Response to Referee 2

Albertus J. Smit[1,2] and Neville Sweijd[3]

[1]Department of Biodiversity and Conservation Biology, University of the Western Cape, Bellville, Cape Town 7535, South Africa

[2]South African Environmental Observation Network, Elwandle Coastal Node, Gqeberha 6031, South Africa

[3]School for Climate Studies, Stellenbosch University, Stellenbosch 7599, South Africa

28th July 2026

About this document

This document accompanies our author comment on the manuscript *A segmented-breakpoint sea-surface-temperature upwelling index for the Benguela Upwelling System*, in open discussion at *Ocean Science*. We respond to the general assessment (Section 1), gather the concerns that our reply to Referee 1 already answers (Section 2), address the physical justification of the method (Section 3), correct the description of the wind index (Section 4), give the further Methods passages the referee asked us to rewrite (Section 5), and reply to each detailed question (Section 6). Line numbers refer to the review PDF.

Three kinds of text appear below, and each item is set in the style it describes:

- Explanatory text, addressed to the referee, is in the normal body font, as in this sentence.
- Text proposed for the revised manuscript is in the same font, set in the red of the document style. Everything marked this way is meant to be read as manuscript replacement text rather than as a reply, and we will place it into the manuscript unchanged.
- AS SUBMITTED. Text quoted from the manuscript as submitted is set small and grey, and carries this label.

The two referee reports. The two reports overlap substantially, and Referee 1 asked us for a full rewrite of the statistical section, which we supplied in our reply to that referee. Reproducing that text here would double the length of an already long document. Section 2 therefore lists the concerns this referee raises that are already answered there, and says where. Everything this referee raises that is really new is answered in full in this document.

1 Response to the general assessment

We thank the referee for a report that took evident time and that has changed the manuscript more than any other single document we have received. We accept the two main criticisms; one of them identifies a clear error.

The language

The referee opens by apologising for not being a native English speaker. We would like to turn that around. A reader who has to read a sentence twice has been failed by the sentence, and a reader who reports that the language is *never exactly to the point* has identified the actual problem. Referee 1 made the same criticism in different words, and the two reports agreeing on this point independently settles the matter. So, the fault is ours.

We accept that the manuscript needs a full rewrite for tone and clarity rather than local repairs, and that the rewrite must be done by us line by line. Our reply to Referee 1 contains seven Methods paragraphs already rewritten to that standard, and they show the voice we intend for the whole manuscript. The referee also notes that section titles announce their subject obliquely... yes. So, the revised titles are given in Section 6.

The missing physical justification

The manuscript demonstrates that a two-segment description fits the data and that its acceptance is statistically defensible, then leaves unasked the prior question of why a two-segment shape should describe an upwelling profile at all. Statistical adequacy is not physical justification. Section 3 addresses this, with the suggested starting point of de Szoeke and Richman (1984). Thank you for the pointer.

The wind index

The referee identifies a sentence that is, as written, clearly wrong, and asks what our scripts actually do. So, we checked the code before writing this reply. It turns out that the implementation is correct and the sentence describing it is not. Section 4 gives the details, the corrected text, and the place in the manuscript where the correct description already appears, which is how the inconsistency escaped us.

Authorship and tooling

The error at L. 184 is ours regardless of how the sentence came to be written, and the responsibility for every sentence in the revision is ours. We will not submit the revision until each of us has again read every line of it against the code and the data.

2 Concerns already addressed in our reply to Referee 1

Both referees raised the concerns below. Each is answered in our reply to Referee 1 (already submitted), so we answer and flag where this referee's version of the concern goes beyond Referee 1's.

- **The statistical tests cannot be repeated from the text, and no formulae are given.** Seven Methods passages are rewritten in full in Section 2 of that reply, with the submitted text quoted beneath each for comparison. The rewrites state the acceptance criteria with their formulae, the effective-sample-size correction, and the event-grouping rule. *Additional here:* this referee asks for the correspondence between the cited mathematical objects and the quantities used in this paper, so we will make this explicit at each step.
- **Benjamini–Hochberg is not defined, and it is unclear what controls what.** Answered in Section 2.7 of our reply to Referee 1, which explains the false-discovery-rate procedure before mentioning Benjamini-Hochberg. See also L. 272 in Section 6, below.
- **The second half of the abstract cannot be understood without reading the paper.** Referee 1 made the same objection. A fully revised abstract is given in Section 4 of that reply. But this referee adds a request about the opening sentence, causing us to revise the already revised abstract; see Section 6.
- **Table font sizes vary and some tables are unreadable without zooming.** Answered in Section 4 of the Referee 1 reply. All Supplement tables will be at a single size, with the two widest rotated to landscape.
- **Figure 2 has overlapping headline and colour definition, and stray characters in the panel titles.** Answered in Section 4 of the Referee 1 reply, together with the overplotting at Cape Frio. Figure 4 has the same overlap, and so they will be corrected together.
- **Label and caption sizes should be checked throughout.** Accepted in the Referee 1 reply for the Supplement tables. This will apply to every figure in the main text and Supplement, with a one minimum type size applied across the set.
- **The AR(1) correction and the effective sample size need explaining.** Completely rewritten in Section 2.4 of the Referee 1 reply, including why residual autocorrelation is more important here

than the choice of the information criterion.

- **A walk-through of the method would help more than added prose.** Referee 1 asked for this and it was a great suggestion! Section 3 of that reply contains the proposed five-panel figure. It answers several of this referee's questions about the workflow; it is worth reading it alongside this document.

3 The physical basis of the two-segment model

Why should an upwelling profile be described by two linear segments meeting at a breakpoint?

Is it granted that rejected profiles do not correspond to upwelling?

And why should a piecewise-linear regression be preferred to methods that integrate over areas bounded by isotherms?

Let's look at each of these questions.

3.1 Why two segments

Thinking about it like this, our simple justification is that the two-segment form is a shape "model" rather than a solution of a mechanistic equation. Maybe we should have said so and then justified the shape on physical grounds. We think:

Coastal upwelling produces a surface expression with a characteristic structure, namely recently upwelled water sititng in a band against the coast, separated from warmer offshore surface water by a frontal transition. Within the upwelled band the cross-shore temperature gradient is set by the age and mixing history of the water since it surfaced, and it is comparatively steep. Seaward of the front the gradient reverts to the weak background gradient of the open ocean. A profile with a steep inshore segment, a weak offshore segment, and a transition between them is therefore the simplest description (form/shape?) that has the physical distinction the index is meant to measure, namely where the upwelled band ends.

The theory support for a frontal transition like this is older than our method. de Szoeke and Richman (1984) examine a wind-driven mixed layer with strong horizontal gradients and show that the interaction of wind mixing with an established cross-shore density gradient focusses the gradient into a front at a predictable location. It does not diffuse it. This gives a physical reason to expect a break of slope in the surface temperature field, and it also predicts that the break should sit at a distance that varies with wind history and stratification. It is this variation our fitted breakpoint measures.

We propose to add the following to the Discussion.

The two-segment form is a description of shape and not a solution of the governing equations, so its justification is that the shape reflects the physical distinction measured by the index. Coastal upwelling places recently surfaced water in a band against the coast, separated from offshore surface water by a frontal transition, and the cross-shore temperature gradient is steeper within the band than beyond it. de Szoeke and Richman (1984) show that a wind-driven mixed layer with a strong horizontal gradient enhances rather than diffuses that transition, and that its position depends on wind history and stratification. A fitted breakpoint therefore estimates a quantity that theory predicts should vary from day to day. We do not claim that the two-segment fit recovers the frontal position in a dynamical sense, and establishing that correspondence would require hydrographic sections across the fitted breakpoint.

3.2 What a rejected profile means

The referee asks whether i) it is granted that rejected profiles do not correspond to upwelling, and ii) states that the information lost through rejection cannot be controlled or quantified. The first question is fair and the answer is no, it is not granted. We would like to qualify the second: we do quantify the loss, but admittedly perhaps not where a reader would look for it.

Rejection rates are reported per UC in Results Section 3.1, and they vary by more than an order of magnitude, namely 8.9 % at Cape Frio, 30.5 % at Walvis Bay, 0.7 % at Lüderitz, 3.5 % at Cape Columbine, and

23.4 % at Cape Peninsula when DJF and JJA are pooled in the OSTIA representative-profile population. We also characterise what the rejected profiles look like. At four of the five UCs they retain a gradual cross-shore gradient of ca. +0.004 to +0.011 °C km^{-1}, so the criterion is mostly screening profiles that warm offshore monotonically without an inflexion rather than profiles that are flat. Cape Peninsula is the exception, where rejected profiles have near-zero slope, so a thermally featureless field.

This quantification supports a more limited claim than we make in the manuscript. A rejection shows that the profile lacks a two-segment structure on that day at that transect, but it does not establish the absence of upwelling. A profile whose upwelled band extends past the outer contour, or whose front is too diffuse to be detected as a break of slope, will be rejected while upwelling is occurring. We will state this limitation explicitly, and we will move the rejection statistics from their present position into the same passage so that both ideas appear together.

3.3 Profile method vs. an area method

An aside... We have spent considerable effort on outside this paper on the question the referee asks here. In a parallel line of work we have built and tested an areal alternative on the same coast. There we delineate daily two-dimensional upwelling polygons from the SST field, link them into events through the season, and resolves them into recurrent cell components using mixture modelling over an event feature space. We use the same four products and the same 2015 to 2025 window as the present paper. Three successive versions of it exist, and the most recent (`uc2d_v3`) is in preparation. So we hold our view about the two types of detectors from this response, but it is interesitng nevertheless.

We raise this only so that the referee knows the question has been taken seriously, and we want to be clear that we are not offering it as justification. We cannot use "evidence" held in an unpublished analysis to support a claim in this manuscript. So, our reply below is therefore based on what is in the present paper alone. Should that paper be published first, the cross-reference will follow.

Back to our response... The referee is right that isotherm-area and closed-isotherm methods integrate over a two-dimensional field and discard less information than a one-dimensional profile fit, and that we assert the value of our approach without critisising alternatives. We do not think the breakpoint method is better (apologies if we implied it). The two answer different questions.

In the Benguela the areal approach is the older of the two, and perhaps it is the defining one. Lutjeharms and Meeuwis (1987) defined an upwelling cell wherever the isotherms close onto the coastline, so that the cool patch is bounded on its landward side by the coast itself. A cell in that sense is a polygon closing on the coast, and it is identified by its outline. Hagen et al. (2001) later made the same idea quantitative as the area enclosed between the coast and the 13 °C isotherm. Both definitions *also* work from the shape of the feature (i.e., a polygon is analogous to the breakpoint contrast measured along a line), and the eight-cell Benguela list that the literature still uses descends from that lineage.

Our method belongs to the other lineage, namely the fixed-distance contrasts that reduce the cross-shore field to a single difference between a coastal value and an offshore reference. These methods are prone to the defect our manuscript is about. So, maybe we should have placed our contribution inside this "contrast lineage". We will be sure to update the Discussion to this end. But the two lineages do meet in our own method in this paper... our breakpoint estimates where the outline would cross the transect, so an outline method and a profile method are measuring the same edge by different means.

An area method measures how much water (aerially) is colder than a chosen threshold, so it quantifies extent and intensity together and is sensitive to the threshold choice. Using polygins like this can also then reveal the width of UCs, and, consequently, the location of the UCs (see also Section 3.4). A profile method measures where the cross-shore structure changes character, so it quantifies a length scale directly and is sensitive to the shape criteria instead. The fixed-distance indices our paper is concerned with are profile methods, and the defect we identify, namely a reference point that sits inside the feature it is meant to sit beyond, is a defect of that lineage. Our contribution is best described as repairing a known lineage of indices rather than as displacing the area-based lineage. The Discussion already makes a version of this argument for two-dimensional segmentation methods, and we will extend it to the isotherm-area methods and state that the choice between them follows from the question rather than from any ranking of the methods.

3.4 Prescribed cell locations

The referee expects the upwelling cells to emerge from the analysis rather than be prescribed, and asks for a few sentences where the concept is introduced. Yes, this is possible. Our reason for prescribing them is practical, namely that the five cells are long-established features of the Benguela with an extensive literature, and using the published windows keeps our per-cell results comparable with that literature. A method that located cells from the data (i.e., data-emergent) would be a different and larger contribution, and we agree it is the more satisfying one (and it is in fact part of the `uc2d_v3` method we mentioned earlier).

The five upwelling cells are prescribed from the published literature rather than located from the data. The reason is comparability, since per-cell results reported against long-established cell windows can be viewed beside earlier work. The method itself does not require prescribed windows, because the breakpoint regression is fitted per transect and the transect mesh is continuous along the coast, so a cell-free alongshore treatment is possible and would allow cell boundaries to emerge from the alongshore structure of x_b and $\bar{D}_{\text{in}}$.

4 The wind index

The referee quotes the following from the submitted Methods.

> AS SUBMITTED. “[...] namely when the alongshore-equatorward wind component, rotated to the local coast-normal direction following Lamont et al. (2018), drives net offshore Ekman transport in the southern-hemisphere upwelling-favourable sense.”

The referee is correct. The clause *rotated to the local coast-normal direction* is wrong. It is the alongshore component that drives offshore Ekman transport, and the sentence as written attributes the transport to a projection onto the wrong axis. The referee is also right that *rotated* is the wrong verb, since the wind is projected onto a local coordinate system rather than rotated, and that Lamont et al. (2018) is an analysis rather than a convention.

Because the referee raised doubts about the implementation, we checked the code before drafting this reply. What the script computes is the following. The local coast-normal unit vector is constructed from the coast and outer-contour endpoints of each transect. Its perpendicular is taken, giving the alongshore unit vector, and the sign is fixed so that the vector points equatorward. Wind stress is projected onto that equatorward alongshore vector to give τ_{along}, and the offshore Ekman transport follows as

$$UI_w = -\frac{\tau_{\text{along}}}{\rho_w f}, \tag{1}$$

with f the Coriolis parameter, negative in the southern hemisphere, so that equatorward alongshore stress yields positive offshore transport. That is the standard Bakun construction and it is physically correct. The coast-normal vector appears in the code only as the intermediate step from which the alongshore direction is obtained, and our sentence described that intermediate step as though it were the projection axis.

The manuscript in fact contains the correct description in the caption of Figure 4, which reads *rotated to the local equatorward unit vector*. The two descriptions contradict each other, and we did not notice.

The daily binary wind detector flags a day at a UC as wind-active when the daily-mean Bakun offshore Ekman transport at the UC is positive ($UI_w > 0$). At each transect the wind stress from ERA5 10 m winds is projected onto the local alongshore unit vector, oriented equatorward, following the coordinate construction of Lamont et al. (2018). Writing τ_{along} for that component, the offshore Ekman transport per unit length of coast is $UI_w = -\tau_{\text{along}}/(\rho_w f)$, where ρ_w is the density of sea water and f is the Coriolis parameter. In the southern hemisphere f is negative, so equatorward alongshore stress gives positive UI_w and drives transport away from the coast, which is the upwelling-favourable sense. The alongshore direction is obtained as the perpendicular to the local coast-normal direction defined by the transect mesh, so the wind projection and the SST profiles use the same geometry. This detector uses the sign of the offshore transport rather than the wind-speed threshold of Lamont et al. (2018), so it is not their event criterion. Wind-stress-curl pumping is not represented, since the curl contributes an upwelling term that does not depend on the choice of coordinate system and is outside the local Bakun formulation used here.

The referee suggests projecting the wind into the coordinates defined by the mesh derived from the Laplace potential. This is what our code does, and the replacement text now says so. We are grateful of the suggestion, because it identifies a real strength of the construction that the submitted text concealed, namely that the wind comparison and the SST profiles share one geometry.

Two related wording problems

The referee objects to a detector that *informs* a classifier and to a comparison of NUD with NUE. Sure, both objections are sound. We suggest:

The daily binary wind detector supplies the wind axis of the paired classification of Section 2.7, and the same daily flags are aggregated to the number of upwelling days and the number of upwelling events used in the comparison of Section 3.3.

On the second, the submitted phrase *the NUD / NUE comparison* is silly. We do not compare a count of days against a count of events, which would definitely be meaningless. We compare the SST-derived and wind-derived versions of each quantity, namely NUD_{BR} against wind-derived NUD, and NUE_{BR} against wind-derived NUE. The revised text will be clear about this.

5 Further Methods text portions rewritten

5.1 Section 2.8, the uncertainty estimates

The referee reports not understanding the essence of this section and asks why block bootstrapping is needed at all. Looking at what we have written, it seems the section conflated two different uncertainty questions without separating them, which is the likely source of the difficulty.

Two distinct uncertainty questions arise. The first concerns cell-scale summaries, such as the median breakpoint position at a cell or the fraction of SST-structure-only days. These summaries are computed over many transects, but adjacent transects are spaced at ca. 4 km and sample overlapping water, so they do not contain independent information. Treating each transect as an independent observation would make the summaries appear more precise than they are. We therefore resample in contiguous alongshore blocks of 25 and 50 km rather than resampling individual transects, since a block preserves the alongshore correlation within it and the blocks can be treated as approximately independent of one another. Where the summary is annual or seasonal, years and seasons are resampled as blocks for the same reason. Reported intervals are 1000-replicate percentile intervals.

The second concerns an individual profile. A cell-scale interval says nothing about whether the breakpoint on one profile is well determined, so we estimate that separately. For a stratified sample of 4000 accepted profiles we resample the pixels of each profile with replacement, refit the breakpoint 1000 times, and record the spread of the fitted positions. A narrow spread means the profile determines its breakpoint sharply. We also test whether the fitting objective has more than one minimum, since a profile with two candidate breakpoints of comparable quality has an ambiguous solution even when the fit is otherwise acceptable. Per-cell summaries of both quantities are in Table S6 and Fig. S4 in the Supplement.

> AS SUBMITTED. "We calculate block-bootstrap 95 % intervals for the main statistics. Alongshore blocks are constructed at the UC scale with two block sizes, 25 km and 50 km. [. . .] The block bootstrap above describes uncertainty around UC-scale summaries [. . .] and it is not, on its own, a statement about identifiability of an individual breakpoint."

5.2 Section 2.1, the gap-filled Level-4 inputs

The referee raises a limitation we had not treated, namely that Level-4 analyses fill gaps, that the filling can propagate a stale field through a long cloudy period, and that the northern Benguela is prone to persistent cloud days. The concern is a bother, and it is real; it directly affects a comparison between an SST-derived detector and a wind-derived one, since a stale SST field held over a cloudy interval would come out as a persistent structure while the wind flag varied freely. We have not quantified this and we will not claim to have done so (I can't think of viable solutions to the problem).

All four SST inputs are Level-4 analyses, so they are gap-free by construction and the values in cloud-obscured regions are interpolated. The interpolation draws on neighbouring water and on the recent history of the field, so an extended cloudy interval can produce a surface field that changes less than the true field does. The northern Benguela experiences persistent low cloud, so this is a relevant caution there in particular. The consequence for our analysis is that a run of SST-structure-active days spanning a cloudy interval may reflect the persistence of the analysis rather than the persistence of the feature, and this affects the SST-structure-only class of the paired classification, since the wind field is unaffected by cloud. We have not quantified the effect. The microwave product MW_OI is largely insensitive to cloud and would be the natural control, but its resolution is too coarse to retain accepted profiles under the AR(1)-corrected criterion, so it cannot serve that purpose here. We flag the issue as a limitation and note that per-pixel analysis-uncertainty fields, distributed with the products, would support a quantitative treatment in future work.

We will add a matching caution about ERA5. It is a reanalysis rather than an observational product, and its winds near a coast are subject to the resolution of the model orography and coastline, so agreement or disagreement with a satellite-derived SST detector is a comparison of two model-mediated products rather than a comparison against truth.

5.3 Walvis Bay, where the referee finds the result astonishing

The referee singles out the discrepancy between wind-derived and SST-derived upwelling at Walvis Bay and says the benefit of the new approach is not clear there.

Walvis Bay is the cell where our acceptance criteria are hardest to satisfy, and the relevant quantities are consistent with one another. It has the largest fitted breakpoint of the five cells at 264 to 270 km, the shallowest accepted inshore slope at $+0.014\,^{\circ}\mathrm{C\,km^{-1}}$, the highest rejection rate at 30.5 %, and the lowest number of active days at 118 of 365. It also becomes wind-only dominated in MAM and JJA. The threshold assay shows the mechanism, since raising the offshore-baseline requirement from 100 to 200 km cuts Walvis Bay acceptance from 0.77 to about 0.50, while the other four cells lose at most 6 %.

Together, these suggest that the Walvis Bay cool tongue is broad and its offshore edge is weakly expressed, so a criterion demanding a clear break of slope with a substantial offshore baseline rejects profiles there that a reader looking at the field would call upwelling. The wind-only days are therefore better explained by the conservatism of our acceptance criteria at this cell than by an absence of surface structure. This appears to be a limitation of the method at broad cells, so stating it as such is maybe more useful than calling it a finding. We will rewrite the Walvis Bay text accordingly.

6 Point-by-point replies

Abstract

Begin the first sentence with "Some". Agreed. Isotherm-area and anomaly-event indices do not contrast against a fixed offshore reference and the opening sentence should not imply that they do. This amends the revised abstract given in our reply to Referee 1, whose opening sentence now reads "Some satellite sea-surface-temperature (SST) upwelling indices measure the contrast between coastal water and an offshore reference point, and that reference is conventionally fixed at 50 to 100 km from the coast." The remainder of the revised abstract is unchanged.

Terminology

What is a "Lamont year"? A fair complaint, since we use the term without defining it. It denotes an upwelling year going from July to June, so that a southern-hemisphere upwelling season falls within one year rather than being split across two calendar years, following the convention of Lamont et al. (2018). We will define it at first use—or, thinking about it now, we will replace it with *upwelling year (July to June)*, which needs no citation to be understood.

What is a "Bakun offshore-Ekman-transport detector"? It means a binary daily flag that is positive when the Bakun offshore Ekman transport is positive. But, yes, the phrase compresses too much and is silly. We will write it that way and reserve *detector* for the two binary daily classifiers once each has been defined, since the paired classification does need a short name for each side.

Section titles announce their subject obliquely. Sure, accepted. Section 3.3 becomes "Quantifying SST and wind classifications agreement", Section 2.7 becomes "Agreement between SST and wind classifications", and Section 2.8 becomes "Uncertainty in cell-scale summaries and in individual breakpoints". Or something like that. We will apply the same test to every title, namely that it should state the subject in words before the section rather than after it.

Methods and Results

L. 150, if the two indices are different quantities, what is the purpose of the scatter plots in Fig. 2? The figure exists to show a dependence rather than to calibrate one index against the other. The two quantities are different estimands, so a high correlation would not permit substituting one for the other. What the figure shows is that the correlation between them is not a fixed property of the pair but a function of the reference distance, rising from 0.06 at 50 km to 0.93 at 500 km. This is the paper's main claim... namely that moving the reference changes what the index measures. We will rewrite the caption and the matching text to say this, since the submitted version gives the figure as a comparison and leaves the reader to infer why the comparison is worth making.

L. 272, what controls what in the false-discovery-rate passage? Okay, the sentence is poorly built and the phrase *FDR-controlled negative* is jargon. Forty seasonal correlations are tested, and at that number a few will come out negative by chance alone. The Benjamini–Hochberg procedure converts each p value into an adjusted value that bounds the expected proportion of false positives among those we declare significant. Four remain at the adjusted threshold. The revised text, given in Section 2.7 of our reply to Referee 1, states this in words before mentioning what the method is called.

Fig. 3, what is a data-emergent value of a variable? It means a value fitted from the data rather than fixed in advance, and it is our coinage (well, not really; it has been used in other scientific disciplines, but not in ocenography as far as I can recall). Since the contrast the paper draws is between a fitted reference and a prescribed one, we will replace *data-emergent* with *fitted* throughout and let the contrast with *fixed* convey the meaning.

L. 408, is the closing statement about surface heat content needed? No. We will delete it. No one would attempt a climate-scale heat-content assessment from a time series assembled from different upwelling indices.

Figures and Supplement

Fig. S5 is too small to read even when magnified, and the breakpoint is not visible. Sure. We will redraw it with fewer panels per page across more pages, mark the fitted breakpoint explicitly with a vertical line (and a label rather than leaving it implicit at the junction of the segments), and draw the two fitted segments in a colour that separates them from the observations. The proposed walk-through figure in our reply to Referee 1 shows the panel design we intend to use.

Are the rejected profiles in Fig. S5 rejected because of the coastal edge buffer, given that they still show a coastal temperature drop? Mostly not. The edge buffer rejects a profile only when the best-fitting breakpoint falls within 10 km of either end, which is uncommon. The usual cause of rejection is the slope-contrast criterion, since a profile that warms offshore gradually and monotonically has a coastal drop but no inflexion sharp enough for the inshore segment to be twice as steep as the offshore one. This is consistent with the rejected-profile slopes of ca. $+0.004$ to $+0.011\,^{\circ}\mathrm{C\,km^{-1}}$ reported in Results Section 3.1. We will label each rejected panel in Fig. S5 with the criterion that rejected it, which answers the question on the figure itself.

If upwelling is curl-driven rather than coastal, why should the inshore slope be steepest? It need not be... Curl-driven upwelling acts over a broad offshore band rather than in a narrow coastal strip, so it can raise cool water without producing the steep inshore gradient our criteria require, and our method will reject such a profile or fit a breakpoint far offshore. We will state this as a scope condition, namely that the method detects the surface signature of coastal upwelling and is insensitive by to broad curl-driven upwelling. The point also connects to the Walvis Bay situation discussed in Section 5, where the broadest cell is also the one our criteria have difficulty with.

The GitHub link returns 404. Apologies for the wasted effort. The repository is still private, and GitHub returns 404 rather than a permission error for private repositories, so the link gives no indication that it exists. We will make the repository public at revision, and the archived snapshot with its digital object identifier will replace the bare link in the code and data availability statement, so that the citation resolves permanently rather than depending on the state of a repository.

Figure and table legibility generally. Answered in Section 2. Every figure and table in the main text and Supplement will be given a once-over against a single minimum type size.

We thank the referee for a review that found an error that we missed, and for pressing on the physical justification of the method, which has improved the manuscript.

New reference proposed for the revision

This is the source the referee suggested, which we propose to add in Section 3. The four additional references proposed in our reply to Referee 1 are listed there.

de Szoeke, R. A. and Richman, J. G.: On wind-driven mixed layers with strong horizontal gradients: a theory with application to coastal upwelling, Journal of Physical Oceanography, 14, 364–377, 1984.