

Authors' response to reviews of the manuscript "Tracking water masses using passive-tracer transport in NEMO v3.4 with NEMOTAM: application to North Atlantic Deep Water and North Atlantic Subtropical Mode Water" by D. Stephenson, S. Müller, and F. Sévellec

We are grateful to the editor and two referees for their positive and constructive feedback, which has been of great benefit to the work presented in this manuscript. We address their specific comments in detail below

Font legend:

1. Reviewers' comments are in black and italic.
2. Our responses are in black and normal font.
3. Text which has been removed from the manuscript is in red and normal font.
4. Text which has been added to the manuscript is in blue and normal font.
5. Unchanged text in the manuscript is in green and normal font.

Page references (page X, line Y) refer to the marked-up version appended to this document.

Anonymous Referee 1

General comments:

*"Stephenson et al. describe a promising and interesting addition to the toolbox for visualizing flow in ocean models. This tool is available only within the context of a tangential and adjoint model but it has an advantage over the more commonly used method of Lagrangian particle tracking: insensitivity to the number of particles simulated. The authors then illustrate the types of results possible with their tool with two well known water masses: North Atlantic Subtropical Mode Water and North Atlantic Deep Water. These results are broadly consistent with results from other models and observations. Finally, as an oceanographer whose expertise is in Lagrangian methods and the water masses discussed in this paper but ****not**** with tangent-linear and adjoint models (TAM), I found the casting of the TAM approach into a purely passive tracer framework to be exceptionally illuminating; Section 2.2 in particular helped me understand the physicality behind TAMs. In my opinion, this paper contributes both a new method for tracking water masses as well as a special, simplifying case of the TAM approach; both will be useful. I recommend this paper for publication in GMD with some minor clarifications listed below."*

We thank the referee for this positive and encouraging response. We are pleased to find that our treatment of the methodology was insightful to a modeller outside of the versatile but admittedly niche field of tangent-linear and adjoint modelling.

Specific comments

*1. "This paper does not explicitly compare this new method to an established baseline (Lagrangian particle tracking, standard tracer methods). In particular computational demands are cited as one reason for using this method over Lagrangian particles yet there are no run time statistics in the manuscript. In my opinion, explicit quantification of the additional computational resource overhead necessary for running the TAM and this method, and insight into how this overhead scales with model resolution, is necessary in order for other scientists to fully evaluate the merits of this method relative to other methods. In the conclusion, the authors suggest it could be used with higher resolution models (i.e. ORCA12) as an "off-the-shelf" addition. That's definitely intriguing and I think more basic information about resources is necessary to support interest along these lines. Note that here I am ***not*** suggesting that the authors do any particle tracking. Rather, I think it's important to document what kinds of resources are necessary to run this method. Describing these run times will also help illustrate the actual TAM tracer workflow"*

A variation on this comment regarding runtime statistics was also brought to our attention by the second reviewer, and is an important point which we had missed. Thank you to both reviewers for

the recommendation, which we have now addressed using a number of efficiency experiments at laminar and eddying resolutions. This has been added to the manuscript as an additional section:

2.4 Performance

In order to test performance, the model was run for one-hundred days at 2° resolution (the ORCA2 configuration used in our case studies) and five days at 0.25° resolution (the ORCA025 configuration, Madec et al. 2012). Each configuration was run in four different modes (nonlinear, nonlinear with TAM trajectory production, tangent-linear with passive tracer, adjoint with passive tracer). Trajectory files were written once per day, and the linear advection scheme was the weighted-mean scheme described in Section [development_advection_schemes](#). Each test was conducted with a range of parallelisation arrangements (16, 32, 64, 128, 256 and 512 CPU cores) and repeated 10 times to account for system variability (Table 1). The tests were conducted on a local HPC system, with 72 compute nodes, each offering two eight-core 2.6GHz processors (Intel Xeon E5-2650 v2) and 64GiB, connected by an InfiniBand QDR network.

The general order of time efficiency is consistent across all tests: the linear forward model is

Table 1. Results of performance tests for two model configurations (ORCA2, upper section and ORCA025, lower section). Top rows: trajectory storage requirements per output. Lower rows: runtime per model day for four model modes: Nonlinear with trajectory output (NLT), nonlinear without trajectory output (NL), tangent-linear (TL) and adjoint (AD). The standard deviation of ten runs is given as a \pm uncertainty. Dashes indicate tests which failed due to insufficient memory.

Cores (nodes)		16 (1)	32 (2)	64 (4)	128 (8)	256 (16)	512 (32)
ORCA2							
Trajectory size (MB/d)		43.909	44.418	43.978	44.412	44.848	46.351
Mode	NLT	2.09 \pm 0.20	1.54 \pm 0.17	1.23 \pm 0.17	1.39 \pm 0.26	2.53 \pm 0.27	4.08 \pm 0.35
	NL	1.98 \pm 0.17	1.37 \pm 0.10	1.10 \pm 0.09	0.99 \pm 0.08	1.85 \pm 0.36	2.18 \pm 0.70
	TL	0.78 \pm 0.01	0.50 \pm 0.01	0.36 \pm 0.09	0.36 \pm 0.13	0.45 \pm 0.18	0.88 \pm 0.10
	AD	0.89 \pm 0.01	0.61 \pm 0.02	0.53 \pm 0.04	0.50 \pm 11	0.59 \pm 0.05	1.99 \pm 0.61
ORCA025							
Trajectory size (MB/d)		-	5217.427	5218.495	5220.824	5217.658	5218.448
Mode	NLT	-	1091.55 \pm 69.38	692.98 \pm 21.38	428.36 \pm 13.76	302.18 \pm 30.55	143.60 \pm 12.73
	NL	-	1065.48 \pm 54.56	656.94 \pm 13.74	386.34 \pm 13.89	262.72 \pm 34.56	131.70 \pm 13.28
	TL	-	-	315.26 \pm 26.17	207.82 \pm 14.43	136.42 \pm 13.00	64.70 \pm 4.20
	AD	-	-	403.74 \pm 36.18	252.36 \pm 14.17	181.12 \pm 7.70	92.40 \pm 6.26

between 2.5 (16 cores) and 3.1 (64 cores) times as fast as the nonlinear model without trajectory output in ORCA2, and 1.9 (64 cores) to 2.0 (512 cores) times as fast in ORCA025. The adjoint is slightly less efficient. The nonlinear runs which produce trajectory output are the slowest across the board, due to the high level of output. Memory use appears higher in the linear model than the nonlinear, such that the linear model could not be run on two 64GiB compute nodes alone in ORCA025 (column two).

The model shows good scaling in the ORCA025 configuration at all CPU arrangements tested here, but begins to worsen in ORCA2 beyond 128 CPU cores. Further, the added time required for trajectory output in nonlinear ORCA2 runs on 128 CPU cores can be considerable. This is possibly due to the generation of a very large number of files during the model run, as the number of files generated is proportional to the number of CPU cores, the trajectory write frequency, and the run length. The required storage size remains relatively constant for runs with larger number of CPU cores, despite the greater number of files (Table 1, first row).

An additional limitation which we discovered during our longest experiments comes from system file-number limits, which can readily be reached for typical systems for very long runs using many cores.

2. Page 2: near line 25: “along with the ability to re-use a single “trajectory” run of a nonlinear model, offer an advantage over passive tracer tracking in a nonlinear OGCM.” How exactly is this re-use different from re-using existing output from an eddy-resolving OGCM for different particle tracking studies? (For example, already cited in the manuscript,

Gary et al. (2011, 2014) and Burkholder and Lozier (2011) used different particle tracks from the same model output in three separate papers.)

We agree with the reviewer that similar offline methods such as Lagrangian particle tracking can indeed reuse identical model outputs for distinct particle release experiments. Our comment is perhaps unclearly worded (having followed a comment regarding Lagrangian methods) but refers instead to online passive tracer methods. We have implemented the following change:

However, due to its probabilistic nature, it offers an advantage over the more conventional tracking technique of Lagrangian particle modelling. Within the Lagrangian framework, ocean sensitivity to initial conditions means that a very large ensemble of initially close particle deployments may be required to representatively sample the full space of particle trajectories. A TAM can be exploited to bypass this requirement, producing a continuous probability distribution of all possible particle trajectories. As such, the method does not describe particle locations and deterministic trajectories, as a Lagrangian approach would, but tracer concentrations and probabilistic pathways. The adjoint and tangent-linear of the model can respectively track the origins and fate of passive-tracer-tagged water in this manner. This native ability to track water both forward and backward in time, along with the ability to re-use a single "trajectory" run of a nonlinear model, offer an advantage over passive tracer tracking in a nonlinear OGCM.

>

However, they offer two distinct advantages over online passive-tracer tracking in a nonlinear OGCM. Firstly, as with many Lagrangian methods, the TAM propagates tracer fields offline, meaning only one (more computationally demanding) simulation in the nonlinear model is required to conduct many experiments under similar conditions. Secondly, the TAM propagates probabilities both backwards and forwards in time (detailed further in Section 2.1). The adjoint and tangent-linear of the model can respectively track the origins and fate of passive-tracer-tagged water in this manner. (page 2, line 28)

3. Page 4, near line 15: The Bower et al. 2009 and 2011 studies did not use profiling floats, only RAFOS floats, which did not profile

Thank you to the referee for bringing this to our attention. We have corrected as follows:

However, new data from profiling floats and modelled Lagrangian drifters

>

However, new data from floats and modelled Lagrangian drifters (page 3, line 54)

4. Page 6, near line 10: "Below the surface, it occupies a narrow latitudinal band at depths of up to 240 m (Fig. 2, red shading)." This maximum depth of 240 m is inconsistent with the greater maximum depths reported on page 9 near line 30. Perhaps the 240 m is the time mean maximum depth or a thickness? Please clarify.

Thanks to the reviewer for highlighting this inconsistency. The value quoted does indeed refer to thickness. The correct value is in fact 310m, and has been corrected below. This also revealed an error in the y-axis tick labels of Figure 2, which has now also been corrected.

Below the surface, it occupies a narrow latitudinal band at depths of up to 240m

>

Below the surface, it occupies a narrow latitudinal band at depths of up to 310m (page 4, line 73)

5. Page 6, near line 15, "Additionally, eddy-induced advective velocities, although present in the nonlinear trajectory, were not included for our passive tracers in NEMOTAM." Are these velocities the same as a Gent-McWilliams (GM90) eddy parameterization? If they were used in the model trajectory, why were they not used to move the tracers as well?

As the referee supposes, these are the GM90 bolus velocities. They are by default in NEMOTAM 3.4 not used for tracer advection. In particular, the nonlinear model calculates them online, and does not explicitly store them in the trajectory, so that they are not seen by NEMOTAM. Instead, the iso-neutral slopes used for their computation are stored, (so EIV can, in fact, be retrieved). We have since, for a different adjoint study, enabled online EIV calculation using these slopes in NEMOTAM by calling the corresponding routine from the nonlinear model. This has allowed us to test the impact on passive-tracer evolution. Preliminary tests reveal little difference for the case studies described in the manuscript (Figs. A,B,C,D,E), but we have added the option of EIV to the source code for further studies.

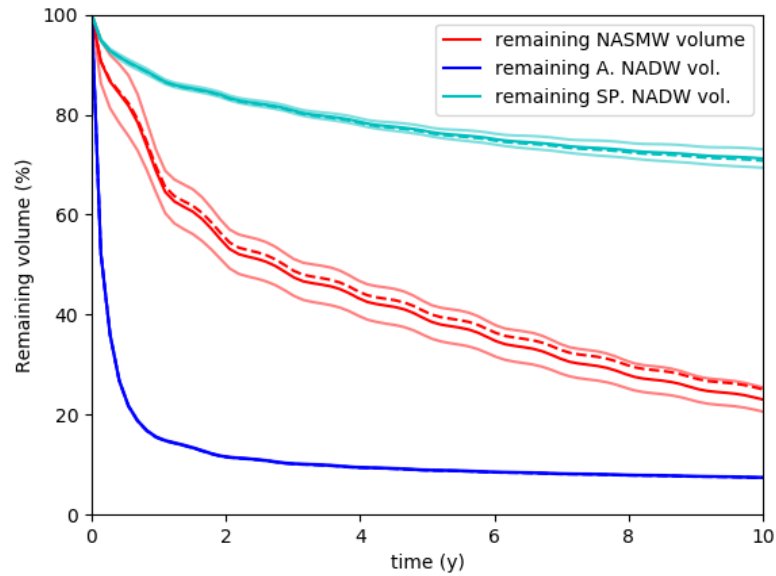


Fig A: Remaining volume of tangent-linear experiments additionally showing normalised absolute upper limit of error. Dark, dashed lines show the volume of tracer remaining without EIV. Dark solid lines are the equivalent with EIV. Faint lines show worst-case limits: EIV-on volume, +/- the globally integrated absolute volume error in every grid cell. All curves are normalised by the injected tracer volume. (Note that blue lines overlap, because of the small sensitivity of this experiment to EIV.)

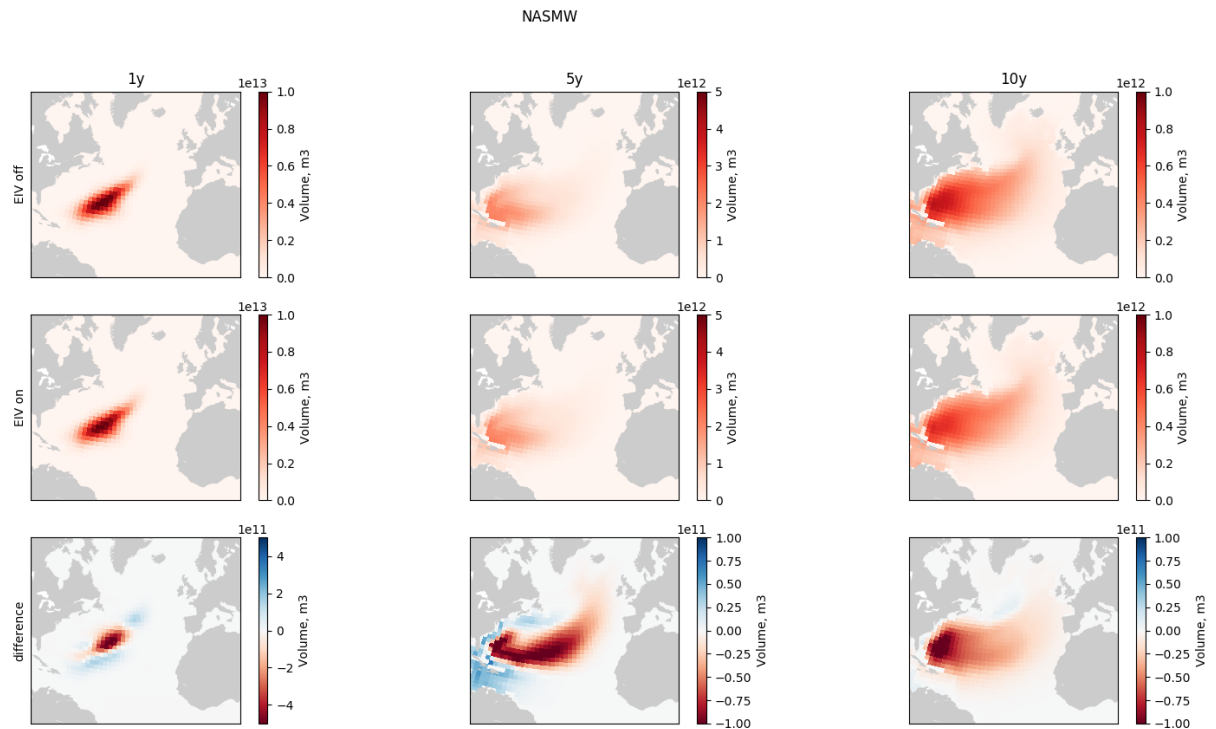


Fig B: Snapshots of depth-integrated volume for tangent-linear NASMW propagation without (top row) and with (middle row) eddy-induced velocities at 1y (left column) 5y (centre column) and 10y (right column). The bottom row shows the difference between the two. (Note the smaller colorbar extent of the last column, demonstrating that the difference is roughly one order of magnitude smaller than the results.)

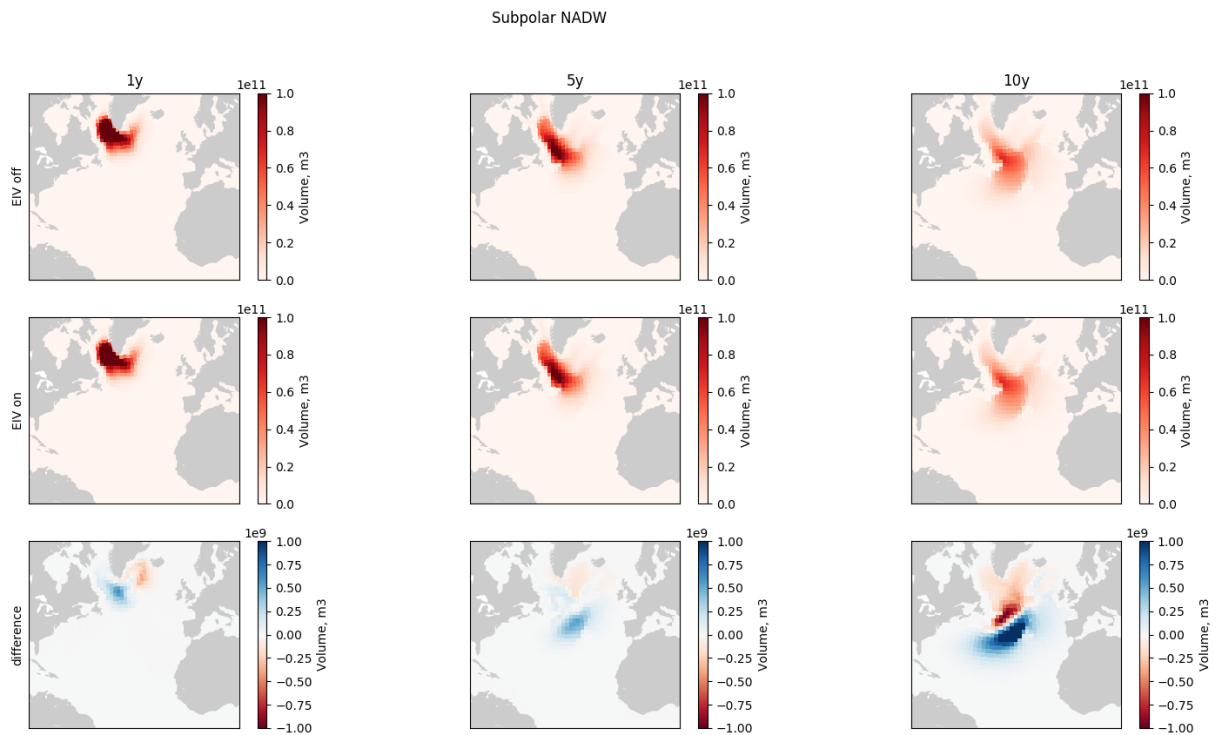


Fig C: As in Figure B, but for subpolar-outcropping NADW. (Note smaller colorbar extent of the last column, demonstrating that the difference is two orders of magnitude as small as the results.)

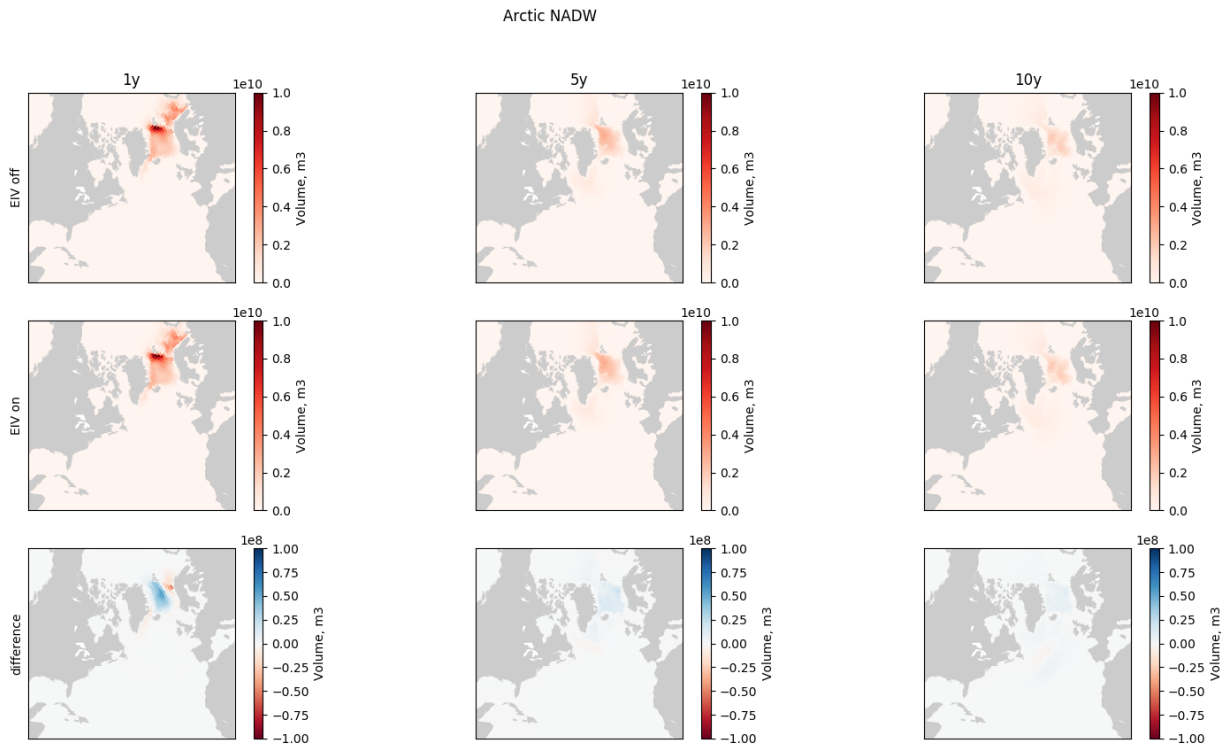


Fig. D: As in Fig. C, but for Arctic-outcropping NADW. (Note smaller colorbar extent of the last column, demonstrating that the difference is two orders of magnitude as small as the results.)

Formation volume, 0-10 years

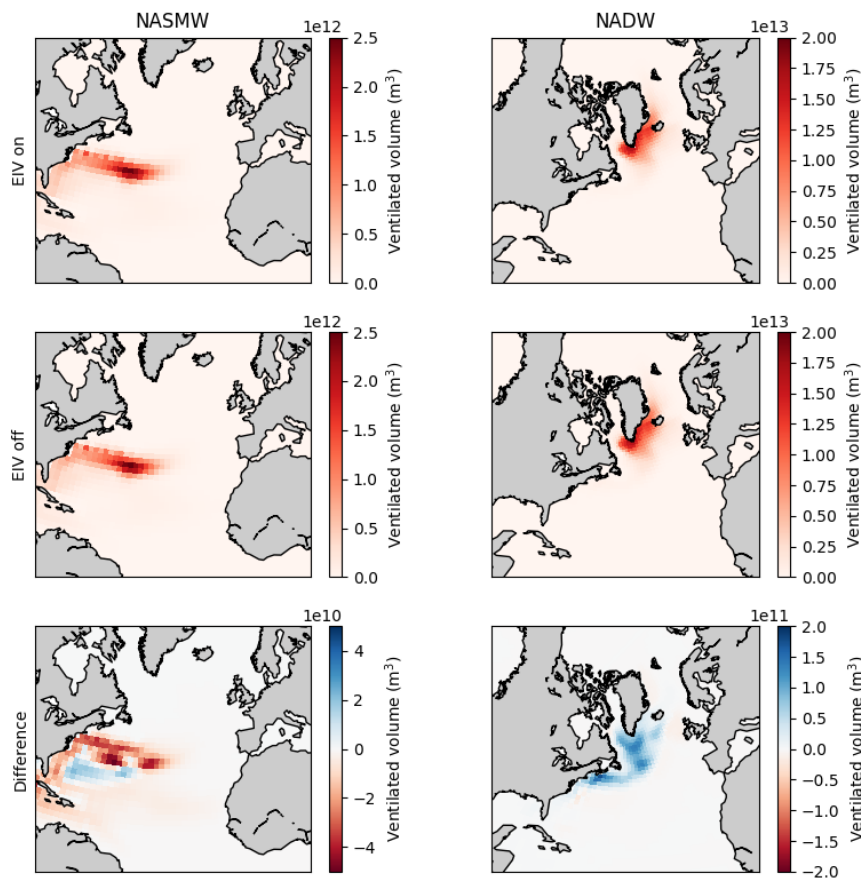


Fig E: Ventilated tracer volume of NASMW (left) and NADW (right) over the first ten years of case study experiments, with (top row) and without (middle row) eddy-induced velocities. The bottom row shows the difference. (Note smaller colorbar extent of the last column, demonstrating that the difference is two orders of magnitude as small as the results.)

We have also implemented the following change to the manuscript:

Additionally, eddy-induced advective velocities, although present in the nonlinear trajectory, were not included for our passive tracers in NEMOTAM.

>

Additionally, eddy-induced advective velocities are computed online by the nonlinear model and not stored explicitly in the trajectory. They may be recomputed from other trajectory variables, but this is not the default behaviour of NEMOTAM. As such, they were not included for the passive-tracer experiments in this study. They have been since been enabled in the passive-tracer source code as an optional online calculation. A preliminary comparative test suggests that errors are minor for our case studies (positive and negative concentration differences tend to cancel out, but even summing absolute differences, we do not exceed ~6% relative to the injection volume). We expect the difference to be greater in steep isopycnal regions such as the Southern Ocean. (page 5, line 8)

6. Page 12, near line 5, "This suggests that the NASMW surface outcrop is not, in fact, the dominant origin of NASMW...". I am confused by the use of "dominant" in this statement because it seems to me that it is not consistent with Fig. 8a which shows that a significant fraction of the NASMW (up to 70%) comes from the outcrop region. Granted, Fig. 8a is on a 0-5 year time scale and other panels in Fig. 8 clearly show sources from outside the

outcrop region on longer time scales, but the contributions of those other sources outside the outcrop region is still less than half of all the NASMW. Please clarify

We thank the referee for highlighting this. We agree that Fig. 8a gives the impression that the majority of NASMW comes from the outcrop region, and have now clarified this inside the text. Unlike in the forward model, where spatial maps represent snapshots, the adjoint requires time-integrated fields be shown. As such we did not include a single outcrop contour for this figure as we thought that this would be misleading given the strong seasonal behaviour of the outcrop (it is non-existent for half of the year). It is thus difficult to convey whether or not tracer ventilated within or outside of the outcrop with this figure, for which a TS histogram is more appropriate (Fig. 9). In our TS budget, it is more clear that the dominant origin indeed has a warmer signature than surface water within the outcrop.

Our modifications are as follows:

Surface origins of NASMW as determined by the backtracking budget analysis (adjoint model simulation). Shading indicates probability density. This corresponds to the likelihood that NASMW is formed in a given region during the time periods [0,5]~yr, (5,10]~yr, (10,30]~yr and (30,50]~yr. Inset percentages show the global integral (that is, the total proportion of the budget which is formed during each time period).

Inset percentages show the global integral (that is, the total proportion of the budget which is formed during each time period).

>

Inset percentages show the global integral (that is, the total proportion of the budget which is formed during each time period). Note that, contrary to Fig. 6, which displays instantaneous fields, here time-integrated fields are displayed. Due to the large variability of its extent over the integration window, the outcrop region is not shown. (Fig. 8, caption)

–

This suggests that the NASMW surface outcrop is not, in fact, the dominant origin of NASMW, contrary to the intuition provided by laminar models of the ventilated thermocline (Luyten et al., 1983), and that ocean dynamics play a fundamental role in mode water formation.

>

While the surface distribution suggests that a relatively small neighbourhood dominates the formation of model NASMW (Fig. 8), we remind that the seasonal variability of surface properties (particularly the outcrop area) reflects strongly on the thermohaline properties of NASMW at formation (Fig. 9). This suggests that the NASMW surface outcrop is not, in fact, the dominant origin of NASMW in the model, contrary to the intuition provided by laminar models of the ventilated thermocline (Luyten et al., 1983). The implication is that mode water formation is not a passive process, and that ocean dynamics must play a fundamental role. (page 9, line 62)

Anonymous Referee 2

General comments:

"In this manuscript, the authors present an analysis of the pathways of NADW and NASMW in the NEMOTAM model, focussing on the long time scales and basin-scale pathways in their 2degree resolution simulation. While the manuscript is well-written and easy to follow, I am not entirely sure it fits the scope of GMD in its present form. Most of the manuscript is about the physical interpretation of the water mass pathways, and there is relatively little information about the technical details of the implementation. On the other hand, the low resolution (even for climate models) means that it is unclear how representative/relevant the results are to the physical oceanography community. Since there is hardly any comparison to observational evidence of the pathways (CFCs? C14 dating?), I suspect the

manuscript in its present form would raise questions in a physical oceanography journal too. I think, however that it should be feasible to rewrite the manuscript to a more traditional GMD manuscript for the audience here to find it interesting. In particular, I think the authors would then need to:

We thank the reviewer for their suggestions, which we address individually below. We remark that our decision to submit this manuscript to GMD was based on perceived benefit to the community of the tool (especially given the existence of the NEMO special issue), rather than to establish new ground in the analysis of either of the considered water masses. These case studies were chosen to showcase the versatility of the technique, and highlight the possibilities available to a scientist undertaking a dedicated study at higher resolution.

1. Add much more information about the technical details of the model. What is its memory/cpu usage? How does it scale?

We are glad that this has been brought to our attention by both reviewers and express our gratitude again for highlighting this oversight. We have prepared an additional section to be added to the manuscript:

2.4 Performance

In order to test performance, the model was run for one-hundred days at 2° resolution (the ORCA2 configuration used in our case studies) and five days at 0.25° resolution (the ORCA025 configuration, Madec et al. 2012). Each configuration was run in four different modes (nonlinear, nonlinear with TAM trajectory production, tangent-linear with passive tracer, adjoint with passive tracer). Trajectory files were written once per day, and the linear advection scheme was the weighted-mean scheme described in Section 2.3. Each test was conducted with a range of parallelisation arrangements (16, 32, 64, 128, 256 and 512 CPU cores) and repeated 10 times to account for system variability (Table 1). The tests were conducted on a local HPC system, with 72 compute nodes, each offering two eight-core 2.6GHz processors (Intel Xeon E5-2650 v2) and 64GiB, connected by an InfiniBand QDR network.

The general order of time efficiency is consistent across all tests: the linear forward model is

Table 1. Results of performance tests for two model configurations (ORCA2, upper section and ORCA025, lower section). Top rows: trajectory storage requirements per output. Lower rows: runtime per model day for four model modes: Nonlinear with trajectory output (NLT), nonlinear without trajectory output (NL), tangent-linear (TL) and adjoint (AD). The standard deviation of ten runs is given as a \pm uncertainty. Dashes indicate tests which failed due to insufficient memory.

Cores (nodes)			16 (1)	32 (2)	64 (4)	128 (8)	256 (16)	512 (32)
ORCA2								
Trajectory size (MB/d)			43.909	44.418	43.978	44.412	44.848	46.351
Mode	NLT	Run time (s/d)	2.09±0.20	1.54±0.17	1.23±0.17	1.39±0.26	2.53±0.27	4.08±0.35
	NL		1.98±0.17	1.37±0.10	1.10±0.09	0.99±0.08	1.85±0.36	2.18±0.70
	TL		0.78±0.01	0.50±0.01	0.36±0.09	0.36±0.13	0.45±0.18	0.88±0.10
	AD		0.89±0.01	0.61±0.02	0.53±0.04	0.50±0.11	0.59±0.05	1.99±0.61
ORCA025								
Trajectory size (MB/d)			-	5217.427	5218.495	5220.824	5217.658	5218.448
Mode	NLT	Run time (s/d)	-	1091.55±69.38	692.98±21.38	428.36±13.76	302.18±30.55	143.60±12.73
	NL		-	1065.48±54.56	656.94±13.74	386.34±13.89	262.72±34.56	131.70±13.28
	TL		-	-	315.26±26.17	207.82±14.43	136.42±13.00	64.70±4.20
	AD		-	-	403.74±36.18	252.36±14.17	181.12±7.70	92.40±6.26

between 2.5 (16 cores) and 3.1 (64 cores) times as fast as the nonlinear model without trajectory output in ORCA2, and 1.9 (64 cores) to 2.0 (512 cores) times as fast in ORCA025. The adjoint is slightly less efficient. The nonlinear runs which produce trajectory output are the slowest across the board, due to the high level of output. Memory use appears higher in the linear model than the nonlinear, such that the linear model could not be run on two 64GiB compute nodes alone in ORCA025 (column two).

The model shows good scaling in the ORCA025 configuration at all CPU arrangements tested here, but begins to worsen in ORCA2 beyond 128 CPU cores. Further, the added time required for trajectory output in nonlinear ORCA2 runs on 128 CPU cores can be considerable. This is possibly due to the generation of a very large number of files during the model run, as the number of files generated is proportional to the number of CPU cores, the trajectory write frequency, and the run length. The required storage size remains relatively constant for runs with larger number of CPU cores, despite the greater number of files (Table 1, first row).

An additional limitation which we discovered during our longest experiments comes from system file-number limits, which can readily be reached for typical systems for very long runs using many cores.

2. Add more information about the implementation of the model. What exactly is meant with a 'perturbation' (page 4, line 28)? A perturbation to what? How does the result depend on the choice of non-perturbed state? Is that irrelevant because of the assumption of linearity? How good is this assumption of linearity anyways? When and where does it break down? How big are the resulting errors?

It was not our intention to provide an in-depth description of the TAM framework in a general context outside of our application. The passive tracer approach is a simplification of an otherwise quite technical method (TAM) for which there is an extensive literature. Our hope was to focus on this simplification motivated by its potential for accessibility and interdisciplinary use. We have now clarified this from the outset, and thank the reviewer for bringing this to our attention. Our modification is as follows:

These relationships are derived in full by Errico (1997). Furthermore, we have(...)

>

These relationships are derived in full by Errico (1997), who provides a thorough treatment of the linear method and its limitations, beyond the simplified, inherently linear use-case of the present study. Finally, we have(...) (page 4, line 4)

Regarding the dependency on the non-perturbed state (the trajectory), in the passive tracer context the perturbation (or as we prefer, "injection") does not actually change anything, and so can be seen as an analysis of the trajectory. In this sense, it is entirely dependent on it. The assumption of linearity and window of validity for a TAM is an example of a typical complication (c.f. Errico et al., 1993, Tellus), but one which is not an issue in the passive tracer case, as there are no feedbacks. For instance, static instability induced by a decrease in temperature, due to the high degree of nonlinearity, would not be triggered by an active linear perturbation, which is an error of the active linear model. However, in a passive linear model, no response is ever triggered by the tracer, and so these errors cannot arise. As these problems are not present in the simplified case of passive tracer tracking, we choose not to describe them, as they lie outside of the realm of our study.

3. Add much more validation of the model implementation. Do the TAM results indeed agree qualitatively with a full (nonlinear) tracer experiment in the 'normal' model? And how would this change when changing resolution?

This is an interesting question in that, despite there being no effective difference between a linear and a nonlinear passive tracer, there remain some important simplifications in the TAM. In particular, the advection scheme is linearised (such that the adjoint is consistent). The nonlinear advection scheme (TVD) is considered in Section 2.3. As described (page 5, lines 4-23) some additional parameterisations are also absent in the linear advection-diffusion routines, particularly EIV. This is the native behaviour of NEMOTAM, but we have since changed it for another adjoint study, allowing us to conduct some tests on passive-tracer propagation. Preliminary tests reveal little difference for the case studies described in the manuscript, but we have added EIV as a switch in the source code for subsequent studies. (Note that we anticipate that in specific regions such as the Antarctic Circumpolar Current, which is much more intense in terms of mesoscale

eddy turbulence, this will make a larger difference.) Regarding this limitation, we have modified the manuscript to discuss it:

Additionally, eddy-induced advective velocities, although present in the nonlinear trajectory, were not included for our passive tracers in NEMOTAM.

>

Additionally, eddy-induced advective velocities are computed online by the nonlinear model and not stored explicitly in the trajectory. They may be recomputed from other trajectory variables, but this is not the default behaviour of NEMOTAM. As such, they were not included for the passive-tracer experiments in this study. They have been since been enabled in the passive-tracer source code as an optional online calculation. A preliminary comparative test suggests that errors are minor for our case studies (positive and negative concentration differences tend to cancel out, but even summing absolute differences, we do not exceed ~6% relative to the injection volume). We expect the difference to be greater in steep isopycnal regions such as the Southern Ocean. (page 5, line 8)

Specific comments

1. page 1, line 2: Be more specific what is 'probabilistic' about the tool. Is it the diffusive component?

This question stimulated an interesting discussion for which we thank the reviewer. In the broadest sense, the Eulerian frame of reference is what determines the probabilistic nature of the approach. Parameterised diffusion, offering sub-grid-scale closure, is an important component of this, but does not determine the probabilistic nature of the approach alone. Rather, in terms of the equations themselves, the advection and diffusion of a concentration in the Eulerian framework rather than of a particle in the Lagrangian perspective, is what allows us to produce a continuous probability distribution natively. This is in opposition to reconstructing one from a series of discrete trajectories, as in the Lagrangian framework. While we choose to maintain the current wording in the abstract for brevity, we have elaborated further in the Introduction:

Lagrangian particles, conversely, are indivisible nodules which may only be advected by the immediate flow. As such, an infinite number of them is required to fully represent an equivalent dye concentration.

>

Lagrangian particles, conversely, are indivisible nodules which may only be advected by the immediate flow. The Eulerian perspective of the passive tracer method makes it inherently probabilistic - the continuous field (tracer concentration) propagated by the model corresponds to a continuous probability distribution. Lagrangian particle trajectories can be seen as discrete samples from this distribution, and so a very large (theoretically infinite) number of them is required to adequately reconstruct it. (page 2, line 14)

2. page 1, line 5: Be clear that the fact that tracer is removed upon contact with the surface is a choice?

Thanks to the referee. This was previously unclear and has been rewritten.

Upon contact with the surface, the tracer is removed from the system, and a record of ventilation is produced.

>

To represent surface (re-)ventilation, we optionally decrease the tracer concentration in the surface layer, and track this concentration removal to produce a ventilation record. (Abstract)

3. page 2, line 1: 'bijective' is a not very common word. Explain?

While we agree that the word is perhaps uncommon in oceanography, it has a strict mathematical definition (a one-to-one relationship) which is an important aspect of the early mathematical diagnoses of subduction described in this line. As such, we choose to maintain its use.

4. page 2, line 8: *How many of these floats (order of magnitude) have been deployed?*

This sentence has been modified for clarity:

Despite the number of these floats, these trajectories remain under-sampled.

>

Despite their number (typically several tens of floats for dedicated pathway tracking studies), these trajectories remain under-sampled. (page 1, line 27)

5. page 2, line 15 (and later): *it is not true that an 'infinite' number of virtual particles are needed. Depending on machine precision, the tracer concentration is also not simulated to infinite accuracy.*

We thank the referee for this important technical point, which we had overlooked. Our intention was to highlight the continuous vs. discrete nature of tracer vs. particle tracking. The manuscript has been revised accordingly:

Lagrangian particles, conversely, are indivisible nodules which may only be advected by the immediate flow. As such, an infinite number of them is required to fully represent an equivalent dye concentration.

>

Lagrangian particles, conversely, are indivisible nodules which may only be advected by the immediate flow. The Eulerian perspective of the passive tracer method makes it inherently probabilistic - the continuous field (tracer concentration) propagated by the model corresponds to a continuous probability distribution. Lagrangian particle trajectories can be seen as discrete samples from this distribution, and so a very large (theoretically infinite) number of them is required to adequately reconstruct it. (page 2, line 14)

–

Our framework is rooted in concentrations and probabilities, comparable to the statistical properties of a high-resolution Lagrangian approach with infinitely many particles.

>

Our framework is rooted in concentrations and probabilities, comparable to the statistical properties of a high-resolution Lagrangian approach with a large (theoretically infinite) number of particles. (page 16, line 29)

6. page 2, line 24: *what exactly is meant with a 'probabilistic pathway'?*

This line was modified in response to the first reviewer's second comment:

However, due to its probabilistic nature, it offers an advantage over the more conventional tracking technique of Lagrangian particle modelling. Within the Lagrangian framework, ocean sensitivity to initial conditions means that a very large ensemble of initially close particle deployments may be required to representatively sample the full space of particle trajectories. A TAM can be exploited to bypass this requirement, producing a continuous probability distribution of all possible particle trajectories. As such, the method does not describe particle locations and deterministic trajectories, as a Lagrangian approach would, but tracer concentrations and probabilistic pathways. The adjoint and tangent-linear of the model can respectively track the origins and fate of passive-tracer-tagged water in this manner. This native ability to track water both forward and backward in time, along with the ability to re-use a single "trajectory" run of a nonlinear model, offer an advantage over passive tracer tracking in a nonlinear OGCM.

>

However, they offer two distinct advantages over online passive-tracer tracking in a nonlinear OGCM. Firstly, as with many Lagrangian methods, the TAM propagates tracer fields offline, meaning only one (more computationally demanding) simulation in the nonlinear model is required to conduct many experiments under similar conditions. Secondly, the TAM propagates probabilities both backwards and forwards in time (detailed further in Section 2.1). The adjoint and tangent-linear of the model can respectively track the origins and fate of passive-tracer-tagged water in this manner. (page 2, line 28)

7. page 3, line 9: The recent Bower et al review paper on Atlantic pathways (<https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2019JC015014>) is not mentioned here?

While this paper is referenced elsewhere (Page 2, line 8), it describes Lagrangian pathways rather than attempts to linearly model passive tracers, which is the focus of this line.

8. page 5, line 2: perhaps very briefly explain bra-ket notation to oceanographers?

We agree that the notation is uncommon, and have added the following short explanation:

(...)we have used the "bra-ket" notation of Dirac (1939).

(...)we have used the "bra-ket" notation of Dirac (1939) in which row and column vectors are respectively written as bras ($\langle \mathbf{a} |$) and kets ($| \mathbf{b} \rangle$) such that closed bra-ket terms become scalar (through the scalar product $\langle \mathbf{a} | \mathbf{b} \rangle = c$). (page 3, line 96)

9. page 6, line 4: what is the effect of the non-resolvment of the ice layers in NEMOTAM? Why are these not implemented? Is it technically impossible?

This is a well-raised point which requires further elaboration. In the context of an active tracer, ice formation and melt is highly nonlinear, and represents an error in NEMOTAM's linearization approach. However, a passive tracer would not interact with these processes. Furthermore, any change in volume is handled in the trajectory, so concentration changes due to, for example, dilution, are pre-emptively addressed by the trajectory.

We have modified the manuscript as follows:

There are two ice layers Fichet (1997) in the background state, although these are not incorporated into NEMOTAM.

>

There are two ice layers Fichet (1997) in the background state. The ice model is not linearised or coupled to NEMOTAM, although in the context of this study, ice dynamics are inherently unaffected (as a tracer which is passive cannot affect ice) (page 4, line 59)

10. page 9, line 23: why not impose an Ertel PV criterion, as is very often done? Is that possible within the TAM framework?

We had considered several criteria when preparing the experiments, and PV would have indeed been possible (any property of the background state can be used to determine where to inject the tracer). Given the lack of a universally accepted NASMW definition and in the interests of simplicity, we used a thickness criterion, as in Gary et al. (2014, J. Phys. Ocean.).

Here, we define NASMW(...)

>

While the method allows us to define water masses in terms of any model variable, we choose for simplicity to utilise the common approach of a temperature-salinity range. In particular, we define model NASMW(...) (page 7, line 23)

11. page 10, line 12: I don't understand the meaning of the word 'mechanically' here

This appears to be a typo and has been corrected. We thank the referee for bringing this to our attention.

(...)remaining tracer **mechanically** tends to reach deeper, colder, fresher waters.

>

(...)remaining tracer tends to reach deeper, colder, fresher waters. (page 8, line 39)

12. page 13, line 5: How sure are the authors that this indeed is a lognormal distribution? I would have liked to see a goodness-of-fit analysis. There are other distributions that produce similar-looking PDFs

We agree with the referee that this description is too specific and have removed it.

This distribution resembles a log-normal PDF with a lower skewness for NASMW lying below the MLD

>

There is a visibly lower skewness for NASMW lying below the MLD (page 9, line 71)

13. Figure 8: I am somewhat surprised that some of the NASMW originates near Greenland? Wasn't one of the conditions that the temperature was higher than 17C? Does that occur within the model in the region near Greenland??

After injection, the tracer is not tied to any particular temperature range, or indeed water mass. Under the assumption that the pathway of the tracer-tagged water is representative of the history of NASMW, it must be concluded that waters at the surface have subducted and warmed, eventually meeting the definition of NASMW. It should of course be borne in mind that subpolar-origin water represents a very small proportion of the NASMW make-up

We have clarified this in the manuscript:

Our findings suggest a high-latitude source makes up a large fraction of this water

>

Our findings suggest a high-latitude source makes up a large fraction of this water, **having followed a pathway from outside of our defined thermohaline range from the surface to eventually contribute to the make-up of NASMW** (page 9, line 88)

14. Figure 10: This bar chart is somewhat confusing because the blue bars raise higher than the striped ones. I assume they are individually normalised? Would it not make more sense to put them on the same y-scale?

We agree that the bar chart seems unintuitive as it is currently captioned, and so have modified the caption to explain more clearly. Normalising by the total does not show the distinct behaviour of the two reservoirs as clearly, as the volume of NASMW below the mixed layer depth is comparably small.

Probability distribution of NASMW age (hatched bars) and NASMW age below the mixed layer depth (blue bars).

>

Probability distribution of NASMW age (hatched bars) and NASMW age of NASMW restricted to be found below the mixed layer depth (blue bars). (Fig. 10, caption)

15. page 16, line 3: *There is plenty of evidence that NADW is not formed annually, but only in some years. That the model does have annual formation clearly is a bias. This should be clearly stated*

These experiments do not display this behaviour due to the forcing applied, which is repeated, rather than historical, which allowed for general behaviours to be studied, and for long (multi-century) experiments to be conducted. We have made this more explicit in the text:

The outcrop oscillates between the two regions with the seasonal cycle.

>

The outcrop oscillates between the two regions with the seasonal cycle. While observed NADW is known to form only in extreme winters and has a strong interannual signature (Avsic et al., 2006), our use of repeated forcing implies that formation has little year-to-year variation. (page 10, line 9)

16. page 23, line 26: *So where is the rest of the NADW, if it doesn't account for 100%?*

The rest of NADW remains in the ocean. The overall budget is close, only the ventilation source are not entirely determine. For any water mass studied in this way, it is not possible to close the source budget entirely (although we are of course closer with NASMW). We have accounted for a vast majority of NADW source and chose to cap our run at 400 years due to technical constraints. It may take many more centuries to approach the proximity to closure exhibited by the source NASMW. We propose the addition of a comment in the manuscript to stress this:

It should be noted that the mean age of model NADW is likely slightly higher, as the results detailed here do not quite account for 100% of the total budget

>

It should be noted that the mean age of model NADW is likely slightly higher. As with any water mass studied in this manner, a proportion will always remain in the ocean (closing the overall budget). Due to this, ventilation does not quite account for 100% of the total budget during our 400-year run. We limit our runs to this length due to technical constraints. (page 16, line 8)

17. page 25, line 16 (and other places): *here, the authors present their model results as if they realistically describe the ocean physics. However, with such a low resolution and known biases, I would prefer to see many more placements of caveats like 'within this model's context' etc*

We agree with the referee's comment and seek to enforce again our position of demonstrating a use-case of the tool, which replicates well observed behaviour of two interesting water masses. We have taken care to refer to "model NASMW" throughout, for example, and will revisit parts of the manuscript where this is not clear. We have edited the manuscript as follows:

Here, we define NASMW as > We define model NASMW as (page 7, line 27)

The outcrop area of NASMW > The outcrop area of NASMW in the model (page 8, line 3)

As outlined in Sect. 2.2, we begin by identifying all NASMW > As outlined in Sect. 2.2, we begin by identifying all water matching our NASMW definition (page 8, line 18)

Newly ventilated NASMW > Newly ventilated model NASMW (page 8, line 23)

NASMW is short-lived. > Our NASMW is short-lived. (page 8, line 33)

Of that which remains in the ocean, 90% is transformed and no longer fits **the** definition of NASMW > Of that which remains in the ocean, 90% is transformed and no longer fits **our** definition of NASMW (page 8, line 36)

However, it can be observed that only around 5% of all NASMW re-ventilation occurs within its surface outcrop > However, it can be observed that only around 5% of all **model** NASMW re-ventilation occurs within its surface outcrop (page 8, line 42)

To track existing NASMW > To track existing **model** NASMW (page 9, line 13)

when NASMW volume is at its maximum > when **model** NASMW volume is at its maximum (page 9, line 19)

neighbourhood of the NASMW outcrop > neighbourhood of **the model's initial** NASMW outcrop (page 9, line 32)

The 60-year mean formation location **of NASMW** is also within this region > The 60-year mean formation location is also within this region (page 9, line 33)

corresponding to the likelihood that NASMW is found in that region > corresponding to the likelihood that **model** NASMW is found in that region (Figure 6, caption)

For **NASMW** over 5 years old, current patterns begin to have a distinct influence on formation. There is a clear signature of the Gulf Stream on the youngest NASMW > For **water** over 5 years old, current patterns begin to have a distinct influence on formation. There is a clear signature of the Gulf Stream on the youngest **model** NASMW (page 9, line 38)

The signature of the Mediterranean outflow is particularly strong on the very oldest NASMW > The signature of the Mediterranean outflow is particularly strong on the very oldest **model** NASMW (page 9, line 45)

when the sources of NASMW are viewed in TS space > when the **model's** sources of NASMW are viewed in TS space (page 9, line 49)

is not, in fact, the dominant origin of NASMW > is not, in fact, the dominant origin of NASMW **in the model** (page 9, line 58)

Evolution of NASMW in TS space > Evolution of **model** NASMW in TS space (Figure 7, caption)

We also consider the time scales involved with NASMW formation > We also consider the time scales involved with NASMW formation **in the model** (page 9, 66)

the expected age of the NASMW > the expected age of NASMW **in the model** (page 9, line 79)

NASMW as a whole > **model** NASMW as a whole (page 9, line 83)

Surface origins of NASMW > Surface origins of **model** NASMW (Figure 8, caption)

Surface origins of NASMW > Surface origins of **model** NASMW (Figure 9, caption)

of NASMW age > of **model** NASMW age (Figure 10, caption)

water of this TS signature persistently outcrops > water of this TS signature persistently outcrops **in the simulation** (page 10, line 4)

in the NADW TS class > **in our** NADW TS class (page 10, line 16)

The Greenland Sea NADW outcrop > The [model's](#) Greenland Sea NADW outcrop (page 10, 18)

subpolar-outcropping NADW > subpolar-outcropping [model](#) NADW (Figure 11,12,13,14 caption)

[It](#) is known to mix into Circumpolar Deep Water > [observed](#) NADW is known to mix into Circumpolar Deep Water (page 11, line 36)

surface-borne NADW > surface-borne [model](#) NADW (page 13, line 24)

with those of the NADW > with those of the [model](#) NADW (page 13, line 27)

Arctic-outcropping NADW > Arctic-outcropping [model](#) NADW (Figure 15,16, caption)

part of subsurface NADW. > part of subsurface NADW [in the simulation](#). (page 13, line 32)

The little tracer which stays in the NADW class > The little tracer which stays in [our defined](#) NADW class (page 13, line 35)

(following water from each of the two distinct NADW outcrop regions) > (following water from each of the two distinct NADW outcrop regions [in the simulation](#)) (page 13, line 39)

Hence, any surface origins of NADW in the Arctic are likely found outside of the NADW TS class > Hence, any surface origins of NADW in the [model's](#) Arctic are likely found outside of the NADW TS class (page 13, line 43)

Using this approach, 86% of the NADW budget > Using this approach, 86% of the [model's](#) NADW budget (page 13, line 47)

total volume of NADW > total volume of [model](#) NADW (page 13, line 52)

NADW spanning the entire 400-year run > [Modelled](#) NADW spanning the entire 400-year run (page 14, line 8)

NADW source waters > [model](#) NADW source waters (page 15, line 8)

The dominant TS class associated with NADW formation > The dominant TS class associated with NADW formation [in the model](#) (page 15, line 15)

of NADW formation in 400 years > of [modelled](#) NADW formation in 400 years (page 15, line 22)

Arctic NADW originates > [modelled](#) Arctic NADW originates (page 15, line 27)

and fresher than NADW > and fresher than [we define NADW to be](#). (page 15, line 30)

of NADW formation > of [simulated](#) NADW formation (page 16, line 4)

in the Labrador and Irminger seas > in the [model's](#) Labrador and Irminger seas (page 16, line 21)

We have shown, for example, that on average > We have shown, for example, that, [in our simulation](#), on average (page 16, line 59)

We have estimated the average age of NASMW > We have estimated the average age of [model](#) NASMW (page 16, line 65)

For NADW > For [simulated](#) NADW (page 16, line 70)

to NADW formation > to [simulated](#) NADW formation (page 16, line 75)

For NADW formed in the subpolar outcrop of the Labrador Sea > For NADW formed in the subpolar outcrop of the model's Labrador Sea (page 16, line 86)

the NADW annual outcrop > the simulated NADW annual outcrop (page 17, line 12)

18. page 26, line 15: Long runs are not necessary for NEMOTAM to work, but only for the applications chosen here right? If the authors would have focussed on other applications, they could have used shorter runs?

We agree that this is a poorly-written line, and have revised it for clarity:

The main barrier to higher resolution for users of our development is the necessity of long trajectory outputs, which are required for NEMOTAM operation (Vidard et al., 2015).

>

The main barrier to higher resolution for users of our development is the necessity of frequent output and storage of trajectory snapshots from the nonlinear model, which are required for NEMOTAM operation (Vidard et al., 2015). Future versions of our tool will allow regional trajectory storage to overcome this barrier. (page 17, line 26)

19. page 26, line 18: How about 'looping' the NEMO fields? Would that work?

This is a very interesting question and one we have considered several times. It could both work and be easily implemented (this technique has in fact been used in other studies in the NEMOTAM predecessor, c.f. Sévellec and Fedorov, 2016, J. Clim.). We chose at this resolution to run for 400 continuous years as it was cleaner and we had the necessary storage available. However, with looped fields, particularly at higher resolution, shocks (for instance from the spurious appearance and disappearance of eddies at the looping point) would introduce physical inconsistencies in the results.

Tracking water masses using passive-tracer transport in NEMO v3.4 with NEMOTAM: application to North Atlantic Deep Water and North Atlantic Subtropical Mode Water

Dafydd Stephenson¹, Simon A. Müller^{1,2}, and Florian Sévellec^{1,3}

¹Ocean and Earth Science, National Oceanography Centre Southampton, University of Southampton, UK

²National Oceanography Centre, ~~European Way~~, Southampton, UK

³Laboratoire d'Océanographie Physique et Spatiale, Univ.-Brest CNRS IRD Ifremer, Brest, France

Correspondence: Dafydd Stephenson (D.Stephenson@noc.soton.ac.uk)

Abstract. Water mass ventilation provides an important link between the atmosphere and the global ocean circulation. In this study, we present a newly developed, probabilistic tool for offline water mass tracking. In particular, NEMOTAM, the tangent-linear and adjoint counterpart to the NEMO ocean general circulation model, is modified to allow passive-tracer transport. By terminating dynamic feedbacks in NEMOTAM, tagged water can be tracked forward and backwards in time as a passive dye, producing a probability distribution of pathways and origins, respectively. ~~Upon contact with the surface, the tracer is removed from the system, and a record of ventilation is produced.~~ To represent surface (re-)ventilation, we optionally decrease the tracer concentration in the surface layer, and track this concentration removal to produce a ventilation record.

Two test cases are detailed, examining the creation and fate of North Atlantic Subtropical Mode Water (NASMW) and North Atlantic Deep Water (NADW) in a 2° configuration of NEMO run with repeated annual forcing for up to 400 years. Model NASMW is shown to have an expected age of 4.5 years, and is predominantly eradicated by internal processes. A bed of more persistent NASMW is detected below the mixed layer with an expected age of 8.7 years. It is shown that while model NADW has two distinct outcrops (in the Arctic and North Atlantic), its formation primarily takes place in the subpolar Labrador and Irminger seas. Its expected age is 112 years.

Copyright statement. *To be completed by publisher*

1 Introduction

The intricate process by which the atmosphere and ocean exchange properties has decisive effects on oceanic and atmospheric circulation, biochemistry, and climate. Tracing the pathways of newly ventilated ocean water masses is a multi-disciplinary pursuit, for which an entire toolbox of methods has been developed.

Thermocline ventilation and water mass formation is an area with a long history of inquiry. From charts of hydrographic sections, Iselin (1939) inferred communication between isopycnals at depth and their outcrops. This idea was developed theoretically over the following decades in analytical models (Welander, 1959, 1971; Luyten et al., 1983). In

this vertically-discretised, laminar framework, potential vorticity (PV) conservation facilitates a bijective relationship between the origin and destination of a water parcel. In reality, however, turbulent mixing means that surface-to-deep pathways have a pronounced probabilistic character.

Observational studies have been able to capture snapshots of this turbulent behaviour locally in passive-tracer dye-release experiments (e.g. Schuert, 1970). However, the effect of turbulence on water mass pathways has only been observed on a large scale since the widespread adoption of cost-effective Lagrangian profiling floats, which have been observed to follow chaotic trajectories (e.g. Fischer and Schott, 2002; Fratantoni et al., 2013; Bower et al., 2019). ~~Despite the number of these floats~~ Despite their number (typically several tens of floats for dedicated pathway tracking studies), these trajectories remain under-sampled.

15

20

25

30

Simulations of larger numbers of Lagrangian particles in sophisticated eddy-resolving ocean general circulation models (OGCMs, e.g. Blanke and Raynaud, 1997; Gary et al., 2014) are reworking long-held assumptions about the routes taken by newly formed water masses. Despite these new developments, such studies are limited by the large computational expense of eddy-resolving models, and the fact that sub-grid-scale dispersion of Lagrangian particles is typically not parameterised in laminar models. An alternative approach is to use a simulated passive tracer (c.f. England and Maier-Reimer, 2001). Such tracers are spatial distributions of dye concentrations which undergo diffusion. Lagrangian particles, conversely, are indivisible nodules which may only be advected by the immediate flow. ~~As such, an infinite number of them is required to fully represent an equivalent dye concentration.~~ The Eulerian perspective of the passive tracer method makes it inherently probabilistic - the continuous field (tracer concentration) propagated by the model corresponds to a continuous probability distribution. Lagrangian particle trajectories can be seen as discrete samples from this distribution, and so a very large (theoretically infinite) number of them is required to adequately reconstruct it. This study presents a method for tracking water masses by means of passive-tracer deployment in the tangent-linear and adjoint model (TAM) developed for the NEMO OGCM (Madec, 2012; Vidard et al., 2015). TAMs are typically used for sensitivity studies and data assimilation (Errico, 1997). However, ~~due to its probabilistic nature, it offers an advantage over the more conventional tracking technique of Lagrangian particle modelling. Within the Lagrangian framework, ocean sensitivity to initial conditions means that a very large ensemble of initially close particle deployments may be required to representatively sample the full space of particle trajectories. A TAM can be exploited to bypass this requirement, producing a continuous probability distribution of all possible particle trajectories. As such, the method does not describe particle locations and deterministic trajectories, as a Lagrangian approach would, but tracer concentrations and probabilistic pathways.~~ they offer two distinct advantages over online passive-tracer tracking in a nonlinear OGCM. Firstly, as with many Lagrangian methods, the TAM propagates tracer fields offline, meaning only one (more computationally demanding) simulation in the nonlinear model is required to conduct many experiments under similar conditions. Secondly, the TAM propagates probabilities both backwards and forwards in time (detailed further in Section 2.1). The adjoint and tangent-linear of the model can respectively track the origins and fate of passive-tracer-tagged water in this manner. ~~This native ability to track water both forward and backward in time, along with the ability to re-use a single “trajectory” run of a nonlinear model, offer an advantage over passive tracer tracking in a nonlinear OGCM.~~

TAM construction is typically a laborious process (e.g. Giering and Kaminski, 1998). The Jacobian of a highly involved nonlinear model (often with millions of degrees of

freedom) must be computed with respect to the ocean state. This provides a linear function mapping small perturbations to the ocean state to their future outcome (the tangent-linear). The adjoint of this linear operator provides the sensitivity of the ocean state to earlier perturbations.

Several studies have bypassed such complications by computing approximations to the true adjoint of tracer transport in an OGCM. An early application of an adjoint approximation to the tracking of water masses was developed by Fukumori et al. (2004), who tracked eastern equatorial Pacific water with the Massachusetts Institute of Technology GCM. Their approximation capitalises on the comparatively simple nature of passive-tracer transport in isolation. For passive tracers, advection and diffusion are respectively skew-symmetric and self-adjoint. They were thus able to map possible past locations of the water mass by reversing the velocity field without changing the diffusion tensor, similarly to a true passive tracer adjoint. The method is further used by Qu et al. (2009), Gao et al. (2011), and Gao et al. (2012) to track Pacific waters and Qu et al. (2013) to track the salinity maximum of the North Atlantic subtropical gyre.

Another pseudo-adjoint passive tracer approach is presented by Khatiwala et al. (2005). Rather than deriving the Jacobian of an OGCM, they derive a tangent-linear approximation empirically. The procedure constructs a “transport matrix” row-by-row by repeatedly perturbing the nonlinear model along basis vectors of the ocean state. The responses following a single time step can be synthesised to derive the elements of the matrix. While more approximate than the Fukumori et al. (2004) approach, it is also more computationally efficient by design. This allows the transport matrix method to be used for long-term passive tracer experiments, for example simulations of long half-life radioisotopes.

The work presented herein is of a similar nature to the above studies, but repurposes an existing TAM for passive tracer tracking. In this sense its adjoint is the true adjoint of the nonlinear model, rather than a bespoke approximation.

We demonstrate the efficacy of the development through case studies of two climatically important water masses of the North Atlantic, whose formation regions are closely aligned with major components of the Atlantic Meridional Overturning Circulation (AMOC). These are North Atlantic Subtropical Mode Water (NASMW) or Eighteen Degree Water, formed in the vicinity of the Gulf Stream, and North Atlantic Deep Water (NADW), formed in the subpolar North Atlantic.

NASMW is the name given to a homogeneous body of water found in the upper region of the western subtropical gyre (STG). It was first identified as early as the Challenger expedition, where repeat soundings revealed its unusual uniformity (Thomson, 1877, their plate XI). It was later named Eighteen Degree Water for its characteristic temperature by Worthington (1959), who provided a first formal definition of the water mass.

NASMW formation has historically been attributed to surface heat loss during winter (e.g. Worthington, 1972; Speer and Tziperman, 1992) although mechanical forcing has been proposed to have an effect on formation by destroying PV (e.g. Thomas, 2005). Formation rate estimates by varying methods have produced substantially different results (addressed by Marshall et al., 2009), and it has recently been suggested that the nature of formation is location-dependent (Joyce, 2011).

Newly ventilated NASMW follows the subtropical gyre circulation, travelling north with the Gulf Stream (Klein and Hogg, 1996). Data from profiling floats (Kwon and Riser, 2005; Fratantoni et al., 2013) and Lagrangian particle modelling (Gary et al., 2014) suggest that a minority of this water should then be exported to the subpolar gyre.

Multiple studies have sought to quantify or describe the fate of NASMW. Forget et al. (2011) estimate that atmospheric exchange is roughly twice as effective at destroying existing NASMW as ocean internal mixing. Davis et al. (2013) discern between NASMW removal by air–sea exchange and internal mixing as respectively fast and slow processes, which occur at different stages of its seasonal cycle. They argue for the existence of a persistent reservoir of NASMW below the mixed layer, which is shielded from destruction by a layer of high PV. Gary et al. (2014) use a Lagrangian modelling approach to analyse the final stages of the NASMW life cycle, and show that the nature of NASMW destruction depends on how the water mass is defined - the lack of a universally accepted definition of NASMW commonly leads to such conflicts between studies (Joyce, 2011).

A quite different water mass in nature, NADW is one of the two primary high-density water masses formed in polar waters which act as pre-conditioners of the thermohaline component of the AMOC (Sévellec and Fedorov, 2011). NADW formation involves a range of processes and source water types. The dominant contributors are Labrador Sea Water (LSW, e.g. Talley and McCartney, 1982) and denser Overflow Waters (OW, e.g. Swift et al., 1980; Hansen et al., 2001), which respectively contribute to lower and upper NADW. LSW is created by cooling and convection of mode waters (McCartney and Talley, 1982) in the Labrador and Irminger Seas (Pickart et al., 2003a, b; Jong and Steur, 2016) during sufficiently severe winters (Clarke and Gascard, 1983). Meanwhile, OW forms as warm North Atlantic water crosses the Greenland-Scotland ridge, cools and sinks (Quadfasel and Käse, 2007). Resulting pressure gradients drive OW back over the ridge into the Atlantic at depth (Hansen and Østerhus, 2000).

Both source waters are exported southward following formation. The classical view from in-situ current measurements (e.g. Leaman and Harris, 1990; Molinari et al., 1992) is that the main southward pathway is the deep western boundary current (DWBC). However, new data from profiling floats and modelled Lagrangian drifters (Bower et al., 2009, 2011) suggest that this is not the complete story.

Gary et al. (2011) investigate alternative export pathways in a hierarchy of model resolutions. They find a second, slower pathway driven by a series of deep, eddy-driven anticyclonic gyres. These facilitate the southward transport of water parcels which have become detached from the DWBC, but are only found in models which are of eddy-resolving resolution. At lower-resolution, Lagrangian particles tend to follow the DWBC (Straneo et al., 2003).

Beyond the Equator, NADW mixes laterally and vertically with ambient water masses, eventually entering the Indian and Pacific basins via the Antarctic Circumpolar Current (Reid and Lynn, 1971). In these basins, it retains a thickness of several hundred metres (Johnson, 2008), and reaches an estimated age of several centuries (e.g. Broecker, 1979; Hirst, 1999).

Here, we present a new tool which allows us to track water masses to and from their point of formation, and apply it to NASMW and NADW. We demonstrate the use of this development by considering the life cycles of the above water masses in a new framework. The paper is set out as follows. In Sect. 2, we summarise the mathematical theory of the TAM approach. We then describe the model, our modifications to it, and how these may be applied to water mass tracking. The application of this development to NASMW is discussed in Sect. 3 with similar considerations for NADW in Sect. 4. We conclude in Sect. 5 with a summary of our method and findings, highlighting any recommendations for future work.

2 Development of the passive tracer module

2.1 Mathematical background

The tangent-linear method describes the evolution of perturbations to the ocean state in a linear framework. Perturbations are described by vectors following the structure of the ocean state vector, containing values for all prognostic variables at all locations. The linear evolution of an initial perturbation $|\mathbf{u}_0\rangle$ to the ocean state vector is described by the equation

$$|\mathbf{u}(t)\rangle = \Psi(t, t_0)|\mathbf{u}_0\rangle \quad (1)$$

where $|\mathbf{u}(t)\rangle$ provides the condition of the perturbation at time t , $\Psi(t, t_0)$ is known as the propagator matrix and we have used the “bra-ket” notation of Dirac (1939) in which row and column vectors are respectively written as bras $\langle\mathbf{a}|$ and kets $|\mathbf{b}\rangle$ such that closed bra-ket terms become scalar (through the scalar product $\langle\mathbf{a}|\mathbf{b}\rangle = c$).

The adjoint approach considers the sensitivity of properties of interest to earlier perturbations. For a “cost function” J (mapping the state vector to such a property) which is scalar-valued and linear, that is $J(\mathbf{u}) = \langle\mathbf{F}|\mathbf{u}\rangle$ (for some \mathbf{F}), the sensitivity is given by

$$\frac{\partial J}{\partial|\mathbf{u}_0\rangle} = \Psi^\dagger(t_0, t)|\mathbf{F}\rangle \quad (2)$$

where $\Psi^\dagger(t_0, t)$ is the adjoint of the propagator matrix, correspondent to its transpose in a Euclidean inner product space. These relationships are derived in full by Errico (1997), who provides a thorough treatment of the linear method and its limitations, beyond the simplified, inherently linear use-case of the present study. Furthermore, Finally, we have

$$\langle \mathbf{F} | \Psi(t, t_0) | \mathbf{u}_0 \rangle = \langle \mathbf{u}_0 | \Psi^\dagger(t_0, t) | \mathbf{F} \rangle \quad (3)$$

following the definition of the adjoint as a linear operator.

2.2 Passive tracer implementation and model description

Perturbations $|\mathbf{u}_0\rangle$ such as in Eq. 1 typically induce an active, or dynamic response (Marotzke et al., 1999). For temperature or salinity perturbations, this corresponds to a modified density field evoking changes in streamlines by modifying pressure gradients. This occurs in tandem with the passive advection and diffusion of the perturbed tracer. We may represent this by splitting the propagator into active (A) and purely passive (P) components $\Psi(t, t_0) = \Psi_A(t, t_0) + \Psi_P(t, t_0)$. Note that we use the term “purely passive” for Ψ_P because the active component Ψ_A encompasses all dynamic interaction between the unperturbed ocean state and the perturbed field. This itself includes a partially passive element (for example passive transport of dynamically modified tracers). The purely passive component $\Psi_P(t, t_0)$, on the other hand, involves no such interactions.

In the case of propagating a passive tracer, $\Psi_A(t, t_0)$ vanishes. Our perturbation is now merely an initial injection of dye of a specified concentration, $|c_0\rangle$, and so Eq. 1 collapses to

$$|c(t)\rangle = \Psi_P(t, t_0) |c_0\rangle \quad (4)$$

This dye is unable to feed back on the ocean state, and so as a perturbation is trivial. We will hence refer to initial perturbations as “injections” and evolved perturbations as dye concentrations, as is more appropriate for the passive tracer case. The common term “cost function”, which relates to optimisation problems, is also a misnomer in the more simple context of propagating passive tracers. In our analysis, $\langle \mathbf{F} |$ acts on the concentration vector $|c(t)\rangle$ to produce a tracer budget. As such we will more suitably refer to cost functions as budget co-vectors, $\langle \mathbf{B} |$.

Hence, in summary, the tangent-linear model Ψ_P describes the time evolution of dye concentrations $|c(t)\rangle$ in response to initial dye injections $|c_0\rangle$. Accordingly, the adjoint model Ψ_P^\dagger describes the sensitivity of tracer budget co-vectors $\langle \mathbf{B} |$ to earlier dye injections $|c_0\rangle$.

The model used herein is the NEMO 3.4 OGCM (Madec, 2012) with the tangent-linear and adjoint package NEMOTAM (Vidard et al., 2015). In particular, the former (nonlinear model) provides the unperturbed, nonlinear model trajectory. Meanwhile the latter (linear model) provides a model

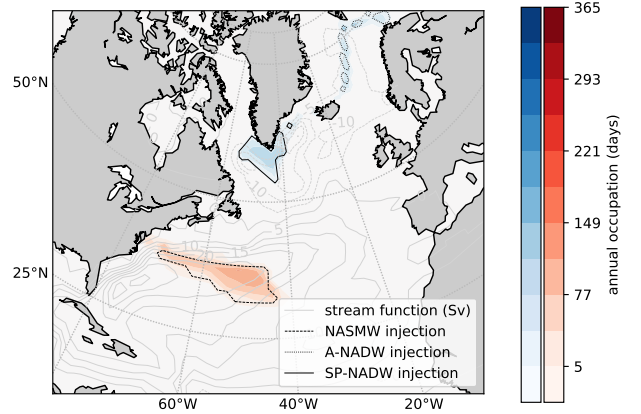



Figure 1. Outcrop locations of NASMW (red) and NADW (blue) from a 60 year climatology of the nonlinear model, following their respective definitions in Sections 3 and 4. Shading indicates the number of days of the year when the water mass is found at a given latitude and longitude. Light grey contours show the time-mean barotropic stream function for the North Atlantic. The dashed (solid, dotted) contour shows the distribution of the dye injection which initialises the tangent-linear run for NASMW (SP-NADW, A-NADW) at the annual maximum (see Section 3.2 (4.2.1, 4.2.2))

implementation of the propagator and its adjoint. Both models are used in the ORCA2-LIM configuration. This consists of a 2° grid, with refinement of the equatorial and Mediterranean regions (Madec and Imbard, 1996). There are 31 vertical levels, whose height follows a hyperbolic tangent function of depth, ranging from 10 m at the surface to 1000 m. There are two ice layers (Fichefet and Maqueda, 1997) in the background state, although these are not incorporated into NEMOTAM. The ice model is not linearised or coupled to NEMOTAM, although in the context of this study, ice dynamics are inherently unaffected (as a tracer which is passive cannot affect ice). The model is forced using the single, repeating normal year of the CORE forcing package (Large and Yeager, 2004). We therefore consider our results in a climatological, as opposed to historical, context. In this configuration, model NASMW (see Section 3 for more details on the definition and properties of NASMW in the nonlinear model) most persistently outcrops in the neighbourhood of the time-mean North Atlantic barotropic stream function maximum, which is 38.8 Sv (Fig. 1, red shading and grey contours). Below the surface, it occupies a narrow latitudinal band at depths of up to 240 m–310 m (Fig. 2, red shading). Conversely, model NADW (as defined and described in Section 4) most often surfaces in the region close to the barotropic stream function minimum of -25.3 Sv (Fig. 1, blue shading and grey contours). At most latitudes, NADW occupies depths below the local time-mean meridional stream function maximum (Fig. 2, blue shading and grey contours). The overall time-mean North Atlantic meridional stream function maximum is 16.1 Sv, occurring at 42°N at a depth of 870 m.

Figure 2. Zonally averaged distribution of NASMW (red) and NADW (blue) in a 60 year climatology of the nonlinear model, following their respective definitions in Sections 3 and 4. Shading indicates the number of days of the year when the water mass is found at a given latitude and depth. Light grey contours show the meridional stream function of the Atlantic. Dashed contours show the distribution at the point of the year when the adjoint run is started **Note: mistake in y-axis ticks corrected**



To isolate the purely passive response detailed in Eq. 4, we set velocity and sea surface height (SSH) modifications to zero in the NEMOTAM time stepping procedure (Fig. 3). Further feedbacks involving parameterisations such as vertical and eddy-induced mixing are already absent in NEMOTAM due to approximations in linearisation (Vidard et al., 2015) and so did not require further modification. ~~Additionally, eddy-induced advective velocities, although present in the nonlinear trajectory, were not included for our passive tracers in NEMOTAM.~~ Additionally, eddy-induced advective velocities are computed online by the nonlinear model and not stored explicitly in the trajectory. They may be recomputed from other trajectory variables, but this is not the default behaviour of NEMOTAM. As such, they were not included for the passive-tracer experiments in this study. They have been since been enabled in the passive-tracer source code as an optional online calculation. A preliminary comparative test suggests that errors are minor for our case studies (positive and negative concentration differences tend to cancel out, but even summing absolute differences, we do not exceed $\sim 6\%$ relative to the injection volume). We expect the difference to be greater in steep isopycnal regions such as the Southern Ocean.

Our water mass tracking procedure is as follows. We begin by defining the water which we intend to track, for example using a temperature-salinity (TS) range. To track the evolution of newly formed water of this type, we identify all surface locations where T and S properties meet our definition in the nonlinear model trajectory. We then inject into the tangent-linear model a dye concentration of 1 in these locations, and run the model forward. To identify the origins of existing water using the adjoint model, the procedure is similar. We begin by identifying model grid cells in the nonlinear trajectory (at all depths) which satisfy our TS definition. The budget vector is populated using the volume of these grid cells, with zero values elsewhere. This budget vector

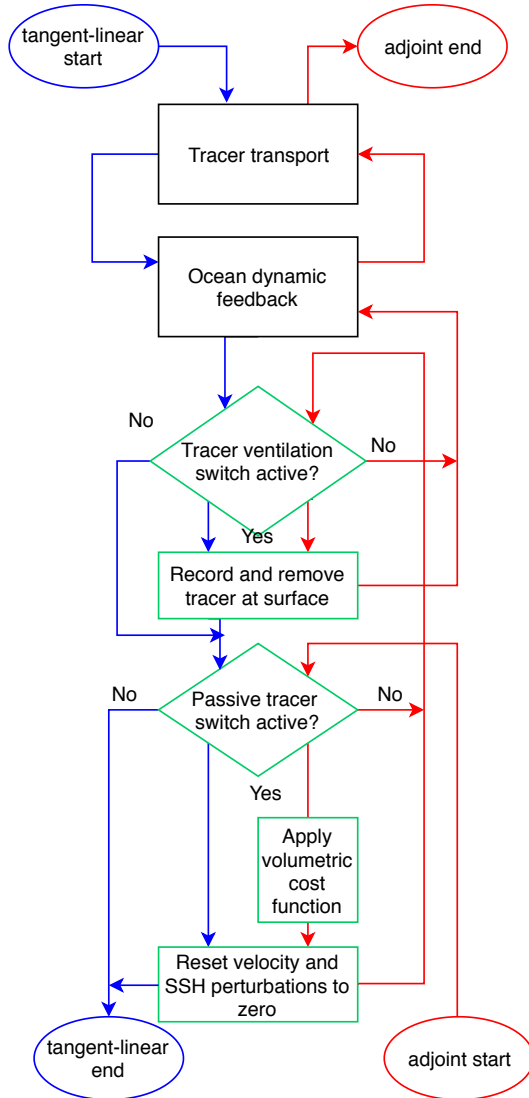
is then propagated back in time using the adjoint model to produce a sensitivity of the budget to earlier concentrations. In both cases, we remove any tracer which reaches the surface using the model's surface temperature restoring scheme. This scheme restores concentrations towards zero with a time scale of 60 days for a 50 m mixed layer (Madec, 2012). In the adjoint case, a record of this restoring leads to a spatiotemporal probability distribution of the surface origins of the tracer.

2.3 Advection schemes

The default advection scheme of the nonlinear model, used to compute the background state around which the model is linearised, is a total variation diminishing (TVD) scheme. This is a flux-corrected transport scheme (FCT, Lévy et al., 2001), which balances an upstream scheme with a centred scheme using a nonlinear weighting parameter described by Zalesak (1979). This scheme is preferable to others available in the model for non-eddy configurations (Lévy et al., 2001), but its nonlinearity means that it is not suitable for a TAM. As such, the standard advection scheme of NEMOTAM is a linearised counterpart in the form of a second-order forward-time-centred-space (FTCS) finite difference scheme. A caveat of such schemes is the presence of a negative artificial diffusion coefficient, which can lead to negative quantities of positive-definite tracers (Owen, 1984). This encourages down-gradient diffusion of otherwise zero concentrations, acting as a source of further negative tracer when combined with the surface restoring scheme. This is particularly problematic in the case of passive-tracer advection as, unlike in the fully active TAM, negative quantities are unphysical.

An alternate scheme which is first-order linear (thus compatible with tangent-linear models and added to NEMOTAM for this study) is the trajectory-upstream scheme. This is an approximation to the classical upstream scheme, whereby

Figure 3. Adaptation of NEMOTAM time-stepping procedure for the tangent-linear (blue) and adjoint (red) cases. Steps outlined in green are new additions to the scheme.



the upstream direction is inferred from the trajectory velocity field (avoiding nonlinear sign functions). As the velocity field cannot be modified in the passive tracer case, however, the method is everywhere equivalent to the classical scheme. The scheme is forward in both time and space, and offers the advantage of not producing negative tracer values. Nevertheless, this scheme also possesses an artificial (positive-valued) diffusion term, leading to a disproportionate spread of tracer, particularly in the vertical.

To strike a balance between the relative merits and caveats of the linear schemes described above, an approximation to the weighted-mean scheme of Fiadeiro and Veronis (1977) was coded into NEMOTAM. Like the TVD scheme, the true weighted-mean scheme is a linear combination of the classical upstream and FTCS methods. As before, we approxi-

mate the classical upstream method by using the trajectory-upstream scheme. The weighting parameter is determined dynamically from the **relative strength of local transport by advection vs. diffusion in the trajectory-trajectory**. Strongly advective regions lean towards upstream advection, while strongly diffusive regions favour the centred formulation. Despite inheriting issues from both FTCS and the trajectory-upstream schemes, it is preferable to either individually. Its negative tracer concentrations are less persistent than those of the TVD and FTCS schemes, and its anomalous vertical spread was found to be less than half that of the upstream scheme (Fig. 4). Some issues have been raised with multivariate and time-dependent extrapolations of the weighted-mean scheme (e.g. Gresho and Lee, 1981). However, nonlinear flux-corrected transport schemes are required in order to improve on the weighted-mean scheme (Gerdes et al., 1991), and these are not compatible with the linear model. We thus proceed using our approximation to the weighted-mean scheme (in both the horizontal and vertical) for the water mass case studies detailed in this paper. We remark that it is entirely possible to incorporate a nonlinear advection scheme into NEMOTAM, but that the resulting quasi-linear model would not have a true adjoint, which we require for our study.

2.4 Performance

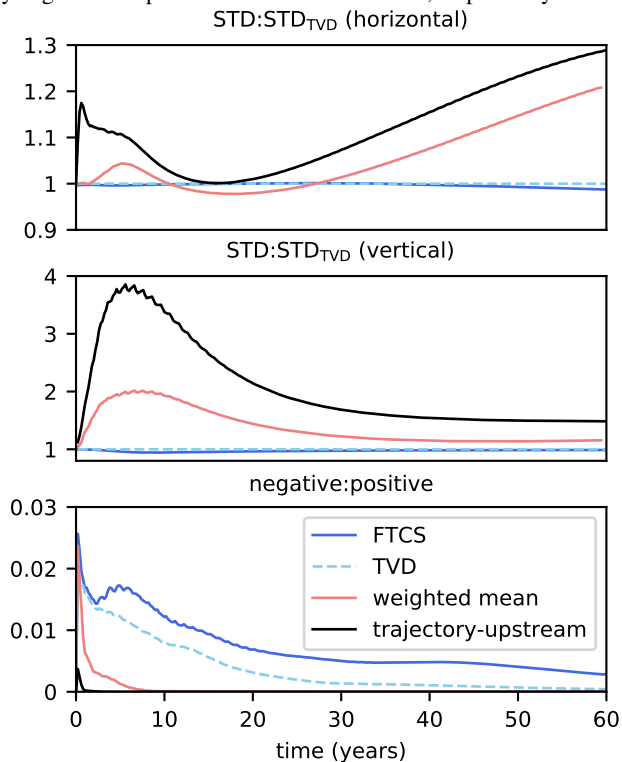
In order to test performance, the model was run for one-hundred days at 2° resolution (the ORCA2 configuration used in our case studies) and five days at 0.25° resolution (the ORCA025 configuration, Madec, 2012). Each configuration was run in four different modes (nonlinear, nonlinear with TAM trajectory production, tangent-linear with passive tracer, adjoint with passive tracer). Trajectory files were written once per day, and the linear advection scheme was the weighted-mean scheme described in Section 2.3. Each test was conducted with a range of parallelisation arrangements (16, 32, 64, 128, 256 and 512 CPU cores) and repeated 10 times to account for system variability (Table 1). The tests were conducted on a local HPC system, with 72 compute nodes, each offering two eight-core 2.6GHz processors (Intel Xeon E5-2650 v2) and 64GiB, connected by an InfiniBand QDR network.

The general order of time efficiency is consistent across all tests: the linear forward model is between 2.5 (16 cores) and 3.1 (64 cores) times as fast as the nonlinear model without trajectory output in ORCA2, and 1.9 (64 cores) to 2.0 (512 cores) times as fast in ORCA025. The adjoint is slightly less efficient. The nonlinear runs which produce trajectory output are the slowest across the board, due to the high level of output. Memory use appears higher in the linear model than the nonlinear, such that the linear model could not be run on two 64GiB compute nodes alone in ORCA025 (column two).

Table 1. Results of performance tests for two model configurations (ORCA2, upper section and ORCA025, lower section). Top rows: trajectory storage requirements per output. Lower rows: runtime per model day for four model modes: Nonlinear with trajectory output (NLT), nonlinear without trajectory output (NL), tangentlinear (TL) and adjoint (AD). The standard deviation of ten runs is given as a \pm uncertainty. Dashes indicate tests which failed due to insufficient memory.

Cores (nodes)			16 (1)	32 (2)	64 (4)	128 (8)	256 (16)	512 (32)
ORCA2								
Trajectory size (MB/d)			43.909	44.417	43.978	44.412	44.848	46.351
Mode	NLT	Run time (s/d)	2.09 \pm 0.20	1.54 \pm 0.17	1.23 \pm 0.17	1.39 \pm 0.26	2.53 \pm 0.27	4.08 \pm 0.35
	NL		1.98 \pm 0.17	1.37 \pm 0.10	1.10 \pm 0.09	0.99 \pm 0.08	1.85 \pm 0.36	2.18 \pm 0.70
	TL		0.78 \pm 0.01	0.50 \pm 0.01	0.36 \pm 0.09	0.36 \pm 0.13	0.45 \pm 0.18	0.88 \pm 0.10
	AD		0.89 \pm 0.01	0.61 \pm 0.02	0.53 \pm 0.04	0.50 \pm 0.11	0.59 \pm 0.05	1.99 \pm 0.61
ORCA025								
Trajectory size (MB/d)			-	5217.427	5218.495	5220.824	5217.658	5218.448
Mode	NLT	Run time (s/d)	-	1091.55 \pm 69.38	692.98 \pm 21.38	428.36 \pm 13.76	302.18 \pm 30.55	143.60 \pm 12.73
	NL		-	1065.48 \pm 54.56	656.94 \pm 13.74	386.34 \pm 13.89	262.72 \pm 34.56	131.70 \pm 13.28
	TL		-	-	315.26 \pm 26.17	207.82 \pm 14.43	136.42 \pm 13.00	64.70 \pm 4.20
	AD		-	-	403.74 \pm 36.18	252.36 \pm 14.17	181.12 \pm 7.70	92.40 \pm 6.26

Figure 4. Comparison of advection schemes in NEMOTAM. To assess the schemes, one of the runs of this study (the tangent-linear run of Sect. 3.2) was repeated with four different schemes and the tracer distributions were compared. Top panel: Lateral spread (spatial standard deviation around mean tracer position) as a ratio to the corresponding spread in the nonlinear TVD scheme, Middle panel: vertical spread (as above). Bottom panel: ratio of volumes occupied by negative and positive tracer concentrations, respectively.



The model shows good scaling in the ORCA025 configuration at all CPU arrangements tested here, but begins to

worsen in ORCA2 beyond 128 CPU cores. Further, the added time required for trajectory output in nonlinear ORCA2 runs on 128 CPU cores can be considerable. This is possibly due to the generation of a very large number of files during the model run, as the number of files generated is proportional to the number of CPU cores, the trajectory write frequency, and the run length. The required storage size remains relatively constant for runs with larger number of CPU cores, despite the greater number of files (Table 1, first row).


An additional limitation which we discovered during our longest experiments comes from system file-number limits, which can readily be reached for typical systems for very long runs using many cores.

3 Application to North Atlantic Subtropical Mode Water

3.1 NASMW definition and properties

We now apply the developments of Sect. 2 to the problem of tracking the origins and fate of NASMW in the ORCA2-LIM simulation. There is no universally accepted definition of NASMW, with differing definitions leading to conflicting results between studies (Joyce, 2011). Here, While the method allows us to define water masses in terms of any model variable, we choose for simplicity to utilise the common approach of a temperature-salinity range. In particular, we define model NASMW as water lying in the temperature band $[17, 19]^\circ \text{C}$, as in the original definition of Worthington (1959), with salinity in the range of $[36.4, 36.6]$ psu. As in some other studies (e.g. Kwon and Riser, 2004; Gary et al., 2014), we impose a third criterion, in our case that NASMW has a thickness of at least 125 m. This condition ensures homogeneity and isolates the water mass from the adjacent Madeira Mode Water (MMW) to its east (Siedler et al., 1987).

Figure 5. Mean-year time series of NADW (blue) and NASMW (red) volume (top panel) and outcrop surface area (bottom panel) including those of the Arctic (dashed line) and subpolar (dotted line) NADW outcrops from the nonlinear trajectory. Circles indicate the injection time of NASMW (red), NADW (blue) and A-NADW (black) in the forward simulations (see Sections 3.2 and 4.2)

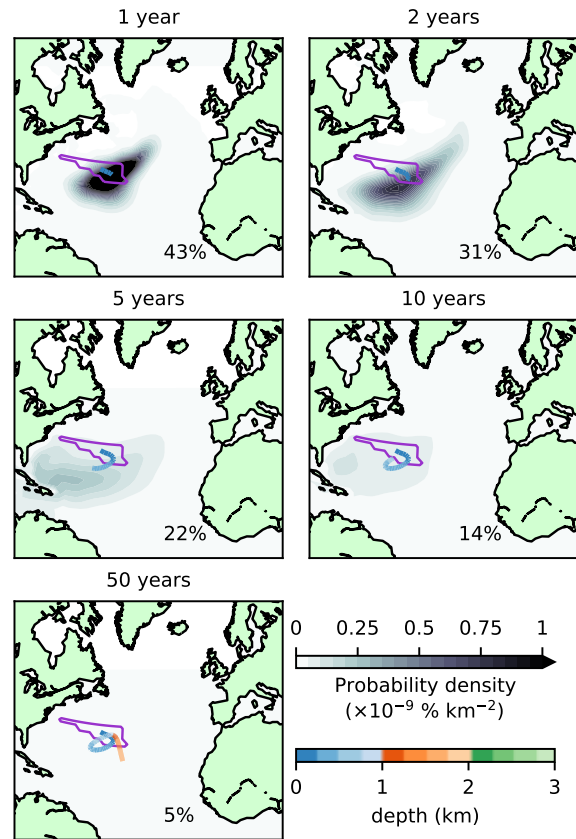


Model NASMW has a consistent outcrop in a single area (Fig. 1) and does not extend far south of its outcrop region below the surface (Fig. 2). The outcrop area of NASMW in the model peaks at the beginning of April at $1.3 \times 10^6 \text{ km}^2$ (Fig. 5), coinciding with peaks in volume ($4.2 \times 10^5 \text{ km}^3$) and thickness (310 m) of the water mass. These maxima occur a month before the annual maximum local MLD of 430 m. Following this, rapid stratification due to summer warming shoals the mixed layer. The outcrop area concurrently diminishes until the water mass is completely sheltered from air–sea exchange. This period of submergence extends from early June to early December, as the water mass’ upper surface progressively deepens to a maximum of 101 m. The volume also decreases, towards $1.3 \times 10^5 \text{ km}^3$ at the annual minimum.

3.2 Tangent-linear run

As outlined in Sect. 2.2, we begin by identifying all NASMW water matching our NASMW definition in the surface layer of the nonlinear model at a given time. We choose this to be the time when the outcrop is maximal (Fig. 5). We then inject a concentration of 1 into the tangent-linear model at these locations, and propagate it for 60 years. Newly-ventilated NASMW—Newly-ventilated model NASMW initially resides close to the surface (Fig. 6). Here, its behaviour is governed by surface currents, which fractionate the water mass. At the point of application, 37% of the tracer lies in the Gulf Stream. Most of this water follows the recirculation of the subtropical gyre, while a minority is exported to the subpolar region (2.5% over four years). Our quantification of subpolar export is on the same order of that found under similar conditions by Burkholder and Lozier (2011), who analyse inter-gyre exchange in a purpose-built Lagrangian study.

Figure 6. Evolution of NASMW-tagged tracer in the tangent-linear model. Black and white shading indicates depth-integrated probability density (corresponding to the likelihood that NASMW is found in that region) at 1, 2, 5, 10, and 50 yr. Inset percentages show the global integral of this field, i.e. the total proportion of tracer not yet re-ventilated at the surface. The coloured line tracks the centre of mass of the tracer from initialisation to its current position, with colours indicating the mean depth of the tracer. The purple contour shows the distribution of the original tracer injection.



Our NASMW is short-lived. Persistent proximity to the surface leads to the re-ventilation of 95% of the initialised tracer over the 60 year run. Of that which remains in the ocean, 90% is transformed and no longer fits the definition of NASMW (Fig. 7). As time progresses, because of vertical mixing together with near-surface tracer removal by the restoring scheme, remaining tracer mechanically tends to reach deeper, colder, fresher waters. Using the passive tracer approach, it is not possible to establish a full, qualitative description of the fate of NASMW. However, it can be observed that only around 5% of all model NASMW re-ventilation occurs within its surface outcrop. The rest is mixed out of the TS class, either remaining in the system for the remainder

of the 60-year simulation, or (more commonly) resurfacing elsewhere.

Due to rapid re-exposure, the e-folding time of our NASMW is just 60 days. This is shorter than the estimations of Gary et al. (2014) (1 yr) and Kwon et al. (2015) (3 yr), but this is likely due to inherent differences in formulation. In the above studies, Lagrangian particles released in NASMW are not removed upon contact with the surface, so as to account for re-ventilation. Accounting for re-exposure, Fratantoni et al. (2013) suggest around 75 days, more closely aligned with our own findings.

3.3 Adjoint runs

To track existing model NASMW back to its source locations, we again follow the procedure outlined in Sect. 2.2. This consists of taking a budget of NASMW in the final year of a nonlinear simulation and propagating this budget backward using the adjoint model. We begin the adjoint run at the same point in the annual cycle as the tangent-linear run, when model NASMW volume is at its maximum (Fig. 5). We bisect the mode water budget co-vector $\langle \mathbf{B}_{MW} \rangle$ into tracer above the mixed layer depth ($\langle \mathbf{B}_{MW}^U \rangle$) and tracer below ($\langle \mathbf{B}_{MW}^L \rangle$), and propagate each part separately. This allows us to explore the potential for a resilient layer of older NASMW below the MLD (Davis et al., 2013). By linearity, the union of the propagated budget vectors is equivalent to the propagated budget vector of NASMW as a whole:

$$\Psi^\dagger(t_0, t)(\langle \mathbf{B}_{MW}^U \rangle + \langle \mathbf{B}_{MW}^L \rangle) = \Psi^\dagger(t_0, t)\langle \mathbf{B}_{MW} \rangle$$

Most NASMW propagated with the adjoint model reaches the surface quickly; 70% of the tracer-tagged water mass is under 5 years old. During this early stage, ventilation occurs predominantly within the subtropical gyre recirculation, in the neighbourhood of the model's initial NASMW outcrop (Fig. 8). The 60-year mean formation location of NASMW is also within this region, at 32° N, 58° W. This is almost coincident with the core of NASMW formation determined by Warren (1972), and agrees with the air–sea exchange-based estimate of Worthington (1959).

For NASMW-water over 5 years old, current patterns begin to have a distinct influence on formation. There is a clear signature of the Gulf Stream on the youngest model NASMW. This evolves backwards as the adjoint propagates, eventually leaving an imprint of the entire subtropical gyre (Fig. 8). Simultaneously, newly formed mode water in the eastern North Atlantic (MMW) begins to cross into NASMW within 10 years. The signature of the Mediterranean outflow is particularly strong on the very oldest model NASMW, which also contains contributions from the subpolar region. The culmination of all of these water types throughout the 60-year evolution is evident when the model's sources of NASMW are viewed in TS space (Fig. 9). Also apparent is that the primary source of NASMW has a warmer signature, which reflects the advection of warmer waters from the south, which cool and subduct to form the water mass. While the surface

distribution suggests that a relatively small neighbourhood dominates the formation of model NASMW (Fig. 8), we remind that the seasonal variability of surface properties (particularly the outcrop area) reflects strongly on the thermohaline properties of NASMW at formation (Fig. 9). This suggests that the NASMW surface outcrop is not, in fact, the dominant origin of NASMW in the model, contrary to the intuition provided by laminar models of the ventilated thermocline (e.g. Luyten et al., 1983), and that ocean dynamics play a fundamental role in mode water formation. The implication is that mode water formation is not a passive process, and that ocean dynamics must play a fundamental role.

We also consider the time scales involved with NASMW formation in the model. By recording the time at which tracer is removed from the budget by the surface restoring scheme, we may construct a probability distribution function (PDF) of water mass age (Fig. 10). This distribution resembles a log-normal PDF with a There is a visibly lower skewness for NASMW lying below the MLD (Fig. 10, blue bars). Also clear is the seasonal cycle of shielding brought on by summer stratification, with less NASMW reaching the surface during these periods. Due to the simulation beginning at the annual maximum in NASMW production, peak NASMW formation occurs almost instantly. For tracer below the mixed layer (3.78% of the total) the peak does not occur until the outcrop maximum of the third year. From the PDFs, the expected age of the model NASMW constituents above and below the MLD are, respectively, 4.70 yr and 8.79 yr. It is evident that NASMW located below the MLD is shielded from renewal, with nearly double the expected age of model NASMW as a whole.

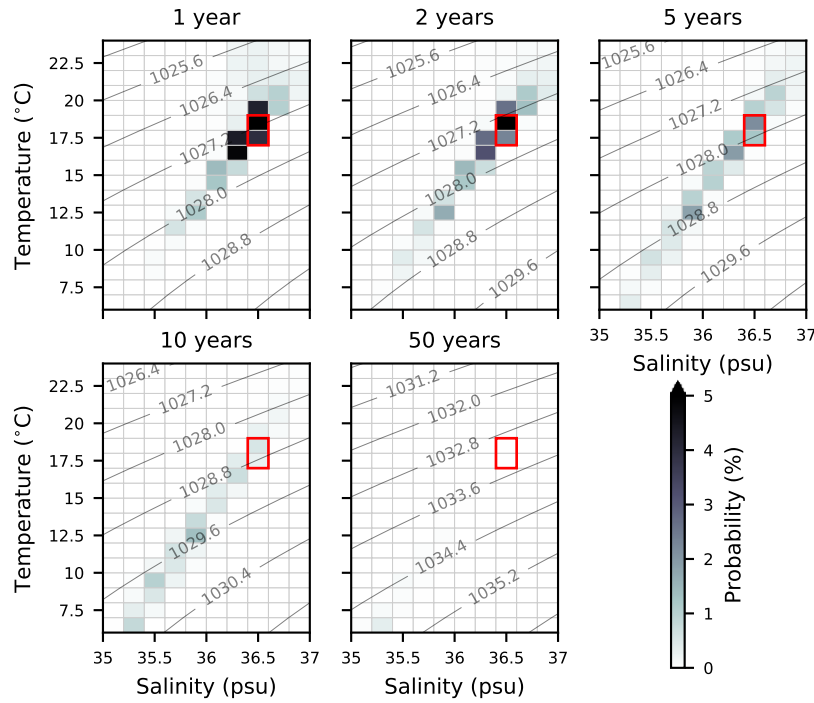
We finally address the asymptotic tail of the PDF, which represents the oldest waters found within NASMW. Our findings suggest a high-latitude source makes up a large fraction of this water, having followed a pathway from outside of our defined thermohaline range from the surface to eventually contribute to the make-up of NASMW (Fig. 8). Indeed, we find that for the 30–50 yr period, some 68.3 % of NASMW originates at latitudes of 45° or higher. The idea of a distant source of the oldest NASMW is concurrent with the few studies which have considered it. Douglass et al. (2013) use an ideal-age tracer to construct a histogram similar to our own, acknowledging a remote source of the very oldest NASMW. Kwon et al. (2015) find that at least 20% of NASMW stems from other regions, and highlight the potential importance of subpolar latitudes.

4 Application to North Atlantic Deep Water

4.1 NADW definition and properties

As in Sect. 3.1, we now consider the properties and behaviour of NADW in the ORCA2-LIM simulation. We define NADW to fall within a temperature–salinity band of [2, 4] °C, [34.9, 35.0] psu, in close alignment with the definition of

Figure 7. Evolution of model NASMW in TS space. Shading indicates likelihood that the water mass found in a particular TS class after 1, 2, 5, 10 and 50 yr. The red box marks the TS range used to define the water mass in this study. Contours show the density at the average depth level of the tracer.




Worthington and Wright (1970). Analysis of the nonlinear model trajectory (about which the model is linearised) shows that there are two distinct latitudes at which water of this TS signature persistently outcrops in the simulation (Fig. 2, blue shading). These correspond to the (subpolar) Labrador-Irminger Sea region, southwest of the Greenland-Scotland ridge, and the (Arctic) Greenland Sea region, northeast of the ridge (Fig. 1, blue shading). The outcrop oscillates between the two regions with the seasonal cycle. We note that, while observed NADW is known to form only in extreme winters and has a strong interannual signature (Avsic et al., 2006), our use of repeated forcing implies that formation has little year-to-year variation. The Labrador-Irminger outcrop peaks at the end of March, around a month after the annual local mixed layer depth maximum of ~ 1 km. At this peak, water in the our NADW TS class occupies 5.3×10^5 km² at the surface (Fig. 5), before the area diminishes with the shoaling mixed layer. The model's Greenland Sea NADW outcrop peaks in November, with a maximal extent of 2.3×10^5 km². At this time, surface NADW is exclusively found northeast of the Greenland-Scotland ridge, and the water mass is shielded from ventilation elsewhere. On average, 66% of the total annual outcrop area is southwest of the ridge, and 34% is to its northeast.

Model NADW volume is almost constant year-round at 4.9×10^7 km³ (18.7% of the model North Atlantic), deviating by less than 0.5% annually.

4.2 Tangent-linear runs

As with NASMW, the dye injection $|c_0|$ coincides with the time step corresponding to the annual maximum outcrop extent, which here falls in April. It should be noted that all of the locations corresponding to this maximum are southwest of the Greenland-Scotland ridge, in the subpolar Labrador-Irminger region. To investigate the nature of the Arctic outcrop, we follow a second injection of dye, in November, when its own outcrop extent is maximal. At this point, the water mass exclusively surfaces within the Arctic circle. We thus refer to these distinct surface waters as SP-NADW (for the southwestern, subpolar outcrop) and A-NADW (for the northeastern, Arctic outcrop). The tangent-linear model was run for 400 years with the SP-NADW dye injection, and 60 years with the A-NADW dye injection. The run lengths were determined by the rate of surface tracer re-ventilation in each case.

Figure 8. Surface origins of **model** NASMW as determined by the backtracking budget analysis (adjoint model simulation). Shading indicates probability density. This corresponds to the likelihood that **model** NASMW is formed in a given region during the time periods [0,5] yr, (5,10] yr, (10,30] yr and (30,50] yr. Inset percentages show the global integral (that is, the total proportion of the budget which is formed during each time period). **Note that, contrary to Fig. 6, which displays instantaneous fields, here time-integrated fields are displayed. Due to the large variability of its extent over the integration window, the outcrop region is not shown.**




4.2.1 SP-NADW

SP-NADW rapidly sinks, with tracer reaching an average depth of 1235 m after 0.2 yr. It initially moves quickly westward, departing the surface region around Cape Farewell. It then spreads throughout the Labrador basin at depth and extends into the Irminger Sea (Fig. 11) at a temperature of 3.5° C (Fig. 12). This initial behaviour is in broad agreement with temperature data from hydrographic sections (e.g. McCartney and Talley, 1982) and spatial distributions captured by CFC measurements (e.g. Rhein et al., 2002), **profiling** floats (e.g. Bower et al., 2009) and models (e.g. Bower et al., 2011; Gary et al., 2011).

The tracer patch then moves southward, steadily deepening. Its deepest average point, 2466 m, is reached after some 22 yr. Its mean position is at first closely tied to the DWBC. However, beyond the Flemish Cap, it takes a more central course through the basin interior. While interior southward

Figure 9. Surface origins of **model** NASMW as determined by the backtracking budget analysis (adjoint model simulation) in TS space. Shading indicates the proportion of the budget originating from a particular TS class at the surface within 50 yr. The red box marks the TS definition of NASMW used in this study. Contours show surface density.



routes generated by deep eddies have been found parallel to the DWBC in recent profiling float data (Lavender et al., 2000; Bower et al., 2009), our model configuration is non-eddying. As such, these pathways are represented by parameterised turbulent diffusion of the tracer, in a manner which would not be captured by Lagrangian drifters in our model (Gary et al., 2011).

The tracer initialised in the model is quickly sequestered, and is thus not vulnerable to re-exposure. Indeed, while 27% of the initial volume of SP-NADW is re-ventilated within the first decade, only a further 24% is removed during the rest of the century (Fig. 13). The tracer's homogeneity is also preserved, with limited mixing into neighbouring TS classes on decadal time scales (48% remains in the original TS class after 20 years, Fig. 12). There is a tendency over hundreds of years for the remaining water to cool and freshen to temperatures lower than 1.25° C and salinities below 34.65 psu (Fig. 14). This follows the well-described behaviour of NADW following intrusion into the Southern Ocean. **Observed** NADW is known to mix into Circumpolar Deep Water, eventually transforming into bottom water with similar thermohaline characteristics to these (e.g. Orsi et al., 1999).

4.2.2 A-NADW

The forward evolution of A-NADW (Fig. 15) is altogether different to SP-NADW, categorised by rapid diffusion and

Figure 10. Probability distribution of **model** NASMW age (hatched bars) and **NASMW age-age of NASMW restricted to be found below the mixed layer depth** (blue bars). The percentage of the NASMW budget formed in a given 0.5 yr bin is indicated by its associated bar. The expected value of the distribution is marked by a solid line.




Figure 11. As in Fig. 6 but for subpolar-outcropping **model** NADW, and times 2, 10, 20 and 50 yr.

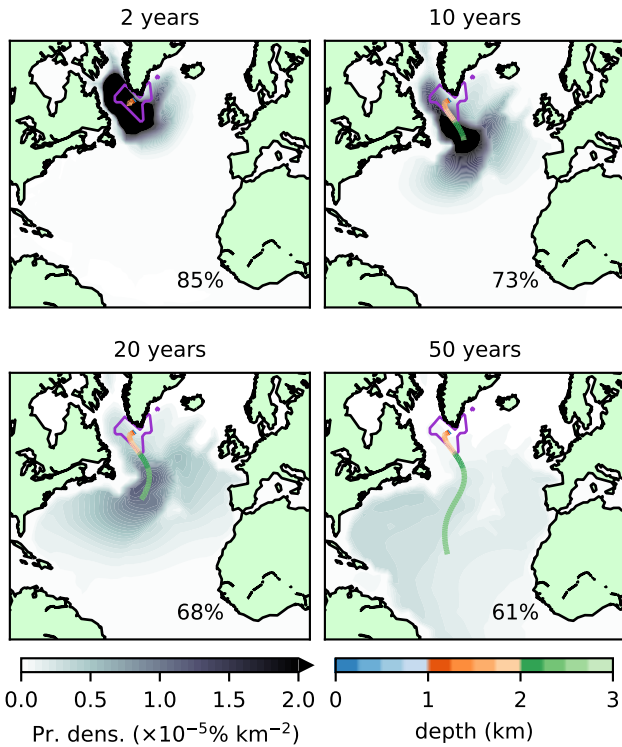
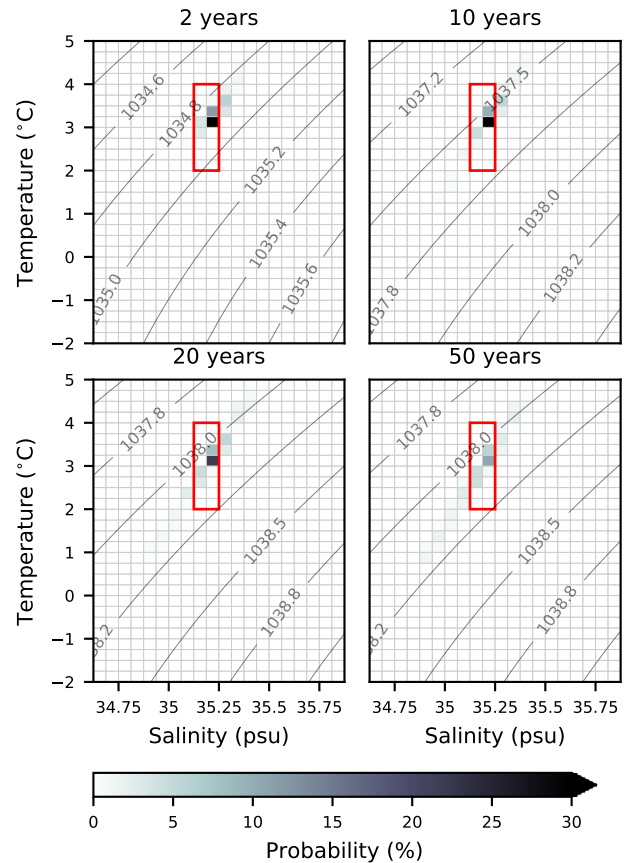



Figure 12. As in Fig. 7, but for subpolar-outcropping **model** NADW at 2, 10, 20 and 50 yr. Note that the region enclosed by the red box is the same in both **subpolar- and Arctic-outcropping** (Fig. 16) cases.



depletion. Unlike SP-NADW, which reaches great depths quickly following formation, A-NADW remains close to the surface on the Greenland and Barents shelves, with an average depth of only 502 m after 1 yr. Due to this surface proximity throughout the run, the majority of tracer is re-exposed to the atmosphere. It is accordingly removed by the

Figure 13. As in Fig. 6 but for subpolar-outcropping **model** NADW, and times 100, 200 and 400 yr.



model’s surface restoring scheme, with 95% of the initial volume removed after 20 yr. That which remains either spreads throughout the Arctic ocean, or spills over the shelf to join its SP-NADW counterpart in the Labrador Sea.

5 Of all of the initialised A-NADW, that destined for the Atlantic basin represents just 3.8%. We may use the velocity fields of the nonlinear model to estimate transport pathways of this passive tracer into the basin. Consider an idealised case with two openings into the basin (here taken to represent the Denmark Strait and west of the Reykjanes Ridge),

at points y_1 and y_2 . The total volume flux into the basin between two times t_0 and t_1 will be approximately

$$\Delta V_{t_0, t_1} = \int_{t_0}^{t_1} ((u_1(\tau)c_1(\tau)\Delta y_1\Delta z_1) + (u_2(\tau)c_2(\tau)\Delta y_2\Delta z_2))d\tau \quad (5)$$

where c_i is the concentration of tracer at opening i , Δy_i is the width of the opening, Δz_i is the depth, u_i is the velocity normal to $\Delta y_i\Delta z_i$, t_0 is the injection time and t_1 is 50 years later. As all of the above quantities are present in the model output, we may scale this approach up to the model grid and estimate each contribution. In this manner, we find that the vast majority of North-Atlantic-destined A-NADW (88%) enters through the Denmark Strait.

Passive tracer injected at the A-NADW outcrop rapidly moves through TS space (Fig. 16. A colder, fresher water type splinters away from the surface-borne **model** NADW, and as a result, the tracer occupies two distinct regions in TS space for the remainder of the run. The warmer of these two water types shares most of its TS properties with those of the **model** NADW at the surface. On the other hand, the colder water type, at $(-0.5, 1)^\circ$ C, bears a similar temperature signature to observed OW (LeBel et al., 2008).

Our simulation shows that surface-borne A-NADW does not proceed to form a substantial part of subsurface NADW **in the simulation**. Within 0.6 yr, 78% of the tracer has been re-exposed to the atmosphere. Of that which remains, some 98% has left the NADW TS class. The little tracer which stays in **the-our defined** NADW class is persistent, eventually following a similar trajectory to SP-NADW through TS space.

The two tangent-linear experiments (following water from each of the two distinct NADW outcrop regions) suggest that the more northeasterly outcrop contributes quantitatively little to the NADW bulk. Hence, any surface origins of NADW in the **model’s** Arctic are likely found outside of the NADW TS class.

4.3 Adjoint run

As before, we take a water mass budget at the end of the nonlinear model simulation (in this case for NADW after 400 model years) and provide it to the adjoint model. Using this approach, 86% of the **model’s** NADW budget can be traced to its creation within 400 years (Fig. 17), **with the rest remaining in the ocean interior**.

During the first year, only 0.13% of the total volume of **model** NADW is tracked back to the surface. The strong presence of shallower NADW during this period leads to a clean signature of the two outcrops (Fig. 17, top panel). On multidecadal time scales, tracer from NADW can be traced back to locations spanning most of the northern North Atlantic (Fig. 17, centre panel), dominated still by the region

Figure 14. As in Fig. 7, but for subpolar-outcropping model NADW at 100, 200 and 400 yr. Note that the region enclosed by the red box is the same in both the subpolar- and Arctic-outcropping (Fig. 16) cases.




Figure 15. As in Fig. 11, but for Arctic-outcropping model NADW.

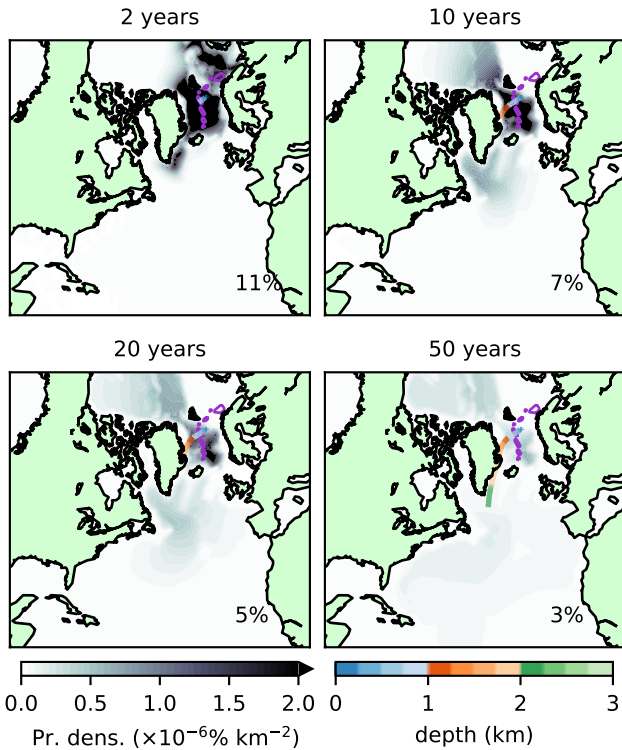
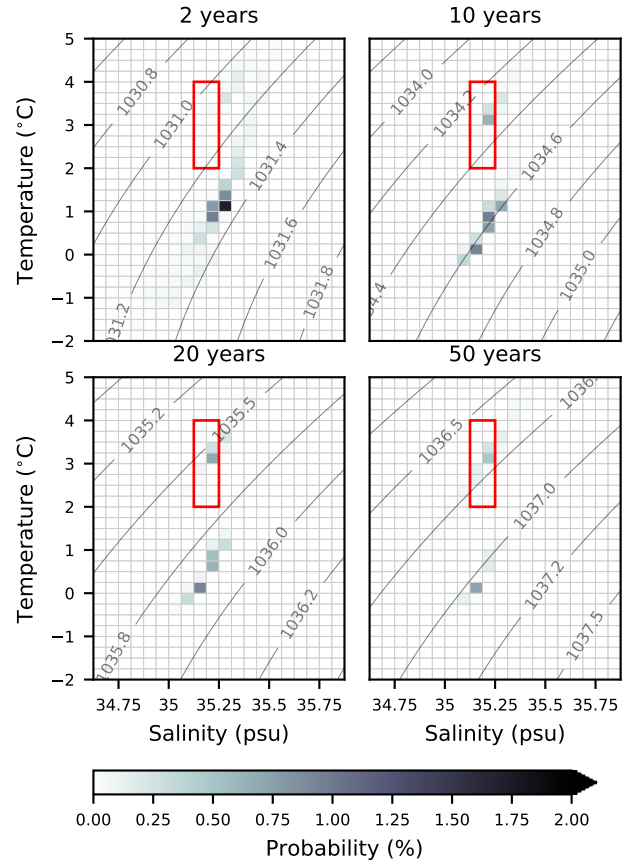


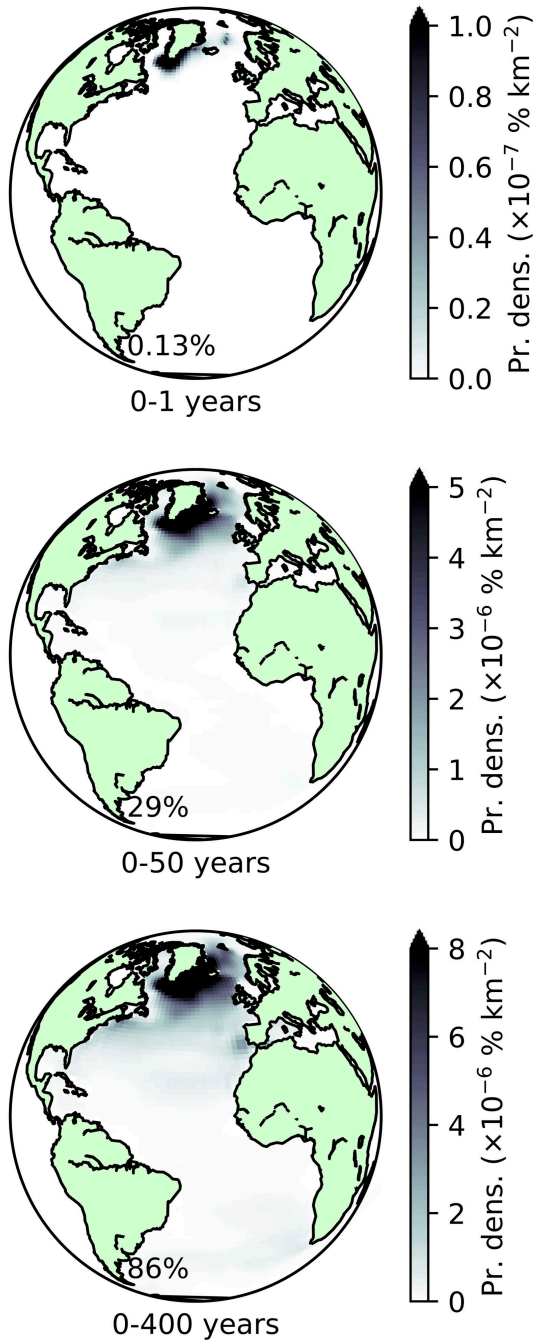
Figure 16. As in Fig. 12, but for Arctic-outcropping NADW. Note that the region enclosed by the red box is the same in both subpolar- (Figs. 12 and 14) and Arctic-outcropping cases.



surrounding its subpolar outcrop. Of particular interest on these time scales is the signature of the Mediterranean Outflow. It has been proposed from hydrographic data that the northward penetration of Mediterranean Water can intermittently reach the subpolar gyre, influencing LSW (Lozier and Stewart, 2008). However, we cannot may only speculate as to whether this mechanism is present in our simulation.

Modelled NADW spanning the entire 400-year run (Fig. 17, bottom panel) has sources throughout the Atlantic basin, with a notable contribution from the eastern boundary of the

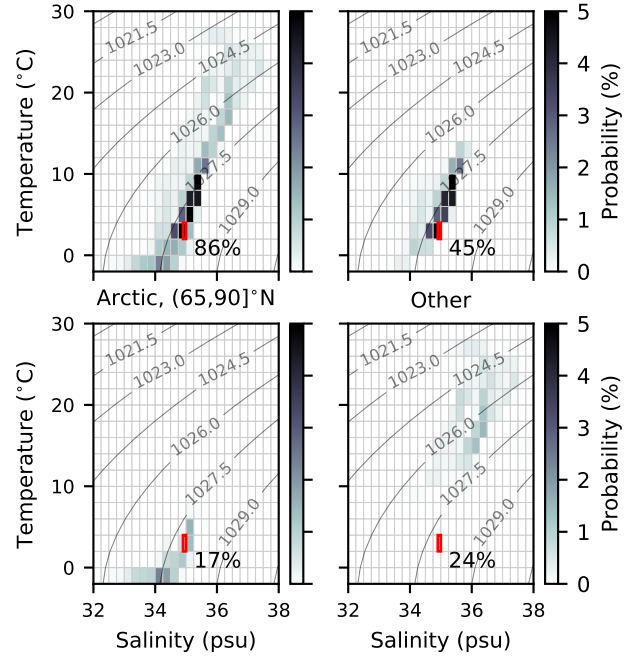
Figure 17. As in Fig. 8, but for NADW and time periods [0,1] yr (top), [0,50] yr (center) and [0,400] yr (bottom)



South Atlantic, local to the Benguela current. The Labrador-Irminger sector remains the primary source of NADW at all considered time scales, however.

The propagated budget vector can be separated into different source regions and signatures. For instance, 31% of the ventilated NADW can be traced back to the model's

Figure 18. As in Fig. 9, but for NADW and time period [0,400] yr. Top left: NADW total volume. Top right: NADW volume traced to the subpolar North Atlantic. Bottom left: NADW volume traced to the Arctic. Bottom right: volume of NADW originating elsewhere.



Irminger Sea, versus just 14% to the Labrador Sea. We may also construct a volumetric census of model NADW source waters in thermohaline space (Fig. 18, upper left panel). This may further be partitioned into waters originating from the two climatic regions (i.e. subpolar and Arctic) associated with the distinct outcrops of NADW (Fig. 18, upper right and lower left panels, respectively). 45 % of the total NADW formed during the 400-year run may be attributed to the subpolar region. The dominant TS class associated with NADW formation in the model matches that of our definition ($[2, 4]^\circ\text{C}$, $[34.9, 35.0]$ psu), though there are contributions from a cluster of subpolar water types in the range $[2, 10]^\circ\text{C}$, $[34.5, 35.5]$ psu. Despite the substantial seasonal surface exposure of NADW in the Arctic (34% of the annual outcrop area, see Sect. 4.1), this region ultimately accounts for only 17% of modelled NADW formation in 400 years. This agrees with our findings in the tangent-linear model (Sect. 4.2) that surface-borne NADW in the Arctic is subject to rapid re-ventilation and does not contribute to the NADW bulk. It further suggests that there is no other narrowly defined TS band which outcrops in the Arctic from which modelled Arctic NADW originates. As such, NADW from this region generally derives from a broad range of waters colder and fresher than NADW we define NADW to be.

NADW which does not originate from either of these two regions of the North Atlantic makes up a substantial proportion (24%, Fig. 18, lower right panel) of the budget. Of this,

10
15
20
25
30

4% is formed at high latitudes in the Southern hemisphere. 5% of the total originates from outside the Atlantic.

As in Sect. 3, we may also consider the temporal distribution of simulated NADW formation (Fig. 19). Age peaks in the 13th year, during which 0.6% of the total budget can be traced back to the surface. The expected age is 112 yr. It should be noted that the mean age of model NADW is likely slightly higher, ~~as the results detailed here do not quite account for 100% of the total budget.~~ As with any water mass studied in this manner, a proportion will always remain in the ocean (closing the overall budget). Due to this, surface ventilation does not quite account for 100% of the total budget during our 400-year run. We limit our runs to this length due to technical constraints. In Sect. 3.3, we remarked on the unusual formation regions associated with the very oldest North Atlantic Subtropical Mode Water. This was associated with its much shorter life cycle. As such, the corresponding tail of the PDF of NADW age holds fewer surprises. Much like the youngest NADW considered here, the oldest NADW originates predominantly from winter convection in the model's Labrador and Irminger seas.

5 Discussion and conclusions

We have presented a newly-developed addition to the NEMOTAM tangent-linear and adjoint modelling framework for the NEMO OGCM. This package allows tangent-linear and adjoint tracking of passive-tracer transport. Our framework is rooted in concentrations and probabilities, comparable to the statistical properties of a high-resolution Lagrangian approach with ~~infinitely many particles.~~ a large (theoretically infinite) number of particles. The development was achieved by deactivating dynamic feedbacks in the time-stepping routine of NEMOTAM, and embedding an alternative advection scheme suitable for passive-tracer transport. This advection scheme, proposed by Fiadeiro and Veronis (1977), is constructed as a linear combination of the upstream and centred schemes, so as to reduce spurious numerical diffusion. Performance tests show good scaling up to 128 cores at 2° resolution and up to at least 512 cores at 0.25° resolution, with much shorter run times (but higher memory demand) in the linear model with passive-tracer transport when compared with the fully-active nonlinear model.

We have exhibited the use of this tool in two case studies concerning the tracking of North Atlantic-borne water masses, North Atlantic Subtropical Mode Water and North Atlantic Deep Water. The method's versatility has been demonstrated through the calculation of several quantities pertaining to these water masses. If a sufficiently long trajectory is used, a near-complete statistical distribution of surface formation can be constructed. This allows diagnosis of expected age and expected origin location, as well as the rate of eradication and re-ventilation of newly ventilated waters.


The linearity of the method ensures that the water being tracked can be partitioned into several components. When these components are propagated separately, the tracking of the whole is equivalent to the tracking of their union.

The results of our case studies show good agreement between the model configuration and common aspects of prior observational and computational studies. We have shown, for example, that, in our simulation, on average, the expected surface origin of tracer initialised within NASMW is 32° N, 58° W (in line with Worthington, 1959; Warren, 1972), that its decay time is around two months (as in Fratantoni et al., 2013), and that a small minority (2.5% in four years) is exported to the subpolar gyre (concurrent with Burkholder and Lozier, 2011). We have estimated the average age of model NASMW as 4.5 yr, also close to observations (e.g. Kwon and Riser, 2004), and shown that a more persistent NASMW subset (with an expected age of over 8.5 yr) underlies the bulk of the water mass (as suggested by Davis et al., 2013).

For simulated NADW, we have shown (in the tangent-linear) that an Arctic outcrop of water with its signature to the northeast of the Greenland-Scotland ridge contributes little to its final form. However, we have found (by backtracking) that Arctic surface waters still make a contribution to simulated NADW formation (17% here), but from a broad range in thermohaline space. It is understood that OW is produced by the cooling and freshening of North Atlantic inflow to the Greenland Sea (e.g. Quadfasel and Käse, 2007; Dickson et al., 2002). The finer details of this transformation, and subsequent resupply of the transformed water into the North Atlantic are not well-captured by the method. The relative importance of pathways into the Atlantic for shelf water are still poorly known (Macrander et al., 2005), but we find using a broad transport estimate that the Denmark Strait is dominant in the model. For NADW formed in the subpolar outcrop of the model's Labrador Sea, our tracer spread reflects well that observed using CFCs (e.g. Rhein et al., 2002). The diffusive behaviour of the passive tracer also means that, despite the non-eddy nature of our model, eddy-driven southward export pathways are better represented than they would be by Lagrangian drifters (e.g. Gary et al., 2011).

Nevertheless, there are several more intricate details of North Atlantic water mass formation and transport which are the subject of ongoing investigation, and as such deserve a dedicated study at higher resolution. For example, the importance of fine-scale bathymetry for an accurate description of on-shelf NADW formation is well-described (Hansen et al., 2001; LeBel et al., 2008; Smethie Jr and Fine, 2001; Dickson et al., 2002). From a selection of models, Chang et al. (2009) find that the contribution of overflow waters to NADW is misrepresented by models coarser than $\frac{1}{12}^\circ$ in their study. In such models, the Faroe Bank Channel is typically unresolved and the Denmark Strait Overflow resultantly dominates, as is the case here. Furthermore, although our tracer qualitatively reproduces southward export of LSW more effectively than a Lagrangian approach would in a laminar model, diffusion

Figure 19. As in blue shading of Fig. 10, but for NADW. The black line is the expected age.



is still parameterised. It is unknown how well this parameterisation represents sub-grid-scale mixing for such processes.

Despite good broad agreement between the passive-tracer pathways and those noted in previous studies of these water masses, there is an interesting disparity between the forward and backward modes within the TAM itself. This originates from using a TS-based description to inform the initial tracer distribution. For example, the backtracking method suggests that NASMW predominantly originates in slightly warmer surface waters than those of the outcrop used to inform the forward model. Meanwhile, A-NADW, while occupying over a third of the simulated NADW annual outcrop, ultimately contributes almost nothing to the subsurface water mass. These deviations highlight the approximation used by many water-mass-tracking model studies - thermohaline characteristics are not a purely passive tracer. Water parcels experience changes in their thermohaline properties, and so water in a particular TS class at depth is not exclusively related to the same TS class at the surface through a passive advection pathway.

TAM use at high resolution is typically limited, due to baroclinic instability. However, this is not detrimental for passive tracer tracking, due to lack of dynamic feedbacks. As such, our tool may be used in conjunction with higher-resolution configurations of NEMO (e.g. ORCA12, Treguier et al., 2017). The main barrier to higher resolution for users of our development is the necessity of ~~long trajectory outputs, frequent output and storage of trajectory snapshots from the nonlinear model,~~ which are required for NEMO-TAM operation (Vidard et al., 2015). ~~Future versions of our tool will allow regional trajectory storage to overcome this barrier. For the water masses discussed here, we used a 400-year trajectory with repeated forcing. This captured the formation of at least 86% of the water masses considered, leaving a small portion of the budget unaccounted for.~~

Although we have presented the development in the context of water mass tracking, there are many potential further applications. Ocean heat uptake pathways and carbon sequestration have been studied by means of modelled passive tracers (e.g. Banks and Gregory, 2006; Xie and Vallis, 2012), and adjoint models (e.g. Hill et al., 2004). A slight modification to ignore vertical velocities read from the trajectory would force positive buoyancy on the tracer. This could allow buoyant anthropogenic pollutants to be tracked, with potential application to ocean plastic tracking and oil spill drift prediction. Water mass tracking may itself be complemented by considering continuous (rather than instantaneous) inputs of tracer at the surface. It is hoped that an off-the-shelf tool, bolted onto an existing OGCM will stimulate further research in these areas.

Despite TAM use for sensitivity analysis being traceable to the 1940's (Park and Xu, 2013), its application to ocean science is still in its infancy. The TAM approach is highly versatile, and its application to state-of-the-art OGCMs permits a great many new insights into ocean dynamics. Our development demonstrates the ability of a tangent-linear and adjoint model among the suite of existing water mass tracking methods. It is hoped that this novel tool will encourage new users to realise the potential of this powerful branch of ocean modelling.

Code availability. NEMO v3.4 is available from <https://forge.ipsl.jussieu.fr/nemo/svn/NEMO/releases/release-3.4> under the CeCILL licence. The configuration presented here is archived on Zenodo for general use (Stephenson, 2019a). The forcing files used in the simulations detailed here are also archived on Zenodo (NEMO Consortium, 2013), as are exact scripts and input files to reproduce our experiments and diagnostics (Stephenson, 2019b)

Author contributions. Dafydd Stephenson co-developed the source code, ran simulations and wrote this manuscript. The development of the source code was overseen by Simon Müller, who provided further suggestions and additional technical assistance. Florian Sévellec proposed and supervised the project, and assisted with experiment design. All authors discussed and contributed to the final manuscript.

Competing interests. No competing interests are present

Acknowledgements. This research was supported by the Natural and Environmental Research Council UK (SMURPHS, NE/N005767/1 and MESO-CLIP, NE/K005928/1 and the SPIT-FIRE DTP) and by the DECLIC and Meso-Var-Clim projects funded through the French CNRS/INSU/LEFE program.

References

- Aysic, T., Karstensen, J., Send, U., and Fischer, J.: Interannual variability of newly formed Labrador Sea Water from 1994 to 2005, *Geophysical Research Letters*, 33, 2006.
- Banks, H. T. and Gregory, J. M.: Mechanisms of ocean heat uptake in a coupled climate model and the implications for tracer based predictions of ocean heat uptake, *Geophysical Research Letters*, 33, L07 608, 2006.
- Blanke, B. and Raynaud, S.: Kinematics of the Pacific equatorial undercurrent: An Eulerian and Lagrangian approach from GCM results, *Journal of Physical Oceanography*, 27, 1038–1053, 1997.
- Bower, A., Lozier, S., and Gary, S.: Export of Labrador Sea water from the subpolar North Atlantic: a Lagrangian perspective, *Deep Sea Research Part II: Topical Studies in Oceanography*, 58, 1798–1818, 2011.
- Bower, A., Lozier, S., Biastoch, A., Drouin, K., Foukal, N., Furey, H., Lankhorst, M., Rühls, S., and Zou, S.: Lagrangian Views of the Pathways of the Atlantic Meridional Overturning Circulation, *Journal of Geophysical Research: Oceans*, 2019.
- Bower, A. S., Lozier, M. S., Gary, S. F., and Böning, C. W.: Interior pathways of the North Atlantic meridional overturning circulation, *Nature*, 459, 243, 2009.
- Broecker, W. S.: A revised estimate for the radiocarbon age of North Atlantic Deep Water, *Journal of Geophysical Research: Oceans*, 84, 3218–3226, 1979.
- Burkholder, K. C. and Lozier, M. S.: Subtropical to subpolar pathways in the North Atlantic: Deductions from Lagrangian trajectories, *Journal of Geophysical Research: Oceans*, 116, C07 017, 2011.
- Chang, Y. S., Garraffo, Z. D., Peters, H., and Özgökmen, T. M.: Pathways of Nordic Overflows from climate model scale and eddy resolving simulations, *Ocean Modelling*, 29, 66–84, 2009.
- Clarke, R. A. and Gascard, J.-C.: The formation of Labrador Sea water. Part I: Large-scale processes, *Journal of Physical Oceanography*, 13, 1764–1778, 1983.
- Davis, X. J., Straneo, F., Kwon, Y. O., Kelly, K. A., and Toole, J. M.: Evolution and formation of North Atlantic Eighteen Degree Water in the Sargasso Sea from moored data, *Deep-Sea Research Part II: Topical Studies in Oceanography*, 91, 11–24, <https://doi.org/10.1016/j.dsr2.2013.02.024>, 2013.
- Dickson, B., Yashayaev, I., Meincke, J., Turrell, B., Dye, S., and Holfort, J.: Rapid freshening of the deep North Atlantic Ocean over the past four decades, *Nature*, 416, 832–837, 2002.
- Dirac, P. A. M.: A new notation for quantum mechanics, in: *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 35, pp. 416–418, Cambridge Univ Press, 1939.
- Douglass, E. M., Kwon, Y.-O., and Jayne, S. R.: A comparison of North Pacific and North Atlantic subtropical mode waters in a climatologically-forced model, *Deep Sea Research Part II: Topical Studies in Oceanography*, 91, 139–151, 2013.
- England, M. H. and Maier-Reimer, E.: Using chemical tracers to assess ocean models, *Reviews of Geophysics*, 39, 29–70, 2001.
- Errico, R. M.: What Is an Adjoint Model?, *Bulletin of the American Meteorological Society*, 78, 2577–2591, [https://doi.org/10.1175/1520-0477\(1997\)078<2577:WIAAM>2.0.CO;2](https://doi.org/10.1175/1520-0477(1997)078<2577:WIAAM>2.0.CO;2), 1997.
- Fiadeiro, M. E. and Veronis, G.: On weighted-mean schemes for the finite-difference approximation to the advection-diffusion equation, *Tellus*, 29, 512–522, 1977.
- Fichefet, T. and Maqueda, M. A. M.: Sensitivity of a global sea ice model to the treatment of ice thermodynamics and dynamics, *Journal of Geophysical Research*, 102, 12 609, <https://doi.org/10.1029/97JC00480>, 1997.
- Fischer, J. and Schott, F. A.: Labrador Sea Water tracked by profiling floats—From the boundary current into the open North Atlantic, *Journal of Physical Oceanography*, 32, 573–584, 2002.
- Forget, G., Maze, G., Buckley, M., and Marshall, J.: Estimated Seasonal Cycle of North Atlantic Eighteen Degree Water Volume, *Journal of Physical Oceanography*, 41, 269–286, <https://doi.org/10.1175/2010JPO4257.1>, 2011.
- Fratantoni, D. M., Kwon, Y.-O., and Hodges, B. A.: Direct observation of subtropical mode water circulation in the western North Atlantic Ocean, *Deep Sea Research Part II: Topical Studies in Oceanography*, 91, 35–56, <https://doi.org/10.1016/j.dsr2.2013.02.027>, 2013.
- Fukumori, I., Lee, T., Cheng, B., and Menemenlis, D.: The origin, pathway, and destination of Niño-3 water estimated by a simulated passive tracer and its adjoint, *Journal of Physical Oceanography*, 34, 582–604, 2004.
- Gao, S., Qu, T., and Fukumori, I.: Effects of mixing on the subduction of South Pacific waters identified by a simulated passive tracer and its adjoint, *Dynamics of Atmospheres and Oceans*, 51, 45–54, 2011.
- Gao, S., Qu, T., and Hu, D.: Origin and pathway of the Luzon Undercurrent identified by a simulated adjoint tracer, *Journal of Geophysical Research: Oceans*, 117, 2012.
- Gary, S. F., Lozier, M. S., Böning, C. W., and Biastoch, A.: Deciphering the pathways for the deep limb of the Meridional Overturning Circulation, *Deep Sea Research Part II: Topical Studies in Oceanography*, 58, 1781–1797, 2011.
- Gary, S. F., Lozier, M. S., Kwon, Y.-O., and Park, J. J.: The fate of North Atlantic subtropical mode water in the FLAME model, *Journal of Physical Oceanography*, 44, 1354–1371, 2014.
- Gerdes, R., Köberle, C., and Willebrand, J.: The influence of numerical advection schemes on the results of ocean general circulation models, *Climate Dynamics*, 5, 211–226, 1991.
- Giering, R. and Kaminski, T.: Recipes for adjoint code construction, *ACM Transactions on Mathematical Software (TOMS)*, 24, 437–474, 1998.

- Gresho, P. M. and Lee, R. L.: Don't suppress the wiggles—they're telling you something!, *Computers & Fluids*, 9, 223–253, 1981.
- Hansen, B. and Østerhus, S.: North Atlantic–Nordic seas exchanges, *Progress in oceanography*, 45, 109–208, 2000.
- 5 Hansen, B., Turrell, W. R., and Osterhus, S.: Decreasing overflow from the Nordic seas into the Atlantic Ocean through the Faroe Bank channel since 1950, *Nature*, 411, 927, 2001.
- Hill, C., Bugnion, V., Follows, M., and Marshall, J.: Evaluating carbon sequestration efficiency in an ocean circulation model
10 by adjoint sensitivity analysis, *Journal of Geophysical Research: Oceans*, 109, 2004.
- Hirst, A. C.: Determination of water component age in ocean models: Application to the fate of North Atlantic Deep Water, *Ocean Modelling*, 1, 81–94, 1999.
- 15 Iselin, C. O.: The influence of vertical and lateral turbulence on the characteristics of the waters at mid-depths, *Eos, Transactions American Geophysical Union*, 20, 414–417, 1939.
- Johnson, G. C.: Quantifying Antarctic bottom water and North Atlantic deep water volumes, *Journal of Geophysical Research: Oceans*, 113, 2008.
- 20 Jong, M. F. and Stur, L.: Strong winter cooling over the Irminger Sea in winter 2014–2015, exceptional deep convection, and the emergence of anomalously low SST, *Geophysical Research Letters*, 43, 7106–7113, 2016.
- 25 Joyce, T. M.: New perspectives on eighteen-degree water formation in the North Atlantic, *Journal of Oceanography*, 68, 45–52, <https://doi.org/DOI.10.1007/s10872-011-0029-0>, 2011.
- Khatiwala, S., Visbeck, M., and Cane, M. A.: Accelerated simulation of passive tracers in ocean circulation models, *Ocean Modelling*, 9, 51–69, 2005.
- 30 Klein, B. and Hogg, N.: On the variability of 18 degree water formation as observed from moored instruments at 55°W, *Deep Sea Research Part I: Oceanographic Research Papers*, 43, 1777–1806, 1996.
- 35 Kwon, Y. O. and Riser, S. C.: North Atlantic subtropical mode water: A history of ocean-atmosphere interaction 1961–2000, *Geophysical Research Letters*, 31, L19 307, <https://doi.org/10.1029/2004GL021116>, 2004.
- Kwon, Y. O. and Riser, S. C.: General circulation of the
40 western subtropical North Atlantic observed using profiling floats, *Journal of Geophysical Research C: Oceans*, 110, 1–22, <https://doi.org/10.1029/2005JC002909>, 2005.
- Kwon, Y.-O., Park, J.-J., Gary, S. F., and Lozier, M. S.: Year-to-year reoutcropping of eighteen degree water in an eddy-resolving ocean simulation, *Journal of Physical Oceanography*, 45, 1189–1204, 2015.
- Large, W. G. and Yeager, S. G.: Diurnal to decadal global forcing for ocean and sea-ice models: the data sets and flux climatologies, *NCAR Technical Note, NCAR/TN-460+STR*, 2004.
- 50 Lavender, K. L., Davis, R. E., and Owens, W. B.: Mid-depth recirculation observed in the interior Labrador and Irminger seas by direct velocity measurements, *Nature*, 407, 66, 2000.
- Leaman, K. D. and Harris, J. E.: On the average absolute transport of the deep western boundary currents east of Abaco Island, the Bahamas, *Journal of Physical Oceanography*, 20, 467–475, 1990.
- LeBel, D. A., Smethie, W. M., Rhein, M., Kieke, D., Fine, R. A., Bullister, J. L., Min, D.-H., Roether, W., Weiss, R. F., Andrié, C., et al.: The formation rate of North Atlantic Deep Water and Eighteen Degree Water calculated from CFC-11 inventories observed during WOCE, *Deep Sea Research Part I: Oceanographic Research Papers*, 55, 891–910, 2008.
- Lévy, M., Estublier, A., and Madec, G.: Choice of an advection scheme for biogeochemical models, *Geophysical Research Letters*, 28, 3725–3728, 2001.
- Lozier, M. S. and Stewart, N. M.: On the temporally varying northward penetration of Mediterranean Overflow Water and eastward penetration of Labrador Sea Water, *Journal of Physical Oceanography*, 38, 2097–2103, 2008.
- Luyten, J. R., Pedlosky, J., and Stommel, H.: The Ventilated Thermocline, *Journal of Physical Oceanography*, 13, 292–309, [https://doi.org/10.1175/1520-0485\(1983\)013<0292:TVT>2.0.CO;2](https://doi.org/10.1175/1520-0485(1983)013<0292:TVT>2.0.CO;2), 1983.
- Macrander, A., Send, U., Valdimarsson, H., Jónsson, S., and Käse, R. H.: Interannual changes in the overflow from the Nordic Seas into the Atlantic Ocean through Denmark Strait, *Geophysical Research Letters*, 32, 2005.
- Madec, G.: the NEMO team: Nemo ocean engine–Version 3.4, Note du Pôle de modélisation. Institut Pierre-Simon Laplace (IPSL), France, 2012.
- Madec, G. and Imbard, M.: A global ocean mesh to overcome the North Pole singularity, *Climate Dynamics*, 12, 381–388, <https://doi.org/10.1007/BF00211684>, 1996.
- Marotzke, J., Giering, R., Zhang, K. Q., Stammer, D., Hill, C., and Lee, T.: Construction of the adjoint MIT ocean general circulation model and application to Atlantic heat transport sensitivity, *Journal of Geophysical Research: Oceans*, 104, 29 529–29 547, 1999.
- Marshall, J., Ferrari, R., Forget, G., Maze, G., Andersson, A., Bates, N., Dewar, W., Doney, S., Fratantoni, D., Joyce, T., et al.: The CLIMODE field campaign: Observing the cycle of convection and restratification over the Gulf Stream, *Bulletin of the American Meteorological Society*, 90, 1337–1350, 2009.
- McCartney, M. S. and Talley, L. D.: The subpolar mode water of the North Atlantic Ocean, *Journal of Physical Oceanography*, 12, 1169–1188, 1982.
- 95 Molinari, R. L., Fine, R. A., and Johns, E.: The deep western boundary current in the tropical North Atlantic Ocean, *Deep Sea Research Part A. Oceanographic Research Papers*, 39, 1967–1984, 1992.
- NEMO Consortium: NEMO Reference configurations inputs, <https://doi.org/10.5281/zenodo.1471702>, 2013.
- Orsi, A., Johnson, G., and Bullister, J.: Circulation, mixing, and production of Antarctic Bottom Water, *Progress in Oceanography*, 43, 55–109, 1999.
- Owen, A.: Artificial diffusion in the numerical modelling of the advective transport of salinity, *Applied mathematical modelling*, 8, 116–120, 1984.
- Park, S. K. and Xu, L.: Data assimilation for atmospheric, oceanic and hydrologic applications, vol. 2, Springer Science & Business Media, 2013.
- 110 Pickart, R. S., Spall, M. A., Ribergaard, M. H., Moore, G., and Milliff, R. F.: Deep convection in the Irminger Sea forced by the Greenland tip jet, *Nature*, 424, 152, 2003a.
- Pickart, R. S., Straneo, F., and Moore, G.: Is Labrador Sea water formed in the Irminger basin?, *Deep Sea Research Part I: Oceanographic Research Papers*, 50, 23–52, 2003b.
- 115 Qu, T., Gao, S., Fukumori, I., Fine, R. A., and Lindstrom, E. J.: Origin and pathway of equatorial 13°C water in the Pacific identified

- by a simulated passive tracer and its adjoint, *Journal of Physical Oceanography*, 39, 1836–1853, 2009.
- Qu, T., Gao, S., and Fukumori, I.: Formation of salinity maximum water and its contribution to the overturning circulation in the North Atlantic as revealed by a global general circulation model, *Journal of Geophysical Research: Oceans*, 118, 1982–1994, 2013.
- Quadfasel, D. and Käse, R. H.: Present-day manifestation of the Nordic Seas overflows, in: *Natural Gas Hydrates-Occurrence, Distribution and Detection* (ed. CK Paull and WP Dillon) American Geophysical Union, vol. 173, pp. 75–89, American Geophysical Union, 2007.
- Reid, J. L. and Lynn, R. J.: On the influence of the Norwegian-Greenland and Weddell seas upon the bottom waters of the Indian and Pacific oceans, *Deep Sea Research*, 18, 1063–1088, 1971.
- Rhein, M., Fischer, J., Smethie, W., Smythe-Wright, D., Weiss, R., Mertens, C., Min, D.-H., Fleischmann, U., and Putzka, A.: Labrador Sea Water: Pathways, CFC inventory, and formation rates, *Journal of Physical Oceanography*, 32, 648–665, 2002.
- Schuert, E. A.: Turbulent diffusion in the intermediate waters of the North Pacific Ocean, *Journal of Geophysical Research*, 75, 673–682, 1970.
- Sévellec, F. and Fedorov, A. V.: Stability of the Atlantic meridional overturning circulation and stratification in a zonally averaged ocean model: Effects of freshwater flux, Southern Ocean winds, and diapycnal diffusion, *Deep Sea Research Part II: Topical Studies in Oceanography*, 58, 1927–1943, 2011.
- Siedler, G., Kuhl, A., and Zenk, W.: The Madeira mode water, *Journal of Physical Oceanography*, 17, 1561–1570, 1987.
- Smethie Jr, W. M. and Fine, R. A.: Rates of North Atlantic Deep Water formation calculated from chlorofluorocarbon inventories, *Deep Sea Research Part I: Oceanographic Research Papers*, 48, 189–215, 2001.
- Speer, K. and Tziperman, E.: Rates of water mass formation in the North Atlantic Ocean, *Journal of Physical Oceanography*, 22, 93–104, 1992.
- Stephenson, D.: ds4g15/PT_TAM_ORCA2 v1.0.3, <https://doi.org/10.5281/zenodo.3635919>, 2019a.
- Stephenson, D.: ds4g15/passive_experiments v1.0.1, <https://doi.org/10.5281/zenodo.3471041>, 2019b.
- Straneo, F., Pickart, R. S., and Lavender, K.: Spreading of Labrador sea water: an advective-diffusive study based on Lagrangian data, *Deep Sea Research Part I: Oceanographic Research Papers*, 50, 701–719, 2003.
- Swift, J. H., Aagaard, K., and Malmberg, S.-A.: The contribution of the Denmark Strait overflow to the deep North Atlantic, *Deep Sea Research Part A. Oceanographic Research Papers*, 27, 29–42, 1980.
- Talley, L. and McCartney, M.: Distribution and circulation of Labrador Sea water, *Journal of Physical Oceanography*, 12, 1189–1205, 1982.
- Thomas, L. N.: Destruction of potential vorticity by winds, *Journal of Physical Oceanography*, 35, 2457–2466, 2005.
- Thomson, C. W.: *The Voyage of the “Challenger”: The Atlantic; a Preliminary Account of the General Results of the Exploring Voyage of HMS “Challenger” During the Year 1873 and the Early Part of the Year 1876*, vol. 1, Macmillan and Company, 1877.
- Treguier, A. M., Barnier, B., Albert, A., Deshayes, J., leSommer, J., Lique, C., Molines, J.-M., Penduff, T., and Talandier, C.: ORCA12, a global ocean-ice model at $\frac{1}{12}^\circ$: successes, shortcomings and their impact on ocean forecasting, in: *EGU General Assembly Conference Abstracts*, vol. 19, p. 17950, 2017.
- Vidard, A., Bouttier, P. A., and Vigilant, F.: NEMOTAM: Tangent and adjoint models for the ocean modelling platform NEMO, *Geoscientific Model Development*, 8, 1245–1257, <https://doi.org/10.5194/gmd-8-1245-2015>, 2015.
- Warren, B. A.: Insensitivity of subtropical mode water characteristics to meteorological fluctuations, *Deep Sea Research*, 19, 1–19, 1972.
- Welander, P.: An advective model of the ocean thermocline, *Tellus A: Dynamic Meteorology and Oceanography*, 11, 309–318, 1959.
- Welander, P.: The thermocline problem, *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 270, 415–421, 1971.
- Worthington, L. and Wright, W. R.: North Atlantic Ocean atlas of potential temperature and salinity in the deep water including temperature, salinity and oxygen profiles from the Erika Dan Cruise of 1962, Tech. rep., Woods Hole Oceanographic Institution, 1970.
- Worthington, L. V.: The 18° water in the Sargasso Sea, *Deep Sea Research*, 5, 297–305, 1959.
- Worthington, L. V.: Negative oceanic heat flux as a cause of water-mass formation, *Journal of Physical Oceanography*, 2, 205–211, 1972.
- Xie, P. and Vallis, G. K.: The passive and active nature of ocean heat uptake in idealized climate change experiments, *Climate Dynamics*, 38, 667–684, 2012.
- Zalesak, S. T.: Fully multidimensional flux-corrected transport algorithms for fluids, *Journal of Computational Physics*, 31, 335–362, 1979.