The Journal of Open Source Software

OceanBioME.jl: A flexible environment for modelling the coupled interactions between ocean biogeochemistry and physics

Jago Strong-Wright [1,2¶], **Si Chen** [1,2*], **Navid C Constantinou** [3,4*], **Simone Silvestri** [5*], **Gregory LeClaire Wagner** [5*], **and John R Taylor** [1,2]

1 Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, United Kingdom **2** Centre for Climate Repair, Cambridge, United Kingdom **3** Australian National University, Australia **4** Australian Research Council Centre for Climate Extremes **5** Massachusetts Institute of Technology, USA ¶ Corresponding author * These authors contributed equally.

DOI: 10.21105/joss.05669

Software

- Review
- Repository
- Archive

Editor: Kristen Thyng

Reviewers:

- @francispoulin
- @iuryt

Submitted: 24 June 2023
Published: 05 October 2023

Statement of Need

To date, about 25% of anthropogenic carbon emissions have been taken up by the ocean (Friedlingstein et al., 2022). This occurs through complex interactions between physics, chemistry, and biology, much of which is poorly understood. Due to the vast size of the ocean and the sparsity of data; modelling and data assimilation play a vital role in quantifying the ocean carbon cycle. Traditionally ocean biogeochemical (BGC) modelling involves large and inflexible code bases written in high-performance but low-level languages. Using and modifying these models usually requires knowledge of the details of the numerical implementation (e.g. the grid layout and time-stepping schemes). Most traditional BGC models also require CPU-based supercomputers to run quickly, which presents a barrier to experimentation and innovation.

One area where novel ideas must be explored with BGC codes is assessing ocean carbon dioxide removal (OCDR) strategies. Quantifying the effectiveness and identifying the impacts of OCDR is challenging due to the aforementioned complexity of the ocean BGC system. Moreover, field trials of OCDR interventions are generally small-scale and targeted, while the intervention required to have a climate-scale impact is regional or global. This necessitates adaptable, easy-to-use, and verifiable BGC modelling tools which can be used to assess OCDR strategies at the fast pace with which they are being developed (National Academies of Sciences & Medicine, 2022). We have built `OceanBioME.jl` to meet these challenges by creating a tool that provides a modular interface to the different components, within the ocean modelling framework provided by `Oceananigans.jl` (Ramadhan et al., 2020). Previously implementing biogeochemical models in `Oceananigans.jl` required the user to add forcing terms and boundary conditions to generic tracers (e.g. Simoes-Sousa et al., 2022). `OceanBioME.jl` provides a suite of biogeochemical models ranging from simple idealized to full-complexity models and models for boundary fluxes (e.g. air-sea gas exchange). Facilitated by `KernelAbstractions.jl` and `CUDA.jl` (Besard et al., 2019; Churavy et al., 2023), `Oceananignas.jl` and `OceanBioME.jl` are built from the ground-up to exploit the power of graphical processor units (GPUs) while also retaining the ability to run on CPUs. The flexibility of the `Oceananigans.jl` framework allows `OceanBioME.jl` to be applied across a wide range of scales and use cases, including small-scale large-eddy simulations and regional and global models.

The Journal of Open Source Software

Summary

`OceanBioME.jl` is a flexible modelling environment written in Julia (Bezanson et al., 2017) for simulating the coupled interactions between ocean biogeochemistry, carbonate chemistry, and physics. `OceanBioME.jl` can be used as a stand-alone box model, or integrated into `Oceananigans.jl` for coupled physical-biogeochemical simulations in one, two, or three dimensions. As a result, `OceanBioME.jl` and `Oceananigans.jl` can be used to simulate the biogeochemical response across an enormous range of scales: from surface boundary layer turbulence at the sub-meter scale to eddying global ocean simulations at the planetary scale, and on computational systems ranging from laptops to supercomputers. An example of a problem involving small-scale flow features is showcased in Figure 1, which shows a simulation of a sub-mesoscale eddy in a 1km x 1km horizontal domain with an intermediate complexity biogeochemical model and a kelp growth model solved along the trajectories of drifting buoys (a list of of examples shown in this paper and links to source code are given at the end of the paper). `OceanBioME.jl` leverages Julia's multiple dispatch and effective inline capabilities to fuse its computations directly into existing `Oceananigans.jl` kernels, thus maintaining `Oceananigans.jl`'s bespoke performance, memory- and cost-efficiency on GPUs in `OceanBioME.jl`-augmented simulations.

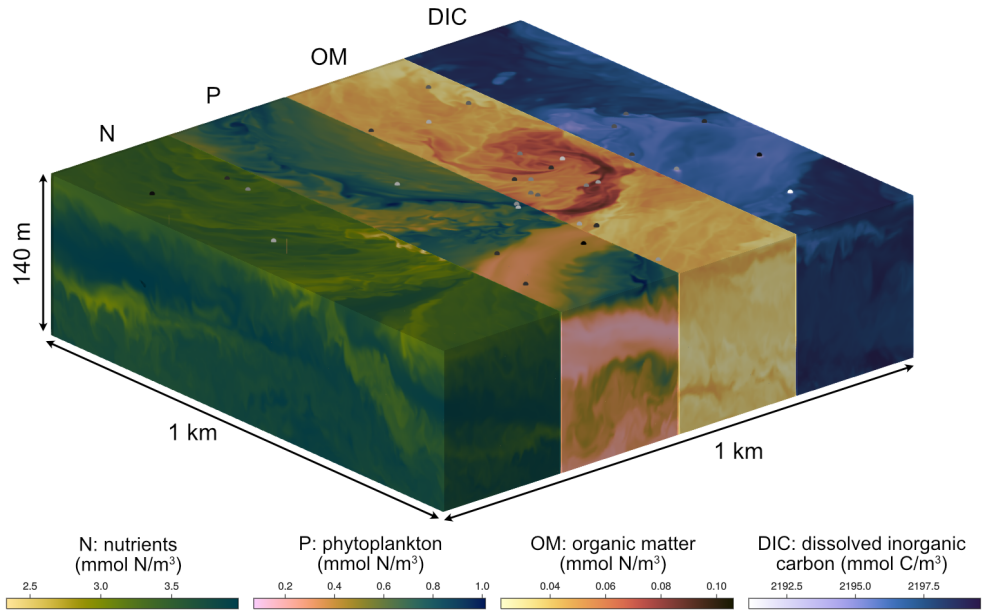

Figure 1: In this simulation of baroclinic instability in the Eady problem, a background buoyancy gradient and corresponding thermal wind generates a sub-mesoscale eddy, roughly following the setup of Taylor (2016). A submesoscale front develops on the periphery of the eddy where intense three-dimensional turbulence develops with inherently non-hydrostatic dynamics. To this physical setup, we added a medium complexity (9 tracers) biogeochemical model, some components of which are shown above. On top of this, we added particles modelling the growth of sugar kelp, which are free-floating and advected by the flow, and carbon dioxide exchange from the air. Thanks to Julia's speed and efficiency the above model (1 km $\times$ 1 km $\times$ 140 m with 512 $\times$ 512 $\times$ 64 grid points) took about 2 hours of computing time to simulate 10 days of evolution on an Nvidia A100 GPU. The figure shows the domain with the colour representing the concentration of various biogeochemical tracer fields: nutrients (nitrate and ammonia), phytoplankton, organic matter (dissolved and particulate), and dissolved inorganic carbon. Darker colours represent higher values. High concentrations of organic matter develop near the centre of the eddy. Small-scale variability and enhanced subduction can also be seen along a sub-mesoscale front. Points on the surface represent the kelp particle positions, with the colour representing the range of frond size. Figure made with `Makie.jl` (Danisch & Krumbiegel, 2021).

`OceanBioME.jl` is built with a highly modular design that allows user control and customization. There are three distinct module types implemented in `OceanBioME.jl`:

- First, we provide tracer-based ecosystem modules in `AdvectedPopulations` as a set of coupled ordinary differential equations (ODEs) which evolve the concentration of the tracer. These equations can be solved by `OceanBioME.jl` as box models. This is useful for both testing and running simple educational models and for understanding the behaviour of biogeochemical models without the effects of the physical dynamics. The same tracer equations can then be seamlessly integrated into `Oceananigans.jl` to further include the effects of advection and diffusion.
- The second module type is `Individual` "biologically active" particles. These consist of individual-based models solved along particle paths, which can be coupled with the tracer-based modules and physics from `Oceananigans.jl`. The biologically active particles can be advected by the currents, and/or they can move according to prescribed dynamics. For example, migrating zooplankton or fish can be modelled with biologically active particles and `OceanBioME.jl` allows these to interact with tracer-based components such as phytoplankton or oxygen.
- The `AdvectedPopulations` are supported by `Boundaries` modules which are easy to apply and provide information at the top and bottom of the ocean. We have implemented comprehensive air-sea flux models (e.g. Wanninkhof, 1992) within the `GasExchange` submodule to calculate carbon dioxide and oxygen flux at the sea surface, and sediment models (e.g. Soetaert et al., 2000) which calculate fluxes of carbon and oxygen at the seafloor.

We currently provide a simple Nutrient-Phytoplankton-Zooplankton-Detritus (NPZD) model (Kuhn et al., 2015), and an intermediate complexity model, LOBSTER (Lévy et al., 2005) and we have created a straightforward "plug and play" framework to add additional tracers such as carbonate and oxygen chemistry systems and additional forcing. A key feature of this package is the ability to easily modify the model equations or add different formulations, allowing exploration and experimentation. If a user wanted to implement a different model they could use the existing ones as a template and modify only a few lines of code where the ODEs are defined as functions. The user can then insert their model into our abstracted framework to couple the model with the other components such as light attenuation and sediments. We provide a detailed tutorial describing how to do this, which also serves as a description of how our models are created. This framework is made possible by our contributions to `Oceananigans.jl`, adding a streamlined user interface to swap biogeochemical models with no modification to other model configurations. Our interface also facilitates rapid prototyping, as models can be implemented and swapped easily by just extending a few key functions. This flexibility and ease-of-use is unmatched in existing biogeochemical models.

`Oceananigans.jl` includes several dynamical cores which include a fully non-hydrostatic model capable of large-eddy simulations (LES) and a free-surface hydrostatic model. This allows `OceanBioME.jl` to be used to simulate biogeochemistry across a vast range of scales using the same BGC model formulation. As an example, Figure 2 shows the annual average surface phytoplankton concentration from a near-global model NPZD model.

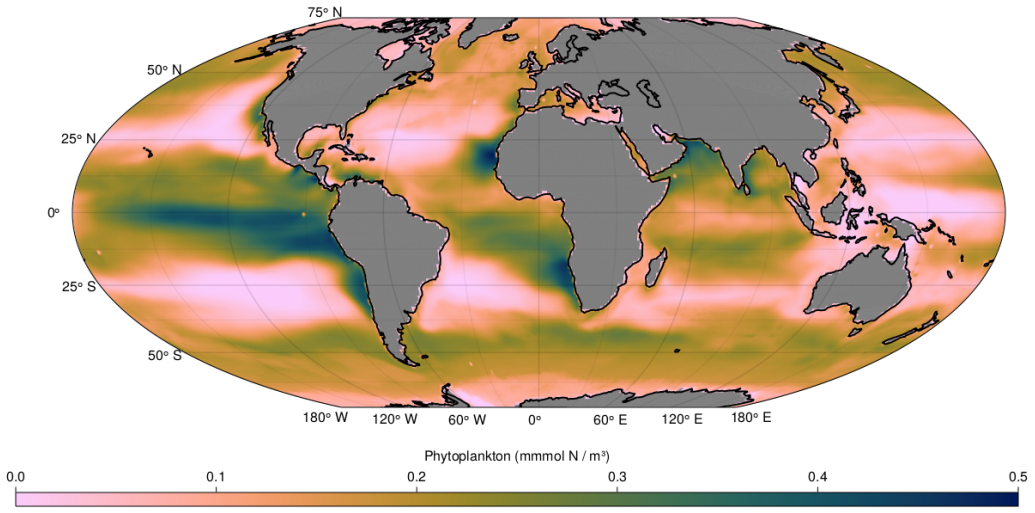

Figure 2: The annual average surface phytoplankton concentration from a near-global NPZD model run. Although the model is uncalibrated and is missing some important processes (e.g. river input), it reproduces the large-scale patterns reasonably well. This simulation used 1° horizontal resolution and 48 (irregularly spaced) vertical points. It took around 45 minutes per year to run on an Nvidia A100 GPU when integrating the physics, or less than 5 minutes per year when using pre-calculated velocity fields. Figure made with `Makie.jl` (Danisch & Krumbiegel, 2021).

The biologically active particles built into `OceanBioME.jl` are particularly useful for OCDR applications. Accurate carbon accounting is essential for assessing the effectiveness of OCDR strategies. Biologically active particles can be used to track carbon from a particular source while accounting for interactions with its surroundings. Biologically active particles can also be used to model OCDR deployment strategies including seaweed cultivation, alkalinity enhancement, and marine biomass regeneration. `OceanBioME.jl` currently includes an extended version of the sugar kelp model presented by Broch & Slagstad (2012) as an example of the utility and implementation of these features. Figure 3 shows a simple column model with an OCDR intervention (macroalgae growth) added after a warm-up period, which increases the carbon export of the system.

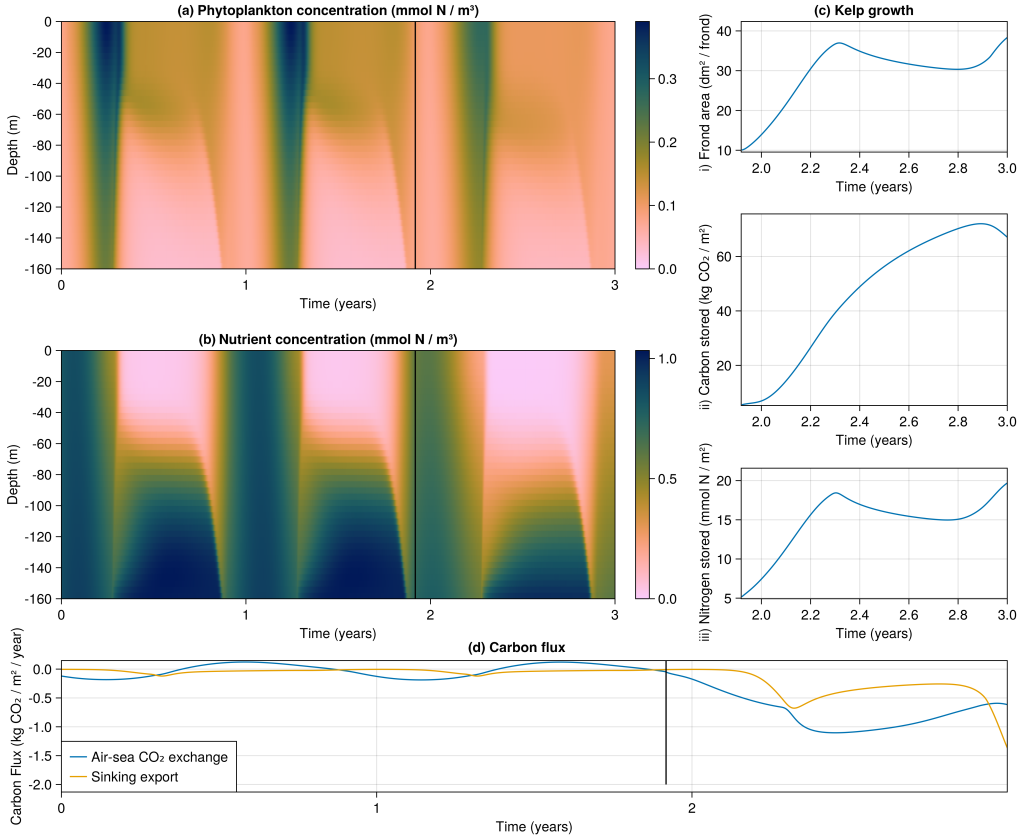

Figure 3: Results of a 1D model, forced by idealised light and mixing, which qualitatively reproduces the biogeochemical cycles in the North Atlantic. Kelp (500 frond / m^2 in the top 50 m of water) is added in December of the 2nd year (black vertical line) which causes an increase in air-sea carbon dioxide exchange and sinking export, as shown in panel (d). Panel (a) shows the phytoplankton growth cycle which also changes in response to the altered nutrient dynamics visible in panel (b). Panel (c) shows the kelp growth with the front size in (i), the carbon storage in (ii), and the nitrogen storage in (iii). Figure made with `Makie.jl` (Danisch & Krumbiegel, 2021).

The implementation of `OceanBioME.jl` models allows for seamless integration with data assimilation packages, such as `EnsembleKalmanProcesses.jl` (Dunbar et al., 2022). This enables rapid calibration of model parameters and provides a powerful utility for integrating observations and models, with the potential to improve model skill and identify key sources of uncertainty.

A key metric for the validity of biogeochemical systems is the conservation of elements such as carbon and nitrogen in the system. We therefore continuously test the implemented models in a variety of simple scenarios (i.e. isolated, with/without air-sea flux, with/without sediment) to ensure that conservation conditions are met, and we will continue to add tests for any models. Additionally, we check `OceanBioME.jl` utilities through standard tests such as comparison to analytical solutions for light attenuation, and conservation of tracers for active particle exudation and sinking. We plan to expand the features of `OceanBioME.jl` in the future and very much welcome user contributions.

Finally, this software is currently facilitating multiple research projects into ocean CDR which would have been significantly harder with other models. For example, Chen et al. (in prep.) uses the active particle coupling provided to investigate the effects of location and planting density of kelp in the open ocean on their carbon drawdown effect, as in the example above. Strong-Wright et al. (in prep.) uses the coupling of both the biogeochemistry and easy interface

to couple the physics to study flow interactions with a fully resolved giant kelp forest model including the effects on nutrient transport and distribution.

Examples

Example	OceanBioME features utilised	Code location
Sub-mesoscale eddy (Figure 1)	LOBSTER biogeochemical model[1] with carbonate model active, CO_2 exchange with the air[2], Light attenuation[3], mass conserving negativity protection[4]	`paper/figures/eady.jl`, similar to `examples/eady.jl`
Near-global proof of concept (Figure 2)	Light attenuation[3], NPZD model[5]	github.com/Ocean-BioME/GlobalOcean-BioME.jl/releases/tag/v0.0.1
Idealised 1D model with kelp individuals (Figure 3)	LOBSTER biogeochemical model[1] with carbonate model and variable Redfield ratio for organic components active, CO_2 exchange with the air[2], light attenuation[3], mass conserving negativity protection[4], and Saccharina Latissima (sugar kelp) model[6]	`paper/figures/column.jl`, similar to `examples/column.jl` and `examples/kelp.jl`

[1] `LOBSTER`
[2] `GasExchange`
[3] `TwoBandPhotosyntheticallyActiveRadiation`
[4] `ScaleNegativeTracers`
[5] `NutrientPhytoplanktonZooplanktonDetritus`
[6] `SLatissima`

Acknowledgements

We want to thank the Climate Modeling Alliance team and `Oceananigans.jl` contributors for their fantastic project. We also thank `JuliaGPU` contributors for `KernelAbstractions.jl` and `CUDA.jl` which make this code possible. Additionally, we are very grateful for the support and funding of the Centre for Climate Repair, Cambridge and the Gordon and Betty Moore Foundation. The authors would also like to thank Samantha Deane and the Kelp Forest Foundation for providing invaluable support, connections, and many helpful discussions. N.C.C. acknowledges support from the Australian Research Council DECRA Fellowship DE210100749. Finally, we would like to thank the reviewers Iury Simoes-Sousa and Francis Poulin for their very helpful comments.

References

Besard, T., Foket, C., & De Sutter, B. (2019). Effective extensible programming: Unleashing julia on GPUs. *IEEE Transactions on Parallel and Distributed Systems*, *30*(4), 827–841.

https://doi.org/10.1109/TPDS.2018.2872064

Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A fresh approach to numerical computing. *SIAM Review*, *59*(1), 65–98. https://doi.org/10.1137/141000671

Broch, O. J., & Slagstad, D. (2012). Modelling seasonal growth and composition of the kelp Saccharina latissima. *Journal of Applied Phycology*, *24*, 759–776. https://doi.org/10.1007/s10811-011-9695-y

Churavy, V., Aluthge, D., Smirnov, A., Samaroo, J., Schloss, J., Wilcox, L. C., Byrne, S., Besard, T., Waruszewski, M., Ramadhan, A., & al., et. (2023). *JuliaGPU/KernelAbstractions.jl: v0.9.8*. https://doi.org/10.5281/zenodo.8176369

Danisch, S., & Krumbiegel, J. (2021). Makie.jl: Flexible high-performance data visualization for Julia. *Journal of Open Source Software*, *6*(65), 3349. https://doi.org/10.21105/joss.03349

Dunbar, O. R. a., Lopez-Gomez, I., Garbuno-Iñigo, A., Huang, D. Z., Bach, E., & Wu, J. (2022). EnsembleKalmanProcesses.jl: Derivative-free ensemble-based model calibration. *Journal of Open Source Software*, *7*(80), 4869. https://doi.org/10.21105/joss.04869

Friedlingstein, P., Jones, M. W., O'Sullivan, M., Andrew, R. M., Bakker, D. C. E., Hauck, J., Le Quéré, C., Peters, G. P., Peters, W., Pongratz, J., Sitch, S., Canadell, J. G., Ciais, P., Jackson, R. B., Alin, S. R., Anthoni, P., Bates, N. R., Becker, M., Bellouin, N., … Zeng, J. (2022). Global Carbon Budget 2021. *Earth System Science Data*, *14*(4), 1917–2005. https://doi.org/10.5194/essd-14-1917-2022

Kuhn, A. M., Fennel, K., & Mattern, J. P. (2015). Model investigations of the North Atlantic spring bloom initiation. *Progress in Oceanography*, *138*, 176–193. https://doi.org/10.1016/j.pocean.2015.07.004

Lévy, M., Gavart, M., Mémery, L., Caniaux, G., & Paci, A. (2005). A four-dimensional mesoscale map of the spring bloom in the northeast Atlantic (POMME experiment): Results of a prognostic model. *Journal of Geophysical Research C: Oceans*, *110*, 1–23. https://doi.org/10.1029/2004JC002588

National Academies of Sciences, Engineering, & Medicine. (2022). *A research strategy for ocean-based carbon dioxide removal and sequestration*. The National Academies Press. https://doi.org/10.17226/26278

Ramadhan, A., Wagner, G. L., Hill, C., Jean-Michel, C., Churavy, V., Souza, A., Edelman, A., Ferrari, R., & Marshall, J. (2020). Oceananigans.jl: Fast and friendly geophysical fluid dynamics on GPUs. *Journal of Open Source Software*, *5*(53), 2018. https://doi.org/10.21105/joss.02018

Simoes-Sousa, I. T., Tandon, A., Pereira, F., Lazaneo, C. Z., & Mahadevan, A. (2022). Mixed layer eddies supply nutrients to enhance the spring phytoplankton bloom. *Frontiers in Marine Science*, *9*, 825027. https://doi.org/10.3389/fmars.2022.825027

Soetaert, K., Middelburg, J. J., Herman, P. M. J., & Buis, K. (2000). On the coupling of benthic and pelagic biogeochemical models. In *Earth-Science Reviews* (Vol. 51, pp. 173–201). https://doi.org/10.1016/S0012-8252(00)00004-0

Taylor, J. R. (2016). Turbulent mixing, restratification, and phytoplankton growth at a submesoscale eddy. *Geophysical Research Letters*, *43*, 5784–5792. https://doi.org/10.1002/2016GL069106

Wanninkhof, R. (1992). Relationship between wind speed and gas exchange over the ocean. *Journal of Geophysical Research*, *97*, 7373–7382. https://doi.org/10.1029/92JC00188