

THE UNIVERSITY OF CHICAGO

GRANULAR MATERIALS AS A MODEL FOR HOW DISORDERED SYSTEMS
ADAPT

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
KIERAN ALEXANDER MURPHY

CHICAGO, ILLINOIS

AUGUST 2019

Copyright © 2019 by Kieran Alexander Murphy
All Rights Reserved

For Braden, Lila, and Melanie

The ship wherein Theseus and the youth of Athens returned from Crete had thirty oars, and was preserved by the Athenians down even to the time of Demetrius Phalereus, for they took away the old planks as they decayed, putting in new and stronger timber in their places, insomuch that this ship became a standing example among the philosophers, for the logical question of things that grow; one side holding that the ship remained the same, and the other contending that it was not the same. – Plutarch, *Theseus*

TABLE OF CONTENTS

LIST OF FIGURES	vii
LIST OF TABLES	ix
ACKNOWLEDGMENTS	x
ABSTRACT	xii
1 INTRODUCTION	1
1.1 Quakes abound	2
1.2 Slow relaxation	4
1.3 Granular materials	5
2 EXPERIMENTAL METHODS	7
2.1 Instronduction	7
2.2 Particles	10
2.3 Plasticity tests	10
2.4 Strain-controlled versus stress-controlled compression	12
2.5 Stress relaxation tests	13
2.6 X-ray radiographs and computed tomography	14
2.7 Surface metrology	14
3 PLASTIC DEFORMATION	16
3.1 Critical state	17
3.1.1 Diffuse failure, not shear banding, due to size	18
3.2 Angle of internal friction	19
3.3 Volatility	21
3.4 A new phase space for plasticity data	23
3.5 Summary	24
4 STRESS DROPS AND REARRANGEMENT CASCADES	26
4.1 Introduction	26
4.2 X-rays link stress drops to restructuring of the packing	28
4.3 Drop-finding with a high pass filter on the first difference	31
4.3.1 Checking the method on synthetic data	33
4.4 Distribution fitting and uncertainty estimation	39
4.5 (τ, s^*) for each particle shape	42
4.5.1 Shapes present a new way to vary s^* and the distance to criticality	45
4.5.2 Microslips constitute the cascades of rearrangement	46
4.5.3 Volatility as a method to approximate s^*	48
4.5.4 Implications of τ : How do cascades proceed?	50
4.6 Strain rate independence of τ	51
4.7 The other component of plastic deformation: sliding	54
4.8 Summary	57

5	STRESS RELAXATION AND MEMORIES	59
5.1	Introduction	59
5.2	Single step stress relaxation	59
5.2.1	Measuring $\dot{\epsilon}_{\text{plastic}}(q)$ to fit stress recharges in plastic deformation . .	61
5.3	Distribution of relaxation times and the Amir-Oreg-Imry model	63
5.4	Two-step relaxation	66
5.4.1	Explanation with AOI DRT	69
5.5	N-step relaxation	71
5.5.1	Three-step stress relaxation	73
5.5.2	Experimental details that make more relaxation steps difficult	74
5.6	Outgrowing the model	76
5.6.1	Finite duration loading	76
5.6.2	The extra fitting parameter for N-step relaxation: Expedited aging .	80
5.7	Complex inner state	83
5.7.1	Particle sliding as a mechanism for relaxation	84
5.7.2	Connections to plastic deformation	87
5.8	Self-erasing memory	88
5.9	The mysterious absence of self-organized criticality in Chapter 4	89
5.10	Summary	90
6	PARTICLE SHAPE, SURFACE ROUGHNESS, AND THE DEGRADATION OF SHEAR STRENGTH	92
6.1	Degradation of aggregate strength with repeated plastic deformation	93
6.2	Brick-laying configuration and orientational ordering	94
6.3	Surface smoothing	99
6.4	Opening the door to granular surface forensics	102
7	CONCLUSIONS	104
	REFERENCES	106

LIST OF FIGURES

2.1	The Instron and the stress state of the packing	8
2.2	The Instron's finite response time	9
2.3	Scanning the surfaces of a statistically relevant sample of lens particles	15
3.1	Raw plasticity data for five shapes	17
3.2	Shear strength versus deviations from a sphere	20
3.3	Plastic deformation phase space	24
4.1	Localized slips revealed by X-ray videography	29
4.2	The need for a better drop-finding method, seen in disk data	32
4.3	Drop magnitude correction for the effect of the high pass filter	34
4.4	Synthetic time series used to test the drop-finding methods	35
4.5	Drop-finding method comparison: accuracy of individual drops	37
4.6	Drop-finding method comparison: τ measurements on synthetic data	38
4.7	Confidence region estimated through bootstrapping	40
4.8	Fitted stress drop magnitude distributions	43
4.9	(τ, s^*) parameter space for the truncated power law fits to $D(s)$	44
4.10	Volatility as an event-agnostic method for measuring s^*	47
4.11	Approximating τ with the moments of the distribution of fractional changes in $q(\epsilon)$	49
4.12	(τ, s^*) fits for over 100 datasets	52
4.13	Strain rate dependence of τ	53
4.14	Detectable event rate for different particle shapes	55
4.15	Comparison of raw $q(\epsilon)$ data for 4:1 disks and 45° lenses	56
5.1	Stress relaxation in a packing of glass beads	60
5.2	Stress recharge after drop events, fitted two ways	62
5.3	Two-step relaxation	67
5.4	Memory in the turnaround of a two-step relaxation	69
5.5	Snapshots from a simulated AOI DRT model exhibiting two-step relaxation	70
5.6	Three-step relaxation in a simulated AOI DRT system	72
5.7	Three-step relaxation in a packing of glass spheres	73
5.8	Memories in the turnaround points of a three-step relaxation	74
5.9	Lab temperature swings of 1-2°F limit the duration of relaxation experiments	75
5.10	Two-step relaxation with $t_w = 0$	77
5.11	Fitting N-step relaxation with a series of logs	80
5.12	Identical start, different finish in two experiments	81
5.13	The granular nature of the relaxation	86
5.14	Self-organized criticality in the Random Field Ising Model with relaxation	89
6.1	Rapid loss of shear strength for lens particles	93
6.2	Fitted tomography of lens particles, cross section	94
6.3	Finding the actual particle size in a reconstructed tomography through dilation	96
6.4	Spatial contact density distributions for lens assemblies	98
6.5	Surface scan of a freshly 3D-printed lens particle	99

6.6	Evolution of RMS roughness S_q for body and edge of lens particles	101
-----	--	-----

LIST OF TABLES

2.1	The particle shapes used in this thesis	11
-----	---	----

ACKNOWLEDGMENTS

Of greater value to me than any impact in physics this research might have is the growth as a scientist and as a human that accompanied it. These acknowledgements are for the people who left their mark on the process that brought me to where I am today. The list will be incomplete, but I'll do what I can in a page or two.

First, graduate school was often tough and the most significant changes were seldom pleasant to undergo. Sharing the process with the fun people of my cohort made it all much more enjoyable. Bike physics with Claiborne, weekly squash matches with Tom, journal club with Aaron, high rise living with Gus and Patrick, and Forget the Year planning with nearly the whole class were highlights I will remember fondly.

Next is the research environment of the James Franck Institute, and specifically Heinrich Jaeger's lab. The older students, especially Marc Miskin, taught me how to be a soft matter physicist through patient mentorship and by example. Only recently have I begun to understand what defines soft condensed matter physics, whose purview appears so broad as to inhibit a clear definition. However, I didn't need this understanding to be pulled in by the excitement and inquisitiveness of the students in the Irvine, Nagel, and Jaeger labs. There were, fortunately for me and unfortunately for these acknowledgements, too many to even attempt a comprehensive list. Leah Roth deserves special mention for a sense of humor that left me gasping for air on countless occasions. On the professor side of things, Sid Nagel and Tom Witten, through salons and bag lunches, helped create an atmosphere where it was fine to appreciate both beauty and ridiculousness in what you do, and to ask 'dumb' questions.

My advisor, Heinrich, had a profound influence on my grad school experience. He gave me room to explore, with a quiet, implicit faith that things would work out, even when I would quite regularly come to his office, panicked, with some perceived flaw in my data or analysis. To the degree I found my scientific identity, I owe the majority of it to this freedom of exploration. He pushed me through my first project, on the cool properties of granular packings of Z-shaped particles (not featured in this thesis), which helped get the

ball rolling on everything that followed. Heinrich also kept research vibrant through mixing with outside fields, which gave me my best research stories of the past six years: a six foot tall arch out of twisted Z particles cut from recycled plastic planks, a virtual reality project in collaboration with an artist at the School of the Art Institute of Chicago, and a swarm robotics collaboration called Jamoeba. Heinrich maintains a balance of impressive breadth and depth in science, and I'm still not quite sure how he manages it all.

Finally, there are the loved ones who shared in my joys and successes, and bolstered me in times of discouragement. I was exceedingly fortunate to have two parents who also earned PhDs and could speak to the hardships of research that seem to be independent of time period and field of science. My siblings, to whom this thesis is dedicated, have been growing into adults at the same time I have, and have each left their mark on my identity as a person and as a scientist. Then there is the wonderful girl I met after a year of living in Chicago. Pearl has been there, through the good times and the bad, with enough warmth to survive even the coldest of Chicago winters. She has helped me learn how to grab the reins, so to speak, and make life an adventure. I am so much more excited for the future after meeting her.

Thank you to everyone who added color to this chapter of my life.

ABSTRACT

Plastic deformation is an irreversible change in form, occurring when a material internally reconfigures to relieve mechanical stress. We study plastic flow in granular materials, where the change in state proceeds via particle rearrangements that occur over a broad range of length and time scales. When subjecting a granular material to a constant rate of deformation, it restructures via rapid cascades of particle movement thought to be universal with quake events in other amorphous materials. We 3D-print granular packings composed of different particle shapes, and measure the distribution of cascade magnitudes during uniaxial compression. We find that the characteristic size of the rearrangement cascades can be varied by over two orders of magnitude with particle shape. The underlying physics of the cascades, as measured by the power law exponent of their magnitude distribution, remains invariant. Its value is measured to be more compatible with a mean-field theory of plastic deformation than with one which accounts for the Eshelby treatment of stress redistribution in an elastic medium. When deformation stops and a granular system is held at a constant strain, it is found to relax approximately logarithmically in time, suggesting a complex inner state capable of storing memory. Via alternating compression and decompression steps we are able to imprint multiple memories that play themselves out through nonmonotonic stress relaxation. We are able to explain many aspects of the behavior with a model recently applied to the same memory behavior in crumpled elastic sheets. The model, which approximates the system as an ensemble of simple relaxing elements with broadly distributed timescales, becomes inaccurate when fitting to the form of the relaxation. The nature of the inaccuracy has implications for a wide class of systems exhibiting anomalous, non-Debye relaxation. Finally, we observe granular rearrangement events during the stress relaxation, suggesting the complex inner state is that of glacially slow plastic flow continuing hours past the cessation of compression. The various parts to this story show that a granular material adapts in ways large and small, fast and slow, and that the process has fundamental and far-reaching commonalities with many other systems in nature.

CHAPTER 1

INTRODUCTION

$\pi\alpha\nu\tau\alpha\ \rho\epsilon\iota$ (*Panta rhei.*) (*(Everything flows.)*)

– Heraclitus

A person never enters the same river twice, in another phrasing of *Panta rhei*, attributed to Heraclitus. No, the idea is not that everything is physically flowing and yield stresses are relative to the timescale considered, though that may be true [1]. Instead, the idea is that the world is in constant flux. Everything changes.

In this thesis, we study fundamental aspects of change as it pertains to disordered physical systems adapting to external influences. The external influences are applied forces, which cause irreversible changes in the physical form of the system. In the experimental work of this thesis, the deformation is at times smooth and continuous, and at others it surges with short-lived activity. The central questions are, first, what governs the statistics of the surges? Second, how can the internal state of the system be represented during the slow, gradual process of deformation?

Our model system is a granular material — an aggregate of millimeter-scale grains — whose state is held entirely in the spatial configuration of the grains. The system evolves from one configuration to another through grain rearrangements in response to externally imposed mechanical stresses. That the state of a granular system is so cleanly defined in this way allows change to be cast in concrete terms and measured in the laboratory.

The size scale of the grains in a granular system presents unique opportunities when studying the dynamics of a disordered physical system in experiments. We are able to vary the system at the grain scale and study the effects on the rearrangement process. We are also able to track the dynamics of individual grains and measure changes in their properties. A rich picture of how granular systems adapt results from the synthesis of the different insights into the dynamics. Several aspects of the physics are found in many other systems in nature,

making granular materials an excellent test bed for understanding how disordered systems undergo change.

Experimentally, we impose on the granular system global deformation and study its response through the stress it supports. The granular packing deforms erratically, with qualities that are strongly dependent on the geometrical shape of the grains (Chapter 3) [2]. In this process, the grains rearrange in rapid bursts of cascading shifts that span a wide range of length scales, from that of a single particle to large swaths of the system (Chapter 4) [2]. Reconfiguration continues over timescales of seconds to longer than what is able to be probed experimentally (Chapter 5). Throughout the process of deformation, the grains degrade via smoothing of surface features, though interestingly the changes are most easily perceived through the evolution of the overall mechanical properties of the system (Chapter 6) [3]. In each of these studies, we use the aggregate response as a lens into the internal dynamics of a disordered system as it is adapting.

1.1 Quakes abound

Many systems restructure and adapt with relative quiescence punctuated by sudden ‘quake’ events, with earthquakes being the most familiar [4]. Neutron stars emit sudden bursts of energy as their crust reforms, with statistics similar to those of earthquakes here on our planet, though some 20 orders of magnitude larger [5]. The brain is a heavily connected network of neurons which can surge in activity in the form of a seizure, dubbed “quakes of the brain” [6]. In finance, surges of trading activity are conjectured to arise from so-called “change blindness” whereby small changes are ignored and large ones are met with overreaction [7]. Evolution has been found to proceed through spurts of change with much less genetic variation in the periods between, termed *punctuated equilibrium* [8, 9]. Finally, the well-studied Barkhausen noise is a crackling of quake-like events as a magnet adjusts to a changing field through surges of domain flips [10].

Quake-like behavior can be produced at the breakfast table (in the style of [11]) by

slowly adding to a spilled drop of water. The outer edge of the droplet appears static much of the time, pinned by small sticking points on the surface of the table. Every so often, the weight of the water will cause the outer boundary of the droplet to surge forward and assume a new form [12]. This behavior is slightly abstracted in interface depinning, a model system commonly referred to in the study of quake behavior in materials [13, 14, 15]. In interface depinning, a d -dimensional interface is pulled through a $d + 1$ -dimensional medium containing some form of pinning [16]. The presence of disorder in the pinning causes the interface to become rugged as it is held back in some places more strongly than in others. Elastic interactions in the interface try to keep it smooth, and surges of motion arise due to the competition between the two. When one part of the interface moves forward, it may set in motion a cascade by freeing neighboring parts to move as well. After such a cascade runs its course, everything halts until energy builds up enough for the next event.

The key ingredients are (1) interactions between the constituent parts of the system that seek to homogenize the response, and (2) disorder in the amount of resistance different parts of the system feel toward changing their state. If (1) dominates over (2), the system responds as a single cohesive unit, and if (2) dominates over (1), the system behavior is that of many constituent pieces. When (1) and (2) compete, the dynamics of the system occur over many length and energy scales. The generality of these conditions suggests an explanation for the ubiquity of quake-like response to change. Studying the surges of activity informs us about the nature of the two ingredients.

A seemingly disparate array of materials undergo plastic deformation similarly, with intermittent dynamics and broadly distributed quake magnitudes. These materials include ice [17, 18], foam [19, 20], bulk metallic glass [21, 22, 23], metal nanocrystals [24], granular materials [25, 26, 27], and even knitted fabric [28]. In all of these systems, there is some form of disorder at the microscale and communication between constituent pieces of the system through mechanical stresses. The constituent pieces range from atoms and molecules to bubbles and droplets all the way up to grains and rocks, with system sizes spanning

an astonishing twelve orders of magnitude in length scale [26]. The physics of the plastic deformation of these systems is thought to be of a dynamical phase transition between a static and a flowing phase as a function of applied stress [29]. The nature of this transition as well as its universality is under debate [30, 15, 31], with its lack of resolution due in large part to the precision necessary in experiments to settle the matter, and to the difficulty in comparing results across different systems.

In this thesis we come at the issue from a different angle: rather than comparing experimental results from different systems, we modify one system significantly and study the differences. Specifically, we vary microscale details of a granular system by changing the shape of the particles, and find it has a strikingly influential effect on the nature of plastic deformation [2]. The wide range of behavior allows us to sift the aspects of granular plasticity that are dependent of the microscale details from those that are not. We measure a power law exponent characterizing the quake magnitude distributions which appears independent of particle shape and enters into the debate about the nature of plasticity.

1.2 Slow relaxation

What happens after the external driving ceases? How does the proverbial dust settle? Relaxation is the means by which a system finds its way to a new state with only the energy it already has – i.e., without continued driving. Dating back to Kohlrausch, who noticed in 1847 that the discharging of a capacitor did not follow a simple exponential [32], it has long been known that relaxation is a complex process which can be very revealing about the internal dynamics of a system.

Change is difficult when there are many possible routes to go about it and too little energy to try them all. In physical systems with so-called *glassy* dynamics, change from state A to state B is a slow, arduous journey because along the way there is no clear direction forward. In contrast with states in thermodynamic equilibrium, definable by a handful of state variables, a system out of equilibrium explores states definable only by more detailed

descriptions. As such, the dynamics can be complex and intricate, capable of retaining memory of the system’s past inside the extra degrees of freedom.

Memory effects in physical systems far from equilibrium showcase the additional intricacy and complexity of relaxation dynamics. A mark is imprinted on the system’s internal state which is internalized and carried into the future. Understanding how the memory is stored by the system and then eventually revealed leads to an understanding of the internal state and its evolution.

As such, memory effects in materials are a rich topic of study and come in many flavors [33]. In Chapter 6 we study a form of memory in granular systems that is able to recall the loading history and play it back, once, through nonmonotonic stress relaxation. A model for the internal state, consisting of a distribution of simple, single-timescale exponential relaxing elements, goes a long way to helping us understand this particular memory effect. We push the limits of the theory through various experimental protocols, and in the process learn more about the evolving inner state of the granular material.

1.3 Granular materials

By one count, granular materials are the second most manipulated material on Earth, after water [34]. As a field, granular physics benefits from the fact that there are so many industrial applications of granular materials. The practical relevance provides an extra source of justification for granular physics research, even if many works’ motivations are to study complexity from a fundamental standpoint [35].

Granular materials present unique opportunities as a model system, because their macroscopic scale allows for direct access to the constituent grains. This allows the tailoring of parameters that govern local stress transmission and redistribution though, e.g., particle shape [36] and surface friction [37]. It also enables tracking the state of individual particles as the system evolves, from their position [38, 39] to the forces they bear [40].

The idea that particle shape is an important driver in a granular material’s plastic

deformation has been pursued in simulations [41, 42, 43], but past experiments have either been mainly limited to two-dimensional systems [44, 27, 45] or confined to spheres [46, 47, 48, 25, 49] or naturally occurring soils and grains [50, 51, 52, 53]. This thesis work takes advantage of 3D-printing in order to create particles whose shape-dependent interactions with contacting neighbors can be tailored with precision, while parameters such as the particles' material stiffness and their surface friction can be kept unchanged.

Specifically, by varying particle shape we are able to change the manner with which particle surfaces meet to support stress, whether by edges, corners, or surfaces with different radii of curvature. Shape also drives where on each particle contacts are likely to occur, leading to different proportions of body forces versus torques. Finally, the resistance particles feel toward reconfiguring along rotational and translational degrees of freedom should change with particle shape (e.g., a flat disk would rather slide than rotate out of plane, and this preference intensifies as the aspect ratio of disks becomes larger). In these ways and others, how stresses are transmitted at the microscale is modifiable through shape to a degree impossible in other plasticity experiments. Granular materials composed of different particle shapes are therefore a powerful system with which to study plasticity.

EXPERIMENTAL METHODS

2.1 Instronduction

The workhorse of this thesis was an Instron 5800 Series universal materials tester (Figure 2.1a). The Instron is able to do only two things, but it does them with very high precision. First, it can control its displacement down to a scale of $2.5\mu\text{m}$, with a possible range of travel of more than a meter. Second, its load cell measures force with a noise level of 0.01N over a span of about 1kN , five orders of magnitude larger. With these capabilities a clever experimentalist can probe any mechanical property of a material, which is why it is called a universal materials tester.

Any force measurement device is ultimately a damped oscillator, and the Instron's load cell is no different. It is described well by the equation for a linear damped harmonic oscillator:

$$m\ddot{z} + b\dot{z} + kz = 0 \quad (2.1)$$

The mass parameter m is about 10^3 kg, incredibly high considering the “mass” at the end of the load cell is only about one kilogram. The other parameters b and k were measured in experiments to be 9×10^4 kg/s and $3 \text{ N}/\mu\text{m}$. The data from these experiments is shown in Figure 2.2, where a piece of acrylic was broken and the force supported by the Instron changed rapidly. The fitted dynamics for a damped harmonic oscillator lie right on top of the measured data. A second ring down occurs in addition to the primary component, and the overall response can be fit remarkably well with the form:

$$F(t) = Ae^{-t/\tau_1}\cos(2\pi f_1 t) + Be^{-t/\tau_2}\cos(2\pi f_2 t) \quad (2.2)$$

with $A \approx \Delta F$ from a sudden force change, $B \sim 1\text{N}$, $(\tau_1, f_1) = (20\text{ms}, 7\text{Hz})$, and $(\tau_2, f_2) =$

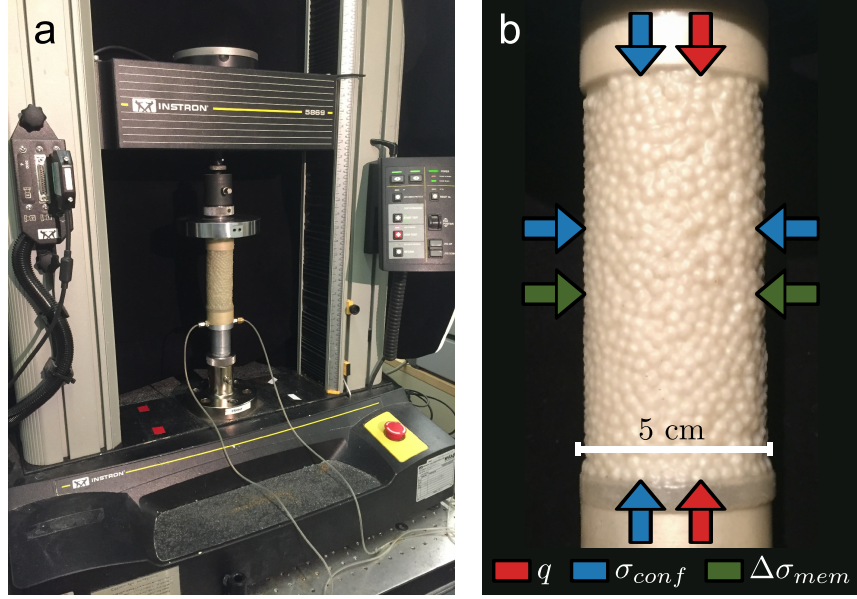


Figure 2.1: The Instron and the stress state of the packing. **(a)** The Instron frame, with a packing of glass spheres under compression. **(b)** The stress state for a packing of icosahedra, with q the deviatoric (axial) stress applied by the Instron, σ_{conf} the isotropic confining pressure applied via vacuum, and $\Delta\sigma_{\text{mem}}$ the additional inward radial stress applied as the membrane expands.

(500ms, 115Hz). The 115Hz tone can be heard during compression tests as a ringing that follows large and sudden changes in force.

Understanding the Instron in this way has been important for the fastest parts of the experiments in this work; namely, the 50ms response time of the load cell after a stress drop (Chapter 4) and for backing out the true force caused by an impacting object (experiments not described here).

The $2.5\mu\text{m}$ step size for vertical motion of the load cell factors into compression experiments. The oscillatory driving causes fluctuations in the data around 0.1N in magnitude, raising the noise floor of the data by an order of magnitude.

For the plasticity experiments of Chapters 3 and 4, data were collected at 40Hz, while for the relaxation experiments of Chapter 5, this rate was increased to 100Hz.

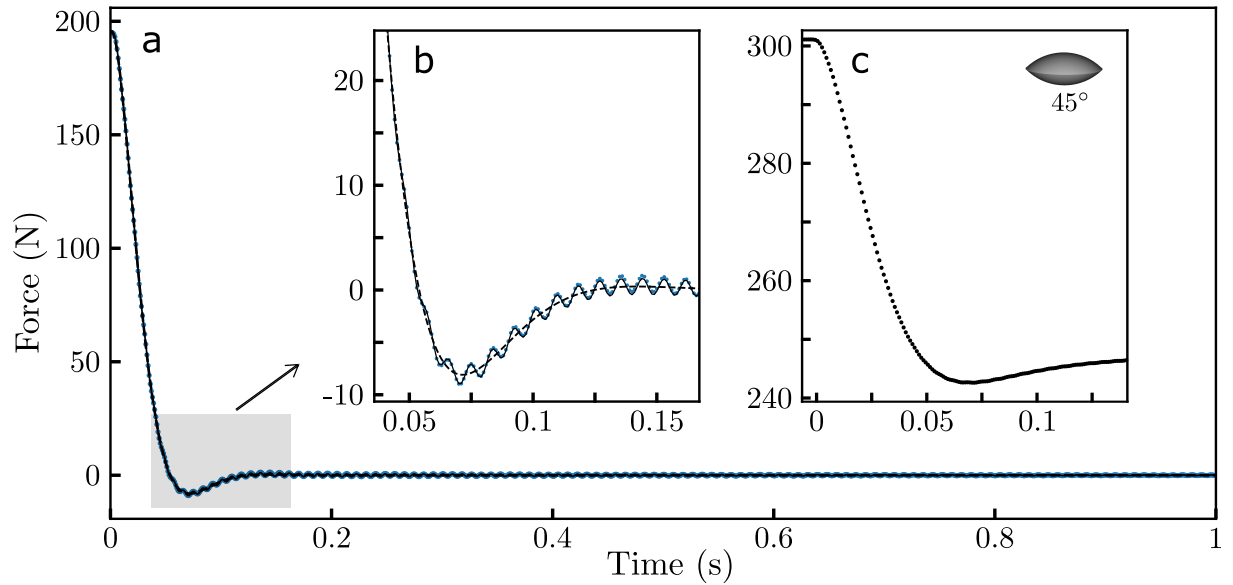


Figure 2.2: The Instron's finite response time. **(a)** The first second of force data after a piece of acrylic broke under the load cell. Note the 50-100ms time for the force to drop, the overshoot, and the rebound to zero. There are two fits to the data, with the first being a single decaying sinusoid (dashed) and the second being the two decaying sinusoids of Equation 2.2 (solid). **(b)** Magnified portion of the data from **(a)**; the single sinusoid fits the overshoot, but the addition of a second fits the 115Hz oscillations remarkably well. **(c)** A stress drop captured during uniaxial compression of a packing of lens particles. The 50-100ms drop time was the same, though the high frequency oscillations are nearly imperceptible (magnify the image to a few hundred percent to see faint oscillations).

2.2 Particles

Table 2.1 lists the particle shapes used in the experiments of this work. In Chapters 3, 4, and 6, all particles were 3D-printed from UV-cured hard plastic (Young’s modulus $E_{\text{mat}} \sim 1\text{GPa}$) using an Objet Connex350 printer with resolution $30\mu\text{m}$. Each particle’s volume was 22.5mm^3 , except for the corner particles, a shape comprised of three adjoined cubes and a total volume of 67.5mm^3 . For Chapters 3 and 4, we controlled for surface properties of the particles by only using particle sets fresh out of the 3D-printer. However, as will be discussed in Chapter 6, the particle surfaces are affected by details of the printing process. Surfaces perpendicular to the printing plane are rougher than those parallel to it as a result of the way the printer deposits the plastic in layers.

For the stress relaxation experiments of Chapter 5, MoSci soda lime glass spheres with diameters 4.7-5.3mm were used. The particles were cleaned with acetone prior to use.

2.3 Plasticity tests

For each experiment about 5,000 particles (a third of that for the corner particles) were poured randomly into a 0.6mm thick, 5.0cm diameter latex membrane to form a column with aspect ratio 2:1 (height to diameter).

Granular materials composed of non-interlocking particles can only support compression; under tension they simply dissociate. We did something akin to prestressing glass by putting the whole granular packing under isotropic compression before applying any shear. This was accomplished via the confining membrane and a vacuum pump which created a pressure differential between the atmosphere and the packing. We then applied a component of pure shear to the packing through uniaxial compression. The force measured by the Instron was the deviatoric stress q of this pure shear component, multiplied by the cross-sectional area of the column.

The membrane provided a small amount of additional confining stress in the lateral






















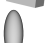
Shape	Symbol	Shape	Symbol
Sphere		Lens, $\gamma = 75^\circ$	
Tetrahedron		Lens, $\gamma = 60^\circ$	
Triangular bipyramid		Lens, $\gamma = 52.5^\circ$	
Cube		Lens, $\gamma = 45^\circ$	
Octahedron		Cone, $\beta = 45^\circ$	
Dodecahedron		Disk, $D/t = 1.5$	
Icosahedron		Disk, $D/t = 3$	
Supercube, $m = 3$		Disk, $D/t = 4$	
Supercube, $m = 4$		Disk, $D/t = 4.5$	
Divot cube		Corner	
Hemisphere		Prolate ellipsoid, 2.5:1:1	

Table 2.1: The particle shapes used in this thesis. Nonstandard shapes: A triangular bipyramid is two tetrahedra joined face to face. A supercube surface is defined by $|x|^m + |y|^m + |z|^m = R^m$, with $m=2$ corresponding to a sphere and $m=\infty$ a cube. A divot cube has a right pyramid removed from each face such that edges sharpen from 90° to 71° . A lens particle is the union of two spherical caps with polar angle γ ; a lens particle with $\gamma = 90^\circ$ is therefore a sphere. The angle β is the aperture of a cone particle; the tip of the cone is removed to prevent puncture of the latex membrane so the more precise shape is a frustum with $r_1/r_2 = 5$. A corner is three adjoining cube particles.

direction due to the elastic energy of its stretching as the packing bulged outward. The stress at which a granular material yields depends on the amount of confinement bolstering the packing, so accounting for the lateral stress correction was necessary for accurately measuring the shear strength. The correction relevant to these experiments is a hoop stress measured by Henkel and Gilbert in 1952 [54] which approximates the lateral expansion of a granular material as that of an incompressible cylinder. The lateral stress correction is given by

$$\Delta\sigma_{\text{mem}} = \frac{2M}{D_0}(1 - \sqrt{1 - \epsilon}) \approx \frac{M\epsilon}{D_0} \quad (2.3)$$

M is the “extension modulus” of the membrane, which scales linearly with its thickness and was measured to be 7 N/cm for the 0.6mm thick membrane. D_0 is the initial diameter of

the membrane, 5cm. At 20% compression, the correction is about 3kPa, a negligible value compared to the large confining pressures used in triaxial tests in geology, but significant for the $\sigma_{\text{conf}} = 20\text{kPa}$ of our experiments.

The stress state is sketched schematically in Figure 2.1b. The full stress tensor on the granular column was

$$\sigma_{ij} = \begin{pmatrix} \sigma_{\text{conf}} + q & 0 & 0 \\ 0 & \sigma_{\text{conf}} + \Delta\sigma_{\text{mem}} & 0 \\ 0 & 0 & \sigma_{\text{conf}} + \Delta\sigma_{\text{mem}} \end{pmatrix} \quad (2.4)$$

The Z component of stress is the first principal stress component, $\sigma_1 = \sigma_{\text{conf}} + q$. The second and third are the X and Y components, degenerate due to symmetry in the transverse plane.

2.4 Strain-controlled versus stress-controlled compression

The loading mode in compression experiments can be parametrized by the ratio of the stiffness of the measurement apparatus to that of the sample, which we call α following the work of Cui *et al.* [55]. This ratio dictates how the system responds to a relaxation event in the sample. Approximating the sample and measurement apparatus as linear springs in series, and therefore in force balance, α is the ratio of the displacement of the sample to that of the apparatus in any force change. Additionally, at any point in time during the compression, the two components have stored elastic energy $U_i = F^2/2k_i$. Thus $U_{\text{Instron}}/U_{\text{sample}} = k_{\text{sample}}/k_{\text{Instron}}$, and α is the ratio of elastic energy in the sample compared to that in the apparatus.

In order to fully resolve individual stress drops, large apparatus stiffness ($\alpha \gg 1$, strain control) and sufficiently slow compression ($\dot{\epsilon} \rightarrow 0$, quasistatic limit) are required [56]. In our experiments $\alpha \approx 3 \times 10^6 \text{Nm}^{-1} / 3 \times 10^5 \text{Nm}^{-1} = 10$. The large value of α manifested in raw plasticity data (Figure 3.1) as the nonlinear recharge of stress after events, which shows that the granular packing reloaded in stress rather than the Instron. This contrasts with

classic stick-slip experiments where $\alpha < 1$ and the reloading is linear [57, 58, 59], showing the force recharge of the measurement apparatus rather than of the sample. If we place a softer spring between the Instron and the sample, we recover linear recharge following slip events (not shown).

The distinction is important: in the low α stick-slip regime, the apparatus compresses during the stick phase and surges forward into the sample during a slip, using its own stored energy to fuel plastic events. In the other extreme, plastic events run on the energy stored in the sample only and the measured stress drops are better indicators of the internal rearrangement process.

Within the sample, elastic energy repeatedly builds enough to overcome local frictional thresholds and then sets off surges of motion, resembling stick-slip. See, e.g., the Burridge-Knopoff model for earthquake dynamics [60] or the elastic network studied by Nguyen and Coppersmith as a progenitor to the q -model [61]. The point is that dynamics *within* the granular system were our object of study, not the interaction of the measurement apparatus with the system. Due to a serendipitous combination of the Instron’s high stiffness and the relatively soft 3D-printed particles, we operated in a large α regime that facilitated studying the rearrangement dynamics.

2.5 Stress relaxation tests

The stress relaxation experiments of Chapter 5 began identically to the plasticity experiments, with compression continuing to 10% strain. At this point, compression halted and the Instron recorded the force over minutes or hours. Since the frame was static, the noise floor in the force data dropped by nearly an order of magnitude compared to the plasticity experiments, i.e., from 0.1N down to 0.01N. To program memories into the packing, the frame was driven forward and backward after various wait times $t_{w,i}$ before holding again. These extra compression steps were run as fast as possible, which we found to be 1s in duration at speeds up to around 2mm/min. If the (de)compression steps were any shorter than

1s in duration or if the speed was any higher, the frame movement was variable from test to test due to the time of the Instron to accelerate and decelerate.

In these experiments we doubled the confining pressure to 40kPa, which roughly doubled the measured forces along with our signal to noise ratio.

2.6 X-ray radiographs and computed tomography

X-ray radiographs were taken with an Orthoscan C-arm fluoroscope at fixed intervals (2s on, 1.5s off) during several compression tests. The beam energy was set at 70kVp for which the machine set the current to $91\mu\text{A}$. The C-arm emits a cone beam from a source to a detector 15 inches away.

The C-arm was also used for obtaining computed tomography (CT) scans of packings, in combination with a homemade rotation stage and reconstruction code described in [62]. 800 X-ray radiographs were taken as the sample was rotated through a full 360° , followed by a filtered backprojection. The process yielded a volume of 20 million voxels whose values were measurements of the x-ray opacity at each point in space. The reconstruction of the total $(6\text{cm})^3$ imaged volume was performed with voxels of $150\mu\text{m}$ side length; as such each 22.5mm^3 particle was represented by several hundred voxels.

2.7 Surface metrology

A LEXT OLS5000 laser confocal microscope was used to measure changes in the micro-scale surface features of the lens particles of Chapter 6. Immediately after 3D-printing and then after every nine compression tests, 50 particles were removed from the granular material to be scanned from the top and 200 were removed to be scanned from the side. The scanned surface area constituted about 0.5% of the total surface area of all particles in the packing, while the scanned edges were about 1% of the total edge length, making the scans a small but statistically representative sample of the granular material.

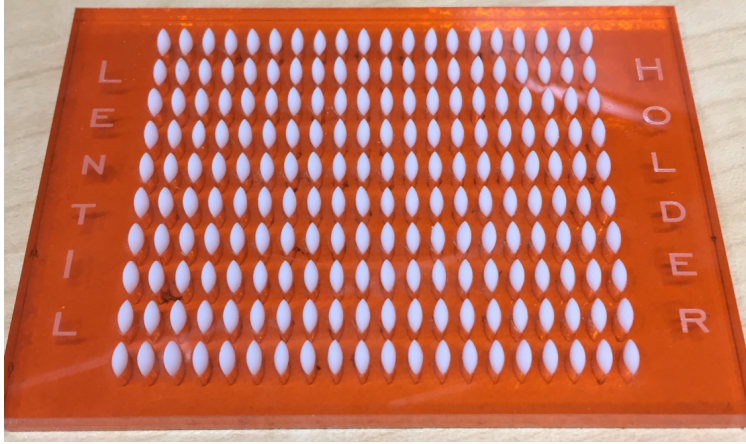


Figure 2.3: Scanning the surfaces of a statistically relevant sample of lens particles. Due to automation in the scanning process, large numbers of particles could be processed at once. 200 particles' edges and 50 particles' faces were scanned after every nine compression tests.

Scans were taken with the $5\times$ objective and $20\mu\text{m}$ vertical steps, yielding height map data with $2.5\mu\text{m}$ spacing between data points in the horizontal plane and a nominal resolution of 30nm in height.

PLASTIC DEFORMATION

A granular material may be thought of as an elastic continuum segmented into grains that are free to rearrange in a nonaffine manner, though to do so they may have to first overcome barriers such as friction and volumetric dilation. With enough applied shear stress the relative rearrangement of the grains becomes less costly in energy than the deformation of the grains individually (or the grains break, which will not be discussed in this thesis). The rearrangement of grains is rich in physics: the system finds new configurations out of countless possibilities in a process that is a compromise between cooperation and individuality. By changing the shape of the grains, we found a remarkable range of behavior during plastic deformation and were granted a foothold into understanding how grain-level details influence the rearrangement process. The deformation was found to be punctuated, in all cases, by sudden drops in stress, which will be the focus of Chapter 4. As a prelude to an in-depth analysis of the drop events, this chapter characterizes the plasticity data in a ‘first-pass’ manner. By first-pass, we mean methods which are ignorant to the existence of drop events.

A sampling of the raw stress-strain data obtained in these experiments is shown in Figure 3.1. Immediately apparent is the large influence particle shape has on the aggregate’s plastic deformation. Some shapes led to deformation with large and erratic fluctuations, like the 3:1 disks and the 45° lenses, while others deformed more quiescently, like the spheres and tetrahedral bipyramids. Additionally, the stress required for plastic deformation changed by nearly a factor of three from the weakest (spheres) to the strongest (bipyramids). By quantitatively measuring these differences, we arrived at shape-dependent trends that shed light on the role particle shape plays in an aggregate’s plastic deformation.

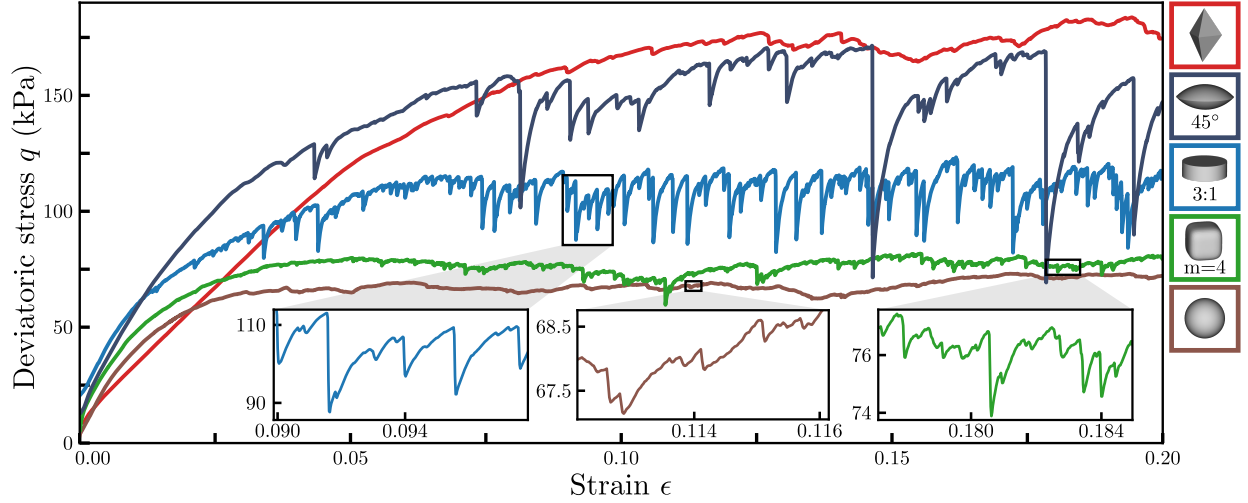


Figure 3.1: Raw plasticity data for five shapes. Raw stress-strain data for a single run each for the five particle shapes shown on the right (see Table 2.1 for shape definitions). The critical state was nominally associated with the range $\epsilon = [0.1, 0.2]$. *Insets:* All data exhibit sudden drops in stress, though at scales that differ by orders of magnitude.

3.1 Critical state

After the first few percent strain, a granular packing enters a steady state for continued plastic deformation. Parameters such as the stress and packing density become roughly constant in a state known to soil mechanicians as the critical¹ state [63].

For the same granular material, a very different yield strength will result if, initially, it is packed densely or loosely. The mechanical properties depend on more than the properties of the grains. However, after a granular system has deformed enough and transitioned into the critical state, it maintains a constant average volume and a constant average stress which depend only on grain properties.

By studying the system while in this well-defined state, we were able to ignore, to a large degree, the initial packing state set when pouring. We used only data from strains $0.1 < \epsilon < 0.2$ for studying plastic deformation.

1. Unrelated to the usage in physics to refer to second order phase transitions

3.1.1 Diffuse failure, not shear banding, due to size

A common failure mode of amorphous systems is shear localization, also known as shear banding [64]. Plastic deformation costs energy to overcome barriers to reconfiguration; in granular materials dilatancy must occur [65] raising the energy density locally. However, if a system can localize all the plastic deformation into a narrow band, it saves energy relative to homogeneous deformation. The thickness of such shear bands is about 10 particles [35, 64], and the relative energy savings is tied to the size of the system relative to a shear band. When the system is small enough, the savings to localize shear are no longer worthwhile, and diffuse failure results. We focus on mesoscale dynamics, where the length scale of the system is an order of magnitude larger than a characteristic rearrangement event but small enough to inhibit shear banding.

Experiments on nanopillars of amorphous metal have demonstrated that when the size of the system is made small enough to inhibit shear banding, the system becomes stronger and fails with more ductility [66, 67]. By inhibiting the least energy expensive failure mode, higher stresses are needed to deform the material, and when the deformation cannot localize, there is no runaway process so ductility results.

We were able to study the same physics at the centimeter scale, keeping the length scales of the granular column in the range of 10-100 particle diameters, but with the added capabilities of modifying the constituent particles of the system and directly imaging the individual rearrangement events with X-rays.

Additionally, the mesoscale is precisely where amorphous plasticity theory struggles. Whether or not shear banding arises is thought to dictate whether mean field approaches to amorphous plasticity are valid. By exploring the territory in the middle, with a system just small enough to prevent shear banding, we probe the crossover between the two regimes. As a result, mesoscale granular systems without shear bands offer an excellent testbed for studying the intermittent dynamics of plastic deformation in amorphous materials.

3.2 Angle of internal friction

Another way to think about granular materials is as volumetric friction, due to the multitude of interfaces in contact within a volume of grains. Indeed, in this light the highly dissipative properties of granular media are completely sensible.

When isotropic pressure is applied to a granular packing, the effect is similar to increasing the normal force on a block on an inclined plane: negligible motion results (as long as the pressure is not enough to cause grain breakage), and the forces needed to overcome friction increase. Under applied shear stress, friction between the grains will resist motion to a point, beyond which Coulombic frictional thresholds are overcome and rearrangement occurs.

One of the early constitutive relations for granular materials is based on this picture, and is called the Mohr-Coulomb failure criterion [68]. It describes how a granular material will flow when

$$\tau = \sigma \tan(\psi) + c, \tag{3.1}$$

where τ and σ are the shear strength and normal stress, respectively, c is the cohesion between grains, and ψ is the angle of internal friction. The compressive stress σ and cohesion c force the grain surfaces together, the analogous coefficient of friction is $\tan(\psi)$, and τ then describes the minimum shear stress required to cause motion.

Cohesion can arise due to, for example, water bridges or van der Waals forces between particles, but we can safely take it to be zero for our dry packings with millimeter-scale particles. Extreme non-convexity of particles, to the point of geometric entanglement, can induce an effective cohesion term [69], but the particle shapes discussed in this work are far from this regime.

The angle of internal friction ψ is a bulk material parameter, dependent on packing density and independent of the stress scale of the experiment. It evolves in a uniaxial compression experiment with packing density until the critical state is entered, beyond which it acquires a well-defined value.

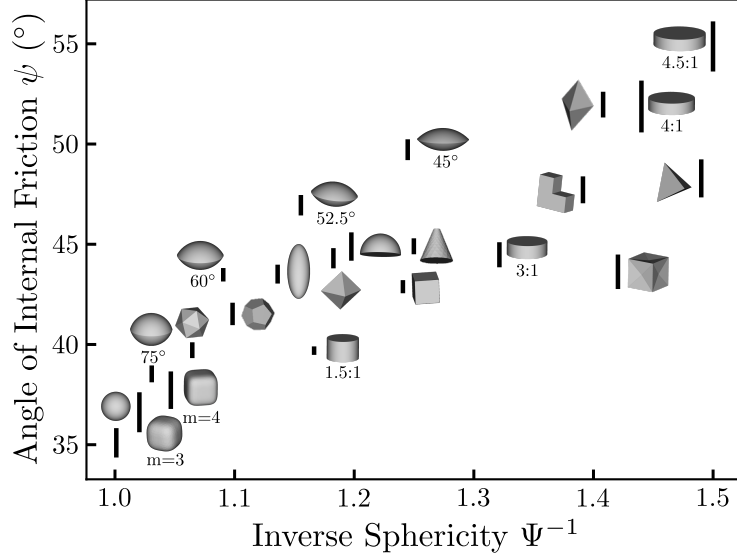


Figure 3.2: Shear strength versus deviations from a sphere. The angle of internal friction plotted against the inverse of sphericity. As a measure of the compactness of a shape, sphericity is the ratio of the surface area of a sphere with the same volume to the surface area of the shape.

The angle is defined for a packing of cohesionless grains as [63]

$$\sin\psi = \frac{\sigma_3 - \sigma_1}{\sigma_3 + \sigma_1} = \frac{\bar{q} - \Delta\sigma_{\text{mem}}}{\bar{q} + 2\sigma_{\text{conf}} + \Delta\sigma_{\text{mem}}}. \quad (3.2)$$

Here σ_3 and σ_1 are the largest and smallest principal stress components, equal to $\bar{q} + \sigma_{\text{conf}}$ and $\sigma_{\text{conf}} + \Delta\sigma_{\text{mem}}$, respectively, and \bar{q} is the average deviatoric stress in the critical state (see Section 2.3 for a discussion of the stresses on the packing). Extracting \bar{q} from measurements of $q(\epsilon)$ involves subtlety if stress drops are ever so large that they correspond to a sizable fraction of the stress plateau reached in the critical state: in these cases a simple average of $q(\epsilon)$ tends to lie below this plateau stress and underestimates the stress needed for plastic deformation. To account for this, we instead defined the plastic stress as the average of all of the stress values immediately preceding a stress drop, similar to what was done in [47].

We can relate ψ directly to the geometry of individual particles in Figure 3.2, where we plot ψ as a function of the inverse particle sphericity Ψ^{-1} . Sphericity is a measure of the compactness of a shape relative to a sphere (whose $\Psi=1$). It is defined as the ratio of

the surface area of a sphere with the same volume to the surface area of the shape [36]. As the data indicate, Ψ^{-1} correlates well with ψ , as would other single-parameter shape descriptors of closeness to a sphere, such as particle roundness [50] or isoperimetric quotient [70], highlighting the particles’ resistance to rotation as the dominant driver in a granular packing’s shear strength [52].

3.3 Volatility

In order to quantify the magnitude of stress fluctuations in the plastic regime, we introduce a measure borrowed from financial mathematics which quantifies the spread of fractional changes that occur in a time series. This measure, volatility, is model-independent and is particularly useful in comparing broadly distributed fluctuations in data with different or changing baselines [71]. Importantly, it sidesteps the issue faced in many plasticity experiments of accurately locating and measuring the magnitude of stress drops [72], giving a robust method to quantify the magnitude of fluctuations in plasticity data across experiments and even across systems.

In the Black-Scholes model for price evolution in financial mathematics, the volatility is defined as the standard deviation of the natural logarithm of fractional changes [71]. It provides a dimensionless, baseline-independent measure of the rate of change of a discretized time series. Treating the stress data as a series q_i , where each “time” step i corresponds to an applied strain value, we define the instantaneous “return”

$$R_i = \ln(q_i/q_{i-1}). \quad (3.3)$$

The volatility V of the series is the sample standard deviation of R , given by

$$V = A \sqrt{\frac{1}{N-1} \sum_i^N (R_i - \bar{R})^2}, \quad (3.4)$$

with N the number of strain intervals in the dataset and A a constant to correct for the effect of data resolution. Finer resolution yields smaller fractional changes in q and therefore a smaller value for V if not corrected. We “annualize”, as done in finance, with the constant A equal to $\sqrt{\epsilon_{\text{win}}/\Delta\epsilon}$, the square root of the number of data points taken in a strain window ϵ_{win} relevant for analysis. For our data, the strain interval $\Delta\epsilon$ between measurements q_{i-1} and q_i is set by the sampling rate of the Instron and strain rate of compression. The strain window ϵ_{win} is simply an overall scaling, its value being more symbolic than anything else as long as it remains constant across the board. We chose $\epsilon_{\text{win}}=0.01$.

Whereas poor time resolution can interfere with the measurement of drop event magnitude distributions [72], the assumptions underlying the calculation of volatility are actually more applicable when the data acquisition timescale is longer than the timescale of a stress drop (but shorter than the timescale separating drops). Specifically, the annualization scaling is based on the assumption that changes in the time series are uncorrelated, which is invalid if the sampling rate is high enough to capture several data points during a stress drop. To this end, we downsampled the data by a factor of two when calculating the volatility. There is correlation in the stress values during the recharge between drops, but their contribution to the volatility is insignificant compared to the large changes in stress due to a drop. The model-independent nature of volatility and its applicability to low resolution data make it a broadly applicable tool for quantifying the magnitude of fluctuations in time series data.

To a first order approximation, the return $R_i \sim \Delta q_i/q_i$ is the fractional change between timesteps, so the volatility can be interpreted as a spread in the fractional changes over the annualization time interval. In other words, a volatility value of 10^{-2} means stress fluctuations are around 1% of the average stress in each $\epsilon_{\text{win}} = 0.01$ strain interval. Parenthetically, the volatility values for many common stock market indices fall in the neighborhood of 10%, a comparable value to that of many of the more volatile shapes shown in Figure 3.3.

3.4 A new phase space for plasticity data

In Figure 3.3 we combine the measurements of the angle of internal friction and volatility for all of the 22 particle shapes tested. Packings composed of spheres offer the least resistance to shear out of all particle shapes, exhibiting the lowest ψ and also low volatility. As ψ increases, the region in phase space of nearly spherical particles splits into two branches, one containing more platy and oblate shapes and one more angular and prolate ones. These branches differ by more than an order of magnitude in V at large ψ .

Packings of platonic solids (in red) grow stronger as the number of particle faces decreases from icosahedra to tetrahedra and as particle corners and edges become more pronounced. However, they generally exhibit quiescent deformation with V values not much larger than those of spheres. Cubes are the exception with a roughly fourfold enhancement in V , providing a first hint about the relative importance of edges and faces. We explored this via several variations on the shape of a cube. Rounding the edges into supercubes drops ψ and fusing three cubes into a corner-shaped particle with larger faces and longer edges increases ψ , but neither has a significant effect on V . On the other hand, while indenting flat cube faces by creating divots enhances ψ slightly (presumably due to the sharper edges), it significantly reduces V , all the way down to the level of spheres.

As the family of lens-shaped particles (in blue) demonstrates, both volatility and angle of internal friction are increased quickly when the particles become more oblate. Disks (in orange) follow a similar trend with increasing ψ as their aspect ratio (diameter to height) increases, though the fractional fluctuations saturate at aspect ratios of 3 and higher.

These findings suggest the presence of competing factors: while shape anisotropy (oblate or prolate) as well as sharp edges or corners were all found to enhance ψ , they do not predict V . Instead, V appears to be more dependent on the degree to which particle contacts involve surfaces with large radius of curvature, such as faces rather than edges. This depends not only the existence of faces but also on the frequency of face-face contacts. With the divot cubes we eliminated flat cube faces, while tetrahedra, a shape with large flat faces, is an example

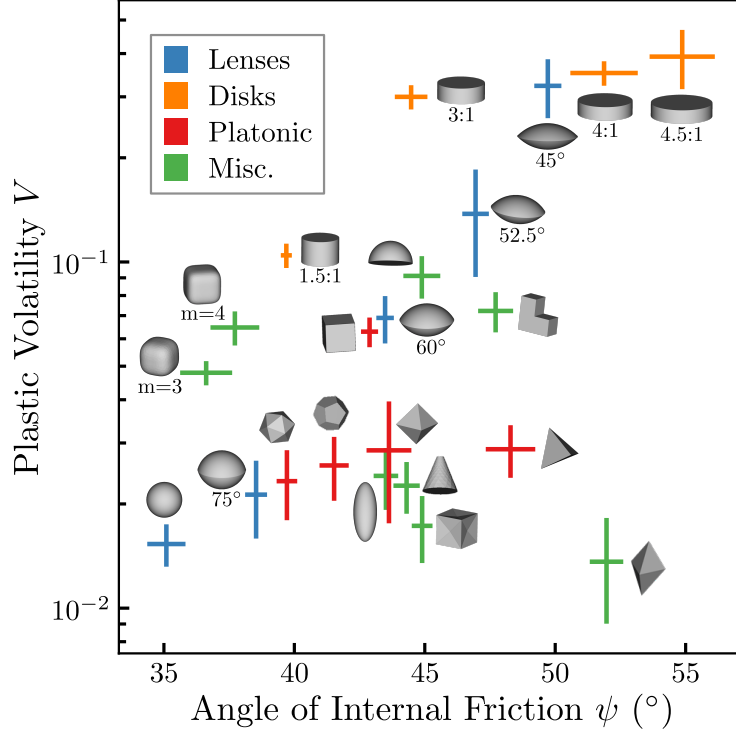


Figure 3.3: Plastic deformation phase space. The magnitude of fluctuations in the plastic regime and the shear strength of the granular packing are mapped out by plotting the annualized plastic volatility V versus the angle of internal friction ψ . Cross bars for each particle shape are the standard deviation of V and ψ across runs.

of a case where the packing structure does not favor face-face contacts [38]. Interestingly, the contacting areas do not have to be flat: increasing the radius of curvature at the contact enhances V . This is demonstrated by the family of lenses and also by comparing hemispheres with cones. The latter two shapes have nearly equal ψ yet differ in the radius of curvature at contact, which mimics the trend in V .

3.5 Summary

In this chapter we have shown that particle shape greatly changes the nature of plastic deformation for a mesoscale granular column. By modifying the manner in which grains meet and resist relative motion, we created materials with a broad range of plastic strength and violence of plastic deformation. We measured the first with a non-dimensionalized

quantity from soil mechanics, the angle of internal friction, and the second with a calculation borrowed from financial mathematics, volatility. The two metrics complement each other well: the magnitude of the baseline is represented in the angle of internal friction, while the volatility ignores everything about the baseline and measures the fractional kicks in the data. Thus the two quantities encapsulate the nature of how a granular packing adapts from one strain state to another.

The broad coverage of the phase space in Figure 3.3 shows that particle shape is a powerful control knob on the nature of granular plasticity. Shape families trace out various arcs in this space, serving as guides to help navigate the infinite space of possible shapes. Plastic strength and volatility can be varied separately, indicating a rich and sensitive dependence of bulk properties on particle geometry.

STRESS DROPS AND REARRANGEMENT CASCADES

4.1 Introduction

Amid the variety of plasticity behavior across shapes seen in the last chapter, every compression test featured sudden drops in stress. The magnified portions of Figure 3.1 show that plastic deformation was punctuated, in all cases, by stress drop events which dissipated energy and prevented the stress from climbing ever higher as compression continued.

The events are natural objects on which to focus, as the markers of rearrangement. The granular system had to restructure under the applied shear, and instead of doing so gradually, it opted for surges of action. These surges had a wide range of magnitudes and were dramatically influenced by particle shape.

In general amorphous plasticity, similarly intermittent deformation is found in a wide range of systems. The individual rearrangement events are unpredictable from one to the next, though the statistics of event magnitudes exhibit similar relations across a wide range of systems and size scales, including metallic glasses [21, 22, 23], foams [19, 20], granular materials [25, 26, 27], ice [17, 18], metals and alloys [59], and even knitted fabric [28]. A common reference scenario for this dynamic response is proximity to a non-equilibrium critical point[15], leading to intermittent dynamics and power law distributions for the event sizes. The physics are of a dynamical phase transition between a static and a flowing phase as a function of applied stress [29]. In this picture, rearrangement events are short-lived instances of the flowing phase that quickly snap back to the static phase following the rapid decrease in stress. When measuring a stress drop, we obtain a proxy for the size of one of these occurrences, and the result is predicted to be power law distributed up to a cutoff set by the proximity to the actual critical point[73]. Above the cutoff the chance of seeing a bigger drop is exponentially suppressed[74, 26].

The specific distribution is termed a truncated power law, and has the form

$$P(s) = Cs^{-\tau}e^{-s/s^*} \quad (4.1)$$

The exponent τ of the power law tells about the mechanism leading to scale-free dynamics in amorphous plasticity. Mean field models predict $\tau = 3/2$, independent of many details of the interactions among the systems components [73]. Simulations and experiments on plastic deformation of amorphous systems currently lack consensus or precision as to whether the exponent is indeed consistent with $3/2$ [25, 26, 75, 23, 59] or not [76, 77, 44, 27, 30, 78].

On the modeling side, much of this can be traced to differences in which the stress redistribution after a local plastic event is treated, i.e., to different treatments of the mesoscale dynamics [29]. When a portion of the system relaxes locally and then offloads stress to all neighboring elements isotropically, as happens in the presence of shear bands in, e.g., bulk metallic glasses and granular materials on larger scales, a mean field exponent $\tau = 3/2$ is found [74, 79]. When, instead, stress redistribution following an event takes the form of a quadrupolar kernel derived from Eshelby's work for localized plastic zones in elastic media [80], the exponent reduces to approximately $\tau = 1.3$ in three-dimensional systems [14, 30].

On the experimental side, discerning between $\tau = 1.3$ and 1.5 requires several decades of event magnitudes to be observed and careful analysis [23, 81, 72]. Most importantly, experimental investigation into the plasticity mechanisms operative at the mesoscale has been scarce due to the difficulty of varying local interactions in systems whose constituent units are atoms, molecules, or even bubbles. The experiments in this thesis offered unique opportunities on this front, as they allowed us to vary the nature of particle-particle interactions via particle shape in order to study the effect on the system's plastic deformation.

4.2 X-rays link stress drops to restructuring of the packing

Something fundamental to the way amorphous systems restructure clearly happens in the packing during stress drop events. Is the process a homogeneous shifting of the packing into a new sheared configuration, or are there localized bursts of activity? If the latter, are the rearrangements of a consistent size scale? We visualized the drop events by repeatedly imaging the sample with X-rays during compression and then subtracting successive images. In this way, fast processes were revealed and anything slower than the timescale between frames was removed in the subtraction.

Drop events last 50-100ms, set by the time the Instron needs to reequilibrate when the force it supports changes (Section 2.1). We X-rayed the sample every 3.5 seconds, meaning any rearrangement events that occur intra- or inter-frame were easily caught by the frame differencing. The entire compression was slowed down to last 40 minutes ($\dot{\epsilon} = 0.5\%/min$) so that any motion slower than drop events was more easily subtracted.

The X-ray difference images were aligned with the stress-strain data for each compression to illuminate the physical process marked by stress drops. Portions of the system, from a single particle to large swaths of the packing, shifted position suddenly — shown as dark blue in Figure 4.1 — whenever the stress dropped in $q(\epsilon)$.

For reference, the double-lobed blobs in many of the X-ray images highlight both where the particles were and to where they moved: the absolute value of the difference image was used in order to make the events show up more clearly. Events 5, 7, 9, and 10 in Figure 4.1 show clear examples of this.

Stress drops with smaller magnitudes, for instance events numbered 5 and 10 in Figure 4.1, showed up as the motion of a single particle in the packing. Intermediate scale drops included the motion of a handful of particles; see events 1 and 3. The largest stress drops were more lively: particles across the packing from one another moved together, as in events 2 and 7.

Originating with Argon’s experiments on sheared bubble rafts to better understand bulk

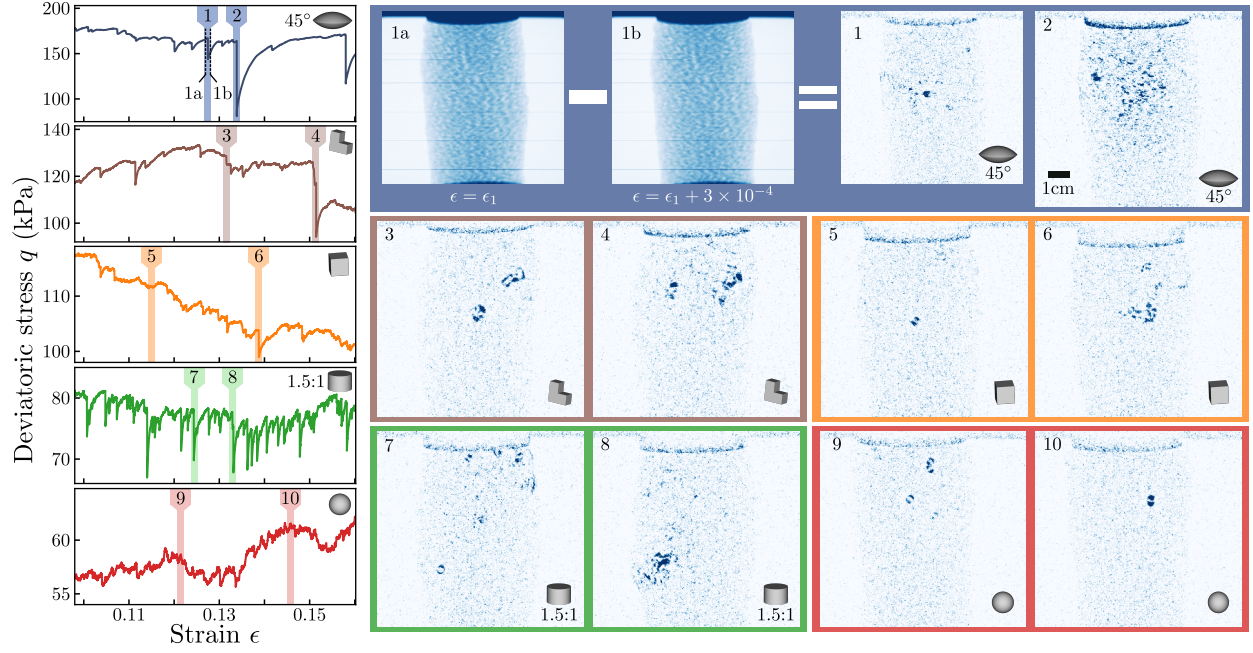


Figure 4.1: Localized slips revealed by X-ray videography. By subtracting successive X-ray radiographs of a packing during compression, we reveal the rapid particle rearrangements that caused the sudden stress drops in $q(\epsilon)$. The shaded vertical bars in the stress-strain data indicate the strains at which the X-ray data on the right were taken. The bars are roughly five times as wide as the strain window captured by an X-ray frame. At the top of each X-ray image the top plate's motion is visible. For a sense of scale, the dark blue regions in 5 and 10 are each a single particle. The difference images highlight both where particles moved to and moved from, causing the double-lobed particle outlines in many of the images.

metallic glass (BMG) plasticity [82, 83], the idea of shear transformation zones (STZs) (or simply shear transformations) has been a fruitful way to understand plastic deformation through units of rearrangement [84]. STZs are regions of $O(10)$ particles that reconfigure when stress becomes large enough to overcome barriers to motion (e.g., friction and dilation in granular systems). They are loci of rearrangement where the shear strain jumps backward to relieve shear stress. STZs can destabilize other regions of the system and set in motion cascades of events; thus in this picture large restructuring events are composed of many STZs instead of a single giant one [85]. The X-ray images show that plastic deformation of the sample was indeed undergone through localized events, in line with the STZ framework. We also found events where the rearrangement in the packing was spread out with far less motion in between, such as Events 3, 4, 7, and 9. These distributed cascades are evidence of the long range connection all particles have with each other in the form of the elastic force network [74, 40].

With these X-rays, we obtain a qualitative picture of the rearrangement process and confirmation that stress drops in $q(\epsilon)$ are signs of the clunky rearrangement of the packing. The magnitude of the stress drop serves as a measure of the scale of the rearrangement that occurred. The packings restructured through bursts of action with quiescence in between, as a large portion of the X-rays taken showed no noticeable movement.

To make quantitative claims with the X-rays, however, is problematic. The X-rays are a two-dimensional projection through the column, meaning we could not detect any component of motion that occurred in the beam direction.

Additionally, and more importantly, the difference imaging revealed the swap of material to void space and vice versa. This means if a particle moved into a previously unoccupied cavity, it would show up in its entirety on the X-ray difference images. If, on the other hand, several particles moved from one relatively dense configuration into another, very little would show in the difference images beyond slices and chops of the particles. This is because the method is agnostic about the positions of the particles. We obtained a qualitative picture

of rearrangement but missed out on anything more quantitative than a rough correlation between length scale of motion and drop event size.

4.3 Drop-finding with a high pass filter on the first difference

Once our focus shifts to the stress drops in $q(\epsilon)$, the natural questions to ask are about their distribution. How often does a huge stress drop, involving large portions of the packing rearranging together, occur relative to a small one? To answer this requires accurately measuring the magnitude of each drop in stress, a subtler task than it first seems. The conventional approach taken in experimental papers is to use the first difference of the data to approximate the first derivative, and then find large negative values indicative of a drop [72, 86]. A threshold is used to ignore noise, and everything exceeding the threshold is a drop. An extra step often applied is Wiener filtering to remove high frequency noise [23, 87, 44], but the rest of the method is the same. Thanks to the low noise floor of the Instron, the high frequencies of $q(\epsilon)$ can be left alone. This is great because the drops, our objects of interest, are among the highest frequencies in the signal and we would prefer leave them as intact as possible. Instead, we find that the low frequency trends in $q(\epsilon)$ are what cause problems.

The bare first difference method — by which we mean any method that finds drop events by scanning for negative values in the first difference, with or without the filtering of high frequency noise — is successful for the largest events, which is why it often escapes scrutiny. Problems arise for the smaller events when there is any trend in stress occurring in the background of a drop. In the drop magnitude distributions common in amorphous plasticity [26], the majority of the events are small, so inaccurately measuring their magnitude greatly skews the found distribution.

As seen in the raw $q(\epsilon)$ data in the top half of Figure 4.2, the stress recharged after a drop with a non-negligible positive slope and continued to do so for a long time relative to the timescale of a drop. In the first difference of the data in the bottom half of the figure, there are clearly events where the stress dropped suddenly from its trend, though not enough to be

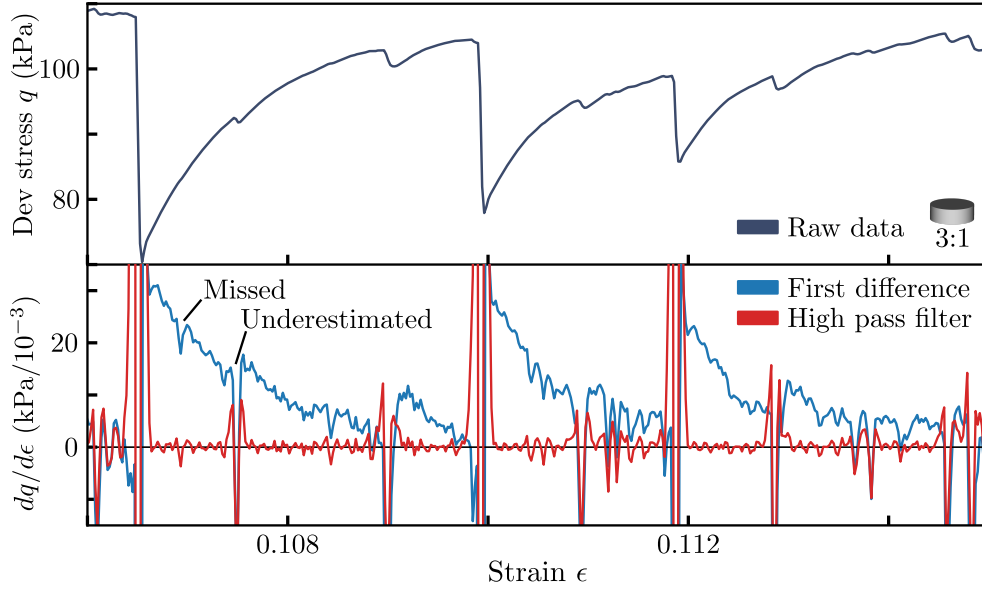


Figure 4.2: The need for a better drop-finding method, seen in disk data. *Top*: Part of a raw data trace for a packing of 3:1 disk particles. *Bottom*: The bare first difference method (blue), with a threshold at some small negative value of $dq/d\epsilon$, would miss and underestimate events that occur during the long stress recharge following a big drop. Passing the first difference data through a high pass filter (HPF) gets around this issue. It introduces a slight artifact, seen as the ‘wings’ bracketing each event in the HPF trace, which needs to be corrected (Figure 4.3).

caught by the threshold. The removal of the background baseline is necessary. This issue was addressed in [72], though for strictly linear recharge, as could arise in a stick-slip experiment. In this case the background trend can be removed with a simple offset to the first derivative. Due to the constantly changing slope of the recharge in our experimental data, we required a more adaptive solution. We found that passing the first difference through a high pass filter (HPF) before proceeding with the normal thresholding worked well. It removed background trends in the first difference of the stress that were of a longer timescale than an event. The timescale of an event is the 50-100ms of the Instron to respond to a change in force (see Figure 2.2) so we used this for the timescale of the filter.

To filter the first difference data, we Gaussian blurred — that is, we convolved with a Gaussian kernel — with $\sigma = 50\text{ms}$, and then subtracted the result off from the unblurred first difference data. The result is the red trace in the bottom of Figure 4.2. All features in the first difference data persisting longer than $\sim 100\text{ms}$ have been removed and the fast timescale drops can now be clearly seen as negative values. Noise in the force data remains as well, but this is ignored with an appropriate threshold.

Each stress drop is slightly reduced in magnitude as an artifact from the Gaussian blurring, though this can be corrected. The result of applying the method to a spike is shown in Figure 4.3. The characteristic ‘wings’ are a result of spreading the drop out with the Gaussian before subtracting it and can be seen in Figure 4.2. These wings allowed us to correct the drop magnitude by averaging their values to infer the size of the blurred drop, which we then add to the found magnitude.

4.3.1 Checking the method on synthetic data

In order to measure the accuracy of various drop-finding methods, we generated synthetic time series that mimic the experimental data. Since the drop magnitudes were known for the synthetic data, we were able to check the ability of different methods to faithfully measure drop magnitudes under different conditions, something which is impossible when working

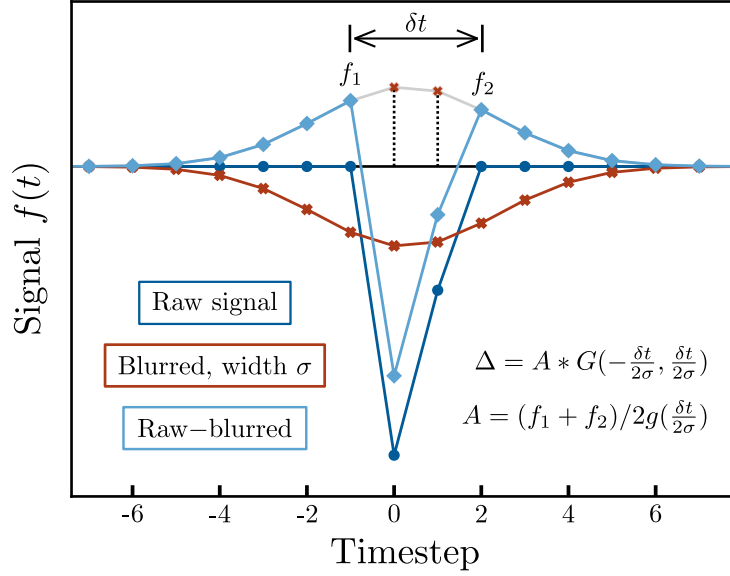


Figure 4.3: Drop magnitude correction for the effect of the high pass filter. The convolution of the raw signal with a Gaussian of width σ removes the baseline but also reduces the magnitude of the original spike. The missing magnitude is the rest of the blurred signal, shown as the two lightly shaded points. These can be estimated using the wing magnitudes f_1 and f_2 and the time interval δt between them. Here $G(a, b)$ is the integral of a standard normal distribution, $g(x)$, between a and b . To interpolate, we average the wing magnitudes and use their time away from the approximate center of the drop $\frac{\delta t}{2}$ to estimate the scaling of the Gaussian, A . We then integrate across δt to approximate the missing magnitude. A simple rescaling would have sufficed to correct the magnitude of drop events which only last a single timestep, but would have erred if events straddle multiple timesteps, as is the case here.

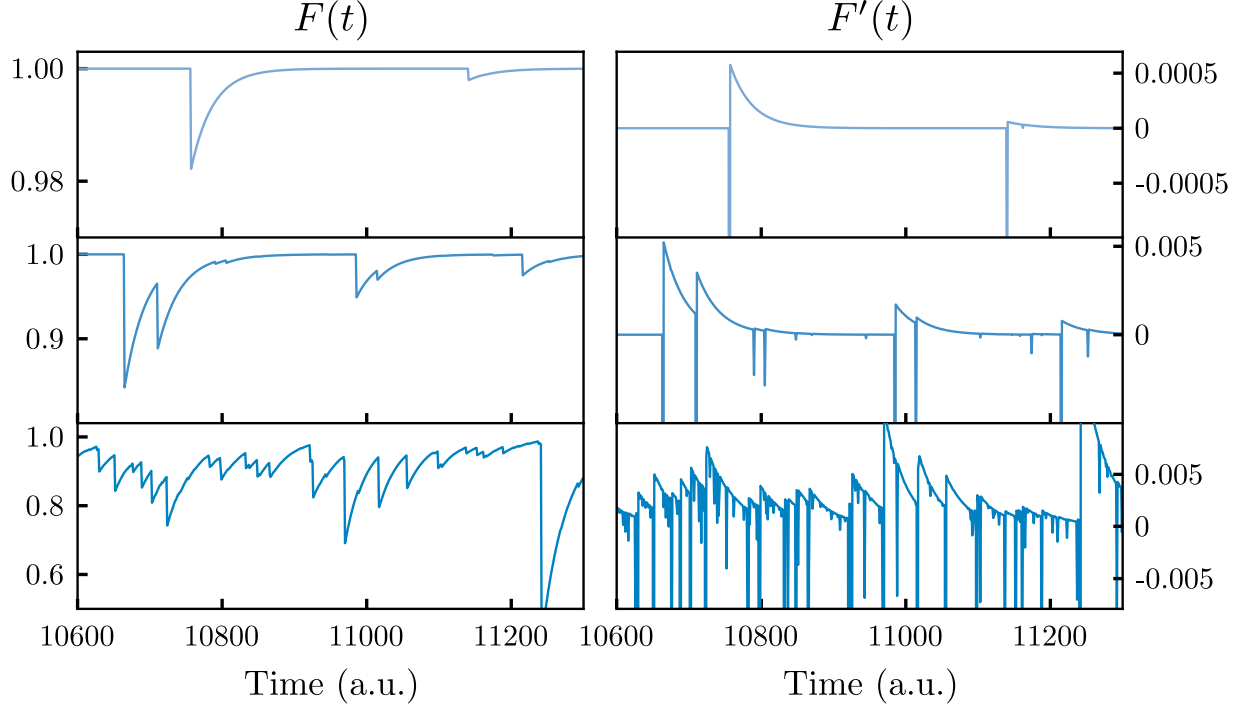


Figure 4.4: Synthetic time series used to test the drop-finding methods. Shown are portions of the synthetic time series on the left and their time derivatives on the right, for event rates of 0.05, 0.8, and 12.8 events per recharge timescale, from top to bottom.

only with experimental data.

Synthetic time series were generated with drop magnitudes sampled from a truncated power law distribution (Equation 4.1) with $\tau=1.2$, $s^*=0.2$. The events occurred randomly in time with a prescribed average rate (i.e., they were a Poisson process), and were followed immediately by an exponential recovery approximating the real stress-strain data (see Section 5.2.1 for the functional form of stress recharges in experimental data). To be specific, when a drop occurred, an exponential was subtracted from the time series:

$$\sigma(t) \rightarrow \sigma(t) - \Delta\sigma e^{-(t-t')/T} \text{ for } t > t'. \quad (4.2)$$

$\Delta\sigma$ is the value that was drawn from the truncated power law distribution. T , the recharge timescale, was set to be 30 times the duration of a drop, a physically realistic value.

Figure 4.4 displays portions of the synthetic data from different event rates, along with

the time derivatives which were processed for drop events. Event rate was the primary focus for the synthetic data, though we also explored the effect of adding various levels of Gaussian noise to the data. Higher levels of noise simply required higher thresholds for event recognition, but did not appear to change the performance of either method.

First we checked the performance of the drop-finding methods on measuring each individual event’s magnitude (Figure 4.5). In line with the intuition of Figure 4.2, the bare first difference missed and underestimated smaller events even when they were relatively infrequent. The application of a high pass filter did much better in measuring the small events. In the most treacherous scenario for drop-finding in the synthetic data, the event rate was more than 10 events per recharge timescale T , and the first difference data became a busy mess (Figure 4.4, bottom right). Both methods missed many events, but the ones that were captured by the method with the HPF are much more accurate, as seen by their proximity to the black dashed line representing perfectly accurate drop-finding (Figure 4.5).

The differences in the compiled distributions are more striking (Figure 4.6). Only for the lowest event rates, where the chance that two occurred within a recharge was negligible, did the bare first difference method yield a drop distribution with the correct power law exponent τ . Quickly τ fell due to small events being disproportionately missed or underestimated. The high pass filter method remained faithful to the actual distribution for event rates more than an order of magnitude larger, showing that the removal of the recharge trends was necessary and the Gaussian blur artifacts were successfully corrected.

At the highest event rates, even the actual drop distribution, though sampled from a truncated power law with $\tau=1.2$, skewed to smaller values of τ due to event overlap. The chance of multiple events occurring simultaneously became non-negligible, and yet again the smallest events were the ones which were missed. This time it was because in an overlap where the magnitudes sum, the larger event masks the smaller.

It should be emphasized that we encountered two forms of event overlap. In general terms, an event is a short timescale drop and then a long timescale perturbation to the time series as

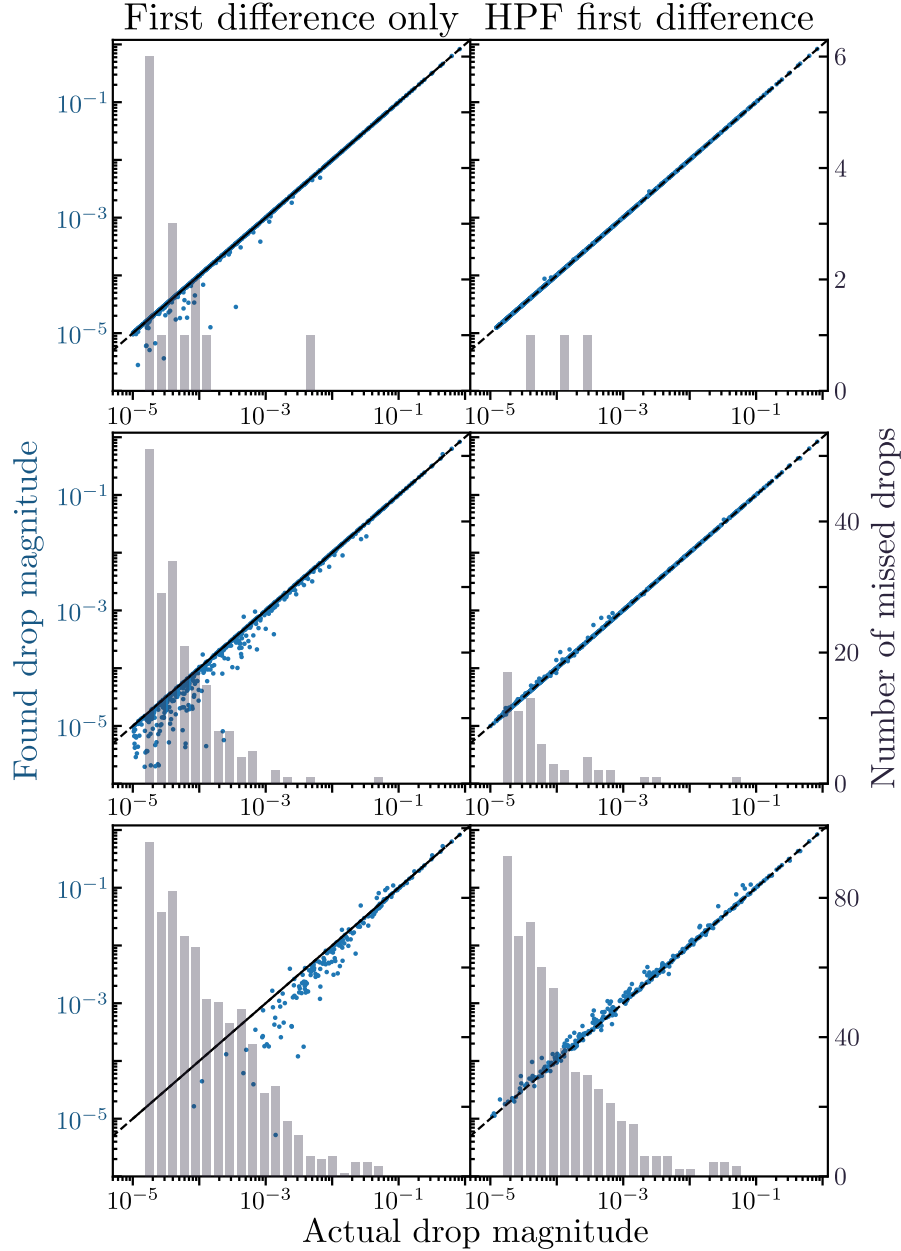


Figure 4.5: Drop-finding method comparison: accuracy of individual drops. For the three synthetic time series shown in Figure 4.4 (event rates 0.05, 0.8, and 12.8 events per recharge timescale, from top to bottom), the accuracy of measuring 1000 drops is displayed for both the bare first difference method and the addition of a high pass filter. The blue scatter points in each plot show the measured magnitude versus the actual (true) magnitude. The black dashed line would be the two matching exactly, and points falling below the line represent drops whose magnitude was underestimated by the method. The transparent gray histogram shows the drops that were not found by the method, binned by their actual magnitude. Both methods are able to measure the largest events. For smaller events, the bare first difference method (left) misses and underestimates many more events than does the method with the high pass filter applied.

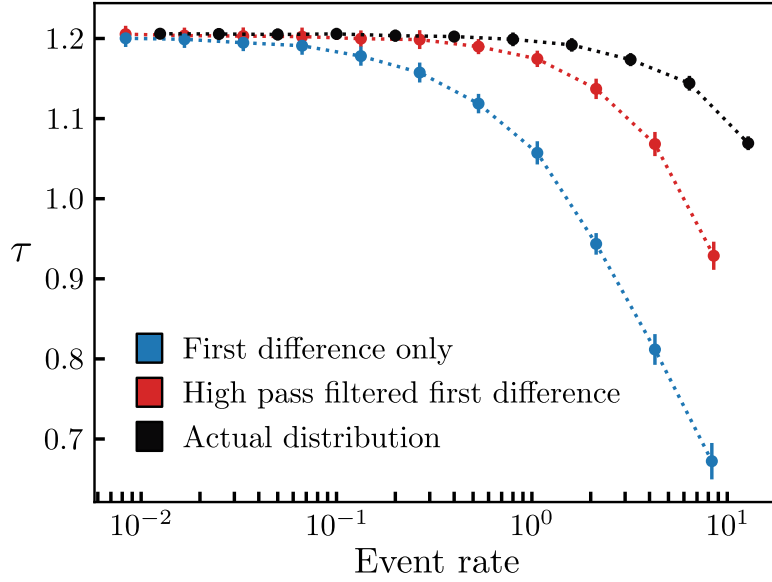


Figure 4.6: Drop-finding method comparison: τ measurements on synthetic data. Shown are the values of τ measured for the synthetic time series data in Figure 4.4. When using only a simple threshold on the first difference, the observed power law exponent quickly deviated from the true value. With the addition of a high pass filter the method performed much closer to optimal. The event rate is the average number of drop events to occur during T , the timescale of the exponential recharge following an event. Error bars are the 2σ uncertainty as measured by nonparametric bootstrap (see Section 4.4).

it reequilibrates. The first overlap is when events are frequent enough to occur during the long timescale perturbation following a previous drop. This is the form of overlap corrected by the addition of a high pass filter to the first difference, which removes the time series perturbation. The second overlap happens when the short timescale drops coincide, necessarily at a much higher event rate. Statistical methods have been proposed to recover the correct stress drop distribution in this regime [72], but lean on extra assumptions. Namely, when all events are recoverable, the degree to which their occurrence is a Poisson process is unimportant. When significant overlap takes place, it becomes necessary to account for exactly how a drop event affects the chance of another occurring at the same time.

4.4 Distribution fitting and uncertainty estimation

Equipped with the drop-finding method of the previous section, we located the stress drops in all of the plasticity data from $0.1 < \epsilon < 0.2$ and compiled a drop magnitude distribution for each particle shape. We normalized the stress drops for each compression test by the average deviatoric stress in the plastic regime, as defined in Section 3.2. This allowed us to directly compare fractional drops for shapes even when the stress scales differed by a factor of three.

As stated in the introduction to this chapter, the fitting form for the distributions was a truncated power law, i.e. $P(s) = Cs^{-\tau}e^{-s/s^*}$. C is a normalization constant so that the probability density integrates to unity and is therefore not a free parameter in the family of distributions. The two free parameters are τ , the power law exponent, and s^* , the upper event size cutoff.

When fitting power laws to data, it is common to use the fact that a power law is a straight line on a log-log plot and measure the exponent from a linear regression on the transformed histogram of the data. This works for a first pass over the data, but to be rigorous, it is problematic for a few reasons. First, binning introduces arbitrariness in the form of the bin width [72] and should not be involved in the fitting process if it can be

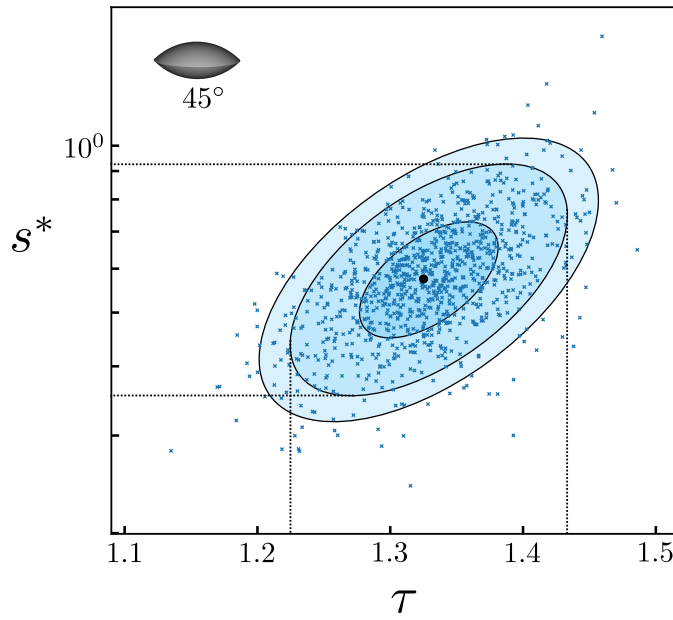


Figure 4.7: Confidence region estimated through bootstrapping. After collecting all of the drops for the 45° lenses, we bootstrapped 1000 times (described in the text) and used the scatter of (τ, s^*) points to estimate our confidence region. Working outward from the best fit parameters (black dot) to the empirical dataset, the 1σ and 2σ regions for the joint distribution and the 95% confidence region are represented by ellipses. The bounds of the confidence region for the joint distribution also conveniently yield the confidence intervals for the marginal (univariate) distributions.

helped. Using the cumulative density function instead of the probability density function (the CDF instead of the PDF) sidesteps histogram binning but adds correlation between the data points ($C(x_i) = C(x_{i-1}) + 1/N$), a problem for linear regression. Second, linear regression assumes independent Gaussian errors on data points, which holds for the raw data but is lost after transforming the data to double logarithmic space. Third, probability distributions have the constraint that they must integrate to 1, something for which standard linear regression does not account. A great source on this subject is [81].

We used maximum likelihood estimation (MLE) on the unbinned event magnitudes to find the best fit distribution [81]. The idea of MLE is that the entire dataset has some probability of occurring under each of the possible truncated power law distributions, and the best fit distribution is the specific one that maximizes the probability (or likelihood) of observing the data we did. This was done by numerical minimization, where the quantity to minimize (over (τ, s^*)) was the negative log likelihood of the event magnitude distribution. Minimizing the negative log likelihood is equivalent to maximizing the log likelihood, which is in turn equivalent to maximizing the likelihood — though opting for the latter is a good way to run into numerical underflow errors. MLE can only select the best out of a family of distributions; it alone cannot tell how good the fit is or if there is a better family of distributions to use.

After an optimal (τ, s^*) is found, there are various ways to measure the uncertainty in these parameters. One is to perturb the dataset in a statistically rigorous manner as outlined below, re-fit each time, and use the spread in (τ, s^*) space to measure the uncertainty. This method is called bootstrapping [88].

When measuring a distribution from observations, we implicitly assume there is a true underlying distribution from which Nature sampled during the physical process. The bootstrap method uses the observed dataset as our best guess for that underlying distribution for the data, and samples from it to generate realistic variants as if the experiment were done many times [88]. Intuitively, the simplest model we have about the true underlying

distribution (i.e., the model with the fewest additional assumptions) is the data itself, so we use the empirical distribution as the underlying distribution and sample new datasets from it. In practice, this means taking the original datapoints and sampling with replacement until the length of the new dataset matches the original. There will be duplicates of some datapoints and others will be left out. The same fitting procedure as before is applied to each of these new datasets, and the spread in the best-fit (τ, s^*) gives a quantitative measure of the uncertainty in these parameters for the original dataset.

We generated 1000 bootstrapped datasets for each shape and obtained a point cloud in (τ, s^*) space, as shown in Figure 4.7. The point clouds were well fit by bivariate Gaussian distributions in $(\tau, \log(s^*))$ space, such that the covariance matrix of the point clouds gave our 1σ and 2σ uncertainties for both the joint and marginal distributions.

4.5 (τ, s^*) for each particle shape

For the experimental drop distributions across all shapes, 16 of which are shown in Figure 4.8, the truncated power law form fit well. In these plots the data was logarithmically binned (blue dots), meaning the bin widths are constant in $\log(s)$ and therefore grow exponentially with s . This was only for display purposes and was done after fitting.

All the distributions in Figure 4.8 have roughly similar shape, with decreasing probability for increasing stress drop size and a drop off for larger drop sizes. The lower end of the accessible range of s in these distributions was given by the experimental noise floor, while the upper end corresponded to catastrophic events with drop magnitudes that were a significant fraction of the average stress.

For each of the 22 particle shapes' drop magnitude distributions, Figure 4.9 shows the best fit (τ, s^*) (black dot), 1σ uncertainty region (dotted ellipse), and 2σ uncertainty region (colored ellipse). Many observations are in order, first about the nature of the uncertainty ellipses and then about the (τ, s^*) values themselves.

The uncertainty shrank in both fit parameters as s^* increased. A larger value of s^* means

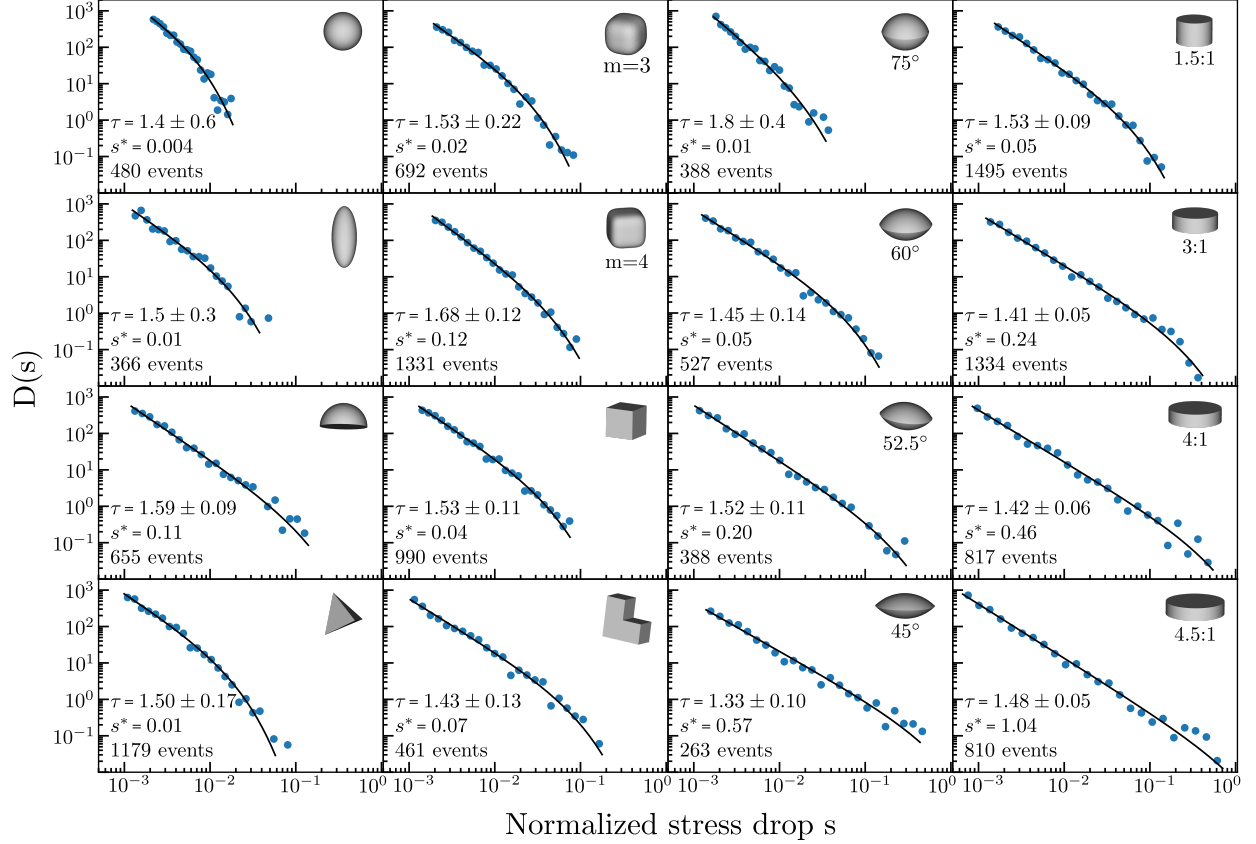


Figure 4.8: Fitted stress drop magnitude distributions. The drop distributions $D(s)$ were fit by Maximum Likelihood Estimation to a truncated power law of the form $P(s) = Cs^{-\tau}\exp(-s/s^*)$. The values cited for the uncertainty on τ are 2σ for the marginal distribution, and only the best fit value of s^* is listed. All plots are over the same range. Data points were binned logarithmically (blue) after fitting (black). The distributions are organized by (from left to right) the first column containing spheres, simple deviations from a sphere, and tetrahedra, the second column cubes and simple deviations from a cube, the third column the lens family, and the fourth column the disk family.

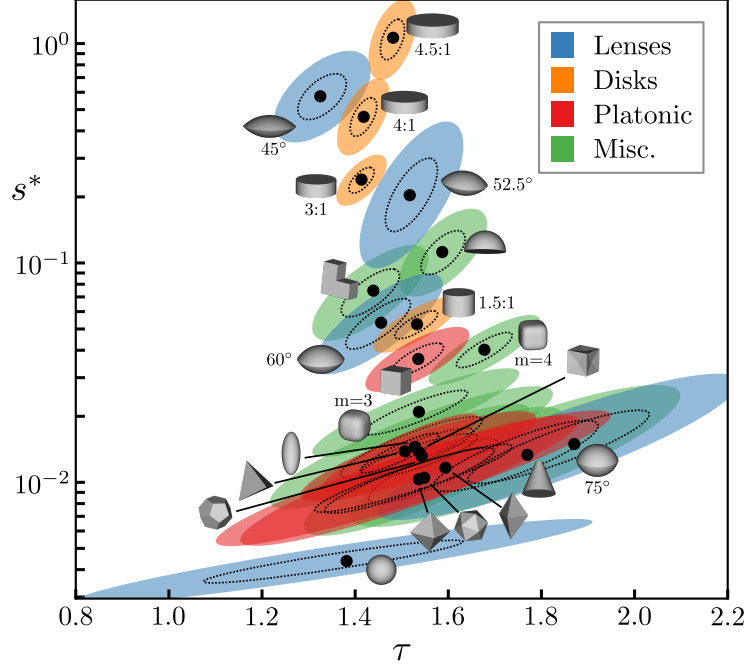


Figure 4.9: (τ, s^*) parameter space for the truncated power law fits to $D(s)$. The dotted (solid) ellipses around each best-fit (τ, s^*) are the 1σ (2σ) confidence regions as determined by the nonparametric bootstrap method.

a greater spread in event sizes to which we could fit the truncated power law distribution, meaning the fitted distribution were more determined. Thus for the best measurement of drop distribution parameters, particle shapes at the top of Figure 4.9 are desirable.

The uncertainty ellipses are skew, with a positive correlation between the fit parameters τ and s^* . We can understand this correlation between the parameters with the following intuition. The event magnitude for the rollover between the power law and the exponential suppression is like a hinge point for the distribution. If this rollover point is underestimated, there will be too many found events with a magnitude larger than the rollover which each have an exponentially suppressed likelihood, greatly lowering the likelihood for the overall dataset. If the rollover magnitude is overestimated, the normalization constant C decreases, lowering the likelihood for every one of the points in the dataset. Comparatively, changes in τ and s^* are less significant as long as the rollover point is accurate.

4.5.1 *Shapes present a new way to vary s^* and the distance to criticality*

In the view that many plastically deforming amorphous systems, including granular materials, operate near a non-equilibrium critical phase transition [89, 74, 15], the value of s^* increases with proximity to the critical point. The clear shape dependence shown in Figure 4.9 indicates, then, that by varying particle shape we are able to create granular aggregates which plastically deform further from or closer to the point in system space where stress drop cascades are entirely scale free and correlated behavior spans the system.

There are various ways to change a plastically deforming system's proximity to criticality. Disorder inhibits coordination and makes a system behave more as a collection of individual units than as a cohesive unit, as in the Random Field Ising Model [90]. In a plasticity model of Dahmen *et al.* [74], the concentration of voids in an abstracted granular system plays the role of disorder and determines s^* . In another [73] it is the amount a relaxing site weakens the surrounding parts of the system. In [56] the stiffness of the driving apparatus limits the maximum size of rearrangement cascades, and in [25] it is the imposed shear strain rate.

In this experimental work we introduce particle shape as a powerful new means to vary s^* . By modifying particle-scale details of the system, we were able to span more than two orders of magnitude in s^* (Figure 4.9). This was nearly the entire range possible in our experiment, considering the sphere data encroached on the detection limit arising from the noise level of the Instron, and the normalization of event magnitudes capped s^* at values only slightly larger than 1.

Thus, to study the nature of the nonequilibrium critical phase transition thought to describe the plasticity of a wide class of amorphous materials, all that is required is a set of spheres, cubes, and lentils. The effect particle shape has on the interactions between particles modifies microscale details in a way that is fundamentally the same as raising the void density in a lattice model [74] or dialing up the strain rate on a packing of glass beads [25].

4.5.2 *Microslips constitute the cascades of rearrangement*

An important point that emerged from Figure 4.9 is that s^* , the characteristic (fractional) stress scale for the largest events, varied over two orders of magnitude while the ratio of system volume to particle volume remained fixed, aside from minor variations in packing fraction. This clearly demonstrates that s^* was not tied to system size. In fact, since in our experiments there were on the order of 10^2 particles within a column cross section (fewer for the oblate shapes since they tended to pack horizontally — see Section 6.2), a cutoff s^* between 10^{-2} and 10^{-1} is of the same order as if a single particle went from load-bearing member of the packing to unstressed rattler.

From this alone we can infer that the power law regime in $D(s)$ is not due to large cascades of many-particle relaxations. Indeed, the x-ray imaging confirms that the vast majority of the measured stress drops correspond to detectable shifts in the position of no more than a couple particles (Figure 4.1). This suggests partial slips at the particle contacts, which relax only a portion of the supported force, as the elementary components of the scale-free cascades constituting stress drops. These would be similar in nature to the partial stress drops thought to occur during earthquakes [91] and the partial stress drops incorporated in some mesoscale models[74]. Supported by qualitative results from x-ray imaging, this suggests that the fluctuation statistics were dominated, in terms of frequency, by cascades of partial stress drop ‘microslips,’ which only occasionally built into larger reorganization events involving groups of particles which moved significantly. Such microslips allow for small shifts in the relative particle positions across a contact, thereby changing the magnitude of the transmitted force without necessarily breaking the contact. In this context platy, oblate particles provided contact geometries well suited to accommodate many microslip events during a given relaxation event, suggesting a reason for their large V and s^* values.

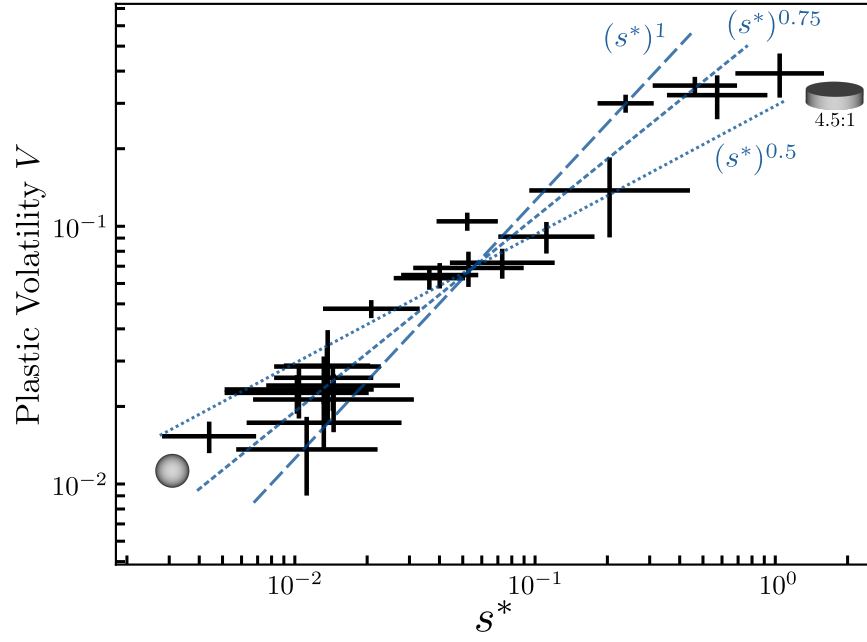


Figure 4.10: Volatility as an event-agnostic method for measuring s^* . Whereas measuring s^* requires all of the drop-finding and distribution-fitting labor of the previous sections, volatility is quickly measured from raw data and can be used as a proxy for distance from criticality. Shown are power law relationships $V \sim (s^*)^1$, $(s^*)^{0.75}$, and $(s^*)^{0.5}$ as guides to the eye.

4.5.3 Volatility as a method to approximate s^*

In Figure 4.10 we show that volatility, beyond serving as an event-agnostic tool for quantifying fluctuations (Section 3.3), can also be used as a method of measuring s^* and the distance to criticality. In the early stages of an experiment, where a wise scientist gets a lay of the land by varying all available parameters as much as possible, s^* can be quickly and easily approximated before having to spend the time to rigorously locate and measure drop event magnitudes.

We can shed light on the clean relationship between the volatility and s^* shown in Figure 4.10. The volatility is approximately the standard deviation of fractional changes in the stress data $q(\epsilon)$. The fractional changes are significantly negative for stress drops and then nearly zero everywhere else. Therefore the distribution of fractional changes in the data is approximately the distribution of fractional drops $D(s)$, combined with a multitude of values near zero which will only show up as a constant factor in the moments of the distribution (especially the higher order moments, which are less affected by values near zero).

$$\text{Var}(s) = \int s^2 P(s) ds = C \int_{s_{\min}}^{\infty} s^{2-\tau} e^{-s/s^*} ds \quad (4.3)$$

If we make the approximation that the exponential suppression kills the probability enough for us to ignore everything above s^* , we can evaluate the integrals. As a reminder, C is the normalization constant of $P(s)$.

$$C^{-1} \sim \int_{s_{\min}}^{s^*} s^{-\tau} ds = \frac{1}{1-\tau} ((s^*)^{(1-\tau)} - s_{\min}^{(1-\tau)}) \quad (4.4)$$

$$\text{Var}(s) \sim \left(\frac{1-\tau}{3-\tau} \right) \frac{(s^*)^{(3-\tau)} - s_{\min}^{(3-\tau)}}{(s^*)^{(1-\tau)} - s_{\min}^{(1-\tau)}} \quad (4.5)$$

If $1 < \tau < 2$, the term with s^* dominates in the numerator and the term with s_{\min}

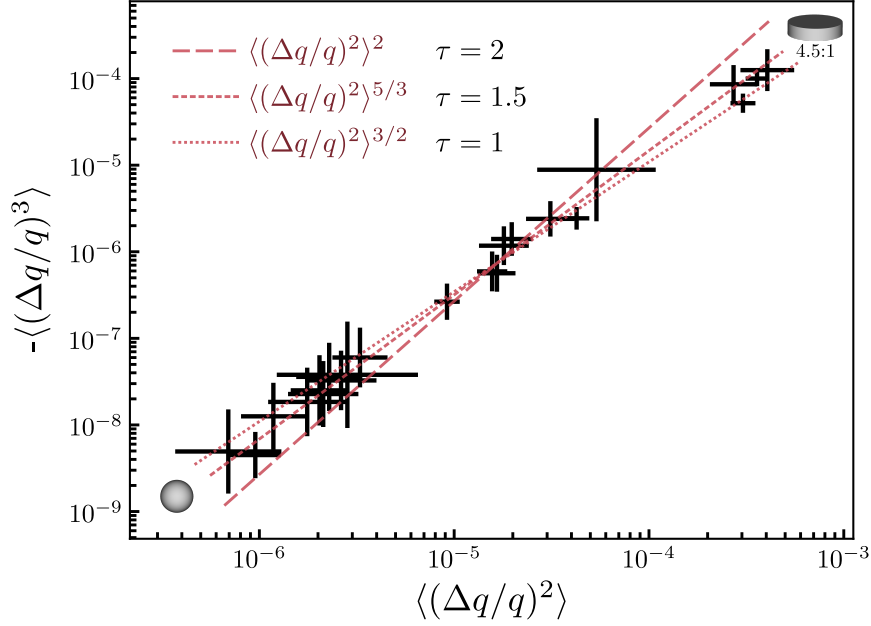


Figure 4.11: Approximating τ with the moments of the distribution of fractional changes in $q(\epsilon)$. Comparing the scaling of the second and third moments of the distribution of fractional stress changes yields a clean relationship. However, as shown by the three power law scalings (as guides to the eye), the discerning power between different values of τ is poor due to the large spread in the values. The error bars displayed are the standard deviation of the spread of logarithmic values of the corresponding moments.

in the denominator. Thus $\text{Var}(s) \sim (s^*)^{3-\tau}$ (with s_{\min} constant for all shapes, set by the noise floor of the Instron), and the volatility grows as s^* to a power $(3 - \tau)/2$ that falls somewhere between 0.5 and 1. Indeed, in Figure 4.10, the volatility grows roughly as $(s^*)^{0.75}$, corresponding to a value of τ equal to 1.5.

This suggests it might be possible to infer the value of τ without actually finding all of the stress drops. We begin with the assumption that stress drops are sampled from a truncated power law with an exponent τ between 1 and 2. The higher moments of $D(\Delta q/q)$ will more successfully ignore the sea of small fractional changes in the time series and emphasize the values of $D(s)$. We therefore compare the third moment of $D(s)$ to the second and find (using the above relationship for s^*)

$$-\langle(\Delta q/q)^3\rangle \sim \langle(\Delta q/q)^2\rangle^{\frac{4-\tau}{3-\tau}} \quad (4.6)$$

We check this scaling in Figure 4.11. The relationship between these moments of the fractional change distribution is quite clean, though its discerning power is not quite strong enough to tell the value of τ between 1.3 and 1.5, given the spread in measurements.

4.5.4 *Implications of τ : How do cascades proceed?*

Regarding the exponent τ , clean trends were harder to isolate. Given the fit uncertainties, most shapes' τ fell in the range 1.3–1.7 and would be compatible with a value of 1.5. As discussed earlier in this chapter, recent mesoscale plasticity models have focused on the kernel used to describe stress redistribution after local yielding. The kernel used in [74] resembles sandpile models where stress is offloaded isotropically to all neighbors, which lends itself to mean field treatment and a value of τ equal to 3/2. Many systems, from nanocrystals to the Earth's crust, have been claimed to exhibit universal deformation with this value of τ [26, 28, 45, 23, 25].

The competing view in the field centers on the need to incorporate the work of Eshelby [80], who found that if a portion of an elastic medium suddenly relaxes its stress by $\delta\sigma_{ij}$, the rest of the medium adjusts its stress in a long-ranged quadrupolar fashion. The form is

$$\Delta\sigma_{ij}(\vec{r}) \propto (\delta\sigma_{ij}) \cos(4\phi)/r^2 \quad (4.7)$$

Importantly, this redistribution is anisotropic — some parts of the system increase in stress and others decrease — and far-field — it decays only as a power law in distance, rather than exponentially. Due to the anisotropy of the redistribution, systems with this stress response cannot be treated in a mean-field manner [15]: they must be simulated explicitly and values for τ are therefore numerical.

Lin *et al.* found $\tau \approx 1.3$ in a three dimensional scalar lattice model [14], while a tensorial approach to more accurately represent stress and strain, observed $\tau=1.28$ in simulations across a wide range of loading conditions [30]. Many other authors of elastoplastic models

find similar values of τ [77, 56, 29].

That the value of τ observed in the particle shape experiments less closely matches that for models with a quadrupolar stress redistribution might indicate either that there is some effective degree of incipient shear banding in the experiments (though no fully formed shear bands were observed in any trials) or that the mesoscale granular characteristics of the aggregate material in our system do not reflect the continuum elastic response encapsulated by a quadrupolar kernel.

The distinction between $\tau = 1.3$ and $\tau = 1.5$ is small in practice, as can be seen by the span of the uncertainty clouds in Figure 4.9. Experimental studies have had difficulty distinguishing between the two, though claims are often made about compatibility with one side while neglecting to mention the other. Our study claims no final verdict on the debate either, though compatibility is slightly higher with $\tau = 1.5$. This compatibility can also be seen when more than 100 other datasets from a wide range of shapes, surface and material properties are plotted together in Figure 4.12. In these datasets are laser-cut acrylic particles, surface-treated particles, non-convex particles like jacks [36], and particles with larger and smaller volumes. It should be cautioned, however, that various aspects of the experiment were less controlled and the data was less scrutinized than for the experiments of Figure 4.9. The same drop-finding and distribution fitting methods were applied.

4.6 Strain rate independence of τ

In simulated plastic systems, it is common to drive by a protocol that is called event-driven [56, 29]. In this scheme, the driving rate is taken to be so slow that events happen in complete isolation of one another. Practically, this means driving the system to its next event and then letting the dynamics run to completion before resuming external driving. There is no overlap between events.

In physical experiments we must have a nonzero driving rate, which leads to strain rate effects. In stick-slip studies, the driving rate determines the regularity of slip events [78],

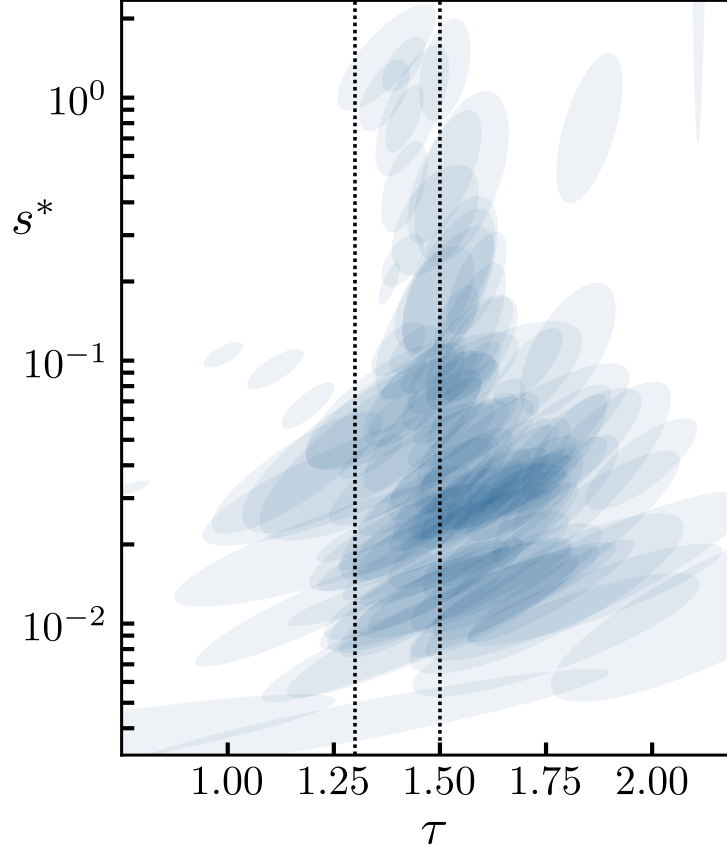


Figure 4.12: (τ, s^*) fits for over 100 datasets. The datasets feature an eclectic mix of particle shapes, materials, size scales, and surface treatments. Included are more than a dozen different laser-cut acrylic particle sets, particle sets that were smoothed by tumbling in polishing media, years-old 3D-printed particles (where the plastic becomes more brittle), and particles with volume larger and smaller than 22.5mm^3 by a factor of 3-4. The dotted lines mark $\tau = 1.3$ and 1.5 . It should be cautioned, however, that various aspects of the experiment were less controlled and the data was less scrutinized than for the experiments of Figure 4.9. The same drop-finding and distribution fitting methods were applied.

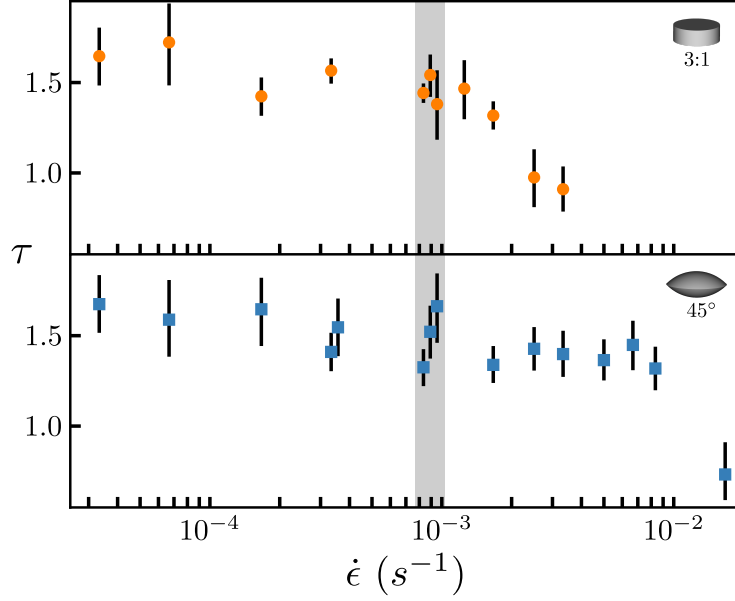


Figure 4.13: Strain rate dependence of τ . Uncertainties were calculated as before with the bootstrap method. The gray region contains the points which were taken at the same strain rate as the rest of the data, $\dot{\epsilon} = 8 \times 10^{-4} \text{ s}^{-1}$. These points for both particle shapes and the $\dot{\epsilon} = 3 \times 10^{-4} \text{ s}^{-1}$ data for the lenses are offset slightly to display multiple sets of runs taken at the same strain rate.

and whether or not they happen at all [57]. When relaxation is present in the system, as we will discuss in Chapter 5, interesting self-organized criticality behavior can arise as the driving rate is lowered to the point that significant relaxation occurs between events [87].

To address the effect of strain rate, we selected two shapes with large s^* and therefore lower uncertainty on the power law exponent τ , and ran dozens of extra compression experiments over more than two orders of magnitude in strain rate. The shapes were 3:1 disks and 45° lenses, and the resulting measurements of τ are shown in Figure 4.13. The measured value of τ remained roughly constant up until a fastest strain rate that is different for the lenses and disks.

The constancy in τ is remarkable considering that the particles wore down over these tests (especially the lenses; see Chapter 6) and new particle sets had to be printed. During the weardown, shear strength steadily lowered to the level of spheres yet τ remained relatively unchanged. This is perhaps unsurprising given the theoretical underpinning for τ previously

discussed, as the nature of the cascades should not change with particle surface smoothing. From an experimental standpoint, however, the constancy is strong evidence for the success of the drop-finding method and for the universality of the dynamics underlying amorphous plasticity.

We can explain the shape-dependent rolloff to smaller τ by referring to Figure 4.14. The 3:1 disks and the 45° lenses differ in the rate of observable drop events by about a factor of 5. The strain rate where τ lowered from 1.5 in Figure 4.6 differs by nearly the same factor of 5 for the two shapes. As has arisen many times in this thesis, the timescale for a single event is constant across the board for all the shapes, 50-100ms, so when average time between events shrinks due to faster strain rates of compression, events begin to overlap. When events overlap, small events are hidden in the shadow of large ones and large events become even larger (as discussed in Section 4.3). This lowers the value of τ , thus explaining the rolloff seen in Figure 4.13 (and reminiscent of Figure 4.6).

We also were safe in our original choice for the rate of compression as $\dot{\epsilon} = 5\%/min$, since the 3:1 disks are still in the constant τ plateau at this strain rate (shown shaded in gray in Figure 4.13) and have nearly the highest event rate (second only to the 1.5:1 disks).

4.7 The other component of plastic deformation: sliding

Figure 4.14 raises important questions about how packings of different shapes navigate plastic deformation. First, particle shapes in the disk family systematically decreased in event rate as their s^* increased. If we assume that there is a fixed amount of nonaffine rearrangement that must happen in a generic granular packing – with constant volume particles – as it deforms from 0 to 0.2 strain, larger events mean that fewer need to occur.

Based off this picture, though, why did some shapes exhibit fewer events while their distribution had the same s^* ? In the raw data of Figure 4.15, it is clearly visible that the 4:1 disks deformed with more stress drops than the 45° lenses, even though everything else about their plasticity data is similar. To answer this, we return to a simple description of

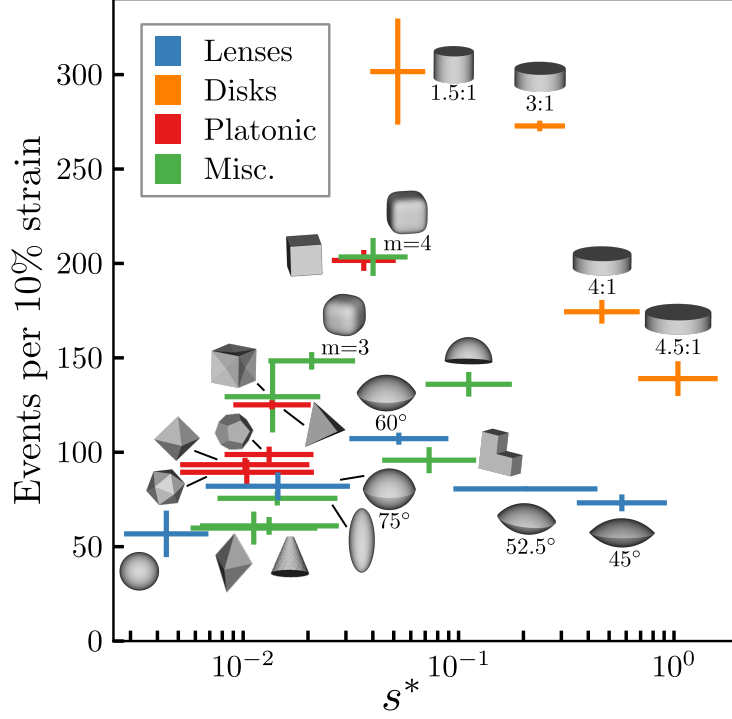


Figure 4.14: Detectable event rate for different particle shapes. Uncertainty for the event rate is calculated as the standard error of the mean count per run.

the deformation of the granular packing.

Though not strictly true, it is helpful to think of the total strain rate of deformation on the sample as the sum of plastic and elastic components,

$$\dot{\epsilon} = \dot{\epsilon}_{\text{elastic}} + \dot{\epsilon}_{\text{plastic}} \quad (4.8)$$

The total strain rate $\dot{\epsilon}$ is a constant: the strain rate imposed by the Instron. The elastic component of strain is proportional to the stress measured, with the Young's modulus as the constant of proportionality. Rearrangements could presumably change the Young's modulus of the packing as new configurations are entered, but we approximate it as a constant between events. Thus we can relate the derivative of the measured stress to the elastic strain rate, $\dot{q} = \frac{dq}{d\epsilon} \dot{\epsilon} = E \dot{\epsilon}_{\text{elastic}}$. Stress drop events, where $\dot{q} < 0$, are moments of high plastic strain rate, momentarily shooting $\dot{\epsilon}_{\text{plastic}}$ past the imposed strain rate $\dot{\epsilon}$ and making $\dot{\epsilon}_{\text{elastic}} < 0$.

There must be an additional component of $\dot{\epsilon}_{\text{plastic}}$, however. As evident in the shape of

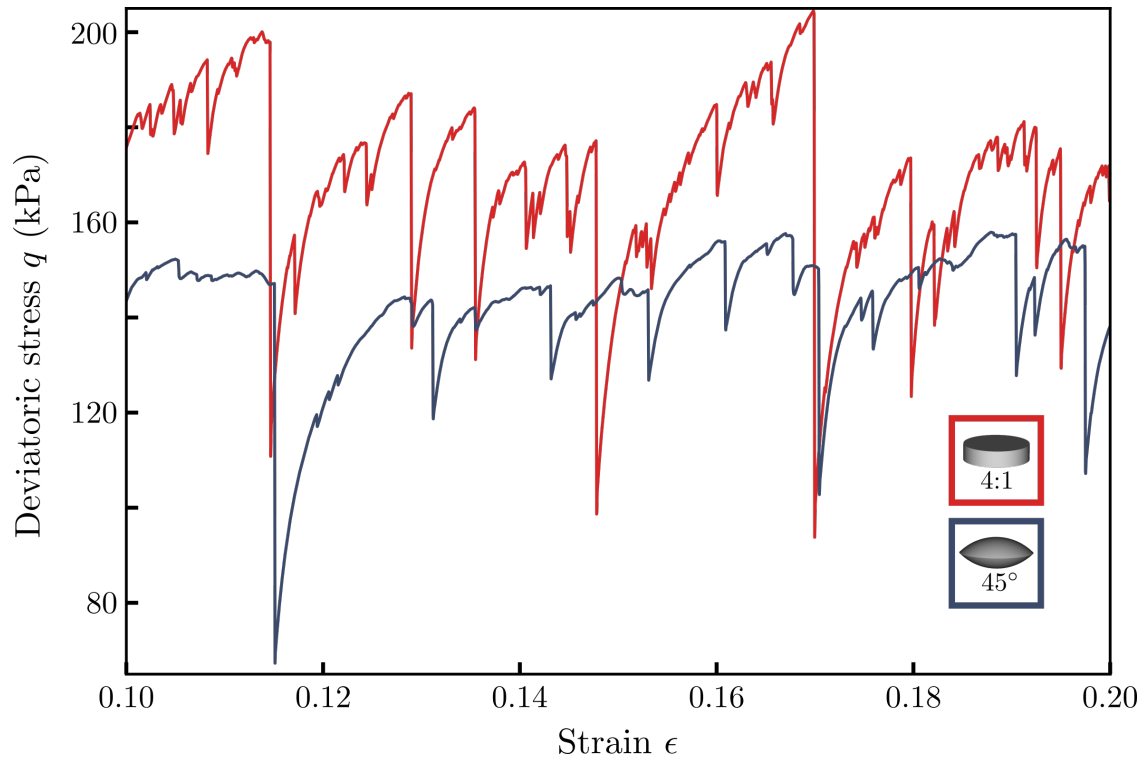


Figure 4.15: Comparison of raw $q(\epsilon)$ data for 4:1 disks and 45° lenses. These particle shapes deformed at a similar plastic stress and their drop distributions are nearly identical, yet the disks plastically deformed with three times as many rearrangement events.

the recharges following events, $\frac{dq}{d\epsilon}$ and therefore $\dot{\epsilon}_{\text{elastic}}$ diminished as the stress increased. This is consistent with a plastic strain rate that is always nonzero, not just during stress drop events, and that increases with stress, leading to plateaus in stress when the plastic strain rate matches the imposed strain rate.

We can now begin to understand the different rates of stress drop events across particle shapes in Figure 4.14. By the end of a compression experiment, nearly all of the applied strain went into plastic deformation. Packings of each particle shape divided the total plastic deformation between sudden events and background sliding. The lens family maintained a low event rate no matter the aspect ratio, indicating a large role in plastic deformation played by sliding and a lesser role played by sudden events than with the disk family. From this observation we might guess that the lens particles wear down more rapidly under repeated compression tests than other particle shapes; this story will be fleshed out in more detail in Chapter 6.

The gradual sliding component of plastic deformation happens continuously while a granular packing supports shear stress, and it is the topic of our next chapter. We will probe it by applying deformation and then stopping, such that $\dot{\epsilon}$ is zero and the changing stress is a direct measure of the plastic strain rate.

4.8 Summary

In this chapter, we focused on the sudden stress drop events that punctuated plastic deformation for every particle shape. Careful analysis of the plasticity data was needed to accurately account for small events and to fit the compiled distribution of event magnitudes. The results were worth the effort, and entered right into a current debate about how detailed a description of stress redistribution is needed to accurately describe the plasticity of amorphous materials. Particle shape was found to be a new means to vary the granular system's proximity to the critical point. In essence, all of the features of shape – the location of corners, the form of the rounded surfaces, the precise mass distribution, etc. – reduce to

a single value, s^* , when it comes to their effect on the rearrangement cascades during plastic deformation. Checking for the effect of strain rate led to the observation that, in addition to the several quantities considered in this chapter and the previous one, shape also dictates the rate of drop events over a compression experiment. We now turn to the form of plastic deformation that is an alternative to the sudden drop events of this chapter.

STRESS RELAXATION AND MEMORIES

5.1 Introduction

The previous chapters have focused on the signs of rearrangement that jump right out of the stress-strain data. We now turn to the background restructuring that occurs whenever stress is supported by the system. It exists as a glacial sliding at all the contacts in the packing, by microns or even nanometers over the course of an experiment, yet these small movements reroute the force network and bolster the packing against further deformation. The experiments are similar to the plasticity ones from Chapters 3 and 4: a granular column is prepared in the same way and then compressed, but now to only an intermediate strain where it is held while the Instron records the force. We find that the slow restructuring process reveals a complex inner state of the packing capable of storing memories which are then autonomously read out over the course of seconds, minutes, or even hours. Further, while a broad class of disordered systems are capable of the complex behavior we find, there are parts to the story that are inherently granular in nature. As in Chapter 4, the insights gained from these experiments add not only to an understanding of how granular materials support stress, but also to the broader story of how disordered systems adapt, slowly and significantly, to external driving.

5.2 Single step stress relaxation

The stress relaxation of soils, and its more relevant cousin, creep, have a rich history of study [92, 93, 94]. During stress relaxation experiments, the sample is held at constant strain while the stress is monitored. In creep, a constant stress is applied and deformation is measured. The latter is important, for example, when subjecting the earth under a structure to a constant load for decades or even centuries. The Leaning Tower of Pisa is probably the

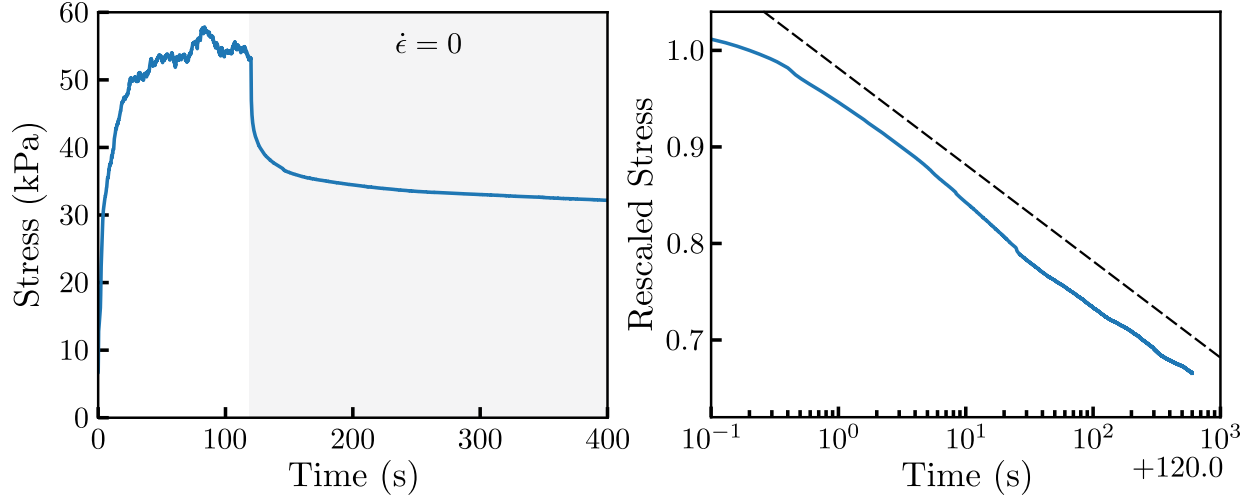


Figure 5.1: Stress relaxation in a packing of glass beads. *Left:* The packing was compressed to $\epsilon = 10\%$ at $\dot{\epsilon} = 5\%/min$ and then held at a constant displacement (shaded region) while the stress was measured. *Right:* We display time logarithmically from the start of the hold and rescale the stress by its value at $t_{\text{Instron}} = 200\text{ms}$. The dashed line is a guide to the eye, showing a logarithmic decay (i.e., $f = A - B \log(t)$).

most famous example of soil creep gone awry, though in actuality only 1.5 of the tower’s 5.5 degrees of tilt (prior to restoration) come from soil creep [95]. In this work we focus exclusively on stress relaxation due to its ease of measurement and because the relaxation dynamics, our object of study, are allowed to progress without any experimental intervention.

A stress relaxation experiment is identical to a plasticity experiment from Chapter 3, except that midway into it we abruptly set $\dot{\epsilon} = 0$. The compression sets up the plastic sliding $\dot{\epsilon}_{\text{plastic}}$ discussed in Section 4.7, which we watch play itself out by suddenly removing its energy source, the uniaxial compression. In other words, we drive the granular packing into a typical state it has during plastic deformation and then study that state by removing the drive and observing how the packing evolves when left to its own devices.

Immediately after halting compression, the stress of a granular packing relaxes approximately logarithmically in time (Figure 5.1), in line with past works [96, 97]. During logarithmic relaxation the same decrement of stress occurs between 1 and 10 seconds as between 1 and 10 years. It is a process that is always slowing down and often shows no sign of stopping on experimental timescales. In a creep experiment on San Francisco bay mud, the

sliding velocity of $200\mu\text{m}$ clay particles was estimated to be somewhere between angstroms per second and angstroms per hour, a truly glacial pace [93].

There is no inherent timescale in logarithmic relaxation, as opposed to many of the canonical examples of relaxation in physics (e.g., the discharging of a capacitor, Debye relaxation of dielectric materials, and radioactive decay). Despite the quirks of logarithmic relaxation, a wide range of systems in nature are found to exhibit it. A few examples are the breaching time for colloidal particles at an interface [98], the current in superconductors [99], the volume of a crumpled piece of paper under weight [100], and the real area of frictional contact between two interfaces [101, 102]. This suggests some universal physics at play in all of these systems, which we can probe with our granular setup.

5.2.1 *Measuring $\dot{\epsilon}_{\text{plastic}}(q)$ to fit stress recharges in plastic deformation*

With just the observation that the stress relaxation of granular materials is logarithmic in time, at least approximately, we can make sense of the shape of stress recharges in the raw compression data of Chapter 4. When $\dot{\epsilon} = 0$ in Equation 4.8, the plastic strain rate is proportional to the rate of stress decrease. Namely, $\dot{\epsilon}_{\text{plastic}} \propto -\dot{q}$. Taking $q = A - B \log(t)$,

$$\dot{\epsilon}_{\text{plastic}} = ae^{q/b} \tag{5.1}$$

In the presence of external driving, this relationship suggests that the plastic sliding in a packing rapidly ramps up when the stress is around the value b . When the plastic strain rate becomes comparable to the compression strain rate, the stress levels off because very little of the applied shear strain results in elastic deformation. We use this form to derive the shape of a stress recharge, by reincorporating the $\dot{\epsilon}$ term.

$$\dot{q} = E(\dot{\epsilon} - \dot{\epsilon}_{\text{plastic}}) = E(\dot{\epsilon} - ae^{q/b}) \tag{5.2}$$

The solution to this equation is, after redefining the stress and strain to be zero at the

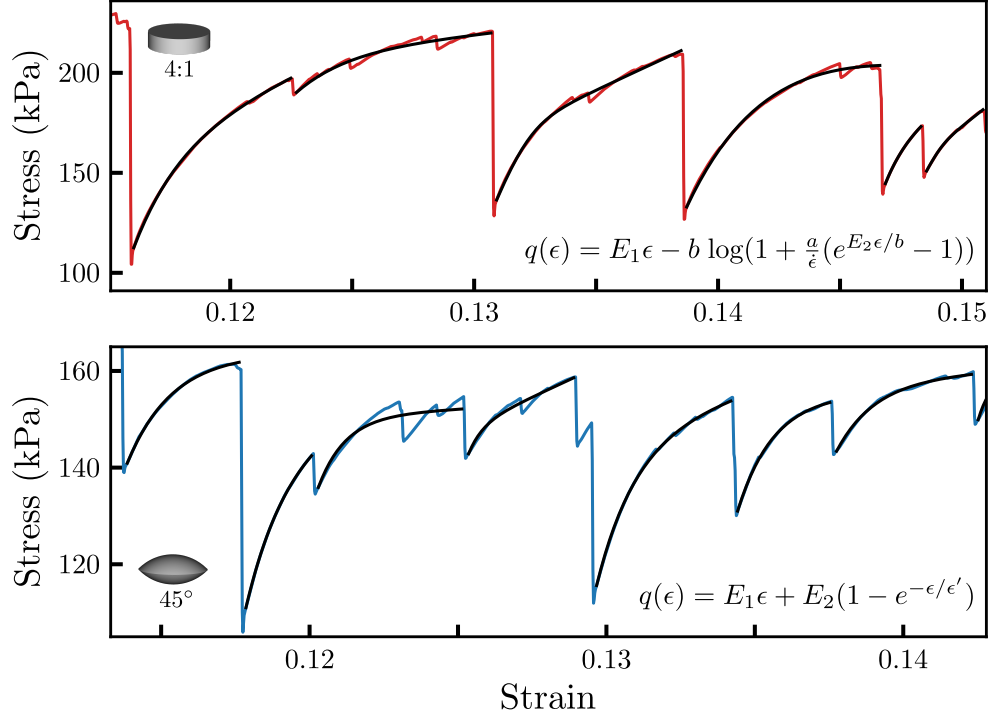


Figure 5.2: Stress recharge after drop events, fitted two ways. In the top plot, we use the functional form derived from stress relaxation that is logarithmic in time (Equation 5.4) to fit stress recharges in plasticity data for a packing of 4:1 disks. In the bottom plot, we use the exponential recovery form (Equation 5.5), like what was used to generate the synthetic data of Section 4.3, to fit the recharges in a packing of 45° lenses. The plasticity data was divided into intervals separated by large stress drop events, and the intervals were each fit (black) with the corresponding functional form. There is no significance in the particle shapes chosen; both functional forms seem to be interchangeable and work well on any plasticity data that features large recharges.

beginning of a recharge (immediately after a stress drop),

$$q(\epsilon) = E\epsilon - b \log\left(1 + \frac{a}{\dot{\epsilon}}(e^{E\epsilon/b} - 1)\right) \quad (5.3)$$

Note that here and elsewhere in this chapter we refer to the natural logarithm (base e) by \log . When $E\epsilon/b \ll 1$, only the $E\epsilon$ term is significant as in purely elastic loading. As $E\epsilon/b$ grows, the stress asymptotes to a value $q = b \log(\dot{\epsilon}/a)$, which is sensibly larger for faster strain rates.

We find that in the data, the stress does not always level off in the late stages of a

recharge. Instead, it frequently maintains a small positive slope, perhaps because minor rearrangements iron out weak points in the packing and steadily increase the aggregate stiffness E . Equivalently, the ‘easy’ plastic sliding could become exhausted, raising b , the stress scale in Equation 5.1 dictating when plastic sliding becomes significant. We introduce an additional fitting degree of freedom for the late-strain modulus,

$$q(\epsilon) = E_1\epsilon - b \log\left(1 + \frac{a}{\dot{\epsilon}}(e^{E_2\epsilon/b} - 1)\right) \quad (5.4)$$

The fits of this form to plasticity data can be seen in the top half of Figure 5.2. In the bottom half of the figure we show that an exponential recovery, as used in the synthetic data of Section 4.3, also works well (with the additional linear term included):

$$q(\epsilon) = E_1\epsilon + E_2(1 - e^{-\epsilon/\epsilon'}) \quad (5.5)$$

In this form, the fitting parameters are E_1 , E_2 , and ϵ' – an early and late strain modulus, and a strain scale for the transition between the two. This is all that seems to be needed to fit the recharge data, though one additional fitting parameter is needed for the fit based on Equation 5.1. The tradeoff for the extra fitting parameter is the ability to directly connect the parameters of the logarithmic stress relaxation to those of the stress recharge (something which is beyond the scope of this work, however).

5.3 Distribution of relaxation times and the Amir-Oreg-Imry model

The widespread occurrence of long time scale relaxation – deemed *anomalous* if different from simple exponential relaxation – has led to many theories and models that strive for simplicity and abstraction in order to account for the ubiquity. One of the more popular splits a system into many uncoupled modes that relax exponentially (i.e., each with a single timescale) in parallel. This idea has been applied to dielectric relaxation as a distribution

of relaxation times (DRT) theory [103, 104] and to the mechanical response of viscoelastic materials in the form of the generalized Maxwell model [105].

The modes each have their own characteristic rate λ and a share Γ of the measured macroscopic quantity, $f(t)$. Without driving, all modes relax to zero according to

$$d\Gamma = -\Gamma\lambda dt \quad (5.6)$$

The measured quantity $f(t)$ is the sum total of the elements, which we will take as an integral over a continuous distribution $P(\lambda)$ of rates

$$f(t) = \int P(\lambda)\Gamma(\lambda, t)d\lambda = \int P(\lambda)\Gamma(\lambda, t=0)e^{-\lambda t}d\lambda \quad (5.7)$$

Presumably any functional form can be fit with the proper distribution of modes and their amplitudes [106], though solving for the distribution given the signal $f(t)$ is an ill-posed problem [107]. The breadth of possibilities deviates from the simplicity just mentioned, so the authors Amir, Oreg, and Imry motivated specific forms for $P(\lambda)$ and $\Gamma(\lambda, t=0)$ that yield logarithmic relaxation [108, 109].

In their extension to DRT (which we will call AOI DRT), the distribution of rates is taken to be $P(\lambda) \sim 1/\lambda$ from some slowest rate in the system λ_{\min} to the fastest λ_{\max} . Such a distribution of relaxation timescales could arise, they argue, in thermal processes with a uniform distribution of energy barriers or multiplicative random processes [108], or more generally from random matrix theory applied to physically relevant interaction matrices [110]. Perhaps surprisingly, we are not required to know or specify λ_{\min} and λ_{\max} , only that they exist and are outside the timescales of the experiment.

A perturbation to the system is assumed to affect the modes uniformly, i.e., $\Gamma(\lambda, t=0) = \Delta$. Under these assumptions, the measured signal will decay as the difference of exponential

integrals of the first order, defined as [111]

$$E_1(x) = \int_x^\infty \frac{e^{-t}}{t} dt \quad (5.8)$$

We can then write the measured signal, with C a constant absorbing both the normalization of $P(\lambda)$ and the magnitude of the excitation, as

$$f(t) = C \int_{\lambda_{\min}}^{\lambda_{\max}} \frac{e^{-\lambda t}}{\lambda} d\lambda = C [E_1(\lambda_{\min} t) - E_1(\lambda_{\max} t)] \quad (5.9)$$

The second term quickly vanishes when $t > \lambda_{\max}^{-1}$ [109], which we assume holds even at the start of the experiment. The first term can be simplified with the expansion of $E_1(x)$ [111], with γ_E the Euler-Mascheroni constant:

$$E_1(x) = -\gamma_E - \log(x) - \sum_{m=1}^{\infty} \frac{(-1)^m x^m}{m * m!} \approx -\gamma_E - \log(x) \text{ for } x \ll 1 \quad (5.10)$$

Thus, as long as the timescales of the experiment are safely between the fastest and slowest processes in the system, this model yields a logarithmic decay. On timescales near λ_{\max}^{-1} and λ_{\min}^{-1} the approximations begin to fail and the logarithmic decay transitions to constant values.

We posit an additional reason to favor a $1/\lambda$ distribution. A maximum entropy distribution over a variable is one into which, after accounting for all known constraints on the distribution, the least additional information is factored [112]. The normal distribution is the maximum entropy distribution where the variance is specified, and the Boltzmann distribution is that for a specified average energy. Originating with Shannon [113], information theory – which has at its core the idea of generalized entropy – has expanded its purview rapidly in recent years, especially into the physics of complex systems ([114, 115, 116], to name a few from the first half of this year). Risking oversimplification, the idea is Occam’s razor for statistical distributions: information tends to wash away in random processes making

distributions in nature as general as possible.

The rates λ exist in the half-space from 0 to ∞ and as such are scale parameters. We know upper and lower bounds on λ must exist – no processes in the system can be infinitely fast or infinitely slow – though we do not know their values. Thus, the distribution with the least additional information is uniform in $\log \lambda$ [117], which is precisely $P(\lambda) \sim 1/\lambda$. Additionally, there is no *a priori* reason to prefer casting the problem in terms of rates λ rather than timescales $\tau = 1/\lambda$, so the scenario should be independent of this choice. This distribution is identical in both rates and timescales, i.e. $P(\lambda) \sim 1/\lambda$ is equivalent to $P(\tau) \sim 1/\tau$ for $\tau = 1/\lambda$, and both are uniform in log space (since $\log \lambda = -\log \tau$, $P(\log \lambda)$ is just the mirror image of $P(\log \tau)$, and the mirroring changes nothing if each is uniform).

In the maximum entropy picture, the distribution of rates for the system is the one with no additional assumptions beyond it being bounded. It is the most general and therefore the simplest. That this distribution yields a logarithmic decay for uniform excitation of the modes makes it a promising explanation for the prevalence of this form of anomalous relaxation in physical systems.

5.4 Two-step relaxation

The relaxation of a disordered system is the means by which it responds to a perturbation. It is how a system adopts a new base state; for a granular packing the unstressed state originally has an axial length L , and at some point in the distant future the new unstressed state will be the one with axial length $L \times (1 - \epsilon_1)$.

What happens when we allow the system to adapt to one state for an amount of time t_w , and then drive it to a new state with axial length $L \times (1 - \epsilon_2)$? Experimentally, this means running a standard stress relaxation experiment for a time t_w and then changing the axial compression from ϵ_1 to ϵ_2 before holding and recording the force relaxation again.

The results of this experiment performed on a column of glass beads are shown in Figure 5.3. The first hold is shaded in red, during which the stress decayed as in Figure 5.1

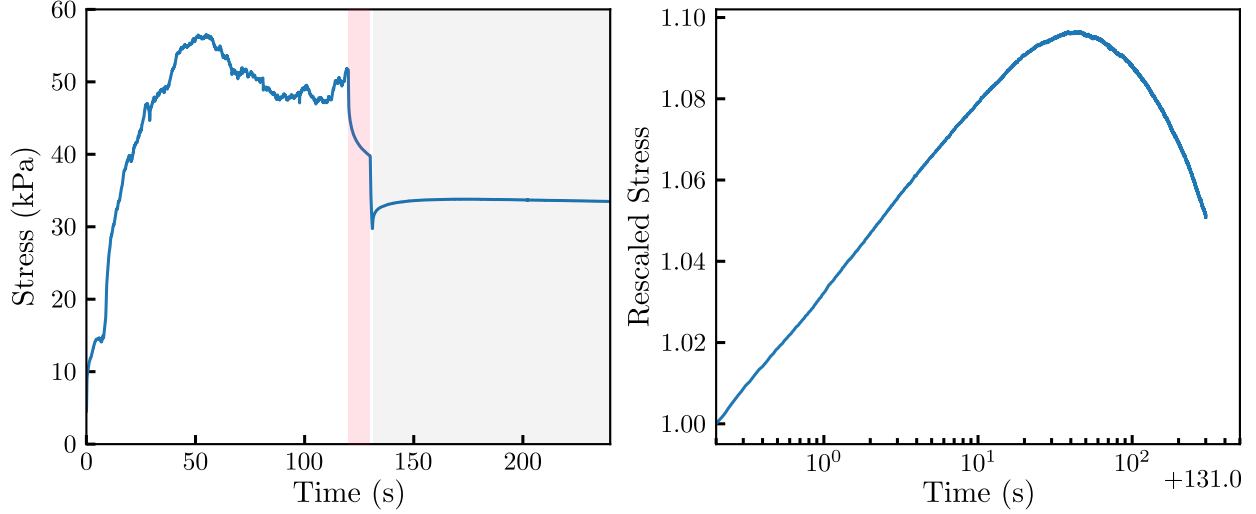


Figure 5.3: Two-step relaxation. *Left:* The test begins as in Figure 5.1 until $t_w =$ seconds into the first hold (shaded, red). The axial strain is then decreased by 1.7×10^{-4} in one second before holding at the new strain (shaded, gray). *Right:* The nonmonotonic stress evolution of the sample is more easily seen when displaying time logarithmically from the start of the second hold.

until it was cut short by a rapid decrease in stress when the Instron decompressed slightly. Remarkably, while the granular column was being held at a fixed strain during the final hold, the stress of the packing *increased* for almost a minute before turning around and decaying toward zero again as it did during the first hold. Nothing external cued the packing to turn its stress around 45 seconds in: the Instron was stationary. The only actor was the internal dynamics of the disordered system, adapting under applied stress and yielding complex behavior.

Two-stepped stress relaxation of crumpled sheets of mylar was studied in Lahini *et al.* [118], and cleverly explained using the AOI flavor of DRT introduced in the previous section. The same nonmonotonic relaxation has since been found in a block of rock salt [119] and in the frictional interface between two pieces of acrylic [102], though never (to our knowledge) in granular materials. Before approaching the explanation of Lahini *et al.*, an attempt to provide intuition is in order.

Very generally, we can think of the properties a system as being dependent on the timescale of probing, just as how in viscoelasticity the storage and loss moduli $G'(\omega)$ and

$G''(\omega)$ of a material depend on the driving frequency or timescale. In this perspective, the logarithmic relaxation of a disordered system showcases the adaptation of different aspects of the system to the new state, across a wide range of timescales. In the AOI DRT, a system has no preferred timescale and is split uniformly among the different relaxation rates (in $\log \lambda$ space).

If, after a time t_w the state is changed from ϵ_1 to ϵ_2 , then the system responds on timescales shorter than t_w as if it were already adjusted to ϵ_1 , while for timescales longer than t_w it responds as if it still preferred the initial (uncompressed) state. If the second strain is less than the first, $\epsilon_2 < \epsilon_1$, then on timescales shorter than t_w the system acts as if it were *stretched* in the process of reducing the axial strain, an odd deformation for a granular system to be able to support. The stress immediately after the decompression is the difference between the stress held by the long and short timescale aspects of the system: the slow part of the system is under compression from $\epsilon = 0$ to ϵ_2 , while the fast is under tension from ϵ_1 to ϵ_2 . The fast part relaxes the tension before the slow relaxes the compression, and the total stress decays upward on timescales shorter than t_w before turning around to relax downward on longer timescales.

Because the timescale for the exchange of the two relaxations depends on the waiting time t_w , the nonmonotonic evolution is evidence of memory in the system. The turnaround point in two-step relaxations, as characterized by the time of the peak t_p in the force, was found to scale with t_w in both our system (Figure 5.4) and in the crumpled mylar [118]. With only the stress relaxation data of the second hold, it is possible to infer the time the system was held in the state with axial strain ϵ_1 ; thus the system retains memory of its past.

The split nature of the granular packing has a broader basis in the physics of systems out of thermodynamic equilibrium. When driven to different states, partial equilibration happens at each and nonmonotonic dynamics can result. The Kovacs effect [120, 121] is where a polymer quenched to one temperature, then held at another, will exhibit nonmonotonic volumetric expansion. There are also many memory effects in disordered systems revealed in

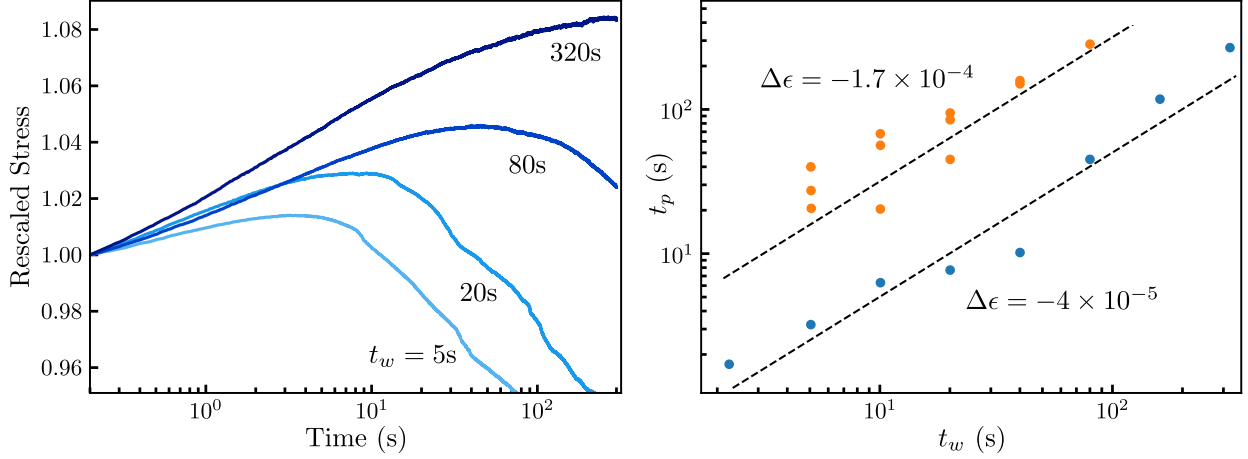


Figure 5.4: Memory in the turnaround of a two-step relaxation. *Left*: Two-step relaxations are shown for different t_w with the the rest of the loading protocol invariant. The shift to later turnaround times t_p demonstrates memory of the first hold was preserved in the packing at the start of the second hold. *Right*: The relationship between t_p and t_w are shown for two different values of the decompression strain $\Delta\epsilon$. The dashed lines are guides for a linear relationship, i.e., $t_p \propto t_w$.

some way by driving [33], which we will not touch on. The nonmonotonic stress relaxation of this chapter is closer to the Kovacs effect in that the memory of the past plays itself out without any need for continued driving; see Section 5.8 for potential applications in cryptography.

5.4.1 Explanation with AOI DRT

Nonmonotonic relaxation can be incorporated naturally into DRT by allowing driving and element amplitudes Γ to take negative values. When driving the system backwards for the second step, the amplitudes of elements of the system which had already relaxed dip below zero, such that at the start of the second hold the system contains counter-relaxing natures.

In the AOI DRT simulation of Figure 5.5, the abstract system composed of 10 relaxing modes is held in three states: up (blue), down (red), and middle (white). In the first part of the test (top panel), the up state is held long enough that it is adapted to, or learned, by all of the modes. The down state is then held for a shorter duration, such that only the faster

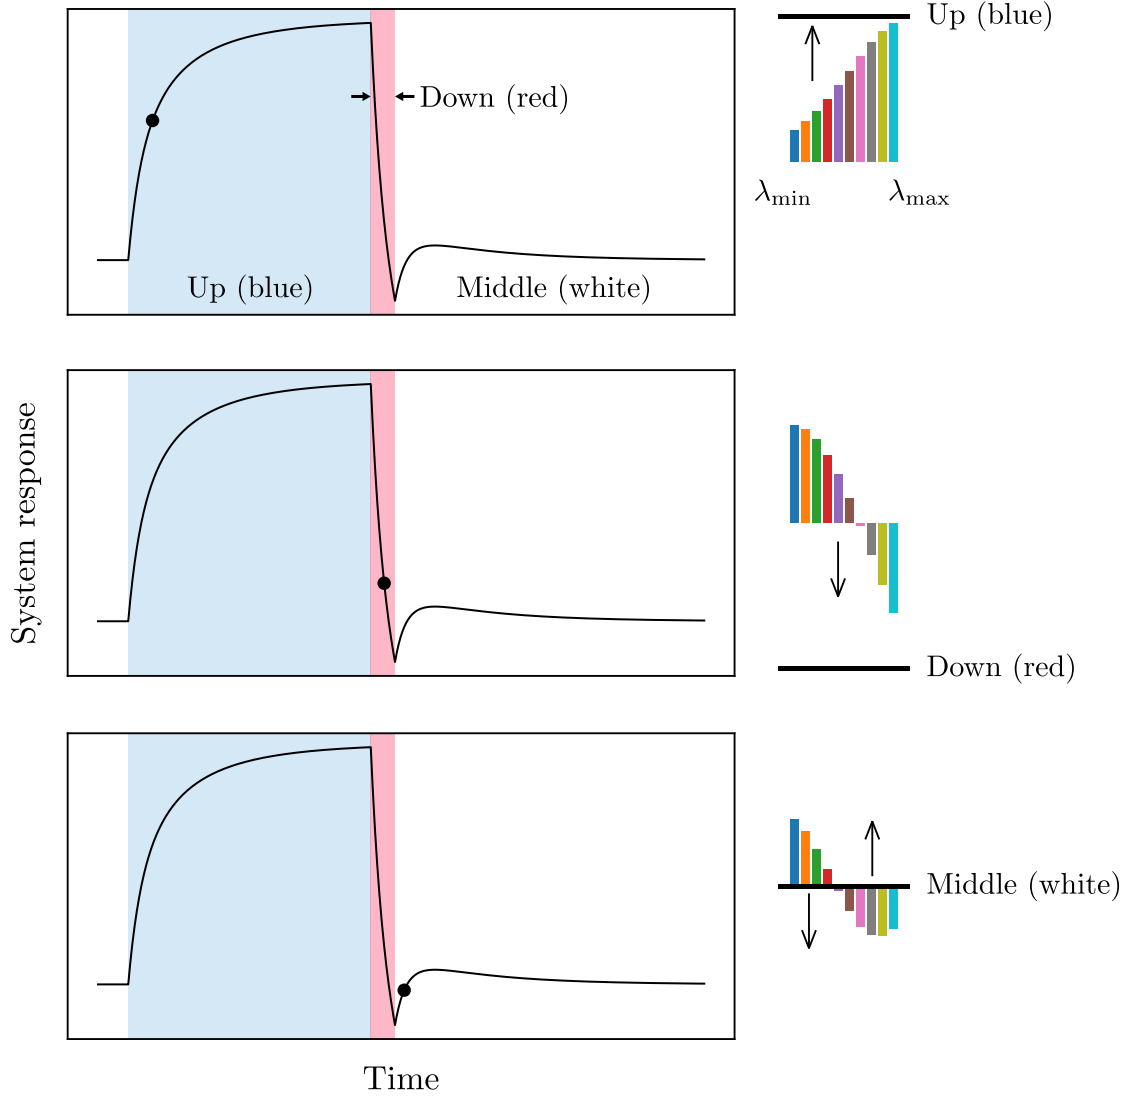


Figure 5.5: Snapshots from a simulated AOI DRT model exhibiting two-step relaxation. In this abstraction of a relaxing system, we hold the system in one of three states and track the evolution of the modes. The instant in time captured in each of the three snapshots is represented by the black dot in the system response on the left. The system is held in a state represented on the left by the shading color and on the right by the position of the black bar (blue shading is the up state, red is the down state, and white is the middle state). The different modes are shown as colored bars on the right, with slowest on the left and fastest on the right. They are distributed uniformly in $\log \lambda$ as in the AOI DRT. A two-step relaxation can be seen in the system response on the left, and the conflicting natures in the distribution of amplitudes on the right. The faster modes learn the down (red) state before the slower modes have time to forget the up (blue) state. The memory then plays itself out in reverse order, with the fast modes relaxing upwards before the slow ones relax downwards.

modes have time to adapt (center panel). At the beginning of the final hold in the middle state, the system is torn between memories of two states, with the slow half preferring the up state and the fast half preferring down. They both come to the middle eventually, but according to their own timescales, causing nonmonotonic relaxation.

Mathematically, we represent the state of the system with $\Gamma(\lambda, t)$ again. At the end of the first hold in a state with excitation magnitude Δ_1 , the system has aged by t_w .

$$\Gamma_{1h}(\lambda) = \Delta_1 e^{-\lambda t_w} \quad (5.11)$$

A decompression of magnitude Δ_2 is then applied, at which point we redefine $t = 0$, and then everything relaxes again. The state of the system during the second hold is given by

$$\Gamma(\lambda, t) = \left[(\Delta_1 e^{-\lambda t_w}) - \Delta_2 \right] e^{-\lambda t} \quad (5.12)$$

The measured signal is again the integral across the rates λ

$$\begin{aligned} f(t) &= C \int_{\lambda_{\min}}^{\lambda_{\max}} \Delta_1 \frac{e^{-\lambda(t+t_w)}}{\lambda} - \Delta_2 \frac{e^{-\lambda t}}{\lambda} d\lambda \\ &\approx C \left(\Delta_1 E_1(\lambda_{\min}(t+t_w)) - \Delta_2 E_1(\lambda_{\min} t) \right) \\ &\approx C \left(\Delta_2 \log(t) - \Delta_1 \log(t+t_w) \right) + \text{const} \end{aligned} \quad (5.13)$$

Thus the relaxation is the difference of two logs, one decreasing that initiated t_w in the past (the original logarithmic decay of the first hold), and one increasing that just started.

5.5 N-step relaxation

In the view that the system has many natures, spread out across a broad range of timescales, what should limit the number of memories held by the system? What loading protocol will allow us to observe N-step relaxation, with $N > 2$?

We turn to the simulated AOI DRT system for guidance. It is clear from observing

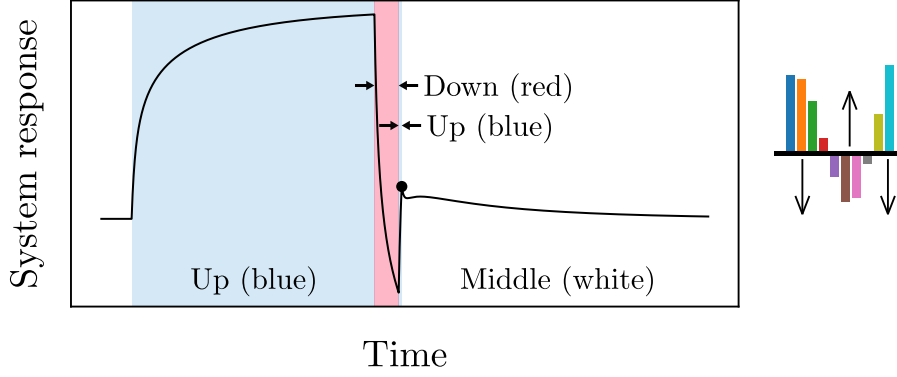


Figure 5.6: Three-step relaxation in a simulated AOI DRT system. The system is briefly held in the up (blue) state just prior to the final hold, which trains the fastest modes with an additional memory and leads to a third step in the relaxation.

the two-step dynamics of Figure 5.5 that the long and short holds had to be imprinted in that order. If the short hold had been first, a subsequent long hold would have erased it. Thus we can see that additional memories must be nestled inside one another, starting with the longest hold and working to the shortest. We can also see that the upper limit on the number of memories is the number of unique timescales in the system. In the simulated system of Figure 5.5, that number is ten, whereas in a disordered system with myriad degrees of freedom, the number would be astronomical.

Capitalizing on this intuition, we create a third step in the simulated system by introducing an extremely short hold in the up state immediately prior to the final hold. Figure 5.6 shows the system response over the entire simulation and a snapshot of the internal state Γ at the start of the final hold. The three memories are evident in the counter-relaxing components: the slowest third of the modes still remember the long first hold in the up (blue) state. The rest adapted to the down (red) state during the second hold, and then only the fastest third had enough time to learn the up (blue) state in the third hold. As the system relaxes, the memories are played out in reverse order: the fastest modes are the first to act and reveal the most recent held state, whereas the slowest ones are the last to share their story, which is of the first held state.

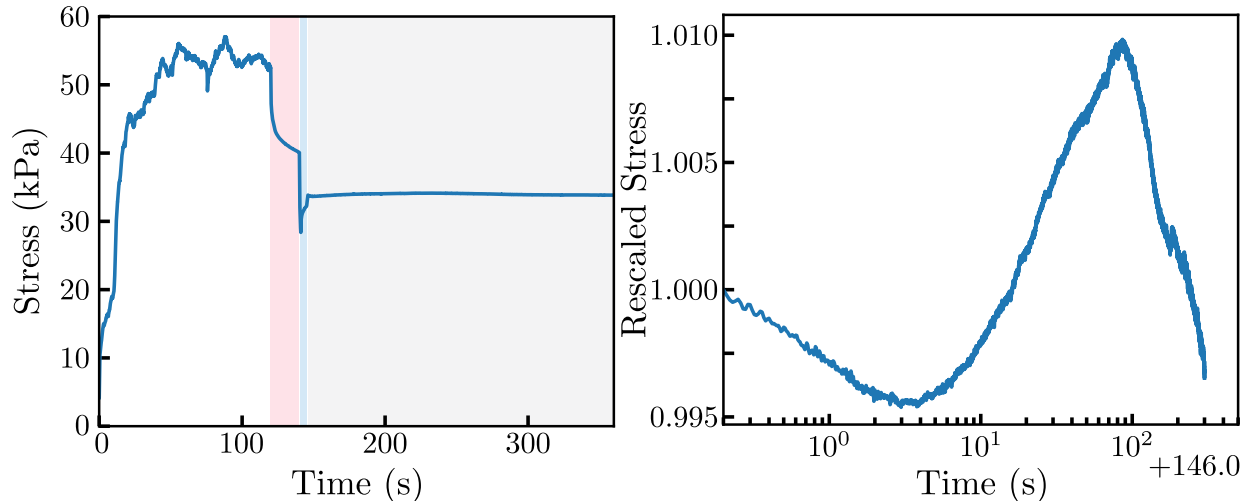


Figure 5.7: Three-step relaxation in a packing of glass spheres. *Left:* We amend the protocol for the two-step relaxation experiments by adding a small compression step partway into the second hold. In this particular experiment, the first hold time (shaded, red) was $t_{w,1} = 20\text{s}$ and the second (shaded, blue) was $t_{w,2} = 4\text{s}$. The first decompression step was $\Delta\epsilon_1 = -2.7 \times 10^{-4}$ over 1s and the second was $\Delta\epsilon_2 = +1.7 \times 10^{-5}$ over 1s. *Right:* The three steps of relaxation with time displayed logarithmically from the start of the final hold (rescaling the stress as before).

5.5.1 Three-step stress relaxation

Figure 5.7 shows three-step stress relaxation, the first case of it being observed in any system, to the best of our knowledge. When displayed with the stress data from the previous steps of driving and holding, the three-step is nearly imperceptible (Figure 5.7, left). Upon magnification, however, the relaxations down, up and finally down again are clearly visible.

In the first 4 seconds of the final hold, the stress of the granular packing decayed toward zero. It then turned around and increased for the next 80 seconds, before changing directions one final time, all the while without any external influence. At the start of the final hold ($t = 146$ seconds in the left plot, $t = 0$ seconds in the right plot of Figure 5.7), the granular system contained competing natures which played themselves out to yield the complicated dance of stress relaxation.

As with the two-step relaxation, we can establish the presence of memory by varying the duration of the holds in the loading history. In Figure 5.8 we keep the same protocol as in

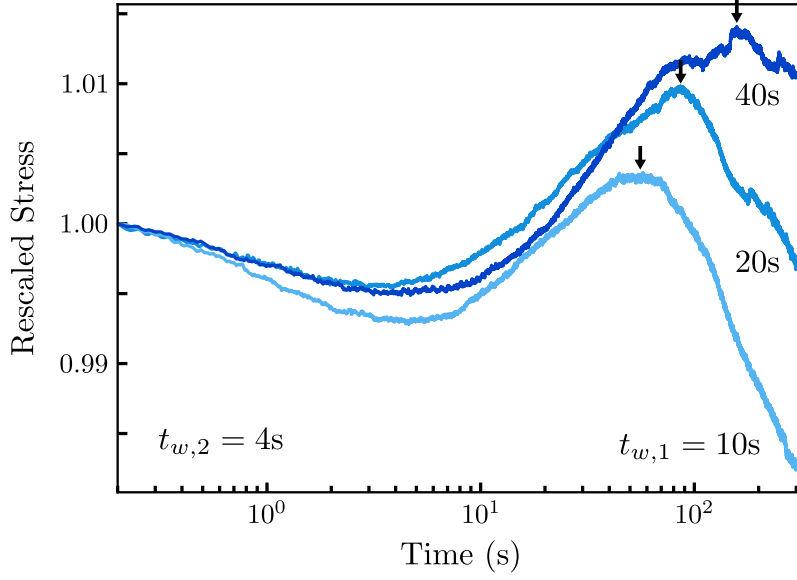


Figure 5.8: Memories in the turnaround points of a three-step relaxation. By varying the first hold time $t_{w,1}$ for the protocol in Figure 5.7, we are able to shift the second peak time. However, due to the variability from run to run (shown in the time of the first valley), this dependence needs many more trials to be run to establish a clear relationship.

Figure 5.7 except for $t_{w,1}$, which is lengthened and shortened by a factor of two. The second turnaround time shifts in step, though the dependence is less clean than with the two-step memories. The variability in three step memories is high between trials, arising in part from the smaller force signal.

The memories are played out in the reverse order to how they were stored, as in the simulated system of Figure 5.6. The first hold plays out last, and the second hold comes first.

5.5.2 *Experimental details that make more relaxation steps difficult*

We now have a formula for N-step relaxation: hold the system in different states for broadly separated timescales, working from longest to shortest. To be able to stash more memories in the system, we should therefore use the largest range of timescales available.

On the short side, we cannot change state faster than the Instron can move, which sets our lower limit at about 1s. The longest experimental stress relaxation data shown thus far

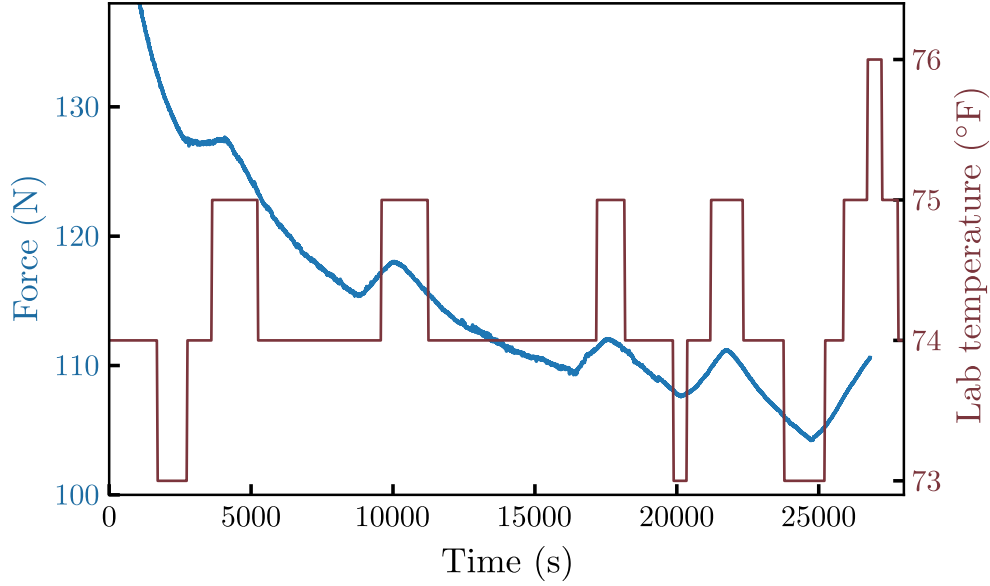


Figure 5.9: Lab temperature swings of 1-2°F limit the duration of relaxation experiments. During a long single step stress relaxation experiment (approx. 8 hours) on a packing of 3D-printed ellipsoid particles, the usual logarithmic decay of force (blue) featured erratic undulations which coincided with small increases in the temperature in the room (red).

has been capped at ten minutes (Figure 5.1). Considering an experiment on the plastic flow of pitch has been continuing for nearly a century [122], we seem to have plenty of room to grow in the long direction. Not so, as it turns out.

Figure 5.9 shows force data from an eight-hour stress relaxation experiment, along with low-resolution data of the temperature in the lab. The heating in the lab turns on every 30-90 minutes, raising the temperature by 1-2°F and causing thermal expansion of everything in the room. The expansion is different for all the various materials in the load frame, the load cell, and the packing, and this differential expansion is enough to cause force changes of a few Newtons. As measured in Section 2.1, changes of a few Newtons for the load cell is equivalent to around a micron of extension, so the differential expansion can be quite small to leave traces in the data. Even if the force undulations were removed from the data, the expansion-contraction cycles of the grains can have a significant effect on the aggregate properties [123, 124], complicating the phenomenon we endeavour to study.

Stress relaxation experimenters in Hernan Makse’s group (CCNY) ran into this difficulty

as well [96] and their resolution matches ours: stick to experiments much shorter than the thermal cycles. Such a resolution is not entirely bad, as it forces the iteration time for data to be shorter.

In terms of precision in the force measurement, by partitioning the system into more counter-relaxing components, we dramatically reduce the signal with each additional partition. In the raw data of Figures 5.1, 5.3, and 5.7 that the measured signal drops by roughly an order of magnitude with each additional step. In Figure 5.7, the noise floor of the Instron is visible in the signal, and the total variation across five minutes is less than 1N (0.5kPa).

For our setup, the combination of temporal limitations and signal resolution make additional steps of relaxation (i.e., beyond three) exceedingly difficult.

5.6 Outgrowing the model

The AOI variant of DRT has taken us far. It granted insight about how to view the counter-relaxing aspects of the system, intuition which was critical to finding 3-step relaxation. It also supplied functional forms for representing the internal state of the system and for fitting N-step experimental data.

After pushing the model to new territory, however, we have encountered issues which require extensions to the AOI DRT, or perhaps an entirely different way to view the problem of relaxation in disordered systems.

5.6.1 *Finite duration loading*

In Figure 5.10, we show two-step relaxation for which there was no first hold (i.e., compression to 10% and then immediate decompression, with $t_w = 0$). In AOI DRT, perturbations are applied instantaneously as a uniform kick to all modes. In this case, $t_w = 0$ would mean exciting all modes by Δ_1 and then immediately back-tracking by Δ_2 , such that the amplitudes are still all uniform at the start of the hold. These dynamics are just the monotonic

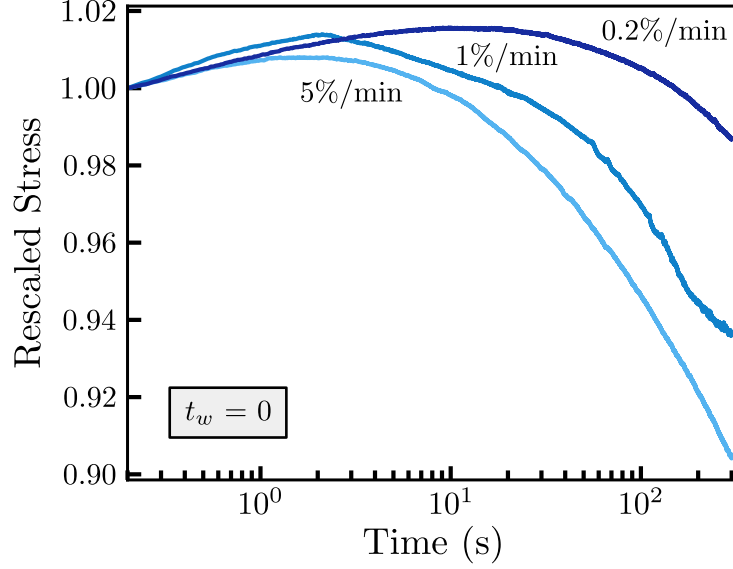


Figure 5.10: Two-step relaxation with $t_w = 0$. In these experiments, the granular packing was compressed to 10% strain at different strain rates, before compression was immediately reduced by $\Delta\epsilon = 8 \times 10^{-5}$ in one second. The two-step relaxation occurs even though $t_w = 0$, something which cannot take place in AOI DRT without accounting for finite driving timescales.

logarithmic relaxation of Section 5.3: there can be no turnaround points.

Instead, we do find two step-relaxation, and that the strain rate of compression seems to insert itself into the state of the system. Slower compression results in a turnaround point later in time.

An attempt to incorporate finite driving was suggested in [109] for conductance relaxation in porous silicon, though in actuality the extension was for a finite duration hold in one state before switching to another. Thus it preceded the two-step analysis of [118] and modified the slowest modes rather than the fastest, as in the case treated here.

We can incorporate finite driving timescales into the AOI DRT by decomposing an instantaneous excitation Δ into pieces distributed in time. We posit that each incremental excitation is proportional to the increment in strain, and modify Equation 5.6 to get

$$d\Gamma = -\Gamma\lambda dt + \alpha d\epsilon = (-\Gamma\lambda + \alpha\dot{\epsilon})dt \quad (5.14)$$

The strain rate $\dot{\epsilon}$ can be positive or negative to reflect driving in either direction, and α is a constant relating the global driving to the effect on each element. During a hold ($\dot{\epsilon} = 0$), the amplitude decays exponentially as before:

$$\Gamma_{\text{hold}}(\lambda, t) = \Gamma_0(\lambda) \frac{e^{-\lambda t}}{\lambda} \quad (5.15)$$

During driving, however, the finite duration allows partial relaxation to occur in the process:

$$\Gamma_{\text{driving}}(\lambda, t) = \Gamma_0(\lambda) e^{-\lambda t} + \frac{\alpha \dot{\epsilon}}{\lambda} (1 - e^{-\lambda t}) \quad (5.16)$$

In both cases, $\Gamma_0(\lambda)$ is the state of the system at the end of the previous step, and t is reset to zero at the beginning of every new step. With Equations 5.15 and 5.16, any sequence of compression steps and holds can be composed.

In the limit of fast driving ($\lambda t \ll 1 \forall \lambda$), we recover the AOI DRT. Equation 5.16 reduces to

$$\Gamma_{\text{driving}}(\lambda, t) \approx \Gamma(\lambda, t = 0) + \alpha \dot{\epsilon} t = \Gamma(\lambda, t = 0) + \alpha \Delta \epsilon \quad (5.17)$$

In this limit each driving step affects all elements with a constant offset proportional to the change in strain, as before. If, instead, the driving timescale falls somewhere in the middle of the range of relaxation timescales, slower modes will be excited uniformly and faster modes will approach an amplitude that scales with $1/\lambda$. The $1/\lambda$ amplitude is the value where the excitation due to driving matches the decay due to relaxation.

For the two-step relaxation with $t_w = 0$ (Figure 5.10), we start the system in the state $\Gamma(\lambda, t = 0) = 0$ and then propagate Γ_{driving} forward in time to find the state of the system at the end of the compression step:

$$\Gamma_{d1}(\lambda) = \frac{\alpha \dot{\epsilon}}{\lambda} (1 - e^{-\lambda t_1}) \quad (5.18)$$

We then drive the in the opposite direction with the same strain rate for time t_2 , such

that the state of the system at the end of the decompression step, and the start of the hold, is

$$\Gamma_{d2}(\lambda) = \frac{\alpha\dot{\epsilon}}{\lambda}(2e^{-\lambda t_2} - e^{-\lambda(t_1+t_2)} - 1) \quad (5.19)$$

The signal $f(t)$ for the relaxation is then given by

$$f(t) = \int_{\lambda_{\min}}^{\lambda_{\max}} P(\lambda) \Gamma_{d2}(\lambda) e^{-\lambda t} d\lambda = C \int_{\lambda_{\min}}^{\lambda_{\max}} 2 \frac{e^{-\lambda(t+t_2)}}{\lambda^2} - \frac{e^{-\lambda t}}{\lambda^2} - \frac{e^{-\lambda(t+t_1+t_2)}}{\lambda^2} d\lambda \quad (5.20)$$

in which C has absorbed all constants. $f(t)$ is the difference of exponential integrals of the second order, which we can sidestep when solving for the turnaround time. The time derivative introduces a factor of λ in each term and turns the integrals into exponential integrals of the first order, as in Equation 5.9. We have

$$\begin{aligned} \frac{df}{dt} &\approx C [E_1(\lambda_{\min} t) + E_1(\lambda_{\min}(t + t_1 + t_2)) - 2E_1(\lambda_{\min}(t + t_2))] \\ &\approx 2\log(t + t_2) - \log(t) - \log(t + t_1 + t_2) \end{aligned} \quad (5.21)$$

This is zero at the turnaround time t_p ,

$$t_p = \frac{t_2^2}{t_1 - t_2} = \frac{(\Delta\epsilon)^2}{\epsilon - \Delta\epsilon} \left(\frac{1}{\dot{\epsilon}} \right) \quad (5.22)$$

In this extension, the turnaround point grows linearly with the inverse strain rate when leaving the final strain ϵ and the decompression strain $\Delta\epsilon$ to be constant. In most experiments, the prefactor on this timescale will be very small, suggesting why, in the slowest compression of Figure 5.10, $1/\dot{\epsilon} = 30,000$ seconds yet $t_p \approx 20$ seconds.

The only additional assumption in this extension was that driving can occur in small steps, synchronous with relaxation of the modes. In the limit where the driving timescale is negligible, AOI DRT is recovered. Otherwise, we find that the driving timescale does indeed insert itself as another timescale in the system.

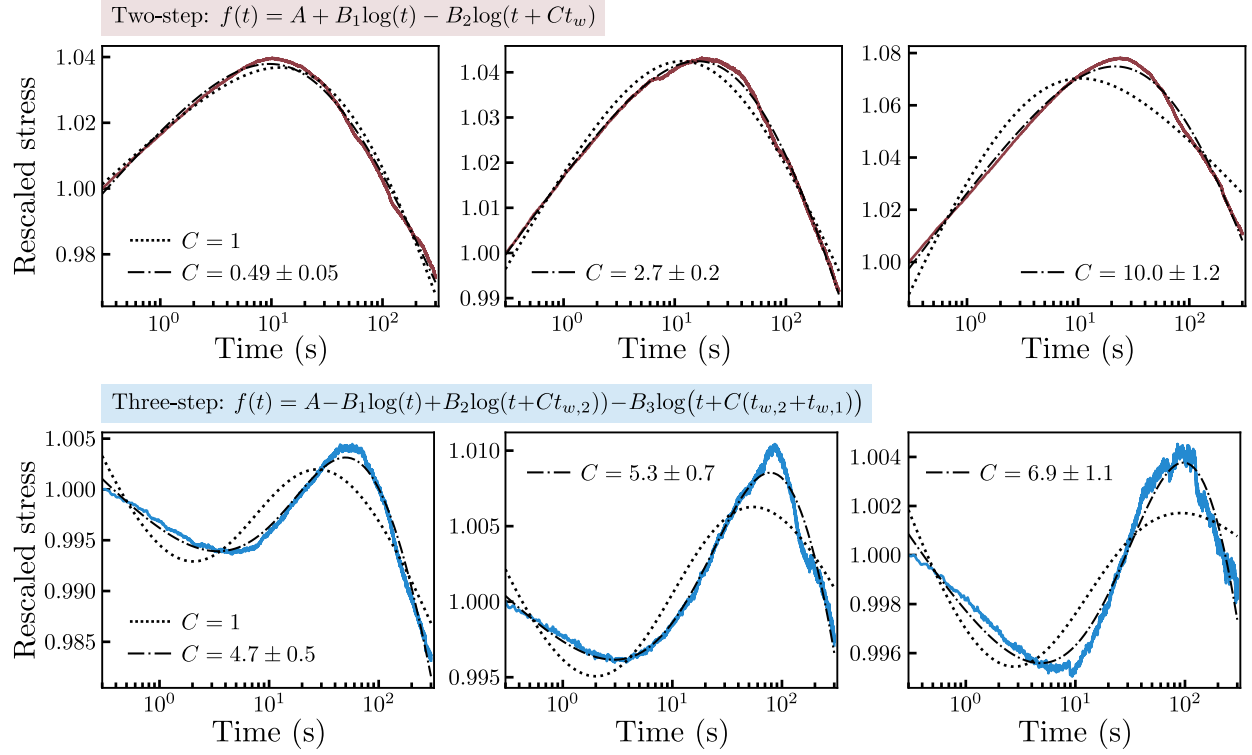


Figure 5.11: Fitting N-step relaxation with a series of logs. When fitting two-step (top) and three-step (bottom) relaxations to a series of alternating logs, as predicted by AOI DRT, a nontrivial constant C is needed to scale the duration of prior aging. The dotted curve in each plot is the best fit to the data with $C = 1$, while the dash-dotted line includes C as an additional fitting parameter.

5.6.2 The extra fitting parameter for N-step relaxation: Expedited aging

To fit two-step relaxation data, Lahini *et al.* [118] use the following form:

$$f(t) = A + B_1 \log(t) - B_2 \log(t + Ct_w) \quad (5.23)$$

In the derivation of the two step relaxation, however, the relaxing signal (Equation 5.13) has no free parameter C to scale the duration of the wait time. It turns out that a value of $C \neq 1$ is necessary to fit the data, both for the crumpled mylar [118] and for our glass beads (Figure 5.11). C was found to be approximately 2.5 across the board in the mylar experiments [118], but in fits to our data on glass beads its value varies.

Lahini *et al.* emphasized the importance of initial conditions for their experiments, and

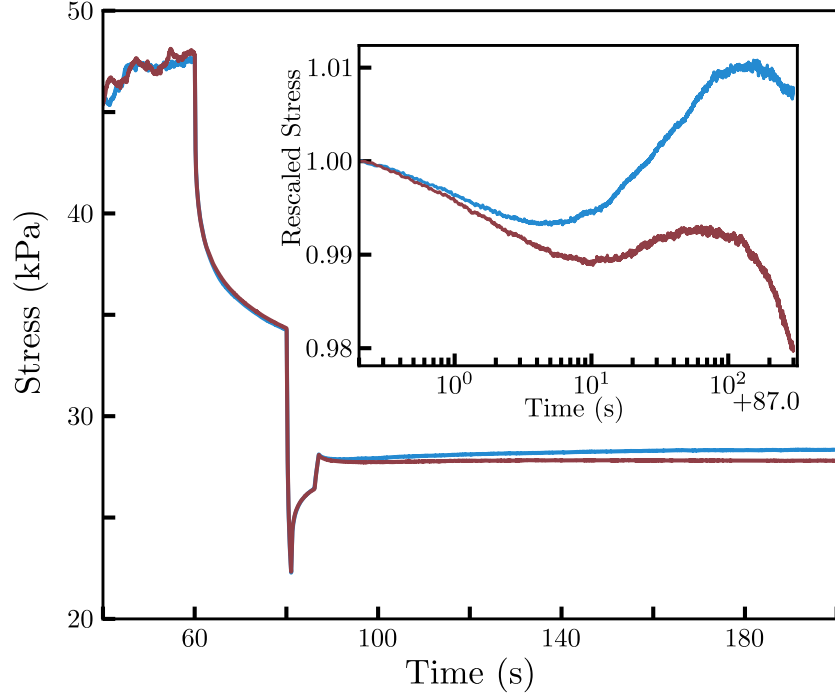


Figure 5.12: Identical start, different finish in two experiments. In these two runs with the same protocol (on a different set of glass spheres – 3mm diameter, unknown supplier), everything until the start of the final hold matches almost perfectly. The three step then quickly differentiates the two runs.

had a special protocol for crumpling the mylar sheets before each compression in order to have reproducible relaxation data [118]. Perhaps the value of C depends sensitively on the initial state of a disordered system, and differences from trial to trial in the poured configurations of the granular packing were enough to cause the spread in values we observe. Indeed, in Figure 5.12, two runs almost perfectly tracked one another in $q(\epsilon)$ leading up to the three step, but then split in the final hold. This suggests a difference in their internal states that laid dormant until the three-step relaxation, and also poses a difficulty for reproducibility.

Investigating the sensitivity of N-step relaxation to initial conditions certainly merits future investigation, since the value of C has important implications for the nature of the underlying relaxation. Here, however, we focus solely on the fact, and its implications, that C does not equal 1.

The two-step relaxation fits in the top row of Figure 5.11 are still reasonable even when

C is forced to have a value of 1 (dotted line). We find that the three-step relaxation data, however, serve as a better test for the fitting form, as shown by the difference between the dotted ($C = 1$) and dash-dotted fits (C free to vary). Among the handful of three-step relaxations taken so far, C varies less than for the two-step, though it is too early to tell if this is just a feature of the small sample size.

Interestingly, $C > 1$ indicates the relaxation acts as if the aging in each prior step lasted longer than it actually did. A logarithm with a time offset $\log(t + T)$ is equivalent to a fresh $\log(t)$ that started T in the past. This means the fitting form Equation 5.13 is simply the superposition of logarithmic decays that start at the beginning of each new hold, and continue into the future. Rescaling the time offset as in Equation 5.23 has the same effect as retroactively shifting the start time of the previous hold. This can be seen in almost all of the dotted $C = 1$ fits, which turn around earlier than the experimental data. The exception is the leftmost two-step relaxation, where the best-fit value of C is actually less than one – an anomaly in our limited dataset.

$C \neq 1$ cannot arise from a simple modification to the AOI DRT. If we allow a different distribution for $P(\lambda)$ and nonuniform excitation $\Delta(\lambda)$ (the two key aspects of the AOI model), both can be subsumed into a single function $g(\lambda)$. We only require that the decompression step affects modes with different λ in the same way as the compression step (i.e., $\Delta(\lambda)$ is only scaled by the magnitude and direction of each step, but does not change form). The measured signal for a two-step will then be

$$f(t) = \int_{\lambda_{\min}}^{\lambda_{\max}} g(\lambda)(e^{-\lambda(t+t_w)} - ce^{-\lambda t})d\lambda = G(t + t_w) - cG(t) \quad (5.24)$$

with the function $G(t)$ given by

$$G(t) = \int_{\lambda_{\min}}^{\lambda_{\max}} g(\lambda)e^{-\lambda t}d\lambda \quad (5.25)$$

and c the ratio of the magnitude of the decompression and compression steps. Importantly,

there is no way for the system to learn a multiple of the wait time t_w , and it arises as a consequence of the complete independence of the modes.

One resolution could be a sort of diffusion between the modes, as discussed in [125] for the α relaxation of glasses. Beyond modifying AOI DRT, sequential relaxation (e.g., hierarchically constrained dynamics [126, 127, 128]) naturally features interaction between different timescales of relaxation. There are numerous routes to explore, with the experimental finding that $C \neq 1$ as a new discriminant for theories about relaxation in disordered systems.

5.7 Complex inner state

The nonmonotonic evolution of the stress supported by the packing indicates a complex inner state that retains an imprint of past driving. Because the value of stress can pass the same point multiple times, there must be extra “hidden” variables to describe the state of the system.

To better understand why, consider the system during a two-step relaxation, at the two moments in time t_1 and t_2 where it supports a stress σ_0 . At both times, the stress and strain are identical, so there must be another variable for which we are not accounting which tells the packing at t_1 to relax upwards in stress and the opposite for t_2 . With a third step, we require at least two additional variables. There are no indications that a granular system is limited in the memories it can store, so there are presumably a multitude of additional variables describing the inner state which do not need to be taken into account for general compression experiments.

Herein lies the value of the relaxation experiments of this chapter: through them we are granted a peek into the evolution of the complex inner state of a granular packing under stress. Upon further reflection, the surprising aspect to this story is not that the aggregate behavior of thousands of athermal particles has complex dynamics capable of storing past history. It is instead surprising that it is ever possible to describe this behavior by constitutive

relationships with one or two variables.

How does the logarithmic relaxation of a single step play into the necessity for a complex inner state? Logarithmic evolution can result from a memoryless process,

$$\frac{dq}{dt} = -ae^{q/b} \quad (5.26)$$

This form (the same we used to fit the stress recharge data in Section 5.2.1) was used to explain logarithmic relaxation of the height of a crumpled sheet of paper under a weight [100]. It was rooted in the idea of a thermally activated process overcoming barriers that grew higher with time. In this relationship, there is no sense of past: the system holds no memory in the same way as a simple Debye relaxation, and it could not lead to two-step relaxation.

The ability to exhibit N-step relaxation requires some sort of inner state. The relaxing rates model, discussed in Section 5.3, stores the state in a distribution of clocks that can record only one timescale each. Models of sequential relaxation, like hierarchically constrained dynamics [126, 127, 128], store the state in the progress made on tasks that continue to increase in difficulty. Models without a complex inner state, such as the one proposed in [129] to explain logarithmic relaxation and based on an ensemble of clocks the system hops between (like a hitchhiker passing between towns, the authors of [129] write) thus need modification to explain multiple steps of relaxation.

5.7.1 Particle sliding as a mechanism for relaxation

The stress relaxation of a granular packing signals its adaptation to a new form. To reconfigure, the packing can do so through grain rearrangements, as in the plasticity experiments of Chapter 4. The material of the grains themselves can also deform, causing shifts in the contact forces as the material slowly flows into a new shape [96]. Both components of deformation will be present in granular relaxation, and we will assume that their contributions

sum to the total relaxation measured.

It should be noted that the distinction between the two cases is a bit blurry: interfacial sliding between particles is also the result of material relaxation at the scale of contacting asperities [93]. However, in the particle-particle sliding case, the material relaxation can result in particle rearrangements that take advantage of the extra configurational degrees of freedom of a granular packing, as compared to a solid block of material.

In the left hand side of Figure 5.13, we show single step relaxation for four different samples. The absolute stress has been scaled out for all, allowing us to compare the fractional decay of the viscoelastic relaxation modulus [105]. Two samples are granular packings of the 5mm glass beads from the rest of this chapter, though one set has been washed in piranha solution to make the particle surfaces more frictional. The relaxation is strikingly different in magnitude. Friction between the particles changed between the two samples, but not the bulk properties of the glass. Because the relaxation decreased significantly when the relative sliding between the particles was inhibited, we infer that the primary component of glass bead relaxation is due to relative particle motion.

The other two samples are made of 3D-printed plastic, with one a packing of ellipsoid grains and the other a solid block. The solid block is taken to represent the limit where 3D-printed grains in a packing are so frictional as to preclude particle rearrangement. That the two relax similarly in magnitude suggests that in a granular packing of 3D-printed particles, the dominant player in stress relaxation is the bulk plastic viscoelasticity.

In the right side of Figure 5.13 we show evidence of the distinctly granular nature of stress relaxation. The events numbered 1-5 occurred during the final hold of various stress relaxation experiments, meaning there was no external driving to have sparked the deviations in the relaxation behavior. Instead, the internal rearrangement of the particles led to multiple spurts of faster relaxation more than a minute into the hold (Event 1), a gradual increase – over several seconds – in the rate of relaxation that culminated in a sudden drop (Event 2), and even a sudden *increase* in stress (Event 3). Event 4 resembles a stress drop and recharge

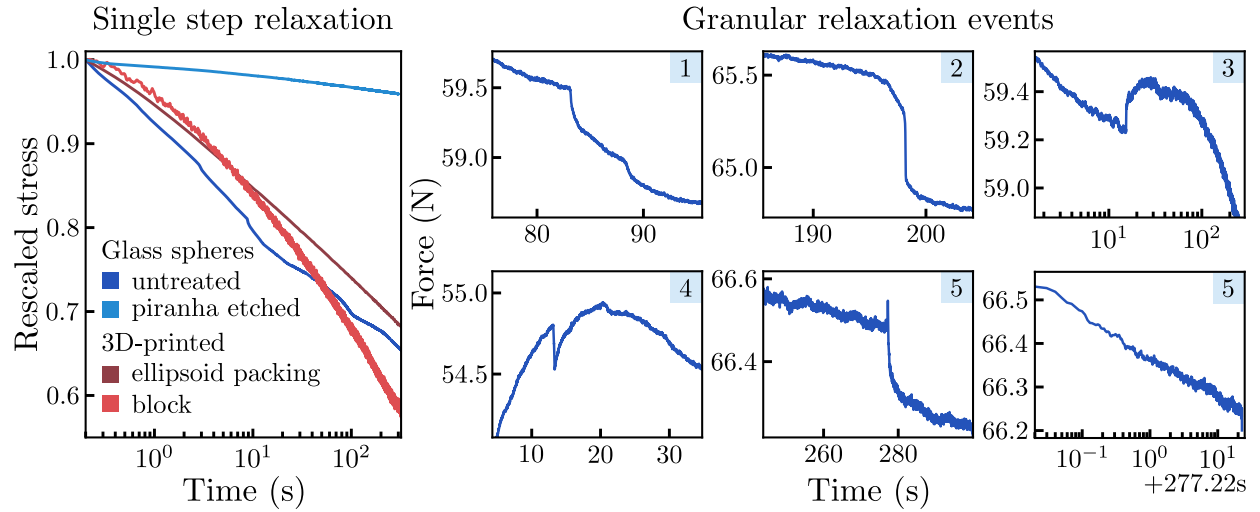


Figure 5.13: The granular nature of the relaxation. *Left:* Single step relaxation for two packings of glass spheres and two samples of 3D-printed Verowhite plastic. Both sets of glass spheres were the 5mm diameter soda lime glass spheres from the rest of this chapter, but one set was washed in piranha solution making the particle surfaces more frictional. The two 3D-printed samples were a packing of ellipsoids and a solid block of plastic.

Right: Events that occurred during the final hold in various relaxation experiments (all with the 5mm glass spheres from the rest of the chapter). For each, the displacement during the event was constant, meaning the internal dynamics were all that could have prompted the event. Event 5 is shown twice, first with a linear time axis and second with time displayed logarithmically from the start of the event, indicating that the event initiated a new logarithmic decay in force.

from Chapter 4, yet there was no external driving to fuel the recharge. Finally, Event 5 is shown twice, the first with time displayed linearly and the second with time growing logarithmically from the moment the event starts. We can see that a brand new logarithmic decay started with the event, suggesting a rearrangement in the packing distributed stress held in a slow mode between a wide range of faster modes (in the DRT framework). Additionally, the logarithmic decay was captured with nearly order of magnitude better temporal resolution because the Instron was not involved in initiating the relaxation.

The two halves of Figure 5.13 suggest the relaxation in the granular system of glass spheres happens primarily through particle-particle sliding rather than bulk material relaxation. The packing reconfigures at an ever-decreasing rate during stress relaxation experiments, though the reconfigurations can still lead to sudden changes in the stress just as they did in the X-rays of Figure 4.2. The plastic deformation does not end when compression ceases; the granular packing continues to adapt.

5.7.2 *Connections to plastic deformation*

Many models for plastic deformation naturally lead to a complex inner state of the granular system which would (1) relax over a broad range of timescales and (2) be filled with degrees of freedom capable of being imprinted with the driving history. STZ theory is a description of plasticity in terms of local configurational jumps blocked by a distribution of energy barriers, and an effective temperature to overcome these barriers as well as create new ones [84]. When driving halts, the system contains a broad distribution of timescales of relaxation and glassy dynamics arise naturally [130]. The different configurational jumps can be signed, and thus different directions of driving could be stored in the distribution of STZ magnitudes and barrier heights.

An interface depinning description of plasticity and relaxation is qualitatively similar though with additional abstraction. Driving the elastic interface through a field of random pinning forces leads to intermittent dynamics relevant to Chapter 4 [16, 89], and a self-affine

roughness of the interface that exists over a broad range of length scales[131]. Equilibration of the interface happens via a growing length scale, such that on scales longer than the equilibration length there is memory of the prior state [132]. Multistep relaxation could therefore result from nested memories, stored across different length scales in a sort of Russian doll picture. In a granular packing this might take the form of shear stress fluctuations existing across the many length scales of the system, which smooth first across several-particle neighborhoods before working up to large chunks of the entire packing.

Theories of amorphous plasticity and glassy relaxation are, individually, incomplete; consequently there is no comprehensive picture yet which successfully unites them. It seems clear on a qualitative level, however, that each is a different but intimately connected reflection of the multitude of metastable configurational states which result from the rampant disorder.

5.8 Self-erasing memory

A granular packing exhibiting multistep relaxation can be thought of as a device that was imprinted with a signal during driving which it then reads off as a one-time message. The signal could be stored in the time of each hold step, for instance, yielding several bits of information per step of driving. The precise amount of information would depend on how reproducible the output signal is, and the precision of the measurement device on the receiving side. Importantly, the time until the message reads itself off would be controllable by the sender. We have been limited in the lab to multistep relaxations that take less than half an hour to complete (Section 5.5.2), but the upper bound on the timescale is presumably much larger. This means there could be a large number of driving steps, and the readout could be made to happen days or even years in the future.

Once the dynamics are complete and the message has been played to the intended audience, there would be no way to recover it again by prying eyes. The message self-destructs.

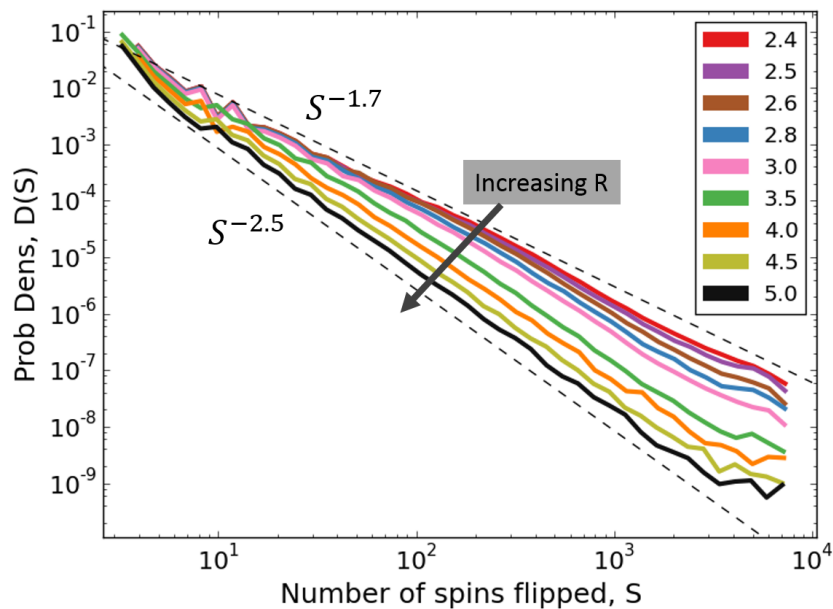


Figure 5.14: Self-organized criticality in the Random Field Ising Model with relaxation. In a tweak to the $T = 0$ Random Field Ising Model, we allow the spins to relax completely between cascades. Rather than the disorder parameter R of the pinning fields controlling the maximum avalanche size as in the standard version of the model [90], we find power law event magnitude distributions with no observable upper cutoff. The relaxation in the system trades a smaller s^* for a larger value of τ , creating results akin to [87] but significantly different from those of the plasticity experiments in Chapter 4.

5.9 The mysterious absence of self-organized criticality in

Chapter 4

An interesting study by Papanikolaou *et al.* [133] found that when plastic deformation was slowed down to the point that significant relaxation occurred between avalanche events, a sort of self-organized criticality (SOC, see [134, 135]) emerged. Similar results were found by adding viscoelastic relaxation to the Burridge-Knopoff block-slider model for earthquakes [136]. We also find the same behavior in the Random Field Ising Model (see, e.g., [90]) when relaxation of the spins occurs on timescales faster than the driving (Figure 5.14, where R is the disorder in the pinning fields on the spins). Indeed, SOC is often claimed or referred to in plasticity experiments due to the slow driving and power law statistics [17, 22, 18, 55].

Signatures of self-organized criticality in these systems are large values of s^* across the

board and exponents τ that increase from the case without relaxation. As we have just seen in this chapter and in Section 4.7, there is relaxation occurring in these granular packings across timescales that should be relevant during compression experiments. Where is the SOC in Chapter 4?

The answer may be connected to the reason the stress drop distributions were more compatible with $\tau = 1.5$ than 1.3, even though the latter is seemingly more applicable to a granular packing without shear bands. The processes underlying the stress relaxation measured in this chapter must occur throughout the continued deformation of Chapter 4 (see, e.g., the amount of relaxation that occurs on a seconds-long timescale in Figure 5.1). If the stress relaxation results from smoothing out fluctuations in the stress field as discussed in Section 5.7.2, it may erase the correlations that arise from the anisotropic stress redistribution and give rise to $\tau=1.3$. If the stress field smooths more than it decreases, the effect after a stress drop would be more similar to a mean-field redistribution as in [74]. In this case, the relaxation of the system is not resetting the stress field between each event, as in our version of the Random Field Ising Model with complete relaxation, and in the work of [133].

5.10 Summary

This chapter began with a simple question about the dependence of the plastic flow rate on stress, with implications for the experiments of Chapter 4. That the stress relaxation of the packing was approximately logarithmic in time suggested that granular packings have a complex inner state capable of being imprinted with memories of past driving. Indeed, with the right driving protocol, a packing can be made to relax by *increasing* stress for some amount of time before turning around and continuing the decay to zero.

The Amir, Oreg, Imry variant of the distribution of relaxation times theory is able to provide intuition for the memory dynamics. The system is modeled as a distribution of simple, single-timescale exponentially relaxing modes, and memories are held in the distribution of their amplitudes. Memories are stashed across different timescales and revealed in the

reverse order in which they were stored. We find that the fitting form of 3-step relaxation seems to require some sort of interaction between the relaxing modes. This idea is corroborated by one of the granular relaxation events in Figure 5.13, where the event sparked a new logarithmic decay in force. Cast within the AOI DRT framework, a new logarithmic decay could only arise from a renewal of the fast rates in the system – the ones which had long before relaxed to zero. Thus as the system reconfigures it may redistribute stress from some modes to others.

The presence of granular events also implicates slow particle-particle sliding as the process underlying stress relaxation. We complement the picture from Chapter 4 of a granular material evolving by particle rearrangements occurring over a wide range of length scales. Under stress, a granular material reconfigures over a wide range of timescales as well.

PARTICLE SHAPE, SURFACE ROUGHNESS, AND THE DEGRADATION OF SHEAR STRENGTH

Extending our focus to even longer timescales, this chapter examines the degradation of the properties of a granular material across multiple compression tests. We find that the wear and tear of plastic deformation permanently alters the surface roughness of the particles and highlights the role of additional features beyond particle shape in mechanical properties of the granular material.

While particle shape drives how contacting surfaces meet, the nature of the surfaces themselves tells the rest of the story. Many mechanical properties of granular materials have been found to depend strongly on the surface roughness of the particles [137, 138, 37, 139, 140, 141]. The process of plastic deformation erodes contacting surfaces and causes the properties of the granular material as a whole to evolve. In geology, faults and the gouge between them scrape against one another, transforming the shape of the grains, their size dispersion, and the roughness of all surfaces over time [142, 143]. In railroad engineering, wear and degradation undergone by the ballast beneath the rails results in an overall loss of shear strength, necessitating frequent maintenance and replacement [144].

In this chapter we study the intertwined roles of particle shape and surface properties in a granular material by way of one particular shape's propensity to wear down under repeated shear. Comparing assemblies of particles, we find a striking change in behavior for convex lens-shaped particles and only marginal evolution for sphere and tetrahedron particles. From this we infer that the macro-scale behavior of the granular material can dramatically magnify changes that occur at the micro-scale, especially for particles whose dominant mechanism of relative motion under shear is sliding. In particular, we show by combining x-ray tomography with high-resolution surface metrology that the degradation of strength of the lens particles is correlated with small changes in the particles surface roughness, while the overall particle

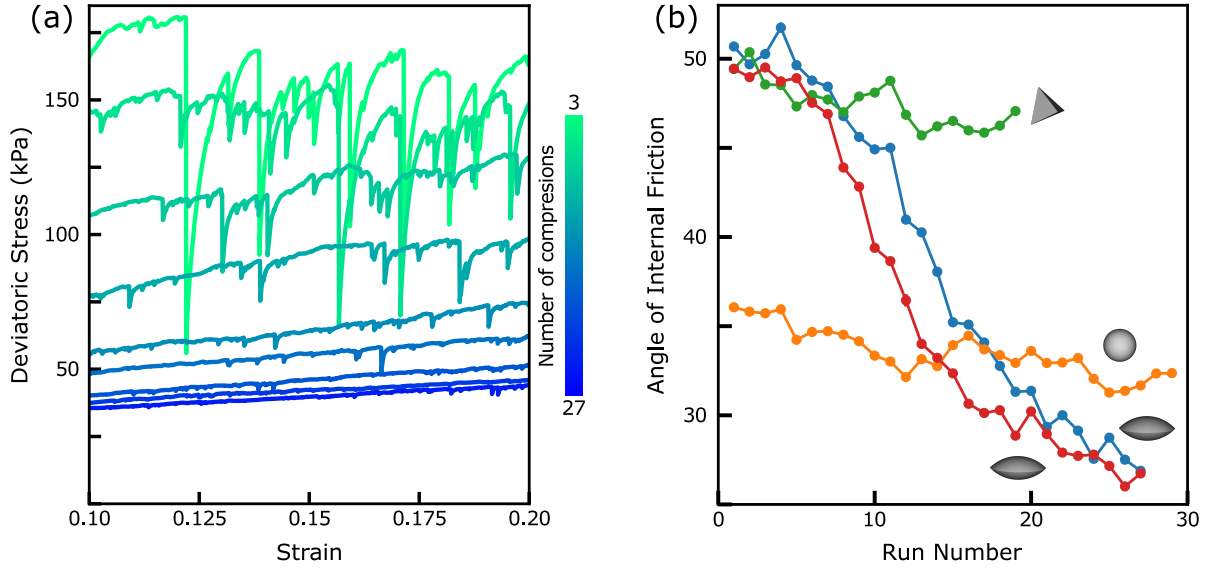


Figure 6.1: Rapid loss of shear strength for lens particles. **(a)** Raw stress-strain data for every third uniaxial compression of an assembly of 45° lenses under fixed confining pressure of 20kPa, showing rapid degradation of the plateau stress in the large strain regime. **(b)** Angle of internal friction versus compression test number for tetrahedra, spheres, and two different sets of 45° lenses.

shape and the distribution of contact locations on the particles surface remain essentially invariant.

6.1 Degradation of aggregate strength with repeated plastic deformation

In Figure 6.1a, the stress needed to plastically deform an assembly of lens particles drops by more than a factor of four after 27 compression tests. The angle of internal friction, shown in Figure 6.1b, decreases by almost 25 degrees, such that by 20 compression cycles the lenses are able support less shear stress than the spheres. This evolution of the lens particles lies in stark contrast to the tetrahedra and spheres, for which after 19 and 30 compression tests, respectively, the angle of internal friction decreases only marginally from its starting value.

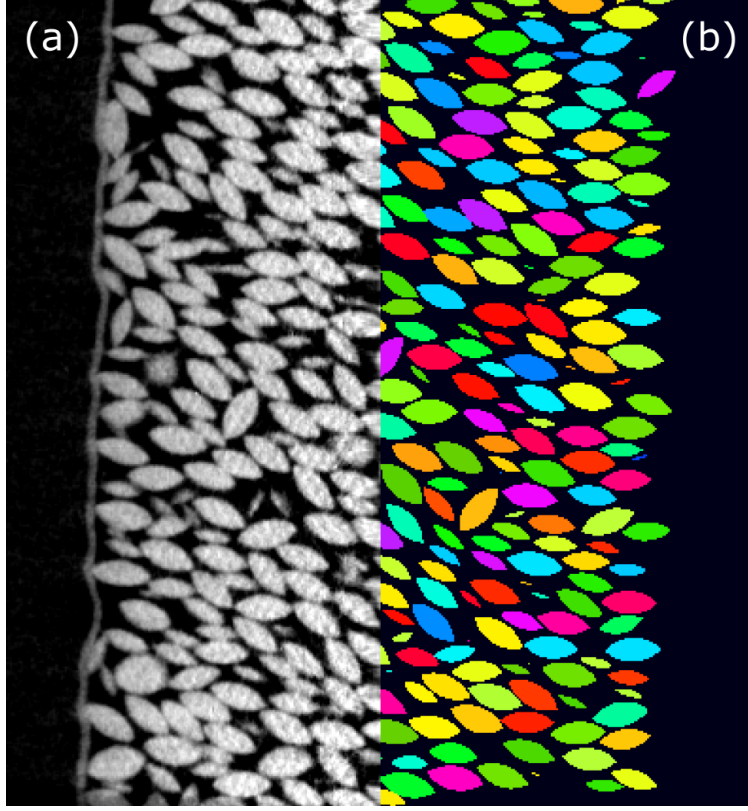


Figure 6.2: Fitted tomography of lens particles, cross section. **(a)** Raw tomography. **(b)** Reconstructed lenses. Panel **(a)** shows the left half of an imaged slice through the center of a packing, while **(b)** shows the right half of the same slice after fitting all of the particles.

6.2 Brick-laying configuration and orientational ordering

As discussed in Chapter 5, the fabric of a granular material is filled with degrees of freedom and thus could be holding the secret to the lens particles' rapid degradation in shear strength. We turn to computed tomography (CT) scans of the packing to check for evolution in the particles' configuration.

Following techniques used to segment CT scans of granular material [145, 39, 38], the raw tomography volume (Fig. 6.2a) was binarized and then segmented using a watershed algorithm. To exclude boundary effects arising from the confining membrane, lenses within two particle widths were excluded from analysis. This first step segmented about 85% of the lens particles, which were then individually fit by numerical optimization that maximized overlap between the segmented volume and an ideal lens shape, as done with tetrahedra in

[38]. The fitted particles from the first pass were subtracted from the tomography so that the remaining particles could be located and fit iteratively using the same method.

The optimization occasionally got stuck in a local minimum and placed a lens particle in the crook between two actual particles. This failure mode was caught by examining the “husk” left behind in the tomography after subtracting a fitted particle. If the fit were correctly centered on a particle, the husk would be a thin shell removable by a binary erosion process. If not, the husk would contain a mass of voxels that survived binary erosion, in which case the fitted particle was removed and the optimization began again with a new random starting seed. The process successfully fit more than 99.5% of the particles, missing only one or two in each packing’s center core with a diameter of approximately 3cm. The reconstructed lenses can be seen in Figure 6.2b, which align seamlessly down the middle with the raw tomography.

Exact determination of contacts from a tomography of a granular packing is rendered impossible by limitations arising from a combination of factors that include particle imperfections, limited imaging precision together with uncertainties in reconstruction, and very generally the small length scales involved in defining an actual contact [38]. The best we can do is to find physically justifiable criteria for identifying contacts which minimize the inclusion of ‘spurious’ contacts and the exclusion of legitimate ones. To this end we applied a procedure previously used with spheres [145], ellipsoids [39, 146], and tetrahedra [38]. All particles were grown by uniform dilation, starting from the fitted centers and orientations, while counting the total number of contacts. In the ideal case, i.e., in the absence of tomography noise, fitting errors, and particle imperfections, the average contact count across all particles in the assembly would jump discontinuously from zero to a nonzero value when the particles have been dilated to their actual size. More realistically, the locations of the contacts can be assumed to be afflicted by noise from a normal distribution whose width folds in the various sources of error in the fitting process [39]. As a result, the measured contact count is expected to grow with the virtual particle length scale as a cumulative normal

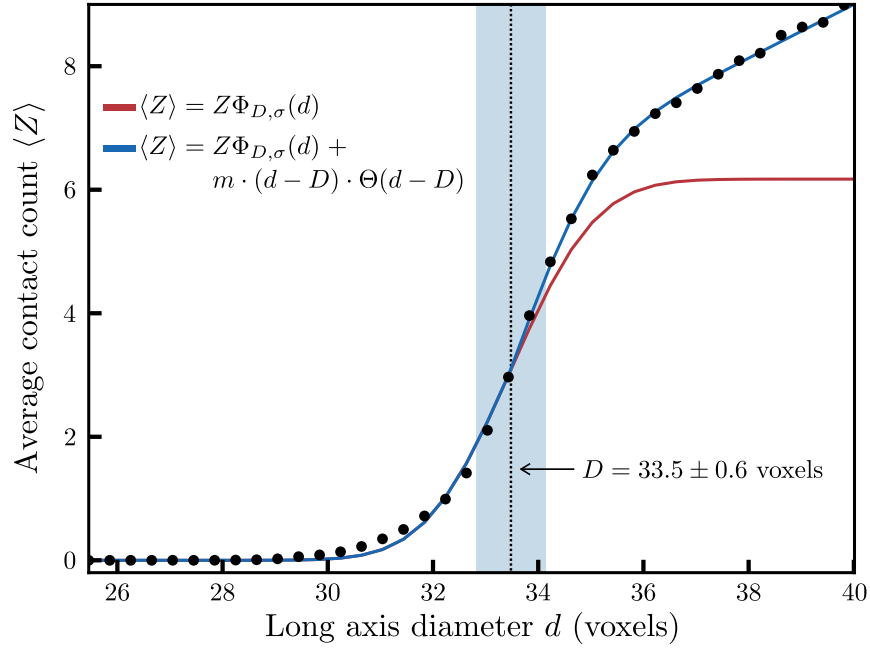


Figure 6.3: Finding the actual particle size in a reconstructed tomography through dilation. After particles have been fit to the tomography, their size is varied as contacts are counted, yielding the black datapoints. These are fit with by a cumulative normal distribution plus a linear term that turns on after the true particle size. Shown is the fit, the fit without the linear term, and the found particle size whose uncertainty is the width of the cumulative normal distribution.

distribution with finite width σ . In Figure 6.3 we show the contact count as a function of the particle scale parameter, fitted to a cumulative normal distribution plus a linear term that results from extra contacts between close particles, after the true size is surpassed. From the fit parameters we obtain the particle scale parameter and the spread σ as a measure of the quality of the fits.

Three full reconstructions were performed, both before and after 27 compression tests were run, and for 0% and 20% strain. For each the value of σ was less than 2% of a particle diameter and approximately the same as in Ref. [38], indicating a satisfactory fit of the tomography data. From these reconstructions, spatial contact density distributions were obtained, shown in Figure 6.4.

The fitted computed tomography scans of the lens particles provide hints at the origin of their rapid evolution with load cycling. Neighboring lens particles frequently contact each other in a brick-laying pattern, seen qualitatively in Figure 6.2 and quantitatively in the excess contacts partway down the side of the particle body in Figure 6.4. Because the smoothly sloping face of the lens lacks obstructing features, contacts midway down the body of the particle draw their resistance to shear predominantly from surface friction between particles. There is an additional abundance of contacts at the edges of particles, suggesting that the sharp equator of the lens shape also plays a significant role in supporting shear stress.

The lens shape's oblateness, which leads to the stacked particle ordering [53], causes a strong resistance to rotation about the particle's major axis, as this would require a high degree of local dilation. Because of this, the particles are more likely to slide relative to one another than to roll [147, 148]. This leads to increased grinding of the surfaces in contact and places more importance on the surface roughness when accounting for the assembly's overall resistance to shear [43]. By contrast, tetrahedra and spheres would be more likely to roll in response to shear since their rotation dilates the surrounding volume less. In turn, for these shapes we expect a slower pace of smoothing on the faces of the particles and a

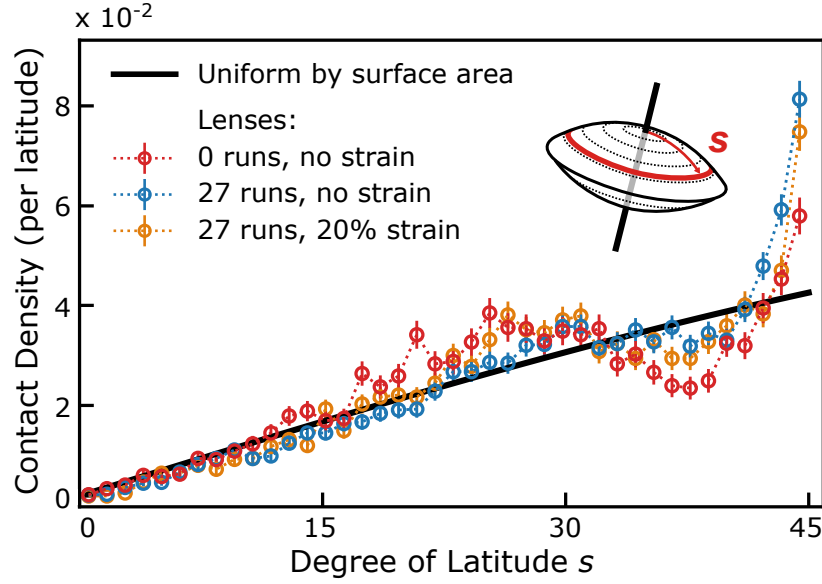


Figure 6.4: Spatial contact density distributions for lens assemblies. Contact points on every particle are binned by degree of latitude, from 0 (the pole) to 45 (the edge). The distributions are normalized to integrate to unity, and the black line is the distribution that would result from contacts distributed uniformly across the surface area of a particle. Data from three separate experiments are shown, with number of compressions and applied strain as indicated.

diminished dependence of shear strength on surface roughness.

No significant changes in the contact distributions are found between particles before and after 27 compression runs, nor between the start and end of a compression test (Figure 6.4). That the contact distribution changes so little between as-poured fresh and worn particles suggests that the evolution of the strength of the lens assembly is not occurring at the millimeter-scale and is not a direct result of how the particles pack together to support shear stresses. Additionally, the lack of significant change in the contact distribution between as-poured assemblies and assemblies after 20% applied strain indicates there is no disruptive restructuring of the contact network during compression, which could cause the drop in shear strength.

While the tomographic data indicate sliding, and thus particle surface friction, as the primary agent of shear strength for lens particles, the limited resolution of the reconstructed assemblies prohibits a direct analysis of the particle surface. To probe these smaller length

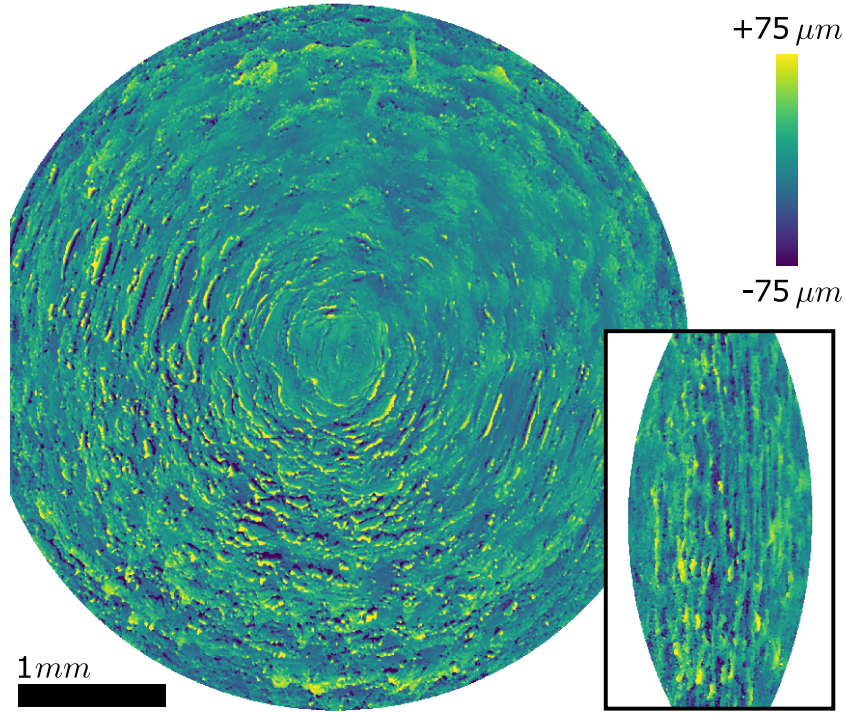


Figure 6.5: Surface scan of a freshly 3D-printed lens particle. Height deviations in the surface normal direction for a lens particle fresh after 3D printing viewed from the top (*main*) and from the side (*inset*). Both are displayed at the same scale.

scales and directly quantify the degree of surface degradation, high-resolution optical surface scans were used.

6.3 Surface smoothing

Surface scans of the particles immediately show imperfections in the 3D printing process that engender a rugged particle surface (Figure 6.5). Clearly apparent on the particle face at this magnification are grooves between print material layers as well as accumulations of material deposited during the printing and curing process. We hypothesize that the jaggedness of these surface features is the main source of friction between the particles, and that any smoothing of these features would lead to lower friction and thus lower shear strength of the assembly as a whole. Further, we reason that the relative sliding between particles, driven by the plastic deformation of the lens assembly under compression, causes these features to

grind and wear down gradually.

To analyze the surface scans, each data point of the height map was placed in its particle coordinate frame by numerically fitting an ideal lens shape to the entire scan, thus removing the overall particle curvature. Displacements normal to the ideal particle surface were calculated for each point, and the root mean square (rms) variation was used as a measure for roughness.

This spatial roughness calculation corresponds to the parameter S_q in surface metrology [149], calculated as

$$S_q = \sqrt{\langle (z - \bar{z})^2 \rangle}. \quad (6.1)$$

Here \bar{z} is the average normal displacement and the brackets indicate an average over the surface. S_q is used for both lens body and edge, where the edge roughness is extracted from a $50\mu\text{m}$ wide strip along the equator in each scan.

As demonstrated in Figure 6.6, the roughness coefficient S_q for lens particles indeed changes systematically with successive compression cycles. S_q on the body of the particles decreases by about 15% after 27 compressions. Near the edges the particle surfaces start out rougher, but smooth to a similar S_q value, undergoing a drop by about 33%. To fully characterize surfaces at contact and link this to a coefficient of friction in a quantitative manner, more comprehensive analysis is required [150, 151, 152]. Nevertheless, the decline in roughness S_q is clearly measurable and correlates with the degradation of shear strength of the assembly.

That edges smooth at a faster rate than the rest of the particle surface is likely due to the preponderance of contacts near the edge (Figure 6.4) as well as the fact that regions with higher curvature tend to erode faster than regions with lower curvature, a phenomenon used to explain pebble smoothing in riverbeds [153, 154]. On the freshly-printed lens particles, the curvature is $\sim (5000\mu\text{m})^{-1}$ and constant everywhere but the edge, where it becomes two orders of magnitude larger to reach $\sim (50\mu\text{m})^{-1}$. Note, however, that similarly sharp edges and corners also exist in the tetrahedra. Any smoothing out of these high-curvature

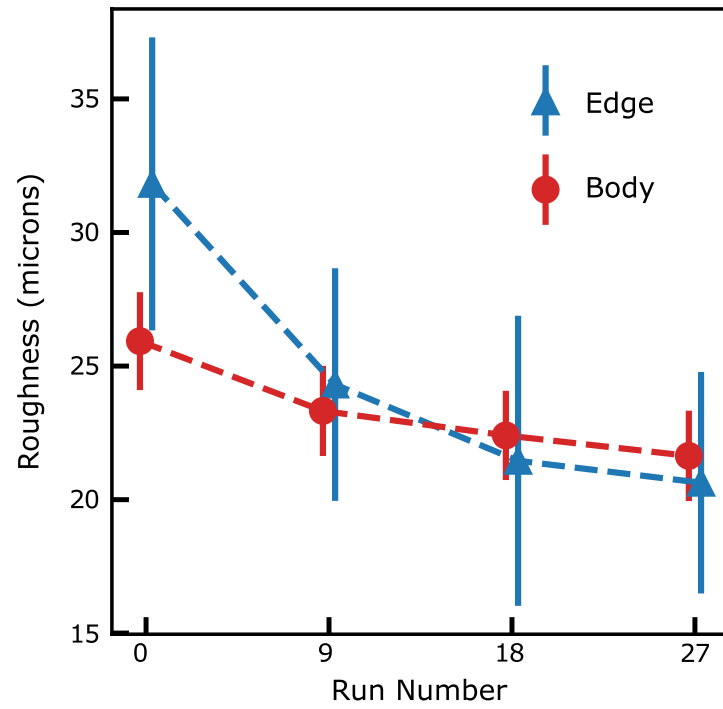


Figure 6.6: Evolution of RMS roughness S_q for body and edge of lens particles. Error bars show the standard deviation of the distribution of values in each case. Data points are offset slightly in the horizontal direction for visual clarity.

features in the particle geometry, therefore, is likely not the primary reason for the loss of shear strength seen in Figure 6.1. Instead, the lenses' fall from high to low angle of internal friction appears to be driven by decreasing surface friction on the body of the particles. This, together with the oblate lens shape, favors sliding, as was hinted in Figure 4.14. As Figure 6.1b demonstrates, once the lens surfaces have become sufficiently smooth, sliding can lead to an angle of internal friction or, equivalently, a shear strength of the granular assembly that is even lower than that for spheres.

6.4 Opening the door to granular surface forensics

The wear experienced by 3D-printed particles under repeated compressive loading is found to drive systematic, controllable changes in their surface roughness. A most striking outcome is how the sensitivity to surface roughness depends on particle shape. The geometry of the convex lenses imposes a degree of global ordering, shown in the tomography data (Figure 6.2), whereby the particles tend to align their minor axes with gravity. This alignment leads to an overabundance of particle contacts midway down the body of the lens and at the edge. The robustness of this spatial contact distribution, combined with the oblique nature of the particle contacts themselves, primes the assembly to deform predominantly by sliding motion. Repeated sliding then smooths out the particle surfaces and enhances the ability to slide. Finally, the finding that assemblies composed of more compact, less oblate particle shapes such as spheres and tetrahedra exhibit significantly lower sensitivity than lenses to repeated loading suggests new strategies by which to design granular materials that either experience wear in a predetermined way, or generate a desired global trend over repeated plastic deformation cycles.

These results highlight the complementary roles of particle shape and particle surface properties in determining the macro- and microscale behavior of granular materials. Combining measurements of overall material properties such as the strength under shear with x-ray tomography and particle surface metrology opens the door to performing detailed gran-

ular forensics. By batch-processing large numbers of particles with laser-scanning confocal microscopy we show that the history of deformation as preserved in the patterns of wear on the particle surfaces can be extracted in a statistically meaningful manner.

CONCLUSIONS

How does a granular system adapt? In this thesis we have observed change on the level of the aggregate (Chapter 3), in rapid cascades of particle rearrangements (Chapter 4), in glacial sliding over a wide range of timescales (Chapter 5), and on the scale of particle surface roughness (Chapter 6). With the cascades and the slow relaxation, the physics of the evolution was fundamentally connected to a broader family of disordered materials. As such, insights and intuitions gained from studying the granular system lead to a better understanding of general amorphous materials in plastic deformation in the former and of the relaxation of disordered systems out of equilibrium in the latter.

In every case study, the granular material was flowing in order to adapt to applied stress. When driven continuously, the granular system did not maintain a steady flow rate and frequently surged forward to keep up with the imposed rate of strain. When the driving ceased, flow continued at an ever decreasing pace, from hundredths of seconds to hours and presumably would have continued much longer. By imprinting a loading protocol and then watching the memories play themselves out, we found the flowing state to be complex enough to require many additional variables of description.

A detailed study of the surges of faster flow rates – the stress drop events – connected with a current debate on the nature of the homogenizing interaction between constituent pieces in a plastically deforming amorphous material. Namely, the exponent τ of the stress drop magnitude distributions equaling 1.3 or 1.5 means the full Eshelby form of the elastic interactions between the grains is required for understanding, or not. Across all the particle shapes, τ was found to be remarkably invariant, and slightly more compatible with a value 1.5. By changing shape we introduced a new way to modify the disorder in the system and the proximity of the plastic deformation to a critical point, as measured by s^* of the drop event magnitude distribution.

It became clear from the different rates of drop events per shape and the arcing recharge

in the stress following a drop event, that plastic deformation occurred continuously as a backdrop to the sudden events. By stopping compression midway into the deformation we were able to study the plastic flow as it progressed under its own devices. The slow logarithmic decline in stress connected the relaxation of the granular system to that of a wide class of other disordered systems showing anomalous relaxation. We described the dynamics in the context of a model with a broad distribution of relaxing elements, and the intuition gained allowed us to imprint multiple memories in the flow state of the relaxing granular column. Eventually, however, we outgrew the model, and found an extension that incorporates relaxation during driving, and discovered that the functional form of the relaxation seems to necessitate interaction between the abstract elements of the model. Events in the relaxation data showed a distinctly granular nature, allowing us to infer that the relaxation occurs via slow sliding at all of the particle-particle interfaces, which occasionally results in a grain rearrangement akin to the larger cascades during continued plastic deformation.

Finally, in work that leaned on more “microscopic” data from the granular packing, we explored the possibilities of granular forensics whereby the details of plastic flow can be studied through volumetric reconstructions of the granular column and surface metrology of statistically significant samples of particles. Ironically, it turned out that the most sensitive measure of wear in a granular material is through the mechanical properties of the aggregate.

As a model system, granular materials allow unique possibilities for the control and observation of the constituent grains, while offering rich behavior with far reaching connections to other disordered systems in nature. To the degree that the granular systems studied in this thesis are representative of disordered systems in general, perhaps the weeping philosopher Heraclitus was correct in a literal sense. On length scales much smaller than a particle up to the system size, and timescales ranging from fractions of a second to longer than we can probe, the granular system flows.

REFERENCES

- [1] Howard A. Barnes. The yield stress – a review or ‘ $\pi\alpha\nu\tau\alpha\ \rho\epsilon\iota$ ’ – everything flows? *Journal of Non-Newtonian Fluid Mechanics*, 81(1):133–178, 1999.
- [2] Kieran A. Murphy, Karin A. Dahmen, and Heinrich M. Jaeger. Transforming mesoscale granular plasticity through particle shape. *Physical Review X*, 9(1):011014, 2019.
- [3] Kieran A. Murphy, Arthur K. MacKeith, Leah K. Roth, and Heinrich M. Jaeger. The intertwined roles of particle shape and surface roughness in controlling the shear strength of a granular material. *arXiv e-prints*, pages arXiv:1902.03280 [cond-mat.soft], Feb 2019.
- [4] Beno Gutenberg and Charles F. Richter. Frequency of earthquakes in california. *Bulletin of the Seismological Society of America*, 34(4):185–188, 1944.
- [5] Baolian Cheng, Richard I. Epstein, Robert A. Guyer, and A. Cody Young. Earthquake-like behaviour of soft γ -ray repeaters. *Nature*, 382(6591):518–520, 1996.
- [6] Ivan Osorio, Mark G. Frei, Didier Sornette, John Milton, and Ying-Cheng Lai. Epileptic seizures: Quakes of the brain? *Physical Review E*, 82(2):021919, 2010.
- [7] Jörgen Vitting Andersen, Andrzej Nowak, Giulia Rotundo, Lael Parrott, and Sebastian Martinez. “Price-quakes” shaking the world’s stock exchanges. *PLOS ONE*, 6(11):e26472, 2011.
- [8] Stephen Jay Gould and Niles Eldredge. Punctuated equilibria: the tempo and mode of evolution reconsidered. *Paleobiology*, 3(2):115151, 1977.
- [9] Per Bak and Kim Sneppen. Punctuated equilibrium and criticality in a simple model of evolution. *Physical Review Letters*, 71(24):4083–4086, 1993.
- [10] Olga Perković, Karin Dahmen, and James P. Sethna. Avalanches, barkhausen noise, and plain old criticality. *Physical Review Letters*, 75(24):4528–4531, 1995.
- [11] Sidney R. Nagel. Klopsteg memorial lecture (august, 1998): Physics at the breakfast table or waking up to physics. *American Journal of Physics*, 67(1):17–25, 1999.
- [12] Deniz Ertar and Mehran Kardar. Critical dynamics of contact line depinning. *Physical Review E*, 49(4):R2532–R2535, 1994.
- [13] Jean-Christophe Baret, Damien Vandembroucq, and Stéphane Roux. Extremal model for amorphous media plasticity. *Physical Review Letters*, 89(19):195506, 2002.
- [14] Jie Lin, Edan Lerner, Alberto Rosso, and Matthieu Wyart. Scaling description of the yielding transition in soft amorphous solids at zero temperature. *Proceedings of the National Academy of Sciences*, 111(40):14382–14387, 2014.

- [15] Misaki Ozawa, Ludovic Berthier, Giulio Biroli, Alberto Rosso, and Gilles Tarjus. Random critical point separates brittle and ductile yielding transitions in amorphous materials. *Proceedings of the National Academy of Sciences*, 115(26):6656, 2018.
- [16] Daniel S. Fisher. Collective transport in random media: from superconductors to earthquakes. *Physics Reports*, 301(1):113–150, 1998.
- [17] Jérôme Weiss and Jean-Robert Grasso. Acoustic emission in single crystals of ice. *The Journal of Physical Chemistry B*, 101(32):6113–6117, 1997.
- [18] M. Carmen Miguel, Alessandro Vespignani, Stefano Zapperi, Jérôme Weiss, and Jean-Robert Grasso. Intermittent dislocation flow in viscoplastic deformation. *Nature*, 410(6829):667–671, 2001.
- [19] D. J. Durian. Bubble-scale model of foam mechanics: Melting, nonlinear behavior, and avalanches. *Physical Review E*, 55(2):1739–1751, 1997.
- [20] I. Cantat and O. Pitois. Stokes experiment in a liquid foam. *Physics of Fluids*, 18(8):083302, 2006.
- [21] A. S. Argon. Strain avalanches in plasticity. *Philosophical Magazine*, 93(28-30):3795–3808, 2013.
- [22] B. A. Sun, H. B. Yu, W. Jiao, H. Y. Bai, D. Q. Zhao, and W. H. Wang. Plasticity of ductile metallic glasses: A self-organized critical state. *Physical Review Letters*, 105(3):035501, 2010.
- [23] James Antonaglia, Wendelin J. Wright, Xiaojun Gu, Rachel R. Byer, Todd C. Hufnagel, Michael LeBlanc, Jonathan T. Uhl, and Karin A. Dahmen. Bulk metallic glasses deform via slip avalanches. *Physical Review Letters*, 112(15):155501, 2014.
- [24] Steffen Brinckmann, Ju-Young Kim, and Julia R. Greer. Fundamental differences in mechanical behavior between two types of crystals at the nanoscale. *Physical Review Letters*, 100(15):155502, 2008.
- [25] D. V. Denisov, K. A. Lörincz, J. T. Uhl, K. A. Dahmen, and P. Schall. Universality of slip avalanches in flowing granular matter. *Nature Communications*, 7:10641, 2016.
- [26] Jonathan T. Uhl, Shivesh Pathak, Danijel Schorlemmer, Xin Liu, Ryan Swindeman, Braden A. W. Brinkman, Michael LeBlanc, Georgios Tsekenis, Nir Friedman, Robert Behringer, Dmitry Denisov, Peter Schall, Xiaojun Gu, Wendelin J. Wright, Todd Hufnagel, Andrew Jennings, Julia R. Greer, P. K. Liaw, Thorsten Becker, Georg Dresen, and Karin A. Dahmen. Universal quake statistics: From compressed nanocrystals to earthquakes. *Scientific Reports*, 5:16493, 2015.
- [27] Jonathan Barés, Dengming Wang, Dong Wang, Thibault Bertrand, Corey S. O’Hern, and Robert P. Behringer. Local and global avalanches in a two-dimensional sheared granular medium. *Physical Review E*, 96(5):052902, 2017.

- [28] Samuel Poincloux, Mokhtar Adda-Bedia, and Frédéric Lechenault. Crackling dynamics in the mechanical response of knitted fabrics. *Physical Review Letters*, 121(5):058002, 2018.
- [29] Alexandre Nicolas, Ezequiel E. Ferrero, Kirsten Martens, and Jean-Louis Barrat. Deformation and flow of amorphous solids: Insights from elastoplastic models. *Reviews of Modern Physics*, 90(4):045006, 2018.
- [30] Zoe Budrikis, David Fernandez Castellanos, Stefan Sandfeld, Michael Zaiser, and Stefano Zapperi. Universal features of amorphous plasticity. *Nature Communications*, 8:15928, 2017.
- [31] J. S. Langer. Shear-transformation-zone theory of yielding in athermal amorphous materials. *Physical Review E*, 92(1):012318, 2015.
- [32] R. Kohlrausch. Nachtrag uber die elastische nachwirkung beim cocon und glasladen. *Ann. Phys. (Leipzig)*, 12(393):7, 1847.
- [33] Nathan C. Keim, Joseph Paulsen, Zorana Zeravcic, Srikanth Sastry, and Sidney R. Nagel. Memory formation in matter. *arXiv e-prints*, pages arXiv:1810.08587 [cond-mat.soft], Oct 2018.
- [34] Patrick Richard, Mario Nicodemi, Renaud Delannay, Philippe Ribière, and Daniel Bideau. Slow relaxation and compaction of granular systems. *Nature Materials*, 4(2):121–128, 2005.
- [35] P. G. de Gennes. Granular matter: a tentative view. *Reviews of Modern Physics*, 71(2):S374–S382, 1999.
- [36] Athanasios G. Athanassiadis, Marc Z. Miskin, Paul Kaplan, Nicholas Rodenberg, Seung Hwan Lee, Jason Merritt, Eric Brown, John Amend, Hod Lipson, and Heinrich M. Jaeger. Particle shape effects on the stress response of granular packings. *Soft Matter*, 10(1):48–59, 2014.
- [37] Sonia Utermann, Philipp Aurin, Markus Benderoth, Cornelius Fischer, and Matthias Schröter. Tailoring the frictional properties of granular media. *Physical Review E*, 84(3):031306, 2011.
- [38] Max Neudecker, Stephan Ulrich, Stephan Herminghaus, and Matthias Schröter. Jammed frictional tetrahedra are hyperstatic. *Physical Review Letters*, 111(2):028001, 2013.
- [39] Fabian M. Schaller, Max Neudecker, Mohammad Saadatfar, Gary Delaney, Klaus Mecke, Gerd E. Schröder-Turk, and Matthias Schröter. Tomographic analysis of jammed ellipsoid packings. *AIP Conference Proceedings*, 1542(1):377–380, 2013.
- [40] T. S. Majmudar and R. P. Behringer. Contact force measurements and stress-induced anisotropy in granular materials. *Nature*, 435(7045):1079–1082, 2005.

- [41] Emilien Azéma, Farhang Radjaï, Robert Peyroux, and Gilles Saussine. Force transmission in a packing of pentagonal particles. *Physical Review E*, 76(1):011301, 2007.
- [42] Emilien Azéma, Nicolas Estrada, and Farhang Radjaï. Nonlinear effects of particle shape angularity in sheared granular media. *Physical Review E*, 86(4):041301, 2012.
- [43] Mauricio Botton, Emilien Azéma, Nicolas Estrada, Farhang Radjaï, and Arcesio Lizcano. Quasistatic rheology and microstructural description of sheared granular materials composed of platy particles. *Physical Review E*, 87(3):032206, 2013.
- [44] Nicholas W. Hayman, Lucie Ducloué, Kate L. Foco, and Karen E. Daniels. Granular controls on periodicity of stick-slip events: Kinematics and force-chains in an experimental fault. *Pure and Applied Geophysics*, 168(12):2239–2257, 2011.
- [45] Matt Harrington and Douglas J. Durian. Anisotropic particles strengthen granular pillars under compression. *Physical Review E*, 97(1):012904, 2018.
- [46] T. Doanh, M. T. Hoang, J.-N. Roux, and C. Dequeker. Stick-slip behaviour of model granular materials in drained triaxial compression. *Granular Matter*, 15(1):1–23, 2013.
- [47] Deshan Cui, Wei Wu, Wei Xiang, Thiep Doanh, Qiong Chen, Shun Wang, Qingbing Liu, and Jinge Wang. Stick-slip behaviours of dry glass beads in triaxial compression. *Granular Matter*, 19(1):1, 2016.
- [48] A Ozbay and AF Cabalar. Effects of triaxial confining pressure and strain rate on stick-slip behavior of a dry granular material. *Granular Matter*, 18(3):1–9, 2016.
- [49] Ignazio Cavarretta, Matthew Coop, and Catherine O’Sullivan. The influence of particle characteristics on the behaviour of coarse grained soils. *Géotechnique*, 60(6):413–423, 2010.
- [50] JC Santamarina and GC Cho. Soil behaviour: the role of particle shape. In *Advances in geotechnical engineering: The skempton conference*, volume 1, pages 604–617. Thomas Telford, London, 2004.
- [51] Jennifer L Anthony and Chris Marone. Influence of particle characteristics on granular friction. *Journal of Geophysical Research: Solid Earth*, 110(B8), 2005.
- [52] Karen Mair, Kevin M. Frye, and Chris Marone. Influence of grain characteristics on the friction of granular shear zones. *Journal of Geophysical Research: Solid Earth*, 107(B10):ECV 4–1–ECV 4–9, 2002.
- [53] A Mahmood and JK Mitchell. Fabric-property relationships in fine granular materials. *Clays and Clay Miner*, 1974.
- [54] D. J. Henkel and G. D. Gilbert. The effect measured of the rubber membrane on the triaxial compression strength of clay samples. *Géotechnique*, 3(1):20–29, 1952.
- [55] Yinan Cui, Giacomo Po, and Nasr Ghoniem. Controlling strain bursts and avalanches at the nano- to micrometer scale. *Physical Review Letters*, 117(15):155502, 2016.

- [56] Mehdi Talamali, Viljo Petäjä, Damien Vandembroucq, and Stéphane Roux. Avalanches, precursors, and finite-size fluctuations in a mesoscopic model of amorphous plasticity. *Physical Review E*, 84(1):016115, 2011.
- [57] S. Nasuno, A. Kudrolli, A. Bak, and J. P. Gollub. Time-resolved studies of stick-slip friction in sheared granular layers. *Physical Review E*, 58(2):2161–2171, 1998.
- [58] Istvan Albert, Pal Tegzes, R Albert, JG Sample, A-L Barabási, T Vicsek, B Kahng, and P Schiffer. Stick-slip fluctuations in granular drag. *Physical Review E*, 64(3):031307, 2001.
- [59] R. Maaß, M. Wraith, J. T. Uhl, J. R. Greer, and K. A. Dahmen. Slip statistics of dislocation avalanches under different loading modes. *Physical Review E*, 91(4):042403, 2015.
- [60] L. Knopoff and R. Burridge. Model and theoretical seismicity. *Bulletin of the Seismological Society of America*, 57(3):341–371, 1967.
- [61] ML Nguyen and SN Coppersmith. Scalar model of inhomogeneous elastic and granular media. *Physical review E*, 62(4):5248, 2000.
- [62] Athanasios G. Athanassiadis, Patrick J. La Rivière, Emil Sidky, Charles Pelizzari, Xiaochuan Pan, and Heinrich M. Jaeger. X-ray tomography system to investigate granular materials during mechanical loading. *Review of Scientific Instruments*, 85(8):083708, 2014.
- [63] Andrew Schofield and Peter Wroth. *Critical state soil mechanics*, volume 310. McGraw-Hill London, 1968.
- [64] A. Tordesillas. Force chain buckling, unjamming transitions and shear banding in dense granular assemblies. *Philosophical Magazine*, 87(32):4987–5016, 2007.
- [65] Osborne Reynolds. LVII. On the dilatancy of media composed of rigid particles in contact. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 20(127):469–481, 1885.
- [66] C. A. Volkert, A. Donohue, and F. Spaepen. Effect of sample size on deformation in amorphous metals. *Journal of Applied Physics*, 103(8):083539, 2008.
- [67] Dongchan Jang and Julia R. Greer. Transition from a strong-yet-brittle to a stronger-and-ductile state by size reduction of metallic glasses. *Nature Materials*, 9:215, 2010.
- [68] Joseph F. Labuz and Arno Zang. Mohr-coulomb failure criterion. *Rock Mechanics and Rock Engineering*, 45(6):975–979, 2012.
- [69] Kieran A Murphy, Nikolaj Reiser, Darius Choksy, Clare E Singer, and Heinrich M Jaeger. Freestanding loadbearing structures with Z-shaped particles. *Granular Matter*, 18(2):26, 2016.

- [70] Pablo F. Damasceno, Michael Engel, and Sharon C. Glotzer. Predictive self-assembly of polyhedra into complex structures. *Science*, 337(6093):453–457, 2012.
- [71] Andrew Adams, Della Bloomfield, Philip Booth, and Peter England. *Investment Mathematics and Statistics*. Graham & Trotman, 1993.
- [72] Michael LeBlanc, Aya Nawano, Wendelin J. Wright, Xiaojun Gu, J. T. Uhl, and Karin A. Dahmen. Avalanche statistics from data with low time resolution. *Physical Review E*, 94(5):052135, 2016.
- [73] Karin A. Dahmen, Yehuda Ben-Zion, and Jonathan T. Uhl. Micromechanical model for deformation in solids with universal predictions for stress-strain curves and slip avalanches. *Physical Review Letters*, 102(17):175501, 2009.
- [74] Karin A. Dahmen, Yehuda Ben-Zion, and Jonathan T. Uhl. A simple analytic theory for the statistics of avalanches in sheared granular materials. *Nat Phys*, 7(7):554–557, 2011.
- [75] Ekhard K. H. Salje and Karin A. Dahmen. Crackling noise in disordered materials. *Annual Review of Condensed Matter Physics*, 5(1):233–254, 2014.
- [76] K. Michael Salerno and Mark O. Robbins. Effect of inertia on sheared disordered solids: Critical scaling of avalanches in two and three dimensions. *Physical Review E*, 88(6):062206, 2013.
- [77] Chen Liu, Ezequiel E. Ferrero, Francesco Puosi, Jean-Louis Barrat, and Kirsten Martens. Driving rate dependence of avalanche statistics and shapes at the yielding transition. *Physical Review Letters*, 116(6):065501, 2016.
- [78] Aghil Abed Zadeh, Jonathan Bars, Joshua E. S. Socolar, and Robert P. Behringer. Seismicity in sheared granular matter. *Physical Review E*, 99(5):052902, 2019.
- [79] Penghui Cao, Karin A. Dahmen, Akihiro Kushima, Wendelin J. Wright, Harold S. Park, Michael P. Short, and Sidney Yip. Nanomechanics of slip avalanches in amorphous plasticity. *Journal of the Mechanics and Physics of Solids*, 114:158–171, 2018.
- [80] J. D. Eshelby. The determination of the elastic field of an ellipsoidal inclusion, and related problems. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 241(1226):376, 1957.
- [81] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.
- [82] A. S. Argon. Plastic deformation in metallic glasses. *Acta Metallurgica*, 27(1):47–58, 1979.
- [83] A. S. Argon and H. Y. Kuo. Plastic flow in a disordered bubble raft (an analog of a metallic glass). *Materials Science and Engineering*, 39(1):101–109, 1979.

- [84] M. L. Falk and J. S. Langer. Dynamics of viscoplastic deformation in amorphous solids. *Physical Review E*, 57(6):7192–7205, 1998.
- [85] Craig Maloney and Anaël Lemaître. Subextensive scaling in the athermal, quasistatic limit of amorphous matter in plastic shear flow. *Physical Review Letters*, 93(1):016001, 2004.
- [86] Aghil Abed Zadeh, Jonathan Barés, and Robert P. Behringer. Avalanches in a granular stick-slip experiment: detection using wavelets. *EPJ Web Conf.*, 140, 2017.
- [87] Stefanos Papanikolaou, Felipe Bohn, Rubem Luis Sommer, Gianfranco Durin, Stefano Zapperi, and James P. Sethna. Universality beyond power laws and the average avalanche shape. *Nature Physics*, 7:316, 2011.
- [88] Bradley Efron and Robert J Tibshirani. *An introduction to the bootstrap*. CRC press, 1994.
- [89] Jie Lin, Thomas Gueudré, Alberto Rosso, and Matthieu Wyart. Criticality in the approach to failure in amorphous solids. *Physical Review Letters*, 115(16):168001, 2015.
- [90] James P. Sethna, Karin A. Dahmen, and Christopher R. Myers. Crackling noise. *Nature*, 410(6825):242–250, 2001.
- [91] Thomas H Heaton. Evidence for and implications of self-healing pulses of slip in earthquake rupture. *Physics of the Earth and Planetary Interiors*, 64(1):1–20, 1990.
- [92] A. W. Bishop. The strength of soils as engineering materials. *Geotechnique*, 16(2):91–130, 1966.
- [93] Matthew R. Kuhn and James K. Mitchell. New perspectives on soil creep. *Journal of Geotechnical Engineering*, 119(3):507–524, 1993.
- [94] G. R. McDowell and J. J. Khan. Creep of granular materials. *Granular Matter*, 5(3):115–120, 2003.
- [95] J. Bai, N. Morgenstern, and D. Chan. Three-dimensional creep analyses of the Leaning Tower of Pisa. *Soils and Foundations*, 48(2):195–205, 2008.
- [96] Jasna Brujić, Ping Wang, Chaoming Song, David L. Johnson, Olivier Sindt, and Hernn A. Makse. Granular dynamics in compaction and stress relaxation. *Physical Review Letters*, 95(12):128001, 2005.
- [97] Amandine Miksic and Mikko J. Alava. Evolution of grain contacts in a granular sample under creep and stress relaxation. *Physical Review E*, 88(3):032207, 2013.
- [98] David M. Kaz, Ryan McGorty, Madhav Mani, Michael P. Brenner, and Vinothan N. Manoharan. Physical ageing of the contact line on colloidal particles at liquid interfaces. *Nature Materials*, 11:138, 2011.

- [99] A. Gurevich and H. K  pfer. Time scales of the flux creep in superconductors. *Physical Review B*, 48(9):6477–6487, 1993.
- [100] Kittiwit Matan, Rachel B. Williams, Thomas A. Witten, and Sidney R. Nagel. Crumpling a thin sheet. *Physical Review Letters*, 88(7):076101, 2002.
- [101] Oded Ben-David, Shmuel M. Rubinstein, and Jay Fineberg. Slip-stick and the evolution of frictional strength. *Nature*, 463:76, 2010.
- [102] Sam Dillavou and Shmuel M. Rubinstein. Nonmonotonic aging and memory in a frictional interface. *Physical Review Letters*, 120(22):224101, 2018.
- [103] T. C. Guo and W. W. Guo. A transient-state theory of dielectric relaxation and the curie-von schweidler law. *Journal of Physics C: Solid State Physics*, 16(10):1955, 1983.
- [104] Andrew K. Jonscher. Dielectric relaxation in solids. *Journal of Physics D: Applied Physics*, 32(14):R57, 1999.
- [105] Nicholas W. Tschoegl. *The phenomenological theory of linear viscoelastic behavior: an introduction*. Springer Science & Business Media, 2012.
- [106] Reiner Zorn. Logarithmic moments of relaxation time distributions. *The Journal of Chemical Physics*, 116(8):3204–3209, 2002.
- [107] H. Sch  fer, E. Sternin, R. Stannarius, M. Arndt, and F. Kremer. Novel approach to the analysis of broadband dielectric spectra. *Physical Review Letters*, 76(12):2177–2180, 1996.
- [108] Ariel Amir, Yuval Oreg, and Yoseph Imry. On relaxations and aging of various glasses. *Proceedings of the National Academy of Sciences*, 109(6):1850–1855, 2012.
- [109] Ariel Amir, Stefano Borini, Yuval Oreg, and Yoseph Imry. Huge (but finite) time scales in slow relaxations: Beyond simple aging. *Physical Review Letters*, 107(18):186407, 2011.
- [110] Ariel Amir, Yuval Oreg, and Yoseph Imry. Localization, anomalous diffusion, and slow relaxations: A random distance matrix approach. *Physical Review Letters*, 105(7):070601, 2010.
- [111] George B. Arfken and Hans J. Weber. *Mathematical Methods for Physicists*. Academic Press, 2011.
- [112] E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957.
- [113] Claude Elwood Shannon. A mathematical theory of communication. *Bell system technical journal*, 27(3):379–423, 1948.
- [114] Stefano Martiniani, Paul M. Chaikin, and Dov Levine. Quantifying hidden order out of equilibrium. *Physical Review X*, 9(1):011031, 2019.

- [115] Christopher W. Lynn, Lia Papadopoulos, Daniel D. Lee, and Danielle S. Bassett. Surges of collective human activity emerge from simple pairwise correlations. *Physical Review X*, 9(1):011022, 2019.
- [116] E. DeGiuli. Random language model. *Physical Review Letters*, 122(12):128301, 2019.
- [117] A. K. Livesey and J. C. Brochon. Analyzing the distribution of decay constants in pulse-fluorimetry using the maximum entropy method. *Biophysical Journal*, 52(5):693–706, 1987.
- [118] Yoav Lahini, Omer Gottesman, Ariel Amir, and Shmuel M. Rubinstein. Nonmonotonic aging and memory retention in disordered mechanical systems. *Physical Review Letters*, 118(8):085501, 2017.
- [119] Yi He, Deyi Jiang, Jie Chen, Rong Liu, Jinyang Fan, and Xiang Jiang. Non-monotonic relaxation and memory effect of rock salt. *Rock Mechanics and Rock Engineering*, 2019.
- [120] AJ Kovacs. Glass transition in amorphous polymers: a phenomenological study. *Adv. Polym. Sci.*, 3(3):394–507, 1963.
- [121] E. M. Bertin, J. P. Bouchaud, J. M. Drouffe, and C. Godrèche. The Kovacs effect in model glasses. *Journal of Physics A: Mathematical and General*, 36(43):10701–10719, 2003.
- [122] R. Edgeworth, B. J. Dalton, and T. Parnell. The pitch drop experiment. *European Journal of Physics*, 5(4):198–200, 1984.
- [123] K. Chen, J. Cole, C. Conger, J. Draskovic, M. Lohr, K. Klein, T. Scheidemantel, and P. Schiffer. Packing grains by thermal cycling. *Nature*, 442(7100):257–257, 2006.
- [124] Thibaut Divoux, Hervé Gayvallet, and Jean-Christophe Gémard. Creep motion of a granular pile induced by thermal cycling. *Physical Review Letters*, 101(14):148303, 2008.
- [125] A. Heuer, M. Wilhelm, H. Zimmermann, and H. W. Spiess. Rate memory of structural relaxation in glasses and its detection by multidimensional NMR. *Physical Review Letters*, 75(15):2851–2854, 1995.
- [126] R. G. Palmer, D. L. Stein, E. Abrahams, and P. W. Anderson. Models of hierarchically constrained dynamics for glassy relaxation. *Physical Review Letters*, 53(10):958–961, 1984.
- [127] Joseph Klafter and Michael F. Shlesinger. On the relationship among three theories of relaxation in disordered systems. *Proceedings of the National Academy of Sciences*, 83(4):848, 1986.
- [128] J. J. Brey and A. Prados. Slow logarithmic relaxation in models with hierarchically constrained dynamics. *Physical Review E*, 63(2):021108, 2001.

- [129] Michael A. Lomholt, Ludvig Lizana, Ralf Metzler, and Tobias Ambjörnsson. Microscopic origin of the logarithmic time evolution of aging processes in complex systems. *Physical Review Letters*, 110(20):208301, 2013.
- [130] Eran Bouchbinder and JS Langer. Shear-transformation-zone theory of linear glassy dynamics. *Physical Review E*, 83(6):061503, 2011.
- [131] E. E. Ferrero, S. Bustingorry, and A. B. Kolton. Nonsteady relaxation and critical exponents at the depinning transition. *Physical Review E*, 87(3):032122, 2013.
- [132] Alejandro B. Kolton, Alberto Rosso, Ezequiel V. Albano, and Thierry Giamarchi. Short-time relaxation of a driven elastic string in a random medium. *Physical Review B*, 74(14):140201, 2006.
- [133] Stefanos Papanikolaou, Dennis M. Dimiduk, Woosong Choi, James P. Sethna, Michael D. Uchic, Christopher F. Woodward, and Stefano Zapperi. Quasi-periodic events in crystal plasticity and the self-organized avalanche oscillator. *Nature*, 490:517, 2012.
- [134] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality. *Physical Review A*, 38(1):364–374, 1988.
- [135] Ronald Dickman, Miguel A. Muñoz, Alessandro Vespignani, and Stefano Zapperi. Paths to self-organized criticality. *Brazilian Journal of Physics*, 30:27–41, 2000.
- [136] E. A Jagla, Francois P. Landes, and Alberto Rosso. Viscoelastic effects in avalanche dynamics: A key to earthquake statistics. *Physical Review Letters*, 112(17):174301, 2014.
- [137] Matthias Schröter, Daniel I. Goldman, and Harry L. Swinney. Stationary state volume fluctuations in a granular medium. *Physical Review E*, 71(3):030301, 2005.
- [138] Nicholas A. Pohlman, Benjamin L. Severson, Julio M. Ottino, and Richard M. Lueptow. Surface roughness effects in granular matter: Influence on angle of repose and the absence of segregation. *Physical Review E*, 73(3):031304, 2006.
- [139] Randy Back. Frictional effects on mass measurements in a column of glass beads. *Granular Matter*, 13(6):723–729, 2011.
- [140] Li-Tsung Sheng, Wei-Ching Chang, and Shu-San Hsiau. Influence of particle surface roughness on creeping granular motion. *Physical Review E*, 94(1):012903, 2016.
- [141] Katalin A. Gillemot, Ellák Somfai, and Tamás Börzsönyi. Shear-driven segregation of dry granular materials with different friction coefficients. *Soft Matter*, 13(2):415–420, 2017.
- [142] Emily E. Brodsky, Jacquelyn J. Gilchrist, Amir Sagy, and Cristiano Collettini. Faults smooth gradually as a function of slip. *Earth and Planetary Science Letters*, 302(1):185–193, 2011.

- [143] François Renard, Karen Mair, and Olav Gundersen. Surface roughness evolution on experimentally simulated faults. *Journal of Structural Geology*, 45:101–112, 2012.
- [144] Pasi Tolppanen, Ove Stephansson, and Lars Stenlid. 3-D degradation analysis of railroad ballast. *Bulletin of Engineering Geology and the Environment*, 61(1):35–42, 2002.
- [145] T. Aste, M. Saadatfar, and T. J. Senden. Geometrical structure of disordered sphere packings. *Physical Review E*, 71(6):061302, 2005.
- [146] Chengjie Xia, Kuan Zhu, Yixin Cao, Haohua Sun, Binqun Kou, and Yujie Wang. X-ray tomography study of the random packing structure of ellipsoids. *Soft Matter*, 10(7):990–996, 2014.
- [147] Richard P. Jensen, Peter J. Bosscher, Michael E. Plesha, and Tuncer B. Edil. Dem simulation of granular mediastructure interface: effects of surface roughness and particle shape. *International Journal for Numerical and Analytical Methods in Geomechanics*, 23(6):531–547, 1999.
- [148] Tamás Börzsönyi and Ralf Stannarius. Granular materials composed of shape-anisotropic grains. *Soft Matter*, 9(31):7401–7418, 2013.
- [149] François Blateyron. The areal field parameters. In Richard Leach, editor, *Characterisation of Areal Surface Texture*, pages 15–43. Springer, Berlin, Heidelberg, 2013.
- [150] M. M. Koura and M. A. Omar. The effect of surface parameters on friction. *Wear*, 73(2):235–246, 1981.
- [151] E. S. Gadelmawla, M. M. Koura, T. M. A. Maksoud, I. M. Elewa, and H. H. Soliman. Roughness parameters. *Journal of Materials Processing Technology*, 123(1):133–145, 2002.
- [152] Pradeep L. Menezes, Kishore, and Satish V. Kailas. Influence of roughness parameters on coefficient of friction under lubricated conditions. *Sadhana*, 33(3):181, 2008.
- [153] D. J. Durian, H. Bideaud, P. Düringer, A. Schröder, F. Thalmann, and C. M. Marques. What is in a pebble shape? *Physical Review Letters*, 97(2):028001, 2006.
- [154] Gabor Domokos, Douglas J. Jerolmack, Andras Á Sipos, and Ákos Török. How river rocks round: Resolving the shape-size paradox. *PLOS ONE*, 9(2):e88657, 2014.