

Exploration of salt wedge dynamics in the Columbia River Estuary using optical measurements of internal ship wakes

Sam Greydanus, REU Intern
Rob Holman, SECNAV/CNO Chair in Oceanography
College of Earth, Ocean, and Atmospheric Sciences, Oregon State University, United States

Dated: August 19, 2014

Abstract

In May of 2013 and beyond, Argus optical measurements of the mouth of the Columbia River (MCR) estuary and plume were collected as part of the RIVET II multi-investigator field experiment. One surprise was the strength of eddy and internal wave signatures observed in movies computed from one-minute averages of high-frequency snapshots (such that gravity waves were averaged out but slicks and variable surface roughness remained). In particular, passing ships left wakes that propagated away at speeds on the order of 0.5 m/s, much slower than gravity waves. They are presumably surface manifestations of internal waves associated with the time-varying salt-wedge. Thus, these internal ship wakes appear to act as probes of internal stratification dynamics. This paper will explore the time variations of these internal waves and relate them to corresponding variations in the estuary salt wedge.

I Introduction

I.I Importance of the MCR

The Columbia River is the largest river in the Pacific Northwest and serves as an artery for commerce and transportation into the region. The Cape Disappointment area at the river mouth is also among the most dangerous regions in the Pacific Northwest, known for its unpredictable currents and massive waves. More than 2,000 vessels and 1,500 people have been lost since the 1790's [Wong, 2002]. Understanding the oceanography of this estuary yields benefits for both commerce and general safety.

I.II Oceanography of the MCR

The mouth of the Columbia River is a dynamic and complex system. Huge waves tend to form at the river mouth where oceangoing waves meet the tidal current and steepen significantly. At the same time, saltwater and freshwater engage in a complicated two-layer exchange to form a 'salt wedge' which moves up and down the estuary according to the tides. The water masses which form this wedge travel large distances - up to two kilometers - over the course of each tidal cycle. The MCR is a large and hostile place for sampling, so little is known of the specifics of these phenomena.

I.III Internal waves

Beginning in May of 2013, the RIVET II multi-investigator field experiment collected Argus optical measurements at the mouth of the Columbia River estuary and plume. Processed to form wave-averaged movies (meaning that any signature moving at the speed of a surface wave was averaged out), these measurements evinced some surprising features. First, the behavior of currents and tidal fronts was far more complex than expected. Second, vessels passing through the estuary frequently left strong wake signatures. These signatures moved slowly, usually with celerities less than a meter per second, consistent with known celerities of internal waves.

Large vessels are known to generate internal waves [Watson et al., 1992] but these are the first observed at the mouth of the Columbia River. The theoretical model of internal waves shows that their celerities are affected by the comparative densities and heights of the two mediums between which they propagate. With this knowledge we hypothesize that the time-varying structure of the estuary's salt wedge is related to the celerities of these internal waves.

I.IV Theory

Manual measurements of ship wake signatures yielded celerity estimates in the range of 0.5 to 1.0 m/s. There were several types of waves which could be responsible for these signatures but only internal wakes propagate with celerities on this order of magnitude.

In general, the behavior of a surface wave is given as:

$$\sigma^2 = (gk) \tanh(kh) \quad (1)$$

where h is the depth of the vertical water column and k is wave number. We can relate k to wavelength by the expression $k = \frac{2\pi}{L}$.

Waves entering the mouth of the Columbia River from the ocean behave as shallow-water waves. Waves earn this classification if $h/L < \frac{1}{20}$. The estuary is dredged to a depth of 12 meters and the wavelength of oceangoing waves is on the order of 160 meters, so $h/L \leq 0.075$, which is almost shallow and we can make the approximation. If we substitute the shallow-water approximation $\tanh(kh) \approx kh$ into Equation 1, we find that $c_s = \frac{\sigma}{k} = \sqrt{gh}$ m/s, so the minimum celerity of the shallow-water waves on the MCR is $\sqrt{9.8 \cdot 12} = 10.8$ m/s. Waves moving at these speeds would have been averaged out from our wave-averaged movies.

Vessels passing through the MCR also generate waves. These waves have wavelengths on the order of 16 meters and thus behave as deep-water waves. If we substitute the approximation $\tanh(kh) \approx 1$ into Equation 1, we find that $c_s = \frac{\sigma}{k} = \frac{gT}{2\pi}$ m/s where T is the period of the wave. We can solve for T using the relationship $L = \frac{gT^2}{2\pi}$, giving us $T = 3.2$ seconds $c_s \approx 5$ m/s. Waves moving at 5 m/s would also be averaged out.

Finally, we should consider internal waves. Figure (1) shows how internal waves occur when there are two layers of water of different densities, one atop the other.




Fig. 1: Internal waves between layers of saltwater and freshwater in an estuary. Taken from [Schilling, 2009]

This situation can be expressed mathematically by Equation 2. We assign the density of the upper layer to ρ_1 and the density of the lower layer to ρ_2 . We let h_1 and h_2 be the heights of the two layers respectively. Finally, k and ω are internal wave number and frequency and g is acceleration due to gravity.

$$\rho_1 \frac{\omega^2}{k} \frac{1}{\tanh(kh_1)} + \rho_2 \frac{\omega^2}{k} \frac{1}{\tanh(kh_2)} - g(\rho_2 - \rho_1) = 0 \quad (2)$$

Once again we can make the approximation $\tanh(kh) \approx kh$ because, in the few ship wakes which create two internal waves, they are on the order of 150 meters apart and they occur at values of $h < 12$, so $h/L < 0.08$, putting them in a range where shallow water approximations are valid.

$$\rho_1 \frac{\omega^2}{h_1 k^2} + \rho_2 \frac{\omega^2}{h_2 k^2} = g(\rho_2 - \rho_1) \quad (3)$$

Now let the celerity of the internal wave be c_s . Recall that $c_s = \frac{\omega}{k}$. This gives us

$$c_s^2 \left(\frac{\rho_1}{h_1} + \frac{\rho_2}{h_2} \right) = g(\rho_2 - \rho_1) \quad (4)$$

Some algebra yields an equation for c_s as a function of the layer heights and densities.

$$c_s^2 = \frac{g(\rho_2 - \rho_1) \cdot h_1 h_2}{\rho_1 h_2 + \rho_2 h_1} \quad (5)$$

This mathematical relationship suggests that the heights of the salt and freshwater layers in the Columbia River estuary will impact the celerity of any internal waves. During high tide the salt wedge moves up the mouth of the river and during low tide the salt wedge retreats. We expect to see a strong correlation between tide level and internal wave celerity. In particular, as h_1/h_2 approaches 1, celerity will increase.

The Columbia River mouth is dredged so that its minimum depth is 12 meters. We can say, then, that $h_1 + h_2 \approx 12$. For our estimates, we will let h_1 vary between 2 and 6 meters. Ocean water has an average density of 1035 kg/m^3 , but near an estuary it is closer to 1030. A reasonable value for the density of estuary water is 1010 kg/m^3 (varies between 1000 and 1020 kg/m^3). Putting these values through Equation 5 and using $g = 9.8$, we predict internal wave celerity in the range $[0.57, 0.76] \text{ m/s}$.

These predictions fit our manual estimates and support our hypothesis that the wake signatures are caused by the propagation of internal waves.

I.V Goals

To test our predictions we will take experimental measurements of internal waves on the MCR. We will then seek a consistent relationship between wake celerity and the position of the salt wedge in the estuary. If there is a plausible connection we will analyze our data to determine what it can reveal about the internal structure of the Columbia River estuary. In particular, we will consider the variation in celerity over time, tide stage, and position in the estuary. Finally, we will discuss the believability of our data and determine whether any correlations are strong enough to use as predictions.

II Methods

II.I Field site

Wave-averaged frames were collected between May and July 2013 from the Cape Disappointment Argus station at the mouth of the Columbia River. After merging and rectification (discussed below), each image was translated to world coordinates such that every pixel covered a 5×5 meter region. Axes were chosen as shown in Figure 2




Fig. 2: Mapping images to world coordinates

When working with image data in MATLAB we always worked in world coordinates (ie meters per second rather than pixels per second).

II.II Argus

Argus stations generally return spatial or temporal samples taken from the total stream of raw camera data. A temporal sample might return the intensity values of one pixel for all time intervals and might be used to find the period of waves at a particular place along the shore [Holman and Stanley, 2007]. A spatial sample might return the average intensities of all pixels over a short span of time and be used to estimate the locations where waves are most likely to break. In the case of the Cape Disappointment Argus station we saved the following image products for every hour of daylight:

1. a single snapshot
2. a 10-minute time exposure
3. a 10-minute variance image
4. a set of wave-averaged images (averaged over 60 seconds) sampled every 30 seconds for the first 17 minutes of every hour

The set of wave-averaged images were used for quantitative analysis of wave celerity whereas the other image products provided qualitative context.

II.III Merging camera views

The mouth of the Columbia River is over two miles wide [Tate, 2006]. To image the entire region with acceptable resolution we merged four separate camera views using mathematical techniques from projective geometry. In essence, we created a fifth camera perspective (a mathematical construction) and then projected each image from each camera into the constructed perspective. Then we merged the camera views into a single continuous image (process taken from [Pennucci et al.,]).


To provide an example, we implemented this process and used it to project and merge three images from Yaquina Head, Oregon.



(a) Camera 1



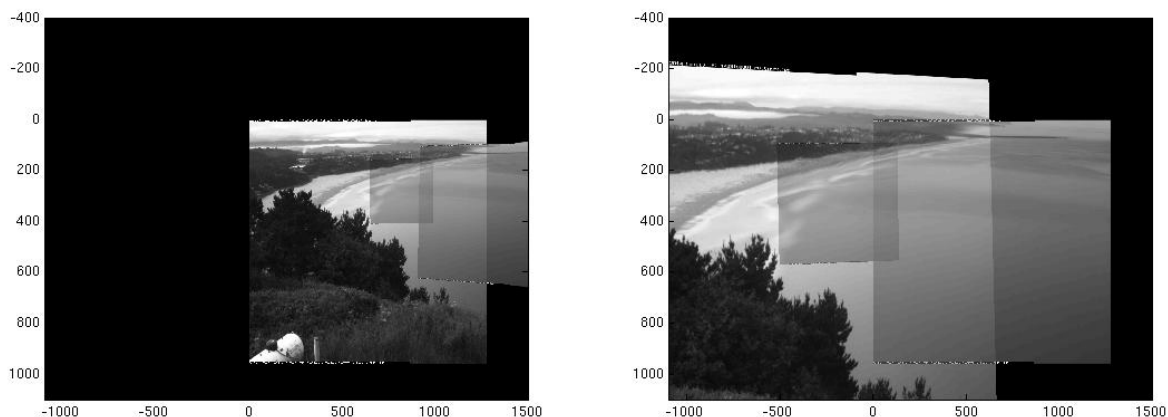
(b) Camera 2



(c) Camera 3

Fig. 3: Original camera perspectives collected by the Argus station at Yaquina Head, Oregon

Shown in Figure 3 are the three separate camera views. Notice that they are taken by different cameras, from different perspectives, and with different degrees of zoom. Using techniques from projective geometry (discussed in the Appendix) we were able to merge them into the continuous camera views shown in Figure 4.



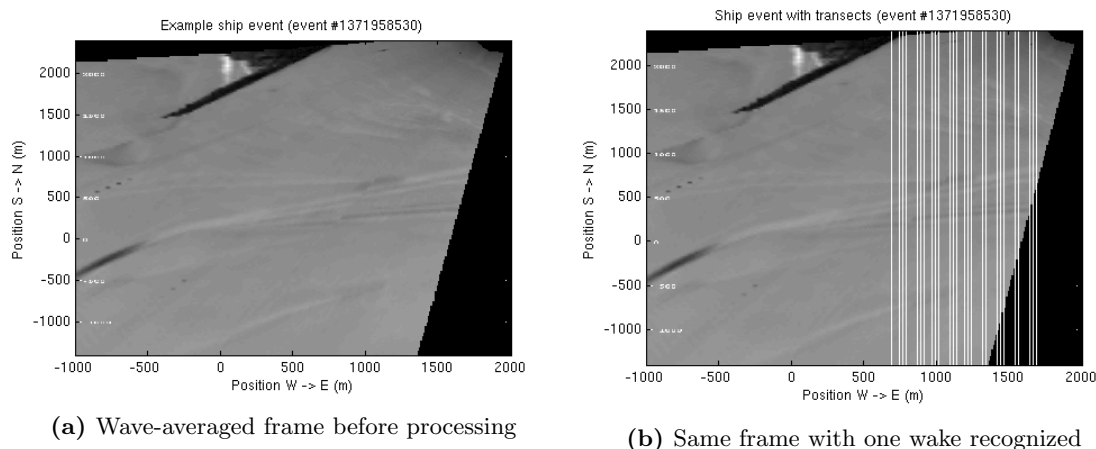
(a) All camera views merged into Camera 1's perspective (b) All camera views merged into another perspective

Fig. 4: Projections of original camera perspectives into single perspectives

II.IV Wake analysis objectives

For each ship event there are two wakes which we must analyze. Ships generally move along the x axis (east-west or west-east directions), so one of the wakes propagates northward and the other propagates southward. The wakes appear in the wave-averaged frames as linear contrast features which propagate at roughly constant velocities. Using only these criteria, we should be able to recognize and track each wake. We chose to process each of the two wakes separately; an example wake signature is shown in Figure 5a and successful recognition of the wake is shown in Figure 5b.

For each of these wakes, our goal is to plot celerity at many different x and y locations. Once we have a grid of celerity vectors we can search for either spatial or temporal patterns. Furthermore, we can compare a large number of these wakes to tidal patterns.



(a) Wave-averaged frame before processing

(b) Same frame with one wake recognized

Fig. 5: Signal processing objective: before and after

Automated wake recognition and measurement is essential under these circumstances because we need

a large number of celerity estimates for every wake and we need confidence intervals. If a user were to manually record wake celerities, they would not be able to take as dense a number of measurements and their confidence estimates would be subjective. Hence we developed an algorithm which takes in between three and thirty (roughly) frames containing a ship event and returns hundreds of celerity magnitudes, each associated with a position, and error estimate. The following section will give a sketch of the principles of this algorithm.

II.V Wake analysis

II.V.1 Preparation

In preparation for analysis, we manually searched through the database of wave-averaged frames and saved the beginning and ending time stamps, t_i and t_f respectively for every ship event. We defined a ship event as a series of wave-averaged frames during which a large (ie cargo) vessel passes through the MCR and creates a measurable wake. With the aid of a MATLAB Graphical User Interface, we sorted through over 60,000 frames and labeled 126 unique ship events. The average ship event contained roughly fifteen frames. For each of the ship wakes, we also tagged the side of the frames (east or west) which should be analyzed.

Next, we generated vectors t , x , and y for time (seconds), west-east position (meters), and south-north position (meters) respectively. These we used to convert pixel data to world coordinates. With the set of two-dimensional images stacked according to time, we obtained a three-dimensional matrix I containing all of the raw pixel information for the ship event.

Finally, we chose the vertical transects along which we would solve for vertical wake celerity. For the purposes of our analysis we chose the set of transects x_{pos} at 20 meter intervals. Henceforth we will let x_{pos} represent the vector of transect positions along the x axis and x_s represent any single transect location along the x axis.

II.V.2 Time stacks

Our objective for this step was to represent the way a particular column of pixels changed over time. In essence, a time stack is a stack of pixels wherein each row is a column taken from a different consecutive frame in the ship event. We can express this in terms of I and x_s as $S = I_{:,x_s,:}$, where S is the 2D matrix of the pixel intensities of a given time stack and the MATLAB convention ":" represents a selection of every index in a given dimension.

Ship wakes often appeared in time stacks as dark slanted lines such as the ones shown in Figure 6. The vertical axis labels the time (in seconds) at which each row was taken relative to the first frame, $I_{:,1}$, of the ship event.

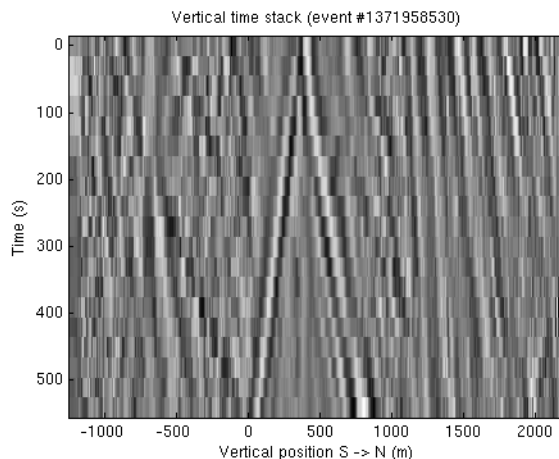


Fig. 6: An example time stack

Before creating time stacks we applied a high-pass filter to the images. Afterwards we removed columns which contained only black pixels and applied a column-wise normalization. We created one time stack for every transect. We took forty transects for every ship wake, giving a total of forty time stacks for each ship event.

II.V.3 Celerity stacks

Our objective for this step was to search a segment of a time stack for contrast features moving at some constant celerity c_y . We accomplished this by generating a vector c_{pos} of possible y-component celerity values (in m/s). Next, we selected several subsections of the time stack for analysis. Mathematically, $S_{sub} = S_{a-b}$, where a is the start time and b is the stop time in seconds. For each possible celerity value c_y , we shifted the rows of the subsection so that signatures moving at that celerity would appear as vertical lines. We then took a columnwise sum of the shifted S_{sub} . The resulting row vector had highest values for y positions where contrast signatures moving at c_y m/s occurred. We stacked these row vectors atop each other according to their c_y values to create a celerity stack C .

The general conversion from time stack to celerity stack is as follows:

$$S(y_s, c_y) = I(y, y_s - c_s(t_s - t_o)) \tag{6}$$

where t_o is the initial time of the ship event. This relationship does not take into account the transformations in pixel intensities described above; it illustrates the spatial rearrangement of pixels.




Fig. 7: An example celerity stack derived from the time stack shown in Figure 6. The two brightest lines at $y_s = 100$ and $y_s = 700$ are caused by the two wake signatures in Figure 6

We created approximately ten celerity stacks from every time stack, resulting in approximately four hundred celerity stacks per ship event. For every time stack we also corrected time and position differences and merged the ten celerity stacks into a single averaged celerity stack, C_{avg} . The averaged celerity stacks were useful because only the peaks moving at constant velocities aligned with each other. Averaged celerity stacks further reduced background noise and enabled us to choose accurate y_s and c_y seeds for peak fitting. Notice that the two peaks of the averaged celerity stack in Figure 8 are stronger and more localized than the two peaks in the example celerity stack in Figure 7.

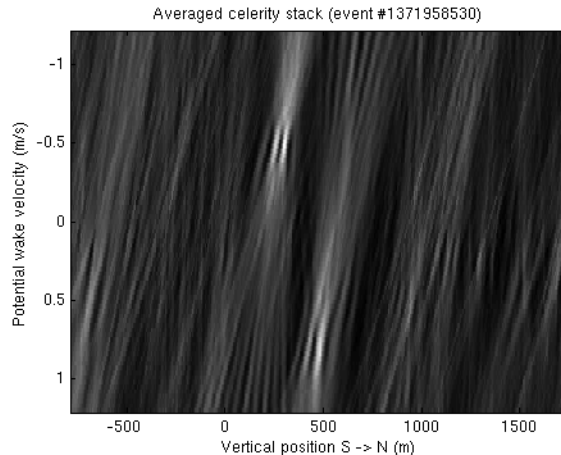


Fig. 8: An average celerity stack. The two peaks occur at approximately (350, -0.6) and (450, 0.6), meaning that the two wakes are both moving at 0.6 m/s and are located at $y = 350$ and $y = 450$ m

II.V.4 Hanning window and Gaussian peak fitting

Our next goal was to automatically recognize the coordinates of the peaks in the celerity stacks for both positive and negative celerities. The simplest way is to simply return the coordinates of the pixel with the highest intensity for both the positive and negative sides. This technique has a very high error rate and does not return an estimate of confidence. Instead, we elected to use one of MATLAB’s built-in optimization functions to fit a 2D Gaussian to the peaks. Often there were several competing peaks and the best overall peak was usually not the strongest peak in every consecutive stack. To select one wake consistently, we used basic kinematics to predict the approximate y -position of the wake and then narrowed our search to that region in particular.

First we generated average celerity stacks for the first five transects in the ship event. Next we found the average locations of the peaks on both positive and negative sides of the averaged celerity stacks. The result was two sets of coordinates or ”seeds”: one for the southerly wake and another for the northerly wake. Using the first five wakes helped to eliminate outliers which would otherwise yield bad seeds.

Next we used these seeds and the estimated speeds of the ship and its wake to greatly narrow our search. We were usually able to guess the position of the next peak with an error of less than ± 30 meters. Using this guess as one of the initial parameters, we then used MATLAB’s nonlinear least-squares optimization function (`lsqcurvefit`) to determine the shape and position of the Gaussian function which best fit the nearby peak. We used the coordinates of the center of the Gaussian as our estimate for the peak location and the parameters for width and length of the Gaussian to estimate confidence (σ).

Finally, we approximated the x -components of the wake celerities using their y -components and their x and y coordinates. The y -components of celerity and the y coordinates we obtained from the Gaussian peak fitting. The x coordinates were simply the x_s values associated with each time stack. Taking advantage of the proportion:

$$\frac{x_s}{y_s} = \frac{\Delta x_s}{\Delta y_s} = \frac{c_x}{c_y} \quad (7)$$

We find that $c_x = \frac{\Delta x_s \cdot c_y}{\Delta y_s}$. Finding Δx_s and Δy_s requires celerity stacks from two separate transects, $C_{i,j}$ and $C_{i+1,j}$ where i is transect number and j is celerity stack number. There is no way to estimate the horizontal celerity components for the last transect and the estimates are subject to significant levels of error because the magnitude of Δy_s can become very large if different signatures are recognized in peaks $C_{i,j}$ and $C_{i+1,j}$.

Fortunately the y -components of wake celerities were always much larger than the x -components, so errors in estimates of the x -components did not significantly skew overall celerity magnitudes. Furthermore, we limited the horizontal celerity estimates to magnitudes less than 0.5 m/s to prevent the massive outliers which occurred when two different signatures were recognized as discussed above.

II.V.5 Celerity products

Our last goal was to take inventory of our output data and sort it according to confidence estimates. Recall that after fitting a Gaussian function to a celerity stack, we were able to approximate x and y positions, x and y celerity components, and two sigma values (the width and length of the Gaussian). For the average wake we generated four hundred celerity stacks and hence obtained four hundred of these celerity products. In many cases, we generated far more. One wake pictured in Figure 9a, contained 735 celerity products. Figure 9b shows those products plotted in world coordinates atop a frame from the original ship event.

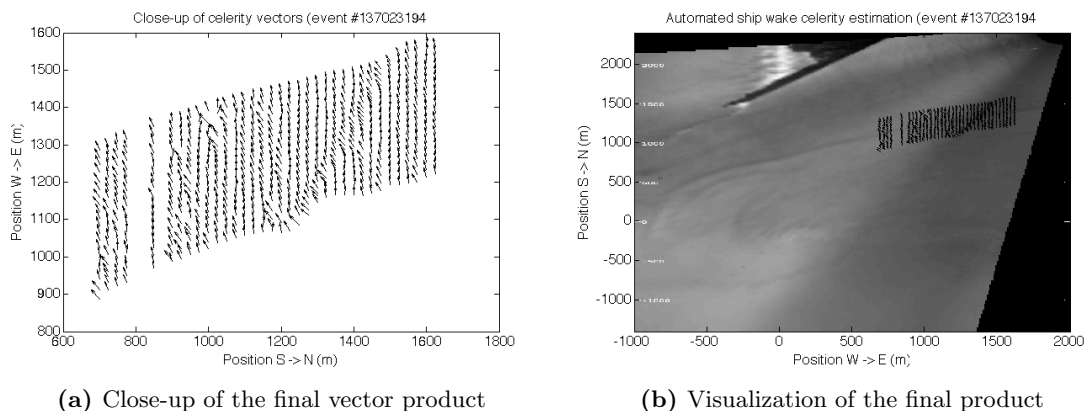


Fig. 9: Signal processing objective: before and after

We determined our confidence estimates (σ scores) using the normalized sum of:

1. the mean column-wise standard deviation of the Gaussian widths
2. the mean column-wise standard deviation of the Gaussian heights
3. the mean magnitude of the Gaussian heights
4. the mean standard deviation of the celerity magnitudes

We used this scoring method because we were satisfied by a qualitative analysis of its performance. Unfortunately, our σ scores have no real-world units and can only be used to rank wakes with respect to each other. Furthermore, these σ scores measure the symptoms of errors in our data rather than the causes (ie variation in celerity magnitudes rather than the problematic Gaussian fitting parameters that caused them). We recognized these issues but chose to continue using our σ scoring method owing to time constraints.

II.VI Filtering data

Despite numerous precautions and sophisticated signal processing, only 40% of all wakes analyzed yielded data of acceptable quality. The others failed due to excessive background noise, nonexistent wakes, or competing signatures. Before analyzing our data, we ranked the 252 wake celerity products by σ and discarded the lowest 60%. We also scored individual transects using our σ scoring system and discarded the lowest-ranked 20%.

After this stringent filtration over a hundred wake products - or 40,000 celerity products - remained. These products were then subjected to random qualitative review to verify their accuracy.

III Results

The strongest correlations in our data were between internal wave celerities and internal wave positions in the MCR (Figure 10). We noticed a general decrease in wave celerity when moving from west to east across the estuary (Figure 10a) and a strong sinusoidal or high-order quadratic relationship when moving

from south to west across the estuary (Figure 10b). The y-value of each data point plotted in figures 10a and 10b is an average of hundreds of individual celerity products with positions in a ± 10 meter range of its corresponding x-value. Considering the large pool of data from which we derived these trends, we are confident in their accuracy.

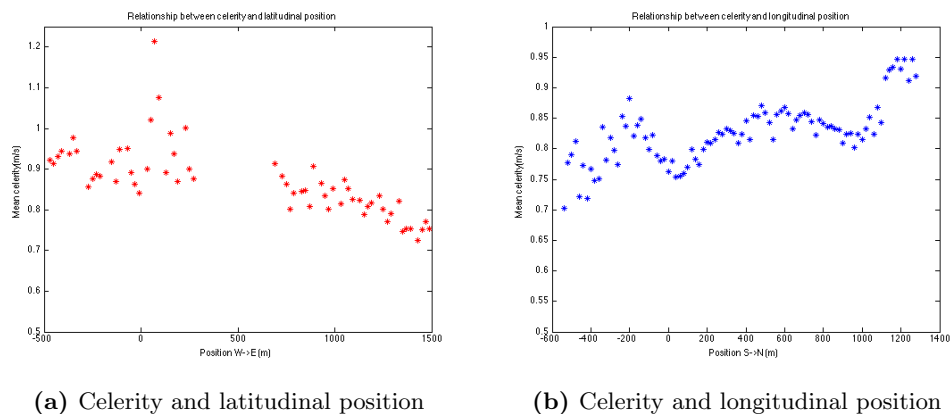


Fig. 10: Averaged correlations between position and wake celerity

We also considered these trends on a case-wise basis by performing linear regression on data from each wake and then plotting the resulting slope values. Moving west to east across the estuary, 72% of wakes evinced a negative correlation as shown in Figure 11a. Moving south to north, this correlation was far less pronounced. Considering the complex cyclical nature of the data in Figure 10b, this is not surprising.

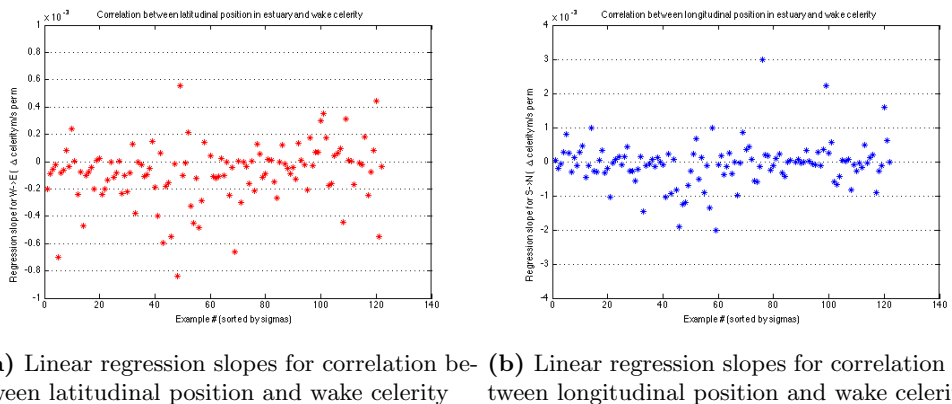
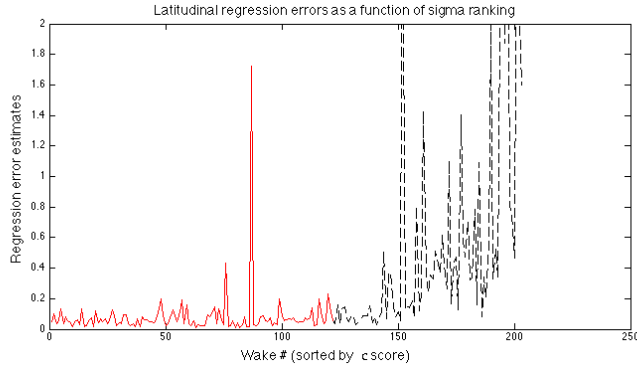



Fig. 11: Slope values returned by linear regression of relationships between position and celerity

MATLAB'S robust linear regression also enabled us to use least-squares error estimates to find the strength of the linear correlations in each wake. In both the south-north and west-east directions we found that linear correlations were strongest for the wakes with the highest σ scores. This is an encouraging finding because it tells us that the wakes we had the highest confidence in were the wakes which exhibited the strongest spatial correlations.



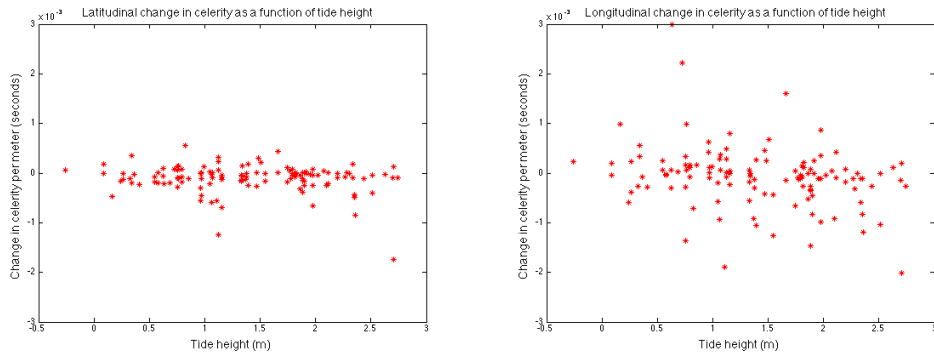
(a) Latitudinal regressions



(b) Longitudinal regressions

Fig. 12: Robust regression least squares error rate for each wake (wakes sorted by σ score). Dashed line portions omitted from previous figures owing to low σ scores

Though we observed strong correlations between internal wave celerity and internal wave position, we were not able to extend these to find relationships between wake celerity and tide phase. First, we found no correlation between mean wake celerity (per ship event) to tide level and tidal current. Second, we found no correlation between linear regression slopes and tide level or tidal current. Figure 13 shows plots of the linear regression slopes and tide level. Neither of the plots shows a significant relationship. We did notice, however, that disproportionately large numbers of ship events occurred at tide heights between [1.5] and [1.75 2.25]. Correlation does not imply causation, but this pattern is worthy of further investigation.



(a) East-west change in celerity as a function of tide (b) North-south change in celerity as a function of tide

Fig. 13: Relationship between celerity regression slopes and tide height

We were also surprised to observe sinusoidal patterns in the relationship between south-north position and celerity magnitude in many individual wake products (four examples shown in Figure 14). These patterns had wavelengths in the range [200 350] meters and did not show up in the averaged longitudinal celerity plot (Figure 10b). However, we suspect that some of the noise in the large sinusoidal plot is caused by these higher-frequency relationships.

More than 75% of the north-south wake product plots evinced these these cyclical patterns and we do not yet know what causes them. We suspect that they are caused by the salt wedge, regional bathymetry, or some combination of the two. The fact that they are averaged out of Figure 14 suggests that their peaks and valleys occur at different locations and therefore cancel each other out. This makes the bathymetry theory less likely because variations in estuary depth should cause the peaks and valleys of these sinusoidal patterns to align and reinforce each other. For this reason, we suspect that some aspect of the internal structure of the MCR estuary causes of these patterns.

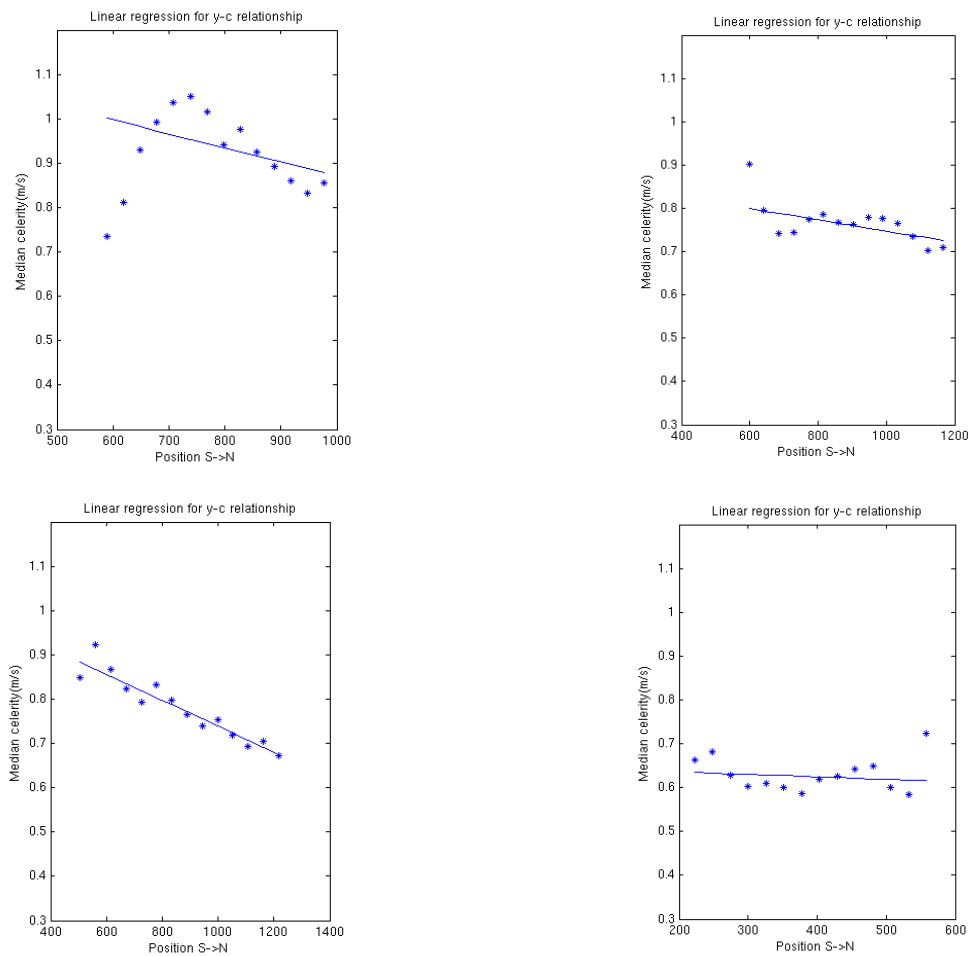


Fig. 14: Example sinusoidal patterns in celerity as a function of south-north position for individual wake products

IV Discussion

IV.I Rationale for two-layer system approximation

Much of this paper is grounded on the assumption that the MCR estuary is a two-layer system. The real-world system is significantly more complex. The minute density differences between freshwater and saltwater means that the two layers frequently mix. Indeed, the transition between freshwater and saltwater

on the MCR is probably more of a gradient than a well-defined boundary.

Even so, the two-layer model is a reasonable approximation of internal structure on the MCR. As shown in the Theory section, we have a strong case for arguing that the ship wake signatures are internal waves. Supposing this is the case, there necessarily must be a two-layer system. Furthermore, if h_1 and h_2 are significantly larger than height of the region where freshwater and saltwater layers mix to form a gradient, we can still approximate that gradient as a well-defined boundary. For these reasons we maintain that the MCR estuary can be reasonably modeled as a two-layer system.

IV.II Contrast signature

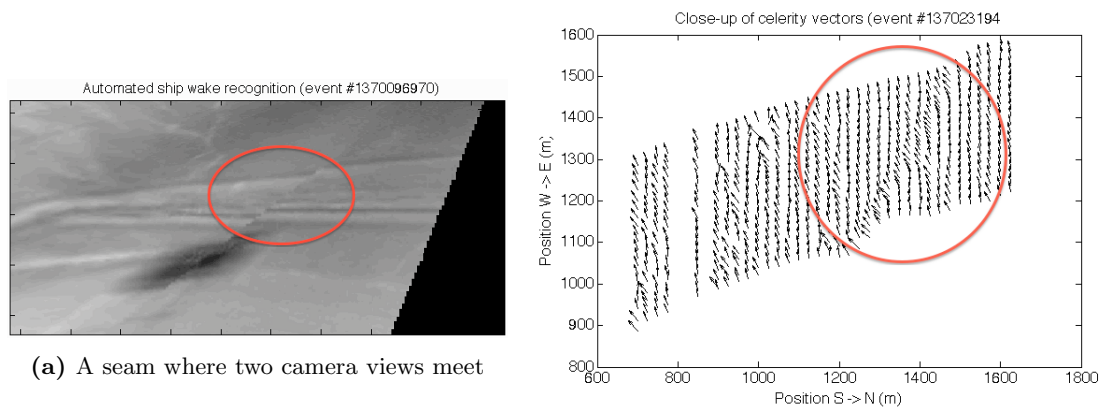
We are still unsure of how internal waves on the MCR manifest themselves as contrast signatures in our wave-averaged movies. One theory is that changes in vertical pressure induced by internal waves alters the wavelengths and amplitudes of small surface waves. These small surface waves reflect light according to their wavelengths and amplitudes, so in wave-averaged Argus images some regions would appear lighter and some would appear darker depending on the position if the internal wave.

IV.III Celerity is double-valued for a ship event

Each ship event generates two wakes. One might wonder that it was necessary to analyze each of these wakes separately. We chose to do this because often one wake was far stronger than the other. Thus by processing them separately we were able to calculate a σ score for each and rate them individually. One consequence is that for some ship events we had two wake products and for others we had only one (wherein the other wake's σ score was too low for consideration).

IV.IV Imperfect rectifications

The Argus images analyzed in this study were merged products of four separate camera views. In the case of the MCR Argus station, these views were imperfectly merged to create the first set of wave-averaged frames. In fact, one of the artifacts of this rectification is visible in ship wakes signatures (Figure 15a) and significantly affected our wake products (Figure 15b).



(a) A seam where two camera views meet

(b) Large horizontal vector components along an artifact

Most of our data was unaffected, but when we eliminate this artifact, errors such as the one shown in Figure 15a will not occur and our results will be yet more precise.

V Conclusion

Having explored the time and spatial variations of internal waves on the MCR, we have discovered that their behavior is closely linked to their location in the Columbia estuary. We found no correlations between

tide level and internal wave celerity, but we suspect that a link between internal wave celerity and salt wedge position does exist. Tide level and salt wedge position are not perfectly correlated; indeed, there are several other measures of salt wedge position. Tidal current, vertical water column salinity, and location of tidal front are all indicators of salt wedge positions which we would consider in future work. Variations in MCR bathymetry which we did not account for in our theoretical equations further obfuscate our findings. Changes in bathymetry alter the ratio h_1/h_2 of the heights of the layers of freshwater and saltwater and greatly complicate our model. In future work, we hope to account for bathymetry in our equations.

Developing an algorithm to automatically detect wave signatures consumed the majority of the time devoted to this study. Now that this algorithm is functional, we have only begun exploring its full potential. Instead of the three months of Argus data used in this study, we could easily scale up our data set to three years of Argus data. This would help us to better understand the nuances of internal waves and draw stronger and more general conclusions.

Our findings indicate that internal waves formed by ships passing through the MCR can be automatically tracked and analyzed to yield useful information about the estuary. For example, we found that waves generally slow as they move upriver but that they evince much more complicated behavior when they move from bank to bank.

No other study has related surface manifestations of internal waves to the internal structure of a water column. This technique requires a large amount of optical data sampled over a large amount of time. It also involves specialized signal processing because good data is difficult to extract and confidence intervals are challenging to compute. These two challenges have prevented science from making use of this technique until now. The preliminary success of our approach shows that surface signatures of internal waves contain useful scientific information.

Our novel technique has expanded the scientific capabilities of Argus and, more generally, helped develop a novel method for modeling estuary dynamics.

VI Acknowledgements

Emphatic thanks to Manuel Castano and Dylan Anderson for their helpful advice and informal MATLAB tutorials. Further thanks to John Stanley for Argus support and computing advice.

REFERENCES

- [Chickadel et al., 2003] Chickadel, C. C., Holman, R. A., and Freilich, M. H. (2003). An optical technique for the measurement of longshore currents. *Journal of Geophysical Research*, 108:1–15.
- [Fitzpatrick, 2002] Fitzpatrick, R. (2002). Gravity waves at an interface. *University of Texas at Austin*.
- [Hartley and Zisserman, 2003] Hartley, R. and Zisserman, A. (2003). *Multiple View Geometry in Computer Vision*. Cambridge University Press.
- [Holman et al., 2013] Holman, R., Plant, N., and Holland, T. (2013). cbathy: A robust algorithm for estimating nearshore bathymetry. *Journal of Geophysical Research*, 118(5):25952609.
- [Holman and Stanley, 2007] Holman, R. and Stanley, J. (2007). The history and technical capabilities of argus. *Coastal Engineering*, 54:477–491.
- [Long and Ozkan-Haller, 2004] Long, J. and Ozkan-Haller, H. (2004). Modeling of the wave circulation field at the nearshore canyon experiment (ncex). *29th International Conference on Coastal Engineering*.
- [Pennucci et al.,] Pennucci, G., Conley, D. C., and Holman, R. Airborne forward eyes final report. *NATO Unclassified*, pages 1–3 and 8–12.
- [Schilling, 2009] Schilling, J. (2009). Mobile platforms measure internal waves. *Center for Coastal Margin Observation and Prediction News*.

- [Suanda et al., 2013] Suanda, S. H., Barth, J. A., Holman, R. A., and Stanley, J. (2013). Shore-based video observations of nonlinear internal waves across the inner shelf. *Journal of Atmospheric and Ocean Technology*, 31:714728.
- [Tate, 2006] Tate, C. (2006). Cape disappointment state park. *HistoryLink: The Free Online Encyclopedia of Washington State History*.
- [Watson et al., 1992] Watson, G., Chapman, R. D., and Apel, J. R. (1992). Measurements of the internal wave wake of a ship in a highly stratified sea loch. *Journal of Geophysical Research*, 9689-9703.
- [Wong, 2002] Wong, L. (2002). Open wide: Examining the mouth of the columbia river. *Oregon State University Extended Campus*.
- [Chickadel et al., 2003, Fitzpatrick, 2002, Holman et al., 2013, Holman and Stanley, 2007, Long and Ozkan-Haller, 2004, Pennucci et al., , Schilling, 2009, Suanda et al., 2013, Tate, 2006, Watson et al., 1992, Wong, 2002, Hartley and Zisserman, 2003]

VII Appendix

VII.I Image transformation geometry

On the Argus system, each image is associated with a database of metadata including camera tilt, roll, and azimuth. The pixels per degree in the field of view and camera distortion factors are taken into account as well. All of this information is combined in a matrix known as the \mathbf{P} matrix. The \mathbf{P} matrix is useful for transformations because the transformation matrix \mathbf{H} can be defined in terms of two \mathbf{P} matrices as follows:

$$\mathbf{H} = \mathbf{P}_1 \times \mathbf{P}_2^{-1} \quad (8)$$

It is useful to think of the \mathbf{P} matrix as a way of storing all the information about a perspective. Obtaining the \mathbf{H} matrix from an image's \mathbf{P} matrix and a target perspective is a simple matter in MATLAB. Afterwards, we can simply interpolate pixels into the new coordinate system using one of MATLAB's built-in functions.

We take advantage of the characteristics of homogeneous coordinates to transpose images according to camera angle.




Fig. 16: Original image

Consider the image of the tile floor shown in Figure (16). Imagine we want to transform the image to world coordinates so that the tiles are each 1 foot square, the origin is at the bottom rightmost intersection, and all the lines are either horizontal or vertical. This problem can be expressed mathematically as a transformation between two different coordinate systems: the camera coordinates (u, v) and the world coordinates (x, y) .

Homogeneous coordinates are particularly well suited for this sort of geometric work. They are similar to Cartesian coordinates except they have an extra scale coordinate w at the end. In terms of Cartesian coordinates they are defined as:

$$(x, y) \leftrightarrow (xw, yw, w) \quad (9)$$

The scale coordinate w gives homogenous coordinates some properties that will be useful later on. For example, multiplying homogeneous coordinates by a scalar does not change their location.

In terms of homogeneous coordinates, the transformation between camera coordinates and world coordinates can be expressed by the matrix equation:

$$\mathbf{Y} = \mathbf{H} \times \mathbf{X} \quad (10)$$

In expanded form it looks like this [Hartley and Zisserman, 2003].:

$$\begin{bmatrix} u_1 & u_2 & u_3 & u_4 \\ v_1 & v_2 & v_3 & v_4 \\ 1 & 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \times \begin{bmatrix} x_1 & x_2 & x_3 & x_4 \\ y_1 & y_2 & y_3 & y_4 \\ 1 & 1 & 1 & 1 \end{bmatrix} \quad (11)$$

Matrices \mathbf{Y} and \mathbf{X} are sets of homogeneous coordinates where each column is a coordinate and each row is an element of that coordinate. Matrix \mathbf{H} is the transformation matrix which exists such that for any point p in world coordinates, $H \times p$ returns the corresponding point q in image coordinates.

Before we transform the image in Figure (16), we need to solve for the values in \mathbf{H} using some algebra. We need to solve for the first eight h values (the ninth one only effects scale) so we need at least four points for which both the world coordinates and image coordinates are known. We obtain these points via simple user input. Next, we can evaluate the cross product in Equation (11) and rewrite the resulting system of equations as shown in Equation (12) where \mathbf{h} is a 9x1 vector [Hartley and Zisserman, 2003].

$$\begin{bmatrix} -x_i & -y_i & -1 & 0 & 0 & 0 & x_i u & y_i u & u \\ 0 & 0 & 0 & -x_i & -y_i & -1 & x_i v & y_i v & v \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} \times \mathbf{h} = \mathbf{0} \quad (12)$$

A nullspace operation in MATLAB yeilds the values of \mathbf{h} . Once we have solved for \mathbf{h} , it is a simple matter to construct a new matrix for the image in world coordinates, cross the indices of that matrix with the 3x3 transformation matrix \mathbf{H} to get image coordinates, and interpolate the pixel intensities at each of those image coordinates into world coordinates.




Fig. 17: Translated image

Figure (17) shows the result of this translation and interpolation when run on Figure (16) using four manually-selected reference points.