

SYMBIOSES: A Modular Framework to Assess Oil Spill Impacts on Commercial Fish Stocks

Achim Randelhoff^{1*}, Mathias Bockwoldt⁴, Håvard G. Frøysa², Daniel Howell², Raymond Nepstad³, Ole Jacob Broch³, JoLynn Carroll¹, Gro Harlaug Refseth¹, Frode B. Vikebø^{2,5} & Geir Morten Skeie¹

¹ Akvaplan-Niva, FRAM – High North Research Centre for Climate and the Environment, 9296 Tromsø, Norway

² Institute of Marine Research, PO Box 1870 Nordnes, 5817 Bergen, Norway

³ SINTEF Ocean, PO Box 4762 Torgarden, 7465 Trondheim, Norway

⁴ Department of Geosciences, UiT The Arctic University of Norway, 9037 Tromsø, Norway

⁵ Geophysical Institute, University of Bergen, PO Box 7830, 5020 Bergen, Norway

*Correspondence:

Achim Randelhoff

ara@akvaplan.niva.no

Keywords:

Eulerian, Lagrangian particle tracking, hydrodynamics, dispersion model, oil spill, fish early life stages, fish spawning stock biomass

ABSTRACT

Abstract number: PAPER333s1

The SYMBIOSES (“System for biology based assessments”) project’s consortium has more than ten years’ experience with simulating impacts of large oil spills on the recruitment and biomass of commercially important fish stocks in the Northern North Atlantic. Its backbone is a framework that couples a high resolution 3-D biological-physical ocean circulation model to an oil module and a model for fish early life stages (ELS), including trophic, oil, and copepod mortality effects using a Dynamic Energy Budget. A second module is applied offline (post-hoc) and evaluates impacts of larval mortalities on harvestable biomass for six commercial fish species in the North Atlantic. While the SYMBIOSES project so far has focused on the North Atlantic, the framework is completely modular and can be applied to any region where suitable models exist. It defines interfaces for arbitrary models both in Eulerian (grid-based) and Lagrangian (particle-based) reference frames, providing a default grid that is independent from the underlying grids of each model plugin. SYMBIOSES is further parallelized across multiple CPU (Central Processing Unit) cores, which can drastically speed up calculation in High Performance Computing settings. Finally, we showcase the SYMBIOSES framework using model results from parameter sweeps, demonstrating sensitivity of fish recruitment and Spawning Stock Biomass to oil spill type, location, and release rate.

INTRODUCTION

Ocean areas are subject to multiple industrial scale uses, such as offshore petroleum/gas, fisheries, aquaculture, or shipping, each with their own associated set of risks. Separating the impacts of each individual use on the marine ecosystem quickly becomes intractable, and we instead need a holistic approach to support ecosystem-based management. Coupling existing numerical spatially explicit models in a common framework can be an effective avenue to achieving such holistic modelling capacity, going beyond the capabilities of each individual

component model, as we have done in SYMBIOSES (“System for biology based assessments”) (Carroll et al., 2022).

The guiding design principle for SYMBIOSES was that the method for linking models should interfere as little as possible with the internal architecture of the individual environmental models. Otherwise, there is a risk that further external development of these models may be onerous to integrate into SYMBIOSES. The framework is therefore comprised of independent software modules coupled via a modular framework. Strict modularity is achieved through well-defined interfaces, exposing generic model features, rather than the underlying model software. This allows software modules to be interchanged transparently, without any modifications to the rest of the system.

METHODS

Before explaining the general abstractions made in the SYMBIOSES framework, we will first outline how it presently works in practice in order to motivate some of the design choices and terminology. The various independent models (e.g., ocean, atmosphere, oil, larvae, and ecotoxicology) are held together by the framework manager (Figure 1). An adult fish population model is run offline based on model outputs like larval mortalities and is thus not managed by the framework manager. Table 1 details the parameters that need to be passed between the models, which is exactly the main task of the framework manager, alongside advancing models in time.

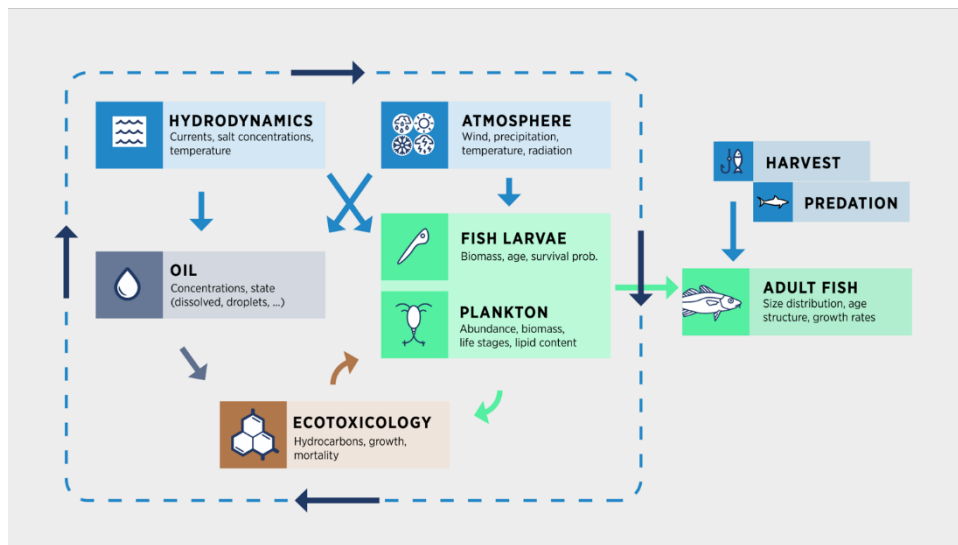


Figure 1: The key components of the system architecture are five independent models (to the left), a software library and the driver program, which initiates and directs computations. GADGET, an adult fish population structure model (Howell and Bogstad, 2010) is run post-hoc based on overall larval mortalities, usually as part of assessments of oil spill scenarios within SYMBIOSES.

Framework

The central component in the SYMBIOSES framework is a class library called `libsymbioses`, which acts as a junction box between models. The library provides abstract interfaces to model properties and functionality. Coupling a new model to SYMBIOSES entails implementing the relevant interfaces for the model. The framework library also handles the registration of models through a plugin system. All registered models, other plugins and globally needed data are maintained by the framework manager component. All communication between models must happen through the framework interfaces for rigorous model pluggability and interchange to hold.

The SYMBIOSES model interface design relies on an object-oriented class structure, defining the communication protocols between components. This allows factoring out common functionality in a flexible and clean manner. The interfaces are exposed to the programmer through libraries for Fortran, C, C++ and Python. This allows existing models to be easily coupled to SYMBIOSES. The Python layer provides the programmer with a convenient scripting interface to control SYMBIOSES workflow, input and output processing, visualization, communication and more.

Plugins

A plugin in the SYMBIOSES framework is a pluggable data provider. Plugins are registered with a feature that uniquely identifies the plugin in the system. Plugins also have an optional name that identifies the underlying software component providing the feature to the system. Every plugin also follows semantic versioning (major.minor.patchlevel) (Preston-Werner, 2013) to ensure compatibility between components. Plugins must also report their git revision, so that the exact state of the entire software system can be determined from outputs.

Models

Models are central data providers in the SYMBIOSES framework. Models represent aspects of the physical world (e.g. hydrodynamics, biology, chemistry). All models are time dependent and can be advanced with a specified time step. Models are plugins with a concept of time: They can be advanced using the *next()* class method. All models are plugins, but not all plugins are models.

By requesting access to registered features, models can depend on other models and plugins. However, models cannot request access to any particular implementation of a feature. This is strictly necessary to make the whole system model implementation agnostic.

The *SymbiosesModel* class defines properties and methods common to all models. The most important of these is the *next()* iterator, which advances the model by a specified time step. Models are either grid based (Eulerian) or particle based (Lagrangian). Since a particular model can provide both a Lagrangian and Eulerian interface to its data, the class specifies methods relevant for both.

The class also keeps track of model time parameters. Each model has its own notion of time, i.e. for how long it has been running. Models can specify start and end times relative to the global framework time, which specifies when the model should run (e.g. start/stop of an oil release).

Framework manager

The framework manager is the main SYMBIOSES component and primary entry point. The *SymbiosesFramework* class tracks registered plugins and models, and other globally defined information. Only one *FrameworkManager* can exist, and it is always globally accessible through the *getInstance()* method which returns a reference to the singleton object.

The SYMBIOSES framework distinguishes between plugins and models. A plugin is a general data provider that can be loaded registered with the *FrameworkManager* dynamically at runtime. Plugins are identified by their type (kind of model) and a valid feature. The model/feature pair defines a unique uniform resource name (URN) (e.g. *hydrodynamic/ocean*). This allows multiple models of the same kind to be registered, as long as the feature is unique (e.g. *aquatic/zooplankton/calanus/copepod* and *aquatic/zooplankton/calanus/nauplii*).

The framework manager keeps track of the global (reference) time and time step. The reference time step is the shortest time step required by any registered model and is usually dictated by the hydrodynamic model.

Model grids

Independent grid-based models rarely use the same grid definitions, making them difficult to couple. To be general, the SYMBIOSES framework only deals with physical properties in point wise locations defined by their coordinates. This ensures that there is no confusion in the interpretation of values, and that models can be coupled without any explicit, detailed information about each other.

However, the framework needs to transparently couple both grid-based (Eulerian) and particle-based (Lagrangian) models. To this end, the framework provides a reference grid for all models (Figure 2). The reference grid is used for plotting and representing particle quantities as concentrations or densities in a common framework. The grid is specified using projected coordinate matrices ($dx \times dy$) to allow for non-Euclidean metrics (e.g. polar stereographic projection).

From oil particles to concentration grids

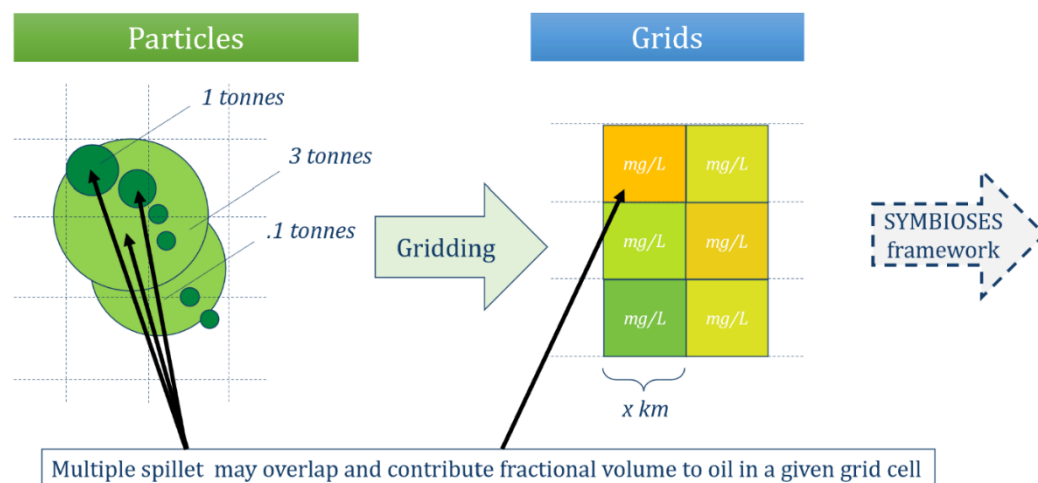


Figure 2: An illustration of how how Lagrangian (particle-based) and Eulerian (grid-based) models are reconciled within the SYMBIOSES framework. In this case, oil particles are

gridded on a pre-defined grid before being reported to the rest of the framework.

Scenario definition file and input

Once the overall system is set up with model input (in particular, ocean model and atmosphere input files), defining a concrete oil spill scenario requires few files. First, the symbioses driver program has a single, generic input file to control aspects of symbioses calculations. The sections in the input file, corresponding to the active, registered model features, are translated into specific model provider inputs by the framework.

Second, there may be configuration files needed by individual models. For instance, in our implementation that relies on OSCAR (“Oil spill contingency and response”) as the oil model (Daae et al., 2018; Nordam et al., 2019; Reed et al., 2004) we need one configuration file that defines the oil spill to OSCAR itself. This file contains parameters such as geographic coordinates, duration, spill depth, oil type, etc.

System outputs

The SYMBIOSES system produces outputs in the standard netCDF4 format, which is based on the HDF5 file format. This makes the outputs high performance and portable over a large variety of platforms. The netCDF outputs are CF-convention compliant (Hassell et al., 2017).

Parallelization

By splitting up the computational problem and utilizing many CPUs simultaneously in parallel, computational times can be significantly reduced. SYMBIOSES has been designed with parallelization in mind. All model interfaces and data accessors are re-entrant and thread safe.

To leverage parallelization to achieve this computational efficiency, the individual software modules must also be parallelized.

Containerization

The project consortium has further installed the SYMBIOSES framework inside a readily available Singularity image file (Kurtzer et al., 2021) using the official Intel oneAPI Docker image. Such an image enables the seamless transfer of workloads from one HPC (High Performance Computing) cluster to another and lessens the overhead of new research groups when setting up SYMBIOSES.

RESULTS

The current iteration of the SYMBIOSES framework has been used for more than 600 oil spill simulations since 2021, both for scenarios nominated by the research team and those nominated by industry partners. So far, SYMBIOSES has been set up for Norwegian use cases (Figure 3). This section showcases some of our findings and highlights SYMBIOSES' capabilities. For instance, Figure 4 shows the oil spill extent resulting from a particular simulated oil spill. Figure 5 shows larval mortality as result of a particular simulated oil exposure.




Figure 3: Map of the SYMBIOSES domain. The orange rectangle is the domain area for SYMBIOSES V1.0, a 1340 km by 460 km region. The entire area on the map is the new domain, covering the North Atlantic, North Sea, and Barents Sea up to the North Pole.

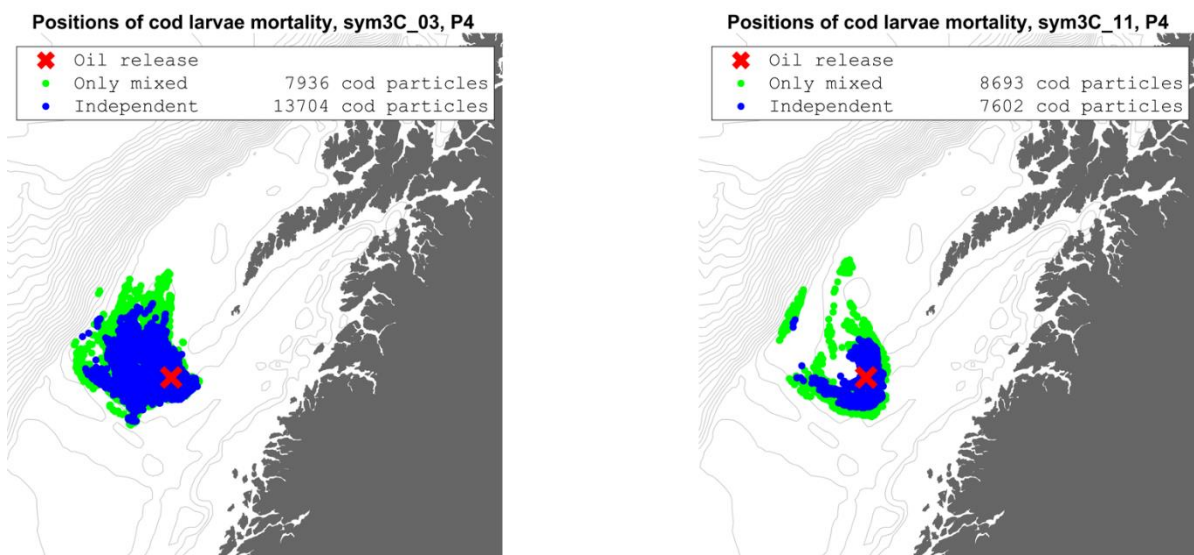


Figure 4. Illustration of region of effect for ELS mortality for two different simulations. The green and blue larvae particles experience mortality for mixed and independent statistical models (the latter from previous versions of SYMBIOSES). Red cross: Oil spill site.




Figure 5. Top: Time-volume of where Polycyclic Aromatic Hydrocarbons exceed a concentration of $1\mu\text{g/l}$ ($\text{km}^3 \cdot \text{days}$). A time-volume is defined as the temporal integral of the affected volume. Bottom: Early life stage mortality for DEB3, a specific set of toxicity parameters for different numbers (hence sizes) of particles.

Sensitivity to particle number

OSCAR, the oil dispersion model used in SYMBIOSES so far, is particle-based. Therefore, a given quantity of released oil must be split up into a given number of particles (called “spillets” in SYMBIOSES). This is relevant because basic parameters such as dissipation and dispersion depend on it. In addition, simulation time will also be longer for a larger number of particles (see next section). Indeed, Figure 5 demonstrates a marked effect of particle size on overall

early life stage (ELS) mortality. This dependency implies that caution must be taken when defining oil spill scenarios and interpreting such results.

Performance benchmarks: Current SYMBIOSES setup on Saga

Saga is a high-performance computing (HPC) cluster located in Trondheim, Norway, and administered by NRIS (Norwegian Research Infrastructure Services, <https://www.sigma2.no/>).

A common SYMBIOSES scenario runs from early spring (just before spring spawning of modelled fish species) and until 45 days after the termination of the modelled oil spill scenario.

On Saga, such a scenario can usually be run on single node (40 CPUs, each with 4GiB memory) within 4-7 days wall time. Figure 6 shows an example for how the time consumption per time step is distributed between the different models.

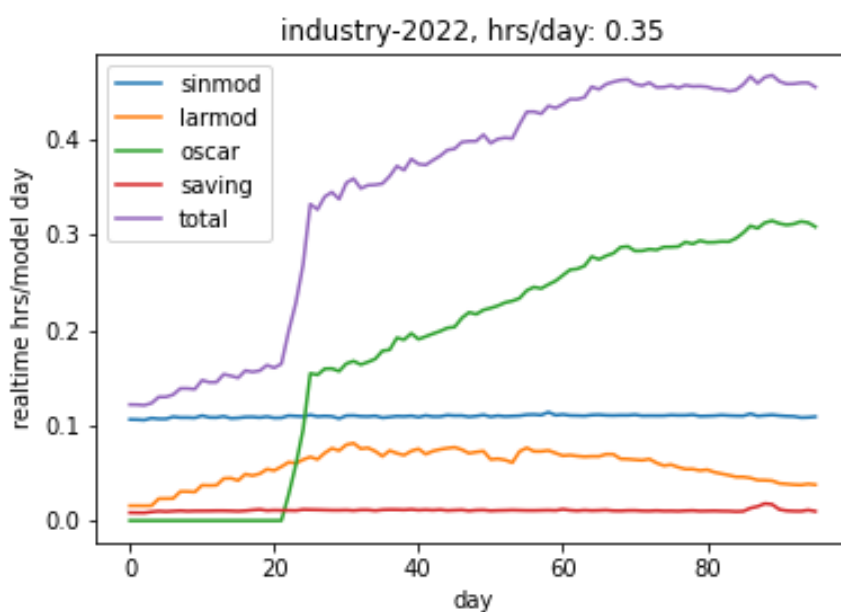


Figure 6: Contribution of individual models to each iteration, in wall time hours per modeled day. Oil is released just after day 20 of the simulation, marked by the rapid increase in OSCAR-related computation time. Note that OSCAR iterations use progressively more computation time as the number of oil particles increases. Also note the bump in LARMOD-related

computing time around days 20-60 (spring spawning). SINMOD, the ocean model (here run without copepods), and file writing (“saving”) operations take approximately constant time for each step. On average, this scenario needed 21 minutes of wall time per simulated day.

In the current setup, the ocean model (SINMOD) can be run both with and without Calanus. These copepods are an important food source to many fish larvae and hence important for ecological modelling as such. Modelling them is computationally expensive: Five replicas each of a specific scenario run with and without copepods took, on average, 107.4 and 36.6 hours to complete, respectively – a 200% overall performance hit. However, further tests showed that simulating Calanus did not offer different results for the larval mortality due to oil itself (Carroll et al., 2023) which factors out baseline mortality for example from food scarcity. Subsequent scenarios were hence modelled without Calanus as a larval food resource. This demonstrates the importance of domain knowledge for tuning computational efficiency for the question at hand when selecting submodules.

DISCUSSION

To our knowledge, SYMBIOSES is presently the only state-of-the-art model connecting oil effects to fisheries. A very early precursor of SYMBIOSES was described in, for example, (Spaulding et al., 1983). Some parts of it survived in OSCAR, which is presently used as the oil model in SYMBIOSES, and other parts have been replaced by newer models.

Future possible improvements can be framework specific or pertain to the underlying models. Among the framework specific improvements, computation may be considerably sped up if the ocean model and potentially larval advection models are run beforehand. Simulating a

given oil spill then only requires solving the oil advection and may be done on much shorter notice. However, this requires large storage space (on the order of tens of terabytes) to allow for sufficient temporal resolution. In addition, it also necessitates neglecting feedbacks of the other models into the ocean and potentially larval advection model, which may be reasonable for specific scenarios.

SYMBIOSES has been extensively used both for basic research (Broch et al., 2020; Carroll et al., 2022) and as part of industry-driven assessments (Carroll et al., 2023) of environmental risks at specific sites. The SYMBIOSES framework is maintained and operational, for application in future assessments when improved data on other fish species become available. As such, we anticipate that SYMBIOSES will be of value to ecosystem-based management, with the possibility of including additional components to assess impacts on marine ecosystems from other industries and activities, such as shipping, ocean wind farms, mariculture and carbon capture and storage.

ACKNOWLEDGMENTS

We thank Radovan Bast, Jonas Juselius, Dan Jonsson, and Ryan James Dillon for previous work on the code base. For the technical aspects, this paper draws on a SYMBIOSES III software documentation (The SYMBIOSES consortium, 2022).

SYMBIOSES (phase III) has been sponsored by Equinor Energy AS, ConocoPhillips Skandinavia AS, Wintershall DEA Norge AS, OMV Norge AS, Lundin Norway AS, Aker BP ASA, Vår Energi AS, with some initial funding provided by the Research Council of Norway.

TABLES

Table 1: Parameters transferred between individual sub-models of the framework.

From	To	Feature	Unit	Purpose
SINMOD	OSCAR	Water depth	m	Oil movement
		3D ocean velocity	m/s	
		2D wind velocity	m/s	Oil surface movement
		Temperature	°C	Oil chemistry; plume behavior
		Salinity	g/kg	
OSCAR	SINMOD	Oil concentration	ppb	Copepod ecotox
OSCAR	LARMOD	Oil concentration	ppb	Fish ELS ecotox
		Chemical properties of oil	MW, K_{ow}	
SINMOD	LARMOD	Water depth	m	ELS movement
		3D ocean velocity	m/s	
		Mixing coefficients	m ² /s	
		2D Wind Field	m/s	Turbulence used in prey calculations
		Temperature	°C	ELS development and movement
		Salinity	g/kg	ELS position (buoyancy)
		2D Calanus abundance	#/m ²	ELS development (prey for larvae)
		2D Calanus stages	#/m ²	
		Total Calanus biomass	g[C]/m ²	

LARMOD	GADGET	Relative reduction of ELS	%	Survival probability
--------	--------	------------------------------	---	----------------------

REFERENCES

Broch, O.J., Nepstad, R., Ellingsen, I., Bast, R., Skeie, G.M., Carroll, J., 2020. Simulating crude oil exposure, uptake and effects in North Atlantic *Calanus finmarchicus* populations. *Marine Environmental Research* 162, 105184. <https://doi.org/10.1016/j.marenvres.2020.105184>

Carroll, J., Frøysa, H.G., Vikebø, F., Broch, O.J., Howell, D., Nepstad, R., Augustine, S., Skeie, G.M., Bockwoldt, M., 2022. An annual profile of the impacts of simulated oil spills on the Northeast Arctic cod and haddock fisheries. *Marine Pollution Bulletin* 184, 114207. <https://doi.org/10.1016/j.marpolbul.2022.114207>

Carroll, J., Skeie, G.M., Vikebø, F., Broch, O.J., Frøysa, H.G., Howell, D., Nepstad, R., Randelhoff, A., Refset, G.H., Bockwoldt, M., Augustine, S., Langangen, Ø., 2023. SYMBIOSES III Final Report: Advancing simulation technologies for ecosystem-based assessment of major oil spills in ecologically sensitive and environmentally challenging areas. https://www.symbioses.no/wp-content/uploads/2023/09/SYMBIOSES-III-Final-Report_010923.pdf

Daae, R.L., Skancke, J., Brandvik, P.J., Faksness, L.-G., 2018. The sensitivity of the surface oil signature to subsurface dispersant injection and weather conditions. *Marine Pollution Bulletin* 127, 175–181. <https://doi.org/10.1016/j.marpolbul.2017.11.067>

Hassell, D., Gregory, J., Blower, J., Lawrence, B.N., Taylor, K.E., 2017. A data model of the Climate and Forecast metadata conventions (CF-1.6) with a software implementation (cf-python v2.1). *Geosci. Model Dev.* 10, 4619–4646. <https://doi.org/10.5194/gmd-10-4619-2017>

- Howell, D., Bogstad, B., 2010. A combined Gadget/FLR model for management strategy evaluations of the Barents Sea fisheries. *ICES Journal of Marine Science* 67, 1998–2004. <https://doi.org/10.1093/icesjms/fsq135>
- Kurtzer, G.M., Cclerget, Bauer, M., Kaneshiro, I., Trudgian, D., Godlove, D., 2021. Hpcng/singularity: Singularity 3.7.3. Zenodo. <https://doi.org/10.5281/ZENODO.1310023>
- Nordam, T., Beegle-Krause, C., Skancke, J., Nepstad, R., Reed, M., 2019. Improving oil spill trajectory modelling in the Arctic. *Marine Pollution Bulletin* 140, 65–74. <https://doi.org/10.1016/j.marpolbul.2019.01.019>
- Preston-Werner, T., 2013. Semantic Versioning 2.0.0. <https://semver.org/>
- Reed, M., Daling, P., Lewis, A., Ditlevsen, M.K., Brørs, B., Clark, J., Aurand, D., 2004. Modelling of dispersant application to oil spills in shallow coastal waters. *Environmental Modelling & Software* 19, 681–690. <https://doi.org/10.1016/j.envsoft.2003.08.014>
- Spaulding, M.L., Saila, S.B., Lorda, E., Walker, H., Anderson, E., Swanson, J.C., 1983. Oil-spill fishery impact assessment model: Application to selected Georges Bank fish species. *Estuarine, Coastal and Shelf Science* 16, 511–541. [https://doi.org/10.1016/0272-7714\(83\)90083-5](https://doi.org/10.1016/0272-7714(83)90083-5)
- The SYMBIOSES consortium, 2022. SYMBIOSES III technical documentation 1.9.0. (Available from authors on request)