Duke et al. Reply on RC1: 'Comment on egusphere-2023-870', Marine Fourrier

General comments

This paper by Duke and co-authors describes a neural network approach to interpolate observations into a gridded pCO2 data product over the period 1998-2019 at 1/12° resolution specifically in the Northeast Pacific. After presenting the specific circulation and hydrographic features of their area, the authors detail the methodological choices they made in developing the NN and advise the reader on specific training steps to improve performance. They then go on to discuss the wider scientific applications of such a product.

Overall, the paper is very well written, nicely structured and easy to read. The figures could be larger to improve readability, but figures and legends are still clear. The level of detail in the methodological and NN development sections is much appreciated, as it is rare to find such detail in oceanographic research papers. In particular, the decision not to normalise all your inputs and to use dynamic provinces is very interesting. Also all the developments you have done on the internal division data ratio, which could be applied to many other uses.

Thank you Marine Fourrier for your time and careful consideration in the assessment of our manuscript. We are glad that the detailed description of the NN approach and transferability of the method development were received well. Below we present a point-by-point response to comments. Our responses are in blue, with manuscript text in quotations and added/revised text italicised.

A few specific comments remain, as detailed below:

Specific comments

It would be nice to further highlight the potential uses of your data product and/or your method outside the traditional straightforward pCO2/FCO2 observing community (i.e. modellers, determination of climate indices, ...).

We have added the following paragraph to section 5 line 483:

"The regional, high-resolution pCO2 product created here could serves as a valuable baseline for regional models (e.g., Pilcher et al., 2018; Hauri et al., 2020). The pCO2 product, and associated air-sea CO2 flux estimates, offers continuous coverage in sparsely sampled regions informed by patterns in well sampled neighbouring waters. The product could be used to aid in model evaluation, use in data assimilation, constrain initial conditions, enhance carbon flux process understanding, and improve regional climate change projections."

As well as additional line to section 4.2 line 351:

"*This relationship supports work showing that the SSH anomaly is an important climate index for the region (Cummins et al., 2005; Di Lorenzo et al., 2008).*"

Have you considered comparing your FCO2 with other products: e.g. SeaFlux (Fay et al., 2021)?

Compared to the averaged SeaFlux FCO_2 product (across 6 different observation-based products, with 5 different wind speed products) our estimated FCO_2 looks similar outside the higher range of outgassing values (positive fluxes > 5 mol m^{-2} yr^{-1}) which exist within the Alaskan Gyre (Figure 5b & 6a).

This comparison has been added in text to section 4 line 290:

"Our ANN-NEP calculated fluxes compare well to air-sea CO_2 fluxes averaged across six unique courser resolution, global observation-based pCO2 products, each using five different wind speed products (r^2=0.81; Fay et al. 2021). However, our work suggests that the global product ensemble may underestimate the outgassing signal from the subpolar Alaskan Gyre (Figure 5b; supplemental Figure 7)."

and included as a new supplemental figure:

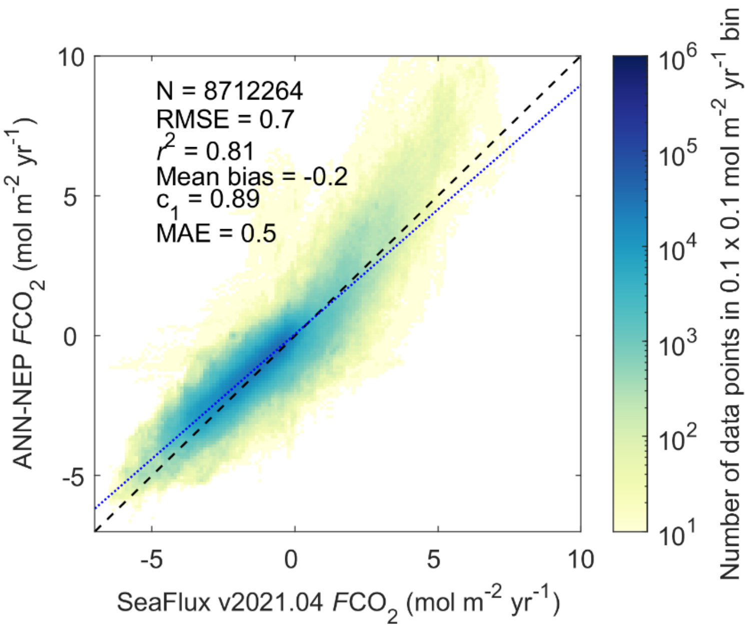

Figure S7 Property to property plot of air-sea CO_2 flux density values calculated from ANN-NEP and from SeaFlux v2021.04 (Fay et al., 2021). The SeaFlux estimates have been interpolated to the 1/12°x1/12° grid of this study. Number of overlapping grid cells withing the study area (N), root mean squared error (RMSE), coefficient of determination (r^2), mean absolute error (MAE), mean bias (calculated as the mean residual), and the slope of the linear regression (c_1). The observed linear relationship is represented by the dotted blue line. Data is binned into 0.1 by 0.1 mol m^{-2} yr^{-1} bins. The dashed black line represents a perfect fit of slope (c_1) = 1 and intercept = 0. Colorbar shows data density on a log scale.

End of section 2.1 and Table 1: there seems to be some confusion here about what you have used for what. You mention the xCO2 data produced by NOAA, but in Table 1 you cite Lanschutzer? As in the next sentence you mention the pCO2 climatology by the same authors, restructuring the sentences if this is not a confusion would be very useful.

We have added to the description at the end of section 2.1 line 112:

"Atmospheric pCO_2 in µatm was downloaded from Landschützer et al. (2020), derived from the National Oceanic and Atmospheric Administration Earth System Research Global Monitoring Laboratory (https://gml.noaa.gov/ccgg/globalview/) atmospheric mole fraction of CO_2 (χCO_2) and SST (Reynolds et al. 2002) as well as sea level pressure (Kalnay et al. 1996) following Dickson et al. (2007)."

This text reflects consistency with the table and alleviates confusion later when in text when χCO_2 should just be referred to as atmospheric pCO_2 in µatm. This was a mistake in the submitted manuscript. Thanks for catching the error.

Table 1: for SST you kept the 1/20° resolution and did not aggregate to 1/12°, unlike the others?

We do aggregate to 1/12°. Table 1 has been corrected to reflect this. Thanks for catching the mistake.

Figure 2c: Is your validation data set representative enough? It does not cover your whole range?

Keeping the independent withheld data randomly selected, yet also representative of the full domain was difficult. The end ranges of the training data (high and low pCO_2 values) are rather scarce in the observational data set, making them important for training. Striking the balance between withholding representative data yet showing critical training data can be tricky. Currently, we are unaware of any community-based recommendations that deal with this issue. However, our final estimate of product uncertainty is not based on the training data but rather the independent withheld data, unlike may other observation-based products.

End of section 2.5: You average the outputs of the 10 NN. You do this directly, but have you tried to give the median +/std as this is also useful information.

The ensemble median looked very similar to the mean (property to property plot below).

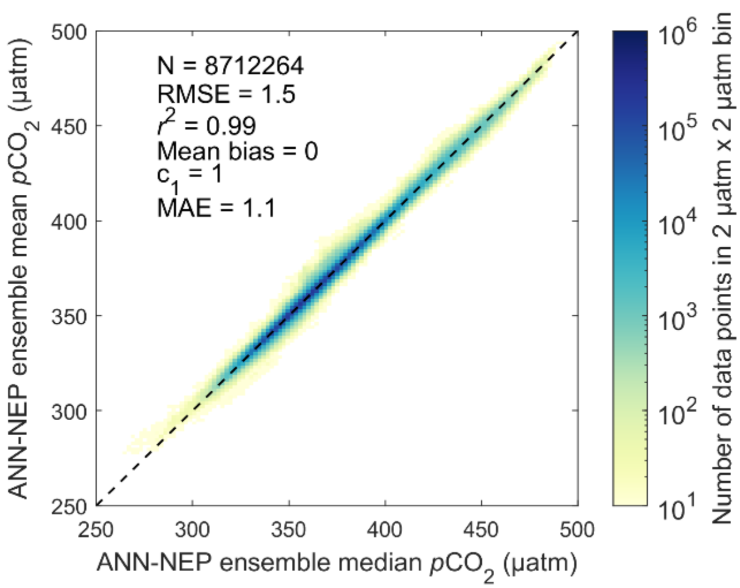

We have added the following to section 3.1 line 210:

"The ensemble median was nearly equivalent to the ensemble mean (r^2 = 0.99; not shown)."

Line 210: "12.9 +:-1.1 μatm" : I'm not sure where this number comes from.

Added in text citation to the plotted data for this value to revised supplemental Figure S4, now including a second panel (b) with each ensemble member plotted against the consistent independent withheld data across all ensemble members.

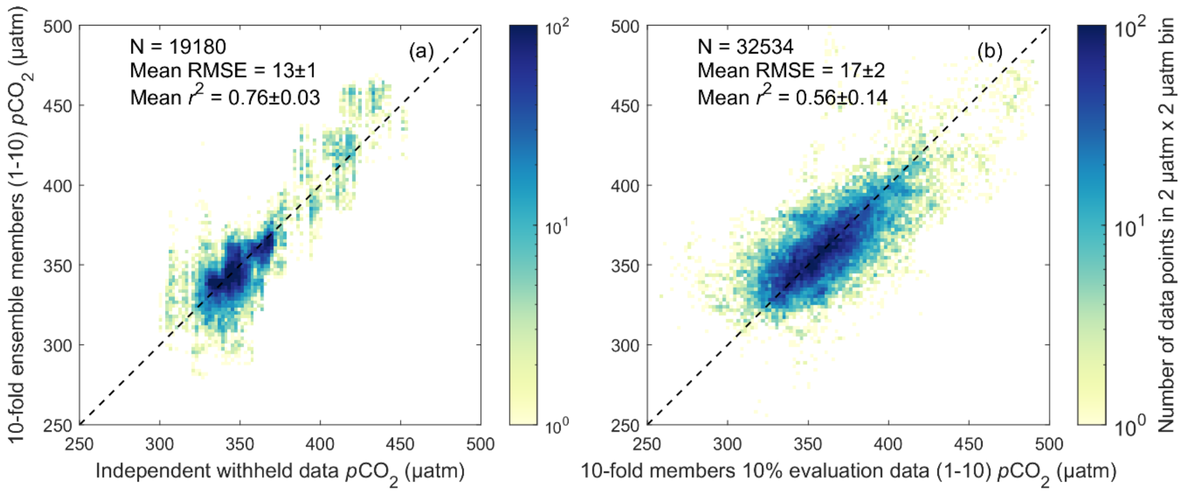

Figure S4 10-fold cross-evaluation (Section 2.4) individual ensemble member estimated pCO_2 against the (a) independent withheld data, and (b) 10% 10-fold evaluation data specific to that ensemble member. Mean root mean squared error (RMSE) and coefficient of determination (r^2) are across all individual ensemble members. Data is binned into 2 µatm by 2 µatm bins. The dashed black line represents a perfect fit of slope (c_1) = 1 and intercept = 0.

Line 332: flux densities as high as 3.6, but your figures end at a maximum of 3?

Figure 7a groups data into seasonal (3 month) bins. The 3.6 value refers to just one month being January of 2000. Removed in text figure citation to avoid confusion.

Lines 394-395: provide some insight into how.

We have added example approaches to section 4.3 line 394 to help disentangle these distinct influences:

"Unravelling the individual influence of these interconnected drivers (i.e., marine heatwaves, sub-decadal variability, and long-term trend) is not possible with this product *alone* but does prompt future inquiry *in combination with regional models and emerging climate analysis tools (e.g., Chapman et al. 2022)."*

In this section you compare an atmospheric xCO2 with an increase in oceanic pCO2. Can't you convert the atmospheric xCO2 to pCO2 to compare things in the same ranges/units?

pCO2ATM = [PT - (RH/100) × PH2O] × xCO2ATM

where PH2O is the water vapour pressure at atmospheric temperature for xCO2ATM (in atm) calculated according to Dickson et al. (2007), RH is the relative humidity (in %) and PT is the total atmospheric pressure (in atm). If some of these are missing, you can get them from products (SeaFlux mentioned above). Otherwise you are comparing things that are not directly comparable.

Thanks for catching this error. It occurred in both in the legend and text. The values we refer to are atmospheric *p*CO_2 in µatm. The legend in figure 9 and text in section 4.4 have been corrected.

Line 437: detail how you got the vertical velocities with Ekman pumping

We have revised our text so it now includes in text citation in section 4.4 line 437 to MATLAB Climate Toolbox *ekman* function:

"We find a strong spatial correlation between the trend in ΔpCO2 and the calculated average vertical velocity associated with Ekman pumping (r^2 = 0.64, $p < 0.01$; Figure 11b). Ekman pumping was calculated from monthly, 1/4° spatial resolution Cross-Calibrated Multiplatform zonal and meridional ocean surface wind speeds (Mears et al. 2019) interpolated to 1/12°, using the MATLAB Climate Toolbox ekman function (Greene et al., 2017, 2019; Kessler et al., 2002)."

Technical corrections

Abstract: For readers not familiar with your particular area, it would be better to introduce the fact that your area is a sink earlier.

Revised line 14 in the abstract to include this regional context:

"Here we use a neural network approach to interpolate sparse observations, creating a monthly gridded seawater partial pressure of CO2 (pCO2) data product from January 1998 to December 2019, at 1/12°x1/12° spatial resolution, in the Northeast Pacific open ocean*, considered a net sink region."*

Line 58: "The estimated long-term trend in surface ocean pCO2 appears to be increasing at less than the atmospheric rate." Less what? rate, amount of increase? Rewrite to clarify: "at a slower rate" or "less than the atmospheric rate of increase".

Revised to:

"The estimated long-term trend in surface ocean pCO2 appears to be increasing at less than the atmospheric rate *of increase (Franco et al., 2021)."*

Line 119: fCO2 has been converted to pCO2: how, give equation.

Revised to include in text citation to text added to the supplementary.

Supplementary Text TXX

"The reported fCO2 estimates were converted to pCO2 using the equation SXX (Körtzinger, 1999):

$$pCO_2 = fCO_2 \times exp[P_{atm}^{surf} \frac{B+(2\delta)}{RT}], \qquad (SXX)$$

where P_{atm}^{surf} is the total atmospheric surface pressure, B and δ are viral coefficients (Weiss, 1974), R is the gas constant and T is the absolute temperature. National Centers for Environmental Prediction (NCEP) monthly mean sea level pressure was used for P_{atm}^{surf} (Kalnay et al., 1996)."

Figure 3: Add some metrics to the graphs (from the text)

Revised figure to include text with mean absolute error values for both products to compare more easily.

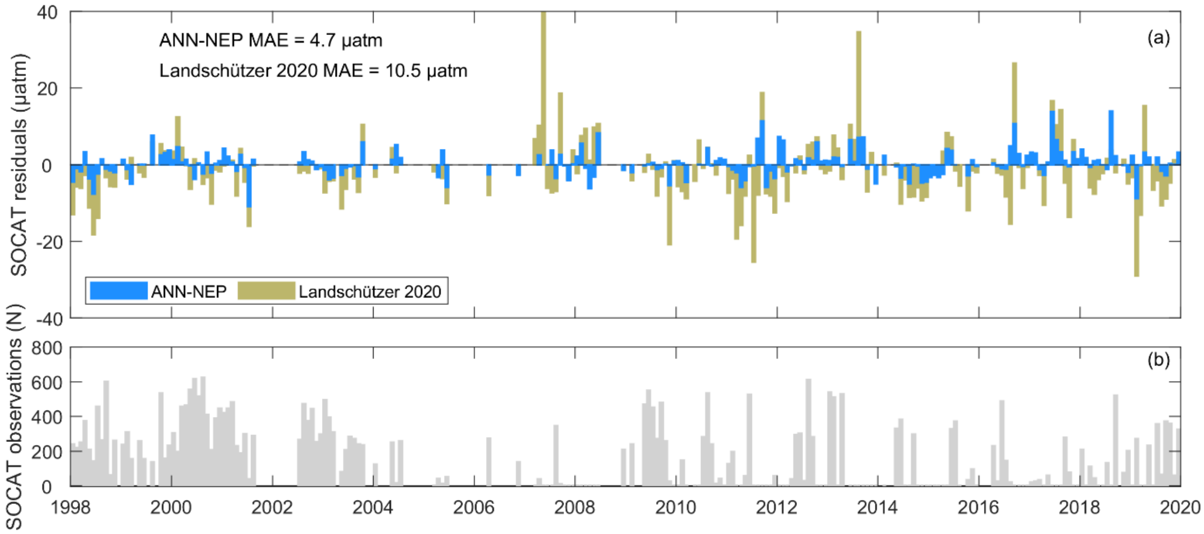

Line 150: "changing the internal division ratio", what did you end up using (see where you give more details).

Revised text to include in-text citation to later section.

"… changing the internal data division ratio (94:6; see Section 3.4 below)."

Line 424: "could be partially tied to if" missing word/rewording

Thanks for catching an error in the description. Revised text to:

"Date combinations resulting in trends exceeding the atmospheric increase could be partly attributed to start and end dates coinciding with periods of weak and strong Alaskan Gyre upwelling, respectively. These upwelling modes induce negative and positive pCO2 anomalies, which further amplify the observed trend."

Figure 7: the black line is not really black.

Revised figure so line is black.

Supplementary Figure 1 a&b: useful to superimpose bar charts of the whole dataset behind to further demonstrate how your subsampled data is representative.

Revised figure to superimpose bar charts of the training dataset behind.

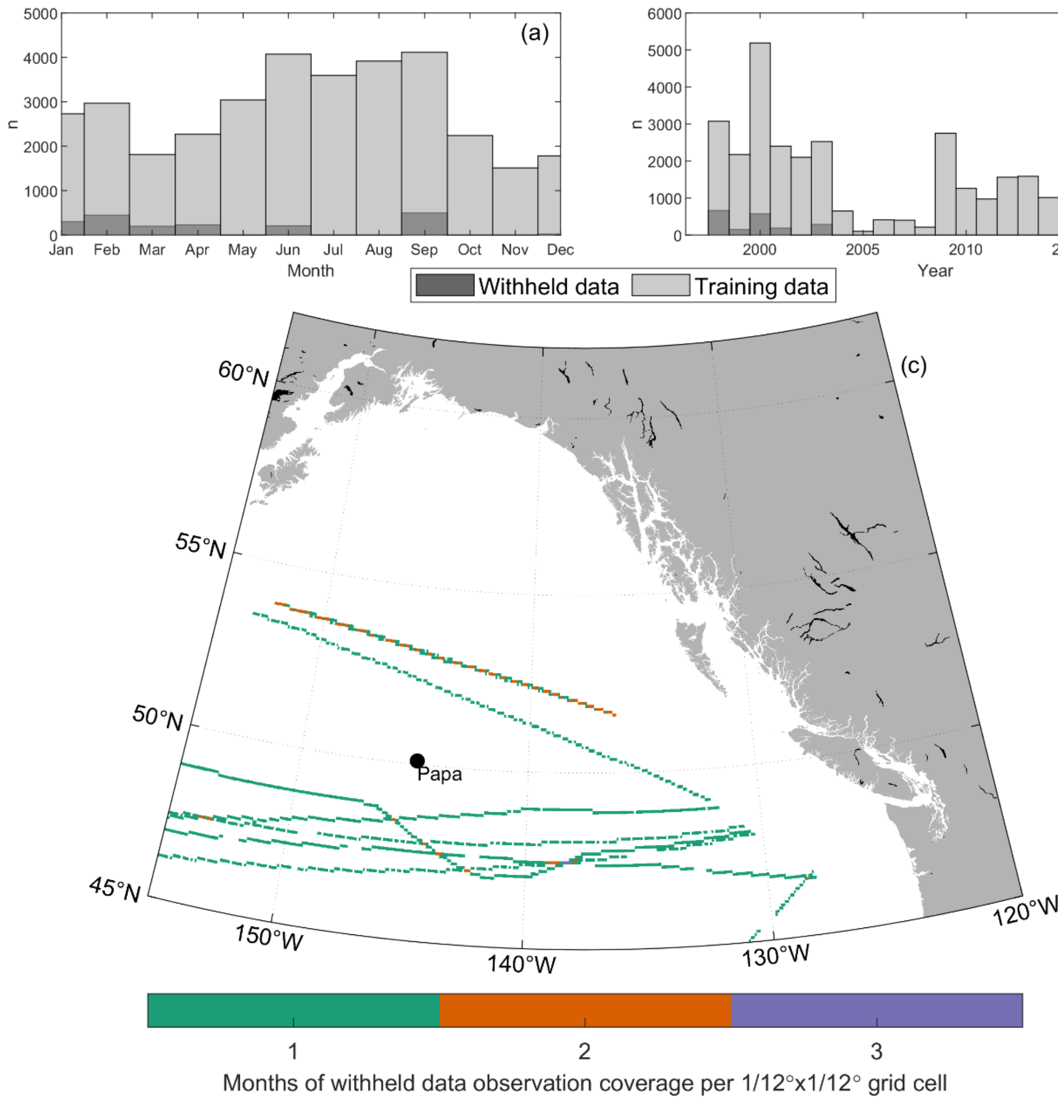
(a)
n
5000
4000
3000
2000
1000
0
Jan Feb Mar Apr May Jun Jul Aug Sep Oct Nov Dec
Month
(b)
n
6000
5000
4000
3000
2000
1000
0
2000 2005 2010 2015 2020
Year
Withheld data
Training data
(c)
60°N
55°N
50°N
45°N
Papa
150°W
140°W
130°W
120°W
1
2
3
Months of withheld data observation coverage per 1/12°x1/12° grid cell