Orientation Tracking of Anisotropic Particles Using the Hough Transform

Erika MacDonald
Stanford University
473 Via Ortega, Stanford CA
emacd@stanford.edu

Abstract

In this final project, a dataset was created in a settling tank within the Environmental Fluid Mechanics Laboratory. The dataset involved dropping unit-aspect-ratio cylinders and capturing images as they settled in water. A detection algorithm was modified to fit the lighting conditions of the experimental setup. Two methods were employed to estimate the orientation of these particles. The first method involved fitting an ellipse with the same second moments as the particle in the image. The second method used a Hough transform to fit a line, to provide an estimate of the particle orientation. Then, the distribution of orientations from those two method was investigated. There are significant differences in the distribution of orientations for these two methods.

1. Introduction

Whether it be ember transport in wildfires, sediment transport, or dynamics of microorganisms, the dynamics of anisotropic particles in fluids are virtually everywhere. Anisotropic particles' rotation and alignment of rotation to the flow is something that is not well understood and is important for those particles' dynamics. Experiments looking at the orientation of anisotropic particles exist, but often use multiple cameras [1, 3].

In the case of a single camera, it is often challenging to be able to accurately extract the orientation of these anisotropic particles. In this project, I create a dataset where unit-aspect-ratio cylinders are settling within a tank. Then, I pre-process the dataset and use a pre-existing particle tracking algorithm. I change the particle tracking algorithm in its detection and compare two methods to estimate orientation. The first method is to estimate orientation by fitting an ellipse with the same second moments as the particle. The second method fits a line using the Hough transform and uses that line to estimate orientation.

2. Background

In the Environmental Fluid Mechanics Laboratory, we use Particle Tracking Velocimetry (PTV) which is a Lagrangian method that tracks particle trajectories. This often is confused with but is distinct from Particle Image Velocimetry (PIV), another method in fluid mechanics experiments. PIV, an application of optical flow, involves a higher seeding density of neutrally buoyant tracer particles. In contrast, PTV does not require particles to be neutrally buoyant or as densley seeded, and it tracks each particle from frame to frame.

In previous work in my lab, I have used a pre-existing PTV code [2] to track particles' position over time. In this PTV code, particles are identified by inputting a minimum area for the particles. The PTV code then looks for particles that are brighter than their four nearest neighbors and also brighter than an inputted threshold value. The found particles are then located to sub-pixel accuracy by applying a Gaussian fit to each spatial direction.

After finding the particle location, the particle is tracked across frames using a kinematic prediction. These tracks are broken when no particle lies within a specified maximum dispersion. We then have information about the particle's position and velocity with time. The PTV code also estimates the particle's orientation for anisotropic particles. The code finds an ellipse that has the same second moments as the particle. Then, it calculates the major and minor axis of the ellipse and the angle between the major axis and the x-axis. This angle is used as the estimated orientation of the particle.

Another method frequently used in Particle Tracking literature to estimate orientation is the Hough transform [4]. As discussed in class, the Hough transform is a voting procedure in which image space is first converted into a parameter space (Hough space). In this Hough space, the parameters with the highest count is considered the best fit.

In this final project, I will compare these two methods to estimate orientation and investigate the differences in the orientations that they track. Certainly, the best way to get a better estimate of the orientation would be to add more cameras. However, certain experiments constrain the number of cameras to one. Given this constraint, I would like to investigate different methods to estimate orientation.

3. Approach

For this problem, I created a dataset using laboratory equipment in the Environmental Fluid Mechanics Laboratory at Stanford. I used a settling tank that I built, which is 10”x10”x50” in dimension and is shown in Figure 1. In this settling tank filled with water, 3D-printed unit-aspect-ratio cylinders made from Polyamide-12 were released. These cylinders are 1 cm in diameter and height and have a specific gravity of around 1.02. In this dataset, the light source was 90 degrees from the camera. The camera that was used was a FLIR Flea 3 monochromatic camera and the series of images that I took was with a framerate of 120 fps. A schematic of the experimental setup can be seen in Figure 2.

Using the data from this experimental setup, the orientation of the particle in the image plane can be estimated. The light source was positioned 90 degrees from the camera, rather than in line with it (which would have backlit the particles), for future research purposes. As mentioned in the conclusion, this unique lighting condition might allow for estimating an additional orientation: the in-and-out orientation of the image plane. However, this was beyond the scope of the current project.

This project was divided into three parts. The first part involved particle detection. Due to side illumination, the particle displayed a gradient, necessitating detection across this gradient. The PTV code I was working with initially used a single threshold value for particle detection, which was inadequate under the given lighting conditions. The second part focused on estimating the particle’s orientation once detected. Two methods were used: ellipse fitting and the Hough transform. The distribution of orientations derived from these methods will be discussed.

4. Analysis

This section discusses particle detection and the two methods used for estimating orientation: ellipse fitting and the Hough transform. Finally, a comparison of these two methods will be provided.

4.1. Particle Detection

The first step was to create the dataset for the project. Figure 3 and Figure 4 show one frame from a set of images taken for the dataset. As seen in the figure, parts of each particle are brighter than others. In this image, the lighting comes from the left-hand side, causing the parts of the particle closer to the light source to be brighter, while the parts occluded from the light source are darker.

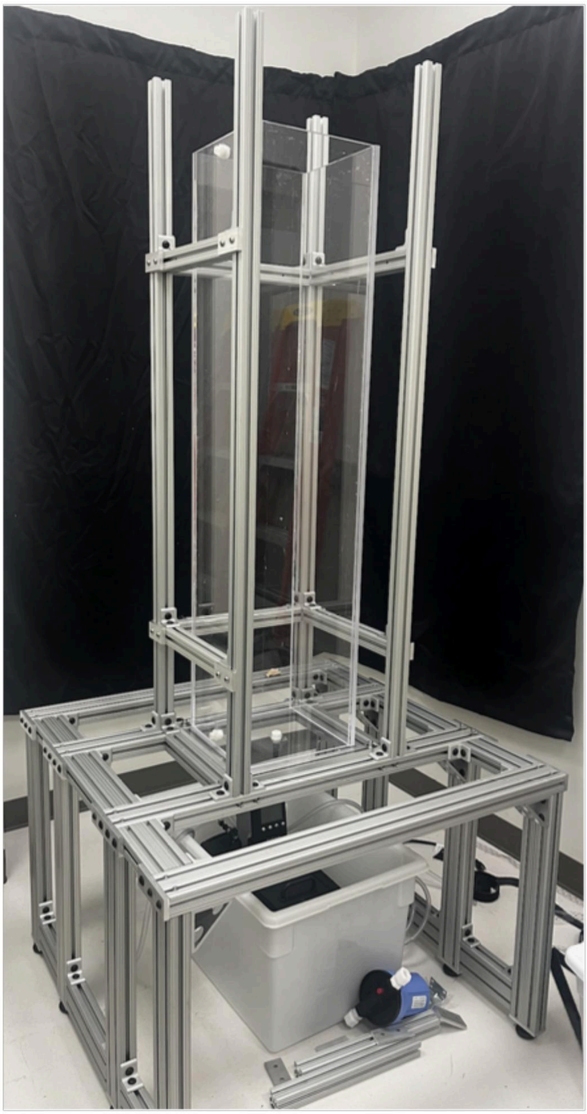

Figure 1. Settling chamber. The dimensions of the tank are 10”x10”x50”

After collecting the data for the project, the first task was particle detection. It became apparent that significant time would need to be spent on the detection process only after starting to work with the data. The initial plan was to use the Particle Tracking algorithm commonly used in the Environmental Fluid Mechanics Laboratory for detection. However, this algorithm detects particles by identifying image regions where the value exceeds a specified threshold, making them brighter than the background. In the dataset,

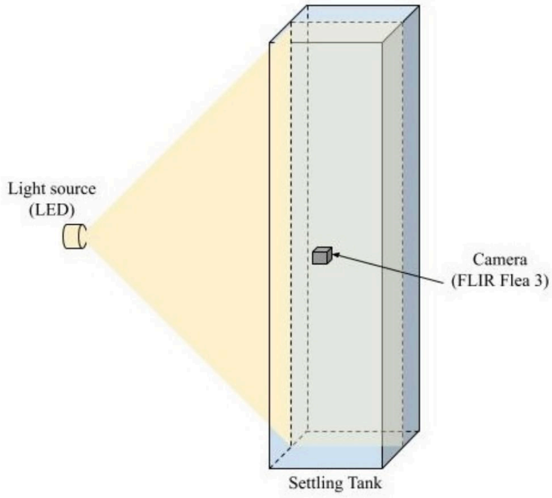

Figure 2. Experimental setup. The camera is 90 degrees from the light source

parts of the particle were darker than the background, while other parts of the same particle were brighter.

To address this, the detection code was modified to identify regions of the image that are both brighter and darker than the background using separate threshold values. Although there may be a better approach for this in the future, this method proved to be fairly effective.

This method first determined the gradient of the grayscale image using the Prewitt gradient operator. A threshold for the magnitude of the gradient was then used to identify regions. These regions were filtered to find areas smaller than a specified maximum area. However, this process sometimes left 'holes' in the particles, meaning the entire particle was not identified. An instance of these 'holes' is shown in Figure 5. To address this, the final step was to fill in the perimeter of the identified region, considering the filled region as the identified particle.

After modifying the particle detection, the particle tracking was performed. Figure 6 shows the trajectories of the particles for one run.

4.2. Ellipse Fitting

The first method used to estimate orientation was ellipse fitting. In this method, an ellipse with the same second moments as the particle was fitted. The angle between the major axis of this ellipse and the x-axis was used to estimate the particle's orientation. As shown in Figure 7, which depicts one frame of one particle, the orientation does not visually match the estimated orientation. A challenge with using ellipse fitting in this scenario is that these particles have a unit

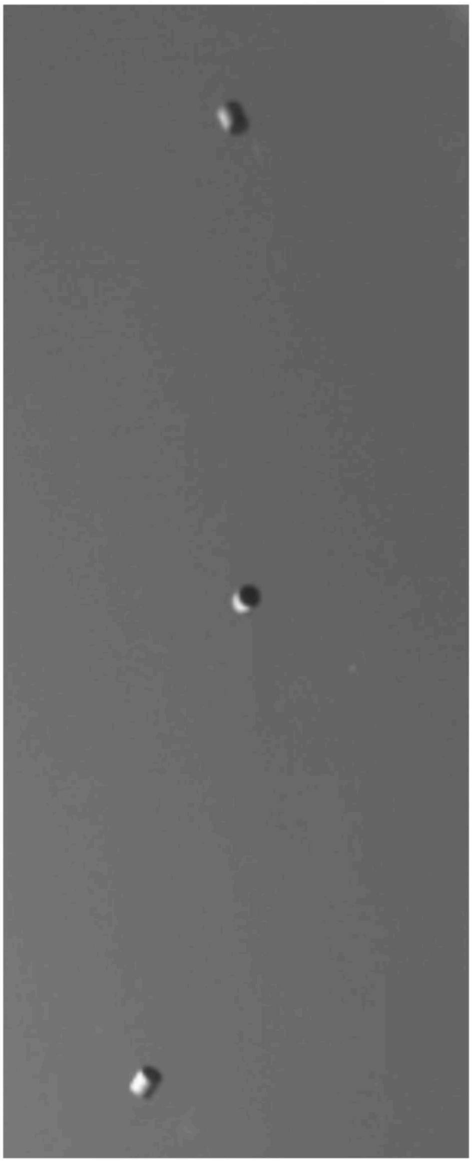

Figure 3. Example of one frame of experimental results

aspect ratio. For particles with larger or smaller aspect ratios, such as rods or disks, ellipse fitting might be a more suitable option.

4.3. Hough Transform

The second method used to estimate orientation was the Hough Transform. For the Hough Tranform, edges for each particle region were detected using the Canny method. In the Canny method, local maxima of the gradient of the im-

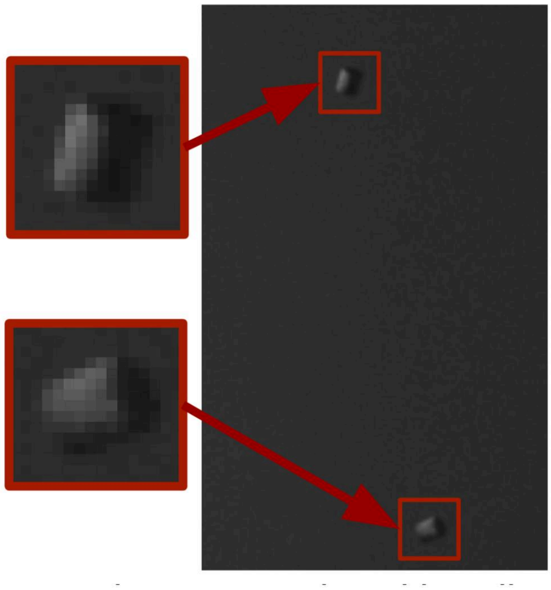

Figure 4. Example of one frame of the images taken.

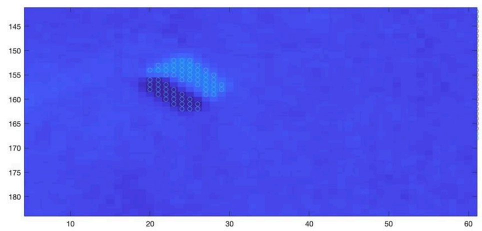

Figure 5. Particle detection results. The circles are detected pixels. As we can see, there is an undetected region within the particle.

age is found. There are two thresholds in this method to detect edges that are stronger and weaker. Having two thresholds should help reduce the effect of noisy measurements.

After detecting the edges, the Hough transform was performed on those edges. A visualization of the Hough transform matrix can be seen in Figure 8. The five highest peaks in the Hough transform matrix were identified. Using the average of the five highest peaks, rather than a single peak, was chosen to reduce the likelihood of error in the orientation estimate. In Figure 8, the red boxes represent the five highest peaks in the Hough transform matrix.

In Figure 9, the Hough lines corresponding to those five highest peaks and the average of those lines are displayed. As we can see, for this particular frame for this particular particle, it seems as though the Hough transform does a better job than the ellipse fitting. The next section, however,

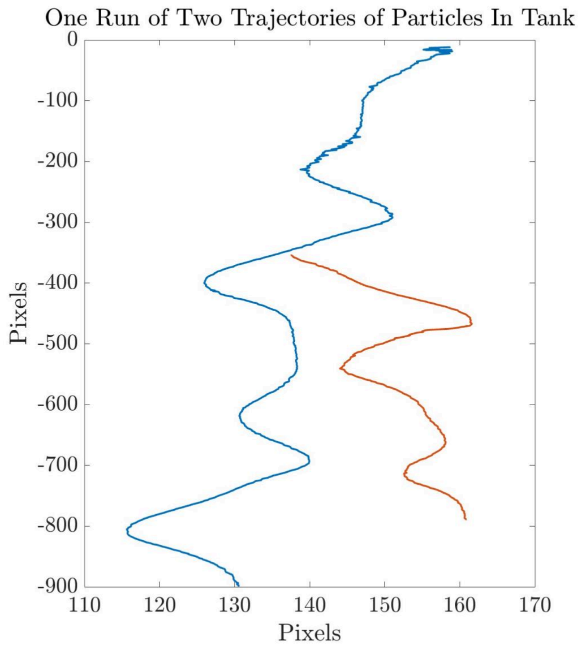

Figure 6. Example of trajectories of position of particles

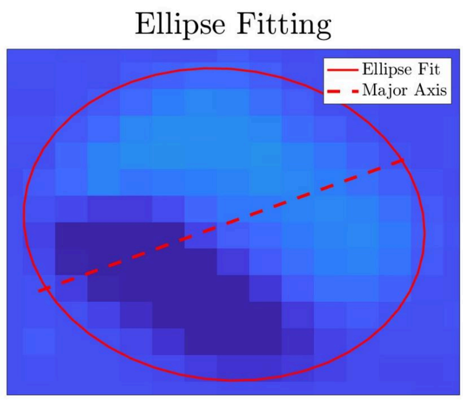

Figure 7. One frame of one particle that was fit using an ellipse.

will quantitatively discuss the differences.

4.4. Comparison of the Two Methods

In order to quantitatively look at the difference for these two methods, the normalized probability density function was found for the orientations under the two methods.

In Figure 10, we can see these normalized probability density functions. First off, we can note that the distribution

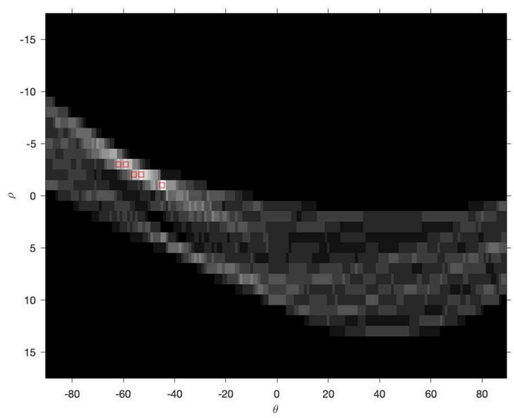

Figure 8. A visualization of the Hough transform matrix.

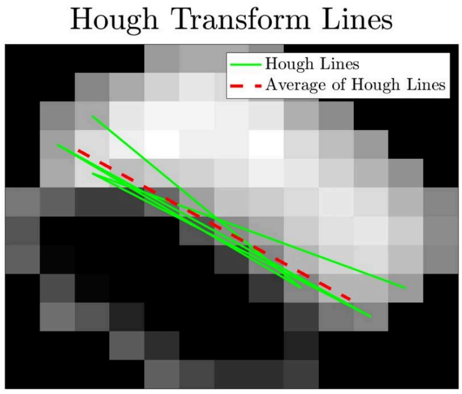

Figure 9. Hough lines overlaid on a particle.

of orientations under the two methods are distinct. The ellipse fitting distribution looks almost bimodal, with the the peaks around -90 and 90 degrees. This could have arisen from the ambiguity of what direction the ellipse is 'facing', and perhaps with knowledge of time history this distribution would no longer be bimodal.

For the Hough Transform case, we can see that is more uniformily distributed compared to the ellipse fit. In the best case scenario, we would have ground truth data to compare to. However, since the data is experimental, it would be hard to get any sort of ground truth data other than by labeling the data frame by frame.

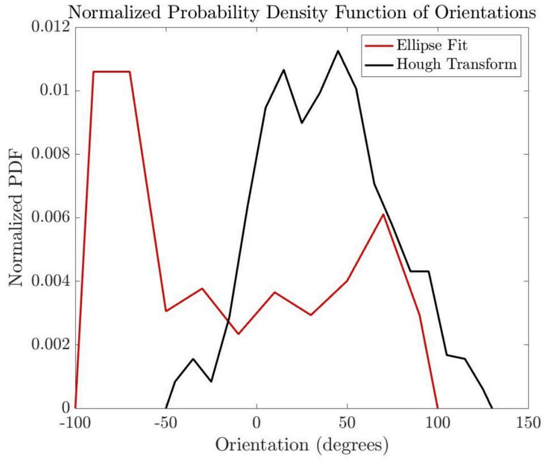

Figure 10. Normalized Probability Density Function of Orientations.

5. Conclusion

5.1. Completed Work

For this project, I created real data in a laboratory to analyze. I calibrated images and preprocessed the data. Then, I modified a detection algorithm to find the particles frame by frame, and then found the orientation using two methods: fitting an ellipse and by using the Hough transform. I compared the two methods by looking at the normalized probability density function of the orientations under each scenario.

5.2. Future Work

There are many possibilities for future work. First, one way of improving the orientation estimates would be to use the particle's time history to make the orientation data less noisy. Since the particle's trajectories and orientations should be smooth, the orientation frame by frame should not change dramatically.

To make the findings more robust, another experiment could be run with multiple cameras. We could continue to do the same thing where we measure orientation with one camera, but we could compare those results with the results from multiple cameras, which would be closer to a ground truth result.

Finally, another possibility for future work would be to exploit this lighting condition to get an additional orientation: the orientation in and out of the image plane. By looking at the gradient across the particle, it seems like an additional orientation could be estimated.

References

[1] Stefan Kramel Rui Ni Guy G. Marcus, Shima Parsa and Greg A. Voth. Measurements of the solid-body rotation of anisotropic particles in 3d turbulence, 2014. 1

[2] Haitao Xu Nicholas T. Ouellette and Eberhard Bodenschatz. A quantitative study of three-dimensional lagrangian particle tracking algorithms, 2006. 1

[3] Federico Toschi Shima Parsa, Enrico Calzavarini and Greg A. Voth. Rotation rate of rods in turbulent fluid flow, 2012. 1

[4] Monica Kishore Nicholas T. Ouellette J.P. Gollub Shima Parsa, Jeffrey S. Guasto and Greg A. Voth. Rotation and alignment of rods in two-dimensional chaotic flow., 2011. 1

6. Supplementary Material

Link to Git repository (Shared with Krishnan): https://github.com/erikamacd/cs231AFinalProject

Video of a particle falling:

https://drive.google.com/file/d/1t85HxU_pssGlSrwtF0SEjniyzKI5Ijis/view?usp=sharing