

Discretization of the vorticity field of a planar jet

Natalie Ross

Department of Computer Science
University of Colorado
Boulder CO 80309-0430 USA

Jean Hertzberg

Department of Mechanical Engineering
University of Colorado
Boulder CO 80309-0427 USA

Elizabeth Bradley

Department of Computer Science
University of Colorado
Boulder CO 80309-0430 USA

Submitted to *J. Computational Physics*.

Abstract

In data assimilation, information from sensors is used to correct the state variables of a numerical model. This has been used to great advantage by the weather prediction community in the context of direct numerical simulation (DNS) models. The motivation of the work described in this paper was to extend this idea to point-vortex models, with the ultimate goal of moving to real-time simulation of a laboratory-scale vortex-dominated planar jet. This entails extracting and discretizing the vortex structures from velocity-field data in a computationally efficient fashion—i.e., using as few discrete vortices as possible to model the measured flow. This paper describes a new strategy for identifying large-scale vortex structures that augments traditional methods with a simple computational topology technique, allowing them to find the boundaries of the coherent structures in a manner that naturally follows the geometry of the flow. This strategy was applied to two standard vortex-extraction methods, vorticity thresholding and Okubo-Weiss, and tested upon velocity field data from the experimental fluid flow. The large-scale structures found in this manner were then modelled with collections of discrete vortices, and the effects of the grain size of the discretization and the parameters of the discrete-vortex model were studied. The results were evaluated by comparing the velocity field induced by the discrete vortices to that measured in the jet. These comparisons showed that the two extraction techniques were comparable in terms of sensitivity and error, suggesting that the computationally simpler vorticity thresholding method is more appropriate for this application. Comparisons of different discretization strategies showed that modelling each large-scale vortex structure with a single discrete vortex provided the best compromise between mean-squared error and computational effort.

1 Introduction

The motivation behind the work described in this paper was to use sensor data to correct the state variables of a point-vortex model of a planar jet flow. This kind of procedure, known as data assimilation, has been used extensively for many years by the numerical weather prediction community, but only recently in conjunction with point-vortex models. Because both accuracy *and speed* are important issues in correcting a running solver, especially if the goal is to move to real-time simulations, and because corrections are delivered repeatedly, the type of fine-grained discretization that is typically required for convergence, such

as [8, pp22–31] or [12], is not a good solution here. Rather, this problem demands an intelligent, easily adaptable discretization: one that models the flow with only as many discrete vortices as are demanded by the situation at hand, and one that is guided by the geometry of the flow and not solely by an arbitrary grid. The solutions described in this paper are based on two classic vortex extraction methods: vorticity thresholding and the Okubo-Weiss criterion. A computational topology technique[26, 27, 28] was incorporated into these methods, allowing them to find the vortex boundaries in a principled manner that follows the smoothness properties of the flow. These techniques were evaluated in the context of particle-image velocimetry (PIV) data from a planar air jet, using three different strategies for distributing discrete vortices to match the characteristics of each large-scale vortex. The results proved to be effective for the purposes of data assimilation into a point-vortex model[29].

Extracting vortex positions and strengths from velocity field data is important for a variety of reasons. Vortices are surprisingly hard to define, however, and experimental data are inevitably noisy, so this is not a trivial problem. Some of the methods that have been developed to work around these challenges identify vortices and distinguish them from other types of coherent structures; others simply discretize the vorticity field. The work of Jeong & Hussain[17] provides a cornerstone for much of the debate in the fluids literature about this topic, along with a useful definition that decomposes the velocity gradient tensor $\nabla\vec{v}$ into symmetric (S) and anti-symmetric (Ω) parts:

$$\begin{aligned}\Omega &= \frac{1}{2}[(\nabla\vec{u}) - (\nabla\vec{u})^T] \\ S &= \frac{1}{2}[(\nabla\vec{u}) + (\nabla\vec{u})^T]\end{aligned}$$

Ω is an effective measurement of vorticity in an incompressible flow. Thus, searching for regions in which the norm of Ω dominates the norm of S can be an effective technique for identifying vortices; this is the Q -criterion of [16]. The Okubo-Weiss method[20, 35] is a two-dimensional version of the Q -criterion that identifies regions where the squared rate of rotation, $|\Omega|^2$, dominates the squared rate of strain, $|S|^2$. In these regions, the flow behavior is classified as elliptic in nature; outside them, the fluid motion is hyperbolic. It is worth mentioning that Haller’s method [13] also identifies elliptical regions and is invariant under *any* type of coordinate change. Haller’s method sprinkles passive tracers throughout the flow, defining vortices as regions where all tracer trajectories remain on elliptic paths. This careful distinction between elliptic and hyperbolic properties is important if one wants to distinguish vortices from other vorticity-bearing structures such as shear layers, but that level of effort is not warranted in the application treated here, where the goal is simply to discretize the vorticity field. A variety of other approaches use $\nabla\vec{v}$ in different ways; [1] and [6], for instance, take its imaginary eigenvalues as evidence of local swirling motion. Some vortex-extraction research attempts to distinguish vortices from other structures that have high vorticity, such as shear layers (e.g., [23, 34]). Other methods use the geometry of the flow velocity directly, searching for foci and centers in streamline patterns[19]; minima in local-angle gradients[23]; spanwise-aligned regions that are highly correlated[30]; connecting regions where the velocities are in opposite directions[34]; and regions where the finite-time Lyapunov exponents of the flow suggest the presence of so-called Lagrangian coherent structures[33]; among others. Many groups have worked out ways to fit velocity data to various analytical forms, such as wavelets[5, 9, 32] or orthogonal and Fourier decompositions[10, 21, 24], and then use those decompositions to find the vortices. Other groups use predictor-corrector methods[2] or even neural nets[18] to find coherent structures in velocity data.

The problem treated in this paper involves a planar flow, and the goal is to compute the positions and strengths of a set of discrete vortices whose induced velocity fields match the measured flow. In contrast to some of the methods mentioned above, there is no need to distinguish between vortices and shear layers. The goal is simply to discretize the vorticity field, and in a manner that can easily be adapted to the precision requirements of different modeling situations.




Figure 1: A planar air jet: (a) Sketch of the experimental apparatus. (b) Side view of the jet flow with 16.83Hz forcing. $Re \approx 93$.

2 Methods

2.1 Apparatus

A sketch of the experiment is presented in Figure 1(a). Filtered air entered the base of a plenum, whose internal cross-section was 400mm by 15mm , and was passed through three flow treatment screens of 1mm mesh spacing with 0.2mm diameter wire. A cubic spline contraction of area ratio 6 to 1 formed the top of the plenum. Flow emerged from an exit slit 400mm long and $2.5 \pm 0.01\text{mm}$ wide, giving an aspect ratio of 160:1. The velocity profile at the jet exit was parabolic, and developed into a plane symmetric Bickley jet within a few nozzle widths downstream[22]. The peak jet velocity was 1.36m/s , corresponding to a Reynolds number of 93 based on the nozzle half-width. The downstream direction is referred to as x , transverse as y and spanwise as z . The system was highly sensitive to mechanical vibrations and motion of the ambient air. To minimize these effects, the entire experiment stood on vibration control mounts and was semi-enclosed by a 1m^3 plexiglass box. The resulting jet had a turbulence intensity, in the absence of any forcing, of less than 0.5%[22].

For the study presented here, a single external loudspeaker was used to excite the antisymmetric mode of the jet. This was accomplished by driving the loudspeaker at the natural frequency of the jet (16.83Hz), as described in [22]. A picture of the jet in its antisymmetric mode is displayed in Figure 1(b). Note that vortices appear well-defined in this flow, making it both a good candidate for point-vortex modelling and an ideal test case for the methods presented in this paper.

Velocity data was acquired using particle image velocimetry (PIV)[25]. The jet was seeded with a theater fog, consisting of a water/glycerin mixture condensed into ≈ 1 micron diameter droplets. Illumination was provided by a New Wave Solo PIV Nd:YAG laser at medium power. Images were captured using a TSI camera, model PIVCAM 13-8 with $1,280 \times 1,024$ 12-bit pixels. PIV processing was carried out with TSI Insight software. The field of view spanned 15 to 22 jet widths in the streamwise (x) direction. Data was acquired at 16 specific phases of the loudspeaker excitation, and 420 realizations were averaged for each phase. The resulting velocity field of 61×65 vectors had a uniform grid spacing of $h = 0.87\text{mm}$. The maximum speed in the velocity field analyzed in this paper was 1.16m/s . These values were used for scaling all lengths and velocities that are reported here.

2.2 PIV results

The left-hand column of Figure 2 presents velocity field data for the jet at four different phases in the flow cycle, approximately 15 milliseconds apart, with 1/4 of the vectors shown. The right-hand column of Figure 2 shows the corresponding vorticity fields, computed from the velocity field using the center difference method, with forward and backward difference at the edges of the domain, as needed. The goal of the methods described in this paper is to compute a distribution of discrete vortices that represent this field.

2.3 Analysis

The methods presented here were built upon two well-known vortex extraction techniques: vorticity thresholding and the Okubo-Weiss method. Both of these techniques were designed to identify regions of high vorticity in data like the snapshots in the right-hand column of Figure 2. A computational topology technique was incorporated into each method in order to allow them to find vortex boundaries in a manner that matches the smoothness properties of the flow. In particular, each high-vorticity region was assumed to be a *connected component*, and was found by “growing” the classification outwards from a single starting point. The method that combines vorticity thresholding and this notion of connectedness is covered in Section 2.3.1; Section 2.3.2 describes how to incorporate connectedness into Okubo-Weiss. Once these large-scale high-vorticity regions were identified, the next task was to model them with a collection of discrete vortices whose induced vorticity and velocity fields model the original PIV data. The discretization strategies used to accomplish this are described in Section 2.3.3. The evaluation process is described in Section 2.4.

2.3.1 Connected Vorticity Thresholding

The connected-component/vorticity thresholding method (hereafter “connected vorticity thresholding”) uses the following strategy to find large-scale vortices in gridded vorticity data:

1. Take the magnitude of the vorticity field and find its maximum value. Label this grid point (x_0, y_0) .
2. Let \mathbf{M} represent the set of points in a particular vortex. Initially, set $\mathbf{M} = \{(x_0, y_0)\}$. Then, starting from the point at (x_0, y_0) , construct a “connected component” of neighboring points whose vorticity magnitude is above a threshold T . For the purposes of this method, a point is considered connected to four “neighbors”—one above, one below, one to the left, and one to the right: $(x_0, y_0 + 1)$, $(x_0, y_0 - 1)$, $(x_0 - 1, y_0)$, and $(x_0 + 1, y_0)$. Check each of these points and add it to \mathbf{M} if its vorticity exceeds T . (Issues regarding the choice of T are discussed below.)
3. Repeat Step (2), recursively checking the neighbors of those neighboring points and including them in the component if their vorticity magnitudes exceed the threshold. Terminate when no new neighbor points meet the criterion. (This is the edge of the component.)
4. Compute the average vorticity over all points \mathbf{M} in the connected component. Multiply this by the area covered by those points to get an approximate circulation, or strength, for the associated large-scale vortex structure.
5. Remove the points identified in Steps (2) and (3) from consideration and repeat the process starting from Step (1) to find the next large-scale vortex.

The threshold parameter T in Step (2) controls how much vorticity is lumped into each vortex; it is intended to be used by the modeler to tune the extraction process to the requirements of his or her problem. The general challenge in isolating coherent structures with this method is to choose a T value that causes the connected component algorithm to include as much area as possible for each structure without accidentally grouping two distinct ones together, as described later in this paper. The results presented in Section 3 use $T = 0.14 | \max(\omega) - \min(\omega) |$, where $\max(\omega)$ and $\min(\omega)$ are the maximum and minimum vorticity values over all grid points in the data set. Thresholds are always arbitrary, of course, as are definitions of what




Figure 2: Velocity and vorticity fields of the planar jet under periodic forcing. Images (a)-(d) show PIV data at various phases of the drive period: (a) 22.5° (b) 90° (c) 202.5° and (d) 292.5° . Images (e)-(h) show the corresponding vorticity fields.

constitutes a “distinct” vortex, and computer interpretation of images often falls far short of what expert eyes can do. Tuning T allows a user of this method to coarsen the discretization at will and at need, in a manner that combines the features of simple thresholding and the notion of a connected component.

2.3.2 Connected Okubo-Weiss

The Okubo-Weiss criterion[20, 35] is more complex and more computationally intensive than vorticity thresholding—and arguably more reliable, in the sense that it is invariant in translating frames of reference (Galilean invariance). The original Okubo-Weiss criterion defines a vortex as any region where $W = |\Omega|^2 - |S|^2 > 0$. The method presented here, termed “connected Okubo-Weiss,” modifies that approach in two ways. First, it allows the user to specify a threshold for the inequality—e.g., to consider only points where $|\Omega|^2$ is *significantly* greater than $|S|^2$. Second, it again uses the notion of a connected component to define the boundaries of the vortices:

1. Compute $W = |\Omega|^2 - |S|^2$. Remove from consideration any grid points (a, b) for which

$$W_{(a,b)} \leq k [\max(W) - \min(W)]$$

2. Find the grid point with the maximum positive value of W .
3. Starting from the point identified in Step (2), compute the connected component—the connected set of grid points where $W > k[\max(W) - \min(W)]$ —using the same recursive neighbor-checking algorithm described in Step (3) of the algorithm in Section 2.3.1.
4. Compute the average vorticity for the points identified in Step (3). Multiply this by the area covered by those points to get an approximate circulation, or strength, for the associated vortex.
5. Remove the points identified in Step (3) from consideration and repeat the process starting with Step (2) to find the next vortex.

Like the method of Section 2.3.1, this method has a tuning parameter: the cutoff k . This parameter is used not only to filter out the points where rotation dominates strain, as in classic Okubo-Weiss, but also to ignore points where strain & rotation rates are similar—i.e., points that fall near the sharp threshold that is embodied in the Okubo-Weiss criterion. Setting $k = 0$ is equivalent to the original Okubo-Weiss mathematics, where the smallest difference between rotation and strain defines a vortex. If instead one wants to enforce a larger separation, one can set k to some small fraction of the range of W . Again, no arbitrary fixed threshold is effective for all modelling problems. The results in Section 3 use $k = 0.006$. This means that values in the lowest 0.6% of the observed W range in the field, where $|\Omega|^2$ is only slightly larger than $|S|^2$, are not considered to be conclusive evidence—one way or the other—of a vortex. Because the mathematics behind this tuning parameter are different than the simpler notion of a vorticity threshold, the reasoning involved in its tuning is somewhat different, as described later in this paper.

The Okubo-Weiss method has been shown to be inaccurate if the velocity gradient tensor is time-varying[31], and there are higher-order corrections to the method for this[14, 15]. The basic version of the criterion is used here because the additional computational cost involved in these corrections is not justified for the problem treated in this paper. In other applications, where variations in time and space dominate the modelling requirements, these accuracy issues may be important, warranting the inclusion of higher-order corrections or even different extraction techniques.

2.3.3 Vortex Discretization

Any decomposition of a vorticity field into discrete vortices is an approximation—one whose accuracy depends on the number of discrete vortices involved. And, as in any modeling problem, it involves a tradeoff: one wants to use *only* as many discrete vortices as are necessary for the requirements of the problem, distributing them such that their superposition models the larger-scale vortex structures that are present in the flow, to within some specified accuracy. The methods described here are designed with data

assimilation in mind: they preprocess static snapshots of the velocity field for use in correcting the state variables of a point-vortex solver. As such, the requirements upon them are different than in applications where perfect matching of the vorticity field is critical[3, 4]. Indeed, that kind of fine-grained discretization could pose problems for the assimilation process, even with modern acceleration techniques[8, Appendix B], since each discrete vortex adds at least three state variables¹ to the model.

In order to assess these effects and tradeoffs, three discretization strategies were evaluated:

- (i) a fine-grained method that places discrete vortices of appropriate strength at every grid point, as is common in simulations of turbulence and mixing,
- (ii) a coarse-grained distribution that places a single discrete vortex of appropriate strength at the local maximum of vorticity of each large-scale vortex, and
- (iii) a parametrized method that distributes discrete vortices of *uniform* strength across each large-scale vortex under the control of a discretization parameter, as described in the following paragraphs.

Strategies (ii) and (iii) work with the connected components found by the methods described in Sections 2.3.1 and 2.3.2; strategy (i) works with the raw vorticity field, placing a discrete vortex at every grid point and setting its strength equal to the circulation in the corresponding grid cell. Assigning a single discrete vortex per connected component is straightforward; strategy (ii) does this using the locations and strengths found in Steps (1) and (4) of the method that computed the component. The parametrized discretization strategy (iii) distributes uniform-strength discrete vortices across each connected component, as described in the steps below. The strength quantum γ_0 used in this procedure is determined by the strength Γ_{min} of the weakest vortex in the flow, via a quantization parameter q : $\gamma_0 = \Gamma_{min}/q$.

1. Find the maximum magnitude of the circulation in the connected component.
2. If there is not already a discrete vortex at the location of the maximum found in Step (1)—i.e., if this is the first time that the algorithm has examined this point—place one there, choosing the sign to match the sign of the circulation. If there is, find the closest uncovered point (by Euclidean distance) and place the discrete vortex there.
3. Subtract the quantum of this discrete vortex from the circulation at the location in Step (2) and repeat from Step (1).

The parameter q controls the grain of the discretization in the obvious way. The results in Section 3 use $q = 2$, for instance, which assigns two discrete vortices to the weakest large-scale vortex. The use of uniform strength vortices is in accordance with early random-walk vortex techniques[11], but is in contrast to modern blob methods[8]. Though this approach does not model the flow as accurately as these modern methods, it has some important advantages from the standpoint of the goals of this paper. Using a fixed quantum for vortex strength reduces the computational complexity of the model, which is key in a speed-limited application like data assimilation. Also, dividing the circulation among several discrete vortices allows the model to track vortex core deformation more effectively.

2.4 Evaluation

To evaluate how accurately a given collection of discrete vortices represents the true flow dynamics, the velocity field induced by that collection of discrete vortices was calculated and compared to the original PIV velocity field. The velocity field was reconstructed in the usual way, using Biot-Savart. The approximate diameter of viscous core of the vortices in Figure 2 is 10 times the grid spacing. To avoid the large unphysical velocities that the singularity in the point-vortex model would produce in this situation, a blob model[7] with simple solid-body rotation in the core was used to reconstruct the velocity fields. There are many more-sophisticated models with better convergence that apply in this situation; the choice here was again motivated by the need for speed that arises in a real-time data-assimilation application. For the

¹ xy position and strength




Figure 3: Vortex extraction: (a) reproduces the vorticity field from Figure 2(g). Images (b) and (c) show the large-scale vortex structures found in that data set by the connected vorticity thresholding method of Section 2.3.1 with $T = 0.14$ and the connected Okubo-Weiss method of Section 2.3.2 with $k = 0.006$, respectively. The color scale in (b) and (c) distinguishes different vortices; the small white square inside each indicates the associated vorticity maximum.

results presented in Section 3, the velocity at each grid point was calculated by summing the velocities induced there by each discrete vortex, computed using the blob model equations. The square of pointwise difference between this field and the PIV field was then calculated and averaged across all grid points. The optimal core radius r_c —i.e., the r_c that minimized the mean-squared error between measured and induced velocity fields—was found individually for each vortex configuration. The overall circulation of the induced field was also calculated.

3 Results & Discussion

3.1 Extraction

Figure 3 shows the results of applying the connected vorticity thresholding and connected Okubo-Weiss methods described in Section 2.3 to one of the vorticity fields from Figure 2, with $T = 0.14$ and $k = 0.006$, respectively. These parameter values were chosen to maximize the amount of circulation enclosed by each large-scale vortex without causing “bridges” to form between adjacent ones, and to capture comparable amounts of circulation, so that the two methods could be evaluated against one another. Both methods successfully identify regions of the flow where the vorticity is high, though their determination of the boundaries are somewhat different. The dip in vorticity on the right side of vortex in the middle of the image, for instance, gets resolved as an indentation by connected vorticity thresholding but as an enclosed hole by connected Okubo-Weiss. The total positive and negative circulations of the measured flow, calculated from the vorticity data shown in Figure 3(a), were 0.0563 and -0.0584, respectively, in normalized units². The connected components in Figure 3(b) and (c) capture only part of that amount, of course, since they only include the high-vorticity regions. The large-scale vortex structures found by connected vorticity thresholding capture 67% of the overall positive circulation and 69% of the negative circulation; those extracted by connected Okubo-Weiss capture 67% and 71%, respectively. The thresholds T and k , which these methods use to define the edge of a vortex, will obviously affect these numbers. Lower values for either parameter will increase the size of the connected components and the amount of the total circulation that they capture, but may also create bridges between neighboring structures. The mathematics behind the tuning process is slightly different in the two cases. Both are ratiometric, but they involve different scales: simple vorticity magnitude in the case of T and the difference between the magnitudes of the rate of

²Lengths and velocities are scaled as described at the end of Section 2.1.




Figure 4: Evaluation of fine-grained vortex discretization: flow fields. The image in (a) reproduces the velocity field from Figure 2(c). The image in (b) shows the velocity field induced by a configuration of blob vortices with one at every grid point of Figure 2(g) and a vortex core radius $r_c = 0.6h$.

rotation and the rate of strain in the case of k . Changing the T parameter in connected vorticity thresholding and the k parameter in connected Okubo-Weiss will thus have slightly different effects because the latter discriminates between different kinds of vorticity-bearing structures while the former simply reacts to the magnitude of ω . In the case of Figure 3, this led to the formation of bridges between like-signed vortices in (b) and *unlike*-signed vortices in (c).

3.2 Discretization

As described in Section 2.3.3, three different strategies were evaluated for distributing discrete vortices to match these flows: a fine-grained method (*i*) that placed discrete vortices of appropriate strength at every grid point, a coarse method (*ii*) that used a single discrete vortex to model each large-scale vortex, and a parametrized method (*iii*) that distributed a variable number of uniform-strength discrete vortices across each large-scale vortex. Each strategy was evaluated by comparing the induced velocity field of the vortex configuration to the PIV data.

Figure 4(b) shows the velocity field for strategy (*i*) applied to the vorticity field of Figure 2(g), with the true velocity field reproduced alongside for easy comparison. Recall that strategy (*i*) calculates the circulation in each grid cell, placing a discrete vortex with the corresponding strength at every grid point. The velocity field induced by this configuration of vortices was calculated using the blob model. The core radius (r_c) value was optimized by minimizing the mean-squared error between measured and reproduced fields, as shown in Figure 5. The optimal r_c value for this discretization was found to be on the order of the grid spacing ($r_c/h = 0.6$), a result that is consistent with the underlying mathematics. To the eye, the induced field in Figure 4(b) is a fairly good match to the measured flow in part (a). The vortices and the large-scale mean flow of the jet are similar, though there are some discrepancies within three grid rows of the boundaries because of edge effects³. These edge effects cause a 1% loss of circulation between parts (a) and (b) of the Figure. The mean-squared error between the two fields is $0.0094 \text{ m}^2/\text{s}^2$, which is 0.7% of the square of the maximum velocity in the field.

Figure 6 shows the velocity fields produced by the coarse-grained discretization—strategy (*ii*)—of the connected components in Figure 3. Part (a) of the figure again shows the PIV data for comparison. Parts (b) and (c) show the velocity fields that are induced by single discrete vortices placed at the high- $|\omega|$ seed point of each connected component found by vorticity thresholding and Okubo-Weiss in that PIV data, respectively. These fields were again computed via the blob model with the vortex core radius value that

³i.e., the effects of the vortex structures that are outside the field of observation, which factor into the velocities of Figure 4(a) but not Figure 4(b).




Figure 5: Mean-squared error in induced velocity field of fine-grained vortex discretization as a function of r_c , expressed as a percentage of the square of the maximum velocity.

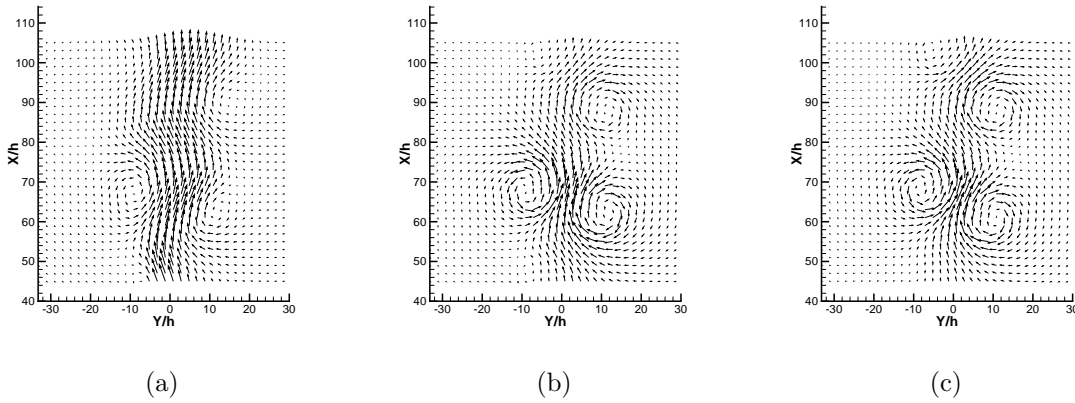


Figure 6: Evaluation of coarse-grained vortex discretization: flow fields. The image in (a) reproduces the velocity field from Figure 2(c). Images (b) and (c) are the velocity fields induced by coarse-grained discretizations of Figures 3(b) and (c), respectively: i.e., a single equivalent-strength blob vortex placed in each connected component. r_c was chosen to minimize the mean-squared error for each case, yielding $r_c = 5h$ in (b) and $r_c = 5.8h$ in (c).




Figure 7: Parametrized vortex discretization: the image in (a) reproduces the vorticity field from Figure 2(g). The grey-scale regions in (b) and (c) show the vortices found in that data set by connected vorticity thresholding and connected Okubo-Weiss, respectively, as in Figure 3. The small white squares inside each of these connected components indicate where the discretization strategy on page 7 placed discrete vortices: 23 in (b) and 46 in (c). The strength quantum chosen for these vortices was chosen to be half of the strength of the weakest vortex in each case (i.e., $q = 2$).

minimized the mean-squared error between measured and induced fields: $r_c = 5h$ for connected vorticity thresholding and $r_c = 5.8h$ for connected Okubo-Weiss. If the extraction and model were perfect, of course, the fields in Figures 6(b) and (c) would be identical to the original velocity field in (a). While both configurations do indeed capture the basic structure of the flow, the induced fields show some error in the upstream regions, where the original vortices are elongated and thus less well-described by the blob model. This disparity is not surprising; reducing the vorticity field of a real-world flow to five discrete vortices is a radical approximation—and one that completely distorts the vorticity distribution by forcing the cores to be circular. The mean-squared errors between the induced fields of Figure 6(b) and (c) and the measured field of (a) were 25% and 21%, respectively. These values are of course significantly larger than the error in the fine-grained discretization of Figure 4(b), where the field was modelled by almost 800 times as many vortices.

Increasing the fidelity of the modelling approximation—using better models of each vortex element, increasing the number of discrete vortices, and/or tailoring their strengths and positions to the flow—will generally improve its accuracy, but at a computational cost that can be prohibitive in a data-assimilation application, where correction frequency and hence data-processing speed are critical. The parametrized approach described in Section 2.3.3—strategy (iii) on page 7—is designed to allow the discretization process to be tuned to fit these requirements. Figure 7 shows how this strategy discretizes the large-scale vortices of Figures 3(b) and (c) with a q value of two (i.e., a vortex strength quantum equal to half of the strength of the weakest vortex in the field). The numbers of discrete vortices used to discretize each connected component reflect the corresponding vortex strengths: $\{2, 4, 7, 8, 2\}$ small vortices in Figure 7(b), from top to bottom, and $\{6, 10, 14, 14, 2\}$ in (c). To evaluate these results, the induced velocity fields were computed for the vortex configurations of Figure 7(b) and (c), again using the blob model with r_c chosen to minimize the mean-squared error in each case, then compared to the PIV velocity field. Visually, these induced fields are indistinguishable from Figures 6(b) and (c) and so are not shown. The mean-squared differences between these fields and the measured field were 25% and 19%, respectively. $r_c = 0.5h$ minimized the mean-squared error in both cases; the effects of changing r_c and q are shown in Figure 8. Overall, the Okubo-Weiss results are slightly better, and there is some improvement in both methods with increasing q , though not as much as one might expect. This is discussed further below. Note that the error values with $q = 1$ are slightly worse than the coarse discretization results of Figure 6, even though the $q = 1$ discretization

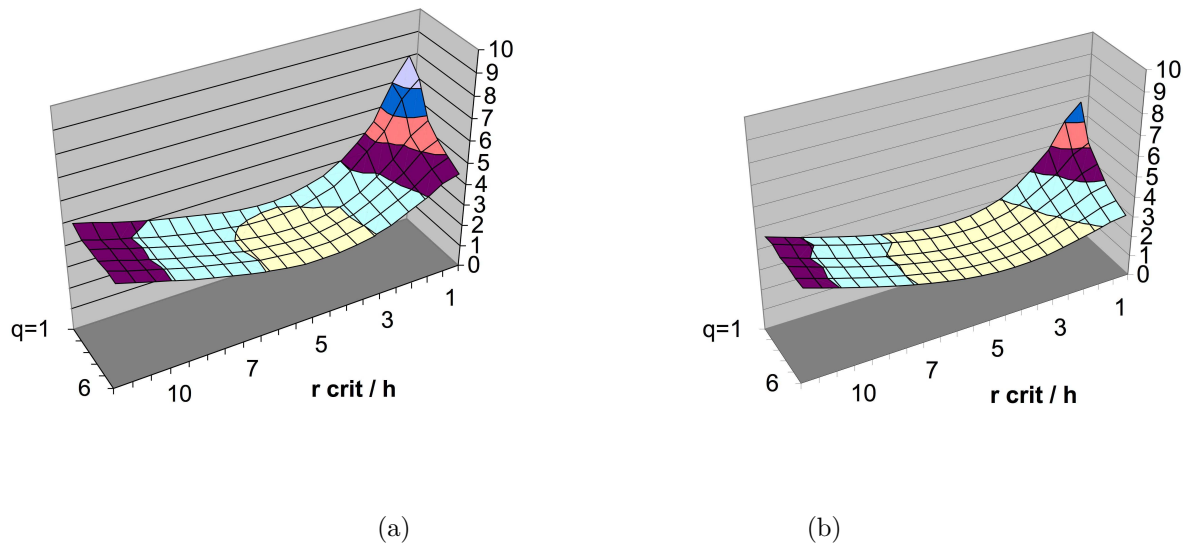


Figure 8: Effects of quantization parameter q and vortex core radius r_c upon mean-squared error of results, expressed in %: (a) connected vorticity thresholding (b) connected Okubo-Weiss.

models the flow with more than twice as many vortices⁴. These additional vortices, though, are clustered around the peak vorticity location and primarily within the core radius, which effectively increases the core radius. At $q = 6$, where the discretization contains 69 and 141 vortices for vorticity thresholding and Okubo-Weiss, respectively, the mean-squared error of the parametrized discretization falls below that of the coarse discretization for both methods.

3.3 Discussion

A full comparison of the two extraction methods across all discretization strategies yields some interesting insights. In terms of the accuracy of the reconstructed velocity fields, there is little difference between connected vorticity thresholding and connected Okubo-Weiss, in either coarse or parametrized discretizations. The mean-squared error and percentage of circulation captured by each are comparable. *For this flow, then, simple vorticity thresholding is just as good as the much more computationally expensive Okubo-Weiss method at extracting regions of high vorticity if both are augmented with connectedness techniques from computational topology to find the boundaries of those regions.*

Comparing discretization strategies is also interesting—and somewhat surprising at first glance, in that using a finer-grained discretization (viz., higher q) did not significantly improve the accuracy of the reconstructed field—for either extraction method. Recall that increasing the discretization parameter (q) in strategy (iii) reduces the size of the discrete vortices that are used to model each large-scale vortex. Clustering additional vortices of uniform value around the local vorticity peaks, however, serves only to diffuse the impact of the vortex on the velocity field—in effect, merely increasing r_c . Note that this is not equivalent, in the limit of large q , to the discretization of strategy (i), which places one discrete vortex at *each grid point*, with strength computed to match the local vorticity. Because that strategy distributes vortices throughout the flow, rather than clustering them near the peaks of the large-scale vortices, and because it tailors their strengths to the data in each region, it does provide increased resolution—particularly in terms of capturing the large-scale flow and the less-regular vortex structures. In terms of accuracy at one instant in time, then, *it does not seem worth the effort to quantize the vorticity via uniform-strength*

⁴11 and 23 vortices for the large-scale structures found by vorticity thresholding and Okubo-Weiss, respectively, compared to five vortices for the coarse discretization.

vortices clustered around the peak of large-scale vortex. A finer discretization will of course allow a model to track the deformation of the cores more effectively; this is discussed in the following section. For the purposes of crafting a traditional model that converges quickly and tracks the flow accurately over time, the traditional full gridded approach is better. For the purposes of data assimilation, however, the requirements are completely different: the measurement and correction dynamics dominate the data-assimilation process, and the preprocessing step must balance speed and accuracy. The methods described here are designed to do exactly that, in an explicit way, via the T , k , and q parameters in the methods in Sections 2.3.1, 2.3.2, and 2.3.3.

4 Conclusion

The vortex extraction and discretization techniques described in this paper were developed for use in data assimilation algorithms, which use information from sensors to correct the state variables of a numerical solver *while it runs*. As such, speed is critical, careful attention to convergence is not necessary, and there is no need to distinguish between vortices and other kinds of vorticity-bearing structures like shear layers. Rather, this application demands a fast, easily adaptable extraction/discretization strategy that models the flow with as few or as many discrete vortices as are demanded by the situation at hand. The methods described in this paper accomplish this by augmenting classic vortex-extraction methods with a simple computational topology technique, defining a single vortex by searching for a connected component of high vorticity. Each of these large-scale structures is then modeled using a collection of uniform-strength discrete vortices distributed around its vorticity peak. These extraction and discretization methods have tuning parameters that allow a user to tailor their actions to the requirements of a particular problem. The threshold parameters T and k control how much vorticity is lumped into each large-scale structure, while the discretization parameter q specifies the strength of the discrete vortices that are used to model each of those structures.

The findings in the previous sections show that the connected vorticity thresholding and connected Okubo-Weiss methods are comparably effective in extracting the large-scale vortex structures in a flow. As the corresponding threshold parameters are varied, each method captures different amounts of the total vorticity of the field, causing the structures to shrink or grow. The results presented here use values for these parameters that maximize the circulation captured without causing “bridging” between neighboring structures. At these values, the circulation captured by the two methods was similar, as was the mean-squared error between the measured and induced flows. In view of the greater computational complexity of Okubo-Weiss, then, simple vorticity thresholding is preferable for the purposes of data assimilation in the flow considered here. More-complex flows may pose different challenges, of course, and the methods here will likely not scale to fully turbulent flows—though they may apply to their smaller scales.

In terms of discretization strategies, the coarse approach appears to provide the best compromise between accuracy and computational complexity. The series of experiments reported in the previous section used between five and 141 vortices to discretize the five large-scale vortex structures, exploring a range of $1 < q < 6$ for the parametrized strategy (iii) and comparing it to the coarse strategy (ii). The latter had lower mean-squared error than all but the finest of the parametrized discretizations ($q = 6$, which uses 69 and 141 discrete vortices, respectively, to model the large-scale vortices found by vorticity thresholding and Okubo-Weiss). In other words, it is better to model each large-scale vortex with a single appropriate-strength discrete vortex placed at its peak, rather than to cluster smaller uniform vortices near that peak—up to a point. This suggests that the coarse discretization strategy is better suited for a data-assimilation application, as it provides the best combination of low error and low computational complexity of the methods studied here.

In general, a finer-grained discretization produces a more-realistic vorticity distribution, but the results here are not consistent with that generalization. This is likely at least partially a consequence of the clustered nature of the distribution used in strategy (iii). The obvious next steps in the future-work plan for this project are to distribute the discrete vortices more realistically across the large-scale vortex structures and/or use non-uniform strengths. This would of course add computational complexity, though, and without full testing in the context of a data-assimilation system, it is difficult to estimate the impact

of these modifications.

A finer-grained discretization also has implications for *running* the model, as it allows that model to advect the vorticity more accurately over time. Here, too, accuracy/complexity tradeoffs must be considered. The standard discretization approach in the fluids literature is to place one discrete vortex at each grid point in the vorticity field and make the grid as dense as feasible. This is one end of the modeling spectrum, where the grain of the approximation is very fine and the model contains a large number of state variables. If one is interested in speed, however, this may not be desirable. The fine-grained discretization of Figure 4(b) has roughly three million state variables⁵, which makes it run tens of thousands times more slowly than the parametrized discretization of Figure 7 and a factor of ≈ 200000 more slowly than the coarse-grained discretization of Figure 6. The discretization strategy described in this paper allows one to optimize between accuracy and computational cost for a given application, choosing the appropriate point on the modelling spectrum via the q parameter in the algorithm on page 7.

Acknowledgements: the authors wish to thank Jeffrey Anderson, Matthew Culbreth, Nathan Farrell, and Mark Rast for their input to this project.

References

- [1] R. Adrian, K. Christensen, and Z. Liu. Analysis and interpretation of instantaneous turbulent velocity fields. *Experiments in Fluids*, 29:275–290, 2000.
- [2] D. Banks and B. Singer. A predictor-corrector technique for visualizing unsteady flow. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):151–163, 1995.
- [3] L. Barba. Spectral-like accuracy in space of a meshless vortex method. In *Advances in Meshfree Techniques*. Springer, 2007.
- [4] L. Barba, A. Leonard, and C. Allen. Advances in viscous vortex methods—Meshless spatial adaption based on radial basis function interpolation. *International Journal for Numerical Methods in Fluids*, 47:387–421, 2005.
- [5] R. Camussi. Coherent structure identification from wavelet analysis of particle image velocimetry data. *Experiments in Fluids*, 32:76–86, 2002.
- [6] M. Chong, A. Perry, and B. Cantwell. A general classification of three-dimensional flow fields. *Physics of Fluids*, 2:765–777, 1990.
- [7] A. Chorin. Numerical study of slightly viscous flow. *Journal of Fluid Mechanics*, 57:785–796, 1973.
- [8] G.-H. Cottet and P. Koumoutsakos. *Vortex Methods: Theory and Practice*. Cambridge University Press, 2000.
- [9] M. Farge, K. Schneider, and N. Kevlahan. Non-gaussianity and coherent vortex simulation for two-dimensional turbulence using an adaptive orthogonal wavelet basis. *Physics of Fluids*, 11(8):2187–2201, August 1999.
- [10] M. Farge, K. Schneider, G. Pellegrino, A. Wray, and R. Rogallo. Coherent vortex extraction in three-dimensional homogeneous turbulence: Comparison between CVS-wavelet and POD-Fourier decompositions. *Physics of Fluids*, 15(10):2886–2896, October 2003.
- [11] K. Gustafson and J. Sethian. *Vortex Methods and Vortex Motion*. SIAM, 1991.
- [12] O. Hald. *Convergence of Vortex Methods*. SIAM, 1991.
- [13] G. Haller. An objective definition of a vortex. *Journal of Fluid Mechanics*, 525:1–26, 2005.

⁵Discrete-vortex positions & strengths at 1024^2 grid points

- [14] B. Hua and P. Klein. An exact criterion for the stirring properties of nearly two-dimensional turbulence. *Physica D*, 113:98–110, 1998.
- [15] B. L. Hua, J. McWilliams, and P. Klein. Lagrangian accelerations in geostrophic turbulence. *Journal of Fluid Mechanics*, 366:87–108, 1998.
- [16] J. Hunt, A. Wray, and P. Moin. Eddies, streams, and convergence zones in turbulent flows. In *Proc. of the Summer Program*, volume CTR-S88, pages 193–208, 1988. Stanford University Center for Turbulence Research Report.
- [17] J. Jeong and F. Hussain. On the identification of a vortex. *J. Fluid. Mech.*, 285:69–94, 1995.
- [18] R. Joseph, S. Viglione, and H. Wolf. Cloud pattern recognition. In *Proceedings of the 1964 19th ACM National Conference*, pages 42.301–42.3017, 1964.
- [19] H. Lugt. The dilemma of defining a vortex. In U. Muller, K. Roesner, and B. Schmidt, editors, *Recent Developments in Theoretical and Experimental Fluid Mechanics*, pages 309–321. Springer, 1979.
- [20] A. Okubo. Horizontal dispersion of floatable trajectories in the vicinity of velocity singularities such as convergencies. *Deep Sea Research*, 17:445–454, 1970.
- [21] A. Palacios, D. Armbruster, E. Kostelich, and E. Stone. Analyzing the dynamics of cellular flames. *Physica D*, 96(1-4):132–161, 1996.
- [22] T. Peacock, J. Hertzberg, Y.-C. Lee, and E. Bradley. Forcing a planar jet flow using MEMS. *Experiments in Fluids*, 37:22–28, 2004.
- [23] R. Pemberton, S. Turnock, T. Dodd, and E. Rogers. A novel method for identifying vortical structures. *Journal of Fluids & Structures*, 16:1051–1057, 2002.
- [24] R. Preisendorfer. *Principal Component Analysis in Meteorology and Oceanography*. Elsevier, 1988.
- [25] M. Raffel, C. Willert, and J. Kompenhans. *Particle Image Velocimetry: A Practical Guide*. Springer, 1998.
- [26] V. Robins, J. Abernethy, N. Rooney, and E. Bradley. Topology and intelligent data analysis. *Intelligent Data Analysis*, 8:505–515, 2004.
- [27] V. Robins, J. Meiss, and E. Bradley. Computing connectedness: An exercise in computational topology. *Nonlinearity*, 11:913–922, 1998.
- [28] V. Robins, N. Rooney, and E. Bradley. Topology-based signal separation. *Chaos*, 14:305–316, 2004.
- [29] N. Ross. *Understanding the Dynamics of Point-Vortex Data Assimilation*. PhD thesis, University of Colorado, 2008.
- [30] F. Scarano and C. B. M. Riethmuller. Pattern recognition analysis of the turbulent flow past a backward facing step. *Physics of Fluids*, 11(12):3808–3818, December 1999.
- [31] H. Segur. Evolution of a tracer gradient in an incompressible, two-dimensional flow. In *IUTAM Symposium on Developments in Geophysical Turbulence*, 1998.
- [32] A. Seigel and J. Weiss. A wavelet-packet census algorithm for calculating vortex statistics. *Physics of Fluids*, 9(7):1988–1999, July 1997.
- [33] S. Shadden, F. Lekien, and J. Marsden. Definition and properties of lagrangian coherent structures from finite-time lyapunov exponents in two-dimensional aperiodic flows. *Physica D*, 212:271–304, 2005.
- [34] H. Vollmers. Detection of vortices and quantitative evaluation of their main parameters from experimental velocity data. *Measurement Science and Technology*, 12:1199–1207, 2001.
- [35] J. Weiss. The dynamics of enstrophy transfer in 2-dimensional hydrodynamics. *Physica D*, 48:273–294, 1991.