Verifying the Correlation between Parachute Surface Area and Terminal Velocity

Darren Apostol, Blake Dee

December 14, 2023

Abstract

This research paper explores the relationship between parachute surface area and terminal velocity, providing empirical evidence to support theoretical models in fluid dynamics. Utilizing a custom setup involving a TF Luna LiDAR sensor, an accelerometer, and an Arduino UNO, the experiment measured the descent of variously sized parachutes. Our findings revealed a clear inverse correlation between parachute surface area and terminal velocity, aligning with the hypothesis that larger parachutes experience greater drag and thus a reduced terminal velocity. The experimentally derived drag coefficient C_d closely matched the theoretical value, further substantiating the theoretical models. Additionally, an intriguing aspect of our research was the examination of spill holes in parachutes. Contrary to initial expectations, the presence of spill holes did not significantly alter the drag coefficient, suggesting their primary role in stabilizing descent rather than affecting drag. However, it should be stated that the conclusions made in this paper regarding spill holes are not conclusive as the analysis of spill holes was a secondary goal. These results not only reinforce fundamental principles of non-conservative forces in fluid dynamics but also offer practical insights for parachute design and aerodynamics, particularly in terms of stability and efficiency.

I. Introduction

Intermediate to even semi-advanced collegiate physics courses focus primarily on analysis of conservative forces such as gravity, spring forces, and electromagnetism. This trend follows as the conservation of mechanical energy within these conservative systems presents a simple yet intuitive foundation for further studies in physics. As such, analysis of non-conservative forces - forces that are path-dependent and often involve energy dissipation rather than conservation - poses a stimulating research topic for undergraduates who likely lack experience with such systems. Thus, this experiment set out to explore the dynamics of a complex non-conservative force - viscous drag - in the context of parachutes.

The objective of the experiment was to verify the established relationship between the coefficient of drag, surface area, and terminal velocity for a falling parachute-mounted object. We hypothesized in agreement with current findings that surface area would inversely correlate with terminal velocity and positively correlate with the force of drag experienced.

To explore this relationship, we delve into the underlying physics governing the motion of parachutes. At the core of our analysis is the concept of terminal velocity, a state where the acceleration of a falling object ceases and it continues to fall at a constant speed. This state occurs when the drag force (F_d), caused by the resistance of air, equals the gravitational force (F_g) acting on the object.

The drag force is described by the equation

$$F_d = \frac{1}{2} \cdot \rho \cdot v^2 \cdot C_d \cdot A \tag{1}$$

where ρ is the air density, v is the velocity of the object, C_d is the drag coefficient, and A is the cross-sectional area. The gravitational force is given by

$$F_g = m \cdot g \tag{2}$$

with m being the mass of the object, and g the acceleration due to gravity. At terminal velocity, these forces are equal, leading to the equation

$$\frac{1}{2} \cdot \rho \cdot v_t^2 \cdot C_d \cdot A = m \cdot g \tag{3}$$

Rearranging this equation for terminal velocity (v_t) yields

$$v_t = \sqrt{\frac{2 \cdot m \cdot g}{\rho \cdot C_d \cdot A}} \tag{4}$$

In our experiment, by altering the surface area A of parachutes, we aim to observe the resultant changes in terminal velocity. The hypothesis suggests that a larger surface area will increase the drag force, thus reducing the terminal velocity. This experiment provides a practical approach to understanding non-conservative forces, specifically viscous drag, and its implications in real-world scenarios, a vital aspect often overlooked in standard physics curricula.

II. Methods

Equipment

Windows laptop with Arduino IDE
TF LUNA LiDAR
HC-06 Bluetooth module
MPU 6050 Accelerometer
Arduino UNO R3 with expansion module
Arduino related cabling & battery pack
Plastic Parachute Material
Wooden Box and Sponge Cushioning

Setup

An ultrasonic sensor and Bluetooth module were wired to the Arduino and integrated into the assembly as shown below. The parachute was then attached to the structure via four symmetrical pulleys.

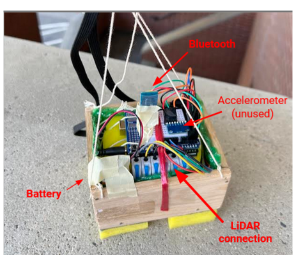

Figure 1: The Arduino and relevant sensors mounted in the cushion box.

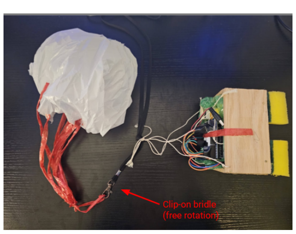

Figure 2: The parachute attachment to the cushioned box.

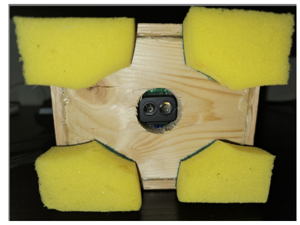

Figure 3: The LiDAR opening saddled within the cushioned box, facing from the bottom of the box.

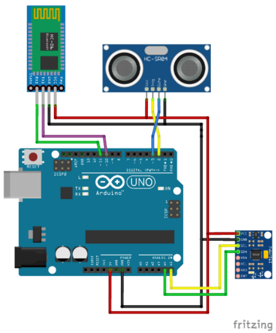

Figure 4: The wiring for the Arduino setup. Note that the sensor in this case is a LiDAR module, not the ultrasonic sensor shown, but the wiring is analogous.

The Bluetooth module allowed for a wireless connection between the Arduino and a nearby Windows laptop. This connection allowed the Arduino to send back raw data that was displayed on the serial monitor of the Arduino IDE. The code uploaded to the Arduino itself allowed for collection and Bluetooth transmission of the accelerometer and LiDAR data.

For the TF Luna LiDAR sensor in our experiment, the Arduino acquired data by sending a laser pulse from the LiDAR sensor and then detecting the reflection of that pulse from an object in front. The LiDAR measures the time delay between the emission of the pulse and its reception after reflecting off the object. This time delay is used to calculate the distance to the object using the equation

$$\Delta x = \frac{t \cdot c}{2} \tag{5}$$

where c is the speed of light in air (approximately 3×10^8 m/s). The sensor took readings at an interval of every 10 milliseconds, and each reading was displayed on the monitor as a pair of values: the time elapsed since the start of measurement (in milliseconds) and the measured distance to the object (in centimeters).

The accelerometer featured a micro electro mechanical system (MEMS) with a spring mass system that, when accelerated, moved the capacitor plate producing a change in voltage. This change was measured in the x-, y-, and z-directions and passed to the Arduino. Calibration for the accelerometer involved using past calibration constants for a previous experiment involving the same setup. However, in this experiment, while accelerometer data was collected, it was not used as all the readings maxed out i.e., the same value which was not useful.

First, square parachutes at 4 varying sized surface areas - 35x35, 50x50, 73x73, and 96x96 cm squared - were cut from plastic and punched with holes around the perimeter of the parachute. The parachutes were then attached to the Arduino setup as per above and released from a 2-story elevation (approximately 6 meters). A control was performed, involving dropping the arduino setup alone with no parachute. LiDAR and accelerometer data was collected and recorded. Next, spill holes were cut into the center of the parachutes to ensure cleaner airflow and a more stable fall. Procedures were repeated for the spill hole epoch and recorded in the same manner.

III. Results

The mass of the setup was $m = 0.397 \pm 0.001$ kg. The distance data (m) from the LiDAR of the setup was plotted against time (s). Because we were attempting to measure terminal velocity, only the last 3 meters of the approximately 7 meter drop were recorded per parachute trial. In the second graph, a fit line was obtained and plotted for each trial using NumPy's `polyfit`. Note that the figures below include both "with spill hole" and "without spill hole" trials for each parachute for completeness.

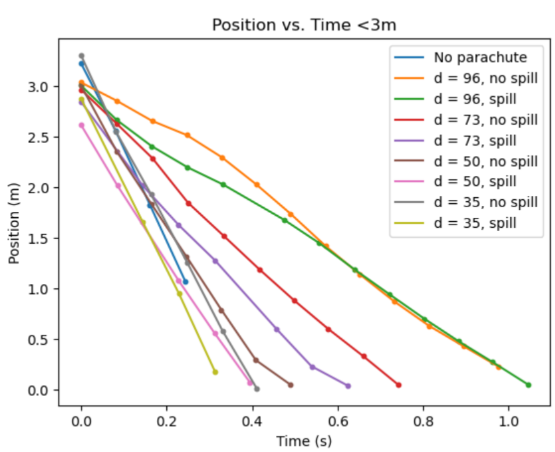

Figure 5: Position data of each parachute data plotted against time. Only the descent approximately at and below 3 meters is shown. d stands for diameter in centimeters.

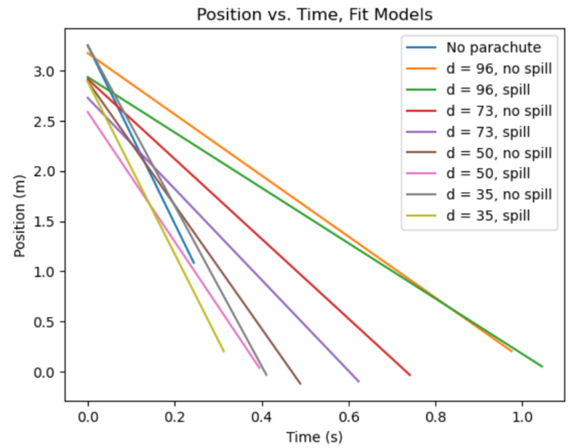

Figure 6: Fit lines for each parachute trial's position data.

The slope of each trial's fit line was obtained to get the velocity.

Diameter (cm)	Velocity (m/s)
96, no spill	3.05
96, spill	2.76
73, no spill	3.98
73, spill	4.54
50, no spill	6.17
50, spill	6.46
35, no spill	8.00
35, spill	8.59
No parachute	8.88

Figure 7: Table of velocities per parachute. Note how the velocity is inversely proportional to parachute diameter.

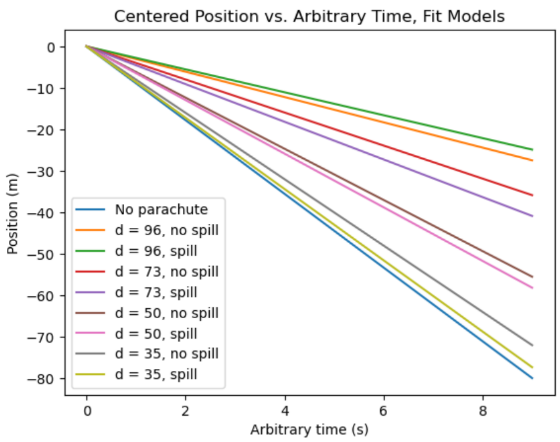

Figure 8: The velocity of each parachute visualized. Note how the plots get steeper the lower the diameter.

For the rest of the main analysis, only the trials without the spill hole i.e., trials with solid parachutes were considered. The spill hole data is discussed in the next section, as it was not our primary goal. Next, the velocity of each parachute was plotted against its area.

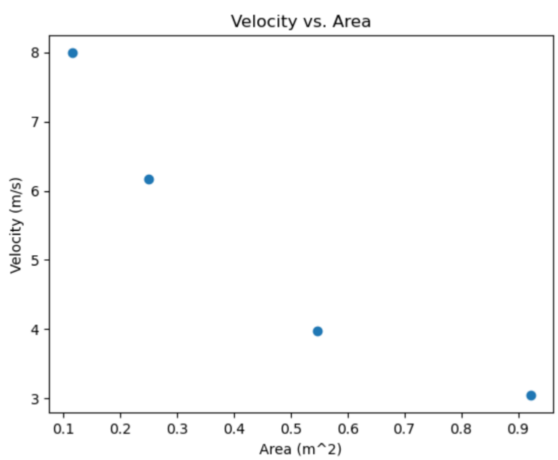

Figure 9: The terminal velocities of each parachute plotted against its area. Note the resemblance to an inverse square root graph.

To compare the data with theoretical values, we had to fit a curve to this experimental data and compare it to a theoretical curve. An experimental curve was generated via SciPy's `curve_fit` using the terminal velocity equation (?) as a model. For this model, we assumed that the density of air r, the acceleration due to gravity g, the opposing drag force, and the mass of the object m all remained constant during the descent of each parachute. We assumed that only the area A changed between different parachutes. The only variable that `curve_fit` could change to fit the experimental data was the coefficient of drag C_d, which we then extracted from the fit curve to obtain our experimental C_d. The theoretical C_d was decided to be 0.775 from the reference, and a curve was generated from that for visual comparison.

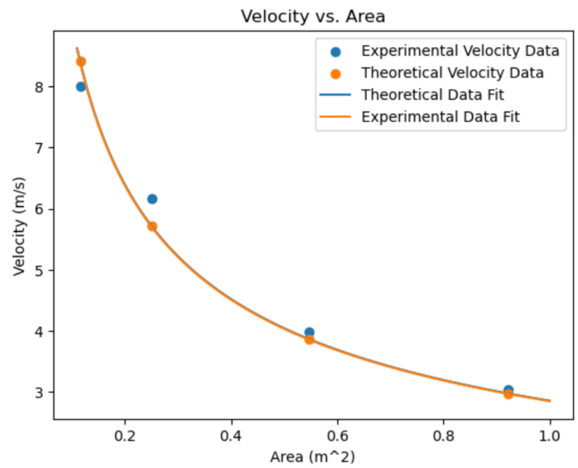

Figure 10: Experimental data and its fit curve along with theoretical "data" and its fit curve. The curves overlap and closely match, suggesting our data is very close to the theoretical.

The experimental C_d was determined to be 0.778 with a standard deviation of 0.050. The limits of acceptability were values within two standard deviations, resulting in an acceptance range of $0.678 < C_d < 0.878$. As the theoretical C_d of 0.775 fell within this range, we determined that the experimental and theoretical C_d matched and that our hypothesis was satisfied.

IV. Conclusion

In this experiment, we tracked the position of an Arduino setup during its descent with various different parachutes to verify the inverse square relationship between parachute area and terminal velocity. To do so quantitatively, we compared an experimentally obtained coefficient of drag C_d to a theoretical, C_d, 0.775, obtained from a reference. Using LiDAR, the distance between the Arduino setup and the ground was measured and plotted over the time of its descent. The terminal velocity per parachute was obtained by finding the slope of the latter half of descent (last 3 meters). These terminal velocities were plotted against each parachute's area, and a curve fit, based on the equation for terminal velocity, was applied. The experimental C_d was determined to be 0.778 ± 0.050, and as the theoretical C_d of 0.775 falls within two standard deviations of the experimental C_d, we determined that they matched. Thus, we quantitatively verified that terminal velocity has a precise inverse square relationship with parachute area.

There were a number of errors in the experiment. One major error is the lack of multiple trials per parachute. This was partially a time issue due to the setbacks we faced in experimental setup, but the integrity of the Arduino setup was also a cause of concern because of the repeated drops; if it broke or had permanent malfunctions, the experiment could have faced even more major setbacks. Because of this concern, only one trial was done per parachute, and after all trials were finished, the experiment was immediately concluded. However, in hindsight, the Arduino setup was remarkably sturdy, with no disconnections occurring at all during the parachute drops. It was certainly possible to collect multiple trials per parachute, even with the smallest parachute, with little to no damage to the Arduino setup as long as the drop and collision angles were consistent. Although the obtained data was sufficient for verification, multiple trials would improve the confidence of verification and the integrity of the experiment. Another source of error is the freedom of motion of the object and parachute during descent. If the Arduino setup, and thus the LiDAR, swayed, the positional measurements would have had undulations inconsistent with free fall. In fact, the full view of the data showed such undulations, particularly in the region slightly above the 3 m mark for most parachute data. These inconsistences were also part of the reason why only the data below the 3 m mark were considered for terminal velocity determination. Human error was also relatively abundant in this experiment, especially with the precise measurement and connection of shroud lines although they should not have a significant effect on parachute action.

Future experiments could explore the effect of spill holes on the coefficient of drag, which we briefly discuss in the next subsection. Spill holes of various sizes could be analyzed with respect to various parachute areas to determine how much the coefficient of drag changes. Other experiments could analyze the acceleration of the object as the parachute unfolds, particularly through an active parachute deployment system that would allow for controlled release of the parachute. This type of control would allow for analysis of accelerations at various velocities which could be helpful in developing safer parachutes. Another experiment could compare various parachute materials, such as fabric or paper, and see how the coefficient of drag or terminal velocity changes under the same parachute area but different material.

Spill Holes

During the course of our experiments, we had issues with the Arduino setup swaying during descent, which distorted positional measurements from the LiDAR as it started to measure diagonal distances. In the end, we solved it via making the setup heavier and making the drop shorter, but we also learned about spill holes. A spill hole is a hole at the center of parachutes which could stabilize an object's descent. They do so by giving turbulent air in the parachute an exit, instead of the air exiting chaotically out the sides. Additionally, spill holes apparently did not have significant effects on the drag coefficients despite being holes in the areas of parachutes. We thought that spill holes would be excellent for our use case, so we cut "pluggable" holes in our parachutes in case sway did occur with our new setup and drop height. In the end, our analysis proceeded with the trails without spill holes, mainly because a hole in the parachute went against the reference manual (flat, solid sheet). However, we decided that it would be an interesting secondary goal to see whether or not spill holes would have a significant effect on drag coefficient, so we collected the data anyway. The spill hole data was then analyzed in the same way as the no spill hole data, with the velocities plotted against area and a `curve_fit` being applied. The curve, as well as the no spill hole and theoretical curves, are plotted below.

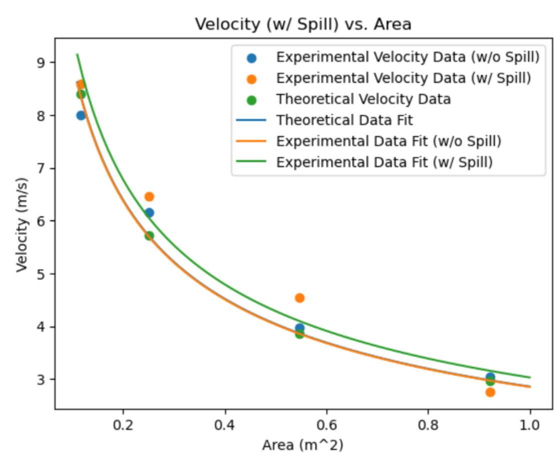

Figure 11: Spill hole velocity vs. area data plotted against no spill hole and theoretical data. Note how the spill hole curve seems shifted up.

The spill hole data had an experimental C_d of 0.689 with a standard deviation of 0.052. Interestingly enough, this falls within the acceptance range of the experimental no spill hole data (in fact, they fall within the acceptance ranges of each other), which makes it appealing to say that spill holes do not have a significant effect on the coefficient of drag. However, this analysis is by no means conclusive, and further experiments should be done to gain more insight on the effect of spill holes.