Ocean Science Discussions

Ocean Sci. Discuss., author comment AC2
https://doi.org/10.5194/os-2021-87-AC2, 2021
© Author(s) 2021. This work is distributed under
the Creative Commons Attribution 4.0 License.

Reply on RC2

David S. Trossman et al.

Author comment on "Tracer and observationally derived constraints on diapycnal diffusivities in an ocean state estimate" by David S. Trossman et al., Ocean Sci. Discuss., https://doi.org/10.5194/os-2021-87-AC2, 2021

Our responses are in boldface:

•In this study, the authors used the ECCO framework to explore a few ways to better constrain the estimates of diapycnal diffusivity. Based on a set of sensitivity analyses, they investigated the impacts of including diapycnal diffusivity estimates that are obtained from microstructure measurements or inferred from CTD measurements, as well as dissolved oxygen measurements. They concluded that both ways could improve the presentation of diapycnal mixing in ECCO. I think this is potentially a very important paper, and I really appreciate the authors put tremendous efforts to address this interesting yet difficult question. The paper should eventually be published. However, before I recommend acceptance, I would like the authors to clarify a few questions and concerns I list below.

• **Thank you for taking your time to provide thoughtful comments.**

• My major concern is the definition of ECCO diapycnal diffusivity used in this study. From what I know about ECCO v4, there are many components involved in the calculation of vertical fluxes as well as diapycnal diffusivity. The diapycnal diffusivity at least consists of three parts: the background diffusivity, which was adjusted through the adjoint process; the parameterized part based on Gaspar et al. (1990); and convective adjustment. In this paper, the authors briefly described some of those terms. But it is still not clear to me what the exact definition of the diapycnal diffusivity the authors analyzed is. The combination of all or some of the components mentioned above? or just the adjusted background diffusivity? This information is critical for the interpretation of almost all the

results presented in this study. And the authors should make that information more explicitly presented.

• The MITgcm uses a quantity they call "diffkr" which is the diapycnal diffusion coefficient in their r-coordinate system and is referred to as a vertical diffusivity (not quite a diapycnal diffusivity). To transform this value into a diapycnal diffusivity that's equivalent to the observational product values, we need to subtract out the (3,3) entry of the along- isopycnal diffusivity tensor, which is not perpendicular to the isopycnal contours but parallel to them. After subtracting this from the vertical diffusivity, we are left with the diapycnal diffusivity. As for what this coefficient physically represents, it is the adjusted background diffusivity via the adjoint estimation process. In other words, it corresponds to the mixing that is not associated with the instabilities determined from Gaspar et al. (1990) or convective adjustment. We have edited this sentence to read: "Vertical mixing–diapycnal plus the vertical component of the along-isopycnal tensor–is determined according to the Gaspar et al. (1990) mixed layer turbulence closure, simple convective adjustment, and estimated background κρ for internal wave-induced mixing." We should clarify, however, that it isn't vital to determine the specific processes the estimated mixing is representing. We have also added the following sentences: "Here, κρ represents a combination of processes." To be clear, the central point of our manuscript is that oxygen observations can potentially help qualitatively infer where there is enhanced turbulence.

•It is good that the authors reminded the readers a couple of times through the text that the "observed" diapycnal diffusivities based on either microstructure measurements or fine-scale parameterization with CTD measurements include uncertainties as well.

•The uncertainties in these observational products are essential to consider to understand the purpose of our manuscript and we're glad our reiterations helped communicate this point.

Detailed comments:

•Lines: 3-5: The authors concluded that "the assimilation of existing in situ temperature,

salinity, and pressure observations is not sufficient to constrain κρ estimated with ECCO". However, since the number of iterations or the adjoint runs in ECCO is limited, is it possible that by running more iterations, the ECCO diffusivities will be better adjusted to the truth, even without including other new datasets? From what has been presented in this paper, I don't believe that possibility has been ruled out, and therefore the statement above is not solid, IMO.

•It's true that running more iterations of ECCO could adjust the diapycnal diffusivity field such that it's closer to that of the real ocean, but we now show in a new figure that this wasn't always true between the first and fifty-ninth iterations for each microstructure campaign. It's possible that the other estimated fields absorb some of the errors in the diapycnal diffusivity field due to the under-determined nature of the ECCO estimation problem. The purpose of our manuscript is not to sort out which would happen with additional iterations of ECCO, but we do show some evidence that your hypothesis may be closer to the truth, in which case it's entirely possible that infinite computational resources are required to converge to the truth. Given that future ECCO optimizations will likely be insufficient in number for the diapycnal diffusivity (or other) field(s) to converge 100% (especially considering their errors), we seek new data sets that could help guide the diapycnal diffusivity field within the number of iterations performed.

• Line 14: I did not find how this conclusion was reached.

•The evidence we have for this is in the correlations between the adjoint sensitivities for the experiments with oxygen in the misfit and those for the experiments with diapycnal diffusivities in the misfit. The adjoint sensitivities tend to agree in sign but are very well correlated, suggesting that the values of the diapycnal diffusivities would be different if one data set were assimilated instead of the other. This part of our abstract now reads: "Information provided by more accurately measured dissolved oxygen concentrations is not equivalent to that from less accurately mea- sured κρ. However, we show that adjoint sensitivities of dissolved oxygen concentration misfits to the state estimate's control space typically direct κρ to improve relative to the Argo-derived and microstructure-inferred values."

•Line 141: The background κρ is time-invariant. But if the parameterized part and/or convective adjustment were included, it should be time-variant.

•That is true, which is why ECCO only estimates the background diapycnal diffusivity field. We tried to make this more clear with the edits to a sentence we quoted above.

•Line 152: What do “14-day adjustments” mean? Could the authors explain that?

•The adjoint averages adjustments to the atmospheric forcing fields, which are re-estimated and then applied over 14-day periods. We have rephrased how we state this sentence: “Average adjustments to the wind stress, wind speed, specific humidity, shortwave downwelling radiation, and surface air temperature are re-estimated and then applied over 14-day periods.”

•Lines 156-158: This information is important but confusing. Could the authors elaborate on how the ECCO κρ used in this study was obtained? Is it based on parameterization (Gaspar et al., 1990)? or is it adjusted through the adjoint? or other ways?

•It is the background vertical diffusivity field that ECCO estimates through its adjoint minus the (3,3) component of the along-isopycnal diffusivity tensor to get the time-invariant background diapycnal diffusivity field. We tried to clarify this in the text.

•Lines 174-176: It seems that the authors were likely asked by other reviewers to add the appendix. I personally think it is not necessary. But I am OK if the authors choose to keep it.

•Because there are two different types of data assimilation systems that have different problems with their diapycnal diffusivity field, we had two manuscripts initially: one that’s the main text and one that was an elaboration of the Appendix. We were unable to perform additional experiments using the data assimilation system of interest (NASA GMAO S2S) in the Appendix because we were not given access. We were able to perform additional simulations using another modeling system (without data assimilation), but exclude these simulations from this manuscript. We simply suggest that it’s worth considering the equivalent problem we’re pointing out in ECCO but with sequential data assimilation systems because there are likely consequences for forecasting systems.

•Line 200: It would be helpful if the authors could explicitly state the differences between the two runs Eκ and Eε.

•The difference is that Eκ includes the diapycnal diffusivities in the misfit function and Eε include the dissipation rates in the misfit function. The values of ε and κ are related to each other through the Osborn (1980) relation so any qualitative differences between comparisons with Eκ and comparisons with Eε are related to the stratification. We edited the bullet points describing exactly how the Eκ and Eε were performed now.

•Lines 214-223: I am confused here. Why does using a previously derived product as an initial condition minimize model drifts? Also, how did the authors conclude that using other products as initial conditions would be worse than using that one?

•The way we described where the initial conditions for oxygen concentrations come from should be improved. The initial conditions come from MITgcm/verification/global_oce_biogeo_bling/input in the model's package, which was derived from World Ocean Atlas (2013). Because the initial conditions are observationally-derived, we concluded that using other products as initial conditions would be worse choices. A figure in our manuscript shows the differences between the model and World Ocean Atlas (2013) product by taking the nearest neighbors to the model grid instead of introducing interpolation errors. Still, this figure shows there are larger errors in some regions than others, which suggests that model is drifting from the initial conditions in many regions.

•Lines 234-237: Since uncertainties associated with undersampling (spatially and temporally) were not considered, the prescribed uncertainties are likely lower bounds. If so, the readers should be reminded about this.

•This was discussed in response to a comment the other reviewer made. Regions where we do not have observations can disagree in the signs of their adjoint sensitivities (or lower the correlation with another experiment's adjoint sensitivities) because we don't have observations there. The sampling issue

when calculating climatological fields could also be an issue because our weights for computing the misfits could be inappropriate if there are, for example, seasonally aliased values. This could also explain some of the disagreements in signs of the adjoint sensitivities (or lower the correlation with another experiment's adjoint sensitivities).

•Lines 296-299: As commented above, is it possible that with more iterations κρ, ECCO will be closer to the observationally-derived κρ?

•Yes, but a new data set could conceivably accelerate the convergence to a more realistic diapycnal diffusivity in ECCO, which could be important because ECCO optimizations will not likely be iterated many more times than the version we're analyzing. However, the new figure we now include showing comparisons with individual microstructure campaigns suggests that more iterations can actually stray further the microstructure-inferred values than the first iteration.

•Lines 356-363: Now I understand the purpose of the run Eε. It would help the readers if some of the information here were added around line 200.

•Okay, we moved our explanation to earlier in the text. Thanks.

•Line 389: What is the amount of the available DO measurements? Is it comparable to or much larger than T/S data? If not, not sure it will be that useful.

•In the World Ocean Database and Argo, there are less dissolved oxygen data relative to T/S, but the reason why oxygen can serve as a useful constraint on the diapycnal diffusivities is that oxygen provides unique information about gradients, as oxygen has a different source function and history compared to T/S. Further, oxygen concentrations can be weighted more than other physical variables (e.g., temperature, salinity, and pressure) to compensate for the relative dearth of oxygen concentration observations. We haven't performed this type of sensitivity analysis because we didn't perform optimization runs, but this has been considered in SOSE.

•Line 420: “microstructure CTD-derived κρ”? Do you mean CTD-derived κρ?

•These were CTD data taken (approximately) concomitant with the microstructure, which is all we were trying to say. We now say: “A preliminary analysis suggests that the percent difference between the full depth-averaged CTD-derived κρ from the finescale parameterization and the microstructure-inferred κρ at the same locations is indistinguishable from zero (1.68%), but the quality of the the CTD data taken concomitantly with microstructure has not been fully assessed.”

•Lines 345-437: The last sentence is disconnected from the sentences before it.

•We have edited the text so that this text reads more smoothly: “A more complete representation and understanding of κρ is possible through these analyses and methods.”

•Figure 4: since there are many data gaps (white area) in the right panels, it might make sense to use a different colorbar.

•Because the purpose of this figure is to show that there are many regions where the difference between the ECCO-estimated diapycnal diffusivities are an order of magnitude different from observationally-derived ones, the only regions that need the reader’s focus are the blue and red ones in panels b,d,f. The difference between the regions with white because there are no data and the regions with white because the ECCO-estimated diapycnal diffusivities are close to the observationally-derived ones is not very important if there are large regions of the ocean with red or blue colors. However, we have replotted this figure using a different colorbar to make it more apparent where there are no data and where there are small differences.