Duke University – Fall 2025

Measuring The Speed of Light

Author: Zachary Maldonado (zim3)
Dr. Huanqian Loh - Department of Physics; Department of Electrical & Computer Engineering
Aug. - Oct., 2025

Executive Summary

This laboratory measures the speed of light c with an experimental design–motivated by a simplified Foucault setup–consisting of a laser, rotating mirror, convex lens, camera sensor, and stationary mirrors. Beam displacement x is recorded vs. motor rotation frequency, yielding a linear relationship whose slope can be used to determine the speed of light, which is measured to be $2.84 \times 10^8(\frac{m}{s}) \pm 0.03 \times 10^8(\frac{m}{s})$. The calculated uncertainty in c is dominated by uncertainty in slope of the laser spot horizontal position, δA_x. The experimental results for c differ by $\sim 5\%$ from the accepted value; an investigation into the effect that uncertainty in frequency δf has on δA_x strongly suggests that residuals are heavily influenced by some other factor contributing to variation in horizontal position. This "other factor" could perhaps explain why we see a 5% deviation from the accepted value of $3.00 \times 10^8(\frac{m}{s})$, and may be attributed to optical/measurement effects such as unclean components, ghost fringes, misalignment, thresholding, or camera noise.

1 Setup and Procedures

1.1 Experimental Setup

This laboratory is based upon a simplified version of Foucault's experiment, which consists of a laser, a rotating mirror, and a return mirror (see figure 1 below for simplified concept). In reality the laser beam is a divergent ray, but for the introduction of this simplified setup it should be considered as ideal. The rotating mirror is attached to a brush-less motor whose frequency can be adjusted from the laboratory computer setup. While the rotating mirror spins, the forward ray is sent and bounces off in different directions depending on the angle of the rotating mirror at the time of reflection. If the return mirror is aligned perpendicularly to the forward ray as it leaves the laser, then the ray can be directed towards the center of the return mirror, which in turn reflects the ray back towards the center of the rotating mirror. Light travels at a constant speed, and like anything else that moves in this universe, it takes time to traverse some distance. That being said, there is a quantifiable change in time as light in the forward ray travels to the return mirror and back to the laser aperture, which allows the motor to change the orientation of the mirror upon it by some small angle θ. Since the return beam encounters a plane rotated θ from its original position when the forward ray came into contact, it is reflected by an angle of $\alpha = 2\theta$ from the original, perpendicular path of the forward beam. See figure 6 and the accompanying equations for a derivation of this relationship. Now the return ray misses the laser aperture by some small distance x, which will depend on the geometry of the experiment, as well as the frequency f of the rotating mirror and, of course, the speed of light c. The true geometry of the experiment will be discussed, however note that the displacement x at the laser aperture is accurate under this simplified version.

Figure 1: Simplified Experimental Setup

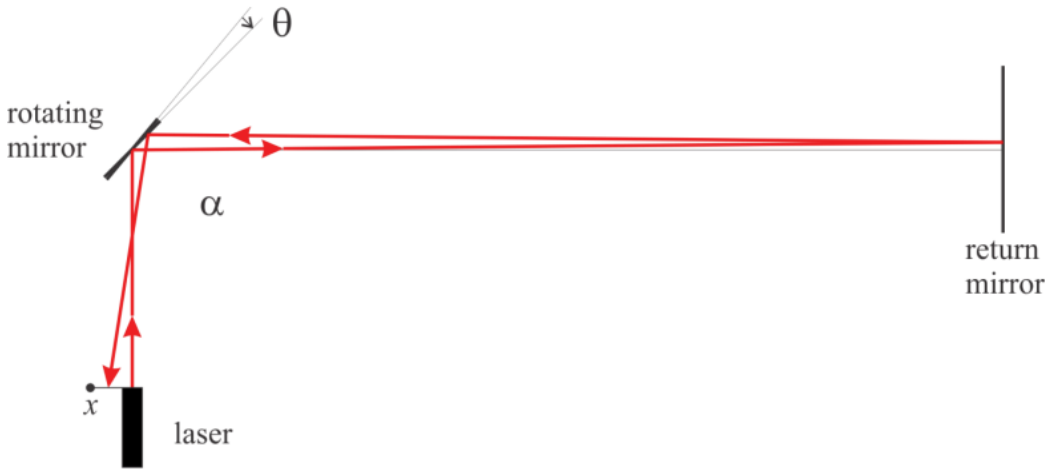

Figure 2: Simplified Experimental Setup at Different Orientation

There are some inherent issues with the simplified setup in Figure 1. Most notable is the issue that this simplified experiment only works for a single orientation of the rotating mirror, that being the orientation where the forward ray is reflected off the rotating mirror to exactly the center of the return mirror. For all other orientations of the rotating mirror, however, we will see a different displacement x since the ray hits a different point on the rotated plane as it returns to the laser aperture (see Figure 2 above). The solution to this problem is to place a convex lens between the rotating mirror and the return mirror such that the rotating mirror is focal length F away from the lens. This ensures that all rays reflected off of the rotating mirror plane originate from the lens's focal point. This is crucial, because it ensures that the ray travels in a straight path towards the return mirror and back, should it pass through the convex lens, ridding the setup of the issues raised in Figure 2 and allowing reliable displacement data.

In Figures 3(a) and 3(b), the benefits of placing the convex lens in such a way are evident. The forward ray, having come from a focal point, is directed in a straight path after crossing the lens, which allows the return ray to hit the center of the rotating mirror again. We can confidently conclude that with the addition of a convex lens, the displacement x does not depend on the initial orientation of the rotating mirror.

In reality, the laser beam is divergent, so we must adjust the diagrams provided in Figures 3(a) and 3(b). Figure 4 displays a ray whose radius increases with distance from the laser aperture. We can treat this divergent ray as a spread of single rays which all hit the rotating mirror. Notice however that the only individual "ray" which directly arrives at the focal point on the rotating mirror is the middle one. This one behaves like the rays described in Figures 3(a) and 3(b), traveling parallel to the optical axis of the lens, while all other rays do not strike the exact focal point on the rotating mirror and as a result will not travel parallel after passing through the lens. We are motivated to resolve this phenomenon by strategically placing the return mirror at the exact distance from the convex lens, d_i, where the image of the laser beam is formed. This way, every ray converges at the return mirror, is then reflected with an identical angle of incidence (diverging back to the lens), and will

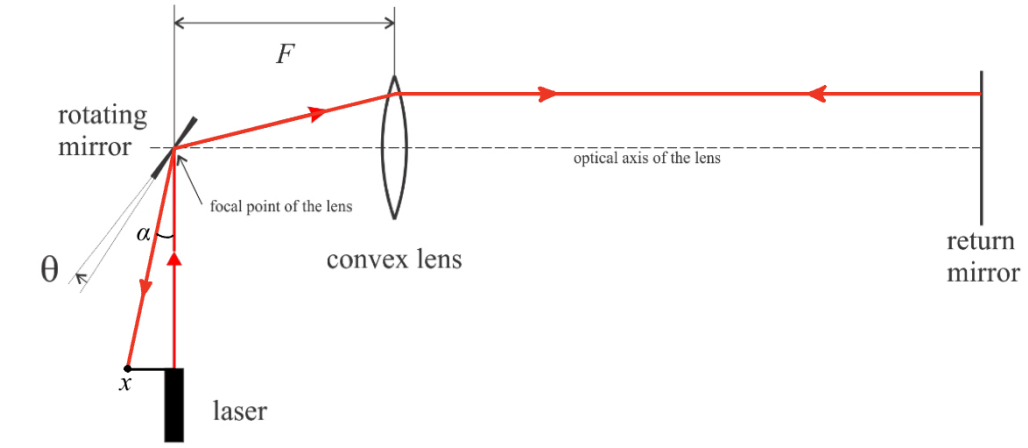

(a) Forward Ray Positive wrt Optical Axis

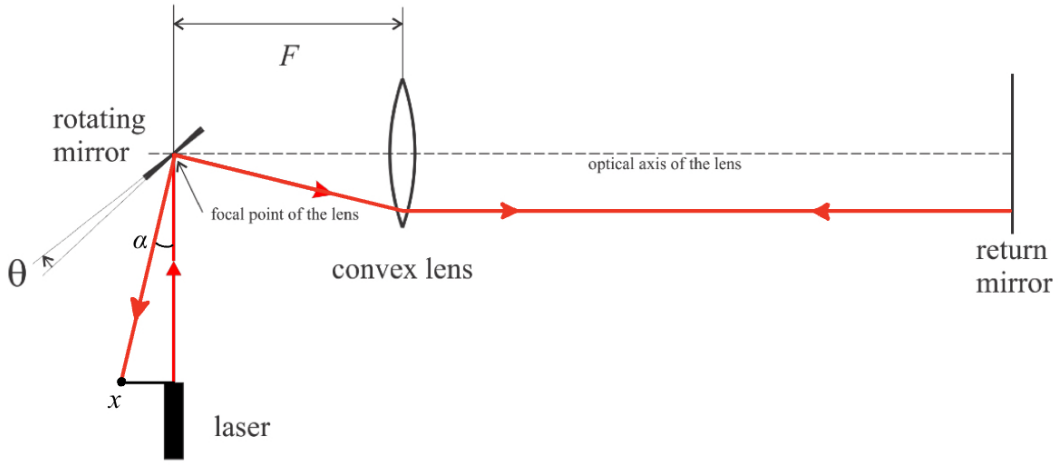

(b) Forward Ray Negative wrt Optical Axis

Figure 3: Alternative Orientations of Rotating Mirror With Lens

converge once more to the original beam radius once having arrived at the laser aperture (or as we will find, at wherever the displacement x is equally observed).

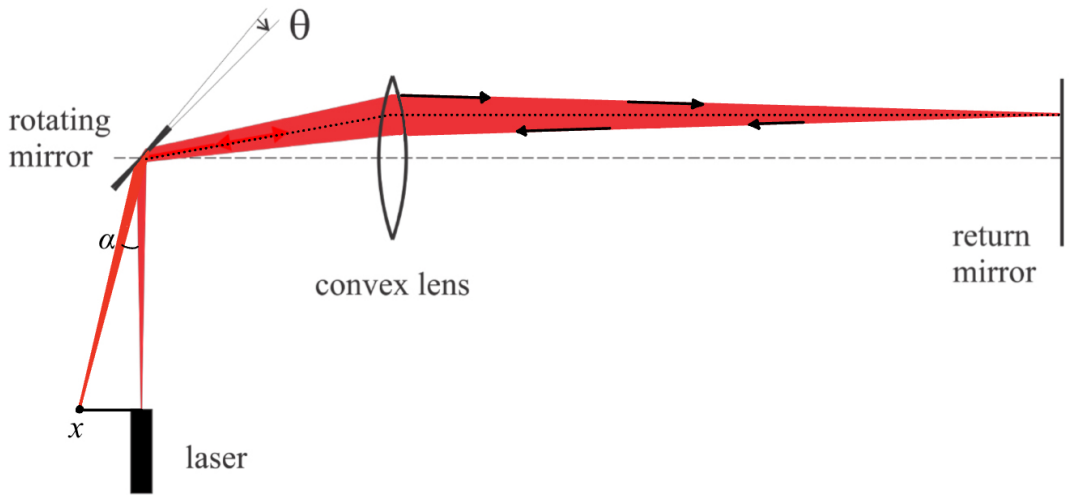

Figure 4: Divergent Laser Beam

With the above considerations in mind, the team developed the experimental setup in Figure 5 to measure the displacement x as the frequency f of the rotating mirror increases. Table 1 provides all relevant measurements and uncertainties recorded for this setup. The focal length of the convex lens used was given as 2 meters. In order to create that separation distance between the rotating mirror and the convex lens, an intermediate mirror was placed which reflects the ray another time, towards the convex lens, for a total ray path of length 2 meters between the rotating mirror and the lens. To determine the optimal position of the return mirror, the thin lens equation was utilized, where in this case d_o is the distance from the laser aperture to the lens and d_i is the distance from the lens to the return mirror:

$$\frac{1}{F} = \frac{1}{d_o} + \frac{1}{d_i} \tag{1}$$

The distance from the laser aperture to the rotating mirror was measured to be roughly 1 meter; adding the focal length F of 2 meters yields d_o. Let us now solve Equation (1) for d_i:

$$\frac{1}{d_i} = \frac{1}{2} - \frac{1}{3} = \frac{1}{6}$$

$$\boxed{d_i = 6 \text{ meters}}$$

Once the positions of the mirrors and lens were secured, a beam splitter lens was placed between the rotating mirror and the laser aperture, angled such that the return ray is redirected at around 90° towards a CCD camera sensor, in Figure 5. Note that the beam splitter lets the forward ray through and reflects it away from the camera, such that the only ray directed into the camera is the return ray. This setup is necessary since the displacement x

is far too small to measure at the laser aperture without blocking the forward ray. See the Appendix for a quick graphic which clarifies why this setup does not distort the displacement x. To prevent the camera sensor from being exposed, a 1000x neutral density filter is attached, which is key for obtaining a more precise measurement of the laser spot position.

To actually measure the frequency of the rotating mirror, a photodiode (not shown in Figure 5) connected to an oscilloscope is used. This apparatus was aligned in-plane with the laser spot on the rotating mirror by observing the thin red line created across the room while the motor is spinning. This "red line" is in fact a rotating beam spot which is how a well aligned photodiode apparatus can record rotations per second, allowing the team to record the corresponding frequency reading to a displacement x.

Distance Measurements

Distance	**Value (cm)**	**Uncertainty (cm)**
Laser to rotating mirror, L_0	100.3	0.2
Rotating mirror to intermediate mirror	102.4	0.2
Intermediate mirror to lens	97.60	0.2
Lens to return mirror	600.0	1.5
Rotating mirror to beam splitter	47.60	0.2
Beam splitter to camera	53.00	0.2

Table 1: Relevant Measurements of Experimental Setup

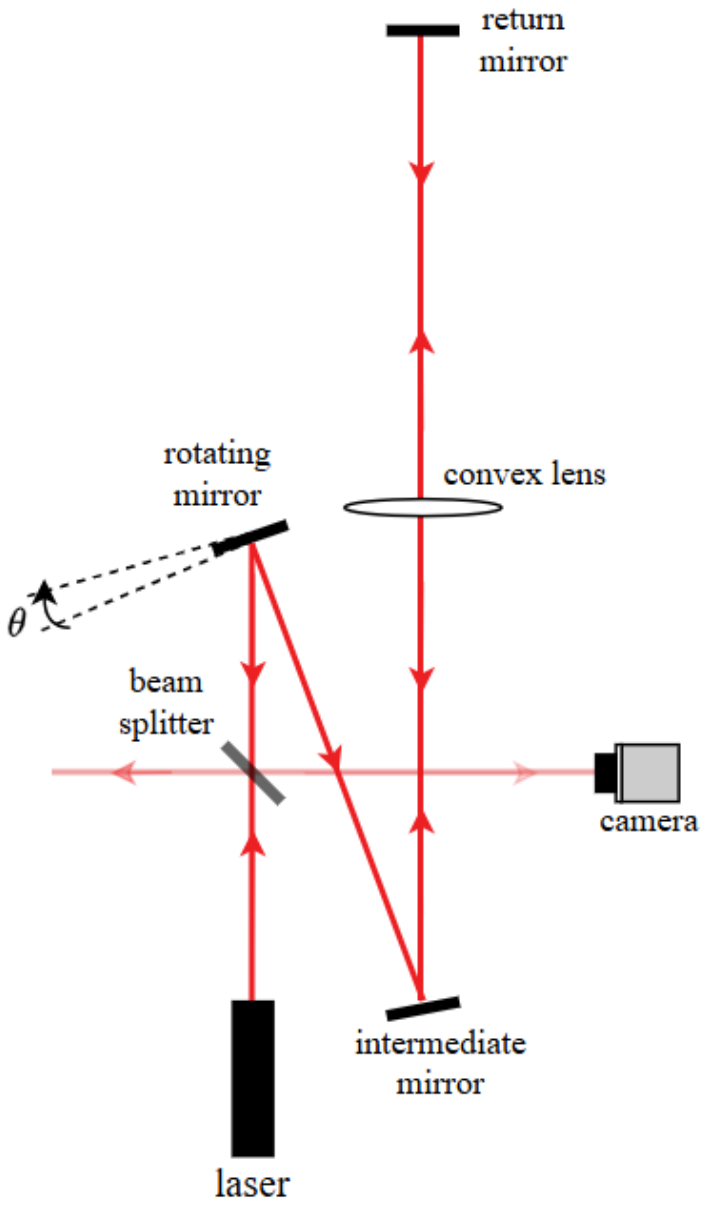

Figure 5: Devised Experimental Setup (not to scale)

1.2 Derivation of $\alpha = 2\theta$ Relationship

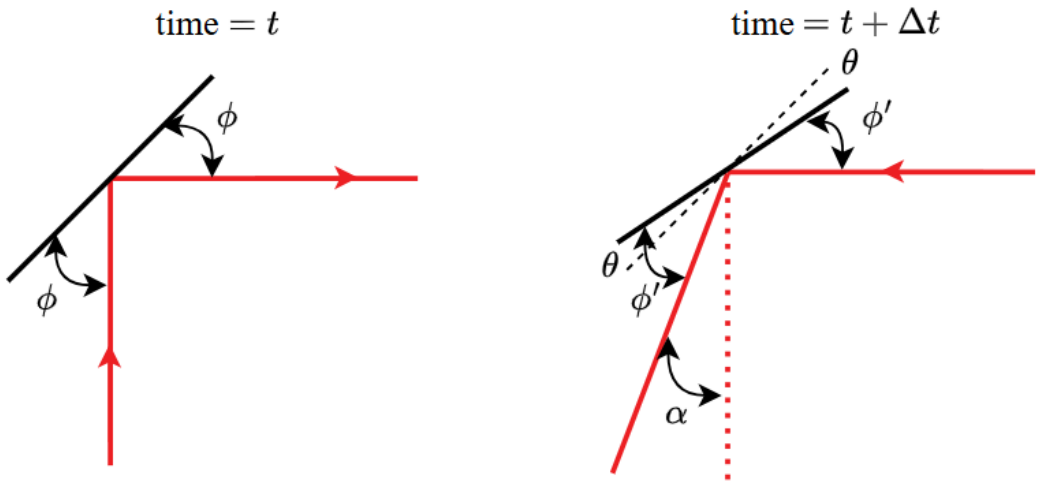

Figure 6: Reflection of Ray after some Δt

Figure 6, shown above, describes how the ray will change once the rotating mirror has changed orientation by some small angle θ from its previous position. We can express the angle that the ray makes with the rotated plane as the difference between the original angle ϕ and the small angle θ:

$$\phi' = \phi - \theta \tag{2}$$

Now, let us observe the portion of the ray $after$ it has been reflected off the rotated plane, this time relating the old angle ϕ to the angle α between the forward and return (reflected) rays:

$$\phi' - \theta + \alpha = \phi \tag{3}$$

We are now able to substitute in for the angle ϕ' by combining equations (2) and (3), which after simplification leaves us with only α in terms of θ:

$$\phi - \theta - \theta + \alpha = \phi$$

$$\boxed{\alpha = 2\theta}$$

1.3 Derivation of Relationship Between x and f

Looking back at Figure 1, let us now call the distance between the rotating mirror and the return mirror L_R, and the distance between the laser aperture and the rotating mirror L_o. Light, traveling at speed c, takes some time Δt to leave the rotating mirror, reflect from the return mirror, and arrive back at the rotating mirror. The following relation results:

$$\Delta t = 2\frac{L_R}{c} \tag{4}$$

The motor is rotating at some constant angular velocity ω, which is known to be expressed as the following:

$$\omega = 2\pi f = \frac{\Delta\theta}{\Delta t} \tag{5}$$

Since Δt is known, we can substitute Equation (4) into Equation (5) and obtain a relationship between θ and f. Note that since Δt is defined as the quantity of time between when the forward and (successive) return rays hit the rotating mirror, $\Delta\theta$ becomes θ, the change in orientation of the rotating mirror during that Δt:

$$\theta = \frac{4\pi f L_R}{c}$$

Let us now look at the right triangle created by the forward ray before the rotating mirror, the return ray after the rotating mirror, and the displacement x at the laser aperture, visible in Figure 1. Using trigonometric relations and substituting our newly defined θ, the following relationship between x and f is produced:

$$tan(\alpha) = tan(2\theta) = \frac{x}{L_o}$$

$$x = L_o tan\left(\frac{8\pi f L_R}{c}\right)$$

Applying the small angle approximation $tan(\alpha) \approx \alpha$, we further reduce x to a linear function of f:

$$\boxed{x(f) \approx \left(\frac{8\pi L_R L_o}{c}\right) f} \tag{6}$$

2 Data

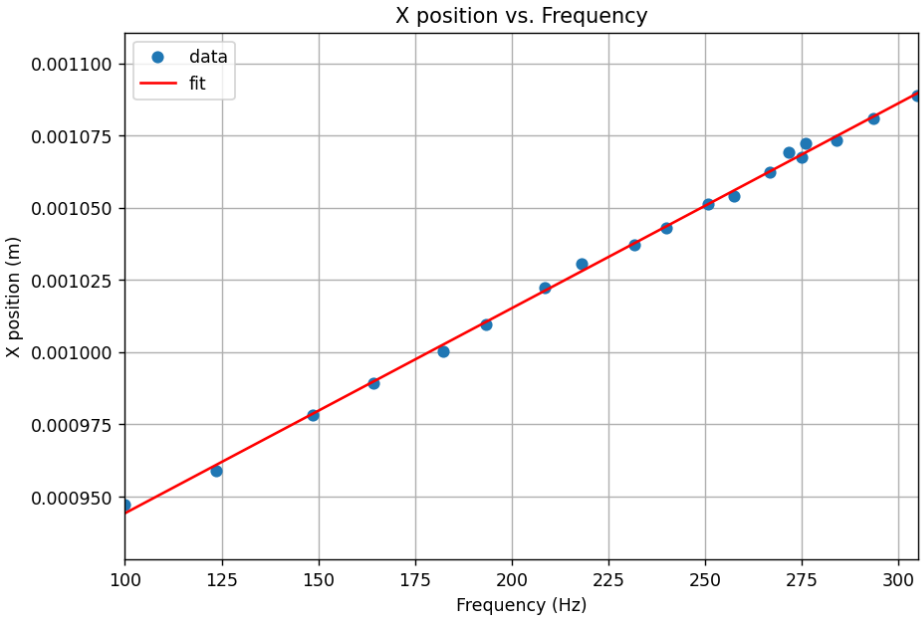

Figure 7: Experimental Relationship Between Laser Spot x Displacement (m) and Frequency (Hz)

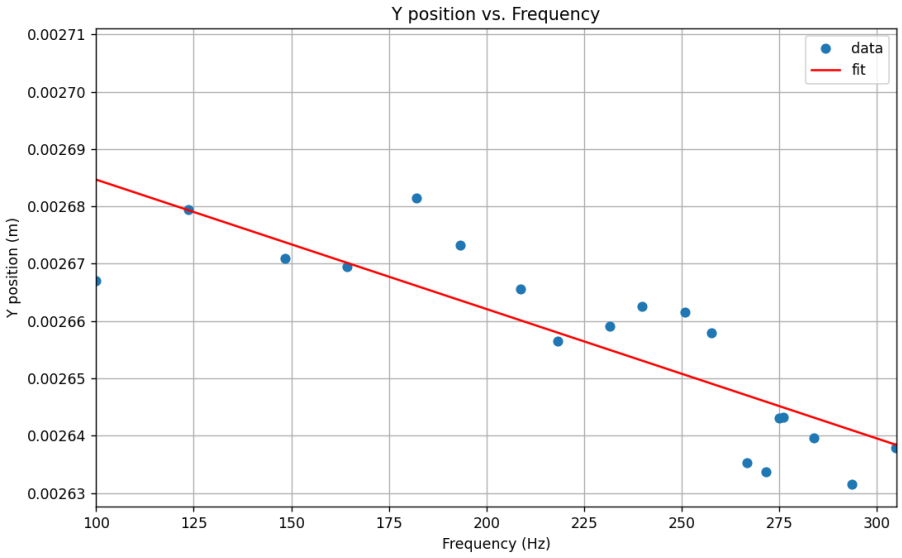

Figure 8: Experimental Relationship Between Laser Spot y Displacement (m) and Frequency (Hz)

3 Discussion of Data and Results

3.1 Processing The Data

Figure 9: Example of Raw Laser Spot Image

As will be discussed in greater detail, the laser spot is not ideal. Vibrations, dust, and positioning of equipment all contribute to imperfections in the laser spot. It does, however, still consist of a spread of pixels each associated with some brightness visible in the camera image, such as the one presented in Figure 9 (shown above). Moreover, the brightness is clearly concentrated around some centroid position $(\bar{x}, \bar{y})$, and generally decreases radially outwards. This trend is the motivation for the creation of a python script which filters through every row of pixels in the image returned from the camera, finding a "center of mass" with a weighted average calculation, and ultimately returning a position for each centroid found for each image taken at its respective frequency.

Here's how it works: the script loads all images taken during experiment, and operates on them one by one. These laser spot images are first converted to grayscale, so that the intensity values associated with each pixel are just integers that can range from 0 to 255. Now each pixel not only has an (x_i, y_i) location, but is also attached to an integer value that can serve as an associated "weight" w_i. As mentioned above, the laser spot has visible imperfections that must be filtered out, which is why the script defines a minimum intensity value to exclude any pixels with an intensity lower than said value from the weighted average calculation:

$$\bar{x} = \sum_{i=1}^{n} \frac{w_i x_i}{w_i} \tag{7}$$

Equation (7) is also used to find $\bar{y}$ in the same manner. A sample post-processing image of the laser spot is provided in Figure 10 below, where the blue dot corresponds to the found centroid position:

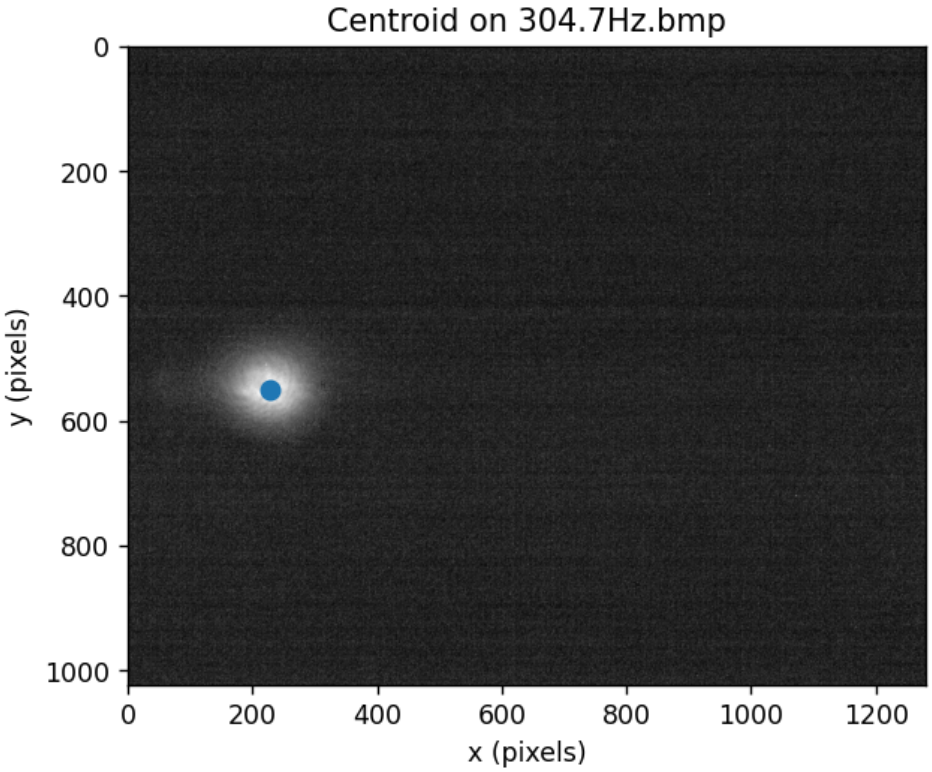

Figure 10: Example of Typical Laser Spot Centroid Post Processing

The centroid positions $\bar{x}$ and $\bar{y}$ are plotted against their corresponding frequencies, producing Figures 7 and 8. A linear fit is also applied to establish a trend line for each, visible in those plots. The slope A_x, then, of this line is the frequency f divided by displacement x. Using this relationship, let us rearrange Equation (6):

$$c = \frac{8\pi L_R L_o}{A_x} \tag{8}$$

Note that in order for Equation (8) to hold, we must report the displacement x in meters, which is done by the relationship that 1 pixel is equivalent to 4.8 μm.

After processing all images, plotting displacement x vs. corresponding frequency f, and evaluating Equation (8), the speed of light was found to be 2.84×10^8 m/s, with an uncertainty of 0.03×10^8 m/s.

3.2 Choosing The Right "Threshold"

As mentioned before, the laser spot as it appears on the camera snapshot suffers from imperfections that will have effects on the calculated position of the $(\bar{x}, \bar{y})$ centroid position. While issues like beam divergence and unwanted reflections are indeed actively addressed, contributions such as small imperfection in lens placement, aberrations in laser internal optics, or dirty optical components (including mirrors, lenses and filters) are much more difficult to account for and do in fact result in noticeable regions of pixels in the laser spot image of higher-than-expected intensity. Hence the need for including a threshold minimum intensity value to exclude these imperfections from the weighted average calculation in order to attain as true a measurement possible.

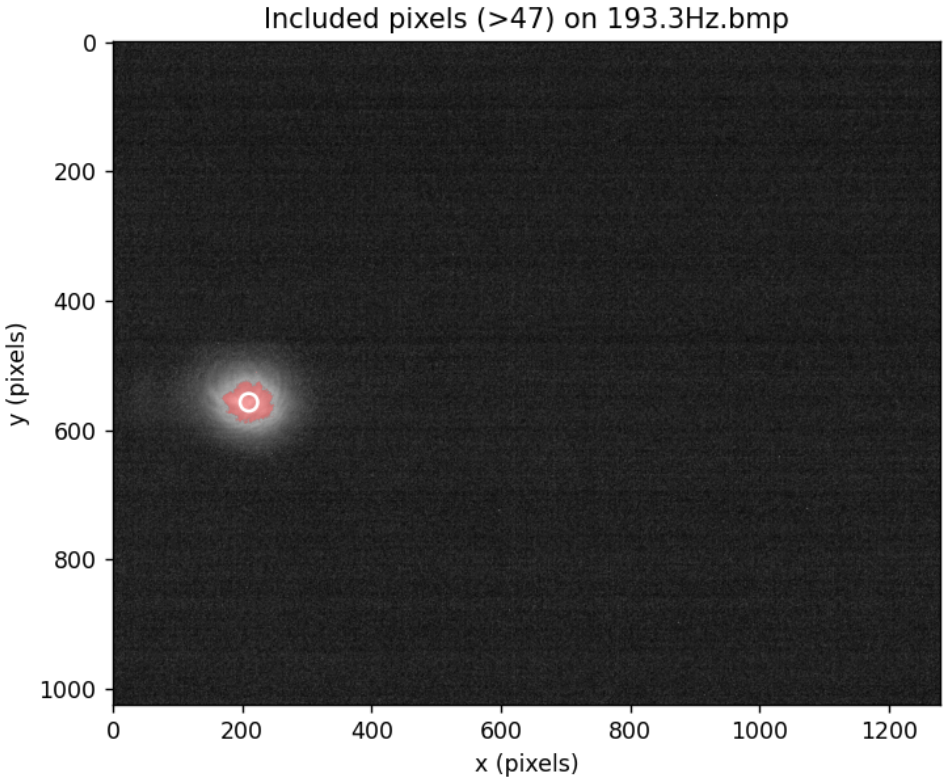

Figure 11: Example of Weighted Average Pixel Threshold Selection

The question then arises: how can we choose a reasonable threshold without bias? The best way to go about this is to judge the results produced via this threshold by (1) its ability to visually exclude all irregularities and (2) its ability to minimize uncertainty in the slope, A_x. Too small a threshold, and the centroid position will be skewed by imperfections. Too large a threshold, and the sample size will not be large enough to assure statistical significance (discrete pixel intensity cutoff will have a significantly higher effect on weighted average). As the threshold approaches either of these cases, uncertainty in slope increases and/or imperfections are included. After trial and error according to these constraints, a grayscale intensity of 47 was accepted to be a most optimal threshold. See figure 11 above for an example range of this threshold at a frequency of 193.3 Hz.

3.3 Error Analysis

Recall that Equation (8) establishes c, the speed of light in meters per second, as a quantity dependent on the distance between the rotating mirror and the return mirror L_R, the distance between the laser aperture and the rotating mirror L_o, and the slope of the fit line A_x. These quantities all have associated uncertainties that each contribute, with varying degrees, to the overall uncertainty in c. Considering c as:

$$c = c(L_R, L_o, A_x)$$

The error propagation formula can be expressed as:

$$\left(\frac{\delta c}{c}\right)^2 = \left(\frac{\delta L_R}{L_R}\right)^2 + \left(\frac{\delta L_o}{L_o}\right)^2 + \left(\frac{\delta A_x}{A_x}\right)^2 \tag{9}$$

And the individual contribution of some parameter, say A_x for example, is:

$$\delta c_{A_x} = \left|\frac{\partial c}{\partial A_x}\delta A_x\right| \tag{10}$$

The uncertainty in the slope, δA_x is obtained by taking the square root of the first entry in the covariance matrix returned by Python's "polyfit" function. Uncertainties δL_R and δL_o are measured and already known. Their values are used in the Python script, in conjunction with the error propagation formula, to calculate the uncertainty in the speed of light δc. Let us also see the individual contributions to this uncertainty, in Table 2 below. Each contribution was calculated in the python script using the above equation for individual contribution:

Individual Contributions to δc

Parameter	**Uncertainty Contribution**
Horizontal position slope (A_x)	2.80×10^6
Rotating mirror to camera sensor distance (L_0)	5.66×10^5
Rotating mirror to return mirror distance (L_R)	5.42×10^5

Table 2: Calculated Individual Contributions to Uncertainty in c

As expected, the slope A_x is the largest contributor to uncertainty. It is unlike the contributions from measured distances in that those quantities have smaller fractional uncertainties and are only concerned with the measurement error for a single recorded data. On the contrary, the slope is a product of the least squares regression line which must balance many residuals altogether. Smaller fluctuations in laser spot position and error in frequency measurement have echoing repercussions for the slope A_x, which also describes why it contributes so dominantly towards overall uncertainty δc. Let us now look to Table 3 for some analysis on this uncertainty in frequency δf, and how it too contributes towards uncertainty in slope δA_x via laser spot fluctuations. During experiment, a maximum uncertainty in frequency of

0.3 Hz was recorded. Contributions to the fluctuations of the laser spot position that come from the uncertainty in the frequency are calculated in the following way:

$$\delta x_f = \left| \frac{d\, x_{\text{model fit}}(f)}{df} \right| \delta f = |A_x|\, \delta f \tag{11}$$

Where $x_{modelfit}(f)$ is the linear fit. Find the resulting value in the second row of Table 3. Let us now calculate the observed fluctuation in the laser spot position:

$$(\delta x_{\text{observed}})^2 = \frac{1}{N-2} \sum_{i=1}^{N} \left[x_i - \left(A_x f_i + B_x \right) \right]^2 \tag{12}$$

Where A_x and B_x are the slope and intercept of the fit line, N is the number of measured positions of the laser spot, and x_i and f_i are the measured spot positions and corresponding frequencies. Find the resulting value in the third row of Table 3. We can also similarly evaluate the observed fluctuations in the vertical position, which produced the value in row 4 of Table 3.

$$(\delta y_{\text{observed}})^2 = \frac{1}{N-2} \sum_{i=1}^{N} \left[y_i - \left(A_y f_i + B_y \right) \right]^2 \tag{13}$$

Laser Spot Fluctuations

Quantity	**Value**
Observed max uncertainty in the frequency (Hz)	0.3
Expected fluctuation based on uncertainty in frequency (m)	2.13×10^{-7}
Observed fluctuations in horizontal position (m)	1.77×10^{-6}
Observed fluctuations in vertical position (m)	9.00×10^{-6}

Table 3: Fluctuations and Uncertainty in The Measured Laser Spot Position

Notice in Table 3 that the horizontal scatter is around 8.3 times larger than what frequency uncertainty alone would cause, and the vertical scatter is around 42 times larger. We can conclude from this phenomenon that frequency uncertainty is not the limiting factor here: something else is driving uncertainty in slope δA_x, δA_y. Optical/measurement effects such as spot shape, ghost fringes, alignment issues, centroid thresholding, or camera noise dominate the actual fluctuations in laser spot position around the least squares regression line.

The uncertainty in vertical position slope A_y was also calculated by the python script, reported in Table 4 along with A_x. Interestingly, the range of uncertainty δA_y is not large enough to encompass the theoretical value of 0. There are a few possible explanations for this trend. Perhaps this particular sample of frequencies consists of some irregular fluctuations in y position of the centroid that pull the slope away from 0 enough to not include it. Indeed, what can be observed in Figure 8 supports this conclusion. Up until about 210 Hz, the data seem to average out to a least squares regression line whose slope would be much closer to 0.

The last four data points however, past 260 Hz, are at considerably lower y positions and do seem to pull the slope to a more negative value. External factors that create asymmetries in the laser spot image like vibrations could be an explanation for this, although this specific example is unlikely. Perhaps mechanical wobble of the motor slightly changes the height of the beam as frequency increases: enough to exclude 0 from the interval of confidence.

Horizontal and Vertical Uncertainties in Slope

Distance	**Value (cm)**	**Uncertainty (cm)**
Horizontal Position Slope A_x	7.1×10^{-7}	7.0×10^{-9}
Vertical Position Slope A_y	-2.3×10^{-7}	3.6×10^{-8}

Table 4: Fitting Results

3.4 Validity of Results

We observe that the measured speed of light for a threshold pixel intensity of 47 is 2.84×10^8 m/s, with an associated uncertainty of 0.03×10^8 m/s. This is roughly a 5% deviation from the accepted value of $\approx 3.00 \times 10^8$ m/s. This deviation would be acceptable should the upper bound of the uncertainty in the experimental result contain the accepted value, but it does not: this upper bound is only 2.87×10^8 m/s.

Interestingly, if we loosen the aforementioned constraints set forth to dictate our threshold choice, we can indeed include the accepted value into the range of uncertainty. Let us now look at a threshold that completely satisfies constraint (1) but violates constraint (2) in that its sample size (viewable in Figure 12) takes on a much-less-inclusive, irregular form that risks pixel biasing and increases uncertainty in slope A_x. Such an intensity threshold would be that of Figure 12, which is 57. Applying this filter yields a measured speed of light: 2.94×10^8 m/s with an uncertainty of 9.81×10^6 m/s. In addition to only exhibiting a 2.0% deviation from the accepted value for the speed of light, c, the upper bound of this range of confidence is 3.04×10^8 m/s; inclusive of the accepted value for c.

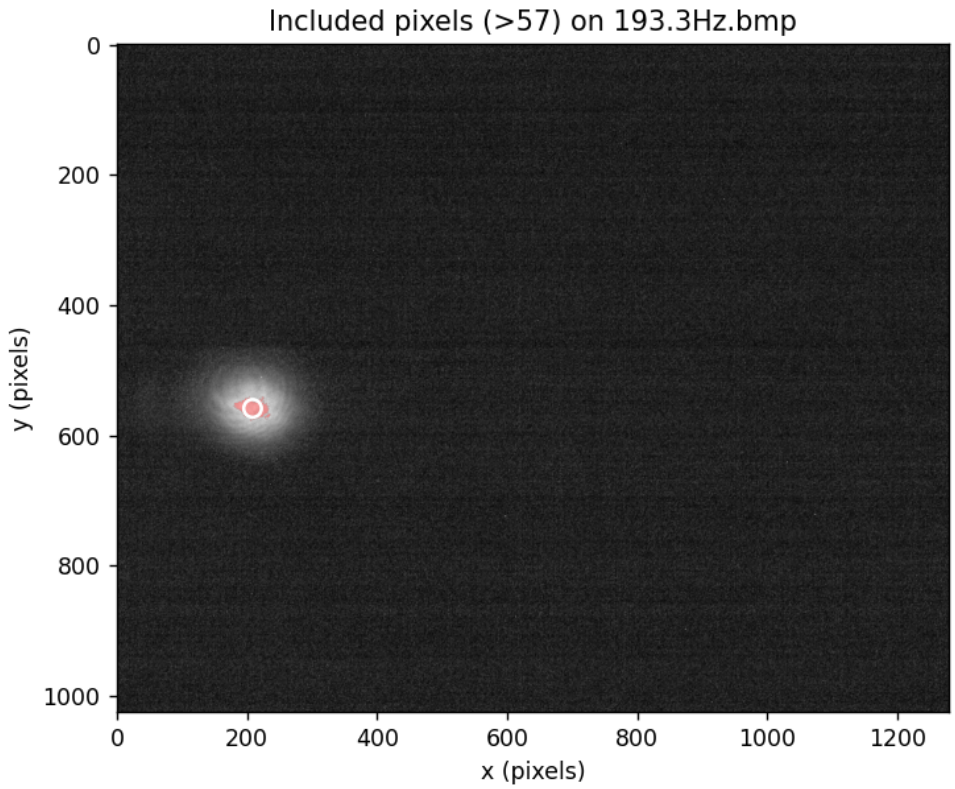

Figure 12: Pixel Sample Size for an Increased Threshold

We are not inclined to accept the results produced by the alternative threshold in Figure 12. Nor are we for any threshold above 47, for that matter. It is most important to note that we are only inclined to increase the threshold having known the accepted value for the speed of light. Should we not have known the true speed of light and first tested the higher intensity threshold of 57, we would instead be inclined to decrease it to include more pixels and assume a sample shape that more accurately reflects an ideal laser beam. Additionally, higher intensity thresholds exclude so much data that the weighted average method of approximating centroid position becomes almost obsolete in that most pixels within the central range are of a similar intensity.

We must look to some explanations for why the measured range of uncertainty for c does not include the accepted value. Limiting errors seem to be systematic (or quasi-systematic) effects that bias the fitted slope, not just random noise captured by δA_x. First, the observed position fluctuations greatly exceed what frequency uncertainty can explain, so δf cannot reconcile the $\sim 5\%$ shortfall. Second, even with the chosen threshold of 47–which was selected to balance artifact rejection and statistical stability–residual ring-like intensity distributions and image irregularities remain and can shift the centroid in a frequency-dependent way, subtly tilting the $x-f$ fit. Third, when the threshold is raised the accepted value falls within the enlarged error bar: but only because the sample becomes sparse/irregular and slope uncertainty inflates, or changes in lone data points pull the fit line. Although this alternative rejects the constraints and processing model, perhaps it may also avoid the irregularities

and off-center fringes that influence the threshold-47 model, allowing it to more closely approximate the speed of light.

4 Conclusion

The linear $x - f$ relationship from which $c = (2.84 \pm 0.03) \times 10^8 (\frac{m}{s})$ is produced is reliable and only suffers a $\sim 5\%$ deviation from the accepted value for c, but also provides insights into errors with the created experimental environment–which are investigated due to the accepted c value not falling within the measured range. Error propagation reveals A_x as the major contributor to uncertainty in speed of light δc, and fluctuation analysis confirms that spot position noise far exceeds the effect of the largest measured frequency uncertainty δf–implicating optical artifacts and/or centroiding choices as the main limiters. While higher intensity thresholds broaden error bars and (sometimes) manipulate slope to better approximate the true speed of light, these models degrade data quality and risk bias. The threshold-47 analysis is evaluated to be most defensible. Future improvements should target optical cleanliness/alignment and spot-model-based centroiding (perhaps Gaussian fits could offer better mitigation) to suppress unwanted shifts in A_x, which will both narrow uncertainties and reduce bias in the estimate of c.

5 Appendix

5.1 Beam Splitter

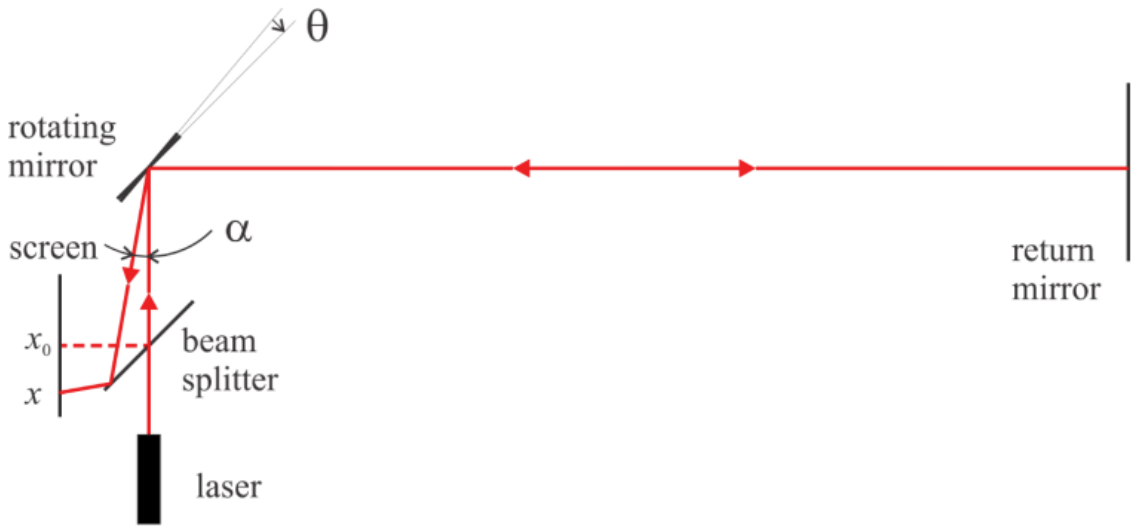

5.2 Python Script

```
"""
PHYSICS 264L - Lab 01
Zachary Maldonado
Due: 09/29/2025
```

```
"""
# Libraries I'll need
from pathlib import Path
import numpy as np
from PIL import Image
import matplotlib.pyplot as plt

THRESHOLD = 47  # ignore dim pixels (0  255   grayscale)

# Fixed list of my image data
image_paths = [
    Path("100Hz.bmp"), Path("123.6Hz.bmp"), Path("148.5Hz.bmp"), Path("164.3.bmp"),
    Path("182.15Hz.bmp"), Path("193.3Hz.bmp"), Path("208.65Hz.bmp"), Path("218.17Hz.bmp"),
    Path("231.58Hz.bmp"), Path("239.89Hz.bmp"), Path("250.85Hz.bmp"), Path("257.5Hz.bmp"),
    Path("266.62Hz.bmp"), Path("271.55Hz.bmp"), Path("275Hz.bmp"), Path("276.04Hz.bmp"),
    Path("283.88Hz.bmp"), Path("293.5Hz.bmp"), Path("304.7Hz.bmp")
]

# Function for the intensity-weighted centroid (center of mass)
def center_of_mass_xy(image_path, threshold=THRESHOLD):
    """
    Return (X_pos, Y_pos) for one image using intensity-weighted centroid.
    Converts to grayscale, thresholds, then uses np.average with weights.
    """
    arr = np.asarray(Image.open(image_path).convert("L"), dtype=float)
    mask = arr > threshold
    w = arr[mask]
    if w.size == 0:
        return np.nan, np.nan
    ys, xs = np.nonzero(mask)
    X_pos = np.average(xs, weights=w)
    Y_pos = np.average(ys, weights=w)
    return X_pos, Y_pos

# Process all images and store the x,y positions of each
x_positions = []
y_positions = []
for p in image_paths:
    X_pos, Y_pos = center_of_mass_xy(p, threshold=THRESHOLD)
    x_positions.append(X_pos)
    y_positions.append(Y_pos)
    print(f"{p.name}: center at (X={X_pos:.3f}, Y={Y_pos:.3f})")

x_positions = np.array(x_positions, dtype=float)
y_positions = np.array(y_positions, dtype=float)

print("\nX positions array:")
print(x_positions)

# Visualize last image with centroid overlay
```

```
last_path = image_paths[-1]
img_last = Image.open(last_path).convert("L")

plt.figure(figsize=(5.5, 5.5))
plt.imshow(img_last, cmap="gray", origin="upper")
plt.scatter([x_positions[-1]], [y_positions[-1]], s=50)
plt.title(f"Centroid on {last_path.name}")
plt.xlabel("x (pixels)")
plt.ylabel("y (pixels)")

# Frequencies (Hz)
freq = np.array(
    [100, 123.6, 148.5, 164.3, 182.15, 193.3, 208.65, 218.17, 231.58, 239.89,
     250.85, 257.5, 266.62, 271.55, 275, 276.04, 283.88, 293.5, 304.7],
    dtype=float
)

# measurement inputs and uncertainties (meters)
# l_m is laser rotating mirror
l_m = 1.003           # m
sigma_lm = 0.002      # m  (0.2 cm)

# l_R is (rot mirror intermediate) + (intermediatelens) + (lens return)
l_R_parts   = np.array([1.024, 0.976, 6.000])   # m
sigma_parts = np.array([0.002, 0.002, 0.015])   # m  (0.2 cm, 0.2 cm, 1.5 cm)
l_R   = l_R_parts.sum()
sigma_lR = np.sqrt((sigma_parts**2).sum())

# Pixel size (m/pixel) and convert positions to meters
px = 4.8e-6
xm_positions = x_positions * px
ym_positions = y_positions * px

# Linear regression (x vs f): slope m_x and covariance matrix
p, cov = np.polyfit(freq, xm_positions, 1, cov=True)
m, b = p
sigma_m = np.sqrt(cov[0, 0])   # uncertainty in slope from covariance

# Predicted fit line for plotting
freqlist = np.linspace(0, 310, 1000)
expfit = m * freqlist + b

# Linear regression (y vs f): slope m_y
p2, cov2 = np.polyfit(freq, ym_positions, 1, cov=True)
m2, b2 = p2
sigma_my = np.sqrt(cov2[0,0])
expfit2 = m2 * freqlist + b2

# Observed fluctuations about best-fit (residual RMS with N-2 in denominator)
N = freq.size
```

```
res_x = xm_positions - (m * freq + b)
delta_x_observed = np.sqrt(np.sum(res_x**2) / (N - 2))

# Same idea in vertical direction
res_y = ym_positions - (m2 * freq + b2)
delta_y_observed = np.sqrt(np.sum(res_y**2) / (N - 2))

# Fluctuation in horizontal position due to frequency uncertainty:
#  x_f  = | d x_model_fit / df |  f  = |A_x|  f   ; use largest observed f
DELTA_F = 0.3
delta_x_from_freq = abs(m) * DELTA_F

# Speed of light and propagated uncertainty
# c = (8*pi*l_m*l_R)/m
K = 8.0 * np.pi * l_m * l_R
c = K / m

# Error propagation for product/quotient:
frac_lm = sigma_lm / l_m
frac_lR = sigma_lR / l_R
frac_m  = sigma_m  / abs(m)
sigma_c = c * np.sqrt(frac_lm**2 + frac_lR**2 + frac_m**2)

print(f"\nMeasured speed of light: c = {c:.6e}    {sigma_c:.2e} m/s")
print(f"Slope A_x (x vs f): m = {m:.6e}    {sigma_m:.2e} m/Hz  (from covariance only)")
print(f"Slope A_y (y vs f): m_y = {m2:.6e}    {sigma_my:.2e} m/Hz")

# Report the lab-specified fluctuation metrics
print("\nLab-instruction fluctuation estimates:")
print(f"Observed uncertainty in the frequency (max) [Hz]: {DELTA_F:.3g}")
print(f"Expected fluctuation based on uncertainty in frequency  x_f  [m]: {delta_x_from_freq:.3e}")
print(f"Observed fluctuations in horizontal position  x_observed  [m]: {delta_x_observed:.3e}")
print(f"Observed fluctuations in vertical position  y_observed  [m]: {delta_y_observed:.3e}")

# Plots
plt.figure(figsize=(6.5, 5.5))
plt.plot(freq, xm_positions, marker="o", linestyle="none", label="data")
plt.plot(freqlist, expfit, 'r-', label="fit")
plt.xlim(100, 305)
plt.ylim(min(xm_positions)*0.98, max(xm_positions)*1.02)
plt.xlabel("Frequency (Hz)")
plt.ylabel("X position (m)")
plt.title("X position vs. Frequency")
plt.grid(True)
plt.legend()

plt.figure(figsize=(10, 6))
plt.plot(freq, ym_positions, marker="o", linestyle="none", label="data")
plt.plot(freqlist, expfit2, 'r-', label="fit")
```

```
plt.xlim(100, 305)
plt.xlabel("Frequency (Hz)")
plt.ylabel("Y position (m)")
plt.title("Y position vs. Frequency")
plt.grid(True)
plt.legend()

# Visualize which pixels are included in the weighted average
MASK_PREVIEW_INDEX = 5  # adjustable
preview_path = image_paths[MASK_PREVIEW_INDEX]
arr = np.asarray(Image.open(preview_path).convert("L"), dtype=float)
mask = arr > THRESHOLD

if np.any(mask):
    ys, xs = np.nonzero(mask)
    w = arr[mask]
    Xp = np.average(xs, weights=w)
    Yp = np.average(ys, weights=w)
else:
    Xp, Yp = np.nan, np.nan

plt.figure(figsize=(6, 6))
plt.imshow(arr, cmap="gray", origin="upper")
overlay = np.zeros((arr.shape[0], arr.shape[1], 4), dtype=float)
overlay[mask] = [1.0, 0.0, 0.0, 0.35]
plt.imshow(overlay, origin="upper")
if np.isfinite(Xp) and np.isfinite(Yp):
    plt.scatter([Xp], [Yp], s=60, edgecolors="white", facecolors="none",
   linewidths=1.5)
plt.title(f"Included pixels (>{THRESHOLD}) on {preview_path.name}")
plt.xlabel("x (pixels)")
plt.ylabel("y (pixels)")
plt.tight_layout()
plt.show()

# Contributions
contrib_lm = c * frac_lm
contrib_lR = c * frac_lR
contrib_m  = c * frac_m

print("\nIndividual uncertainty contributions to  _c  (m/s):")
print(f"  From L0 =  l a s e r rotating -mirror  : {contrib_lm:.2e}")
print(f"  From LR =  r o t a t i n g return  mirror : {contrib_lR:.2e}")
print(f"  From slope m (covariance only)   : {contrib_m:.2e}")
print(f" Quadrature sum                      : {sigma_c:.2e}")

print("\nCovariance matrix from polyfit (x vs f):")
print(cov)
print(cov2)
```