Review of the NEMO Scientific Advisory committee on the first draft of the NEMO development strategy targeting 2023-2027

20/01/2022

Members:

Nils Wedi (chair)
Magdalena Balmaseda
Annalisa Bracco
Hans Burchard
Gokhan Danabasoglu
Sergey Danilov
Oli Fuhrer
Marilaure Grégoire
Stephen Griffies
Jean-Francois Lemieux
Bill Lipscomb
Baylor Fox-Kemper
Stephane Popinet

Secretary:
Mike Bell

Executive Summary

The SAC would like to congratulate the NEMO consortium on the achieved milestones of the past development strategy and the recognition of the NEMO consortium in the international community. Commended efforts include, among others, the developments on time-stepping efficiency, new bulk formulations, the use of OSMOSIS observations and addressing known biases, and the new sea ice model SI3.

The SAC also appreciates the interactive process of the formulation for the next NEMO development strategy.

A "multifarious" community model is mentioned in the document (chapter 9). It is certainly appreciated that NEMO necessarily covers a range of diverse topics to enable research. One question was regarding the scope of NEMO, e.g. to simulate coastal and estuarine marine ecosystems in the future. The SAC appreciates also the thorough overview and detailed discussions of the technical points in the different sections,

- Recognising and welcoming the questionnaire for identification of NEMO development drivers
- Recognising the potential challenges and different demands for a research community model on the one hand and supporting operational marine, weather and climate services on the other hand.
- Recognising the importance of upgrading and developing a shared CI/CD infrastructure, encouraging external software contributions (also from outside the consortium ?!)
- Recognising the importance of the issues with the vertical coordinate in the ocean
- Recognising engagement in new research areas with ML/AI, with a priority to provide enabling infrastructure linking NEMO and AI/ML technologies

- Recognising the importance of being able to flexibly deploy NEMO components on a range of emerging hybrid HPC architectures as well as cloud computing infrastructures.
- Recognising the worthwhile efforts on supporting reduced and mixed precision simulations
- Recognising the more general notion of cost/benefit considerations of the entire NEMO workflow
- Recognising the engagement in community support
- Recognising the desire to analyse and actively improve the Carbon footprint associated with activities involving NEMO

The SAC strongly recommends formulating near term and longer term science development priorities, e.g. a short (2 years), medium (4 years) and a longer term (10 year) horizon and through this clearly connect these development priorities with an overall direction of travel, and potentially measurable milestones. This seems particularly important with available resources in mind.

The SAC also strongly recommends to leverage developments with an external developer community, with specific suggestions on building the CD/CI environment also with external to the consortium developers/contributors in mind, seeking CICE and SI3 collaboration, engaging the science community, e.g. via building interfaces to community parametrisation initiatives (e.g. CVmix, GOTM) for fostering MIP studies and relevant process observations collected as “truth” (e.g. OSMOSIS), engage more on HPC adaptation with others and describe the interaction with international initiatives (e.g. WMO working groups, DestinE, Ocean digital twin), and finally provide a strategic plan how to engage younger university students into NEMO code developments given the trends towards other programming languages and HPC environments.

There are also a range of individual comments on each topic below which the SAC would like the NEMO developers to consider in the next revision.

<u>Main development priorities, trends and risks</u>

- Is the development strategy sufficiently ambitious or should the developers actively consider planning a progressive transition towards an alternative dynamical core ocean kernel? The sea ice discussion is a notable exception and these considerations could extent more widely, e.g. considering modular splitting of physics/thermodynamics and dynamics to prepare for alternative grids and dynamical kernels
- Efficiency of computing and the HPC portability and flexibility in the use of an increasingly diverse HPC landscape is a risk with too strong reliance on “minimal changes” required adaptation strategies.
- It would appear that accelerating the efficiency of tracer transport (e.g. grid coarsening, alternative advection schemes, ML/AI based ?) may be a priority
- Is the questionnaire sufficient for a gap analysis or should there be a more thorough analysis of limitations (structural or missing components)?
- Icebergs from a sealevel point of view have perhaps much lower priority, whereas Antarctica may be a higher priority ?
- Is there a trend towards “stand-alone” rather than “integrative developments” for coupled simulations (e.g. ABL, reduced-complexity ocean mixed layer for sea-ice) ?

<u>Structural comments</u>

- Consider structuring the vision for future developments in terms of a (short) longer term strategy, in the time frame of 10 years, which will be implemented in e.g. rolling (and more detailed) four-year implementation plans & priorities. (this is how ECMWF operates)

- In reading through the 2018-2022 report, I see that many items listed in the 2023-2027 draft were also central to the earlier report. I suggest each chapter of the new report should provide a concise summary of what has been completed (or at least started) based on the previous report. In this way, the new report becomes a vision/wish document that also articulates the progress relative to the previous report. Otherwise, I am unable to know how realistic the goals are. Many groups want many things, but for the report to be useful there needs to be a firm grounding in realism along with specific prioritization with champions on the hook.
- The current version is simply too long. As you also indicated, while some sections seem more aligned with a development strategy, many sections seem to be just providing a wish list, perhaps without much justification. There have to be better articulations of immediate needs vs. medium-term needs vs. nice-to-haves vs. etc.
- I would argue that "less is more" and eliminate many of the listed items in the document and focus on what can really be achieved in the next five years. I liked section 3 which acknowledged what lessons were learnt from the previous exercise and articulates a much reduced scope for this document. Essentially just focusing on 1 or 2 main items. I suggest using this framework for other sections as well. Items not included can be kept in a web-based living document.
- It is unclear how this development strategy addresses the needs and priorities of the science plan that is referred to in a few places. The current document feels too heavy on the individual wish lists from the consortium members. So, more stress on common needs and goals is needed. Tying to the science plan can help remedy this issue.
- I found the document very "code-centric" and in some respects lacking in broader perspective. Following a militaristic analogy it is more about "tactics" than "strategy". By "code-centric" I mean that it is easy to believe that NEMO is essentially the code and the software infrastructure surrounding it. I believe that NEMO (and any other large piece of software) is much more than that: it is the network of developers, users, maintainers, their interactions (including other groups), their history and expertise etc. In many ways, this "people-centric" view is much more important than the code itself (which could/should be seen almost as an accessory).
- Similarly, I found that even technical references to the developments in other models were often short and relatively shallow. For example, I was surprised that no references at all were made to projects aiming for generic modelling of vertical mixing such as CVMix or GOTM. This is only one example.
- There is a lack of an overarching strategic overview (why these choices?) on I/O , AGRIF, HPC adaptation, and interfaces to old as well as novel emerging code components
- AGRIF could be considered an option as part of the dynamical kernel strategic discussion ? Considering also CFL limitations of unstructured grids.

<u>Other general comments</u>

- Note that the monitoring activities, which were considered in the 2018-2022 strategy, were not mentioned in the current one. The monitoring activities are prime applications of NEMO.

Carbon footprint

- I appreciate the commitment to reducing the HPC carbon footprint. My institution hasn't yet given this problem the attention it needs. NEMO could set a positive example for us and other modeling groups.
- In the discussion of carbon footprint in Section 2.2, I suggest some mention be given to the source for the energy to run the computer. An inefficient model running with wind power has a tiny carbon footprint, whereas an efficient model running with coal power has a relatively huge footprint. So in parallel to making a more efficient code, it would be useful for the NEMO community to help drive the use of computers using low carbon fuels. Raising these points among scientific users and developers will eventually bubble up to the managers who are making decisions about computer resources.
- The discussion on energy efficiency or greenness seems rather incomplete. It seems to reflect some numbers from countries that primarily depend on nuclear energy. If a supercomputer depends more on non-renewable resources, then the travel footprint is not really even on the map. For greenness, one needs to consider how to properly dispose of nuclear waste, which is usually not discussed. There are also additional challenges. For example, how do you balance an energy efficient code with a development path that plans to include more features which will likely increase the cost of the model per degrees of freedom?
- Consider also I/O workflows and data storage
- CO2 should not be a driver of developments, but rather computational cost and energy efficiency (time-to-solution, energy-to-solution)

Ocean model kernel

- Another manifestation of this “code-centric” perspective regards the lack of planning and discussion of an eventual (complete) code rewrite of NEMO. It seems that this is not even considered as an option (i.e. this would not be Nemo anymore / this is just too costly). I believe that any code older than, say, 15 years should be considered for an entire rewrite, and Nemo is obviously much older than that. Such a code rewrite is indeed costly, but one needs to put this cost in perspective with the multiple costs induced by an old code base, such as:

    - obsolescence of the language/data structures/programming model and resulting inadaptation to current hardware
    - intrinsic modelling limitations induced by the language and/or poor initial choices
    - accumulation of (often obscure) workarounds designed to avoid the most drastic effects of the two points above
    - The cost of the rewrite can also be mitigated by planning a progressive transition when possible. This is precisely what I would expect a “development strategy” to be about.

- Does it actually make sense to have a model that covers all scales from global climate to estuarine dynamics? I do not know the answer, but this should be discussed.

- It seems to me that more thought should be given to the link between vertical coordinates and bathymetry representation: the Brinkmann penalization method is given too much emphasis.
- I was surprised to see so few references to Hycom (and the papers by Bleck) in discussions of numerical mixing / vertical coordinates.

- Vertical coordinate / vertical mixing / biogeochemical models etc.: given the natural decoupling between vertical and horizontal that this entails, it is a natural area where common libraries should be developed (as aimed for by e.g. CVmix and GOTM). There is very little (no?) discussion of this in the document.
- There is a lack of discussion on connecting vertical and horizontal discretisation.
- Are there methods existing / efforts planned to quantify numerical mixing?
- Why not trying to implement vertically adaptive coordinates?
    - Explain the strategic choice ALE vs fully Lagrangian
    - Open science question if "intermediate" approach viable
    - Cost increase and impacts on parallelisation related to above choices?
- Page 26: Interesting discussion on vertical grids. *What exactly is the vision for a vertical grid? Is it necessary to have a target grid, or is the vertically adaptive strategy more promising?*
- Page 25: "Our view is that code using such grids [i.e., unstructured] would constitute another model." *Does NEMO use curvilinear grids?*
- "There may be merit in using the alternative meshes for estuarine models *(what is this?),* but that is not the main focus of the NEMO consortium members."
- Although non-hydrostaticity is not mentioned as a primary concern, I believe this poses the broader issue of overall model consistency, which is not discussed in the document at the moment (from a general perspective). The following points could/should be discussed:
    - physical/mathematical consistency: across scales (e.g. hydrostaticity etc.), conservation properties (what is conserved, what is not, what should be improved), monotonicity/positivity etc.
    - numerical consistency: what are the main approximations in the numerical schemes? which are appropriate / which should be improved? order of approximation of various terms etc.
    - The only non-hydrostatic option discussed seems to be that of Auclair et al. using a very short (pseudo)-acoustic timestep. This seems to contradict the frugality/energy efficiency primary goal of new developments in Nemo.
- Tracer transport efficiency should be a development priority?
- I like the focus on three main issues (vertical coordinate, bathymetry, energy efficiency) for a 4 year plan, needs more info on what will be done to address ALE and the energy efficiency, It seems to me that for now the benefits of RK3 is the main content of addressing energy efficiency?

<u>Bathymetry and internal waves</u>

- Consider efforts to optimize the bathymetry and treatment of lateral boundary conditions using data-driven methods. [As an example, the Brinkman volume penalization approach for lateral boundary conditions lends itself to use observations to estimate porosity/permeability or for the envelop bathymetry. Along these lines, it would be worth investigating the use of adjoint methods can be used in combination with altimeter data to create a bathymetry/boundary condition configuration with the correct separation of Gulf Stream.]
- Addressing realism of internal wave processes. In this regard I also noticed a lack of interest in BBL parameterizations, but the internal wave field is likely to depend on the bottom topography (and layer) treatment. In my mind some of the mixing problems will not be solved without looking carefully at internal wave-induced upwelling, which cannot be represented correctly unless the BBL is considered. Most attention is at the surface, where of course we have more observations, but may not be enough for climate applications [In

terms of collaborations, the US has a call at the moment for global and regional improvements in internal wave treatment for ocean models, so there would be some movement in that direction here as well, on top of the large program run out of Stanford+ Scripps and others on ML for IW parameterizations]

Parametrisation of eddy closures

- There is a working group on eddy closures but it is not clear what the priorities are and how these fit to the use in the stated ocean model resolutions 1/4/,1/12,1/24/,1/36 as are used e.g. in marine and climate services.
- Here the implementation part should be very straightforward. Are there any indications that biharmonic GM or GEOMETRIC will lead to improvements? To what extent are they needed (NEMO is moving to higher resolutions)?

Surface Fluxes and vertical mixing

- The overall strategy does not become clear. There are many particulate options listed, some of which would probably not really work. There are probably KPP and GLS in there as major lines, and it seems that OSMOSIS will be added as a further line. Which concept will be used for which purpose? For some aspects, it will be difficult to add them to all strategies, e.g., the Langmuir Turbulence cannot easily be added to the GLS is the community experience.
- Encourage to continue efforts on the improved representation of the vertical mixing. Improved parametrizations will be important to capture the impact of unresolved mesoscale eddy fields in the ¼ of configuration, which will likely continue to be the work-horse configuration for many ensemble applications and for model development ?
- The use of the OSMOSIS observations and addressing known biases has helped to drive similar efforts in many other modelling groups.
- Continue with efforts on stochastic parameterization. In particular, consider the stochastic component in the development of new parameterizations such as those resulting from OSMOSIS, or advances in bulk formulations. The stochastic parameterizations are becoming increasing important for scientific and operational activities using ensembles.
- Is there any though to be closer to community efforts (CVMix)?
-

Sea-Ice

- The Sea Ice section describes SI3 as generally state-of-the-art but in need of some upgrades, such as dynamic-thermodynamic splitting, a more efficient tracer advection scheme, and better ice strength parameterizations. Some of these improvements have been present in CICE for a long time, which suggests that it's not straightforward to transfer methods and code from CICE to SI3 (and vice versa), leading to duplication of effort. Are there ways to work more closely with the CICE Consortium, for example by sharing code related to column physics/thermodynamics?
- There may be technical barriers and code diversity is important, but collaboration could be considered on selected aspects of sea-ice modelling between SI3 and CICE, e.g. interfacing to the Icepack physics package?
- Rheology: Results from the Sea ice rheology experiment (SIREX, Bouchat et al. in press, Hutter et al. in press) show that the four rheologies tested (MEB, VP, EVP, EAP) all lead to intersection angles (between fault lines) that are too wide as opposed to observations. I would suggest that the developers implement their code in a way that different yield curves

can be easily implemented so that this problem can be studied and eventually fixed. I would also suggest that they consider the plastic potential approach of Ringeisen et al. 2021.

- Again, Icepack is a community effort providing a lot of physical parameterizations. Can it be that a lot of stuff is already available?
- SI3 already has the grounding scheme (for ice keels) that we developed (Lemieux et al. 2016). We have recently implemented a more sophisticated scheme that uses the ice thickness distribution. As SI3 cube is also a multi-thickness category model, this new scheme could be implemented. There is a paper in review (Dupont et al.) in the cryosphere about it.
- So far our grounding schemes have only been tested assuming the ocean depth is fixed. A nice improvement and a nice study would be to also consider changes in the SSH.

Land ice / Ocean interactions

- The Land Ice section is comprehensive and well written. The Met Office is already a world leader in land ice – ocean coupling and is ideally positioned to build on that effort. I like how the authors have prioritized tasks, and I generally agree with the priorities. I have a few minor suggestions:

    - For ice-sheet and sea-level projections, ocean – ice shelf interactions are critical, whereas icebergs are probably secondary. For this reason, I might be inclined to put ice shelves in category 1 and icebergs in category 2. If iceberg model development (e.g., iceberg dynamics and thermodynamics) is in fact an equal priority, then the authors could better spell out the scientific rationale.
    - Sub-shelf cavity and boundary-layer parameterizations and ice – ocean coupling are still immature, so these could be framed as grand challenges requiring international collaboration. Several co-authors have participated in community projects such as MISOMIP and ISMIP6, which are great opportunities for collaboration. I also suggest working with other ESM groups such as CESM and GFDL, which have similar goals.
    - The plan envisions improvements in two-way nesting (AGRIF) for high-resolution coastal regions such as estuaries. To further motivate AGRIF development, the authors could mention the potential benefits of two-way nesting for resolving small-scale heat and mass exchange around the Antarctic continental shelf and in sub-ice-shelf cavities.

- Not clear what approach will be taken to provide sufficient vertical resolution near the ice-ocean interface of an ice shelf. Coordinates should be surface following similar to the approach of FESOM. Coupling should be done with GLS to get entrainment right at the bottom of the plume.

Tides

- I would have expected more discussion on the links with vertical coordinates / coastline representation / non-hydrostaticity.
- Possibly also discuss tides as part of the ocean model kernel ?

Marine Biogeochemistry

- Is it planned to use the FABM (Framework of Aquatic Biological Models, https://github.com/fabm-model/fabm) interface?

- Coupling between physics and biogeochemistry: is it intended to allow a two-way coupling (with biology influencing the ocean physics/dynamics)?
- Is it intended to develop a spectrally resolved light model? (Radiative Transfer Model) (e.g. to improve optical properties of the water column)
- Suspended particle matter (SPM) dynamics is not mentioned in the evolution. What are the plans? (connected to the question on the ambition/scope for estuarine modelling)
- Is there a fish model, what are the plans (online, offline, both)?

Two-way nesting capability (AGRIF and grid coarsening)

- Only some groups so far interested in AGRIF, it may be a good strategy to address CFL limitations of globally unstructured grids, but this could be elaborated on more if this is the strategic goal and discuss the potential in that context.
- Agrif is mentioned as one of main tools in reducing the carbon footprint and, in principle, as a tool that may incorporate/generalise the coarsening needed for BGC. However, its role in the strategy document is not as central as it needs to be, given the priority given to these tasks.
- Several issues mentioned but it is difficult to judge the severity or importance of these (nice to have, crucial, priority)
- There is a lot of hope on wider use of AGRIF, but I do not see that this challenge is properly reflected. Otherwise, the work discussed in this section is very appropriate.
- Continue grid coarsening efforts. The ability to support dual or multiple (e.g. coarse/fine) grids will benefit not only biochemistry applications, but also efficient model output and interface with assimilation (e.g. multi-resolution variational methods, or ensemble statistics). It is acknowledged that this is not a trivial task, since it is likely to require special efforts for the versatile interpolation of coastlines.
- Is interfacing of NEMO at its boundaries e.g. with unstructured grid models used to simulate river-estuarine systems in scope?
- Mesh adaptive methods: these are probably key to further the goal of reduced com-putational cost / enhanced energy efficiency. Not considering a code rewrite of the Nemo core drastically limits what can be done. As mentioned in the document, considering only AGRIF for this presents serious limitations (load-balancing, model consistency across refinement boundaries, etc.).
- On horizontal grids and local resolution enhancements. Consider exploiting the multi-polar grid capabilities for targeted enhancements in horizontal resolution. Additional poles could be possibly place around coastal areas closed to western boundaries, for instance. This could be an alternative approach to AGRIF? This may interface more easily with data assimilation and coupled ocean-atmosphere models which is a concern with the AGRIF approach. A discussion on this would be appreciated.

High performance computing

- I think the document contains a of important points from a HPC perspective and illustrates some of the important immediate struggles that all software in the Earth science domain has. The HPC strategy reads reasonable and contains - at least from my point of view - very little technological risk. This may be intentional and obviously entails the risk that other ocean models taking more risk may - if they are successful in their technological bets - become more popular. This may be a high-level question of the NEMO partners and steering committee to consider.

- There seem to be two strands for HPC, a more traditional (tiling, multi-core optimization, …) and a more “aggressive” (piping the existing Fortran code through PSyclone as a first step towards a DSL and automatic optimization). While these strands make sense in terms of being short-term and mid-term, they are in principle not very compatible, in the sense that the DSL compiler will have the responsibility to do some of these optimizations and if they are introduced into the code base the parser will actually have a harder time interpreting the existing code base. For a consortium strategy, it would be nice to have a clearer plan in place what achievements need to be met by the PSyclone version and - if they are met - how and when the transition to this version will be done by the consortium to this version.
- MPI + OpenMP (used as multi-core acceleration as well as for GPU offload) seems to be a core strategy, is the addition of openMP also via the Psyclone DSL ? If not how does this work together?
- Very little is said about the ecosystem of larger on-going initiatives and how the NEMO consortium plans to interact, profit or engage with these initiatives / projects (e.g. DestinE, Digital Twin Ocean, etc..)
- There appear to be several strategies to improve performance, but what I am missing as an introduction is an analysis of the computational time spent in the different NEMO components for a range of resolutions to further the argument of the articulated desire to analyse the cost / benefit of developments
- The document is fairly specific about target use-cases and simulation throughputs, but not very specific about what the current status is, so it is hard to judge how feasible / realistic / needed some of the HPC developments outline in the strategy are.
- I am missing a discussion on balancing complexity and compute performance for the range of target scales (e.g. model resolution) and the ability/recommendation to a scientifically sound reduction of complexity, in the light of CO_2 footprint reduction, but also in what can be reasonably initialised from observations
- The chapter seems isolated from the scientific ones, and I am missing a discussion on a strategy to address scalability issues together with domain experts, e.g. the mentioned non-parallel performance of the iceberg module, the (lack of) scaling performance of the sea-ice module, the free surface barotropic mode scalability, a general analysis of each component complexity and its scalability etc.
- With some of the high-resolution simulations that are targeted, output can become one of the dominant bottlenecks. Very little information is given how the NEMO consortium plans to deals with these challenges (on-line compression, in-situ, …)
- What is the cost of I/O at high resolutions and is XIOS a sufficient strategy for IO for both operational service and research environments?
- The use of optimised libraries is not mentioned, is it applicable for parts of NEMO or could NEMO be reformulated such that (e.g. matrix-matrix multiplication libraries may be used ?), this raises the further question on co-development plans of both algorithm and HPC adaptation
- The expected outcomes of some of the adaptations are not clearly formulated, what is the expected acceleration of the NEMO code, say a reference 1/12 global configuration in 2027 compared to 2023 in time to solution for 1 simulation day ?
- It is mentioned that the use of tiling did not bring benefits, yet the long term goal is tiling (I would agree, that this is an important development for increasing the flexibility on a range of architectures, but the text reads odd?).

- Adaptations for HPC: running efficiently on e.g. GPUs is indeed a vital issue for Nemo. As suggested above, I am not convinced that automatic code transformation (using tools such as Psyclone) can achieve this. The 2.3x speedup obtained on a single NVIDIA A100 GPU seems to agree with me. I don't think decent speedups can be obtained without a code rewrite (which would have multiple other benefits). One may also argue that developing a code such as Psyclone is also costly.
- Moving towards DSL --- the reliance on PSyclone. Is it decided, and what are the implications for all users of NEMO?
- Even if not a code rewrite, I am lacking a strategy for restructuring of code components and data structures to be able to map more flexibly to a hierarchy of memory (distributed, e.g. on device, on host, HBM, NVRAM) available on emerging architectures, especially if memory bandwidth is an issue with NEMO.
- I recommend to further pursue and continue to facilitate the reduced / mixed precision approach as an easy route for reducing the computational cost and individual energy footprint
- Other types of parallelism – this is not concrete enough.
- With the convergence of AI and HPC, the HPC landscape is changing again very rapidly and fundamentally. It is not unlikely, that some models running on HPC systems will look more like containerized services / components talking with each other. Also, with the end of Moore's Law specialization will only continue and a lot of the hardware design will revolve around architectures ideal for AI/ML workloads. For a strategy document it would be nice to at least have some of these developments reflected and some ideas of how the NEMO consortium plans to address them.

AI/ML

- Data assimilation, machine learning etc.: again here a code rewrite would allow to benefit from a large number of recent advances in these areas (automatic differentiation, etc.).
- Perhaps developments of the Brinkman / immersed boundary method could be combined with parameter learning in ML/AI, but is there sufficient data to constrain these ?
- Consider a general framework to couple the dynamical and data-driven model components. A framework for adding empirical data-driven terms to the model tendencies will benefit both data assimilation activities (e.g. IAU, bias correction), stochastic physics, and other terms arising from ML applications.
- NEMO could pursue publishing an ocean benchmark dataset for the ML/AI community (see also https://github.com/pangeo-data/WeatherBench)

Verification / Validation

- Page 6: Quite a diversity of views towards NEMO as an estuarine model. Ranging from just employing "estuarine boxes" towards fully estuarine models with drying & flooding, advection of turbulent quantities. Specifically NOC wants to see NEMO being extended towards estuaries.
- Are there any idealised estuarine test cases that are used to assess NEMO's ability to simulate estuarine dynamics? This would be my general question regarding test cases: are there certain standard test suites for specific model problems to see how NEMO performs on certain scales?

Software maintenance / Tools / Infrastructure

- Some of the discussion on git and SETTE in several places
- Consider introducing options for slim-restart files (e.g. a single time step, truncate precision, only independent variables), which can be used in ensemble applications when binary identical results are not needed.

Data assimilation interfaces

- The implications of different horizontal and vertical coordinates should take into account the implications for observation operators.
- There are other users of the observation operator outside the NEMO consortium. Their needs and feedback should be considered when planning major changes to this infrastructure, such as those envisaged for the JEDI compatibility.
- Continue efforts on maintenance and development of the IAU interface. The changes made to this interface should be clearly documented.
- Continue efforts on adjoint and tangent linear developments, which will benefit other applications beyond data assimilation.
- Consider a modular DA interface to minimize intrusion in the code. This may require re-structuring the NEMO code so there is a clear and generic interface for DA, e.g. the NEMO model callable as a subroutine to advance a given time-interval within a DA controlled structure [DA algorithms require a model trajectory, calculate observation departures, and reading and applying assimilation increments].
- Consider any developments on support for ensembles in data assimilation to be applicable or compatible with the use of ensembles in other applications (e.g. coupled forecasts).

NEMO Community & Support

- Consider the establishment of a working group devoted to training of users and developers.
- Consider mechanisms to more actively entrain the expertise from outside the NEMO consortium in the development of NEMO
- In section 17, it is indicated that NEMO is a community model. Yet, the document seems to be based on input solely from the consortium members. Please articulate how the broader community's input has been considered while drafting this document.
- I recognize that some of these aspects are discussed in the document, but I would have expected a much more detailed presentation and discussion of, for example:
    - how is the development effort distributed between different partners/institutions?
    - how does Nemo/its community positions itself relative to other ocean models/communities?
    - are any developments shared with these other communities?
    - should any of the points above be improved in 2023-2027? how?
- Following on the "people-centric" point of view, although this point is discussed in the document, I believe much more thought should be given to developing the user/developer community. This includes tools, events, documentation etc. making it easier for new users/developers to join and contribute. It seems to me that NEMO is still following a somewhat "old-style" centralized development perspective, which could be beneficially replaced with a much more flexible (and welcoming) decentralised perspective.
- In the text itself "a high barrier to entry to include new code in the main NEMO trunk" is mentioned in the text (p.31)

Additional References

Dupont et al., a probabilistic seabed-ice keel interaction model, the cryosphere, under review.

Bouchat et al., SIREX, Part I : scaling and statistical properties of sea-ice deformation fields, JGR Oceans, in press.

Hutter et al., SIREX, Part II : evaluating simulated linear kinematic features in high-resolution sea-ice simulations, JGR Oceans, in press.

Ringeisen et al., Non-normal flow rules affect fracture angles in viscous-plastic rheologies, The cryosphere, 15, 2873-2888, 2021.