

Identifying Lagrangian Coherent Vortices in a Mesoscale Ocean Model

Nathaniel Tarshish^a, Ryan Abernathey^{b,c}, Ci Zhang^{b,c}, Carolina O. Dufour^d,
Ivy Frenger^e, Stephen M. Griffies^{a,f}

^a *Atmospheric and Oceanic Sciences Program, Princeton University, Princeton, New Jersey, USA*

^b *Division of Ocean and Climate Physics, Lamont-Doherty Earth Observatory, USA*

^c *Department of Earth and Environmental Sciences, Columbia University in the City of New York, USA*

^d *Department of Atmospheric and Oceanic Sciences, McGill University, Montreal, Quebec, Canada*

^e *GEOMAR Helmholtz Centre for Ocean Research Kiel, Kiel, Germany*

^f *NOAA/Geophysical Fluid Dynamics Laboratory, Princeton, New Jersey*

Abstract

We identify Lagrangian coherent vortices in a global mesoscale eddy-permitting ocean model using the rotation-based method of Haller et al. (2016). We present an analysis of the acute sensitivity of the identification results to varying the method’s free parameters, and develop physically justified parameter choices that allow for systematic vortex identification. In contrast to prior vortex studies, we probe the broad spectrum of coherency in the ocean by determining free parameter choices that partition the spectrum into distinct *coherency classes*, allowing for the identification of strictly coherent, moderately coherent, and leaky vortices. Our tuning methodology is grounded in a combination of sensitivity analysis, convergence tests, and consideration of the ocean model’s physics. To aid in this process, we introduce the *Coherency Index*, a novel Lagrangian diagnostic for mathematically quantifying the degree of material coherency of a Lagrangian vortex. We aim for this manuscript and the accompanying open-access code to serve as a manual and toolset for the oceanographer interested in harnessing a rigorous Lagrangian method to uncover coherent structures in ocean models and observations.

Keywords: Coherent Structures, Mesoscale Eddies, Vortex Identification, Lagrangian Analysis, RCLV

1. Introduction

Transient ocean mesoscale fluctuations play a central role in the global climate system, transporting climate-relevant tracers such as heat and carbon. In satellite observations and numerical simulations, a portion of these fluctuations take the form of materially coherent vortices (i.e., vortices whose interior material coheres together). Consensus has not emerged, however, as to the abundance of these coherent vortices in the global oceans and the magnitude of the associated transport.

In the literature, coherent transport questions have been investigated primarily through the lens of Eulerian-based vortex identification: a set of techniques that attempt to diagnose coherent structures via anomaly contours of an Eulerian field (e.g., sea level anomaly or Okubo-Weiss parameter of Okubo (1970); Weiss (1991)). To compute transport estimates, these studies make the fundamental assumption that Eulerian "vortex contours" trap and transport their material interior (for recent examples, see Raj et al. (2016); Dong et al. (2014); Zhang et al. (2014a,b)). Many studies justify this assumption by considering U_g/c , where U_g is the maximum geostrophic speed within the vortex and c is the mean translation speed (Chelton et al., 2007, 2011). When this diagnostic parameter is greater than unity, rotation dominates over translation, indicating that the identified feature departs from linear dynamics and "can advect a parcel of trapped fluid" (Chelton et al., 2011). As practitioners of this methodology acknowledge, the mere presence of some ability to trap fluid does not quantitatively guarantee complete material coherence over a vortex's total lifetime (Zhang et al., 2014a; d'Ovidio et al., 2013). Nonetheless, past studies make this assumption to reach estimates suggesting that identified vortices collectively play a significant role in global ocean transport (Dong et al., 2014; Zhang et al., 2014b; Petersen et al., 2013).

On the other hand, several studies have called into question the ability of Eulerian vortex identification methods to correctly diagnose material transport

30 (Abernathey & Haller (2018); Haller et al. (2016); Beron-Vera et al. (2015);
 dOvidio et al. (2013); Froyland et al. (2007) and references therein). The com-
 munity’s development of rigorous Lagrangian-based methods has cast doubt
 on Eulerian vortex methodology, challenging the notion that Eulerian anomaly
 fields and diagnostics such as U_g/c can reliably uncover material transport.
 35 In contrast, Lagrangian techniques exploit material trajectories—the fundamen-
 tal elements of fluid transport—and thus provide direct access to the transport
 structure of ocean flows.

A feature common to both Lagrangian and Eulerian vortex identification
 methods, however, is the prominent role of free parameters, including those
 40 related to numerical implementation. For a Lagrangian method, the number of
 free parameters varies from 3-4 (depending on the implementation) in the case
 of the Lagrangian-Averaged Vorticity Deviation (LAVD) method of Haller et al.
 (2016) and up to 6-7 for the geodesic ”black hole vortex” method of Haller &
 Beron-Vera (2013). An informative discussion of the strengths, weaknesses, and
 45 free parameters of Lagrangian vortex algorithms can be found in Hadjighasem
 et al. (2017), which demonstrates promising results for the LAVD method. For
 reference, the identification of the popular Eulerian sea surface height (SSH)
 eddies of Chelton et al. (2011) required nine free parameters.¹

It may be inevitable that these free parameters exist. After all, a vortex
 50 identification algorithm must take a global field (e.g., sea surface height, velocity,
 LAVD) as input and return a discrete object as output. As in mathematics,
 passing from the continuous to the discrete limit requires the introduction of a
 threshold or tolerance.

¹The free parameters of Chelton et al. (2011) are as follows: minimum SSH threshold for
 all pixels inside anti-cyclonic eddies, maximum SSH threshold for all pixels inside anti-cyclonic
 eddies, minimum SSH threshold for all pixels inside cyclonic eddies, maximum SSH threshold
 for all pixels inside cyclonic eddies, minimum number of pixels within eddy, maximum number
 of pixels within eddy, eddy SSH amplitude threshold, maximum distance between interior
 eddy points, and the choice of the number of pixels within a local neighborhood of an SSH
 extremum.

The presence of free parameters presents a challenge for the oceanographer,
 55 whose desire is to identify physical structures that exist in the ocean independent of the particulars of an identification method. To complicate matters, the relationship between the mathematical parameters of the Lagrangian method and the physics of the identified structures is opaque. For example, distinguishing between coherent versus incoherent sets of trajectories with the transport
 60 operator or spectral clustering approach boils down to a delicate mathematical task: identifying a gap in the eigenvalues of an operator’s spectrum (Froyland et al., 2007; Hadjighasem et al., 2016). Without exploring and understanding the physical significance of varying this parameter, the oceanographer cannot harness the method to objectively shed light on ocean transport.

65 In the applied mathematics literature that introduced these Lagrangian techniques, the methods’ sensitivity and free parameters are not always systematically examined. For example, the authors of several leading vortex methods in the comparison paper of Hadjighasem et al. (2017) omit details on how the method was tuned: “We therefore rely on our expertise and experience to choose
 70 a reasonable set of parameters for each method with the intention that (i) The choice of parameter(s) results in the most favorable outcome for the corresponding method and (ii) The outcome is robust, i.e., small variations in the parameters do not lead to drastic changes in the outcome”. Thus, while the applied mathematics literature provides promising tools, it lacks instructions on how
 75 to physically rationalize a tuning of the sensitive parameters. Furthermore, the literature lacks a clear definition of a “favorable outcome”, which suggests that the final results are determined by the author’s preferences rather reproducible objectives.

Our paper aims to help fill this gap in the existing literature by creating
 80 a bridge that spans from the methods papers of the applied math community to the vortex census studies of oceanography. Our objective is to implement and apply a Lagrangian-based vortex identification method to a high-resolution global ocean model, conduct a rigorous and thorough sensitivity analysis, and arrive at a physically justified tuning that systematically identifies materially

85 coherent vortices. We focus on a particular vortex identification method (the
 LAVD method of Haller et al. (2016)), but, in conducting our analysis, aim to
 provide a tuning methodology that is relevant for other techniques. Wherever
 possible, we strive to lend physical interpretation to the variation of the mathe-
 matical parameters. In doing so, we have taken steps to ensure that the tuning
 90 is robust, in the sense that small adjustments to the free parameters do not
 produce large changes in the results.

Stemming from this effort, we introduce a new Lagrangian-based diagnostic
 of material coherency, which we refer to as the *Coherency Index*. This diagnostic
 allows us to precisely quantify the sensitivity of the identified vortices' coherency
 95 to varying the method's free parameters. Beyond this usage, the diagnostic has
 broad applicability, and could be employed to compare the coherency of vortices
 across data sets, models, and identification methods.

2. Method and Materials

2.1. The LAVD Method

100 Haller et al. (2016) introduced a Lagrangian-based vortex identification (ab-
 breviated as the LAVD method) that exploits rotational coherence as a proxy for
 material coherence. We have implemented the LAVD identification method in
 an open-source Python package **floater** (available at <https://github.com/rabernat/floater>).
 The intuition for the method springs from the observation that fluid trajecto-
 105 ries inside a coherent vortex are spatially and rotationally organized in bands,
 layered concentrically about a rotating near-circular core. This combined spa-
 tial and rotational organization makes it possible to detect materially coherent
 objects via a trajectory-based metric of rotation.

In particular, we consider two dimensional flow and examine the relative
 vorticity defined as $\zeta = \partial_x v - \partial_y u$ and equal to twice the fluid's local rotation
 speed relative to the Earth's rotation. Let a fluid element be labeled by a ma-
 terial label \mathbf{a} , and its position in physical space at time t be notated as $\mathbf{X}(\mathbf{a}, t)$.
 We introduce the Lagrangian relative vorticity field, $\zeta(\mathbf{X}(\mathbf{a}, t); t)$, which takes

as input a particular material trajectory and a specific time. For convenience, we choose to label the material by its position in physical space at some reference time t_0 , so that $\mathbf{a} = \mathbf{x}_0$. We are interested in characterizing not just the instantaneous local rotation, but the average local rotation experienced by a material element over a finite time interval $[t_i, t_f]$. To do so, we define the Lagrangian-Averaged Vorticity Deviation,

$$\text{LAVD}_{t_i}^{t_f}(\mathbf{x}_0) = \frac{1}{t_f - t_i} \int_{t_i}^{t_f} \left| \zeta(\mathbf{X}(\mathbf{x}_0, \tau); \tau) - \overline{\zeta(\tau)} \right| d\tau, \quad (1)$$

where $\overline{\zeta}$ is the domain average of the relative vorticity field (Haller et al., 2016).

110 For a given time interval, LAVD assigns a single scalar to each trajectory, the magnitude of which describes the average amount of local rotation experienced by the material element. To understand the rotational structure of the flow, the convention is to visualize the Lagrangian field, $\text{LAVD}_{t_i}^{t_f}(\mathbf{x}_0)$, by plotting the value for each material element in the element’s initial physical location (i.e.,
 115 $\mathbf{X}(\mathbf{x}_0, t_i)$). An example visualization is shown in Figure 1 for the Kuroshio current in the NOAA/GFDL CM2.6 model (detailed in the following section). The algorithm utilizes this representation of $\text{LAVD}_{t_i}^{t_f}(\mathbf{x}_0)$, which we simply refer to as “the LAVD field for $[t_i, t_f]$ ”, to identify rotationally coherent vortices.

We now return to our earlier observation that coherent vortices are composed
 120 of collectively rotating fluid that is organized into concentric bands around a rotating near-circular core. Interpreting this feature in terms of the LAVD field provides the definition of a rotationally coherent Lagrangian vortex (RCLV) as a nested family of iso-LAVD contours containing an innermost LAVD maximum. The outer boundary of the RCLV is determined by a threshold on the convexity
 125 of the contour.² A sample nested family of LAVD contours is shown for an Agulhas ring in Figure 2. As explored in detail in Section 3.3, the convexity threshold is the primary free parameter of the LAVD method. The coherency

²A contour is convex if the line segment joining any pair of interior points lies entirely within the contour (e.g., squares, ovals, and circles are convex, but crescents or star-shaped polygons are not).




Figure 1: LAVD field computed over a 30-day interval in the Kuroshio region of the Northeast tropical Pacific Ocean. Smooth near-circular regions of LAVD anomaly represent mesoscale regions that rotate coherently over the subsequent 30-day interval.

and population of the identified RCLVs are acutely sensitive to the convexity threshold. The above presentation is oriented towards the oceanographer, but
 130 we refer readers to Haller et al. (2016) and Haller (2016) for a rigorous and mathematical presentation of how LAVD naturally arises from a dynamic polar decomposition of the deformation gradient.

The definition of an RCLV relies on the empirically-supported hypothesis




Figure 2: The contour field (shown in purple to yellow) of the Agulhas ring of Figure 8, identified with a contour increment of $4 \times 10^{-7} \text{ s}^{-1}$. Local maxima are shown as green points. In practice, we use a contour increment of 10^{-8} s^{-1} for RCLV identification, but the resultant density of the contour field is so high that visualizations make the field look nearly continuous and are less pedagogical.

that the more convex an LAVD contour, the more the interior fluid rotates
 135 together, and by extension, the more the fluid coheres. It is worth noting that
 this connection between LAVD contour convexity and material coherency, while
 critical to the method, lacks rigorous mathematical justification. For general
 finite contours, the justification is empirical and “motivated by the near-circular
 cross-section generally observed for stable vortices” Haller et al. (2016). The
 140 majority of Section 3.3 is devoted to investigating this empirical relationship.

As a preliminary demonstration of the method, we compare the evolution of
 vortex material identified by the method to randomly-selected “control” fluid.
 As shown in Figures 3, the rotationally coherent object diagnosed by the method
 exhibits a high degree of material coherence. In this particular case, we find a
 145 downwelling vortex that contracts the surface material into a localized patch. (In
 general, we also find upwelling RCLVs that disperse material and non-divergent
 RCLVs that preserve their area). The control material, however, is widely dis-
 persed and stirred with the environmental fluid.

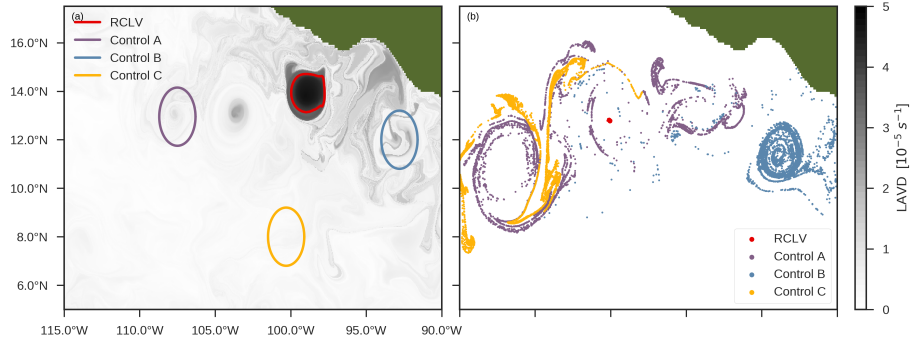


Figure 3: a) An LAVD field computed over a 60-day interval in the North Pacific Ocean. Contours indicate the detected RCLV boundary (shown in red) and randomly selected control volumes (shown in purple, blue, and gold). b) The material interiors after 60 days of advection by the surface flow (numerical details on particle advection are described in the subsequent section). The RCLV material (shown in red) contracts to a localized patch, a signature of a strongly downwelling vortex.

2.1.1. Implementation Details

150 Given Lagrangian trajectories, the LAVD method reduces the problem of identifying coherent vortices to the computational task of finding convex contours in a scalar field (i.e, LAVD). This computational task belongs to the field of *image processing* and is amenable to multi-core parallelization, and less demanding in comparison to other Lagrangian vortex identification methods
155 (as discussed in Hadjighasem et al. (2017)). Contour identification, which also forms the basis of Eulerian vortex algorithms, is a well-studied task that can be accomplished with existing optimized routines. We take advantage of Python’s popular `scikit-image` package described in Van der Walt et al. (2014). The scalability of the underlying computational task allows vortices to be identified
160 not just in the limited domains Lagrangian vortex methods are typically developed and tested on (e.g., cat’s eye flow, Bickley jet, and others in Hadjighasem et al. (2017)), but also in the high-resolution global domains of ocean models. In the past, this global identification was only computationally feasible with Eulerian vortex methods as in Petersen et al. (2013).

165 To apply the image processing methods, we create a two dimensional image from the LAVD scalar field. A pixel gives the value of the corresponding particle’s LAVD. The global domain we employ for this study spans all 360° in longitude and from 80° S to 64° N in latitude. This domain is covered with a Lagrangian particle mesh at a uniform resolution of $1/32^\circ$, producing an LAVD
170 image with dimensions of 11520×4608 pixels. (Note that we respect the spherical geometry when computing areas and distances by using the appropriate local tangent plane projections.)

Identifying local maxima and contours in an image requires the introduction of two free parameters: a parameter to control how many local maxima
175 are extracted from the image and a parameter that sets an LAVD increment for identifying contours. Similar parameters are also found in many Eulerian sea surface anomaly methods (Chelton et al., 2007; Frenger et al., 2015); the fundamental difference is that our underlying field is derived from Lagrangian

quantities. For this study, we inform our determination of these parameters by
 180 consideration of a) the length scale of RCLV permitted by the ocean model, b)
 sensitivity analysis, and c) convergence checks. These efforts are described in
 Sections 3.1 and 3.2. While these free parameters are inherent to the method
 (and equivalent parameters can be found in the original paper’s MATLAB im-
 plementation, publicly available at Hadjighasem (2016)), they are not discussed
 185 in the original manuscript of Haller et al. (2016). The other free parameter,
 the convexity threshold, which is fundamentally related to the coherency of the
 vortex, is discussed in Section 3.3.

2.2. The CM2.6 Coupled Climate Model

We identify RCLVs in the NOAA/GFDL CM2.6 coupled climate model,
 190 which makes use of a 0.1° ocean model component. The ocean component is
 based on the MOM5 code of Griffies (2012) and uses a z^* vertical coordinate
 (quasi-Eulerian). Horizontal grid spacing is roughly 11 km in the equatorial
 region and 5 km in the high latitudes. The atmospheric component makes use of
 a 50 km configuration with a full diurnal cycle of air-sea fluxes communicated to
 195 the ocean every 20 minutes. Hence, the ocean experiences a broad and realistic
 spacetime range of forcing. CM2.6 has generated a growing suite of publications
 considering the role of explicitly represented ocean eddies in the climate system
 (Delworth et al., 2012; Winton et al., 2014; Griffies et al., 2015; Dufour et al.,
 2015; Saba et al., 2016; Goddard et al., 2017; Dufour et al., 2017).

200 To maintain the horizontal isotropy of the grid and account for variations
 in the Rossby deformation radius with latitude, the meridional spacing of the
 model’s grid decreases towards the poles. At its furthest extent near 80°S , the
 meridional spacing reaches a minimum of approximately $1/24^\circ$. The model’s res-
 olution permits organized turbulent structures at the mesoscale, and parametrizes
 205 the effect of the unresolved sub-grid-scale dynamics on the mesoscale features
 as described in Griffies et al. (2015). A theoretical minimum of four grid cells
 is required to construct the basic circular velocity field pattern of a vortex as
 depicted in Figure 4. Accordingly, we intend to identify coherent vortices with

length scales greater than or equal to the length of two grid cells (which varies
 210 from a minimum of 10 km at high latitudes to a maximum of 20 km at the
 equator).

When this computationally demanding simulation was run, a decision was
 made to output horizontal velocities at daily-averaged frequency at the sea sur-
 face and monthly-averaged frequency for the full water column. Because the
 215 monthly-averaged velocities are highly smoothed, they are not suitable for de-
 tecting mesoscale coherent structures. Consequently, this study is limited to
 considering advection by the near surface horizontal flow (top model grid point,
 representing a depth of 10 meters). The Lagrangian trajectories generated by
 this flow can be considered to be the approximate trajectories of inertial parti-
 220 cles whose strong buoyancy effectively constrains them to remain close to the
 sea surface (Haller & Sapsis, 2008; Maxey & Riley, 1983). Examples of such par-
 ticles include surface debris and pollutants, as well as drogued drifters. While
 this flow is two dimensional, three dimensional motions can be inferred from the
 convergence and divergence of trajectories, as is discussed in detail in subsequent
 225 sections.

2.3. Lagrangian Particle Mesh

Because linearly interpolating the velocity field allows us to probe length
 scales slightly finer than the grid-scale (which is greater or equal to $1/24^\circ$), we
 cover the domain with a Lagrangian particle mesh of uniform $1/32^\circ$ resolution.
 230 This Lagrangian particle resolution corresponds to initializing a nominal $4 \times$
 4 array of particles in the interior of each nominal $0.1^\circ \times 0.1^\circ$ velocity cell as
 shown in Figure 4. In this work, we advect the Lagrangian particles in the
 daily mean surface currents from the pre-industrial CM2.6 simulation of 200
 years. Advection using archived velocity fields is conducted with the MITgcm
 235 (Marshall et al., 1997; Adcroft et al., 2014) in "offline mode" as described in
 Abernathey & Marshall (2013). Trajectories are integrated using a 4-th order
 Runge-Kutta scheme with a timestep of 900.0 seconds.

Velocity fields are interpolated from the Arakawa B-grid of MOM5 to an




Figure 4: Schematic depicting a horizontal velocity field containing two neighboring grid-scale vortices (with centers A and B) on a Southwest convention Arakawa C-grid. Each vortex exhibits a four-cell clock-wise circular flow pattern. A sample of the Lagrangian particle mesh (shown as black points) is displayed in the lower leftmost grid cell. The vortex centers are three grid cells from each other, which corresponds to a separation of roughly ten Lagrangian particles.

Arakawa C-grid required for the MITgcm (Arakawa & R. Lamb, 1977). The
 240 MOM5 B-grid is based on the Murray Tripolar grid of Murray (1996), which
 has a standard quasi-mercator latitude by longitude format south of 65° N, but
 contains a non-spherical grid north of 65° N with singularities over Northern
 Canada and Northern Russia (Griffies, 2012). To avoid the complexities as-
 sociated with interpolating from the non-spherical polar region of the B-grid
 245 to a spherical C-grid, the domain is truncated at 65° N. The Lagrangian par-
 ticle mesh covers this near-global domain at $1/32^\circ$ resolution with 37 million
 particles. The LAVD and position of each Lagrangian particle is output and
 saved daily. For example, Figure 3 exhibits particle trajectories after 60 days of
 advection.

250 For a 30-day interval, we distribute the advection of the global particle mesh
 in MITgcm over 128 processors. This computation takes an hour and produces

32 Gb of trajectory data.

3. Results

We proceed to analyze RCLV identified in CM2.6 and walk through how
255 to rigorously determine the method’s free parameters in order to yield robust
results. We will take the trajectory data as given, and thus do not conduct a
sensitivity analysis of the Lagrangian trajectory computation (e.g., we do not
address sensitivity to the advection time-step, spatial resolution of the velocity
field, or vorticity computation). We only concern ourselves with the sensitivity
260 of the free parameters inherent to the LAVD method, which is the objective of
the paper. Without loss of generality, we conduct this free parameter analysis
on RCLV identified in a 30-day window.

3.1. Minimum Distance

Following the RCLV recipe, we seek closed and convex contours that encircle
265 local maximum LAVD pixels. The first step is to identify local maximum LAVD
pixels using the `skimage.feature.peak_local_max` method. By definition, a
local maximum exists with respect to a neighborhood whose extent requires
specification, thus introducing a new parameter, *min_distance*, which is the
minimum allowable number of LAVD pixels separating two local maxima. The
270 largest possible number of local maxima can be extracted with *min_distance* = 1.

We note that the LAVD machinery is designed to identify stand-alone vor-
tices. The problem of vortex merger and detecting multi-core vortex structures
(as discussed in Hughes & Miller (2017)) is beyond the purview of this project.
275 Given our focus on identifying single-core RCLVs, and considering our former
observation that the smallest vortex that is theoretically permitted by the model
occupies 4 grid cells, we expect a minimum separation of at least three grid cells
between LAVD maxima at the respective centers of two neighboring RCLVs as
illustrated in Figure 4. Interpreting this criterion in terms of the nominal $4 \times$

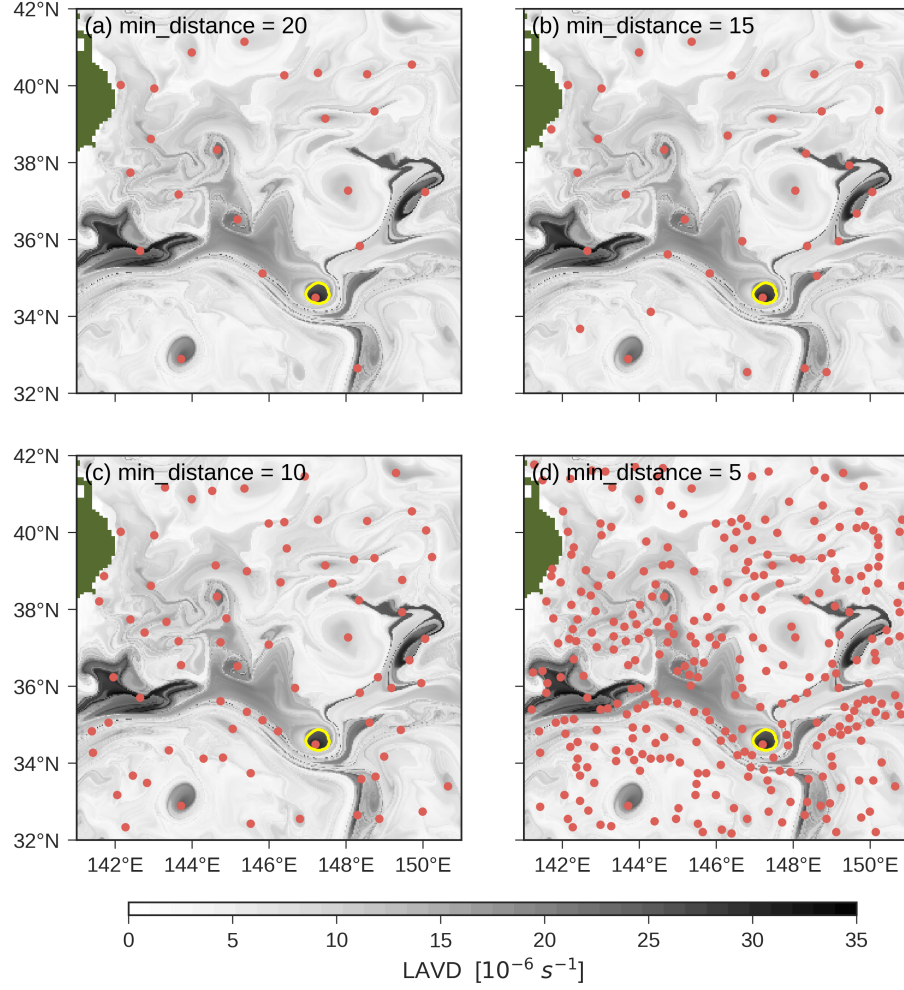


Figure 5: Kuroshio LAVD field with the local maxima (overlain in red scatter points) from identification with *min_distance* choices of a) 20, b) 15, c) 10, d) 5 pixels. The RCLV search in all cases identifies a single vortex whose boundary (shown as yellow contour).

280 4 LAVD pixel array in the interior of each grid cell provides the requirement that LAVD maxima should be at least 12 pixels apart. Sensitivity analysis reveals that identification results largely do not depend on slight variations in *min_distance*. In practice, we find that mesoscale vortex features are almost always separated by at least a hundred kilometers. A demonstration of this characteristic spacing is shown in Figure 5. Even though the number of local maxima 285 strongly varies, the final identified vortex field is identical: the prominent and isolated mesoscale anomaly is separated from its nearest neighbor maxima by at least 20 pixels. Global identification with *min_distance* set to 5, 10, and 15 pixels also yielded identical RCLV fields. We have used *min_distance* = 10 in 290 the following analysis.

3.2. Contour Increment

Given a local LAVD maximum, the next step in the identification process is to find the outermost convex iso-LAVD contour that encloses the maximum. We numerically identify this boundary contour by iterating through the nested 295 contour field, starting from the maximum and expanding outward. The iterative search terminates when a contour is identified that exceeds the convexity threshold. Contours are found using the `skimage.measure.find_contours()` method, which implements the marching cubes algorithm of Lorensen & Cline (1987). To find contours, the method linearly interpolates between the image's 300 discrete pixels. As a result, the contour field is continuous; a distinct contour exists for every value in the interval between the minimum and maximum LAVD pixel values in the image. To make progress, we are required to specify a fixed LAVD increment with which to iterate through the nested contour field. This contour increment introduces another free parameter to the method. Figure 305 2 displays the nested contour field for the Agulhas ring field of Figure 8. A contour increment of $4 \times 10^{-7} \text{ s}^{-1}$ was employed to produce this contour field.

A coarse LAVD contour increment will distort the geometry of the LAVD field, giving rise to a problematic coupling between the convexity deficiency and the contour increment. We desire that the effects of varying the free parameters




Figure 6: Average difference in LAVD between a pixel and its contiguous neighbors, computed for the LAVD field shown in Figure 2. Contours (shown sequentially from red to yellow) are the RCLV identification results for various choices of contour increment. Convergence occurs with a contour increment of 10^{-8} s^{-1} .

310 to be independent, and must ensure that the contour increment is fine enough
to resolve the contour convexity and minimize the coupling with convexity de-
ficiency. On the other hand, we expect that the amortized time complexity of
the vortex search varies inversely with the contour increment. An excessively
fine contour increment will thus require great computational cost, yet minimally
315 impact the grid-scale features of the RCLV boundary.

Arriving at a contour increment that balances the competing interests (re-
solving the contour convexity vs. minimizing computation time) requires un-
derstanding how the LAVD varies from pixel to pixel. We assess this variation
by, for instance, taking the average of the difference between the LAVD value
at a pixel and the LAVD value at its four neighboring pixels:

$$\overline{\Delta \text{LAVD}}[i, j] = \frac{\sum_{(k,l) \in \text{Neighbors}} \left| \text{LAVD}[i, j] - \text{LAVD}[k, l] \right|}{4} \quad (2)$$

where $||$ is the absolute value and the sum runs over the indices of the contiguous
pixels, $\text{Neighbors} = \{(i+1, j), (i-1, j), (i, j+1), (i, j-1)\}$. Figure 6 presents
this averaged difference field for the Agulhas LAVD field previously plotted in
Figure 2. With the exception of fewer than ten pixels in the local 128 by 128
320 pixel domain, the pixel-to-pixel variation in LAVD exceeds 10^{-8} s^{-1} . We thus
expect that a contour increment of 10^{-8} s^{-1} appropriately resolve the grid-scale
RCLV features.

We verify this hypothesis by conducting a sensitivity analysis in which we
compare RCLV boundaries identified with varying choices of contour increment.
325 The resulting contours are shown in Figure 6. Per our expectations, convergence
is achieved with a contour increment between 10^{-7} to 10^{-8} s^{-1} . On a global
scale, we tested the robustness of the contour increment choice by comparing
identification with a contour increment of 10^{-8} s^{-1} to 10^{-9} s^{-1} . Both resultant
datasets contained the same number of vortices and there was less than a 1%
330 difference in cumulative area between the two datasets. As a result, we have
employed a contour increment of 10^{-8} s^{-1} in this paper. For reference, RCLV
detection with `floater` on a global LAVD image with a contour increment of

10^{-8} s^{-1} takes roughly 3 hours.




Figure 7: Schematic of a contour surrounded by its convex hull. The convexity deficiency equals the difference in area between the convex hull and the contour (orange region) divided by the contour’s area (blue region). For the above contour, the convexity deficiency equals 0.06.

3.3. Convexity Deficiency

For the LAVD method, the fundamental tuning question is how to determine the boundary of the RCLV using knowledge of the LAVD contours’ convexity. To this end, we utilize the convexity deficiency, abbreviated as CD, as a measure of the convexity of a 2D LAVD contour. CD is equal to the area between the contour and its convex hull, divided by the area enclosed by the contour (Haller et al., 2016). The convex hull of a curve is the smallest convex set that contains the curve as shown in Figure 7. If a curve has a small value of CD, then the curve is nearly convex.

While plots in the vein of Figure 3 evince a sound connection between LAVD contours and material coherency, they do not directly address the question of how to determine the convexity deficiency parameter. This question is particularly important because the RCLV identification results are highly sensitive




Figure 8: LAVD field computed over a 30-day period in the Agulhas region. Contours (shown in blue) are the RCLV identification results for various choices of CD, all encircling an Agulhas ring (formerly shown in Figure 2).

to variations in CD. This sensitivity is demonstrated in Figure 8 for the RCLV detection of an Agulhas ring.

Varying CD from .01 to .15 results in a concentric set of RCLV boundaries with increasingly large spiral filaments. This spiral shape is ubiquitous to RCLVs and reflects the fact that RCLVs decay from the outside towards the center by “unrolling” their exterior layer-by-layer. This unrolling motion can be seen in Figure 9, which shows the fate of the Agulhas ring material interior. RCLVs were identified using only the first 30 days of trajectory data, but trajectories past 30 days are shown to demonstrate the vortex decay process. All the vortices become less compact with time, but the higher CD cases show a faster deterioration with material globally breaking away from the core vortex and dispersing.

Aside from expanding the boundaries of previously identified RCLVs, increasing the CD also leads to the identification of new structures. This behavior is demonstrated in Figure 10 which displays the detection results for a domain in the South Atlantic with Agulhas ring activity. The number of the identified structures varies sharply with the CD: identification with CD choices of .01, .05, .1, .15 yields 3, 11, 20, and 24 vortices respectively.

As we might expect, the three RCLVs identified with $CD = .01$ remain compact and coherent over the 30-day advection. In comparison, the structures that appear only at higher CD exhibit more filamentary tails and less compact geometry. We emphasize that varying CD changes not only the number, but also the nature of the identified RCLV. In the subsequent section, we introduce a statistical measure to quantify this difference in coherency behavior.

To settle on a value of CD thus requires a subjective judgment of how coherent a feature must be in order to merit distinguishing it from the turbulent environmental fluid as a “coherent structure”. To avoid this subjectivity, the idea of simply setting $CD = 0$ and eliminating it as a free parameter may sound appealing. The consequence is that RCLV boundaries are strictly convex polygons, composed of edges with length equal to the resolution of the Lagrangian particle mesh. Upon closer examination, we find that this choice is unsatisfactory: we have merely transformed the ambiguity in determining CD into an

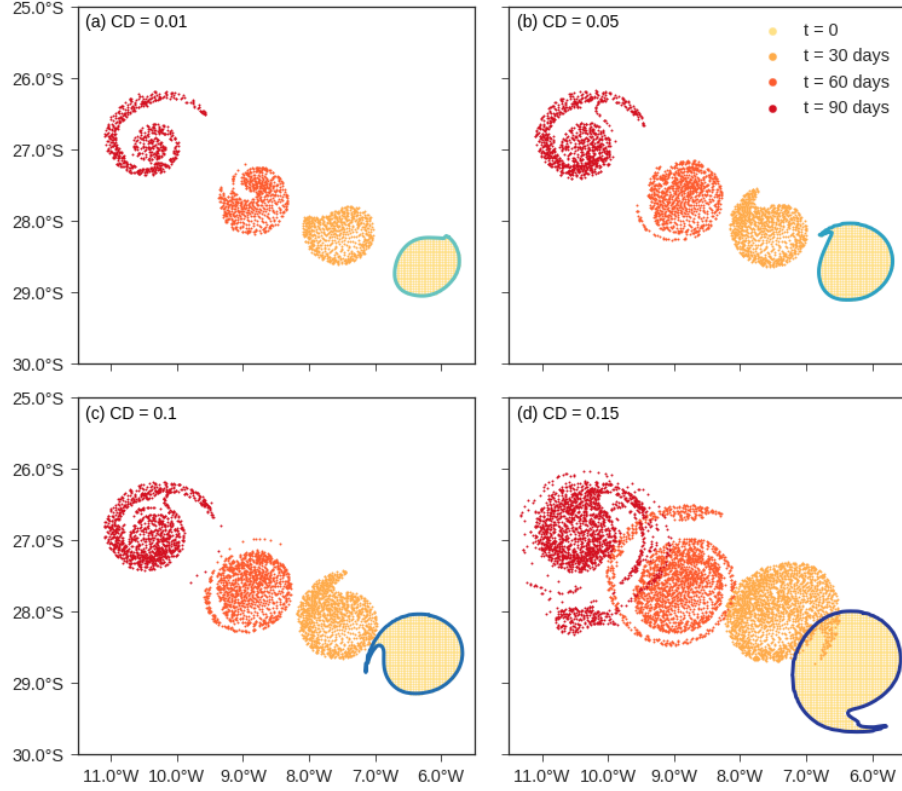


Figure 9: Evolution of RCLVs shown in Figure 8 for parameter choices of a) $CD = .01$, b) $CD = .05$, c) $CD = 0.1$, d) $CD = 0.15$. RCLVs were identified using only the first 30 days of trajectory data. Later trajectories are shown to demonstrate the vortex decay process. See text for discussion.

ambiguity in choosing the Lagrangian mesh resolution.

Perhaps part of the uncertainty here is due to the somewhat ill-defined nature of a coherent structure’s boundary. This ambiguity is even present in
 380 McWilliams’ original work on 2D quasi-geostrophic turbulence, where he steers clear of quantifying the boundary and a precise distinction between vortex and environment is challenging (McWilliams, 1984). Recently, Wang et al. (2016) studied the life cycle of an Agulhas ring using Lagrangian-based vortex detection
 385 and showed that periodic “coherence regain” events, during which the vortex

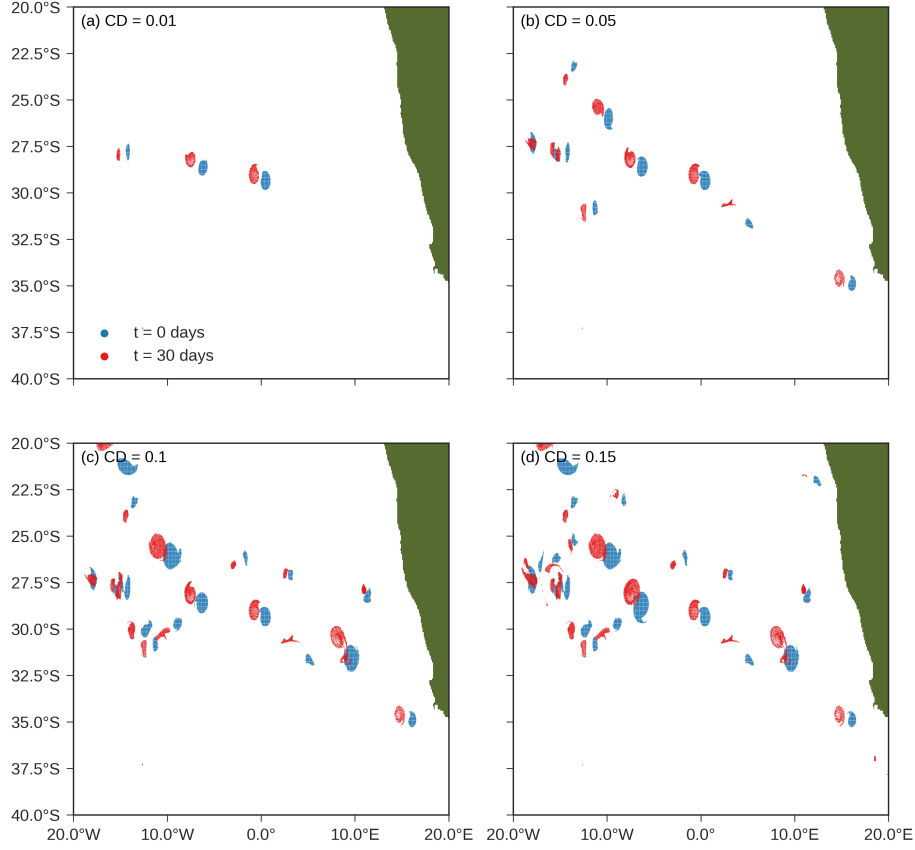


Figure 10: Initial locations (shown in blue) and final locations (shown in red) of 30-day RCLVs in the Agulhas region of the South Atlantic Ocean. RCLVs are identified with parameter choices of a) $CD = .01$, b) $CD = .05$, c) $CD = 0.1$, d) $CD = 0.15$.

entrains surrounding fluid, complicate this distinction between environment and vortex.

This ambiguity in the definition of a coherent structure’s boundary—along with the acute sensitivity of the method to CD —motivates the idea of not settling on a single value of CD . Rather, we argue that it is appropriate and desirable to leverage the sensitivity to CD in order to separate RCLVs into three broad classes: strictly coherent, moderately coherent, and weakly coherent (aka

leaky). In order to accomplish this goal, we need to choose three representative CD values for each coherency class. While there is some subjectivity in choosing representative CD values, this setup is arguably better motivated than the alternative of simply choosing a single CD for all of the analysis. This approach also has the advantage of revealing dynamical insight on the geographical distribution of coherency in the ocean. For instance, it may be valuable to know what regions form leaky vortices exclusively or which regions experience a high frequency of strict coherent vortex formation.

3.4. Coherency Index

To achieve a coherency categorization based on CD, we have to make precise the relationship between material coherency and CD. To this end, we introduce a *Coherency Index*. To quantify an RCLV's spatial compactness, we compute the variance of the particle positions,

$$\sigma^2(t) = \left\langle \left| \mathbf{X}(t) - \langle \mathbf{X}(t) \rangle \right|^2 \right\rangle, \quad (3)$$

where $\langle \rangle$ indicates an average over the set of RCLV particles and $|\cdot|$ is the standard Euclidean distance. A coherent vortex remains spatially compact under advection. To assess the change in spatial compactness over a time interval $[0, t]$, we compute the *Coherency Index* (CI),

$$\text{CI} = \frac{\sigma^2(0) - \sigma^2(t)}{\sigma^2(0)}. \quad (4)$$

The sign and magnitude of CI specify the material coherency of an RCLV. By construction, CI is positive for vortices that are strongly coherent, with a theoretical limit of $\text{CI} = 1$ for vigorously downwelling vortices that contract to a point (as in Figure 3). For a weakly coherent vortex that sheds material, $\sigma^2(t)$ can be arbitrarily larger than $\sigma^2(0)$, and CI becomes arbitrarily negative. Figure 11 presents sample vortex motions and their associated CIs. The vortices were selected from the vortex field discussed earlier in panel d) of Figure 10. As we might expect, the CI metric penalizes a vortex for dispersing and developing filaments, and rewards a vortex for growing more compact. We also observe

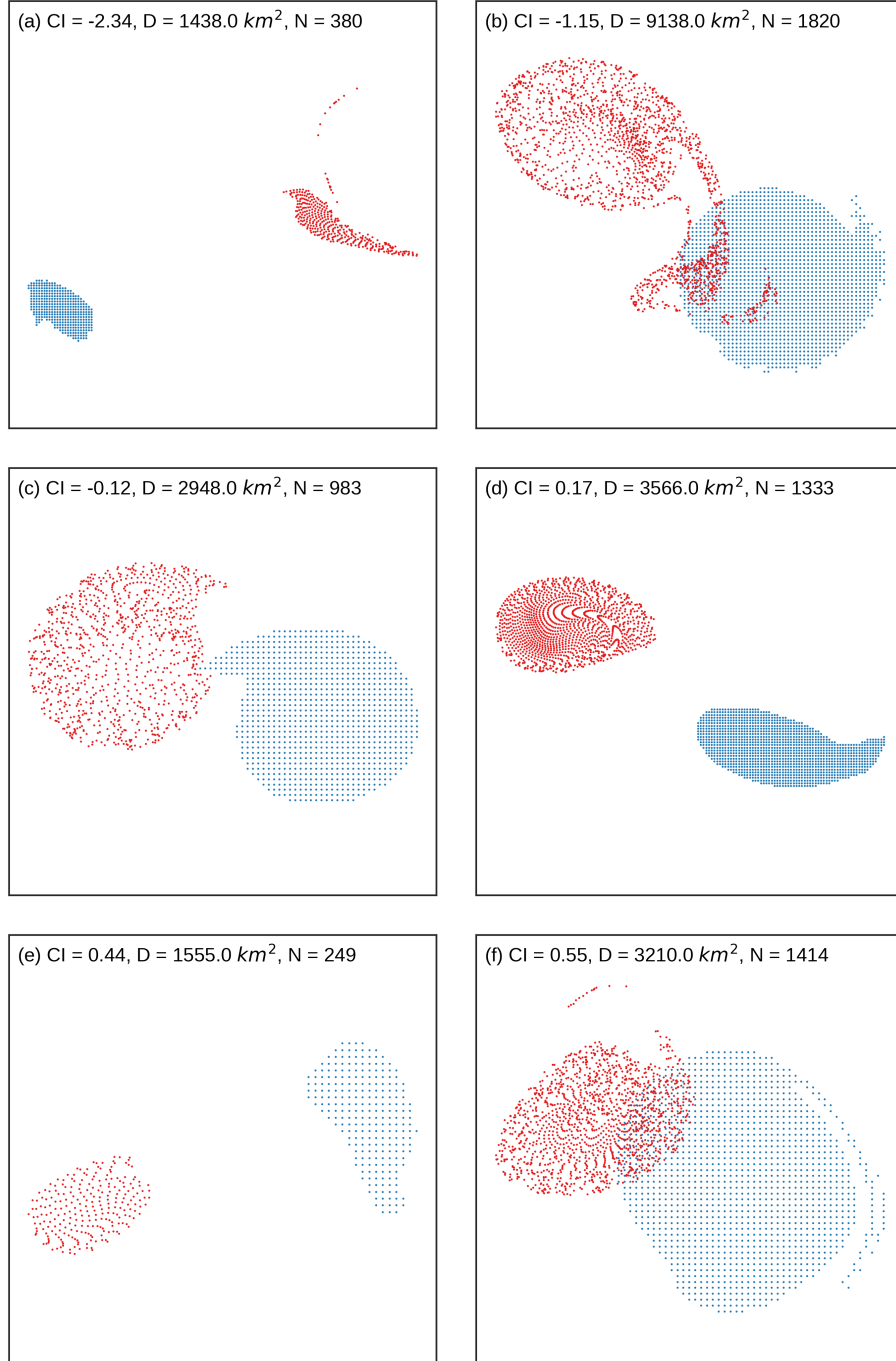


Figure 11: Values of Coherency Index (CI) for representative RCLVs, which were selected from the vortex field shown in panel d) of Figure 10. Rows demonstrate examples of leaky, coherent, and contracting vortices respectively. Initial locations of trajectories are shown in blue and final locations are shown in red. The dispersion (D) and number of particles (N) (directly proportional to the vortex's area) are included to illustrate the explanation of why the dispersion is a poor coherency metric.

that CI values close to zero correspond to vortices whose movements closely resemble solid body motion.

In prior literature, the variance of Lagrangian particle *displacements* has been used to quantify how sets of particles spread out over time (see the review of LaCasce (2008), and references therein). This quantity, known as the dispersion, is given by

$$D(t) = \left\langle \left| \mathbf{X}(t) - \mathbf{X}(0) - \langle \mathbf{X}(t) - \mathbf{X}(0) \rangle \right|^2 \right\rangle. \quad (5)$$

Note that $D(t)$ measures how the particles disperse over the time interval $[0, t]$, while $\sigma^2(t)$ measures only the instantaneous spatial distribution of the particles at a fixed time t . We now explain several reasons why $D(t)$ has limited utility as a coherency metric.

First, we note the problematic dependence of $D(t)$ on the vortex's area, which governs the spatial scale of the displacements. The same value of $D(t)$ can correspond to a fairly coherent large vortex and to a leaky small vortex. This undesirable feature of $D(t)$ is demonstrated in Figure 11. A similar point is illustrated in Figure 2 of LaCasce (2008) with the conclusion being that “[w]hile the dispersion reflects the clouds size, it is fairly insensitive to the clouds distribution in space”.

Another fundamental limitation of $D(t)$ is the quantity's inability to distinguish between radial expansion versus contraction. Consider a particle cloud with circular symmetry that is radially expanded or contracted. Since the vortex maintains its circular shape, the mean displacement ($\langle \mathbf{X}(t) - \mathbf{X}(0) \rangle$) is zero by symmetry. The value of $D(t)$ is then set by only the *magnitude* of the change in the vortex's radius. $D(t)$ is ignorant of whether the change in the vortex's radius is positive or negative, and thus cannot distinguish between expansion and contraction. Therefore, $D(t)$ overlooks a signature difference between upwelling and downwelling vortices. These limitations of $D(t)$ are not shared by CI, which is independent of vortex area and capable of distinguishing expansion from contraction.

To further understand the relationship between CI and the convexity defi-




Figure 12: A box-and-whiskers plot of the Coherency Index distributions for RCLV datasets resulting from identification with various choices of convexity deficiency parameter. For each distribution, the mean is indicated by a black point, the second (third) quartile is contained in the lower (upper) box, and the standard deviation is displayed as a whisker. Data farther than one standard deviation from the mean is plotted discretely. RCLVs were identified across the full global domain using 30-day trajectories. To put the values of Coherency Index in context, see Figure 11. Also see text for further discussion.

ciency, we conduct a global identification of RCLVs for various choices of convexity deficiency. We then compute and compare the Coherency Index distributions. All calculations are done with a global 30-day LAVD field. The results are summarized in the table below and displayed in Figure 12. Note that a

440

limited spatial region of this global data set was previously shown in Figure 10.

	Number	Mean CI	Median CI
CD = .01	109	0.28	0.48
CD = .05	346	-0.18	0.20
CD = 0.1	579	-0.48	0.11
CD = 0.15	789	-0.74	0.03
CD = 0.2	1014	-0.91	-0.03
CD = 0.25	1250	-1.17	-0.08
CD = 0.3	1534	-1.47	-0.16

The clear trend aligns with our expectations: the higher the value of convexity deficiency, the greater the number of identified leaky vortices. This fact is reflected by the Coherency Index distributions medians, means, and quartiles all shifting uniformly down the y-axis towards negatives CI values in Figure 12.

First, we observe that for the strictest convexity deficiency threshold of $CD = .01$, the mean convexity index is positive, indicating that the method is primarily identifying coherent objects that grow more compact with time. At $CD = .05$, the mean value of the convexity index changes sign to $-.18$. This signals a coherency regime change: although the median is still positive, there are now many dispersive RCLVs in the dataset. All subsequent larger values of CD intensify the leakiness and shift the mean further negative. The key point here is that to classify the strictly coherent RCLVs we should choose a $CD < .05$.

The next shift in coherency regime occurs between $CD = 0.15$ and $CD = 0.2$. In this interval, the median changes sign, revealing that at $CD = 0.2$ most identified vortices are leaky. The proximity of the mean to the 1st quartile of data also reflects the presence of very incoherent vortices. Since most vortices leak and some do so to an extreme degree, we conclude that between $CD = 0.15$ and $CD = 0.2$ we transition to the weakly coherent regime. These findings suggest that we employ parameter choices of $CD = .01$, $CD = 0.1$, and $CD = 0.25$ to survey the three different types of vortices.

In summary, we delineated between the leaky and moderately coherent regimes by when the mean Coherency Index is approximately zero. We then

defined the end of the moderately coherent by when the median Coherency
 465 Index is approximately zero. Sample results for the CD values corresponding
 to these definitions are shown in Figure 13 for a domain in the Pacific Ocean.
 The varying dispersive nature of the identified vortices is in agreement with our
 expectations based on the above definitions.

The broad fluctuation in the results emphasizes the importance of under-
 470 standing the method’s sensitivity and how it relates to the physics of the iden-
 tified vortices before applying the method to address vortex census questions.
 Importantly, we find an order of magnitude more leaky vortices than strictly
 coherent vortices. Censuses conducted with $CD = .01$ versus $CD = 0.25$ would
 consequently reach RCLV net transport estimates that differ by at least an order
 475 of magnitude.

For comparison, we have included panel a) of Figure 14 to demonstrate the
 dispersive nature of randomly selected patches of fluid. The control results are
 characterized by a CI distribution with an approximate mean of -7 and median
 of -5.

480 4. Discussion and Conclusions

The above analysis demonstrated that the LAVD method results are highly
 sensitive to variations in the free parameters: convexity deficiency, *min_distance*,
 and contour increment. We also described, however, that this sensitivity can be
 physically rationalized, enabling the oceanographer to utilize the method to re-
 485 veal objective structures. Through convergence analysis, we have demonstrated
 that the full mesoscale vortex field can be detected with $min_distance = 10$,
 and that the nested LAVD contour fields can be fully resolved with a con-
 tour increment of 10^{-8} s^{-1} . Subsequently, we introduced a Coherency Index to
 quantify the coherent versus dispersive nature of the identification results. This
 490 diagnostic enabled us to select tunings of the convexity deficiency that explore
 the entire spectrum of coherency in the ocean. In particular, we showed that
 convexity deficiency choices of .01, 0.1, and 0.25 are representative values for

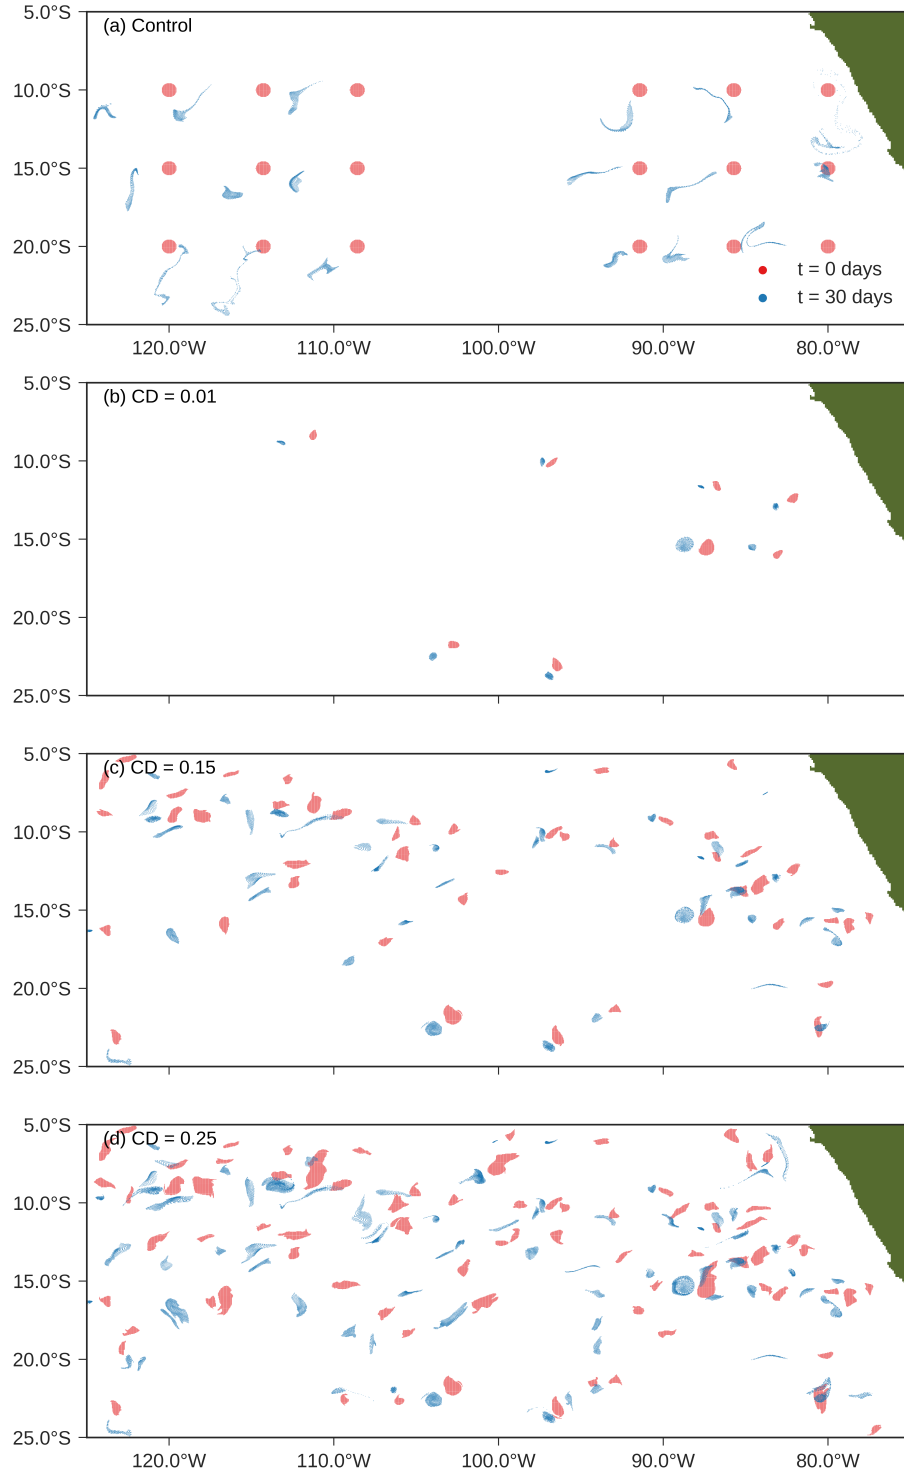


Figure 13: a) trajectories of two randomly selected arrays of control regions, and RCLV results for b) $CD = .01$; c) $CD = 0.1$; d) $CD = 0.25$ on a common subdomain in the Pacific Ocean. See text for discussion.

strictly, moderately, and weakly coherent vortices. Note that because resolution and dynamics vary from model to model, these exact values may only be appropriate for our specific model dataset. However, the tuning methodology developed here should be equally effective and relevant for applying the method to other models and observations. Furthermore, the coherency index provides a means to compare and assess results across models and even identification methods. Forthcoming work builds on the foundation laid in this paper, and presents the results from a global census of vortices for each coherency class in CM2.6.

Our analysis was limited to the two-dimensional surface manifestation of the vortex flow, but coherent vortices are, of course, fully three dimensional structures. The significant computational challenge posed by the 3D problem in conjunction with the absence of subsurface daily velocity field data in our model put 3D considerations beyond the scope of this study. We share this 2D limitation with nearly all other vortex studies. The LAVD methodology, however, naturally extends to three dimensions, in which a three-dimensional array of Lagrangian particles may be employed to search for convex iso-surfaces of LAVD as in Section 11 of Haller et al. (2016). By only studying the surface realization of the vortex, we are neglecting the vertical motion responsible for the upwelling and downwelling. The presence of this vertical motion, however, can still be directly inferred and quantitatively addressed within the context of our 2D results by invoking three-dimensional incompressibility.

We expect that horizontal motion dominates over vertical motion for vortices that show no signs of convergence/divergence (i.e. $CI \approx 0$) as in Figure 11c,d. For downwelling vortices, however, the horizontal motion of a surface-constrained Lagrangian particle may differ significantly from the horizontal motion of the fully 3D fluid trajectory. Consider the downwelling vortex of Figure 3. The vortex has an initial horizontal surface area of $3.5 \times 10^4 \text{ km}^2$, and the depth of the ocean model’s surface grid cell is 10 m, giving a total initial volume of 350 km^3 in the surface layer. After 60 days of advection, this material has been compressed to a narrow core with an area of 430 km^2 . Consequently, we

infer that the vortex downwelled over the 60 day interval with a vertical volume
 525 transport of $6.7 \times 10^4 \text{ m}^3/\text{s}$ (or $6.7 \times 10^{-2} \text{ Sv}$) . Assuming that the downwelling
 occurred in the vortex core, we also estimate a characteristic downwelling ve-
 locity of $1.6 \times 10^{-4} \text{ m/s}$ (or 13.5 m/day). Thus, if the vortex column penetrates
 less than approximately 800 m into the ocean interior, surface fluid may escape
 the coherent structure by being pumped down through the vortex column and
 530 exiting at depth. Future research should aim to explore this 3D structure of
 coherent vortices.

In terms of the detection method, implementing a root finding algorithm in
 the contour search could potentially improve the algorithm. Currently, we con-
 duct the contour search by stepping through the LAVD field at a fixed contour
 535 increment. Root finding algorithms (e.g., the bisection method as described in
 Burden & Faires (1985)) are designed to rapidly converge on the features of
 a function by taking taking variably-sized steps. Consider a family of LAVD
 contours and let $\text{CD}(x)$ give the convexity deficiency of the contour with LAVD
 $= x$. Suppose that the desired CD threshold for RCLV identification is CD_o .
 540 The RCLV boundary is then specified by $\text{CD}(x) = \text{CD}_o$, and x can be found
 using a root finding algorithm. Root finding would eliminate the need to specify
 a fixed contour increment. Efforts are underway to implement this in `floater`
 (available at <https://github.com/rabernat/floater>) and explore it as an alterna-
 tive means of searching the contour field.

545 5. Acknowledgment

NT would like to thank Spencer Clark for his Python expertise and helpful
 discussions. The authors thank Laure Zanna and Sonya Legg for discussion
 and comments. We thank four anonymous reviewers whose input has improved
 the manuscript. NT acknowledges support from U.S. Department of Energy
 550 under Award DE-SC0012457. CD and NT acknowledge support from National
 Aeronautics and Space Administration (NASA) under Award NNX14AL40G.
 RPA acknowledges support from an NSF CAREER award (OCE 1553593) and

NASA award NNX16AJ35G.

References

- 555 Abernathey, R., & Haller, G. (2018). Transport by lagrangian vortices in the eastern pacific. *J. Phys. Oceanogr.*, *48*, 667–685. doi:10.1175/JPO-D-17-0102.1.
- Abernathey, R., & Marshall, J. C. (2013). Global surface eddy diffusivities derived from satellite altimetry. *J. Geophys. Res.*, *118*, 901–916.
- 560 Adcroft, A., Hill, C., Marshall, J. et al. (2014). *MITgcm user manual..* MIT Department of EAPS Cambridge. URL: http://mitgcm.org/public/r2_manual/latest/online_documents/manual.html.
- Arakawa, A., & R. Lamb, V. (1977). Computational design of the basic dynamical processes of the ucla general circulation model, . *17*.
- 565 Beron-Vera, F. J., Olascoaga, M. J., Haller, G., Farazmand, M., Triñanes, J., & Wang, Y. (2015). Dissipative inertial transport patterns near coherent lagrangian eddies in the ocean. *Chaos*, *25*, 087412.
- Burden, R. L., & Faires, J. D. (1985). *Numerical analysis*. PWS Publishers.
- Chelton, D. B., Schlax, M. G., & Samelson, R. M. (2011). Global observations of nonlinear mesoscale eddies. *Prog. Oceanogr.*, *91*, 167–216.
- 570 Chelton, D. B., Schlax, M. G., Samelson, R. M., & de Szoeke, R. A. (2007). Global observations of large oceanic eddies. *Geophys. Res. Lett.*, *34*.
- Delworth, T. L., Rosati, A., Anderson, W., Adcroft, A. J., Balaji, V., Benson, R., Dixon, K., Griffies, S. M., Lee, H.-C., Pacanowski, R. C. et al. (2012). Simulated climate and climate change in the gfdl cm2.5 high-resolution coupled climate model. *J. Clim.*, *25*, 2755–2781.
- 575 Dong, C., McWilliams, J. C., Liu, Y., & Chen, D. (2014). Global heat and salt transports by eddy movement. *Nat. Commun.*, *5*, 3294.

- Dufour, C. O., Griffies, S. M., de Souza, G. F., Frenger, I., Morrison, A. K.,
580 Palter, J. B., Sarmiento, J. L., Galbraith, E. D., Dunne, J. P., Anderson,
W. G., & Slater, R. D. (2015). Role of mesoscale eddies in cross-frontal
transport of heat and biogeochemical tracers in the southern ocean. *J. Phys.*
Oceanogr., *45*, 3057–3081. doi:10.1175/JPO-D-14-0240.1.
- Dufour, C. O., Morrison, A. K., Griffies, S. M., Frenger, I., Zanowski, H., &
585 Winton, M. (2017). Preconditioning of the weddell sea polynya by the ocean
mesoscale and dense water overflows. *J. Clim.*, *30*, 7719–7737.
- dOvidio, F., De Monte, S., Della Penna, A., Cotté, C., & Guinet, C. (2013).
Ecological implications of eddy retention in the open ocean: a lagrangian
approach. *J. Phys. A*, *46*, 254023.
- 590 Frenger, I., Münnich, M., Gruber, N., & Knutti, R. (2015). Southern ocean
eddy phenomenology. *J. Geophys. Res.: Oceans*, *120*, 7413–7449.
- Froyland, G., Padberg, K., England, M. H., & Treguier, A. M. (2007). Detection
of coherent oceanic structures via transfer operators. *Phys. Rev. Lett.*, *98*,
224503.
- 595 Goddard, P. B., Dufour, C. O., Yin, J., Griffies, S. M., & Winton, M. (2017).
Co2-induced ocean warming of the antarctic continental shelf in an eddying
global climate model. *J. Geophys. Res.: Oceans*, .
- Griffies, S. M. (2012). *Elements of the Modular Ocean Model (MOM): 2013*
release (GFDL Ocean Group Technical Report No. 7). Princeton, USA:
600 NOAA/Geophysical Fluid Dynamics Laboratory. 614 + xiii pages.
- Griffies, S. M., Winton, M., Anderson, W. G., Benson, R., Delworth, T. L.,
Dufour, C. O., Dunne, J. P., Goddard, P., Morrison, A. K., Rosati, A. et al.
(2015). Impacts on ocean heat from transient mesoscale eddies in a hierarchy
of climate models. *J. Clim.*, *28*, 952–977.
- 605 Hadjighasem, A. (2016). Lagrangian-averaged-vorticity-deviation-
lavd. Github Repository. URL: <https://github.com/LCSETH/>

Lagrangian-Averaged-Vorticity-Deviation-LAVD/tree/master/2D/
src.

- Hadjighasem, A., Farazmand, M., Blazeovski, D., Froyland, G., & Haller, G.
610 (2017). A critical comparison of lagrangian methods for coherent structure
detection. *Chaos*, 27, 053104.
- Hadjighasem, A., Karrasch, D., Teramoto, H., & Haller, G. (2016). Spectral-
clustering approach to lagrangian vortex detection. *Phys. Rev. E*, 93. doi:10.
1103/physreve.93.063107.
- 615 Haller, G. (2016). Dynamic rotation and stretch tensors from a dynamic polar
decomposition. *J. Mech. and Phys. Solids*, 86, 70–93.
- Haller, G., & Beron-Vera, F. J. (2013). Coherent lagrangian vortices: the black
holes of turbulence. *J. Fluid Mech.*, 731, R4. doi:10.1017/jfm.2013.391.
- Haller, G., Hadjighasem, A., Farazmand, M., & Huhn, F. (2016). Defining
620 coherent vortices objectively from the vorticity. *J. Fluid Mech.*, 795, 136–
173.
- Haller, G., & Sapsis, T. (2008). Where do inertial particles go in fluid flows?
Physica D, 237, 573 – 583. doi:10.1016/j.physd.2007.09.027.
- Hughes, C. W., & Miller, P. I. (2017). Rapid water transport by long-lasting
625 modon eddy pairs in the southern midlatitude oceans. *Geophys. Res. Lett.*,
44. doi:10.1002/2017gl075198.
- LaCasce, J. (2008). Statistics from lagrangian observations. *Prog. Oceanogr.*,
77, 1–29.
- Lorensen, W. E., & Cline, H. E. (1987). Marching cubes: A high resolu-
630 tion 3d surface construction algorithm. *SIGGRAPH Comput. Graph.*, 21,
163–169. URL: <http://doi.acm.org/10.1145/37402.37422>. doi:10.1145/
37402.37422.

- Marshall, J., Adcroft, A., Hill, C., Perelman, L., & Heisey, C. (1997). A finite-volume, incompressible navier stokes model for studies of the ocean on parallel
635 computers. *J. Geophys. Res.: Oceans*, 102, 5753–5766.
- Maxey, M. R., & Riley, J. J. (1983). Equation of motion for a small rigid sphere in a nonuniform flow. *Phys. Fluids*, 26, 883–889. doi:10.1063/1.864230.
- McWilliams, J. C. (1984). The emergence of isolated coherent vortices in turbulent flow. *J. Fluid Mech.*, 146, 21–43.
- 640 Murray, R. J. (1996). Explicit generation of orthogonal grids for ocean models. *J. Comput. Phys.*, 126, 251–273.
- Okubo, A. (1970). Horizontal dispersion of floatable particles in the vicinity of velocity singularities such as convergences. In *Deep sea research and oceanographic abstracts* (pp. 445–454). Elsevier volume 17.
- 645 Petersen, M. R., Williams, S. J., Maltrud, M. E., Hecht, M. W., & Hamann, B. (2013). A three-dimensional eddy census of a high-resolution global ocean simulation. *J. Geophys. Res.: Oceans*, 118, 1759–1774.
- Raj, R., Johannessen, J., Eldevik, T., Nilsen, J., & Halo, I. (2016). Quantifying mesoscale eddies in the lofoten basin. *J. Geophys. Res.: Oceans*, 121, 4503–
650 4521.
- Saba, V. S., Griffies, S. M., Anderson, W. G., Winton, M., Alexander, M. A., Delworth, T. L., Hare, J. A., Harrison, M. J., Rosati, A., Vecchi, G. A. et al. (2016). Enhanced warming of the northwest atlantic ocean under climate change. *J. Geophys. Res.: Oceans*, 121, 118–132.
- 655 Van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Guillard, E., & Yu, T. (2014). scikit-image: image processing in python. *PeerJ*, 2, e453.
- Wang, Y., Beron-Vera, F., & Olascoaga, M. (2016). The life cycle of a coherent lagrangian agulhas ring. *J. Geophys. Res.: Oceans*, 121, 3944–3954.

- 660 Weiss, J. (1991). The dynamics of enstrophy transfer in two-dimensional hydro-
dynamics. *Physica D*, *48*, 273–294.
- Winton, M., Anderson, W. G., Delworth, T. L., Griffies, S. M., Hurlin,
W. J., & Rosati, A. (2014). Has coarse ocean resolution biased simula-
tions of transient climate sensitivity? *Geophys. Res. Lett.*, *41*, 8522–8529.
665 doi:10.1002/2014GL061523.
- Zhang, Y., Liu, Z., Zhao, Y., Wang, W., Li, J., & Xu, J. (2014a). Mesoscale
eddies transport deep-sea sediments. *Sci. Rep.*, *4*.
- Zhang, Z., Wang, W., & Qiu, B. (2014b). Oceanic mass transport by mesoscale
eddies. *Science*, *345*, 322–324.