



OGS

National Institute
of Oceanography
and Applied
Geophysics

Carlos Enmanuel Soto Lopez

Via Beirut 2

Trieste, Italy

I-34014

✉ csoto@ogs.it

Dear Editors,

Please find attached a revised version of the manuscript entitled “Data-Informed Inversion Model (DIIM): a framework to retrieve marine optical constituents in the BOUSSOLE site using a three-stream irradiance model”, submitted for consideration as a regular Article (DOI: <https://doi.org/10.5194/gmd-2024-174>).

We have amended the manuscript according to the comments of the reviewers. We also added a version of the revised manuscript with the main differences in blue and a point-by-point response to their comments, together with a description of the changes made to the manuscript.

Yours sincerely,

Carlos Enmanuel Soto Lopez

Point by point reply to the Executive Editor

Executive Editor comment:

Dear authors,

Unfortunately, after checking your manuscript, it has come to our attention that it does not comply with our “Code and Data Policy” .

https://www.geoscientific-model-development.net/policies/code_and_data_policy.html

In your Code and Data Availability section you provide a handful of links for sites containing information related to your manuscript. However, such sites are not acceptable for scientific publication. They include a GitHub site and Google drives. You have included the DOI for a Zenodo repository containing part of the information, after an initial request from the Topical Editor; however, the Zenodo repository that you provide does not contain the necessary information to replicate the work that you present in your manuscript. For example, when trying to run the Python notebook on model availability, it tries to connect to the external GitHub site (moreover under certain circumstances/configurations it fails to do it).

I have to be clear here. All the models, code and data necessary to replicate your work must be contained in the Zenodo repository and must be standalone, that is, work without need for additional download of software or data from third party sites (excluding common libraries such as pandas, numpy, etc. that should be listed with their versions). Also, the Code and Data Availability section in your manuscript contains several references to sites that do not comply with our policy. This section is not there to advertise the “last version” of a software or the work of the authors, but to provide the specific information that assures the compliance with the principle of scientific reproducibility. Including such unuseful information only introduces unnecessary noise and complexity in the assessment of the compliance of your submitted manuscript. Therefore, please, remove from this section all the information about sites that do not comply with the policy for permanent archival of code and data.

Therefore, the current situation with your manuscript is irregular. Please, publish your code in one of the appropriate repositories and reply to this comment with the relevant information (link and a permanent identifier for it (e.g. DOI)) as soon as possible, as we can not accept manuscripts in Discussions that do not comply with our policy. Also, please include the relevant primary input/output data. Also, you must include a modified “Code and Data Availability” section in a potentially reviewed manuscript, containing the DOI of the new repositories. I note that if you do not fix this problem, we will have to reject your manuscript for publication in our journal.

Juan A. Añel

Geosci. Model Dev. Executive Editor

Author reply: After careful review ing of the code and data availability, we have uploaded in Zenodo a standalone version of the code used to reproduce the results from the article, as well as the data needed to reproduce it [7]. In this new upload, we make no mention of other sites, except the standard python3 libraries.

We have also prepared a README with instructions on how to install the required external packages with a Makefile. We also included a Python script (reproduce.py). By launching reproduce.py, the user is guided to reproduce the calculations and figures in the manuscript, with the option to run all calculations (-a), only the Bayesian computations (-b), only the sensitivity analysis and MCMC analysis (-m), or only the calculations related to training the neural network (-n).

DOI from the new code and data availability: 10.5281/zenodo.14609747.

Link to the code and data availability: <https://doi.org/10.5281/zenodo.14609747>.

Changes made to the original manuscript:

- In the section Code and data availability, we have added in the first paragraph a description of the Zenodo availability, as well as the corresponding DOI.
 - As a consequence of the additional review, we have added the citation for the literature value of some parameters used, citations missed during the first upload, as well as the correction of a typo for the index of one of the parameters.
 - Finally, for clarity at the moment of reading, we have changed the wording of some parts in section Results.
-

Point by point reply to Reviewer 1

Reviewer comment: Section 1. *The authors refer to inversions of semi-analytical expressions of the RTE (l33-l34). As the term is mentioned in the abstract but not used afterwards, the reader could have doubts as to whether this is actually the case in the paper. If the complete RTE model (i.e. discretized on the vertical) is not used afterwards, this should be explained in the introduction. Similarly, it should also be stated from the beginning if the aim of the paper is to estimate IOPs at the sea surface only (assuming vertical homogeneity of marine optical components), and not throughout the water column.*

Author reply: We have added information in the Introduction about the type of model used. The model is based on an expression derived from the Radiative Transfer Equation (RTE), with some terms—such as the chlorophyll-to-carbon ratio (see manuscript, Eq. 3)—estimated empirically. This makes it a semi-analytical expression. We have also clarified that the inversion is intended solely to approximate the inherent optical properties (IOPs) at the sea surface.

Reviewer comment: Section 2. *I suggest moving this section after section 3 and explain more clearly how the different data are related to the RTE model variables, surface boundary conditions, etc. In addition, I would suggest introducing a schematic diagram of the water column, including the different fluxes, irradiances, and variables of interest, data collected, etc, to improve readability.*

Author reply: As suggested, we have moved this section to follow the model description and expanded on the relationship between the OASIM data and the boundary conditions. We also provided more detail on the origin of the data, noting that it was previously used to assess the validity of the OASIM model in the Mediterranean Sea [4]. Finally, we also included a diagram describing the different fluxes involved in the RTE.

Reviewer comment: Section 3. *It looks that the marine optical constituents are assumed to be constant on the vertical. It should be clarified if vertical homogeneity is assumed, as this is far from a trivial assumption. It could significantly limit the applicability of the framework to locations where such an assumption is not verified, e.g. when shallow sub-surface chlorophyll maximum concentration occurs. A discussion should be added explaining how the scheme could be extended for inhomogeneous water columns. As for vertical homogeneity of optical constituents, it seems that temporal persistence is assumed over the daily cycle. Please confirm and provide some justification.*

Author reply: The equations assume a homogeneous, infinitely deep layer. By making this assumption, the system of equations becomes linear and can be solved analytically. The assumption is justified for deep case 1 waters (dominated by phytoplankton), like the one studied in the present

work. During winter, the chlorophyll concentration in the first layer is approximately constant due to mixing (see [5], Fig. 1), while most of the downward irradiance comes from the first 10 to 20 meters (see [6], Fig. 1 and Fig. 2). During summer, there is no mixing, but still there is a region, around 20 to 50 meters, with constant chlorophyll concentrations, making the assumption justified. For coastal areas, we are considering extending the model to include the process of light reflection at the sea floor, which will be the matter for future work. Concerning the temporal persistence, we are computing daily averages of the different concentrations (averaging over hours with light), making it coherent with satellite data retrievals. We added all this information to the modified manuscript.

Reviewer comment: Section 4. *Section 4.4.2 in particular, is very difficult to understand and needs to be clarified. More generally, the sequence of algorithmic operations in the Bayesian inversion framework and associated partitioning of the training data set need to be better explained, so that the reader understands the order in which the latent variables and the model parameters to be optimized are calculated. This clarification is necessary to fully understand the applicability of the method and the results in section 5. The algorithmic complexity of the sequence also needs to be quantified and explained more transparently. The list of optimized parameters is confusing in places. Please clarify the actual list of optimized parameters, referring to the list displayed in Table 1 and 2.*

Author reply: We substantially re-wrote and re-organized section 4 to make the explanation of how the algorithm works clearer, including tables of step by step explanations for the Bayesian inversion and Monte Carlo algorithms (Alg. 1, Alg. 2, and Alg. 3). Concerning the algorithm complexity, the aim is to understand how the computational time increases as a function of the dimensionality of the problem. The general consensus is that under certain conditions, the computational time increases polynomially, but in general, the increase could be exponential [1]. Indeed, the complexity estimation for this kind of algorithms is an active area of research, and could be in itself the subject of another work.

Reviewer comment: *As presented, the overall convergence of the method does not seem very robust and is subject to numerous adjustments depending on the dataset and DYFAMED site considered. Section 4.4.2: The sentence “we choose to perturb our parameters in such a way that we end up with a fourteen-parameter space” is not understandable. What does the hyperparameter alpha q (line 283) represent? The parameter perturbation method needs to be better explained (also related to Table 5). The last sentence of section 4.4 “... and sampled only uncorrelated values,” is rather cryptic.*

Author reply: Motivated by the reviewer’s comment on the non-robustness of the method, we analyzed the two proposed algorithms separately, looking for the underlying reasons explaining their different estimates.

The mean value of the retrieved optical constituents for both methods were consistent with each other, in the sense that the Variational Bayes method gives values within the uncertainty of the Bayesian approach. The robustness of the method is confirmed by the statistics on the test data, which were not used for hyperparameter tuning or neural network training, as both approaches produce values close to the observations. Additional evidence of robustness comes from the comparison with a state-of-the-art algorithm for chlorophyll estimation from satellite observations, presented at the end of the Results section.

On the other hand, the forward model optimization presented differences between the MCMC mean values and those obtained with the SGVB estimator. Upon analyzing the results, we identified what we believe to be the main reasons behind these discrepancies. On one hand, the loss functions used in the two methods are different due to the regularization term used for the neural network training. On the other hand, the algorithm used to train the neural network relies on standardized

data and minibatch minimization, two common practices in neural network training that help improve generalization and prevent overfitting by approximating the full dataset gradient with that of smaller subsets. For this resubmission, we therefore also applied the standardization to the data used in the MCMC algorithm to ensure a fairer comparison. This change made the final parameters estimated by both methods closer to each other with respect to the previous experiments, in a way that 22 of the 24 perturbed parameters were within the uncertainty of the MCMC estimates; Nonetheless, the parameters obtained through simultaneous training of the neural network using the SGVB estimator demonstrated better generalization properties.

In the author's view, the fact that both methods give slightly different results in finding the optimal parameters is not an indicator of the incompatibility or non-robustness of the methods. Indeed, MCMC estimates the probability density of the parameters given the training data, i.e. the posterior. This distribution reflects the region of parameter space where the optimal parameters are likely to lie, while its mean provides a point estimate of those parameters. Instead, the SGVB estimator exploits minibatch minimization and the training of a Neural Network to estimate the MLE of the parameters, a different estimator with slightly different assumptions (we never linearized the Forward function, for example). Both algorithms succeed in finding a satisfactory estimator of the optimal parameters, minimizing the RMSE on the test data. While MCMC methods are more commonly used, our goal was to demonstrate the validity and efficiency of the SGVB estimator in marine inversion problems.

Reviewer comment: *There are no objective elements in the paper to guide the choice between the MCMC vs. SGVB methods. The GMD framework calls for guidelines to direct the user towards the most suitable method (Bayesian approach or neural network) instead of experimenting (as seems to be proposed) until a satisfactory converged solution is found.*

Author reply: We added a few comments in the Conclusion section detailing about the advantages and disadvantages of both methods. In summary, we recommend the MCMC and traditional Bayesian approaches for their theoretical simplicity and interpretability, and their ability to estimate the uncertainty of the results. The Variational Bayes approach is more effective for intractable problems, where the posterior is costly to compute, returning good estimates of the MLE. At the moment, our findings are that the method doesn't return reliable uncertainty quantification, but still offers a good alternative for the latent variable estimation (estimation of the optical constituents), as well as for the model optimization. In the paper, we compared the results of both methods, showing their equivalence, since the dimensionality of the problem allows it, but in most of the state-of-the-art forward models, the latent space, as well as the parameter space, is usually much larger, and the standard methods are not any more computationally feasible.

Reviewer comment: Section 5. *The presentation of the results (section 5) is confusing. It is not clear what the purpose of the sensitivity study introduced at the beginning of the section is, since its main conclusion (the use of a single set of parameters over one year is sub-optimal) is not used thereafter anyway. This part should be removed, or better justified. What does the new `bphy,Int` parameter introduced here (Figure 3 and Table 5)? A new structuring of this section into three parts (results of the MCMC method, results of the SGVB method, comparison of the 2 methods) should be considered, including a more in-depth evaluation and interpretation of the results. For example, the reason why the IOP values found with SGVB differ very slightly from the original values compared with MCMC requires substantiated explanations.*

Author reply: We re-structured the results of the section into four parts: the first one focuses on the Bayesian retrieval of the optically active constituents on the surface of the sea and the uncertainty estimation; the second on the parameter optimization; the third on the comparison between the Bayesian outputs and the Variational Bayes approach; the last one compares our results with a state

of the art algorithm for satellite sea surface chlorophyll a estimation. We expect this rewriting to make the content of the section more clear.

Regarding the sensitivity study, we provided additional motivation by framing it as a measure of the seasonal variability of the parameters. Moreover, it allowed us to analyze the differences between the two optimization methods. In particular, it revealed that among the two parameters showing the greatest discrepancy between methods, only one played a significant role in optimizing the particulate backscattering coefficient. Since these measurements were the noisiest ones, we speculate that overfitting could have misled the MCMC algorithm.

Reviewer comment: *Figure D1 does not include a comparison with the chlorophyll a concentration estimated using conventional ocean colour (OC) algorithms available in the CMEMS catalogue. It is strongly recommended to add these figures in the plot and provide comments about their consistency, as one of the potentially key advantages of the proposed method is to be used in operational setups.*

Author reply: To compare our inversion results with a conventional ocean color algorithm for estimating chlorophyll-a concentration, we added a dedicated subsection in the Results. In this analysis, we performed a comparison over an extended region near the BOUSSOLE buoy, using the MedOC4.2020 algorithm [2]. The reference dataset consists of measurements obtained via High-Performance Liquid Chromatography (HPLC) [3]. The comparison showed the consistency between both methods (See manuscript, Fig. 11).

Reviewer comment: Section 6. *The question of the reproducibility of the inversions in other sites needs to be addressed explicitly in a more convincing manner. The discussion refers to some generic aspects (e.g. use of neural networks in earth sciences) but does not address how useful the proposed tools will be for other sites.*

Author reply: We plan to assess the validity of the inversion on the rest of the Mediterranean in a future work. At the moment, we presented a validation only on a region close to the BOUSSOLE buoy of $4^\circ \times 4^\circ$, in the North West Mediterranean Sea, where conditions are similar to the assumptions made for this work.

Reviewer comment: Specific comments and questions: *Abstract, « we conclude that both methods are consistent with the Radiative Transfer Equation » : what is really meant by this sentence ?*

Author reply: We changed this sentence; what we meant was that the SGVB estimator was estimating an inversion of the RTE more than just adjusting a NN to data. We agreed that our results were not primarily focused on that aspect and therefore replaced it with the comparison between our inversion method and the state-of-the-art algorithm for chlorophyll-a concentration, as we believe this comparison is more relevant and deserves to be highlighted in the abstract.

Reviewer comment: *MCMC = Markov Chain Monte Carlo (abstract,) or Markov State Monte Carlo (section 4) : please unify the nomenclature or explain nuances*

Author reply: The correct name is Markov Chain Monte Carlo.

Reviewer comment: *Line 43-44 and elsewhere: replace « density » of optical constituents by « concentration » ?*

Author reply: Thank you for pointing that out.

Reviewer comment: *Line 55 : It allows for ...*

Author reply: Thank you for bringing that to my attention.

Reviewer comment: *Line 85 « After filtering ... » : please revise and clarify this sentence*

Author reply: First, we removed any data coming from the buoy reporting an absolute tilt higher or lower than 10 degrees. We also removed the data recorded at a depth more than 2 m below the nominal values (4 m and 9 m, depending on the instrument of measurement). Also, the downward light attenuation coefficient data were filtered with an analog high-pass filter, using the package SciPy from the programming language Python, filtering the noise with a frequency less than 4 hours. Finally, we proceeded to average the daily values.

Reviewer comment: *Line 90 : below instead of above ?*

Author reply: Thank you, we meant below.

Reviewer comment: *Line 93-94 : what is the $\ll$ heigh vertical variability $\gg$? Please revise the sentence and better explain the rationale behind the choice of measurements at a depth of 9m*

Author reply: We made it clearer now. The measurements were taken at 9 meters deep. Because there is low variability, the chlorophyll measurements were considered as sea surface measurements. However, the downward light attenuation coefficient has a high variability, so we can not consider it as a measure at the sea surface.

Reviewer comment: *Equation (1) : d/dh missing in second and third equations*

Author reply: I really appreciate all these comments—thank you. Yes, I've made the corrections.

Reviewer comment: *Section 3.1 : an equation is missing to describe how PAR is related to the direct, scattered and upward irradiances.*

Author reply: Thank you for pointing this out. You are right—an equation describing how PAR relates to the direct, scattered, and upward irradiances was missing. We have now included it in Section 3.1.

Reviewer comment: *Section 4.2. It is not clear if the process to adjust the alpha hyperparameter defining the model error is dependent on the accuracy of the surface boundary conditions (OASIM). Please provide comments and clarification.*

Author reply: Alpha is a hyperparameter. In linear regression problems, it typically represents the inverse of a regularization term and is often set to a small value so as to not affect the maximum likelihood estimate (MLE). In this resubmission, we discussed in the Results section about the prior. Since it is an informative prior, without it, the estimated uncertainty increases by one order of magnitude (see manuscript, Results section). Therefore, the choice of alpha must balance two objectives: First, it must ensure robustness (i.e., small variations in alpha do not significantly affect the results) and second, it has to yield realistic uncertainty estimates, meaning that the reported uncertainty should, on average, match the observed discrepancies between model predictions and actual data. We have revised the Appendix to clarify this point.

Reviewer comment: *Line 180 : $x(\lambda)$ is a 5-component vector (including E_u) while it is a 4 component vector in eq(11). Please explain.*

Author reply: The correct number is 4 components.

Reviewer comment: *Line 217: I don't understand why in situ observations are available for 3 wavelengths only. This is not consistent with what is said in section 2.3. Please update section 2.3 accordingly.*

Author reply: My apologies if it was not clear. I added a small clarification at the end of the data acquisition section. Taking into account the assumptions and data availability, the in-situ observations considered are sea surface chlorophyll, 9 meters deep downward light attenuation coefficient in 5 wavelengths, (412.5,442.5,490,510,555) nm, and sea surface particulate backward scattering coefficient

at 3 wavelengths (442,490,510) nm.

Reviewer comment: Line 238: his - $\dot{z}$ the

Author reply: This section was re-worked.

Reviewer comment: Line 256: Please explain the “standard error propagation scheme” used to compute the uncertainties.

Author reply: I expanded on the equations. I was referring to the error propagation equations: $\Delta F(\vec{x})^2 = \nabla_x F(\vec{x}) \Sigma^x \nabla_x F(\vec{x})^T$, where $\Delta F(\vec{x})$ is the error of a function $F(\vec{x})$, $\nabla_x F(\vec{x})$ is the Jacobian, and Σ^x is the covariance matrix of x . In our case, $\Sigma^x = \Sigma_{\vec{z}^{d*}}$. These equations assume that each component of $\vec{x}$ is not correlated with the others, and, in this respect, is only an approximation for nonlinear functions.

Reviewer comment: Figure 1: the Estimate of the optimal parameters (“Lambda tilde”) is not shown in the diagram.

Author reply: They are the parameters from $p_{\hat{\Lambda}}(\vec{y}, \mathcal{H}|\vec{x}, \hat{z})$ which is the likelihood given the estimated latent variable $\hat{z}$ and the boundary conditions $\vec{x}$. In other words it is the probability of observing the Remote Sensing Reflectance, and observations, as a function of the estimated parameters $\hat{\Lambda}$, given the boundary conditions and the state of the estimated state of the ocean.

Reviewer comment: Table 4: why only 9 parameters are shown here ?

Author reply: We perturbed 15 lambda dependent parameters, and 9 non lambda dependent. The table shows the 9 that were not dependent on lambda, and Fig. 10 (in the new manuscript) has the final values for the lambda dependent ones. On the other hand, on other tables, we report the perturbation factors δ_i . We rewrote the section where we explain how the parameters were perturbed. More specifically:

The values of the λ dependent vector of dimension five representing the phytoplacton-specific absorption coefficients $\vec{a}_{\text{phy}}$ were perturbed as: $\vec{a}_{\text{phy}}^* = \delta_{a_{\text{phy}}} \vec{a}_{\text{phy}}^0$ with $\delta_{a_{\text{phy}}}$ a learnable scalar, and $\vec{a}_{\text{phy}}^0$ the literature values. We chose it like this to maintain the shape of the function $a_{\text{phy}}(\lambda)$ unperturbed.

For the carbon-specific scattering and backscattering coefficients $b_{\text{phy}}(\lambda)$ and $b_{b,\text{phy}}(\lambda)$, we first linearly interpolated them with the literature values, and perturbed the tangent and the intercept of the linear interpolations, $b_{\text{phy}}(\lambda)^* = \delta_{b_{\text{phy},\text{int}}} b_{\text{phy},\text{int}}^0 + \delta_{b_{\text{phy},\text{T}}} b_{\text{phy},\text{T}}^0 \lambda$.

The parameters d_{CDOM} , $b_{r,\text{NAP}}$, S_{CDOM} , $\Theta_{\text{chla}}^{\min}$, Θ_{chla}^0 , β , σ , Q_a and Q_b perturbations consisted in per parameter scalar multiplications. All the other parameters were left unperturbed.

In this way, we perturbed 24 parameters, 9 of them by multiplying them for a scalar δ_i , i equal to each of the perturbed parameters, the five components of $\vec{a}_{\text{phy}}$ by multiplying them by the same scalar $\delta_{a_{\text{phy}}}$, and finally, $b_{\text{phy}}(\lambda)$ and $b_{b,\text{phy}}(\lambda)$ by linearly interpolating them, and perturbing the tangent and the intercept of each of them, making a total of 14 perturbation factors.

Reviewer comment: End of Appendix A: “For completeness, ... has to be exchanged ...”

Author reply: We agree, thank you for pointing it out.

Reviewer comment: Line 518: “... these two errors”.

Author reply: We made some adjustments to Appendix B.

Reviewer comment: Figure D1. Why 2012 time series? Do you mean 2005-2013 time series ?

Author reply: Thank you for pointing it out, we fixed it as 2005-2013 time series.

Point by point reply to Reviewer 2

Reviewer comment: Line 89, "data filtered with an analog high pass filter, using the package SciPy". It cannot be an analog filter. Maybe the prototype is analog, but it must be a digital filter. Please explain/correct this

Author reply: We substituted the word 'analog' with the specific name of the filter. Now the text say "...were filtered with a Butterworth high-pass filter, using the package SciPy".

Reviewer comment: Line 130, *respectably* => *respectively*

Author reply: I fixed the typo. Thank you for pointing it out.

Reviewer comment: Line 249, in eq. 19 K is not defined.

Author reply: K is the Jacobian of R with respect to z . We added the information in the new version of the manuscript.

Reviewer comment:

- Line 313, *DK divergence* => *DKL divergence*.
- Line 347, *the DK divergence* => *the DKL divergence*.
- Line 357, *Cholezky* => *Cholesky*.
- Line 371, $q(z,y,x)$ => $q(z|y,x)$

Author reply: We corrected all the typos in the updated manuscript.

Reviewer comment: Fig. 1, What is exactly "model of the measurements". Maybe you can add R_{rs}^{MODEL} and $\hat{H}(Z, X; \Lambda)$.

Author reply: We changed the term to "Forward Model", as this is the name used throughout the rest of the paper.

Reviewer comment: Fig. 2 caption: $q_\phi(z|z,y)$ should be $q_\phi(z|x,y)$, and "dose that learn the Cholezky" => those? that learn the Cholesky.

Author reply: We changed the typos in the updated manuscript.

Reviewer comment: Fig 5. (Related to it) Please give more explanation on: Why MCMC sub or over estimates? Why the SGVB does not have a confidence interval if the method provides it? And why do you think is the cause the SGVB uncertainty fails?

Author reply: Regarding the confidence interval, the Variational Bayes method approximates the posterior of the optical constituents, called Z , or latent variables in the text. To this end, the neural network estimates both the mean and the covariance; however, since the covariance underestimates the true discrepancy between predictions and observations, we relied only on the mean values. For uncertainty estimation, linearizing the Forward Function around the estimated mean and applying standard error propagation methods yielded better performance. On the other had, the Variational Bayes estimates only the MLE of the parameters, as plotted in Fig. 5.

Regarding the different values obtained by the two methods: after standardizing the observations (dividing by the standard deviation) used for the MCMC algorithm, the differences are minimal (see Fig. 10 of the new manuscript). They could be due, e.g, to overfitting, changes in the loss function (the NN loss function used a regularization), or model error, i.e. the error generated by the approximation the posterior with a NN.

Regarding the underestimation of the covariance matrix, I speculate that it is due to training the neural network to match observations directly. The original goal is to invert the Forward Model. If we were to

use the Maximum Likelihood Estimate (MLE), computing the inverse of the Fisher Information Matrix would result in extremely large uncertainty estimates. In contrast, the Maximum A Posteriori (MAP) estimate yields lower uncertainty, as it incorporates prior information. The 'posterior' represented by the neural network includes not only a term for the observations but also a regularization term, which, when viewed from a Bayesian perspective can be interpreted as an additional form of prior information. This added prior information reduces the spread of the posterior, thereby underestimating the true discrepancy between the modeled data and the observations. However, since this remains speculative, in the paper we focused solely on reporting the result that the neural network underestimates the uncertainty.

References

- [1] Alexandre Belloni and Victor Chernozhukov. “On the computational complexity of MCMC-based estimators in large samples”. In: (2009).
- [2] S Colella et al. *EU Copernicus Marine Service Quality Information Document for the Ocean Colour Mediterranean and Black Sea Observation Product, OCEANCOLOUR_MED_BGC_L3_NRT_009_143, Issue 4.1, Mercator Ocean International*. (Accessed on 05-23-2025). 2025. DOI: <https://doi.org/10.48670/moi-00299>.
- [3] V Di Biagio, S Campanella, and G Cossarini. *In situ dataset for initialization and validation of the Copernicus Med-MFC biogeochemical model system (MedBGCins)*. 2025. DOI: <https://doi.org/10.5281/zenodo.15489967>.
- [4] Paolo Lazzari et al. “Assessment of the spectral downward irradiance at the surface of the Mediterranean Sea using the OASIM ocean-atmosphere radiative model”. In: *Ocean Science Discussions* 2020 (2020), pp. 1–39.
- [5] A. Mignot et al. “From the shape of the vertical profile of in vivo fluorescence to Chlorophyll-a concentration”. In: *Biogeosciences* 8.8 (2011), pp. 2391–2406. DOI: 10.5194/bg-8-2391-2011. URL: <https://bg.copernicus.org/articles/8/2391/2011/>.
- [6] JJ Simpson and TD Dickey. “The relationship between downward irradiance and upper ocean structure”. In: *Journal of Physical Oceanography* 11.3 (1981), pp. 309–323.
- [7] C. Soto. *Data-Informed Inversion Model (DIIM)*. Version 1.3. Sept. 2025. DOI: 10.5281/zenodo.14609747.