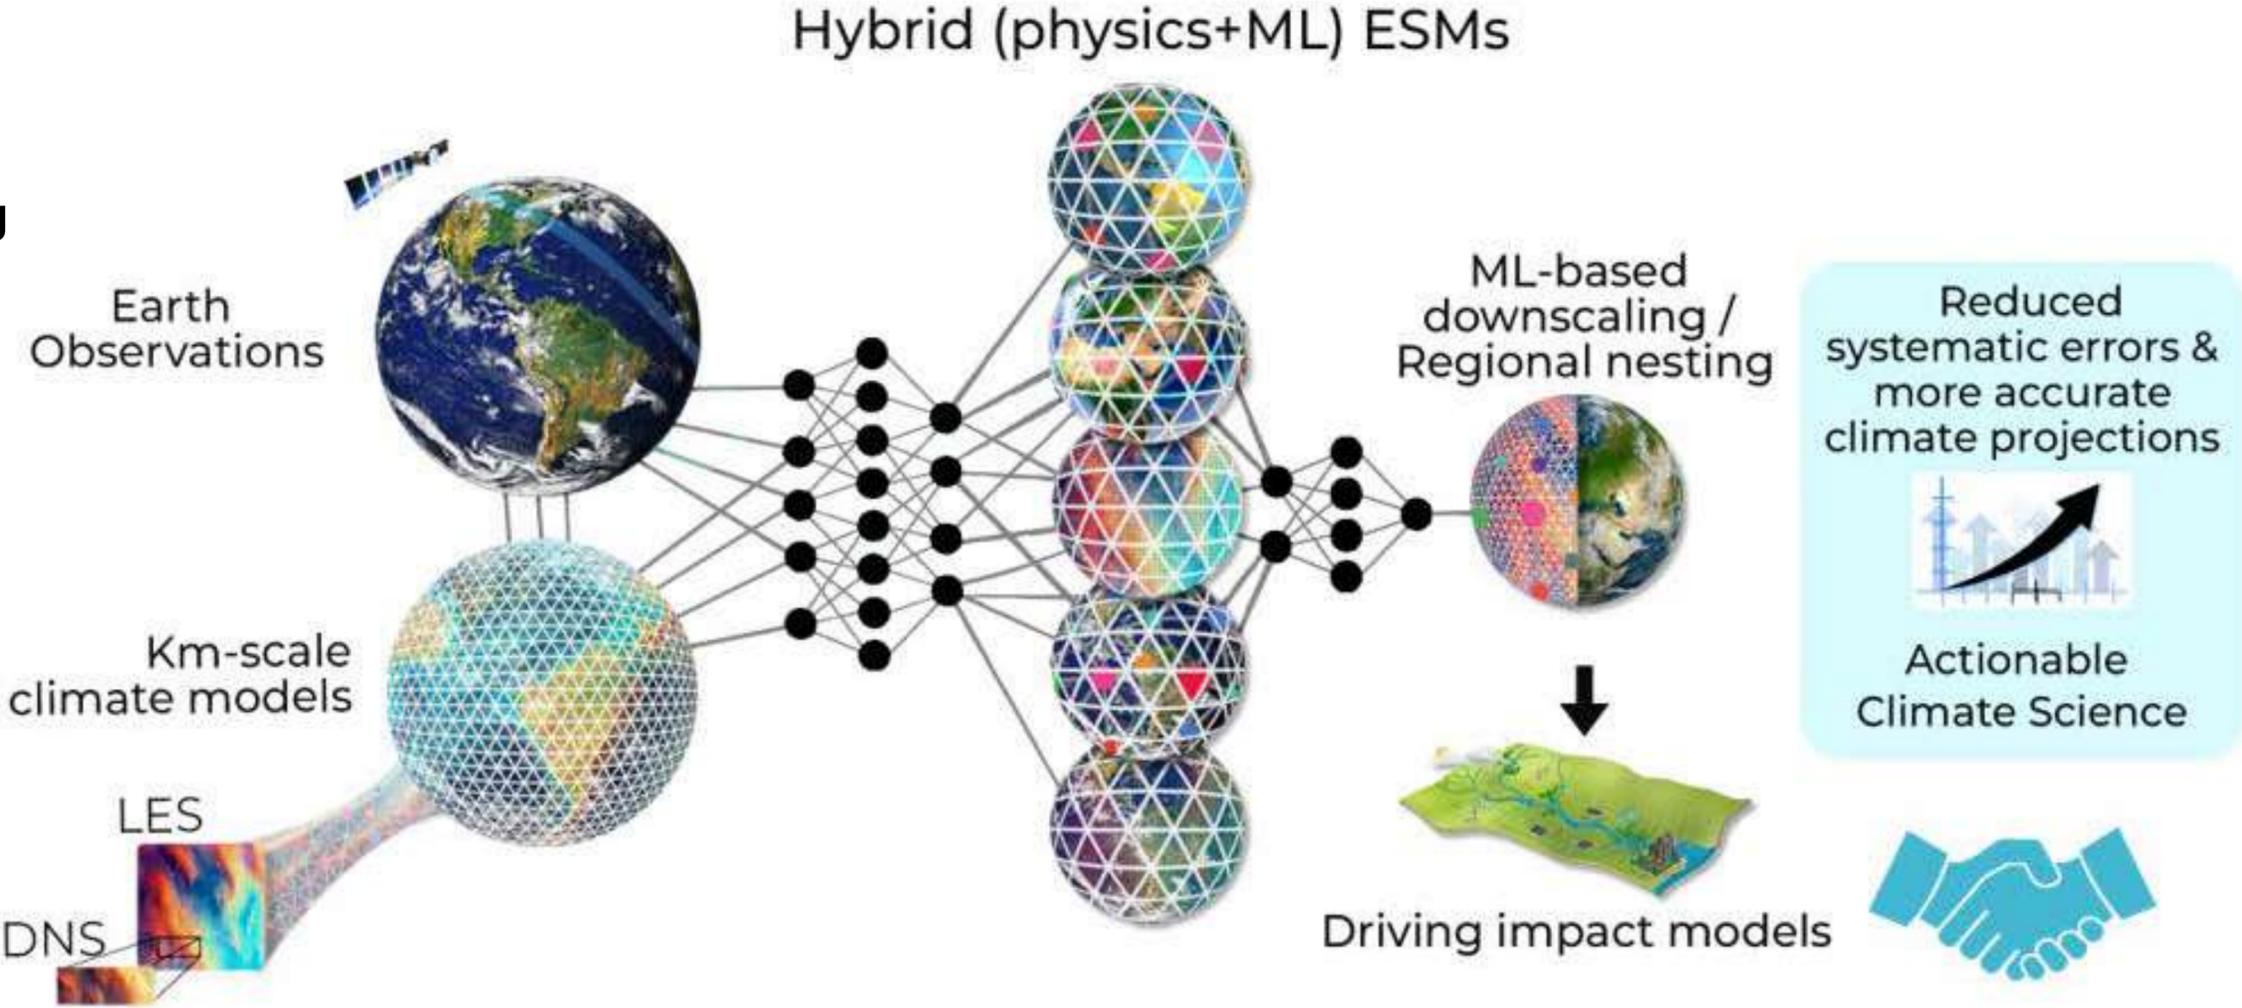
NeuralGCM

**Hybrid w/ online : non-blurry forecast + stable simulators (runs ~10 years)**



# AI-native hybrid climate models ?

See P. Gentine's keynote during TRACCS General Assembly



Eyring, Gentine, et al., 2024 <https://doi.org/10.1038/s41561-024-01527-w>

**LEGO land :**  
Versatile models allowing exploration of advanced ML workflows  
**Harnessing global observations !**

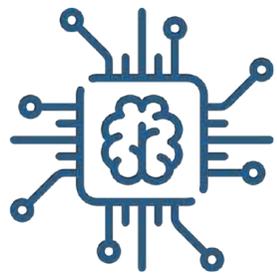
## Differentiable programming in earth system models

# Summary



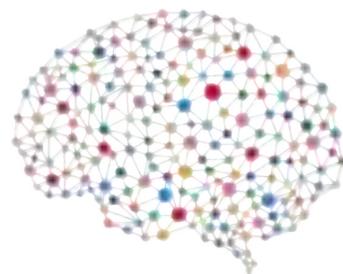
- Why we are **augmenting ocean models** w/ trained comp.

=



- Described how this can be **done in practice** today

+



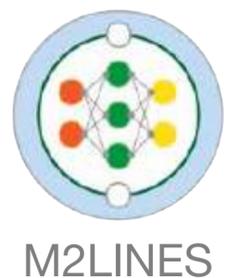
- Further progress require **deeper recast** of our models

+



- Exploration w/ new generation of **AI-native hybrid** models

- Exciting time for **cross-disciplinary** investigations !



# References

- Frezat et al. (2021). Physical invariance in neural networks for subgrid-scale scalar flux modeling. <https://doi.org/10.1103/PhysRevFluids.6.024607>
- Frezat et al. (2022) A posteriori learning for quasi-geostrophic turbulence parametrization, JAMES. 14. <https://doi.org/10.1029/2022MS003124>
- Frezat et al. (2024) Gradient-free online learning of subgrid-scale dynamics with neural emulators, sub., <https://doi.org/10.48550/arXiv.2310.19385>
- Yan (2024). Adjoint-based online learning of two-layer quasi-geostrophic baroclinic turbulence. arXiv. <https://doi.org/10.48550/arXiv.2411.14106>