Pseudo Power Law Statistics in a Jammed, Amorphous Solid

A Senior Project

By

Jacob Brian Hass

Advisor: Dr. Nathan Keim

Department of Physics, California Polytechnic University SLO

June 25, 2018

Approval Page

Title: Pseudo Power Law Statistics in a Jammed, Amorphous Solid

Author: Jacob Brian Hass

Date Submitted: June 25, 2018

Senior Project Advisor: Dr. Nathan Keim

_____

Signature

_____

Date

Contents

List of Figures

1 Introduction

In this paper we will be studying a jammed, amorphous solid. It's important to first distinguish this from crystalline solids which have also been studied in great depth. In crystalline structures, it is clear that the solid deforms based on its structure or geometry [11]. However, amorphous solids are more disordered and the rules that govern how the solid deforms are still being investigated [11]. Although we will not be looking at the underpinnings of how the solid deforms, we will be investigating the statistics behind how the solid deforms under strain.

We will be looking at a jammed solid at strain amplitudes that are below the irreversibility transition. After many cycles of cyclic shearing, a particle's trajectory will start and stop in the same spot over one period of shearing. This creates a limit cycle or closed orbit. In contrast, above the irreversibility transition, a particle's trajectory will generally not start and stop in the same place but will wander; the solid can no longer form limit cycles. Although behavior above this transition is interesting, we will focus on strains below the irreversibility transition.

As our amorphous solid deforms, particles within the solid will rearrange, or slide past each other, so the solid reaches a lower-energy and lower-stress state. There are two types of these rearrangements which are reversible and irreversible rearrangements. Reversible rearrangements can be undone, or returned to their original state, by applying the reverse strain to the system [11]. This implies the system is hysteric, not to be confused with time reversible where we simply need to reverse the strain over time. Irreversible rearrangements are much more interesting in that these rearrangements cannot be undone by applying the reverse strain to the system [11]. These rearrangements cause the system to permanently drop to a lower energy level than before [11]. Above the irreversibility transition, there is an increase in the number of irreversible rearrangements [3].

These rearrangements are an example of crackling noise. Crackling noise is defined by a system relaxing through discrete events [12]. This crackling noise can be seen in numerous natural systems ranging from earthquakes to the polarization of a group of magnetic spins [12]. In our solid, the rearrangements described earlier are composed of a localized group of particles sliding past each other which cause the system to decrease in energy or relax. Therefore, our solid displays crackling noise through these rearrangements.

We specifically investigate the probability distribution of the sizes of events. Previous research has shown that rearrangements within the solid should follow a power law distribution; however, this work was done mainly computationally [2][11]. A power law distribution gives the distribution a number of unique properties most importantly its scale free property. Here, we take an experimental approach to confirm this power law with our data. We also discuss the effects that different driving frequencies and amplitudes have on the distribution. Lastly, since this material is known to reach limit cycles, we explore the effects that many cycles of shearing have on the distribution.

2 Experimental Setup

Our experiment involves a needle shearing a jammed solid, Figure 2, adsorbed at an oil-water interface . A wall mount holds two glass slides to confine the solid while Helmholtz coils move the needle as seen in Figure 1.

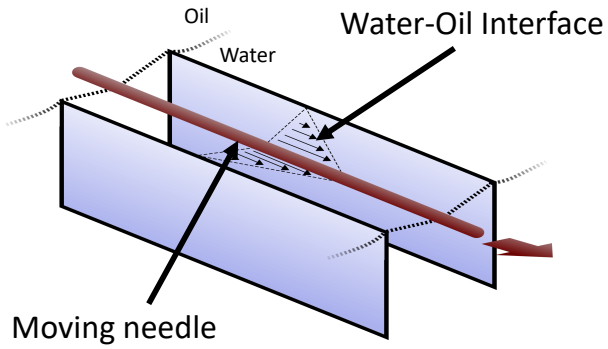

Figure 1: The oil-water interface where the solid is confined along with the needle and glass slides

Our solid consists of 3.4μm and 5.8μm plastic beads adsorbed at an oil-water interface as seen from a top view in Figure 2. The particles are held apart due to an electric dipole forming between each particle and the water [6]. Unfortunately, contaminants introduced surfactant causing the surface tension to lower [9]. This caused the particles to sit lower in the water and since the electrostatic repulsion happens via the oil, the repulsion is weakened [9]. With less repulsion, the particles tend to aggregate.

To get rid of any contaminants, we needed to make sure the particles were cleaned first. To clean the particles we centrifuged and sonicated them repeatedly; however, we later learned that centrifuging them may cause them to bunch together and form more aggregates than if we didn't centrifuge them. Further testing of this may be able to illuminate if centrifuging is beneficial or detrimental. Additionally, we found mixed results with sonicating the particles. Originally, we had sonicated particles multiple times in a ultrasonic bath. This produced varying solids that were somewhat aggregate free to solids with a lot of contamination. We then tested using a sonicating probe and found we could produce solids that were equal to or even better than levels of contamination in an ultrasonic bath. Again, further testing into the effects of using a sonicating probe versus a sonicating bath will be needed. We also found that there may have been freeze damage of the particles. This may have also caused the particles to aggregate more than expected.

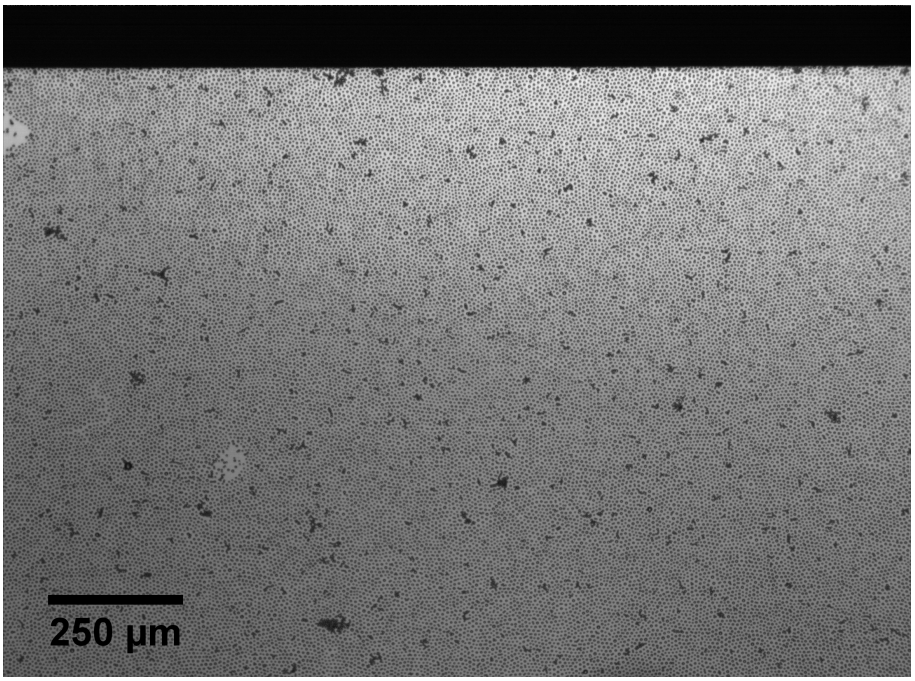

Figure 2: A top view of our substance of 3.4μm and 5.8μm plastic beads adsorbed at an oil-water interface

In order to confine the solid, we use an apparatus that holds two glass slides in the oil water interface seen in Figure 1. These slides act like walls for the solid and confine it to a finite area. We use a needle placed in the middle of these walls to shear the solid. The needle can be seen at the top of Figure 2 as

a black bar. We use Helmholtz coils to hold the needle in place with a strong magnetic field. Two extra wires are wrapped around one of the coils to perturb the magnetic field slightly. Using a DAC wired to an amplifier we can send a current through these extra coils to perturb the magnetic field and shear the material sinusoidally.

3 Analysis

To begin analyzing our data we tracked each movie using Trackpy [8]. This allowed us to find the coordinates for each particle throughout a movie and identify rearrangements. We then used this data to identify different rearrangements within the movie by grouping particles together that had rearranged close to each other in both space and time. We did this for each cycle, or period of the driving force, starting from the first cycle and continuing to the last cycle of each movie. Using these rearrangements, we constructed a histogram of the rearrangement sizes. Relying on the procedures outlined in Clauset 2009, we fit a power law to this histogram to obtain an accurate fit for the relationship between frequency of rearrangements and rearrangement size [1].

3.1 Identifying Rearrangments

In order to identify rearrangements we first needed to classify which particles rearranged and which particles had not. We used D^2_{min} as a measure of how non-affine a particle's movement was [4]. Particles can be classified as rearranged if $D^2_{min} \geq 0.015$ for a particle. We chose this number because it was above noise levels in the solid but low enough to tag particles that had actually rearranged.

Unfortunately, we found there were some spurtious particles being tagged as rearranged. These were particles that were not tracked correctly by Trackpy. Particles that are stuck together or in a glob are hard for Trackpy to track; therefore, the phenomenon is more common in dirty or contaminated solids. When Trackpy doesn't track a particle correctly its motion may look like it jerks from frame to frame in contrast to moving smoothly. For example, if two particles are stuck together Trackpy may track one particle in a frame and then switch to the other particle in the next frame. It then looks like the particle has moved signifigantly between frames. This causes the particle's D^2_{min} value to become elevated. Since D^2_{min} takes into acccount the position of particles around it, all the particles around the problem particle have elevated D^2_{min} values. This phenomenon is shown in Figure 3 where the center particle

was incorrectly tracked and the surrounding particles have elevated D^2_{min} values. Therefore, we had an increased number of rearrangements with size on the order of 10 to 20 particles. We were able to resolve this issue by removing any particle that could not be tracked over the entire movie. Filtering out these particles was able to greatly reduce the amount of noise seen in the affine field calculations.

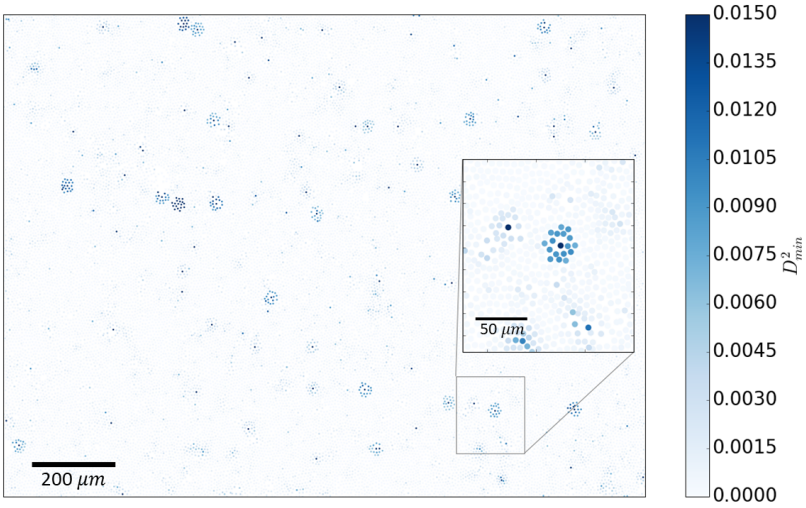

Figure 3: Our substance plotted by D^2_{min} values with a spurious event highlighted.

Now that we were able to classify rearranged particles, we needed to group them spatially and temporally. We began by identifying the time and position when a particle had rearranged. Using a kdtree algorithm, or a nearest neighbor search, we identified pairs of particles that had rearranged within 1.5 times the mean particle spacing and 1.2 s of each other. We chose 1.2 s because we had watched rearrangements occurring within the solid and found it to be roughly this time span. In other words, if r is the distance between two particles and Δt is the time between two particles rearranging, then two particles are grouped together if $r \leq 1.5 \times$ (mean particle spacing) and $\Delta t \leq 1.2$ s. Using the Networkx library in Python we then created a graph of every particle in the solid. Using the particles as nodes, we created edges between particles using the output pairs of the kdtree algorithm. This then created clusters of particles that were connected within the graph. These connected particles became rearrangements within the solid as shown in Figure 4. It appears that in Figure 4 some of the rearrangements overlap. This is because we plotted the particles at the location they rearranged so you could have rearrangements in the same place but at different times. It is important to note that we filtered out any particles that had already

rearranged when creating these clusters. We did this because we expected a particle to only rearrange once every cycle; therefore, we analyzed each data set cycle by cycle, starting at cycle 1, thereby resetting the particles that were filtered out. We were then able to identify the number and size of rearrangements throughout the movie.

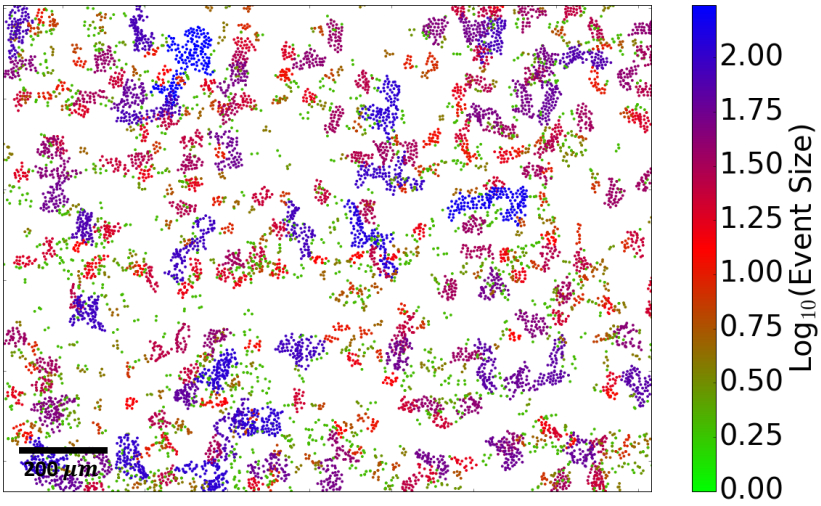

Figure 4: A plot of rearrangements for $f = 0.1$ Hz and $\gamma_0 = 0.06$. The needle sits at the bottom of this plot and shears horizontally. The particles are plotted at their locations when they rearrange which causes some rearrangements to appear to overlap.

3.2 Power Law Statistics

Power law statistics can be seen in a plethora of systems in nature spanning from earthquakes to word frequency [7]. Power-law statistics have the unique attribute that they are scale-free. Scale-free distributions imply a number of unique properties. One property is that it will look the same on any scale. Therefore, looking at a small part of the distribution will have the exact same shape as the overall distribution. Another is that scaling one variable by a factor will scale the other variable by the same factor to some power. We will be discussing a power law governed by

$$f(x) = Cx^{-\alpha} \tag{1}$$

where f is the probability density function (PDF), C and α are constants and x is the independent variable or event size in our case. As we can see there must be a minimum value x_{min} where the PDF follows

a power-law above $x_{\min}$. This is justified because when $x = 0$, the PDF can no longer be normalized. Additionally, the PDF cannot be normalized when $\alpha \leq 1$. We can now normalize the function for $x \geq x_{min}$ and $\alpha > 1$ [1]. For the continuous case, this leads to

$$f(x) = \frac{\alpha - 1}{x_{min}} \Big(\frac{x}{x_{min}} \Big)^{-\alpha} \tag{2}$$

For the discrete case, this leads to

$$f(x) = \frac{x^{-\alpha}}{\zeta(\alpha, x_{min})} \tag{3}$$

where

$$\zeta(\alpha, x_{min}) = \sum_{n=0}^{\infty} (n + x_{min})^{-\alpha} \tag{4}$$

is the Hurwitz zeta function [1]. Note that we will be using the discrete case to fit our function; however, the continuous case is included for completeness. We assume a discrete case since rearrangement sizes can only be integers. Additionally, we found that for our data, both the discrete and continuous PDF fit our data well after after finding α and x_{min} through a discrete fit.

With a theoretical PDF, we now need to find the best values of α and x_{min} to fit the given data. For a continuous PDF, for any given value of x_{min} we can find α using

$$\hat{\alpha} = 1 + n \left[\sum_{i=1}^{n} \ln \frac{x_i}{x_{min}} \right]^{-1} \tag{5}$$

where $\hat{\alpha}$ is an approximate value for α and n is the total number of data points [1]. For the discrete case we can approximate

$$\hat{\alpha} = 1 + n \left[\sum_{i=1}^{n} \ln \frac{x_i}{x_{min} - 0.5} \right]^{-1} \tag{6}$$

Additionally, we can calculate the error in $\hat{\alpha}$ using

$$\sigma = \frac{\hat{\alpha} - 1}{\sqrt{n}} + \mathcal{O}(1/n) \tag{7}$$

when we take the $\mathcal{O}$ term to be positive [1]. Although this is σ for the continuous case, σ for the discrete case is similar for large n and x_{min} [1].

In order to determine the accuracy of the fit, we need to define a function called the complementary cumulative distribution function (CDF). In contrast to the cumulative distribution function, $F(x)$ which is

the probability of a data point taking on a value below some x, the CDF, denoted $\bar{F}(x)$, is the probability of a data point taking on a value above some x. Therefore, the CDF is the integral of the PDF from x to ∞. This allows us to find the CDF for the continuous case to be [1]

$$\bar{F}(x) = \int_x^\infty f(x')\mathrm{d}x' = \left(\frac{x}{x_{min}}\right)^{-\alpha+1} \tag{8}$$

For the discrete case, we find

$$\bar{F}(x) = \frac{\zeta(\alpha, x)}{\zeta(\alpha, x_{min}} \tag{9}$$

Using our definition of the CDF, we can easily calculate empirically what the CDF will be for any given x.

To determine the best fit for our distribution, we need to find the best value of x_{min} to use. We will do this by using a Kolmogorov-Smirnov (KS) test [13]. This test essentially tries to minimize the residual between the empirical and theoretical CDF. Using Equation 9, we can define a residual between the empirical CDF, $\bar{P}(x)$, and the theoretical CDF, $\bar{F}(x)$. Note that for different values of x_{min} we will have a different set of remainders. We will denote the residual as

$$D(x) =\mid \bar{F}(x) - \bar{P}(x) \mid \tag{10}$$

where D is the magnitude of our residual [1]. For each value of x_{min} we will take the maximum of the function $D(x)$, to be as conservative as possible. From this list of values for D, we will find the smallest value, which we call the KS statistic, and its corresponding x_{min} will be our best estimation for x_{min} [1]. From this value of x_{min} we can then calculate the best α that fits the data.

With a best fit distribution for the power law we now need to determine how well the distribution fits our data. We will utilize a technique commonly used in non-parametric statistics of generating a large number of synthetic data sets and seeing if they fit a power-law distribution better or worse than our data set [1]. If we have N total data points we will generate the same number of data points in each synthetic data set. The data set can be broken up into two distinct regions: values that lie below x_{min}, the head, which follow an unknown distribution and values that lie above x_{min}, the tail, which follow our best fit distribution [1]. We can define the number of points in the head as n_{head} and the number of data points in the tail as n_{tail}. We will generate random data points in the head with probability n_{head}/N

and random data points in the tail with probability n_{tail}/N [1]. To generate random data in the head we will randomly choose a value from the actual data set that lies in the head. To generate random points in the tail we will invert the PDF so we have $x(f)$ for our best fit distribution. Then, we will randomly generate data points from this inverted distribution.

We will generate 2500 synthetic data sets using the procedures outlined and fit a new power-law distribution to each data set. We will calculate a KS statistic for each different data set. We can now define p to be the proportion of synthetic data sets with a KS statistic that is larger than the KS statistic from the original data set. When p is close to 1, the fitted power law is a good fit for the data since it is a better fit than most of the synthetic data sets, but when p is small, the power law does not fit the data well. Generally, if a p value is below 0.1 a power law distribution can be ruled out [1]. Even with a large p value, this merely indicates a power law fits the data well; however, there may be alternate distributions such as an exponential that also fit the data. It's important to note that $n_{tail} \geq 100$ to generate an accurate value for p or else there are simply not enough points to distinguish between random data and data fitting a power law [1].

4 Results

We began by fitting several sets of data with different parameters such as driving frequency, driving amplitude and periods of shearing, or cycle number, to see if a power-law distribution fits our data well. We didn't average over multiple cycles of the same dataset because the transient seemed to die off quickly meaning the first cycles were too important to be averaged, and in the steady state the rearrangements were very similar from cycle to cycle meaning averaging wouldn't change the distributions that much. If a majority of our p values with sufficient data were below the threshold of 0.1, then it would be clear a power law distribution does not accurately describe the system. Using the methods as outlined above in the "Power Law Statistics" section, we were able to fit a power law to various probability densities obtained from our experiment. We were able to do so for different driving frequencies, driving amplitudes and cycle numbers indicating a power law distribution was a good fit to our system. A typical example of a fit for a probability density can be seen in Figure 5.

Instead of discussing the probability density, sometimes discussing a histogram of rearrangement

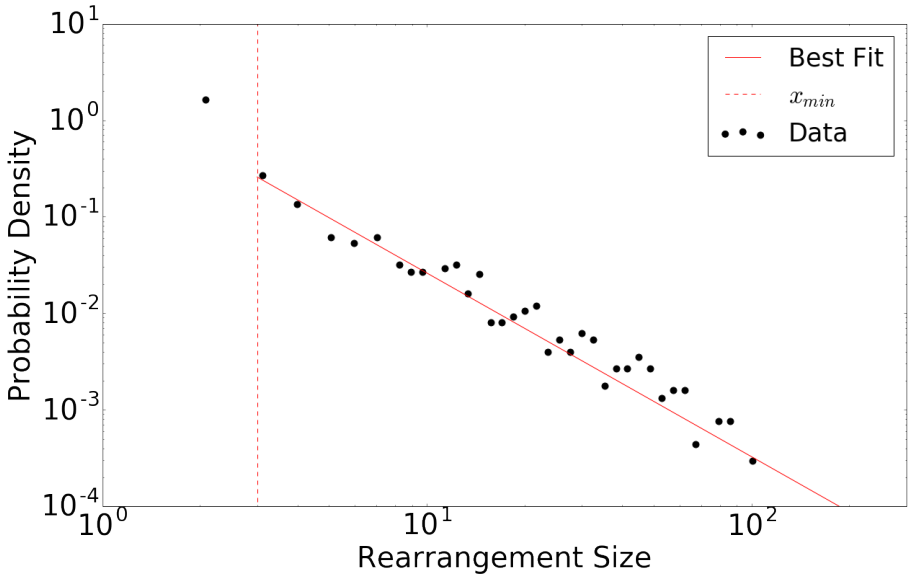

Figure 5: The probability density for cycle three of the data set of $f = 0.1Hz$ and $\gamma_0 = 0.06$. The exponent value was $\alpha = 1.90 \pm 0.04$ along with $x_{min} = 3$. The best fit follows the data well with a value of $p = 0.5$

sizes is more tangible. To transform our fit from a probability density to a histogram of frequency of events, we multiplied the probability density by a normalization factor. This normalization factor was calculated by taking the median of the ratio of the heights of the histogram and the probability density. The normalization factor was arbitrary in that it had little mathematical foundation, but the factor seemed to produce results consistent with the data. We also needed to account for the histogram being dependent on the bin width since the probability density is not. To account for this, we multiplied the probability density by the bin width. The bin width is essentially x or the size of the rearrangement. Recalling that the probability density is defined by $f(x) \propto x^{-\alpha}$, we can multiply by x to obtain $f(x) \propto x^{-\alpha+1}$. Therefore, changing α to $\alpha-1$ will produce the best fit for the histogram. Note that for some distributions this technique left us with a value of $\alpha < 1$. This poses a problem since $\zeta(\alpha, x)$ only converges for $\alpha > 1$. To work around this we reverted back to the continuous equations since they were a very good approximation for the discrete equations. A typical example of a best fit to a histogram can be seen below in Figure 6.

Now that we are confident that a power law distribution accurately describes our data, we can further explore the properties of amorphous solids. We will answer three key questions:

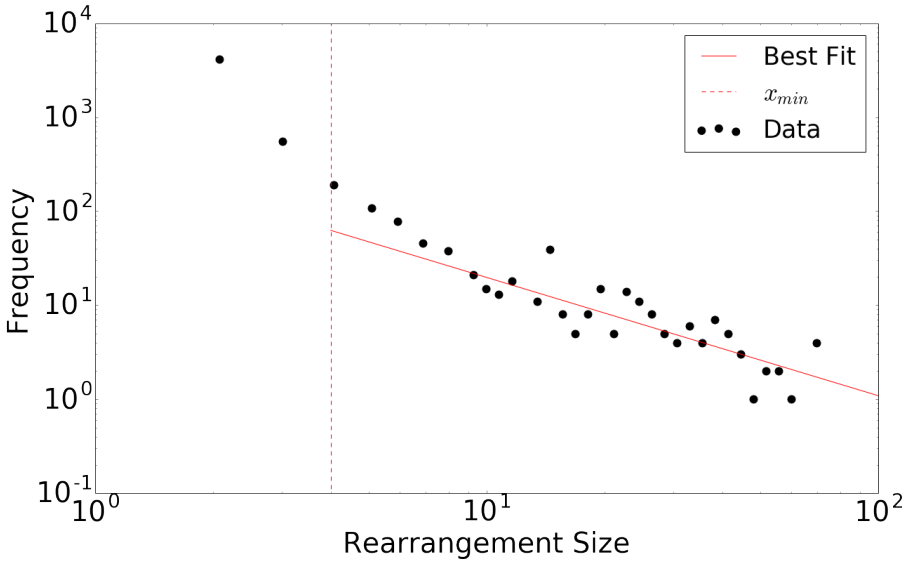

Figure 6: A histogram of rearrangement sizes for cycle one of the data set of $\gamma_0 = 0.06$ and $f = 0.01Hz$. The $x_{min} = 4$ value is shown as a dotted line. The exponent parameter was $\alpha = 2.26 \pm 0.05$ and we determined that $p = 0.7512$.

- How the histogram of rearrangements changes based on cycle number.
- How the histogram of rearrangements changes based on driving amplitude.
- How the power law distribution changes based on driving frequency.

We know the solid eventually reaches a limit cycle after many cycles, or periods of driving. Therefore, there must be a distinction between the transient and steady state. We should be able to see this in the histogram of rearrangements. As we go to a larger number of cycles, we should observe a smaller number of both larger and smaller rearrangements. We have seen in previous experiments that higher driving amplitudes will lead to more rearrangements, and we would like to confirm these results with multiple driving amplitudes [5]. Additionally, it is unknown what effect, if any, the driving frequency has on the system.

4.1 Cycle Number and Driving Amplitude

We began by simply plotting the histogram of rearrangements for the first and last cycle of different data sets. An example of these histograms can be seen in Figure 7. Comparing rearrangements in the

first and last cycle, it appears they follow the same overall distribution, but the later cycle has fewer rearrangements. This implies that there are irreversible rearrangements in the transient of each data set which do not appear in the steady state; however, even with this increase in events, the transient and steady state appear be drawn from similar distribution.

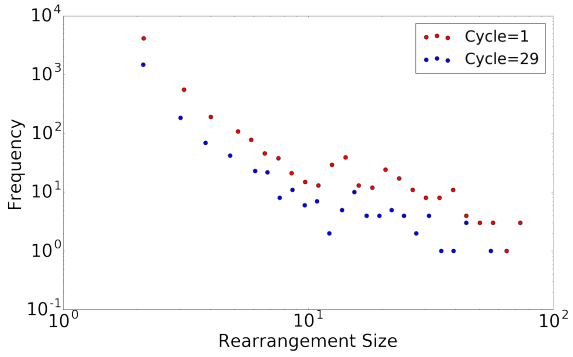

(a) A histogram of rearrangements for the first and last cycle of a data set with $f = 0.01Hz$ and $\gamma_0 = 0.06$.

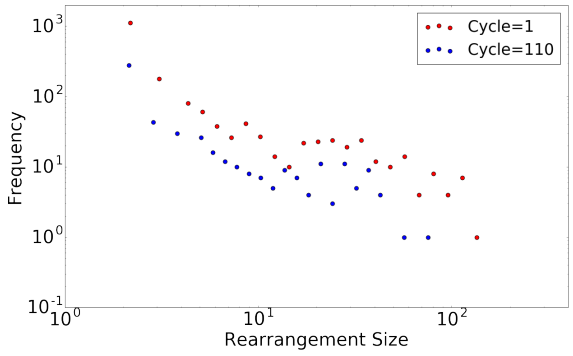

(b) A histogram of rearrangements for the first and last cycle of a data set with $f = 0.1Hz$ and $\gamma_0 = 0.06$.

Figure 7: A comparison of the histogram of rearrangements for different cycle numbers.

An easy way to classify this would be to look at the number of rearrangements larger than ten particles for each cycles. We chose the threshold of ten particles because anything over ten particles is very unlikely to be contaminated by noise and is generally more interesting due to its size. Below ten particles, we find that the number of rearrangements may vary slightly based on the method of identifying rearrangements. Using our analysis techniques, the number rearrangements may vary based on the spacial and temporal thresholds set especially in smaller rearrangements. However, above ten particles we expect these to be nearly the same regardless of analysis techniques used. We then plotted the rearrangements greater than ten particles for various cycles as seen in Figure 8.

From Figure 8, we can make a number of different observations. First, we can see there is a clear relationship between number of events and the cycle number. As we go out to larger cycles we see fewer rearrangements compared to earlier cycles as expected. This is evidence of a transient in both plots which eventually levels out to become a steady state. These results are consistent with previous observations of the system as they rely on the system reaching a limit cycle, or steady state [5].

We can also discuss the effect of driving amplitude on the number of rearrangements. As seen in Figure 8, it appears that higher driving amplitudes produce a greater number of rearrangements. Across

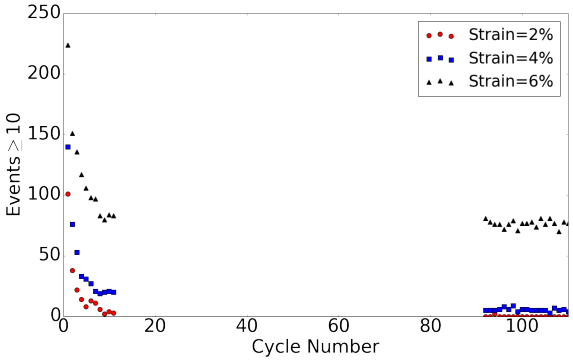

(a) Number of events larger than ten particles versus cycle number for different driving amplitudes at $f = 0.1Hz$

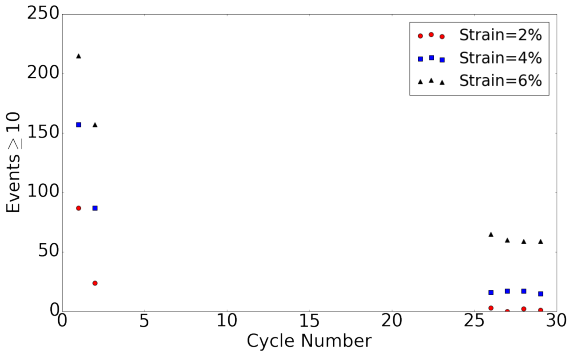

(b) Number of events larger than ten particles versus cycle number for different driving amplitudes at $f = 0.01Hz$.

Figure 8: A comparison of the number of rearrangements greater than ten particles versus cycle number

every single cycle for both frequencies, a higher strain creates a larger number of rearrangements compared to smaller strains. Again, this is consistent with prior results that larger driving amplitudes should produce more rearrangements [5].

4.2 Driving Frequency

To begin analysis of the the effects of the driving frequency on the distribution of rearrangements, we found the exponent value, α, for a variety of different strain amplitudes at $f = 0.1Hz$ and $f = 0.01Hz$ for various cycles within each data set. In order to ensure that α was representative of the data, we excluded any distributions where $p < 0.1$ and the number of rearrangements, n, above x_{min} was less than 100. As stated in the statistics section, distributions with $p < 0.1$ and $n < 100$ did not fit a power law well. We then plotted α versus strain which gave us the graph seen in Figure 9. Each point represents a value of α taken from cycles with strains of γ_0 = 0.2, 0.4, and 0.6. We chose not to distinguish between steady state versus transient values of α because they varied very little and there was no clear trend. Additionally, there are very few points at γ_0 = 0.2 and 0.4 because there were few rearrangements in the system which left most cycles with $n < 100$. There could be future research seeing if there is difference between α in the steady state or transient.

Looking at Figure 9, it appears that for each frequency, α stays in roughly the same range for $\gamma_0 = 0.02, 0.04,$ and 0.06. Therefore, systems with the same driving frequency will have similar probability densities for driving amplitudes below the irreversibility transition. The most interesting feature is that

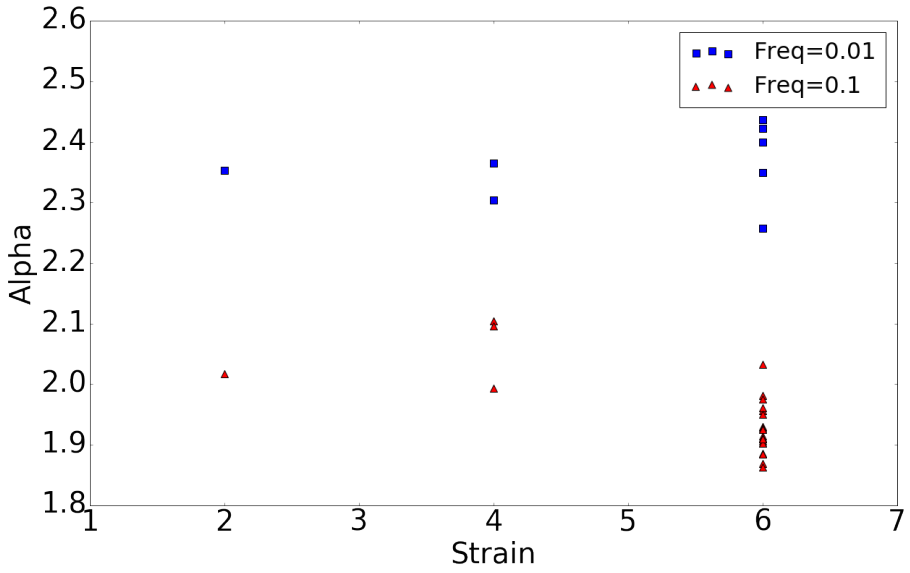

Figure 9: The exponent parameter α plotted as a function of strain amplitude.

for these different frequencies, α seems to be different. This suggests α depends on the driving frequency: as frequency decreases the value of α increases meaning the probability distribution will not extend as far into larger rearrangements.

In order to verify that Figure 9 describes our data, we plotted the rearrangement distribution for $\gamma_0 = 0.06$ with $f = 0.1Hz$ and $f = 0.01Hz$. This can be seen below in Figure 10.

Both distributions seen in Figure 10 look very similar because the difference in α between frequencies was small; however, it is important to remember that the values for α were distinctly different. It seems that for moderately sized rearrangements, roughly size 10 to 60 particles, the distributions are generally the same. The only differences occur at the upper and lower ends of the distributions. The lower end of each distribution are different but this is contaminated by noise so it may not be very significant. The distributions also seem to deviate slightly at the upper end of the distribution at rearrangements of roughly 100 particles. It looks like the distribution for $f = 0.1Hz$ extends further out into larger rearrangements than the distribution for $f = 0.01Hz$. This means that the higher frequency of $f = 0.1Hz$ may have a better chance of seeing larger rearrangements. This trend may hold for even higher or lower frequencies, but more investigation would be needed. Additionally, we plotted the probability density for each data set which can be seen in Figure 11. The probability density is very similar to the histogram,

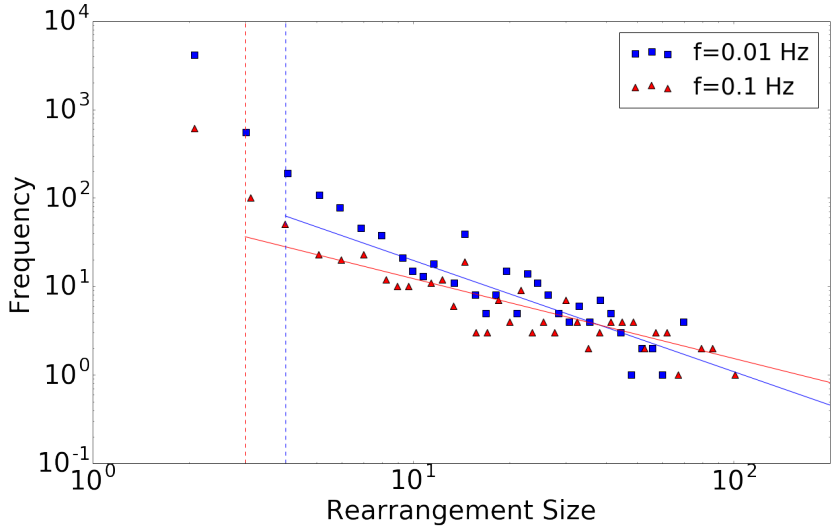

Figure 10: The histogram of rearrangements for $\gamma_0 = 0.06$ at $f = 0.1Hz$ and $f = 0.01Hz$.

but it is more evident that $f = 0.1Hz$ has a larger probability of seeing bigger rearrangements.

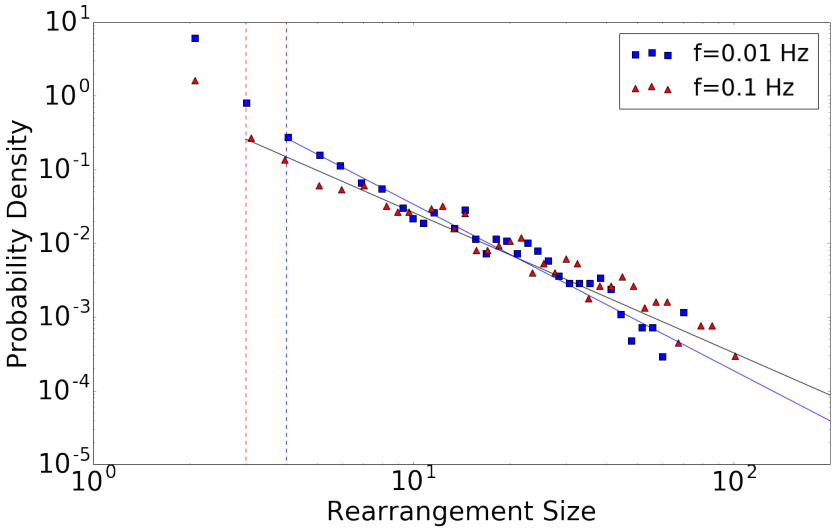

Figure 11: The probability density for $\gamma_0 = 0.06$ at $f = 0.1Hz$ and $f = 0.01Hz$.

5 Conclusion

By studying our jammed, amorphous solid, we were able to confirm the findings of previous studies that a power law distribution is a good fit for rearrangements within the solid. We also found that the

measured exponent value, α, of the power law is dependent on the driving frequency. As the frequency decreases, α appears to increase meaning there will be fewer large rearrangements. We also found the number of rearrangements within the solid depends on both the number of cycles of shearing and the shearing amplitude. As the number of cycles increases there will be fewer rearrangements. Although the transient has a larger number of rearrangements, the rearrangements in the transient and steady state appear to be drawn from the same distribution. It would be a project for future research to see if there is a difference in α in the transient and steady state. We also concluded that higher amplitudes yield more rearrangements.

6 Acknowledgements

I would like to thank Dr. Heather Smith of the Cal Poly Statistics Department for her help clarifying some statistics questions.

References

[1] Clauset, A., Shalizi, C., and M. E. J. Newman. Power-law Distributions in Empirical Data. *SIAM Review* **51**, 2009, 661-703.

[2] Regev, I., *et al*. Reversibility and Criticality in Amorphous Solids. *Nat. Commun*. 6:8805 doi: 10.1038/ncomms9805 (2015).

[3] Keim, N., and P. Arratia. Role of Disorder in Finite-Amplitude Shear of a 2D Jammed Material. *Soft Matter*, 2015, **11**, 1539.

[4] Falk, M., and J. Langer. Dynamics of Viscoplastic Deformation in Amorphous Solids, *Phys. Rev. E: Stat. Phys., Plasmas, Fluids, Relat. Interdiscip. Top.* **57** 7192-7205 (1998).

[5] Keim, N., and P. Arratia. Mechanical and Microscopic Properties of the Reversible Plastic Regime in a 2D Jammed Matterial, *Phys. Rev. Lett.* **112**, 028302 (2014).

[6] Masschaele, K., *et al*. Finite Ion-Size Effects Dominate the Interaction between Charged Colloidal Particles at an Oil-Water Interface, *Phys. Rev. Lett.* **105**, 048303 (2010).

[7] Zipf, George. *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology*. Addison-Wesley, 1949.

[8] Allan, D., *et al*. Trackpy. GitHub repository, (2016). DOI: 10.5281/zenodo.60550.

[9] Park, B. J. et al. Direct Measurements of the Effects of Salt and Surfactant on Interaction Forces between Colloidal Particles at Water-Oil Interfaces. *Langmuir* **24**, 16861694 (2008).

[10] Stamou, D., Duschl, C. and Johannsmann, D. Long-range attraction between colloidal spheres at the air-water interface: The consequence of an irregular meniscus. *Phys. Rev. E* **62**, 52635272 (2000).

[11] Maloney, C., and A. Lemaitre. Amorphous Systems in Athermal, Quasistatic Shear. *Phys. Rev. E* **74**, 016118 (2006).

[12] Sethna, J., Dahmen, K., and C. Meyers. "Crackling Noise." *Nature* **410**, 242-250 (2001).

[13] *NIST/SEMATECH e-Handbook of Statistical Methods*, https://www.itl.nist.gov/div898/handbook/eda/section3/eda35g.htm, 2 June 2018.