Generative Models For Spatial-Based Bottom Hypoxia Forecasting In The Black Sea (GEMSS 2024)

Victor Mangeleer, Marilaure Grégoire and Gilles Louppe

University of Liège

Let's keep in touch !

Introduction

Can we **accurately forecast** oceanic oxygen levels using only **Earth observations**?

Motivation

1. Forecasting with Neural Networks: A Spatial-Based Neural Network can provide near real-time predictions using satellite imagery data (Ex: Sentinel-3, ...).
2. Early detection and prevention of hypoxia: Accurate forecasts can detect hypoxia ($O_2 < 63\ \mathrm{mmol/m^3}$) several days in advance, enabling preventive actions to protect marine and human communities.

Problem Statement

Given a set of Earth Observations X (such as SST, SSH, etc.), can we infer the oxygen concentration levels y at the bottom of the Black Sea continental shelf ? In other words, can we approximate $\mathrm{P(y|X)}$?

Solutions

The oxygen distribution $\mathrm{P(y)}$ is mostly unimodal but becomes bimodal in summer, indicating low oxygen concentration levels. To approximate our likelihood $\mathrm{P(y|X)}$, we choose:

1. Regression (most common approach in physical sciences)

$$\mathrm{P(y|X)} \approx \mathcal{N}(\mu, \sigma^2)$$

2. Mixture Density Network (statistical prediction)

$$\mathrm{P(y|X)} \approx \sum \pi_k \mathcal{N}(\mu_k\ \sigma_k^2)$$

Visualization

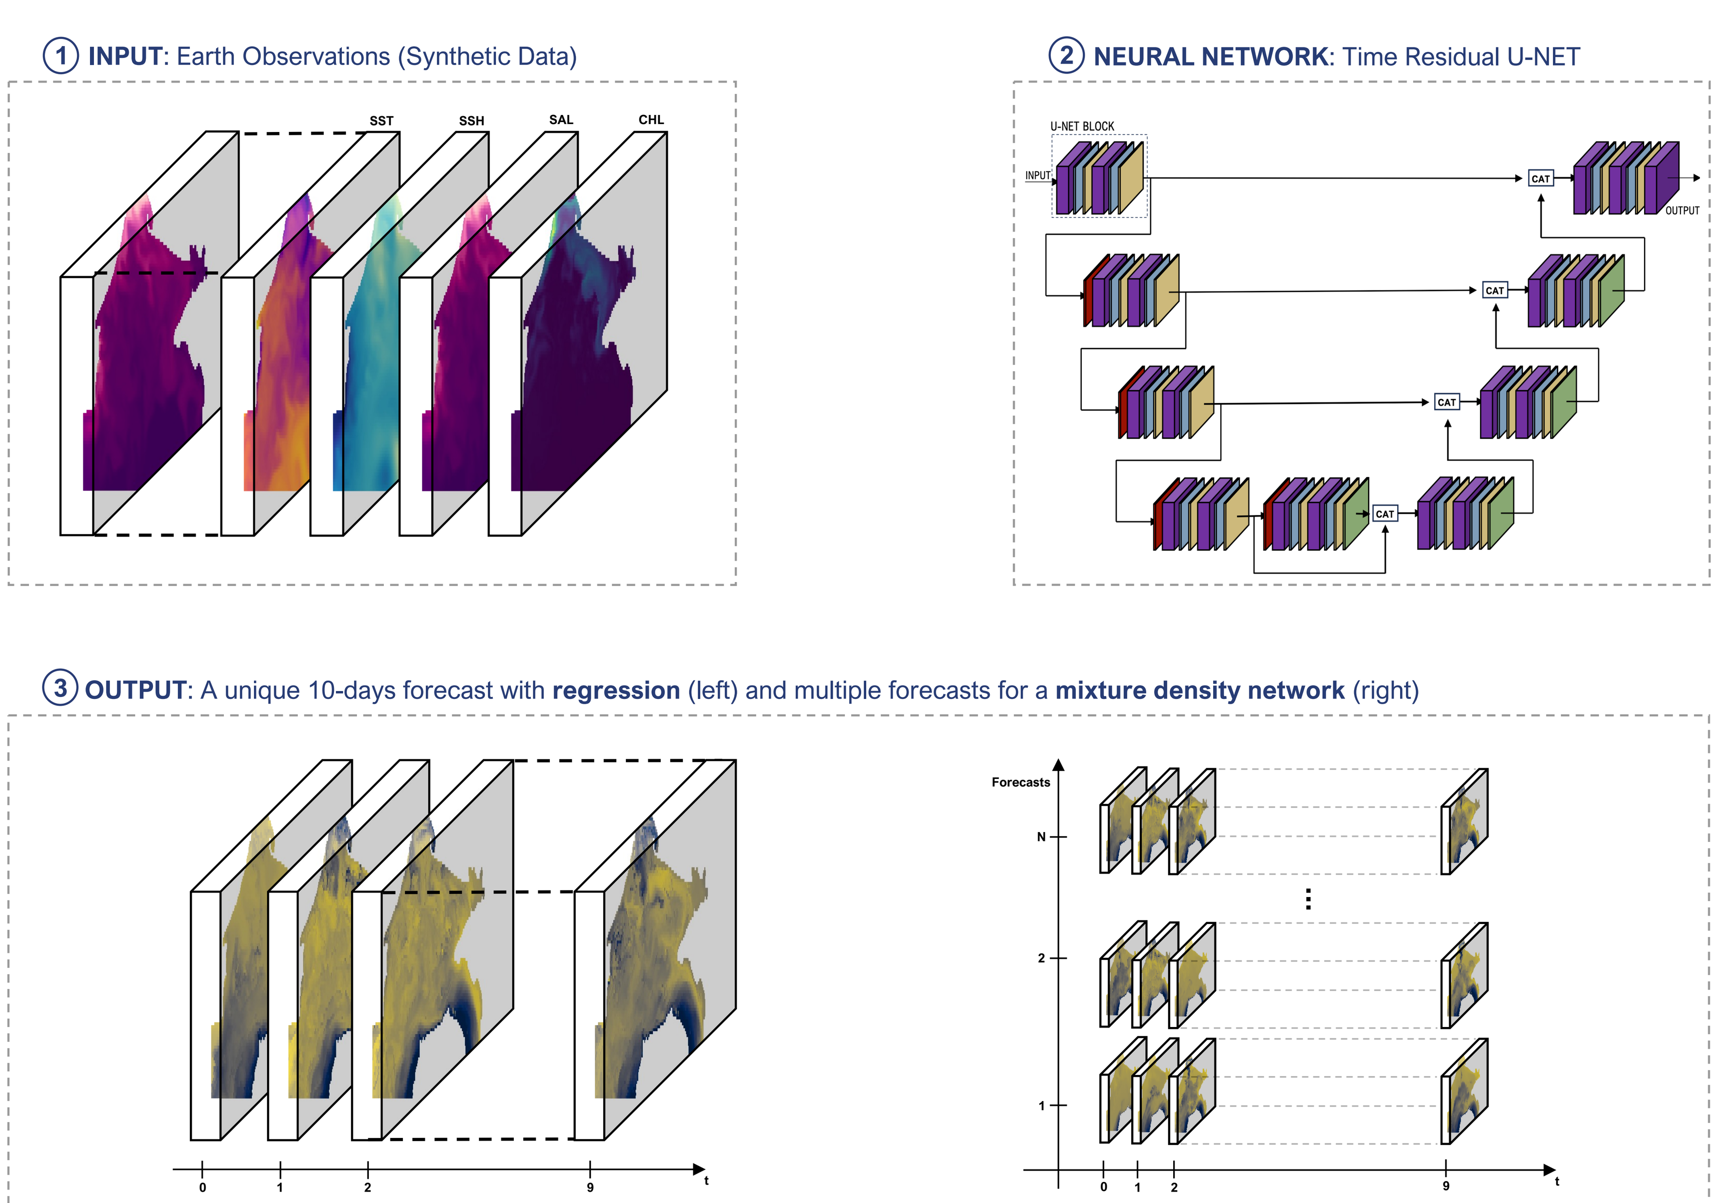

Figure 1: Visual representation of the training pipeline: First, the surface observations are standardized and stacked according to the number of prior days given as input. Next, the data is fed into a time residual U-NET, which adjusts its channels based on the date it must make predictions. Finally, the output is a single forecast for the regression model, while for the mixture density network, multiple forecasts are generated by sampling the approximated likelihood.

Results

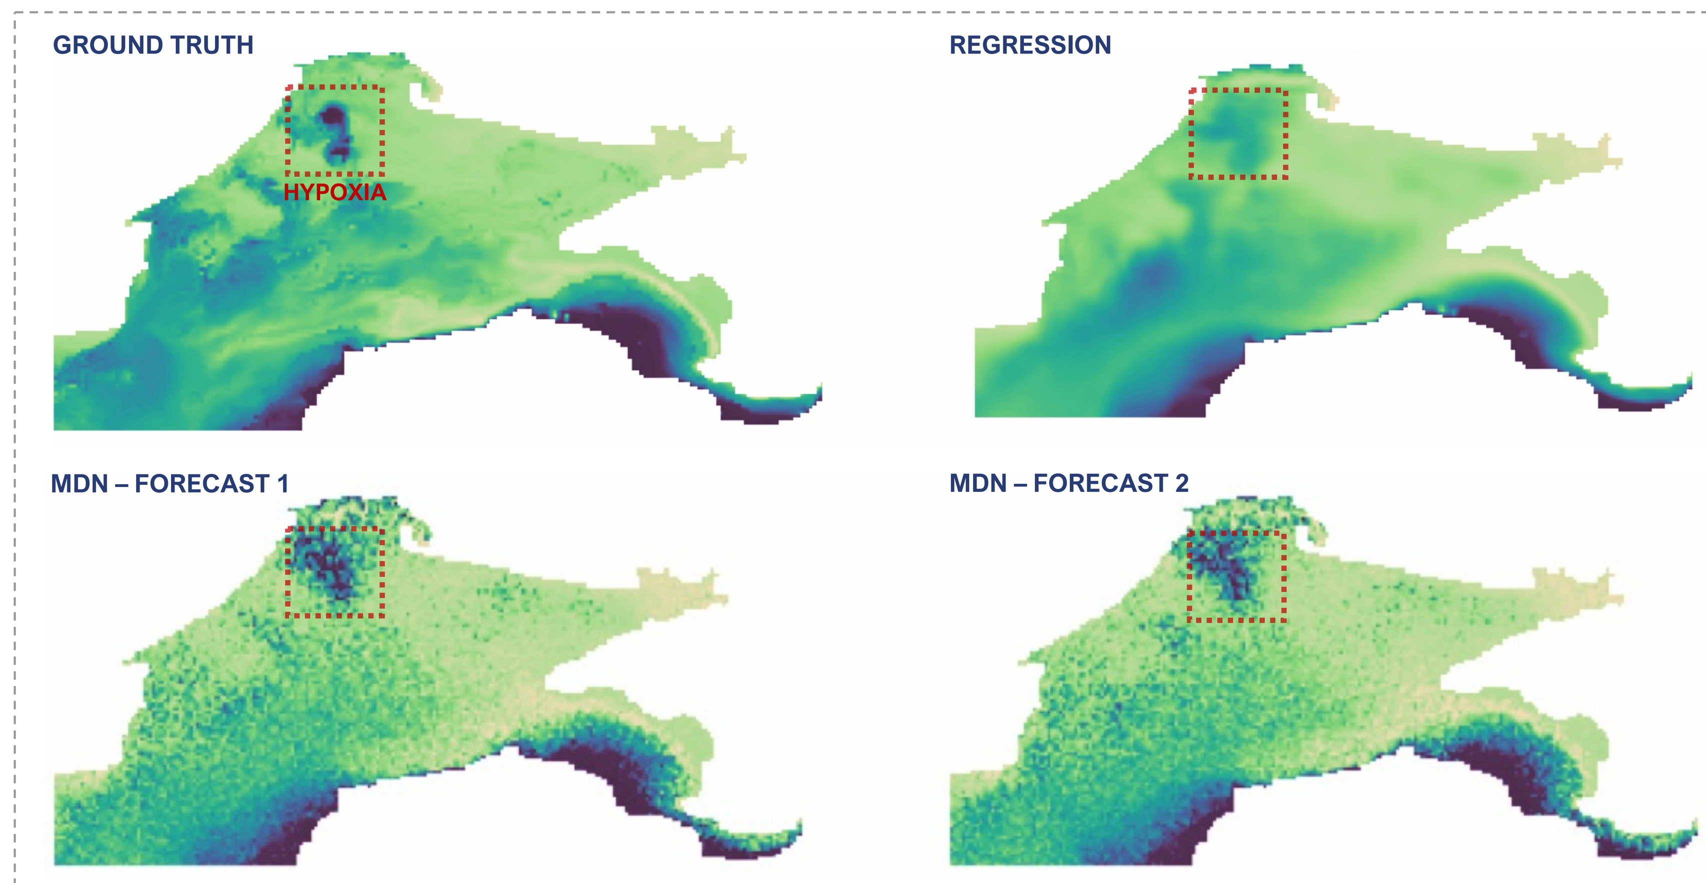

Figure 2: Visual comparison of ground truth (top-left) and forecasted oxygen levels. The darker spot indicates hypoxia. The regression model (top-right) overestimates oxygen levels at the hypoxia location. In contrast, the mixture density network (using two Gaussians) capture the hypoxia event, as shown by the darker spots in its two randomly sampled forecasts.

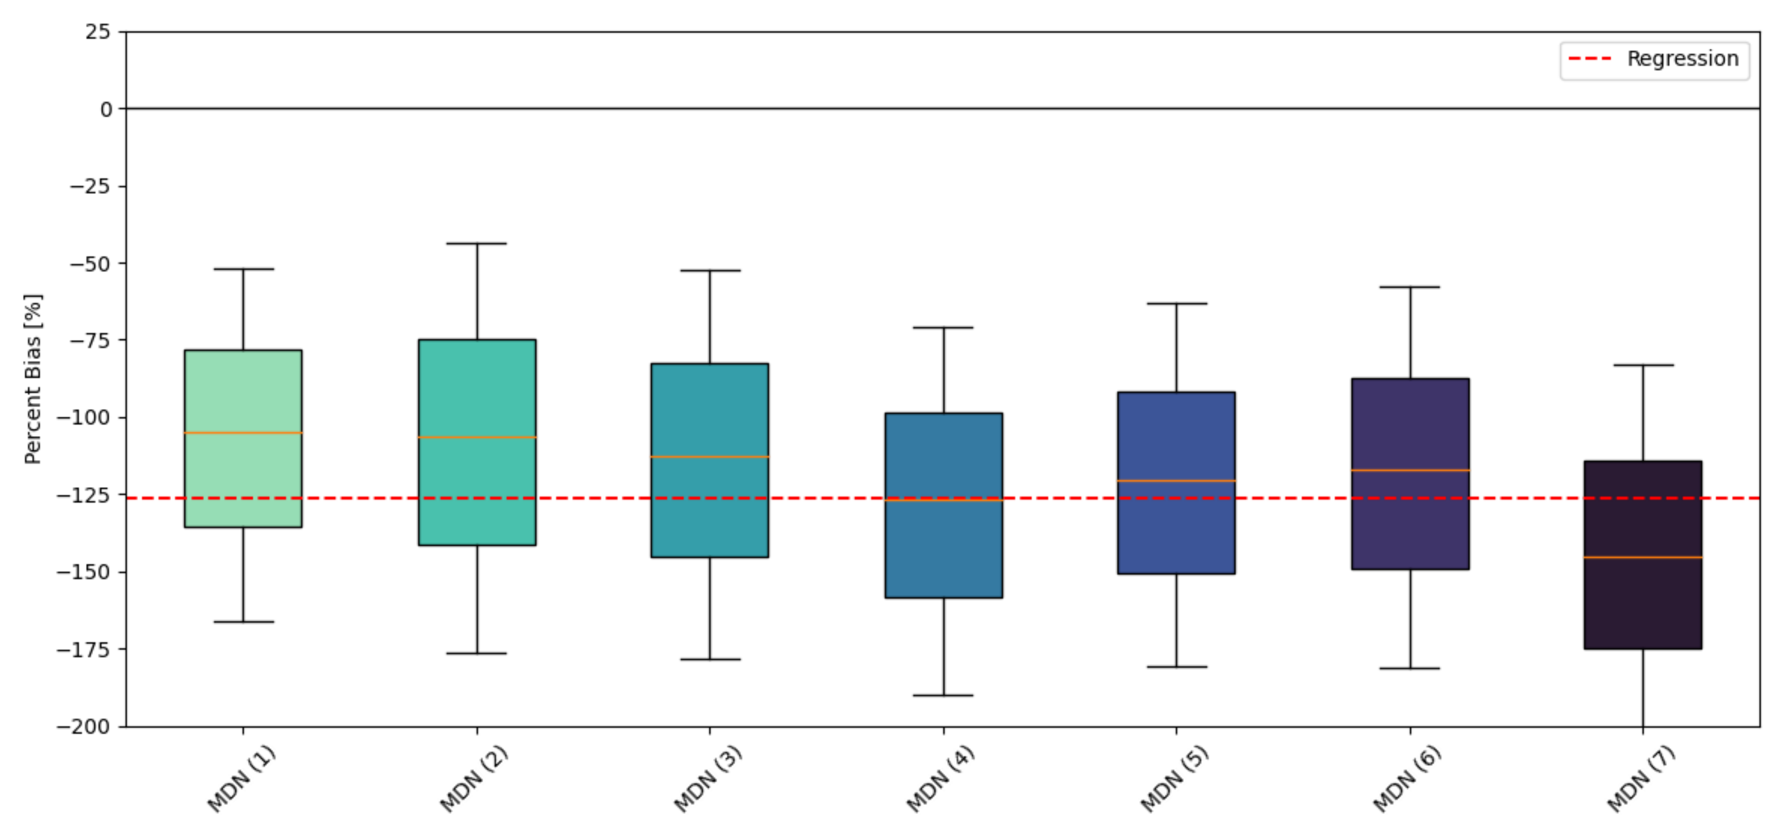

Figure 3: Comparison of Global Percent Bias (PB) between the regression model and mixture density networks on the validation dataset. A positive PB indicates underestimation of oxygen concentration, while a negative PB indicates overestimation. Each MDN model generated 25 forecasts, and we computed the minimum and maximum quantiles at 10% and 90%, respectively. All models tend to overestimate oxygen concentration (PB < 0), but most of the MDN forecasts have a lower bias than the regression model in 50% of their predictions. For the best MDN model (MDN 1), the minimum quantile shows that even 10% of the generated forecasts had a bias 75% lower than that obtained through regression.

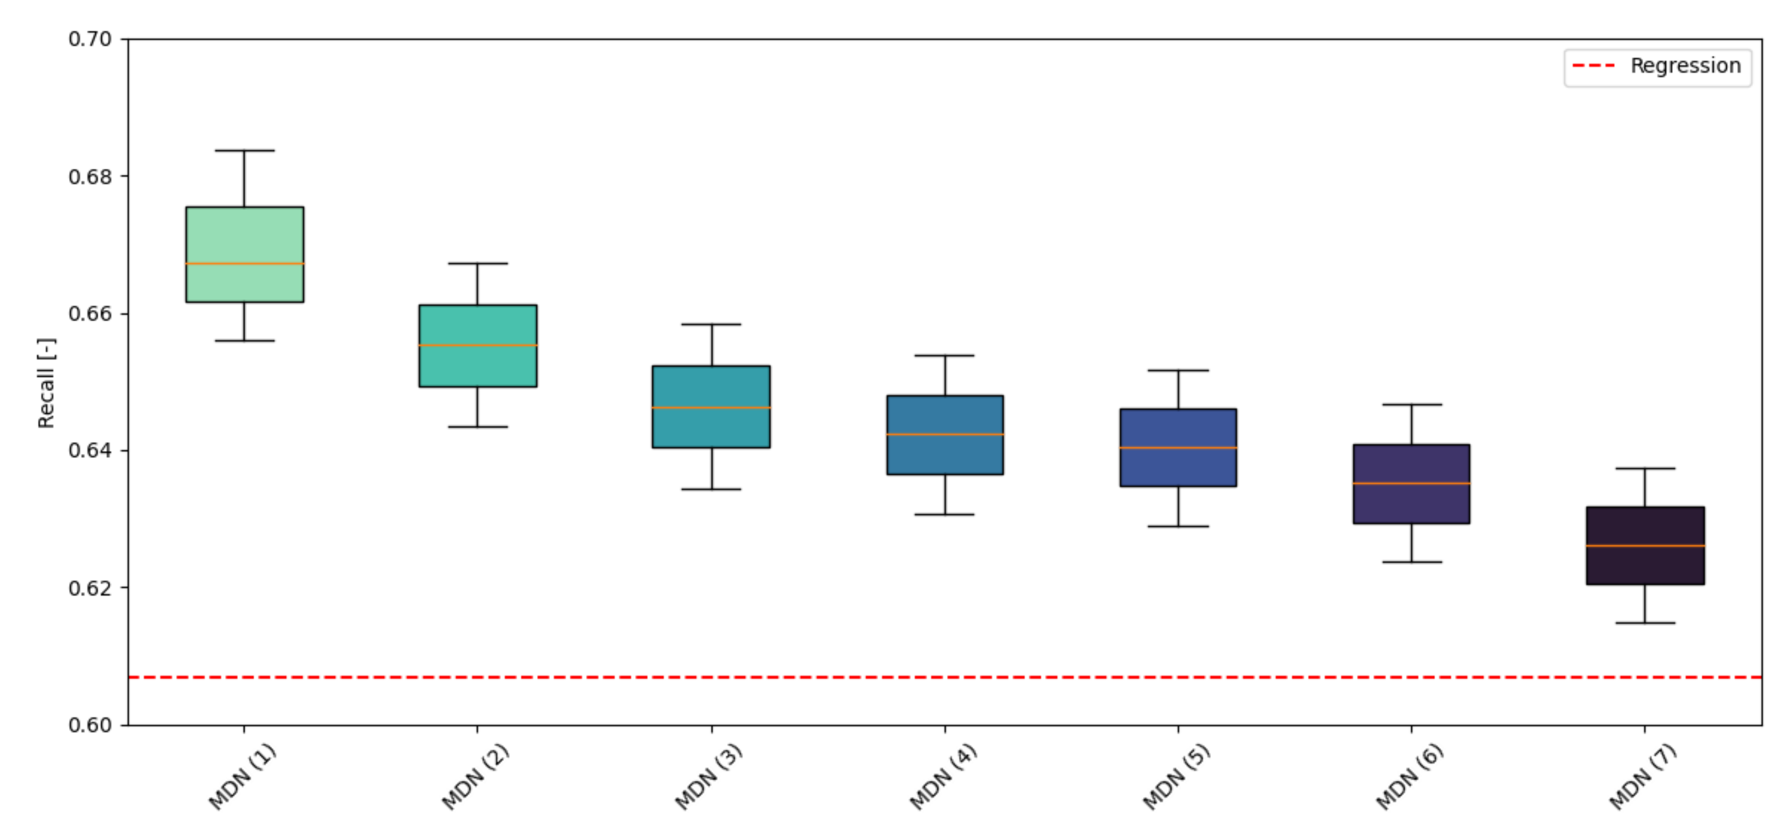

Figure 4: Comparison of Global Recall between the regression model and mixture density networks on the validation dataset. Recall measures the ability to correctly predict hypoxia when it occurs. Each MDN model generated 25 forecasts, and we computed the minimum and maximum quantiles at 10% and 90%, respectively. All MDN forecasts consistently outperformed regression, with the best MDN achieving a 7% higher recall in 50% of its forecasts.

Conclusion

- Forecasting: Regression models are limited to learning average oxygen levels, failing to capture specific sea dynamics over years and future scenarios, resulting in high bias.
- Detection of Hypoxia: Regression models tend to predict higher oxygen levels, which complicates the identification of regions experiencing hypoxia.
- Mixture Density Networks: They offer a more accurate representation of sea dynamics and reduced bias through complex density approximations. However, they face challenges with marginal likelihood approximation, resulting in blurred oxygen profiles spatially.