

# Opening the Black Box of Ocean Wave Learning

Felipe C. Minuzzi and Leandro Farina

Institute of Mathematics and Statistics, Federal University of Rio Grande do Sul (UFRGS), Brazil

felipe.minuzzi@ufrgs.br, farina@alum.mit.edu

## Introduction

- Operational models remain dependent on parametrization to close the wave action balance.
- Data-driven approaches deliver high skill but raise questions about robustness and physical interpretability.
- We introduce a two-stage workflow to balance predictive skill and physical transparency.

## Methodology & Data

**Our goal:** Map oceanographic physical variables states to the dimensionless significant wave height  $y = gH_s/U_{10}^2$ .

- **Data Source:** ERA5 Reanalysis (2018-2022 for training, 2023 for out-of-sample testing).
- **Domain:** South Atlantic.
- **The Black Box (CNN):** A lightweight feature-mixing network using pointwise ( $1 \times 1$ ) convolutions across 21 dimensionless features, with Integrated Gradients (IG) to extract feature attributions and isolate physical drivers.
- **The White Box (SR):** *SymbWaves* (built on PySR) uses genetic algorithms to search mathematical expression space ( $+$ ,  $-$ ,  $\times$ ,  $\div$ , exponentiation,  $\log$ ) for parsimonious physical laws.

## The Black Box: CNN & Explainability

### Predictive Benchmark & Domain Shift

- CNN achieves an overall MAPE of 21.04%.

### Integrated Gradients (IG) Analysis

- **Primary Driver:** Wave age ( $\beta = c_p/U_{10}$ ) dominates attributions.
- **Secondary Drivers:** Directional predictors (wind direction, swell partitions).

*The CNN robustly identifies relevant drivers, but fails to transfer the functional relationships across basins.*

## The White Box: Discovering Symbolic Laws

Guided by the CNN's feature reliance, we segmented the system by wave age ( $\beta$ ) to capture distinct physical regimes.

**1. Wind-Sea: Basin-Stable Linear Law** For young waves ( $\beta \leq 1.3$ ), SR consistently recovered a linear relationship:

$$y = C_w \beta \quad \text{where } C_w \approx 0.20$$

This law showed remarkable stability, independently validating a scaled form of Toba's growth law within the ERA5 data.

**2. Swell: Universal Scaling Law (USSL)** In the stable swell regime ( $2 \leq \beta < 20$ ), SR yielded a near-universal power law:

$$y = C_s \beta^\alpha \quad \text{with } \alpha \approx 1.8, C_s \approx 0.074$$

**3. Directionally-Modulated Swell Law (DMSL)** Allowing interactions with directional features produced a higher-accuracy variant:

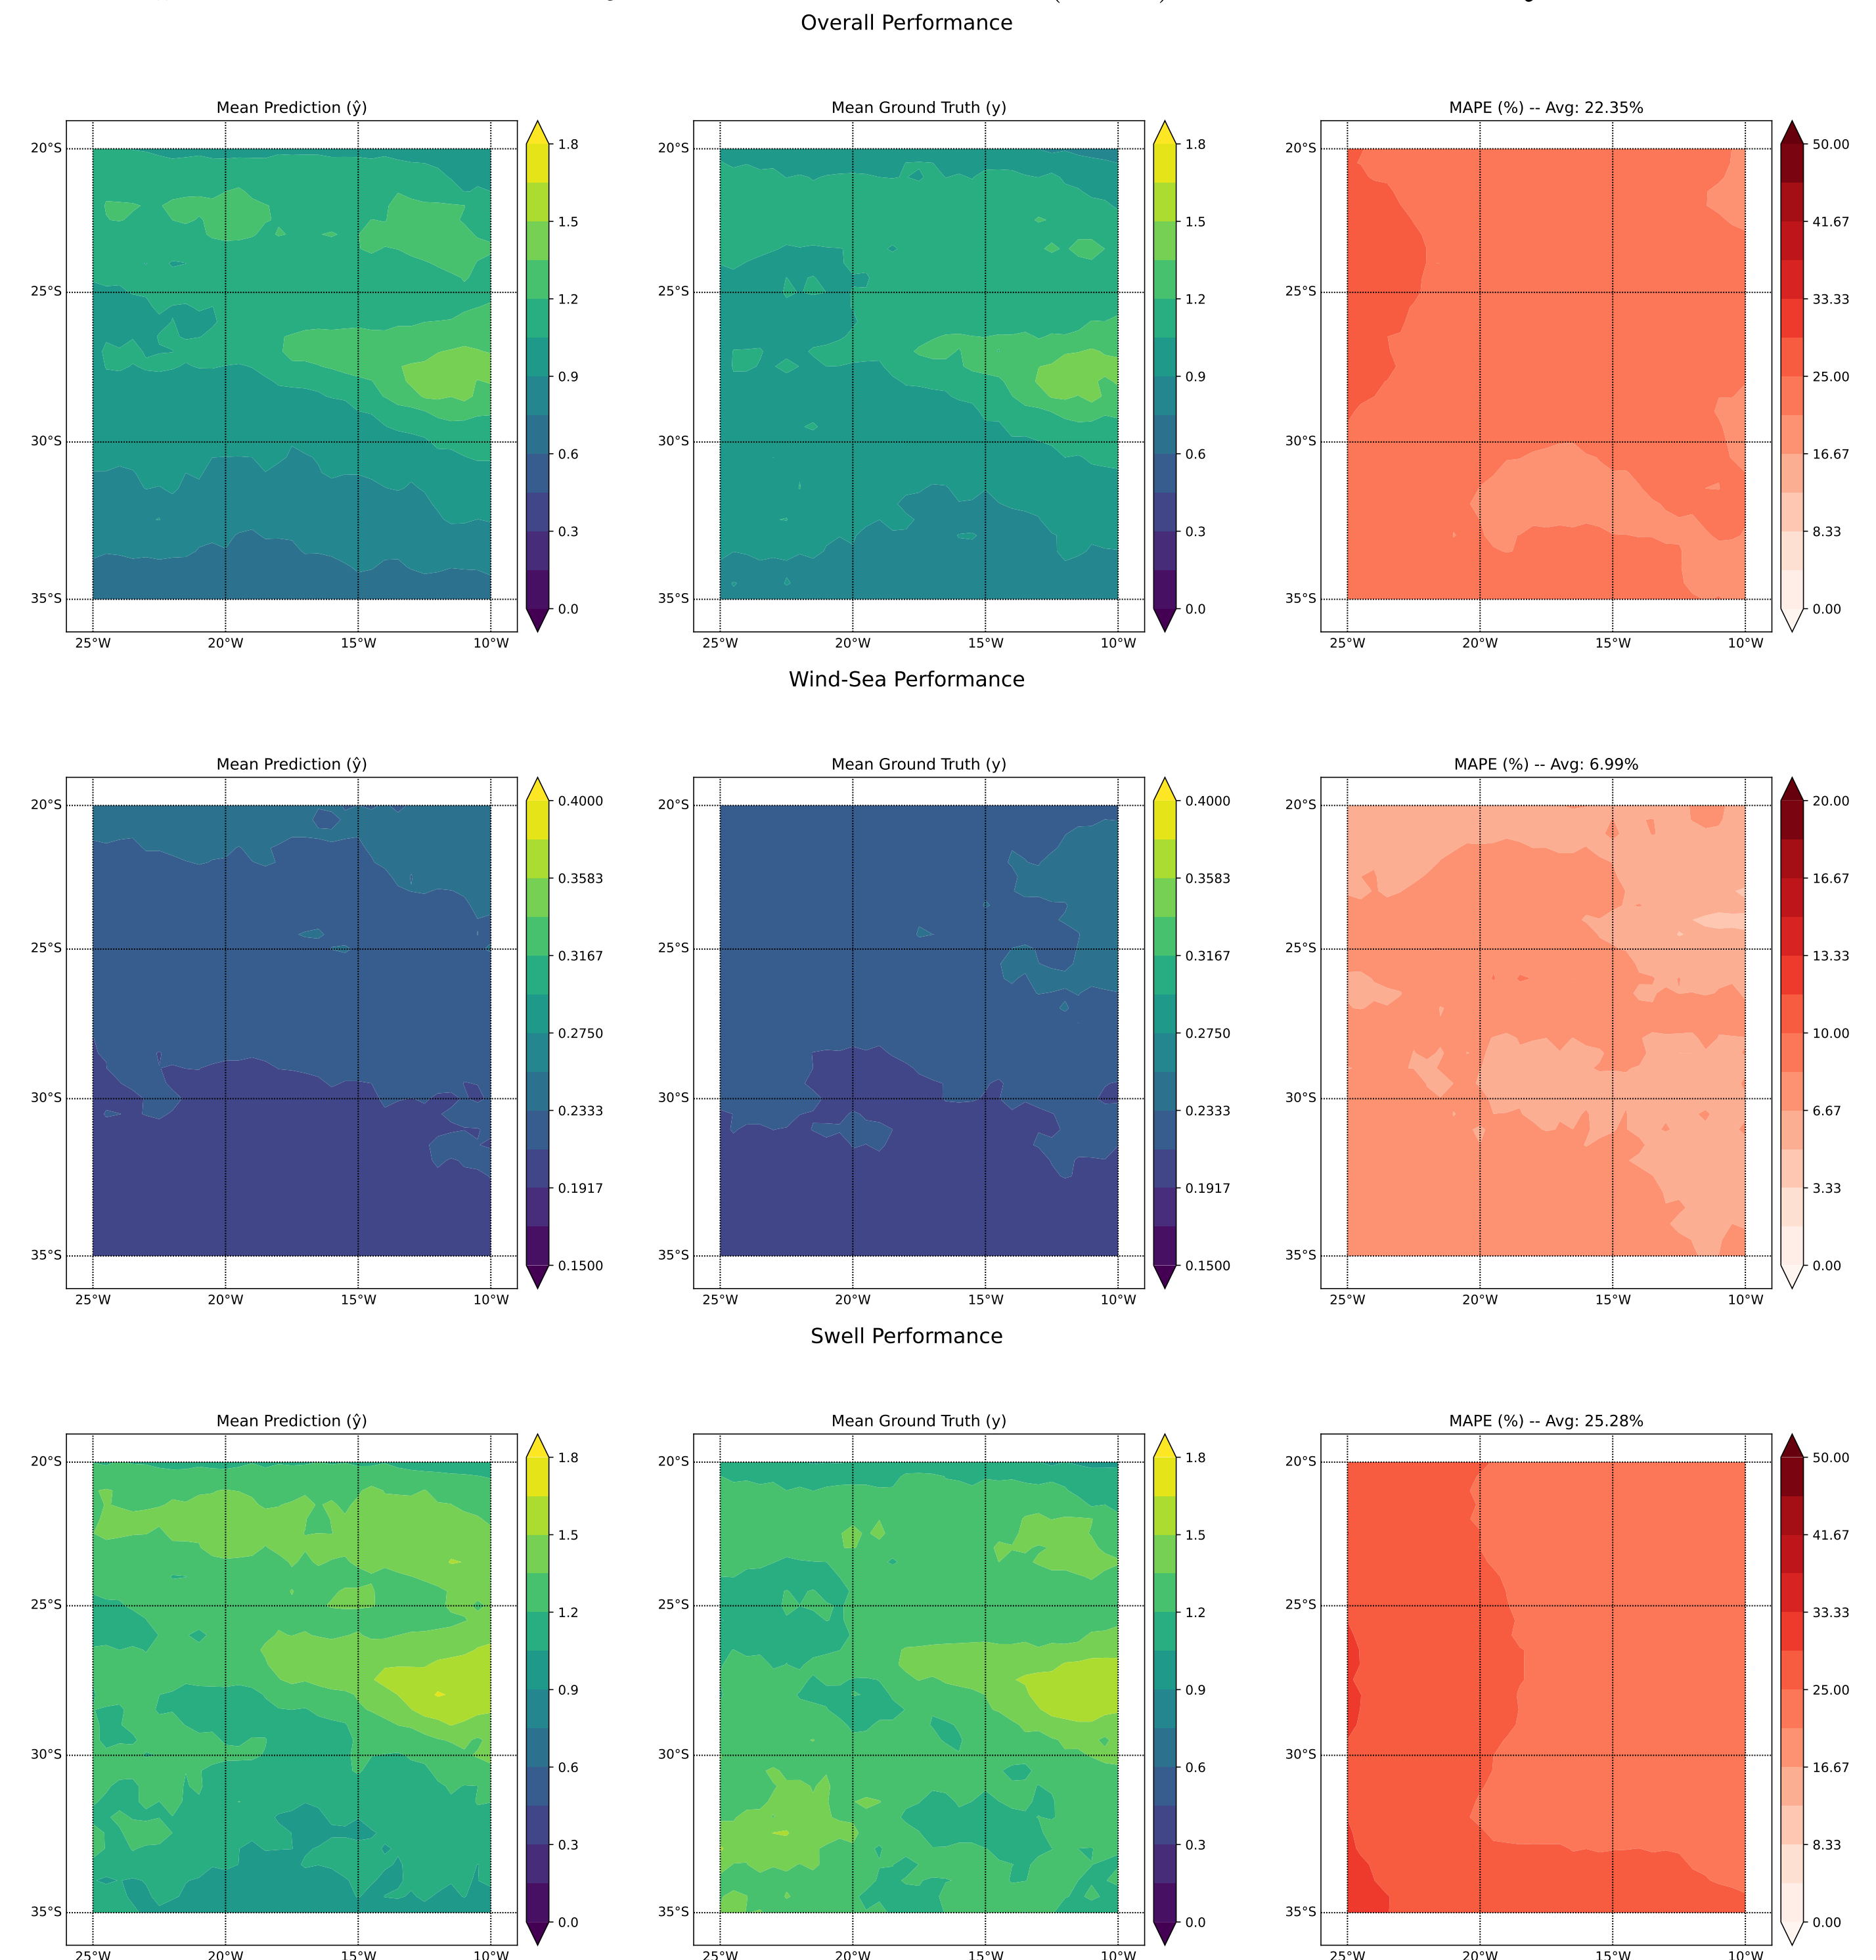
$$y = (C_m \beta) \sqrt{A \pm \cos(\theta_{swell})}$$

**Quantitative Performance:** Overall ( $MAPE_\Omega$ ) 24.32% (Wind-sea: 17.65%, Swell: 25.13%)

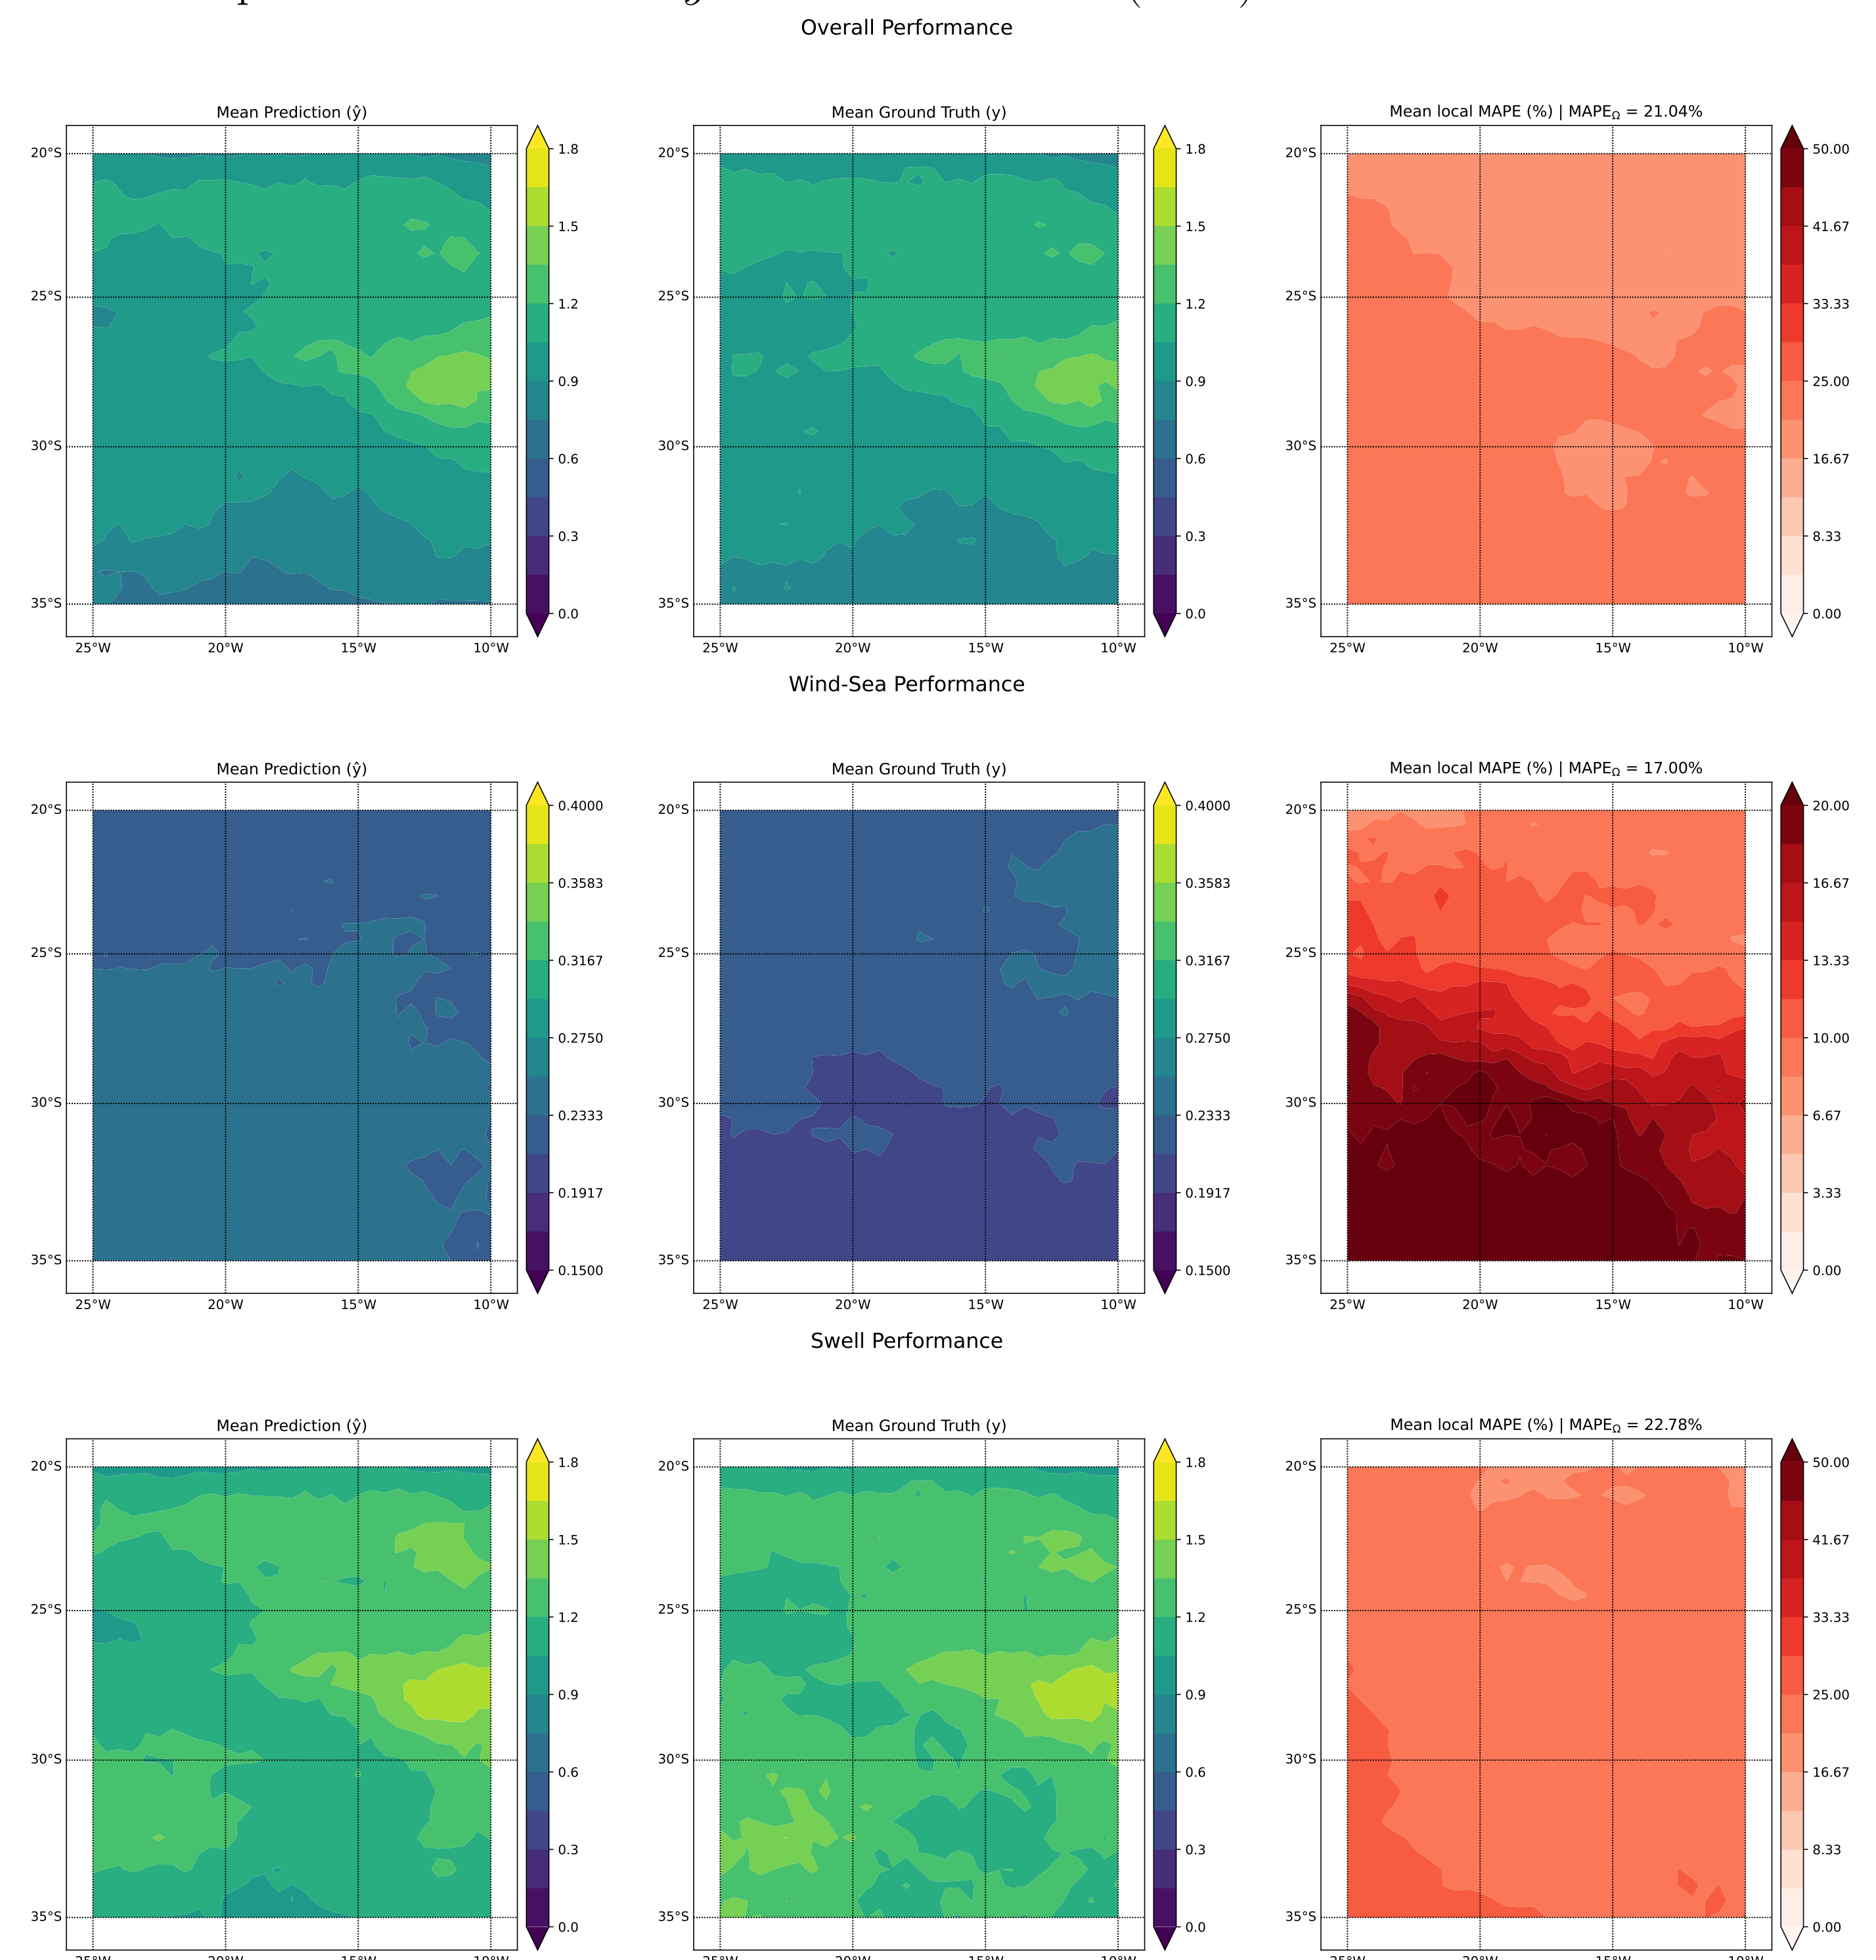
*Symbolic Regression distills universal, transferable laws that outperform deep learning in out-of-basin scenarios.*

## Results

Spatial distribution of  $y$  in the test domain (2023) obtained with *SymbWaves*.



Spatial distribution of  $y$  in the test domain (2023) obtained with CNN.



## Conclusions

- **Trade-off:** Deep learning excels at localized accuracy, but SR provides unparalleled stability and wind-sea generalizability.
- **The Hybrid Advantage:** Using a black-box to identify features followed by SR to distill them creates interpretable and transferable ESM parameterizations.