



HAL
open science

Avalanche dynamics of a model granular fault

Victor Levy Dit Vehel

► **To cite this version:**

Victor Levy Dit Vehel. Avalanche dynamics of a model granular fault. Physics [physics]. Université de Lyon, 2021. English. NNT : 2021LYSE1180 . tel-03566461

HAL Id: tel-03566461

<https://theses.hal.science/tel-03566461>

Submitted on 11 Feb 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



N° d'ordre NNT : 2021LYSE1180

THÈSE de DOCTORAT DE L'UNIVERSITÉ DE LYON

Opérée au sein de

l'Université Claude Bernard Lyon 1

Ecole Doctorale N° 52

Physique et Astrophysique de Lyon

Spécialité de doctorat : Physique

Soutenue et obtenue le 28/09/2021, par :

Victor Levy dit Vehel

Avalanche dynamics of a model granular fault

Devant le jury composé de :

Axelle AMON

Maîtresse de conférences, Université de Rennes

Xiaoping JIA

Professeur des Universités, ESPCI

Elsa BAYART

Chargée de recherche CNRS, ENS de Lyon

Jérôme WEISS

Directeur de recherche CNRS, Institut des Sciences de la Terre

Loic VANEL

Professeur des Universités, Université Lyon 1

Osvanny RAMOS

Maître de conférences, Université Lyon 1

Rapporteuse

Rapporteur

Examinatrice

Examineur

Examineur

Directeur de thèse

Acknowledgements

J'ai tendance à penser que l'on se construit par ce que l'on vit, et ces presque 4 années passées à l'ILM auront été particulièrement formatrices à cet égard – scientifiquement bien sûr, et humainement. Si la lecture d'article aura (un peu) contribué à me faire grandir du point de vue scientifique, c'est bien les personnes que j'ai croisées pendant ma thèse qui ont apportés les plus grandes contributions à mon évolution.

Puisqu'il s'agit des remerciements d'un manuscrit de thèse, je vais commencer par remercier ceux qui ont permis sa soutenance : mon jury de thèse. Merci donc à Loic Vanel, Jérôme Weiss et Elsa Bayart pour avoir accepté de faire parti de mon jury, pour avoir fait le déplacement et pour m'avoir écouté puis posé des questions (sans me la donner). Des remerciements particulier sont à adresser à Axelle Amon et Xiaoping Jia qui auront également joué le rôle de rapportrice et rapporteur.

Pour pouvoir soutenir, il faut d'abord écrire et c'est bien cette étape qui aura été pour moi la plus éprouvante. Osvanny pourrait légitimement être cité dans chaque paragraphe de ces remerciements mais c'est lors de la dernière ligne droite que sa dévotion en tant qu'encadrant de thèse aura été la plus forte. Merci pour avoir été disponible pour des relectures à presque toute heure du jour et de la nuit, pour m'avoir pointé (de nombreuse fois) dans la bonne direction tout en m'ayant laissé la liberté d'explorer mes propres idées, et pour m'avoir fait confiance toute ces années pour travailler sur ton bébé que j'ai maintenant un peu adopté. Ma gratitude ne se limite bien sûr pas au plan professionnel. Osvanny sera vite devenu et j'espère restera longtemps un ami sur lequel j'ai pu compter, pour m'aider a déménager par exemple, ou me prodiguer des conseils plein de bon sens (“viens travailler avant 16h, c'est mieux” par exemple).

L'équipe Liquides **et** Interfaces (et biophysique !) ne se limite pas a Osvanny cependant, et beaucoup d'autre membres de cette grande équipe ont contribué à rendre ces 4 dernières années (dans et hors les murs du labo) particulièrement

agréable. Sans ordre particulier ; les permanents ; Agnès, Anne-Laure, Antoine, Bruno, Catherine, Cécile, Gilles, Jean, Marie, Mathieu, Olivier, Sylvain (les deux) et Rémy, les autres thésards ; Malèke, Akash, Alexis, Antony, Aymeric, Grégoire, Louis, Romain, Ong, Saba et Yedhir, la team stagiaire ; Julie, Frédérick, Gaétan, Guillaume, Camille et Manon. Je ne peux pas non plus oublier l'équipe de 12 heures, 37 minutes et 30 secondes : Arsouille, Clem, Elise, Tinj et Tomtom qui ont toléré mes discussions de nombreux midi, mais aussi pendant de nombreux dîners, soirées, sorties et même apéros zoom ! Merci pour votre bienveillance, votre soutien pendant la thèse, et votre amitié au delà.

Merci également aux amis hors laboratoire et à ma famille, qui auront contribué à ma bonne santé morale. Papa pour avoir été toujours disponible pour des questions techniques et des conseils de parcours, Maman pour son aide logistique permanente, Pierem, Kévin, Jules, Quentin, Romain, Tarik, Sora, Jordan, Fluo, Laurence, Victor, Julien, Nico et Benoit.

Enfin, Veronica, qui m'aura beaucoup aidé moralement pendant des journées, des soirées et des nuits de rédaction éprouvantes. Merci de m'avoir permis de finir ma thèse tout court, et encore plus merci de m'avoir permis de la finir en ta compagnie.

Résumé

Le cadre de cette thèse est une expérience modèle reproduisant des comportements invariants d'échelle présentés par divers phénomènes naturels, les tremblements de terres entre autres.

Le modèle expérimental utilisé est constitué d'un granulaire bidisperse, fait dans une matériau photoélastique, cisailé dans une cellule periodique cylindrique. Lorsque que le granulaire est cisailé, des émissions acoustiques se produisent. Nous appelons ces événements "labquakes". La géométrie de la cellule autorise une cisaillement sans limite, ce qui nous permet d'obtenir une quantité très élevé de données et consituter des statistiques riches. Le suivi du système est assuré par un grand nombre de techniques. 6 capteurs acoustiques enregistrent les sons produit par le granulaire, 2 capteurs de force mesurent les couples résistant au cisaillement, 2 capteurs de position le volume de l'expérience et 24 caméras assurent un suivi de la position des grains, ainsi que des chaines de forces grâce à la photoélasticité du matériau.

Une partie de la thèse se concentre sur les relations entre changements soudains de volume et de couple résistant, deux grandeurs présentant une relation non trivial. Cette relation est d'abord décrite, puis expliquée par des propriétés géométrique des chaines de force. Dans un second temps, une méthode de localisation des sources acoustiques dans le milieu granulaire est présentée. Cette méthode, basée sur la localisation hyperbolique, servira dans de futurs travaux pour corrélér les caractéristiques mesurées d'un événement aux propriétés locales de son lieu d'emission. Finalement, des résultats préliminaires sur la prédiction des événements assisté par machine learning seront présentés.

Abstract

The framework of this thesis is a model experiment which reproduces scale invariant behaviors displayed by various natural phenomena, earthquakes among others.

The experimental setup consists of a bidisperse granular medium made of a photoelastic material, sheared in a cylindrical periodical cell. When the granular is sheared, acoustic emissions occur. We call these events “labquake”. The geometry of the cell allows for unrestricted shearing, which allows us to obtain a very high amount of data and build up rich statistics. The monitoring of the system is ensured by a large number of techniques. 6 acoustic sensors record the sounds produced by the granular material, 2 force sensors measure the resisting torque, 2 position sensors measure the volume of the experiment and 24 cameras monitor the position of the grains, as well as the force chains thanks to the photoelasticity of the material.

A part of the thesis focuses on the relationship between sudden changes in volume and resistive torque, two quantities with a non-trivial relationship. This relationship is first described and then explained by geometrical properties of force chains. In a second step, a method for localizing acoustic sources in granular media is presented. This method, based on hyperbolic localization, will be used in future works to correlate the measured characteristics of an event to the local properties of its emission site. Finally, preliminary results on machine learning assisted event prediction will be presented.

Contents

1	Introduction	1
1.1	Earthquakes and Scale Invariance	2
1.1.1	Avalanches phenomenons	2
1.1.2	Consequences of heavy-tailed distributions	3
1.2	Modeling earthquakes	4
1.2.1	Geo-mechanical approaches	4
1.2.2	The contribution of complexity	7
1.2.3	Analog experiments	8
1.2.4	statistical laws of earthquakes	10
1.2.5	Labquake project and thesis goals	15
2	LabQuake Experimental setup and methods	17
2.1	Experimental apparatus	18
2.2	Acoustic and mechanical measures	21
2.2.1	Acoustic sensors	21
2.2.2	Mechanical sensors	22
2.3	Imaging setup	25
2.3.1	Image fusing techniques	27
2.3.2	Image exploitation	35
3	Relevance to earthquakes	39
3.1	Energy distribution	40
3.1.1	Distribution of the different event types	40
3.1.2	Robustness across acoustical sensors	41
3.1.3	Coherence of the definitions	41
3.2	Omori law	43
3.3	Inter-event time distribution	45
3.4	Changes in experiment behavior	47

4	Probing the structure	51
4.1	Evolution of the volume	52
4.1.1	Height discontinuity events	52
4.1.2	Augmenting labquakes properties	54
4.1.3	Force chain orientation as a likely cause	56
4.2	Evaluation of the force network fractal dimension using images	62
4.2.1	Minkowski–Bouligand dimension	62
4.2.2	Application to thresholded images	62
4.3	Acoustic-based fractal dimension	67
4.3.1	propagation model	67
4.3.2	A simpler case: sound propagation in air	69
4.3.3	Application to our experiment	71
5	Acoustic Location of Events	75
5.1	Acoustic location and hyperbolic methods	76
5.2	Setup and method	78
5.2.1	Experimental design	78
5.2.2	Ballistic propagation model	80
5.2.3	Time-of-Arrival Measurement	82
5.3	Evaluating source likelihood	83
5.3.1	Single branch likelihood	83
5.3.2	Agglomerated likelihood	84
5.4	Event location performance	86
5.4.1	Border impacts	86
5.4.2	Bulk sources	88
5.4.3	Limitations and Future development	91
6	Event Prediction Attempts	93
6.1	Signal-average based precursors	94
6.1.1	Dilation precursors	94
6.1.2	Force precursors	98
6.2	Combining precursors with machine learning	101
6.2.1	Decision tree and random forest	102
6.2.2	Defining and designing goals	104
6.2.3	Performance	107
6.3	Perspectives	112
	Conclusions	113

A	Power-law exponent influence	115
B	Polariscope output intensity	117
C	Random forests parameters	121
C.1	Feature definitions	121
C.2	Feature normalization	122

CONTENTS

In case this document is read on a digital (PDF) version, every citation reference is clickable, pointing the reader to the bibliographic reference.

In turn, every citation in the bibliography contains a hyperlink (“cited on page X”), sending the reader back to where the work was cited.

Chapter 1

Introduction

Contents

1.1 Earthquakes and Scale Invariance	2
1.1.1 Avalanches phenomenons	2
1.1.2 Consequences of heavy-tailed distributions	3
1.2 Modeling earthquakes	4
1.2.1 Geo-mechanical approaches	4
1.2.2 The contribution of complexity	7
1.2.3 Analog experiments	8
1.2.4 statistical laws of earthquakes	10
1.2.5 Labquake project and thesis goals	15

1.1 Earthquakes and Scale Invariance

1.1.1 Avalanches phenomenons



Figure 1.1: From left to right; 2011 Tohoku earthquake aftermath (©37frames), Snow avalanche (©Archive SLF) and 1929 market crash headlines (©The Halifax Chronicle)

Even though the word “avalanches” call to mind disastrous snow slides [1], a broader reading of the term will include many more phenomena, such as market crashes [2–5], species evolution [6], solar eruptions [7, 8], various social interactions that can be modeled by networks [9], granular materials [1, 10–17], sub-critical fracture [18, 19] and earthquakes [20–23]. Avalanches occur in out-of-equilibrium systems where small events (snowfall, stock sell-off) may trigger other events, leading to a chain reaction. Such chain reaction can span over the entire system and lead to catastrophic disasters... or they may not. In some systems, avalanches can stop at any point before their largest possible scale: the skier dropping its glove in the snow will not cause a deadly avalanche. When an avalanche phenomenon covers a large range of (abstract) scales, it is useful to study and characterize its distribution. In many cases, including the examples mentioned above, the probability of stopping at scale s will follow a power law of the form $P(s) \sim s^{-a}$, with a a positive real number. For $a = 1$ for instance, it means scales ten times bigger will be ten times rarer as well. In the case of earthquakes, the distribution of their energies follows a power law and the observed exponent for worldwide statistics is $a \sim 1.66^1$, and is closely related with the Gutenberg-Richter law, linking magnitude and the cumulative quantity of earthquakes [24].

The term avalanche intuitively carries the idea of chain reactions, as domino effects or rolling snowballs. However, a power-law distribution of sudden events, usually called avalanches, may also be a consequence of discrete events taking place in a fractal landscape, as for example, discrete failure events within a power-law

¹This can vary when looking at particular regions or time periods.

distribution of rupture thresholds.

An avalanche process described by a power-law distribution will have the additional property of being scale-invariant: no particular scale of avalanche will come out as particular or typical. This is specific to this distribution, and if we compare it to other ones, like the bell curve describing the average human size, we notice a typical value of 1.65 m.

Let us now discuss the practical effects of power-law distributions.

1.1.2 Consequences of heavy-tailed distributions

To get a feel for this sort of relationship and its consequences, let us imagine the two following probability distributions, a power-law and a normal distribution:

$$P_{PL}(s) \sim s^{-a} \quad (1.1)$$

$$P_N(s) \sim \exp(-(s/\sigma)^2) \quad (1.2)$$

Now, consider how the probability changes with a doubling of the event scale:

$$P_{PL}(2s) \sim 2^{-a} s^{-a} \quad (1.3)$$

$$P_{PL}(2s) \sim 2^{-a} P_{PL}(s) \quad (1.4)$$

and:

$$P_N(2s) = \exp(-4(s/\sigma)^2) \quad (1.5)$$

$$P_N(2s) = (P_N(s))^4 \quad (1.6)$$

While the power-law power probability is simply scaled down by a multiplicative factor, the normal distribution gets raised to the power of four. More generally, power laws converge much slower to 0 for larger and larger event sizes. The impossibility to asymptotically bound power-laws by an exponential makes them “heavy-tailed”, as illustrated in figure 1.2. For comparison purposes, let us give some context to these distributions:

- The Maxwell-Boltzmann distribution ($P(x) \propto x^2 e^{-x^2}$) describes the speed of molecules in an ideal gas,
- the normal distribution ($P(x) \propto e^{-(x-1)^2}$), also called the bell curve, describes many things including average human size,
- the exponential distribution ($P(x) \propto e^{-x}$) describes the time between radioactive disintegration,

- a power law of exponent 1.66 corresponds to the distribution of earthquake energies,
- a power law of exponent 2.3 corresponds to the initial mass function in stellar systems [25].

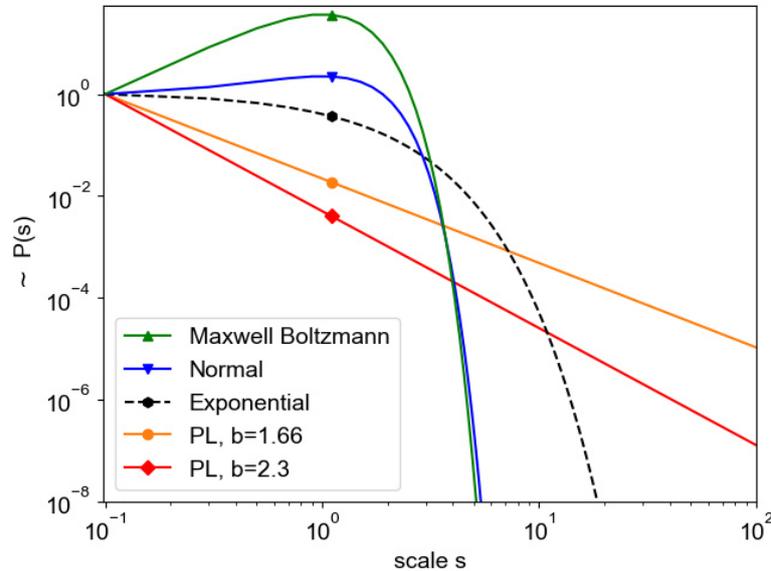


Figure 1.2: Several probability distributions in log-log scale. The dashed blacked line represent the boundary above which distributions are considered heavy tailed.

In all distribution, very high values are very rare. But in heavy-tailed distributions, higher values decrease in probability at a slow rate, making extreme samples more common. In the case of earthquakes, the consequences can be devastating.

1.2 Modeling earthquakes

Earthquakes always have been a part of human life, but their systematic study was propelled after the great 1906 San Francisco earthquake, officially causing over 3000 casualties [26] – about 1% of the city’s population at the time. The terrible consequences of earthquakes are strong incentives to develop our understanding of them. Several different approaches have been developed over the past century, some of which I will now succinctly detail.

1.2.1 Geo-mechanical approaches

One of the earliest theoretical attempts at understanding earthquakes dates back to 1911, by geophysicist H. F. Reid [27]. His *elastic rebound theory* stipulates that,

along a fault, each side is slowly deformed. At some point, the strain is too high and fracture happens. The elastic energy stored in the fault is released causing an earthquake. This process is schematized in figure 1.3.

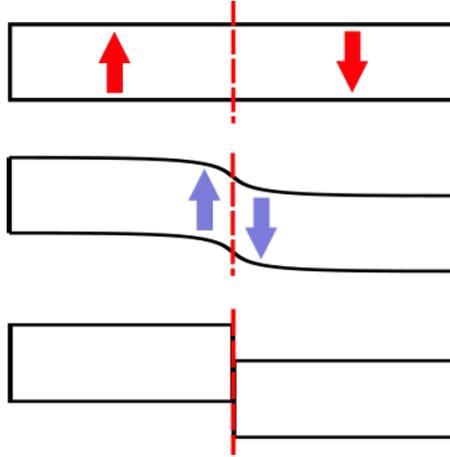


Figure 1.3: *top*: Stress is applied to each side of a block. *mid*: Elastic energy is stored in the deformed block. *bot*: The block breaks and energy is released.

This simple model was coherent with the observed displacement of the ground. However, a more precise analysis of the radiated waves [28], the low amount of stress released by an earthquake in relation to the available one, the high energies needed to shear over a fractured surface, and above all, the lack of healing required to generate a second earthquake at the same location and close in time to the first one, subsequently set “stick-slip” sliding mechanisms as a more plausible explanation of earthquakes [29]. Stick-slip is the name given to the jerking motion that can happen when an object slides against another. The initial steps are similar to what was proposed by Reid; the objects *stick*, elastic energy is accumulated until a limit is reached – but instead of a fracture, *slipping* occurs.

The proposition of stick-slip as a plausible mechanism for earthquakes by Brace and Byerlee in the mid-sixties [29] was quickly followed by friction models. First, analytical ones, like the spring-block model, [30, 31] and later phenomenological ones, like the *Rate and state friction* (RSF) model [32–34].

The most common analytical model is perhaps the *spring-block model*, developed originally by Burridge and Knopoff [30] and illustrated in figure 1.4. A bunch of blocks, sitting on a frictional surface, are linked by springs both to a common plate and to neighboring blocks. As the plate slowly moves tension is stored in the system until one block slips. This will elongate the neighboring spring, possibly causing the second block to slip as well, and eventually generating a chain reaction, or avalanche.

The *Rate and state friction* (RSF) brings a phenomenological description of the

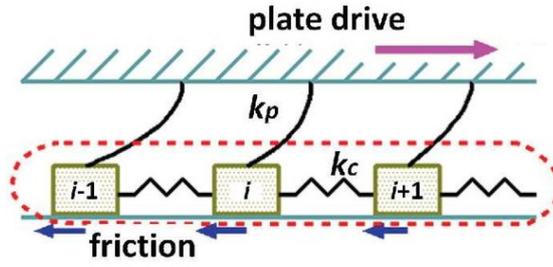


Figure 1.4: Spring-block model. (from [35])

frictional dynamics [32–34]. It can be formulated as follow:

$$\mu = \mu^* + A \ln \frac{V}{V^*} + B \ln \frac{\theta V^*}{D_c} \quad (1.7)$$

with μ and μ^* the current and characteristic coefficient of friction, V and V^* the current and characteristic slip rates, D_c the characteristic slip length, θ a state variable, and A , B coefficient parameters of the model.

RSF has also been used in conjunction with spring-blocks to model the friction of each block [36], successfully capturing the multi-scale behavior of earthquakes and some time-correlation properties like aftershocks and memory effects (figure 1.5).

Besides the facts presented by Brace and Byerlee [29] in favor of a stick-slip description of earthquakes, the fact that subcritical fracture displays a jerky dynamic with a power-law distribution of event sizes has been continuously used as an analog phenomenon. We can cite relevant works on subcritical fracture and earthquakes ranging from the sixties [37, 38] to very recent ones [19, 39, 40]. The also recent works of Jay Fineberg [41] analyzing the onset dynamics of a frictional block that have shown that frictional sliding follows the same theoretical description as

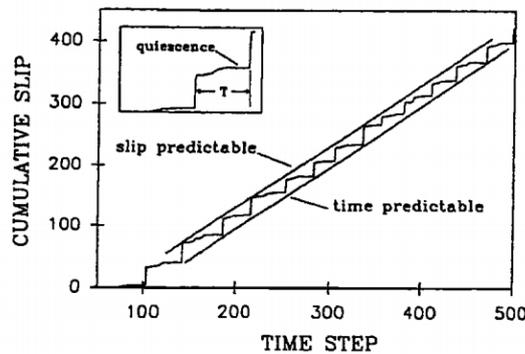


Figure 1.5: Cumulative slip of a spring-block cellular automaton, displaying memory effects and predictability. (from [36]).

a shear fracture. The authors' claim "*friction is fracture*" seems to conciliate both approaches to explain earthquake dynamics.

1.2.2 The contribution of complexity

The rise of complex systems

During the 1970s, the concept of *complex systems* was emerging [42], and took off in the 1980s – the Santa Fe Institute, dedicated to the study of such systems, was founded in 1984. *Complexity* characterizes systems where a large number of simple constituents, with simple interaction rules, are brought together. In such systems, (much) more complex behavior may arise through emergence. Complex systems occur in many fields, such as economics or mathematics. In physics, this concept was initially applied in condensed matter, regarding phase transitions [43, 44].

At the same time, geophysicists started to exploit them as alternative approaches to analyze earthquakes. Renormalization groups theory was used in asperity models [45], while some studies [46, 47] focused on the fractality [48] of seismicity and its relation to the size distributions of earthquakes. In 1989, Rundle proposed a thermodynamical approach of earthquakes [49], framing them as fluctuations around an average slipping dynamic [50]. A year later, the interactions between different faults were interpreted in the framework of chaos theory [51]. However, it was the concept of *Self-organized criticality* (SOC) [52] that drove the interest of the physics community into earthquakes, with the development of a large variety of models, and the interpretation of earthquake dynamics in the light of SOC [21, 22, 53–62].

Self-organized criticality

In 1987, Bak, Tang, and Wiesenfeld proposed the *BTW model* (known also as the sandpile model) [52]. In this toy model, 0 to 3 particles are placed on each tile of a regular, square grid. Particles are then dropped one by one on a random tile of the grid. When a tile has 4 particles, it “topples” and redistributes one particle to each of its neighboring tiles, which can topple as well. Particles falling off the edge of the grid are removed from the system. After a while of repeating this process, avalanches will happen at every step and are distributed following a power law of slope 1.29 [63]

The authors used this simple model to introduce the concept of Self-Organized Criticality (SOC). The existence of power-law distributions of event sizes, with the possibility of an event reaching the system size, was interpreted as a critical dynamic,

making an analogy with phase transitions [64]. However, in contrast to them, the absence of a direct tuning parameter designated these systems as self-organized. SOC had the ambition of explaining practically all phenomena displaying scale invariance, which attracted the interest of a large community and generated a large number of works in diverse disciplines.

The direct relevance of the SOC ideas to earthquake physics is rather limited, which we may say is expected given the simplicity of the ingredients and rules of the original model. However, those ideas sparked the interest of the physics community in earthquake-like dynamics and were a starting point for more suited models, like the Olami-Feder-Christensen (OFC) one, [65]. The OFC model is a cellular automaton translation to the spring-block model [30] and, therefore, it is much closer to the actual earthquakes than the original sandpile model [52]. The emergence of different experiments displaying scale-invariant behavior, often linked to earthquake statistics, is another relevant legacy of the SOC ideas.

1.2.3 Analog experiments

The complexity of earthquakes is such that capturing their dynamics in a theoretical model will always be an oversimplification. Finding a physical system sharing a common behavior is a less drastic simplification to approach earthquake dynamics. With this perspective, a large number of analog experiments have been set. Much like theoretical studies, many efforts were focused on fracture and friction. Early fracture-oriented studies [37, 38] have shown micro-fractures, occurring during rock compression tests, exhibited a multi-scale behavior. More recent works on subcritical fracture have also displayed power laws distributions similar to the G-R law as well as Omori-like relations [19, 39, 40].

Friction experiments have usually focused on shearing two surfaces against one another, usually blocks. Several materials have been used, such as acrylic [41] (more controllable) or actual rocks [66] (closer to geological materials). To further the resemblance with actual faults, granular material has been added in some experiments between the sheared surface. One example, from [67]², is shown in figure 1.6.

In parallel, the importance of other mechanisms was being investigated, such as fault healing [70–73] or water lubrication [74, 75]. Contemporary experiments have included these aspects using ice [76–78]. In these experiments, a solid disk is rotating at the surface of a water bath. As the experiment is conducted under freezing temperature, a layer of ice forms above the water. This provides both

²The setup was first used in [67], but the schematic here has been taken from [68].

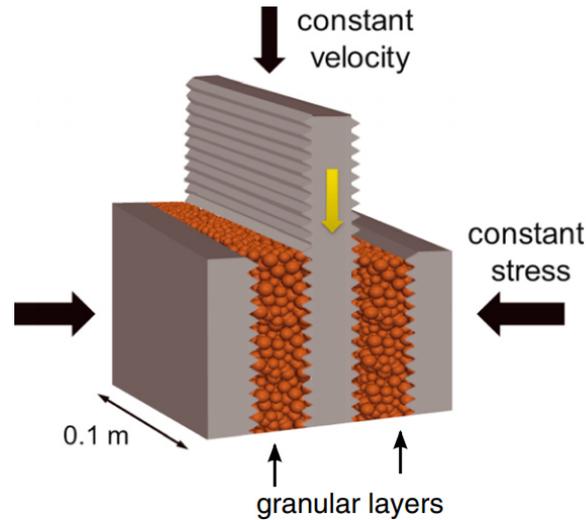


Figure 1.6: Experimental setup used in [67–69], among others works. It consists of a plate slowly driven ($1\text{--}100\ \mu\text{m/s}$ in [67]) between two others, with both gaps filled with a monodisperse granular made of glass beads. (*From [68]*)

lubrication, thanks to water, as well fault healing, when the ice forms back after a rupture event.

It is clear the complexity of earthquakes is unlikely to be capture by a single model or experimental setup. Nucleation of earthquakes, propagation along a fault, interactions between faults, statistical properties are all very different phenomena happening at vastly different scales.

Different complementary models and analog experiments working at different scales and aiming at diverse goals will be necessary to reach a more complete understanding of earthquake physics. Our work focuses primarily on earthquake statistics, with a system comprised of a granular material. Other works have used similar approaches [79–81], with one system in particular [82], a bi-axially sheared 3D granular material which replicates each of the statistical law we have discussed earlier, but with a twist. In this system, the total strain plays the role of time, showing any memory effect must be present in the structure.

The labquake project, the framework of my thesis, falls in this line of research. Our main experiment, which will be detailed in chapter 2, is shown in figure 1.7. Succinctly, a photoelastic granular material arranged in a 2D layer is continuously sheared under constant load. During the shear, acoustic emissions occur and have been shown to exhibit statistical features found in earthquakes [83]. Before presenting this agreement in detail in chapter 3, I will now explain each statistical law our

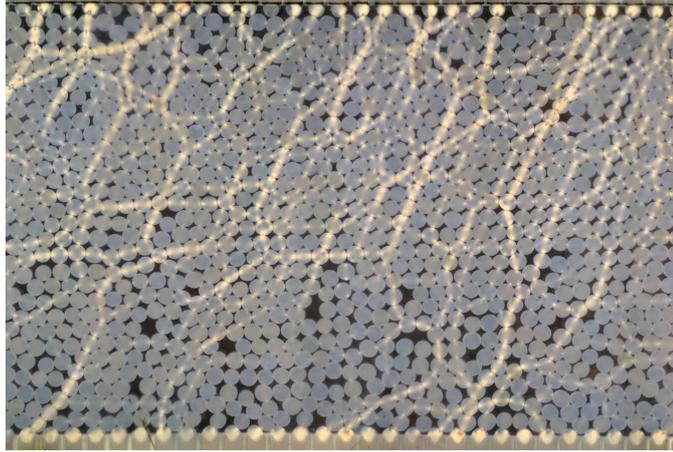


Figure 1.7: The main experimental setup of my thesis. It consists of a sheared 2D layer of disks.

experiment reproduces.

1.2.4 statistical laws of earthquakes

The incursion of many physicists into earthquakes dynamics, propelled by the ideas of Complexity, brought fresh ideas into the field. However, it is quite common to find wrong interpretations about the statistical relation describing seismicity in physics papers. Therefore, as a good practice, we encourage the community to use real earthquake data to directly confront their results on models or analog experiments.

All the statistical analyses in this section are performed on real earthquakes data. A very recent (2019) earthquake database, The QTM Seismicity Catalog, has been used [84]. This catalog range from 2008 to 2017 and focuses on southern California.

Energy distribution

Earthquakes display many statistical features, explained through several phenomenological relations. In our project we have focused on three main relations: The GR law, the Omori law and the distribution of inter-event times. . The first was eluded to at the beginning of this introduction and relates to the distribution of earthquakes energy. the **Gutenberg-Richter law** [24] describes the probability distribution to observe an event of magnitude at least M_0 :

$$P(M \geq M_0) \sim M_0^{-b} \quad (1.8)$$

with b empirically measured to be close to 1 (with some variation [85]). The physicist is however more used to talking in term of non-integrated probability distribution,

and energy instead of magnitude. To convert the GR law into a more familiar form, one can use the following relationships, the first from [86] and the second being the fundamental theorem of calculus:

$$M \sim 2/3 \log(E) \quad (1.9)$$

$$P(X_0) = \frac{d}{dX} P(X \geq X_0) \quad (1.10)$$

One can finally derive a probability density function for the energy:

$$P(M) \sim M^{-5/3} \sim M^{-1.67} \quad (1.11)$$

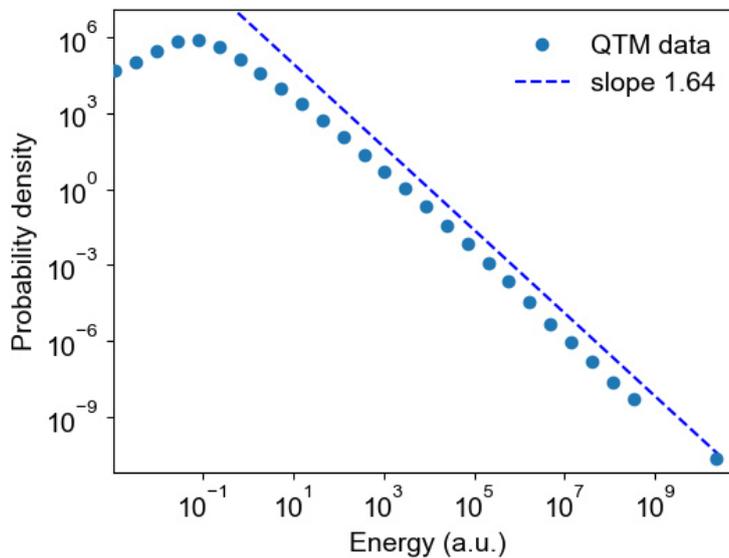


Figure 1.8: Energy distribution of earthquakes, with a linear regression of slope 1.64.

Figure 1.8 represents the energy distribution law under the form of equation 1.11 accompanied with a linear fit of slope 1.64. Given the relevance of the exponent values (explained in appendix A), specific methods have been developed to perform an optimal fit of the data. The method and python package developed by [87] was used for this linear regression. It will also be used throughout this thesis.

Aftershock rate decay

The second³ statistical feature we have focused on is the **Omori law** [88] (sometimes the Utsu-Omori law [89]) and relates to aftershocks. After large earthquakes, the seismic activity is higher than usual, causing additional quakes smaller than the

³Even if listed as second here, it was actually described before the GR law, in 1894.

mainshocks – these are called “aftershocks”. The Omori law stipulates the rate of aftershock n at time t after the mainshock decrease as:

$$n(t) = n_0 + \frac{k}{(c+t)^p} \quad (1.12)$$

with n_0 , k , c and p empirical constants⁴. k is a global activity rate multiplier, n_0 is the background rate and c acts as a time lag. The most important parameter, the p exponent, dictates the decay rate of $n(t)$ and is typically close to 1.

To compute the aftershock rate, a mainshock must first be defined. The earthquake I will present the aftershock sequence of is the 2010 Baja California earthquake, which happened on April 4th, 2010. A section of the QTM catalog after this mainshock is shown in figure 1.9 while the aftershock rate is shown in figure 1.10.

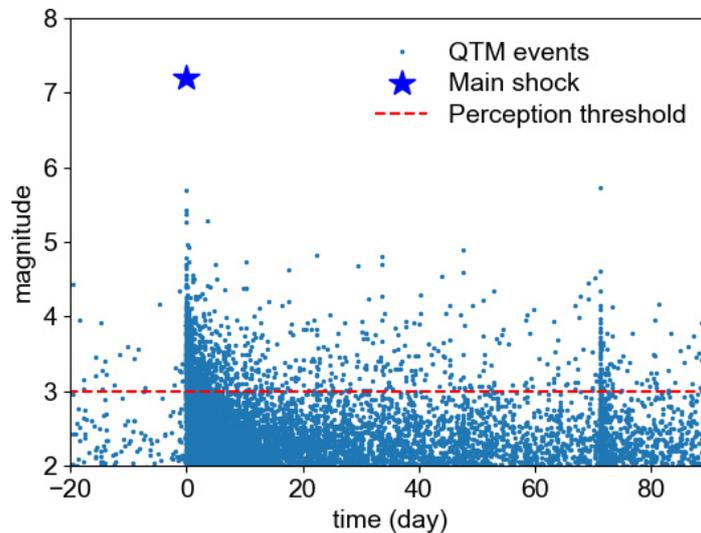


Figure 1.9: An aftershock sequence following a main quake which happened on April 4th, 2010. Note the magnitude axis is cropped to exclude the smallest event for clarity. Another fairly large quake (M6) can be seen around 70 days, with its own, smaller, increase in seismic activity.

If the concept of aftershock may make intuitive sense at first, it is actually far from obvious to properly define. Since earthquakes can *a priori* happen anytime at any size, how can an independent event be differentiated from one that is the consequence of the previous one? Depending on what is or what is not an aftershock, measures and interpretations may vary and several models and techniques have been deployed to tackle this issue [90–95]. A very crude method to separate aftershocks from the background activity is to change the detection threshold. The effect of the threshold

⁴Note that while they are called “constants”, they depend on the period and seismic region studied.

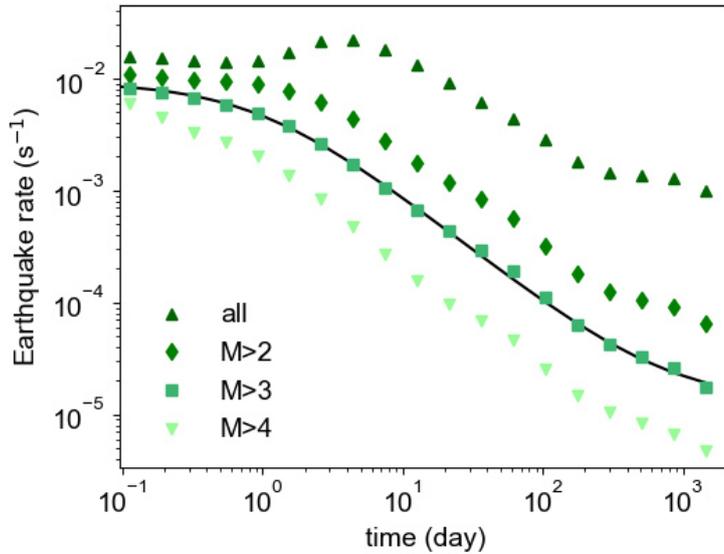


Figure 1.10: Seismic rate decay obtained by selectively counting earthquake above a range of threshold, following the main shock shown in figure 1.9. The black lines (solid and dashed) are not fit, but visual clues. The decay is longer and more pronounced for higher thresholds since the respective baseline activities are lower.

choice is illustrated by figure 1.10. The beginning of the curves (between 10^{-1} and 10^0 days) are less meaningful as the aftershock rate is harder to compute right after the mainshock, smaller events are hidden by larger ones. This is a phenomenon called *catalog incompleteness*. After a day, all curves decay following power laws with more or less well-defined slopes. As expected, the seismic rate decays to lower values the higher the threshold is. The decay rate itself does not change very significantly once any threshold is applied. A striking feature of this seismic rate slowdown is its timescale, ranging in the years.

Waiting times between events

The last property we will discuss is a rather recent one and comes from the physics community. It was described by A. Corral in [96]. To quantify earthquake recurrence at a given region, one can study the series of waiting time τ between events of energy above a given threshold E_0 , noted $\tau_{E \geq E_0}$. The distribution of $\tau_{E \geq E_0}$ can be rescaled for comparison between thresholds. For this purpose, let us call $\tau_{E \geq E_0}^*$ the average of $\tau_{E \geq E_0}$ and $R_{E \geq E_0} = 1/\tau_{E \geq E_0}^*$, respectively representing the average inter-event time and the event rate. We can then define the dimensionless waiting time $\theta = \tau_{E \geq E_0} / \tau_{E \geq E_0}^*$ and compute its rescaled probability distribution $PDF(\theta) \cdot R_{E \geq E_0}$, as is done for the QTM catalog in figure 1.11. At any threshold the data collapses well around a master curve in black of equation $f(\theta) = C \theta^{\gamma-1} \exp(-\theta^\delta / B)$ with

parameters $C=0.5$, $\gamma=0.67$, $B=1.58$ and $\delta=0.98$. This master curve consists of a slowly decreasing power law of exponent -0.33 for values below $\theta_{E \geq E_0} = 1$ and an exponential decay for values above. The later part of the curve indicates the inter-event time is highly unlikely to be very large, as this distribution is not heavy-tailed but exponentially bounded.

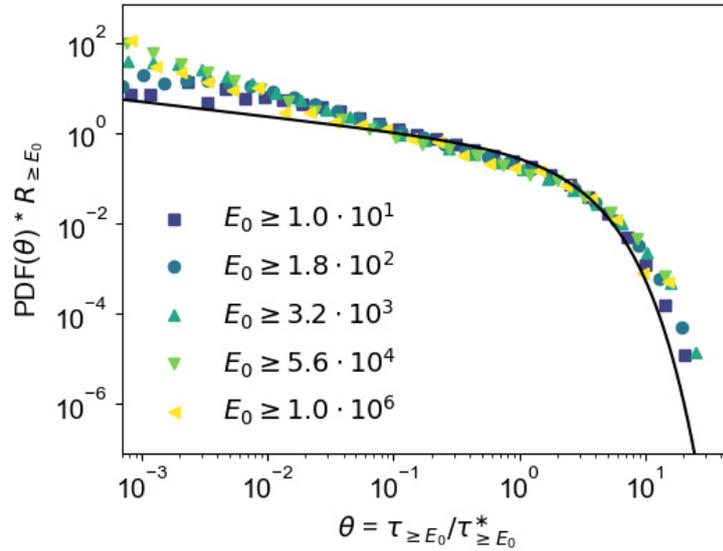


Figure 1.11: Normalized probability distribution of $\theta = \tau_{E \geq E_0} / \tau_{E \geq E_0}^*$ for different threshold E_0 . All distribution collapse perfectly on the universal function $f(\theta) = C \theta^{\gamma-1} \exp(-\theta^\delta / B)$.

Besides the remarkable universality across geographical regions and energy thresholds, an important feature of this master curve is its power-law beginning. If earthquakes were a completely memory-less process, like radioactive decay, we would only expect a decreasing exponential. This indicates there *is* memory in such processes. This weak memory effect is characterized by the fact the system appears to “remember” an event of energy E_0 for a duration equal to the average inter-event times of event larger than E_0 .

1.2.5 Labquake project and thesis goals

The labquake project, led by Osvanny Ramos, aims at developing a fully controllable analog experiment, capable of –with relatively simple ingredients– reproducing the statistical relations describing seismicity. Finding simpler systems experiencing the same dynamics may allow a better understanding of the physics that rules this behavior. More explicitly, long terms goals are tackling questions like the origin and robustness of the dynamics, the analysis of memory effects, and the possibilities of predicting catastrophic events in earthquake-like systems.

This model experiment, a continuously sheared photoelastic granular, will be presented in chapter 2 of this thesis. The experimental method used to make measurements will also be detailed, which includes acoustic, force, volume, and optical measurements.

The statistical relevance, which has now been established [83], will be detailed in chapter 3. Strong agreement with each of the three statistical laws presented before has been observed.

Chapter 4 will focus on structural characterizations of the granular assembly. First, a multi-scale behavior of the volume change of the granular will be presented. This new data augment our previous mechanical definition of labquakes, giving rise to a new classification of events. Quantitative analysis of the force chain network will give insight into this classification. The rest of this chapter will present a novel acoustic-based method to probe the inner structure of this force chain network.

Exploiting both images and acoustic data, chapter 5 will describe a technique developed to locate sound emission sources in small, compact, and 2D granular systems such as is ours. After detailing the theoretical foundation of the method, it will be demonstrated in two different cells.

Finally, preliminary attempts at prediction will be the topic of chapter 6. First, precursors based on distinct data will be listed and evaluated. Using basic machine learning, these precursors will then be merged and exploited to make an agglomerated prediction with hopefully better forecasting capability.

Chapter 2

LabQuake Experimental setup and methods

Contents

2.1	Experimental apparatus	18
2.2	Acoustic and mechanical measures	21
2.2.1	Acoustic sensors	21
2.2.2	Mechanical sensors	22
2.3	Imaging setup	25
2.3.1	Image fusing techniques	27
2.3.2	Image exploitation	35

2.1 Experimental apparatus

The main experiment of the labquake project and the one I studied (see figures 2.1 and 2.2) is a 2D cylindrical granular pile, confined in between two concentric fixed acrylic cylinders (*side walls*) and bounded by two rough circular rings (*top / bottom*). The rings are made using the same material as the grains and have 99 “teeth” made of half-cylinders of diameter $d=6.4$ mm with a periodic spacing of $\sqrt{2}d$ (~ 9.05 mm).

The granular material used in this experiment is a bidisperse mix of plastic disks. They are 4 mm in thickness and the diameters are 6.4 mm and 7 mm. To avoid crystallization, both diameters are used in equal proportions. With the goal to measure the stress in the granular using photoelasticimetry, a photoelastic material was required. Common choices include homemade particles using clear rubber or commercially available materials [97]. However, such material have low Young modulus (1~10 MPa) and would not be able to handle our confining pressure. Fortunately, we had access to an *Objet30* 3D-printer which can use the (proprietary) *Durus White 430* material, a UV-cured plastic with photoelastic properties, and a Young modulus in the range of 100 MPa.

A constant dead load placed over the top ring compresses the granular pile, of weight $7 \sim 40$ kg, depending on the experiment – but typically 27.5 kg. Using heavy weights ensures the vertical compression is well controlled and does not depend on spurious frictional interactions between the granular material and the sidewalls. The top ring is free to move vertically but not to rotate, while the bottom one is slowly rotated with a period of 18.33 h, quasi-statically shearing the granular pile with a linear velocity of 48.84 mm/h. Thanks to a lever and a force sensor, we measure the torque $\Gamma(t)$ applied by the granular pile on the top ring. Six piezoelectric pinducersTM are inserted regularly in the top ring and simultaneously record acoustic emissions. Placed above the top plate, two inductive positions sensors monitor its position $h(t)$. Finally, an array of 24 cameras driven by *Raspberry Pi 3B+* surrounds the setup to take full panoramic pictures. Except for the pictures, taken at 4 seconds intervals, all the other measurements are done at a rate of 100 000 samples per second, using a NI-USB-6366 card. The system is left to evolve for a typical duration of 24h.

Thanks to its particular geometry, the system can be sheared almost indefinitely, with strains larger than 100%. This allows for a very large amount of data to be collected. This is crucial to perform statistical analysis regarding large events, as

they occur very rarely. A shear band divides the structure into a very mobile zone, corresponding to a layer of about 10 grain diameters adjacent to the moving ring at the bottom of the system; and a zone that is almost static in the top part of the pile [83, 98, 99].

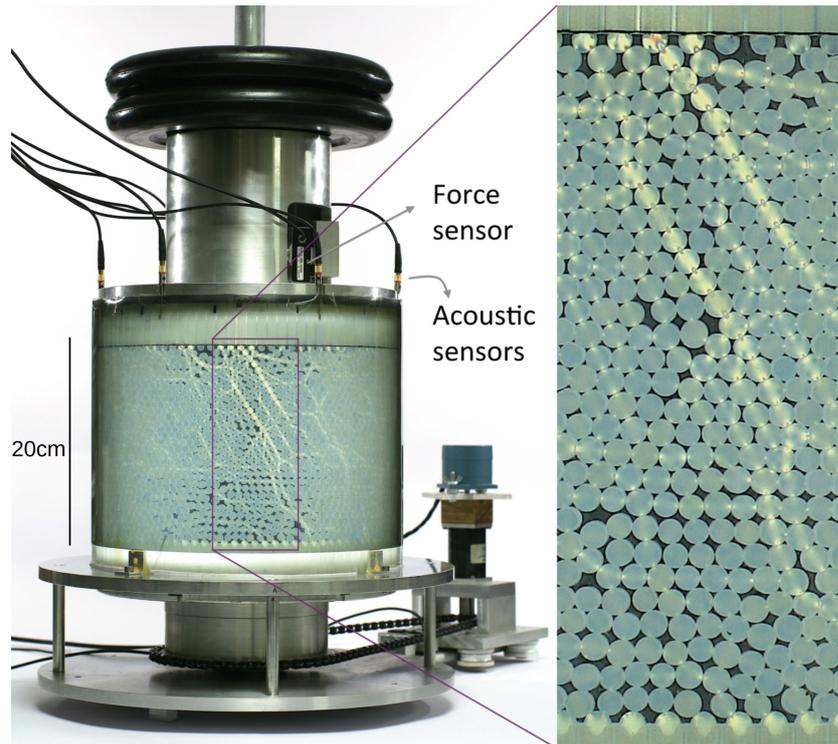


Figure 2.1: The main Labquake experiment. Six piezoelectric sensors record acoustic data, two force sensors measures the resistive torque on the top as well as the tension in the driving chain and two position sensors measure the height of the top plate (not seen in the picture). Force chains can be seen on the zoomed view of the granular.

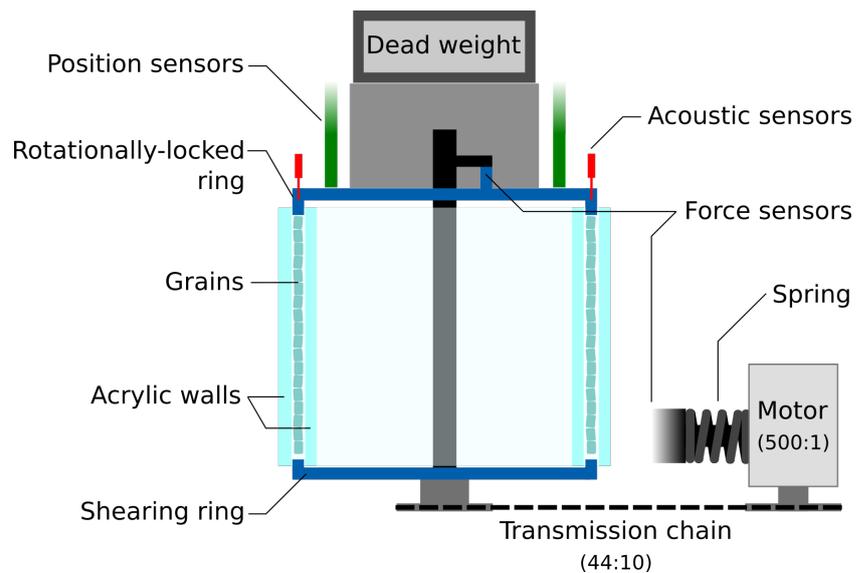


Figure 2.2: Schematic of the cylindrical shearing experiment shown in figure 2.1 as a cut-view.

2.2 Acoustic and mechanical measures

During the shear, grains rub and slide against each other and sudden re-arrangements of the structure occur. All of these processes can create acoustic emissions, as well as torque and volume discontinuities. In this section, I will describe the processing done on the acoustic and torque signals to define the core measurements used in the labquake project, and throughout my thesis. The position signal will be studied in the first half of chapter 4.

2.2.1 Acoustic sensors

Acoustic emissions are recorded by six piezoelectric pinducersTM (*VP-1.5 from CTS Valpey Corp*) regularly inserted in the top ring closing the experiment. These sensors have a rather flat response over a large band of frequencies (1kHz to 10MHz), containing typical frequencies produced by our system. To ensure good and constant acoustic coupling, silicon oil is added to the hole before inserting each sensor. Silicon oil has two advantages: it does not dry quickly and spreads very easily, ensuring the holes are well filled throughout an experiment.

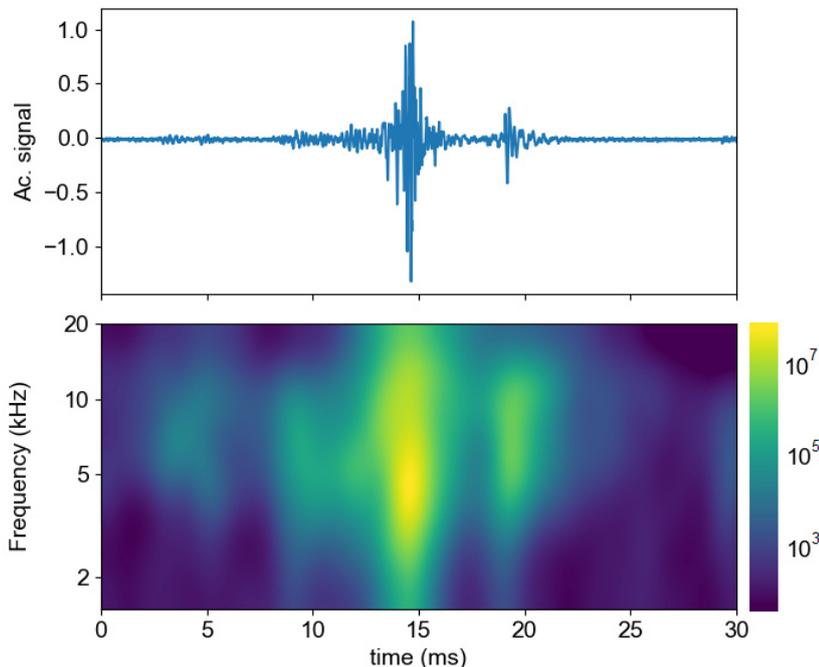


Figure 2.3: *Top*: recorded acoustic signal displaying a large acoustic event around $t = 15\text{ ms}$. *Bottom*: Smoothed scalogram of the above signal, with lighter shade indicating higher energy. The color is in log-scale.

To detect events in the acoustic signal, we perform a Discrete Wavelet Transform (*DWT*) on instants separated by $50\mu\text{s}$. 128 frequencies, logarithmically spaced

between $f_{min} = 1\text{kHz}$ and $f_{max} = 20\text{kHz}$, are analyzed. A python implementation (*tftb* on pip) of the Matlab Time-Frequency Toolbox was used to perform this spectral analysis. The obtained scalogram is then smoothed using a Gaussian blur of parameters $\sigma_t = 0.6\text{ ms}$ and $\sigma_f = 12\text{ bins}$ (respectively for the time and frequency axis). The resulting scalogram is shown in figure 2.3. Events are then detected on this smoothed time-frequency representation using a basic peak detection algorithm. Detecting a peak provides its timing, and the energy is defined as the peak value.

2.2.2 Mechanical sensors

Both mechanical signals, torque and position, consist of a smoothly varying part with short discontinuities, as illustrated on the top of figure 2.5 (in the next subsection). Discontinuities are detected and characterized through convolution with a normalized derivation filter D of the form:

$$\begin{aligned} D(t) &= -1/K && \text{for } t \text{ in } [-t_{f,2}, -t_{f,1}] \\ D(t) &= +1/K && \text{for } t \text{ in } [t_{f,1}, t_{f,2}] \\ D(t) &= 0 && \text{otherwise.} \end{aligned} \tag{2.1}$$

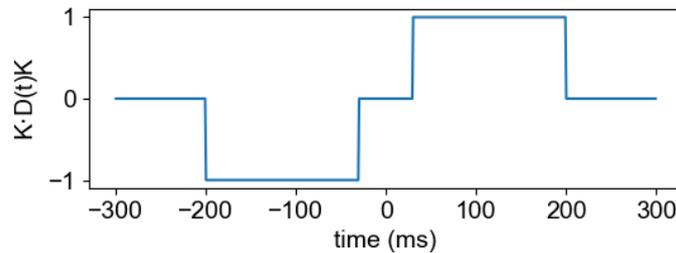


Figure 2.4: Shape of the filter D used to detect discontinuities in mechanical signals.

with $t_{f,1} = 30\text{ ms}$ and $t_{f,2} = 200\text{ ms}$ and $K = t_{f,2} - t_{f,1}$ a normalization constant. Choosing $t_{f,1} \neq 0$ creates a “gap” in the middle of the filter, as illustrated by figure 2.4. The reason for this is to account for the finite speed of the discontinuities. With this gap, we can measure only the change between before and after an event, and ignore the transient displacement.

Torque and mechanical energy

The torque signals are measured with *Interface SML-900N* force sensors. To obtain the torque, we multiply the force value by the lever length. By assimilating the setup as a torsion spring, the mechanical energy can be written as $E_{m,tot} = \frac{1}{2}K \cdot \Gamma^2$. The spring constant K is the composition of several different parts of the experimental setup; $K^{-1} = K_{gears}^{-1} + K_{chain}^{-1} + K_{sensors}^{-1} + K_{lever}^{-1} + K_{granular}^{-1}$. Since everything but the grains is made out of large pieces of steel or aluminium, we can approximate $K^{-1} = K_{grains}^{-1}$. Mechanical energy release events are then defined as $E_m \sim \Gamma_i^2 - \Gamma_f^2$ with Γ_i and Γ_f the initial and final torque values.

In figure 2.5, a torque signal (top) and its convolution with the previously described filter (bottom) is shown. 1D peak detection is then used to detect the energy and timing of events.

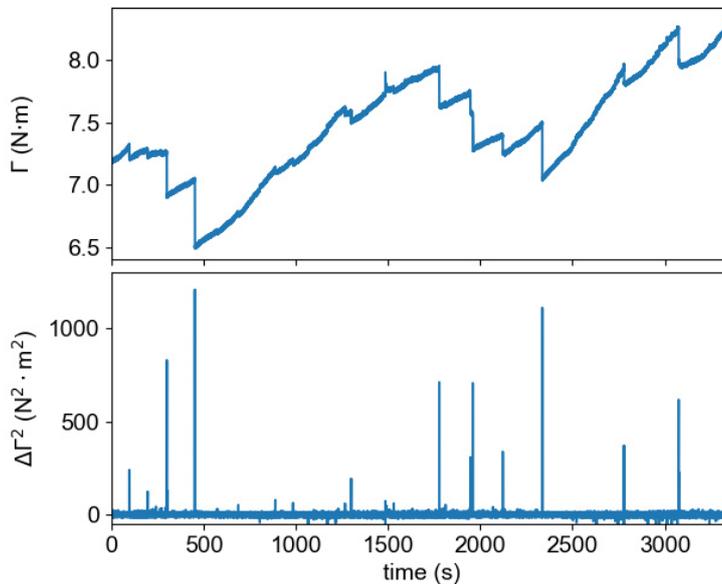


Figure 2.5: *Top*: the torque signal Γ . *Bottom*: smoothed derivative of Γ^2 .

Dilation

The two position sensors (*Baumer IPRM 12I9505/S14*) are suspended above the experiment, attached to controlled stepping motors. When a sensor is getting too close (or too far), its motor brings it back into an optimal distance range.

Figure 2.6 shows the first hour of a position signal alongside a zoom around a large drop. On the zoomed part, oscillations can be observed. These are due to the way the sensor is installed, suspended above the experiment. Large events shake the structure and the sensor can oscillate for a short while, acting as a dampened pendulum. However, these oscillations are easy to filter using the previously detailed filter.

Chapter 4 is dedicated to the study of the position signal.

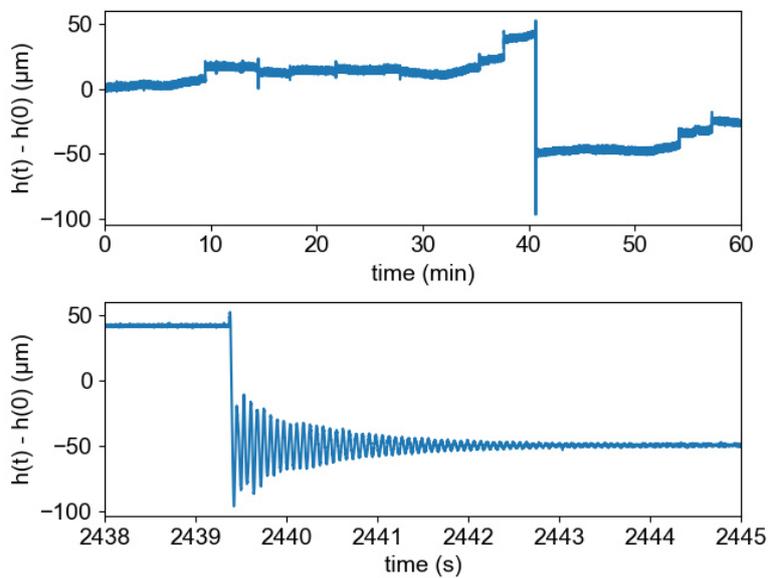


Figure 2.6: *Top*: the first hour of the position measurement. *Bottom*: a zoom of the large event right after 40 minutes.

2.3 Imaging setup

When joining the *LabQuakes* project during my Master’s internship and for the better part of the following year, my first task was to completely redesign the imaging system, which included both the instruments and the methods. The original setup, shown in figure 2.7, was comprised of 12 HD (1920×1080p) *Logitech* USB webcams. With 12 cameras, each has to observe at the very least a 12th of the cylinder, which is 30°. Since they were placed around 30 cm away from the surface, curvatures effects were strong and the sides of each camera region of interest (ROI) were very blurry. This setup was previously used for qualitative analysis, but it was not adapted for quantitative measurements.

I replaced the old optical system with 24 *Raspberry Pi 3B+*, each equipped with V2.1 camera modules (3280×2464 px), which can be seen in figure 2.8. The increased number of sensors reduces the angle each unit must record by two, to 15°. In addition, the new cameras are placed almost twice further as the previous ones. Finally, in the latest version of the setup, the polarizers used to measure photoelasticity are now placed in front of each camera, covering their sight, as illustrated in figure 2.9a. The previous solution, a large sheet covering the outer cylinder, introduced some defects (such as where the edges of the sheet meets), and was much more fragile. Thanks to these changes, current images (see figure 2.9b) are now of good enough quality for quantitative analysis.

Contrary to the torque sensor (for instance), which measures a property of the whole system, each image provides data related to only a fraction of the system. From the set of 24 independent images, we want to be able to measure system-wide properties. Schematically, two approaches can be considered:

- *measuring on individual images* → *fusing data*
- *fusing images* → *measuring on a fused panorama*

The second approach has several advantages: faster (smaller total pixel area), the curvature is corrected prior to measurements, and some analyses are easier to do on panoramic images (on force chains, for instance). For these reasons, this approach was chosen. Two methods were considered to fuse images, which I will now explain.

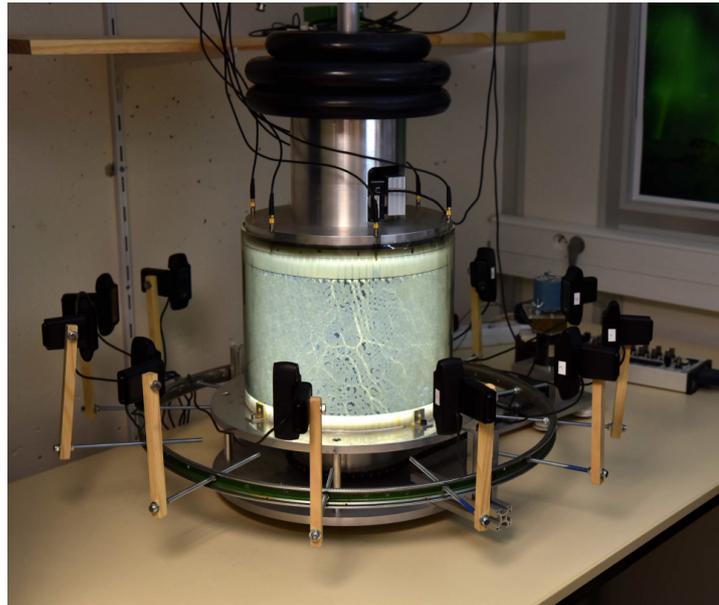


Figure 2.7: Former setup with 12 USB webcams placed at a distance of roughly 30cm.

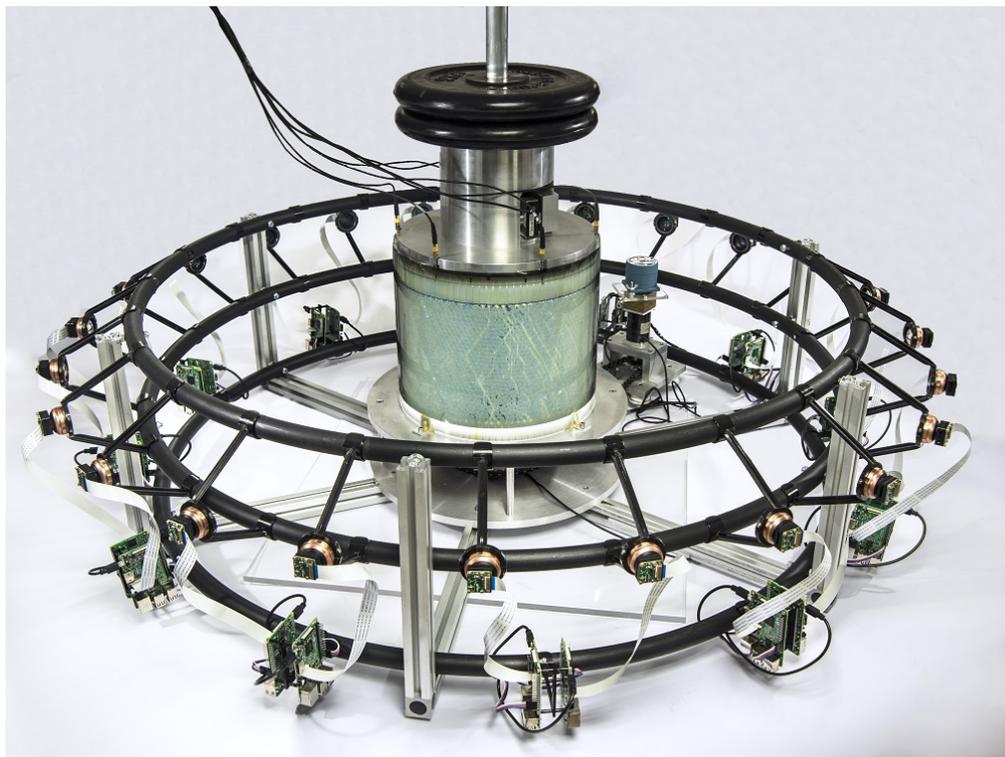
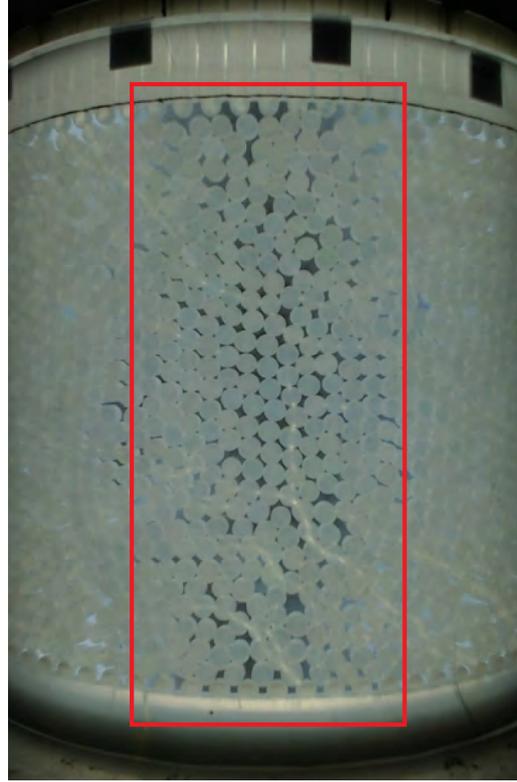


Figure 2.8: Imaging setup, forming a ring of 24 independent cameras placed at 50cm.



(a) Analyzing polarizers are now held by a frame, in front of each camera.



(b) Typical picture taken by a camera. The region of interest is boxed in red.

Figure 2.9

2.3.1 Image fusing techniques

The geometry of the optical system is fairly simple: 24 cameras are located on a circle, perpendicular to the setup axis. For this reason, the first approach to fuse neighboring images was to use our knowledge of the system geometry and use explicit theoretical formulas to straighten the images.

Theoretical approach

We observe a cylinder of know radius R , from a distance d , as schematized in figure 2.10. On the picture (*in the focal plane*) we observe a point \mathbf{P} of coordinates (x, y) , which corresponds to a real point P_m located on the cylinder surface of coordinates (x_m, y_m) . We want to compute P_m as a function of P . Note that the x_m coordinate is defined along the surface of the cylinder, such that $x_m = R\alpha$.

Using the angles and length as defined on 2.10, we can show that:

$$\alpha = -\beta + \arcsin\left(\frac{R+d}{R}\sin(\beta)\right) \quad (2.2)$$

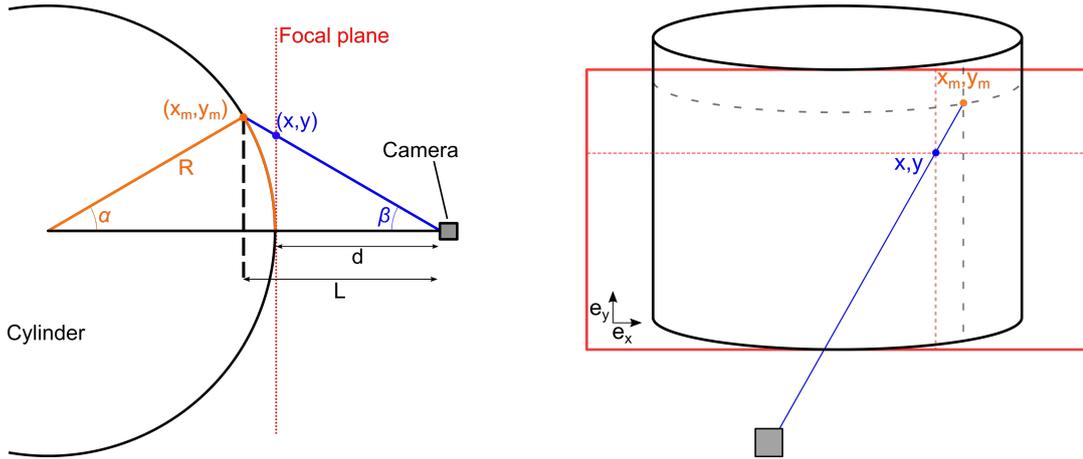


Figure 2.10: A top view of the model is represented on the left side, with (x_m, y_m) denoting the coordinates in curved geometry and (x, y) the coordinates within the picture. On the right side, a front view of the same model.

Where $\beta = \arctan(x/d)$. Let us call L to the distance between the intersection of the dashed and horizontal lines on the left part of figure 2.10. Two applications of the Thales theorem gives us:

$$y_m = \frac{yL}{d} \quad (2.3)$$

We can finally express both x_m and y_m as:

$$x_m = R\alpha \quad (2.4)$$

$$y_m = y \cdot \left(1 + \frac{R}{d}(1 - \cos(\alpha))\right) \quad (2.5)$$

A curved grid of observed points (x, y) (*dark blue*) is overlaid on a corresponding grid of real points (x_m, y_m) (*light blue*) on figure 2.11, alongside a vector field to show the transformation.

By applying this transform to our pictures, we can now “flatten” them and remove curvature effects. However, each image must still be aligned and fused with its neighbors. This procedure of fusing different data sets with different coordinates systems (here, neighboring images) is called *registration*. For this process to be accurate, it helps to have at least some overlap: points in each set describing the same object, to serve as a reference. An example is shown in figure 2.12, where two orthogonal laser scans are fused to form a more complete 3D view. In this example, a vertical edge of the building (among other details) can serve as a reference point for the alignment.

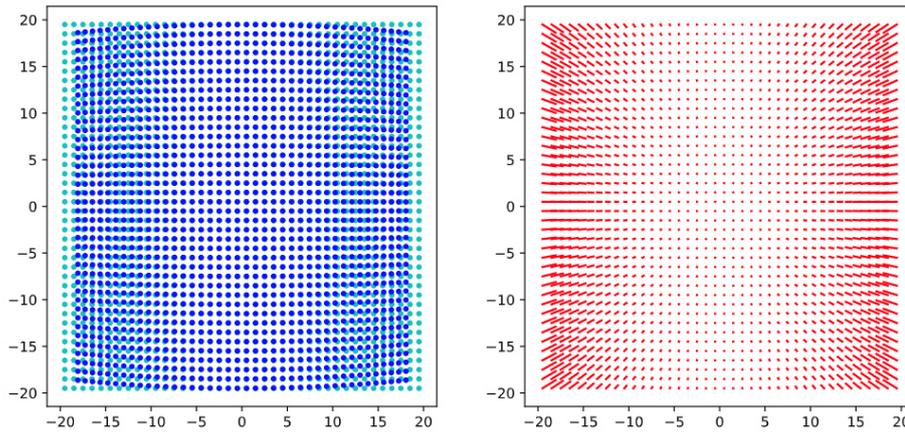


Figure 2.11: On the left, the viewed, deformed grid is represented in dark blue, with the real grid overlaid in light blue. The point-wise transformation is represented on the right.

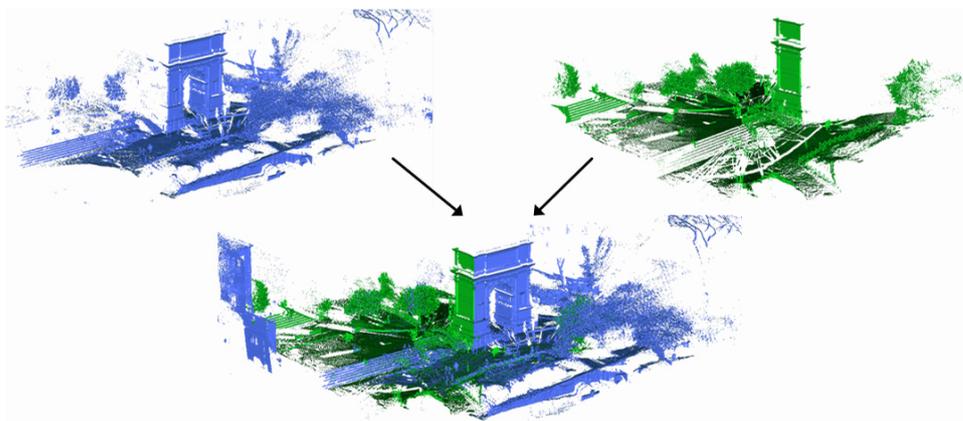


Figure 2.12: An example of registration of laser scans of an unknown monument A side view is matched with a front view. *Sourced from [100].*

In our case, two options are available. Either we can directly use the images themselves for alignment, using cross-correlation, or we can use indirect data, such as tracking output. The first option may seem to be better at first sight since raw images are not subject to processing-induced errors, such as possibly poorly tracked grains. However, the area where we aim to fuse neighboring images is along their edges. This is where they are the most visually dissimilar since it is observed at a high but opposite angle from each camera. It was thus chosen to use tracked grain coordinates as reference points for alignment.

The experimental design detailed above implies that the granular pile is not crystallized. However, the grains are relatively close in diameter ($\sim 9\%$ difference) and the local ordering can still resemble hexagons. Because of this, using a correlation

between sets of tracked grains may present a lot of local minimums which would be hard to discriminate between. To prevent this issue, we only used a given category of grain (small or large) as reference points, since such a subset is likely to have much less regularity. The geometric model we used to flatten our picture has some limitations, in the sense that it does not include any rotation (roll, pan, and tilt) of either the observed object or the camera. These can be accounted for by giving some leeway to the cross-correlation process. Instead of comparing two sets \mathcal{P}_1 and \mathcal{P}_2 , we can compare the first one to a transformed version of the second, $T(\mathcal{P}_2)$. The flexibility of the matching process is determined by the extent of the transformation allowed in T . In our case, T was only able to rotate and shear the points. In figure 2.13, two sets of neighboring points, comprised only of small grains, have been aligned. This alignment required a rotation.

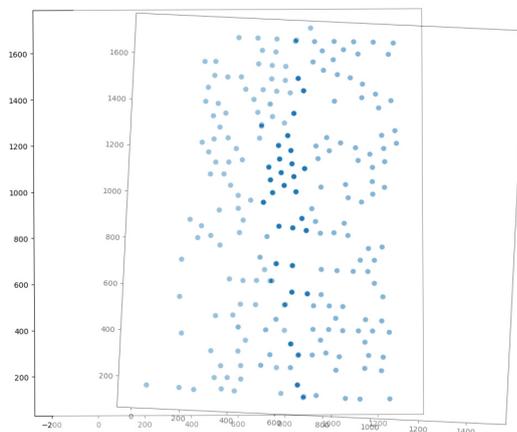


Figure 2.13: Two matched subsets of points, comprised of small grain position data. The overlapping points appear darker. Initial figure frames are left to show the data has been, among other things, rotated.

One big issue with this approach however is error accumulation. We have 24 cameras, producing 24 sets of points. Using point registration, we transform data to perform a local match. This means \mathcal{P}_2 will be transformed to match \mathcal{P}_1 , but \mathcal{P}_3 must be transformed *to match the transformation of \mathcal{P}_2* . Small errors can get magnified after 23 steps. In practice, local geometry is conserved in all cases but global geometry is not, as illustrated by figure 2.14 where the detected grains end up curving downward. Such drift can be fitted and corrected for, but that would add yet another ad-hoc transformation to our data.

This geometrical method was used for qualitative measurements back when the setup had 12 cameras. Due to this limitation, however, it was not sufficient anymore for quantitative analysis. I thus devised a completely different method based on

empirically measuring the transformation to apply, which I will now explain.

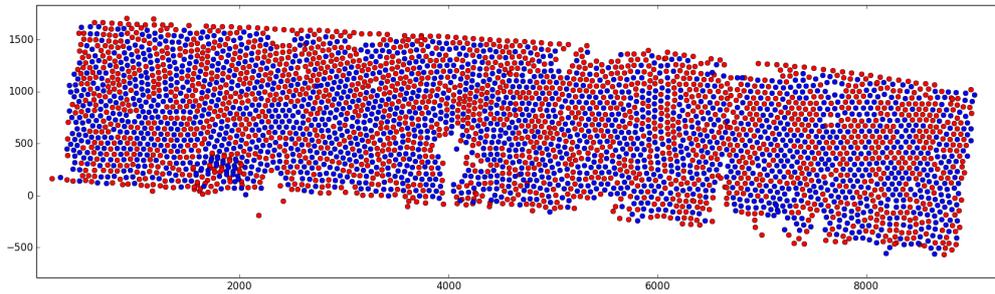


Figure 2.14: A panoramic detection reconstructed with iterative registration. Local geometry looks fine, but a global drift downward can be observed. Color codes grain size: red for larger grain, blue for smaller ones. Holes were caused by cracks in the cylindrical outer wall, as well as tears in the polarizing sheets. Both were eventually changed, the consequences of which are discussed at the end of chapter 3.

Empirical approach

This method is inspired by the use of *fiducial markers* in computer vision [101–104]: these markers serve as references for the imaging system to better interpret its environment, by providing information such as the location and orientation of the marked objects. The obtained data can then be used for various applications, such as augmented reality (see figure 2.15), motion capture, or deformation mapping - our objective here.

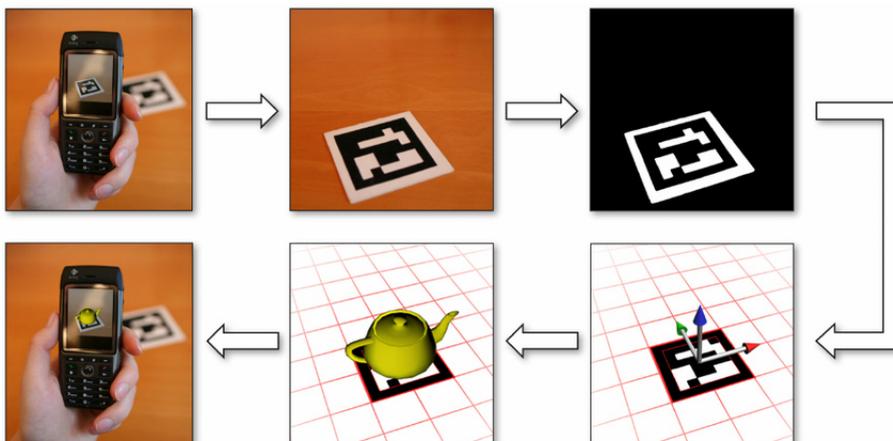


Figure 2.15: A fiducial marker is used to transmit position and orientation data to the smartphone so it can correctly overlay a 3D kettle on top of the table. Sourced from [102].

The marker I designed consists of a long rectangular grid of red dots, with unique green round marks regularly placed along the central line of the grid (see figure 2.16). Once printed on a large sheet of paper, we can wrap it around the cylindrical wall of our experiment. The green marks are similar to 5bits QR-code, with a known absolute position within the grid. The position of each red dot can be known relative to its nearest green marks and thus, absolutely located as well.

We can then take a picture¹ of the wrapped cylinder, and assign to each red dot seen in each camera an absolute coordinate in the flat plane represented on the printed paper. This defines a function, taking pixels coordinates and image number as input and providing absolute coordinates as outputs. This function can then be interpolated and applied to every pixel in our images where the marked paper has been removed as if we were unwrapping the image like the label on a can of food.

Practically, this transformation is done using the *LinearNDInterpolator* function in the scipy [105] python library, which performs a Delaunay triangulation followed by a barycentric weighting. The computation of the interpolation parameters is very expensive and takes roughly 3 hours on a modern laptop to run. However, since these parameters do not change during an experiment, they can be saved and directly applied afterward. Experimentally, this calibration is done once per experiment at the start. This is sufficient as cameras do not move during an experiment, but also necessary as they may move during setup and cleanup.

Once the pixel positions have been interpolated to the absolute frame of reference, they will no longer be aligned on a regular square grid. An additional step of interpolation is required to cast the data into the shape of a raster picture. This is done using the same scipy linear interpolation function. Finally, images are stitched together using a procedure based called “seam carving”². Seam carving [106] was originally developed to reduce the size of an image without cropping, which may lead to loss of data, or squeezing, which will change the aspect ratio. The procedure revolves around computing a path (or a *seam*) that will minimize information loss, usually by minimizing the image gradient along the path. The path (or *seam*) is then “carved out” of the image.

This new method has several key advantages over the previous one. Since the deformation is measured instead of computed, there are essentially no approximations and everything should be accounted for. It also makes it less crucial to set the camera perfectly horizontal. However, the most important advantage over the first

¹The green and red colors are chosen for easier separation during image processing since we capture RGB pictures.

²(now called *liquid rescaling* in GIMP)

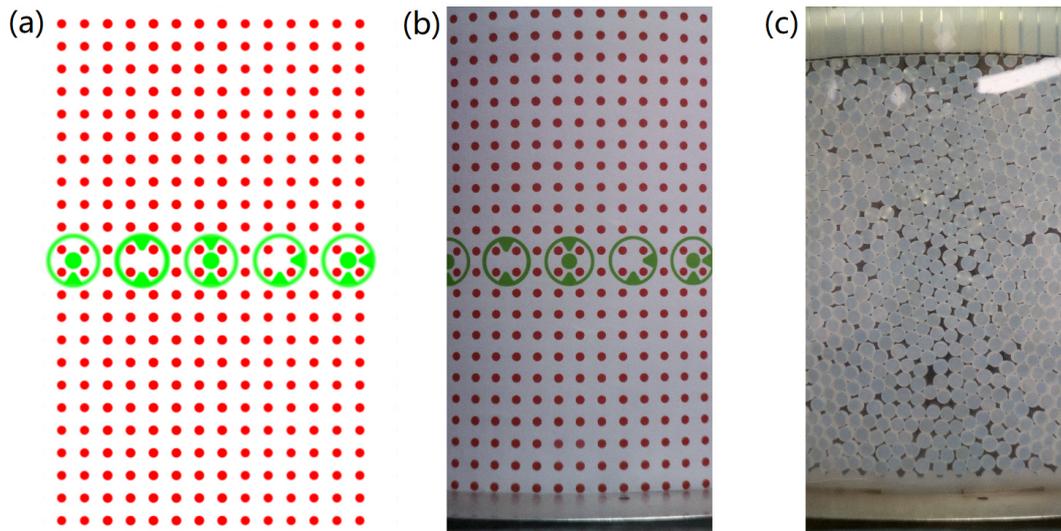
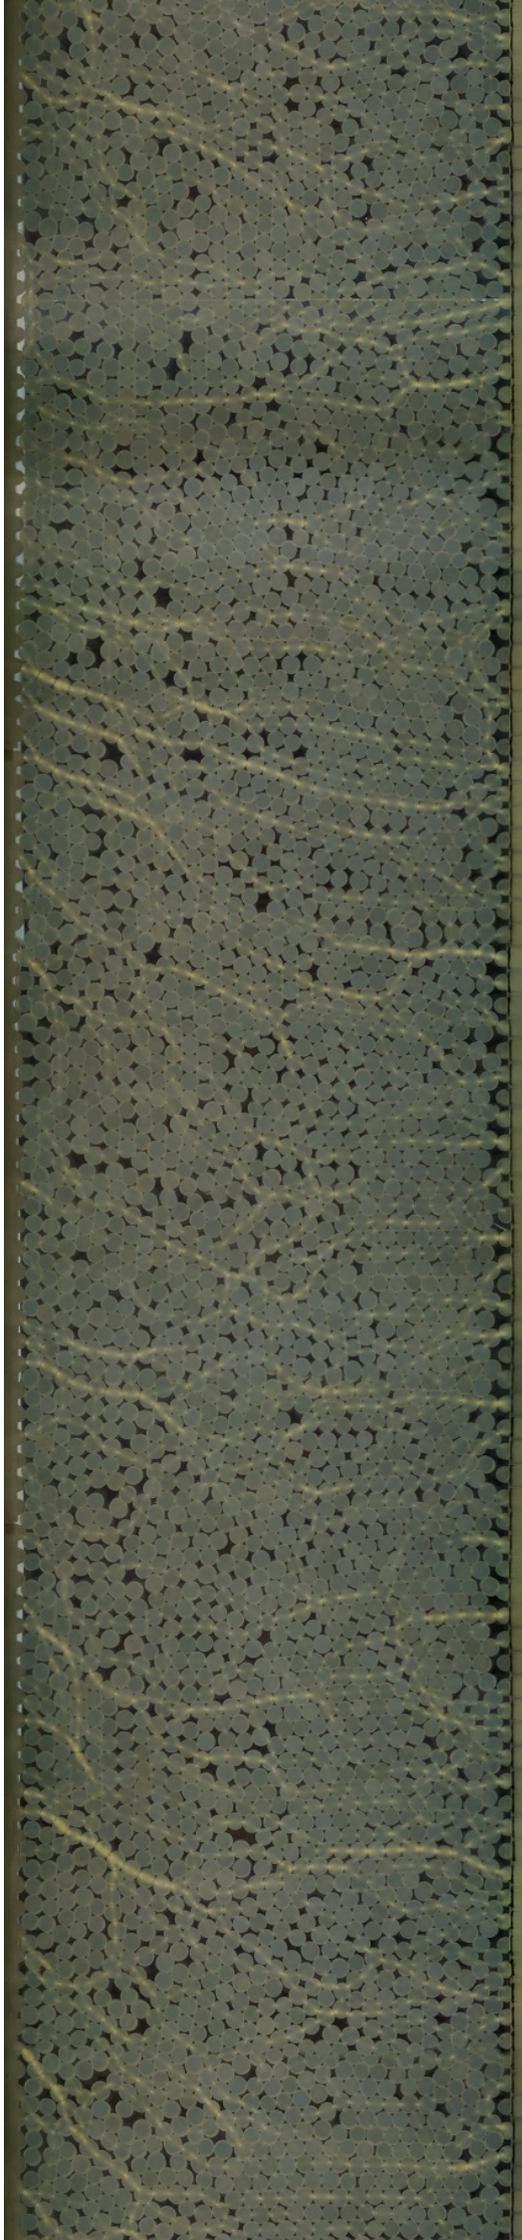


Figure 2.16: (a): Flat, original picture of the reference grid. Green symbols are unique and provide global position information within the reference while red dots provide local geometry information in relation the green symbols. (b) Picture of the printed reference grid wrapped around the outer cylinder of the experiment. (c) Picture taken without the printed reference grid. (*With a light reflections, which are not present during experiments.*)

method is the absolute frame of reference common to all transformations. Since the reference grid location information is relative to the whole paper sheet, every pixel is not only flattened but placed in that absolute frame as well. Thanks to this property, the reconstruction error between data from camera 1 and camera 2 is expected to be the same as between camera 1 and camera 24: there is no accumulation of errors.

A panoramic image created using this technique is shown on the next page.



2.3.2 Image exploitation

Tracking

In the procedures previously described, we focused on unwrapping the data from a collection of 24 independent and curved images to exploitable, globally positioned data. From there, two major measurements can be performed: grains position tracking and force network extraction. The tracking is done using an implementation of the Hough algorithm [107] found in OpenCV for Python [108] which shows great performances; the estimated number of missed particles is $< 0.2\%$ and the mean error on the position is $\sim 5\%$ of a grain radius.

Birefringence and photoelasticity

Materials that have two refractive indexes are called *birefringent*. When light goes through such a material, the relevant index will depend on the direction of polarization of the light. Some materials are naturally birefringent, such as calcite or ruby. This property can also be achieved by some material under strain, in which case it will be called *photoelastic*. This phenomenon is illustrated in figure 2.17.

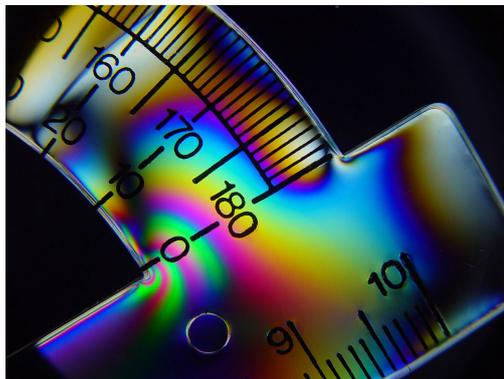


Figure 2.17: A protractor showing tension lines thanks to polarized light. *Picture by Nevit Dilmen, licence CC BY-SA 3.0*

Let us consider a photoelastic material at rest (without birefringence) with a refractive index n_0 and thickness h . If a stress of σ_1, σ_2 (along its principal directions) is applied, an index difference will appear:

$$\Delta n = n_2 - n_1 = C(\sigma_2 - \sigma_1) \quad (2.6)$$

where C is a material-specific constant expressed in *Brewster*, with $1 \text{ Bw} = 10^{-12} \text{ Pa}^{-1}$. It encompasses both the photoelastic response to strain, and the strain response to the stress. If two coherent monochromatic waves of wavelength λ go through the

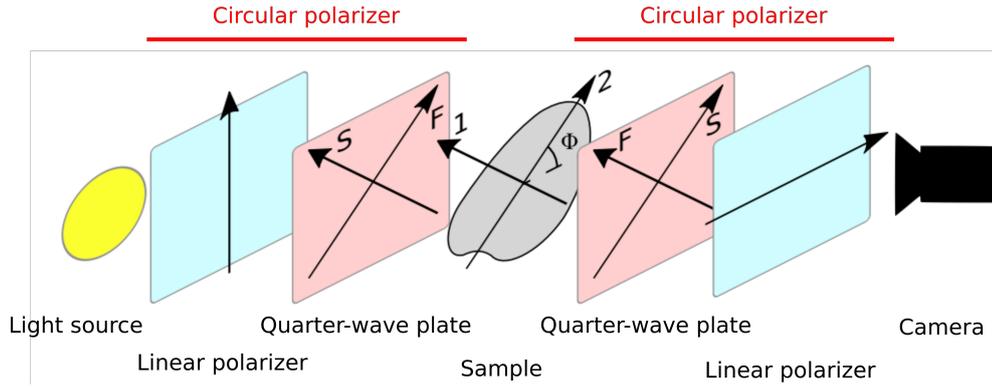


Figure 2.18: Schematic of a polariscope. The sample is placed between two circular polarizers (made of a linear polarizer and a quarter-wave plate). The “S” and “F” denote the slow and fast axis of the quarter-wave plates, respectively.

material (polarized respectively along axis 1 and 2) a phase shift $2\pi p$ will be induced, where p is given by:

$$p = \frac{h C(\sigma_2 - \sigma_1)}{\lambda} \quad (2.7)$$

To visualize the stress using this phase shift, polarized light is needed. A schematic of the optical setup is shown in figure 2.18. By using circularly polarized light, the signal observed by the camera is made independent of ϕ , the orientation of the stress relative to the individual (linear and quarter-wave) polarizers.

After a wave has gone through the setup, the light intensity is:

$$I \sim \sin^2(\pi p) = \sin^2\left(\frac{\pi h C(\sigma_2 - \sigma_1)}{\lambda}\right) \quad (2.8)$$

The details of this calculation are found in annex B. Under no stress, $p = 0$ and the intensity is null: this is a *dark field* configuration where small stress will locally increase luminosity. Turning either circular polarizers by 90° will reverse this behavior: stress will darken an otherwise bright sample.

Force estimation

For the relatively soft grains commonly used in photoelastic studies (cite K. Daniels...) and for the typical stressed applied... For soft grains, the strain is larger for a given stress and the range of value covered by the phase shift p can be large. In this case, several interference fringes can appear. For hard grains like ours, no fringe appears the vast majority of the time. Depending on the case, different approaches exist based either on total luminosity, gradient integration [109] or fringe counting. The

photoelasticity wiki [110] is a great source of information concerning experimental and analysis methods regarding photoelasticity.

In our case, we used the *gradient-square method*, or G^2 method. This method is well explained in the *photoelasticity wiki*, but I will quickly recall it below.

First, we define the gradient squared as:

$$|\nabla|^2 = \frac{1}{4} \left[\left(\frac{I_{i-1,j} - I_{i+1,j}}{2} \right)^2 + \left(\frac{I_{i-1,j-1} - I_{i+1,j+1}}{2\sqrt{2}} \right)^2 + \left(\frac{I_{i,j-1} - I_{i,j+1}}{2} \right)^2 + \left(\frac{I_{i+1,j-1} - I_{i-1,j+1}}{2\sqrt{2}} \right)^2 \right] \quad (2.9)$$

Where $I_{i,j}$ is the intensity of the image at coordinate (i, j) . The G^2 value is then defined as the normalized sum of $|\nabla|^2$ over the region of interest (ROI):

$$G^2 = \frac{1}{N} \sum_{i,j \in ROI} |\nabla|^2 \quad (2.10)$$

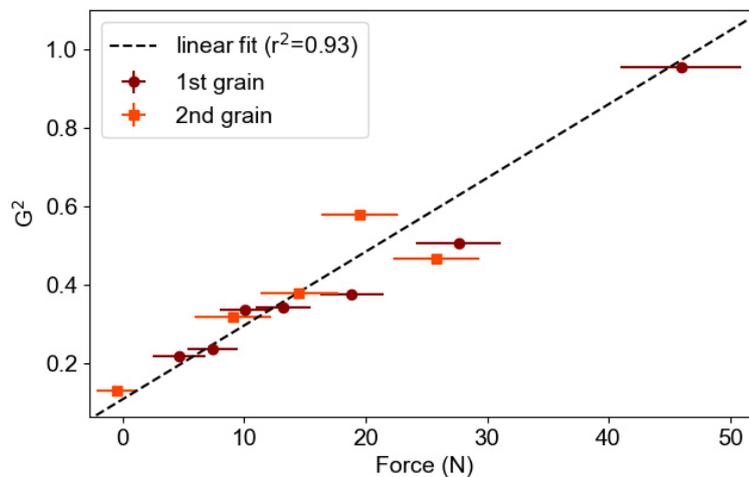


Figure 2.19: Calibration of the G^2 value against a vertically applied force on a single grain, for two grains.

This G^2 method is not applied directly to the raw images: our images are colored, and the G^2 is defined on single-channel images (greyscale). One could simply use the luminosity of the images – however, this is not adapted since the photoelasticity-induced phase shift depends on the wavelength. Hence, mixing colors through luminosity (which is a weighted sum of the red, green, and blue channels) could attenuate and complicate the gradient behavior. As it happens, the intensity change in the red channel is the most correlated with applied force, and all the following analyses will be performed on this channel. Figure 2.20 and 2.21 show eight pictures of our grains

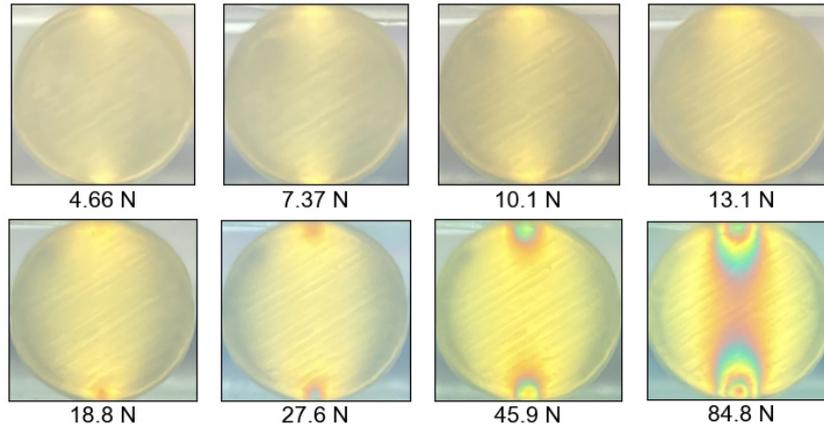


Figure 2.20: Eight pictures of the grain during the uni-axial stress test realized for figure 2.19.

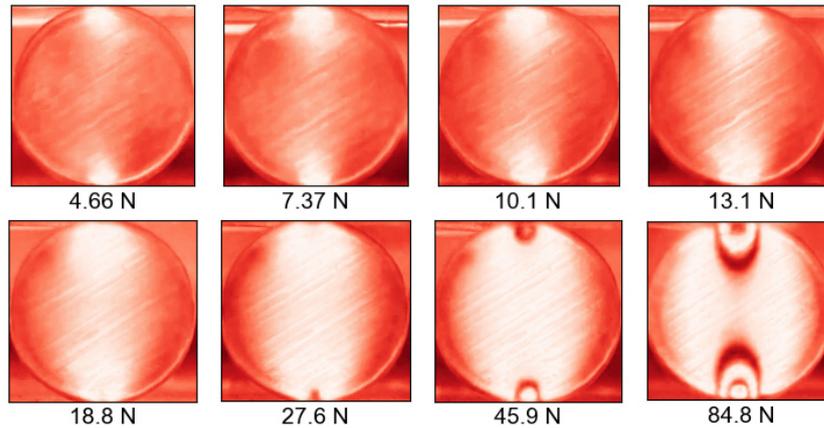


Figure 2.21: Red channel of the eight pictures shown above.

(large size, $d = 7$ mm) under different loads. The first set of images are standard RGB images and the second set is only comprised of the red channel.

In figure 2.19, the G^2 value is compared to a vertically applied load. We can see that there is a good linear correlation. The large horizontal error bar on the force is explained by the test methodology. First, the load is applied to the grain, a picture is taken, and the load removed. The force is measured during the whole process. Since the loading is not done with a deadweight but a press, the grain can undergo a small quasi-plastic relaxation. Thus, the measured force drops a little during the process. This range of force is the dominant cause of uncertainty in figure 2.19.

Chapter 3

Relevance to earthquakes

Contents

3.1	Energy distribution	40
3.1.1	Distribution of the different event types	40
3.1.2	Robustness across acoustical sensors	41
3.1.3	Coherence of the definitions	41
3.2	Omori law	43
3.3	Inter-event time distribution	45
3.4	Changes in experiment behavior	47

This rather short chapter mostly corresponds to the results published in 2019 [83].

In the previous material and methods chapter, I explained the two ways we used to define labquakes (and their associated energy), either using the acoustic signal or the force signal. Both types of events will be analyzed here, and will respectively be referred to as *acoustic event* and *mechanical event*, with their respective energy denoted as E_a and E_m .

Let us now detail the statistical similarities between labquakes and earthquakes, with a focus on the laws that were detailed at the end of the introduction.

As most interpretations focus on power-law exponents, the energies and the associated probability distributions will sometimes be multiplied to align elements on some figures. This will be indicated by a “ \propto ” symbol on the relevant axis labels.

3.1 Energy distribution

3.1.1 Distribution of the different event types

Figure 3.1 presents the probability distribution of energy for earthquakes and both types of labquake events (acoustic and mechanical). *The distributions have been shifted along both axes for easier comparison.* The distributions shown were computed on around $\sim 5.5 \times 10^5$ earthquakes, $\sim 2 \times 10^6$ acoustic labquake event, and $\sim 5 \times 10^3$ mechanical events. The earthquake data presented here comes from the same catalog as the one used in the introduction [84].

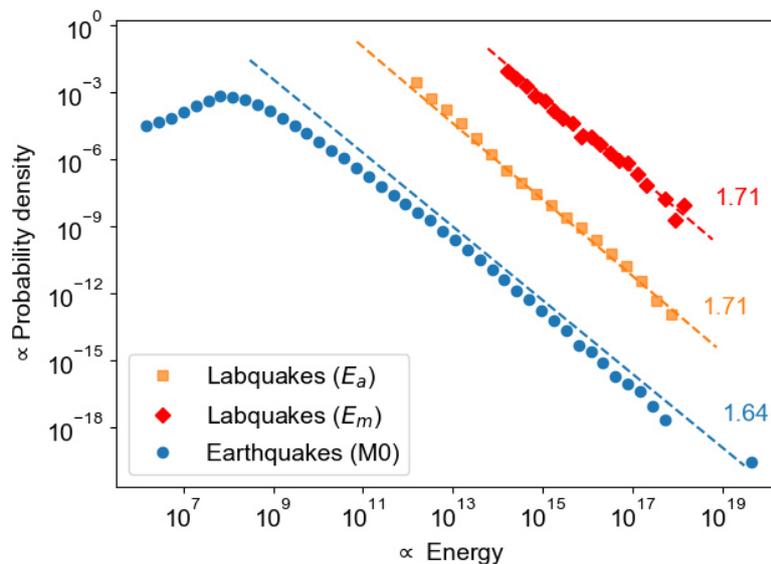


Figure 3.1: Energy distributions of earthquakes, acoustic and mechanical labquakes. All are well described by power laws of comparable slopes.

Both E_a and E_m distributions are compatible with a power-law description of

the form $P(E) \sim E^{-\beta}$, with $\beta_a = 1.71$ and $\beta_m = 1.71$. These values were obtained by the maximum likelihood method described in [87]. The exponents are coherent between both kinds of artificial events, and very comparable with the one describing this earthquake catalog.

3.1.2 Robustness across acoustical sensors

The distribution of acoustic energies E_a shown in figure 3.1 is based on a single sensor. To check the results do not depend on the sensor, we can compare their distributions, as shown in figure 3.2.

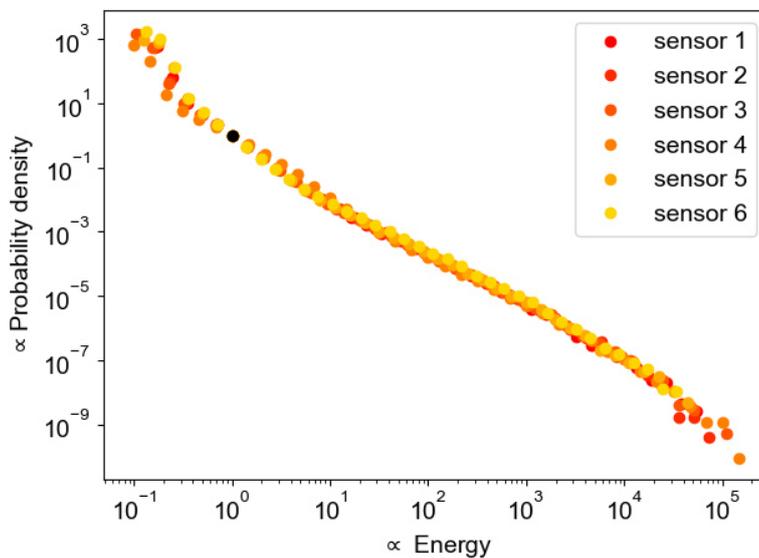


Figure 3.2: Energy distribution of the sensors, shifted so they all share a common point at (1, 1), highlighted in black.

All distributions in figure 3.2 have been shifted, along both axes, to share a common point at the beginning of the power law, highlighted in black on the figure.

3.1.3 Coherence of the definitions

The acoustic and mechanical definition of events produces the same distributions, but we can verify their agreement event-wise as well. Since mechanical events are far less numerous than acoustic events, we will look for acoustic emissions matching our torque drop, rather than the contrary.

To associate an acoustic energy to a torque drop, we find the acoustic event that is closest in time, for each acoustic sensor. As some sensors will be closer the event than others, we will take the average energy from all the sensors. Before doing

so, however, the data from each sensor is rescaled the same way they have been to make figure 3.2. This way, the average is not dominated by the pinducerTM with the highest coupling.

Finally, the acoustic energy is binned by their associated mechanical energy. The median of E_a is computed for each bin, and these finally serve to perform a linear regression.

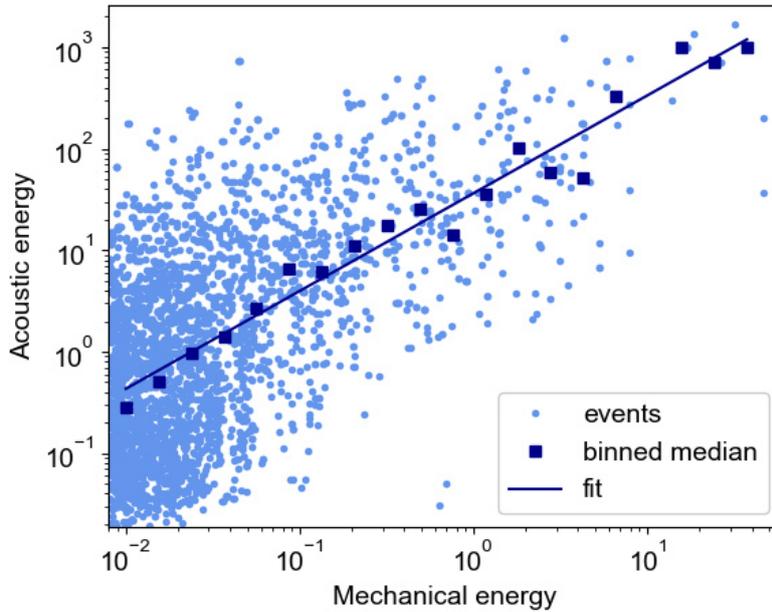


Figure 3.3: Equivalence of the mechanical and acoustic energies. The regression is performed on binned medians of all paired events.

This process results in figure 3.3, with the event data in light blue, the binned median in dark blue, and the fit in solid dark blue. The slope of the fit is 0.96, indicating an agreement on average. However, the spread of the point cloud shows E_a and E_m are not trivially linked either.

3.2 Omori law

To study aftershock and foreshock sequences, I started by investigating the activity after the largest event of a dataset, with the idea to fit the evolution of the event rate using the equation presented in the introduction:

$$n(t) = n_0 + \frac{k}{(c + t)^p}$$

Let us remind the physical meaning of each parameter. k is a global activity rate multiplier, n_0 is the background rate and c acts as a time lag which helps to account for the plateau due to catalog incompleteness. Finally, the Omori exponent p dictates the decay rate and is typically close to 1.

However, as shown in figure 3.4, the activity curves obtained are jittery and subject to “catalog” incompleteness, particularly in the cases of foreshocks. This makes the estimation of the parameters difficult.

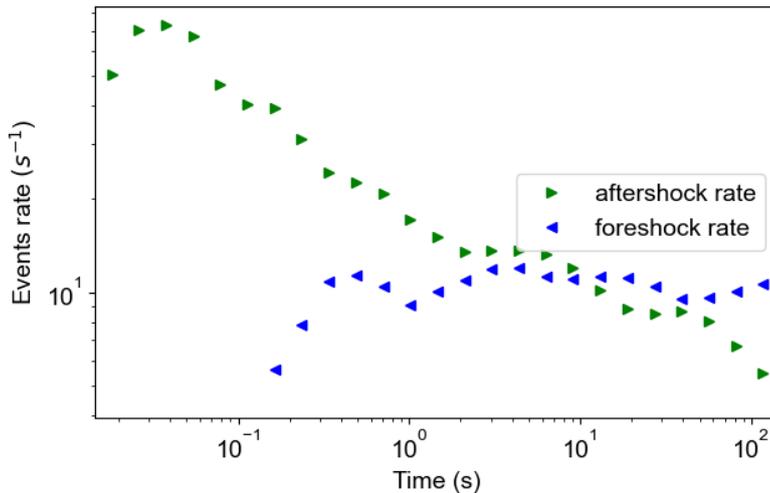


Figure 3.4: Foreshock and aftershock sequence for the largest event.

To solve this issue, we can average the activity rate over many events. Figure 3.5 display the same analysis as done in figure 3.4, but average over the 4500 largest events. The trends appear much more clearly, and the resulting fits are:

$$n_{\text{aftershock}}^{4500}(t_+) = 8.75 + \frac{5.54}{(0.011 + t_+)^{0.8}}$$

$$n_{\text{foreshock}}^{4500}(t_-) = 9.35 + \frac{2.13}{(0.0029 + t_-)^{0.6}}$$

where $t_- = t_e - t$ and $t_+ = t - t_e$, with t_e the event timing. The Omori exponent is noticeably smaller for the foreshock rate, as well as the rate increase prior to the event. The plateau effect is less pronounced as well.

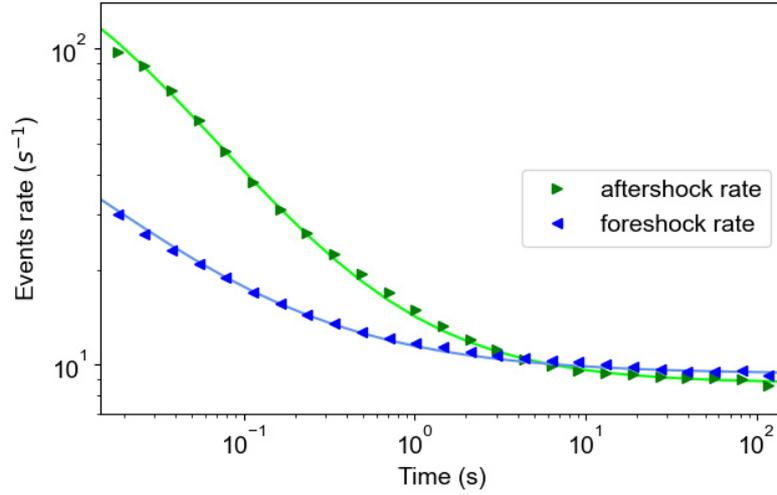


Figure 3.5: Foreshock and aftershock sequence average over the largest 4500 E_a events.

Changing the threshold does not change the behavior, however, the parameters are slightly different. An example of a larger threshold is shown in figure 3.6, selecting only 1000 events, with the associated fit parameters below it.

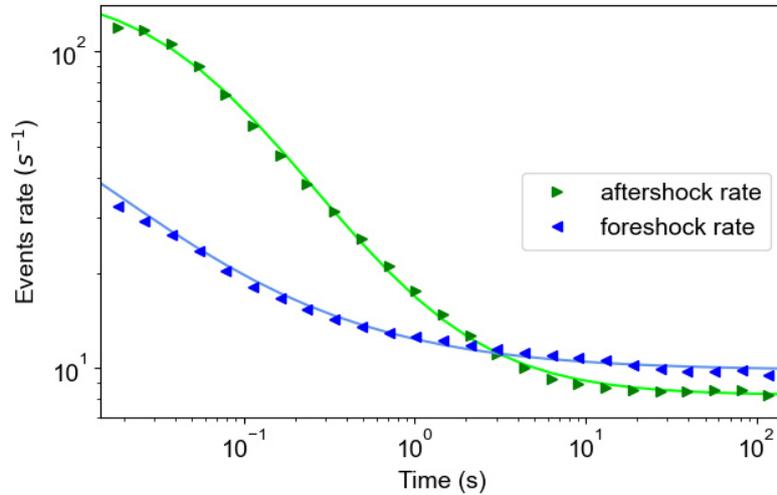


Figure 3.6: Foreshock and aftershock sequence average over the largest 1000 E_a events.

$$n_{\text{aftershock}}^{1000}(t_+) = 8.22 + \frac{9.18}{(0.061 + t_+)^{1.0}}$$

$$n_{\text{foreshock}}^{1000}(t_-) = 9.85 + \frac{2.53}{(0.029 + t_-)^{0.65}}$$

The background rates are further apart, the time offset c is larger (as proportionally more events may be missed), and the decay rates are larger as well. This is coherent with the fact that the effects of these 1000 events are not diluted with the next 3500 largest events, as they were in the previous figure (3.5).

3.3 Inter-event time distribution

Let us quickly remind the quantities defined in the introduction to study the inter-event time distribution. We first define the time series of waiting times between events of energy above a given threshold E_0 , noted $\tau_{E \geq E_0}$. This time series is first rescaled by its average $\tau_{\geq E_0}^*$, to make $\theta = \tau_{\geq E_0} / \tau_{\geq E_0}^*$. The distribution of θ is then itself rescaled by the inverse of $\tau_{\geq E_0}^*$, denoted $R_{E \geq E_0}$.

Several distributions $PDF(\theta) \cdot R_{\geq E_0}$, with different thresholds E_0 , are represented in figure 3.7. The chosen thresholds are compared against the initial energy distribution in figure 3.8. Much like in the case of earthquakes, the data collapses well around the master curve proposed by Corral [96]. The equation of the master curve has not been fit to our data, instead, the same parameters have been kept to underline the resemblance. The equation of the master curve is $f(\theta) = 0.5 \theta^{-0.33} \exp(-\theta^{0.98}/1.58)$.

The conclusions that were drawn for earthquakes apply here as well. There are weak memory effects in the system, which appears to “remember” an event of energy E_0 for a duration equal to the average inter-event times of event larger than E_0 .

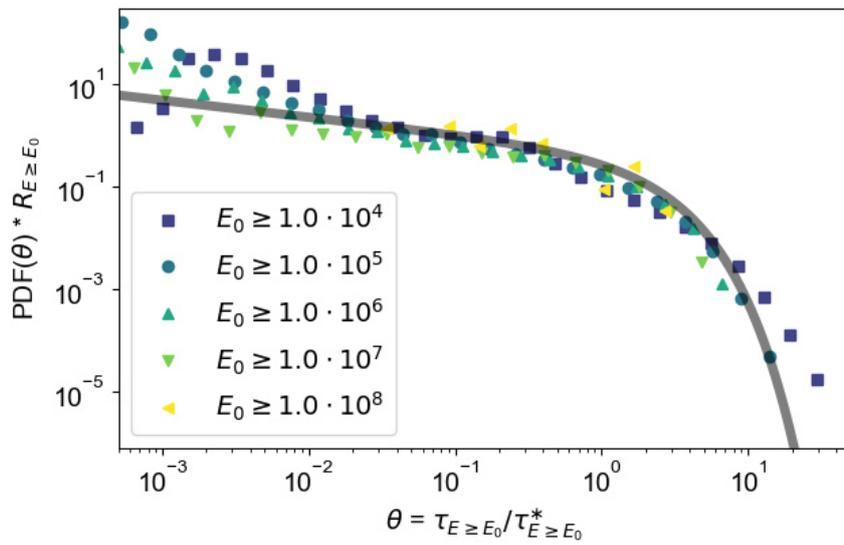


Figure 3.7: Normalized probability distribution of the dimensionless waiting times $\theta = \tau_{\geq E_0} / \tau_{\geq E_0}^*$.

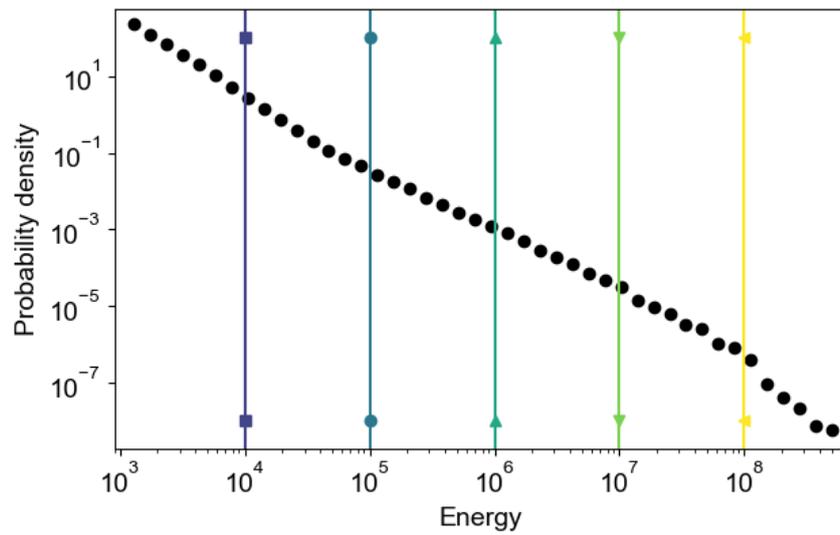


Figure 3.8: Energy distribution of the data set used to make figure 3.7

3.4 Changes in experiment behavior

In the imaging section of chapter 2, I explained parts of the experiments were changed to improve image quality: the cylindrical sidewalls, as well as the bottom ring shearing the grain which was made shorter (to augment the vertical field of view) and partly replaced with metal. This had several effects on the statistical features presented in this chapter.

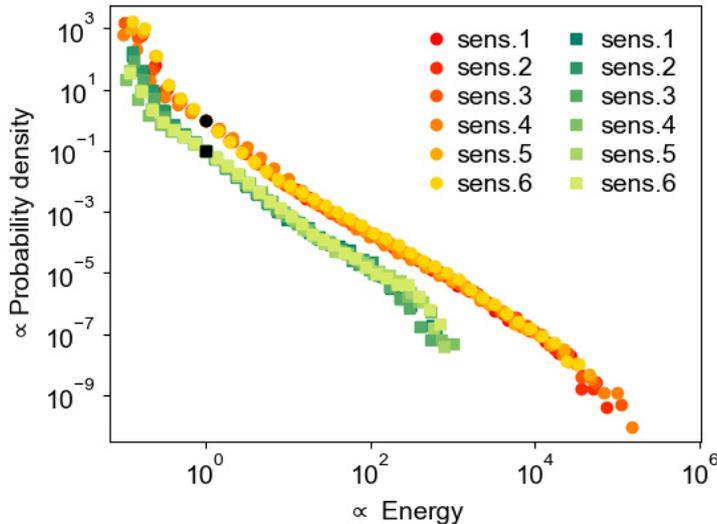


Figure 3.9: Comparison of aligned acoustic energy distributions. The recent data set is indicated by squares, the older one by circles.

In figure 3.9, we can see the acoustic energy distribution has significantly shrunk, from roughly 6 decades to 4. In addition, the exponent of the best fitting power law is slightly higher, at $\beta = 1.85$, indicating larger events are slightly less common.

In addition, the analysis of after and foreshock sequences has also revealed some differences, as illustrated by figure 3.10 which reproduces the method done for figure 3.5, averaging over 4500 events. The activity returns to its background rate faster, in about 0.5 seconds compared to the 3 seconds in older data sets. A quiescence period of around 5 seconds can also be observed in both foreshock and aftershock sequences, where the event rate dips roughly 21% below its asymptotic value (at 2.19 s^{-1} , compared to 2.78 s^{-1}). The fitted law, matching well both curves, has the equation: $n(t) = 8.22 + \frac{2.77}{(0.073+t)^2}$. The most striking (and still unexplained) feature is the almost perfect overlap of both curves. The analysis has been checked multiple times and this is not due to a coding mistake, indicating a change in the experiment behavior.

Finally, the inter-event time distribution presents deviations that are coherent with the ones observed in the Omori law: There is a dip below the master curve

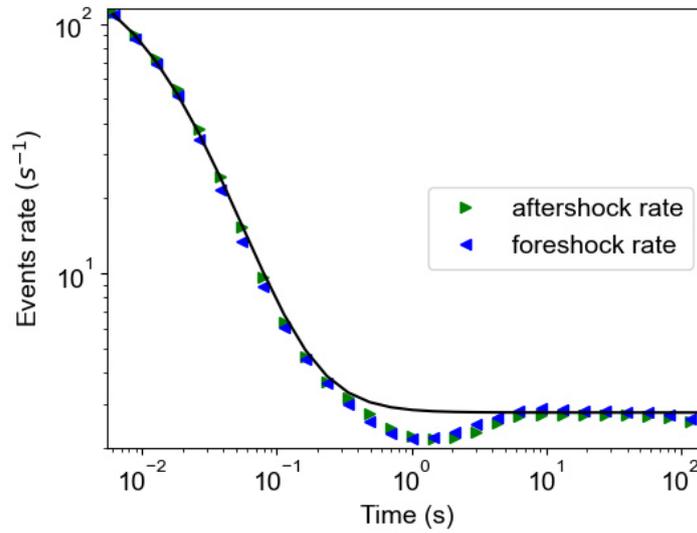


Figure 3.10: Aftershock and foreshock sequences for recent data sets.

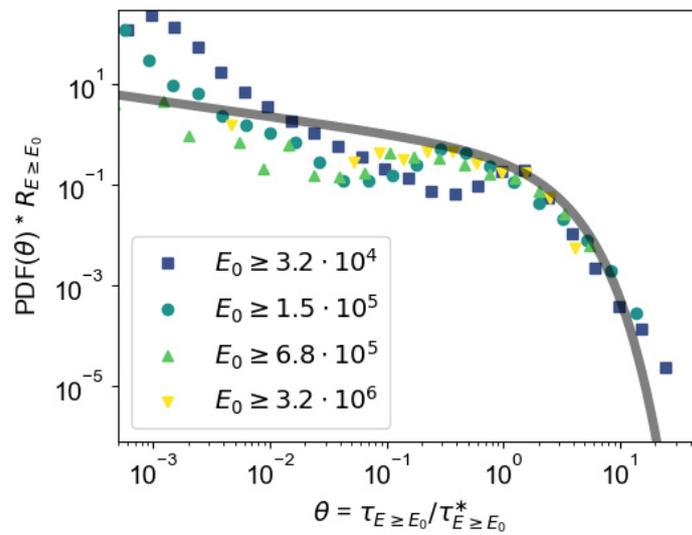


Figure 3.11: Normalized probability distribution of the dimensionless waiting times for recent data sets.

around $\theta = 10^{-1}$, and the beginning of the distributions have steeper slopes.

Considering the deviations observed and the nature of the work done on the experiment, we expect an increased stiffness of the system is responsible for these changes. While this evolution of the setup was unexpected, exploring what dictates the statistical law's exponents is one of the long-term goals of this project. New parts are being made to investigate the influence of the changes done on the experiment.

Chapter 4

Probing the structure

Contents

4.1	Evolution of the volume	52
4.1.1	Height discontinuity events	52
4.1.2	Augmenting labquakes properties	54
4.1.3	Force chain orientation as a likely cause	56
4.2	Evaluation of the force network fractal dimension using images	62
4.2.1	Minkowski–Bouligand dimension	62
4.2.2	Application to thresholded images	62
4.3	Acoustic-based fractal dimension	67
4.3.1	propagation model	67
4.3.2	A simpler case: sound propagation in air	69
4.3.3	Application to our experiment	71

4.1 Evolution of the volume

The dilation of a granular layer under shear is a classical and a rather intuitive phenomenon reported more than a century ago [111]. To study this behavior, two inductive sensors measuring the vertical movement of the upper plate were added to the setup (*Baumer IPRM 12I9505/S14*). In figure 4.1, one can see the signal $h(t)$ seems to evolve relatively smoothly with a few discontinuities. These discontinuities are denoted as Δh^1 . On the zoomed event (bottom), oscillations can be observed. As explained at the end of section 2.2.2, these oscillations are easy to filter out and do not hinder any of the following analyses. Controlled stepping motors compensate for the movement of the plate to guarantee the sensors are always in their range of measurements.

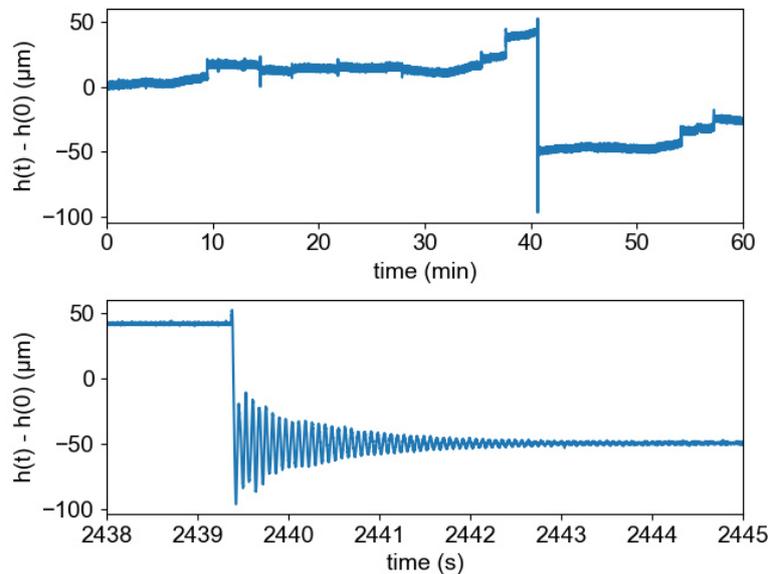


Figure 4.1: *Top*: the first hour of the position measurement. *Bottom*: a zoom of the large event right after 40 minutes.

4.1.1 Height discontinuity events

The expected behavior of a granular under shear is to dilate until the grains rearrange suddenly, causing a collapse in the structure. In our experiment, this manifests as a drop ($\Delta h < 0$) of the top plate. This view corresponds well with the large event observed around 40 minutes in figure 4.1. However, we can also observe some sudden *jumps* ($\Delta h > 0$) around 10 or 55 minutes.

¹This notation **does not** represent the signal $h(t) - h(0)$

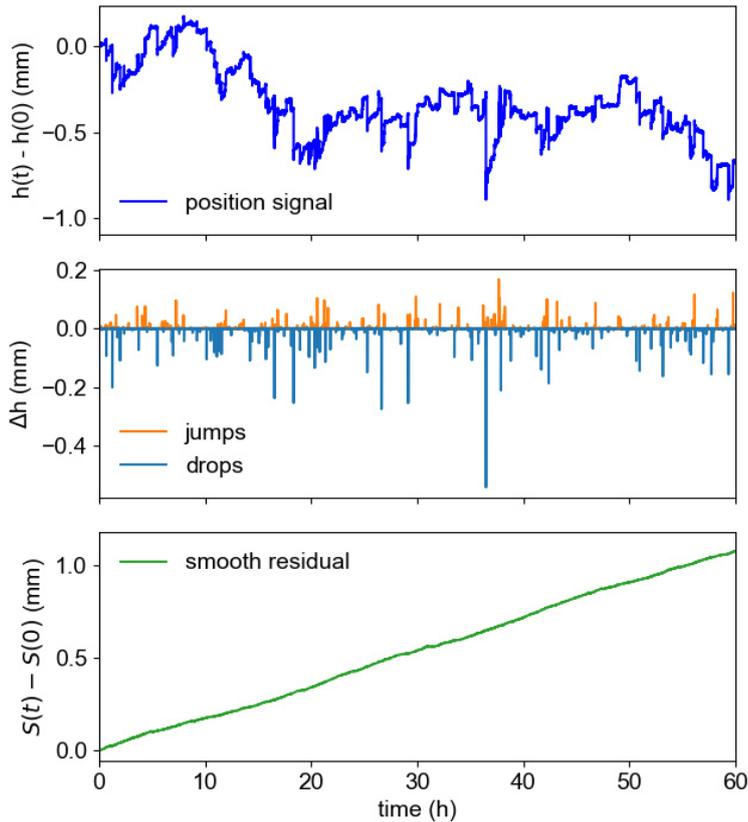


Figure 4.2: From the raw position signal (*top*), a series of jumps and drops (*middle*) can be extracted. By removing these discontinuities from the raw signal, we can obtain a smooth, continuous dilation (*bottom*).

By looking for positive and negative peaks in the convolution, we can measure the timing and displacement of events with a precision as low as $0.1 \mu\text{m}$. 60 hours of the raw position signal is shown on the top of figure 4.2. We can extract the time series of jumps and drops (*same figure, middle*) from this signal using the described filter. Finally, a smooth dilation $S(t) = h(t) - \Delta h(t)$ (*bottom*) is obtained by subtracting the discontinuities from the raw signal. This smooth residual can be interpreted as a signature of the expected granular behavior under shear. Overall, the discontinuities dominate the dynamic as the overall signal $h(t)$ trends downward, due to very small but accumulating plastic deformations of the grains.

In addition, it seems from figure 4.2 the drops are generally larger in size than the jumps, with no jumps larger than 0.2 mm while several drops pass this threshold. To verify this, we can examine their respective probability distribution, shown in figure 4.3. We can see the distribution of drops indeed reaches higher values than the jumps. This is expected as such events are more energetically favorable. Another difference can be seen slightly between 2×10^{-4} and $2 \times 10^{-3} \text{ mm}$: jumps are more numerous in this range. Finally, the two distributions come together at the lowest

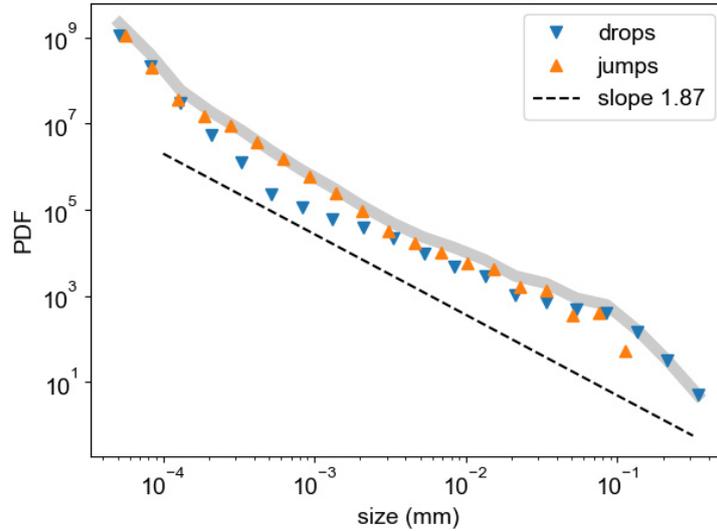


Figure 4.3: distribution of the jumps and drops absolute sizes. The dashed power-law is not a fit of this data, but corresponds to the fitted value for acoustic events for comparison purposes.

values. This can be understood as the threshold is near noise level, which is *a priori* symmetrically distributed in sign. Finally, a dashed line corresponding to a power law of slope -1.85, the value obtained by fitting the distribution of acoustic event size, is added for comparison purposes.

4.1.2 Augmenting labquakes properties

The good agreement between the power-law slope obtained from acoustic events and the distribution of position events Δh is shared with the definition of events based on torque, $E_m \sim \Delta(\Gamma^2)$. We showed acoustic event E_a and position events Δh are similarly distributed, both well described by a power-law with a slope of -1.85. As shown in chapter 2, acoustic energies E_a and mechanical energies E_m (defined using the torque Γ) are also strongly correlated. Since both torque events and position events are mechanical in nature, one can wonder about the possible links between them.

The resisting Γ torque is measured orthoradially and the Δh events are related to vertical force changes (under constant pressure and free volume condition). It seems reasonable to assume that, since torque and vertical forces are simply orthogonal projections of the same force being applied at a given angle by the force chains, $\Delta\Gamma$ and Δh should vary the same way. For instance, when a force chain breaks, it is intuitive to think that the orthoradial force will decrease ($\Delta\Gamma < 0$) and the now less supported top plate will drop ($\Delta h < 0$). However, figure 4.4 shows counter-examples

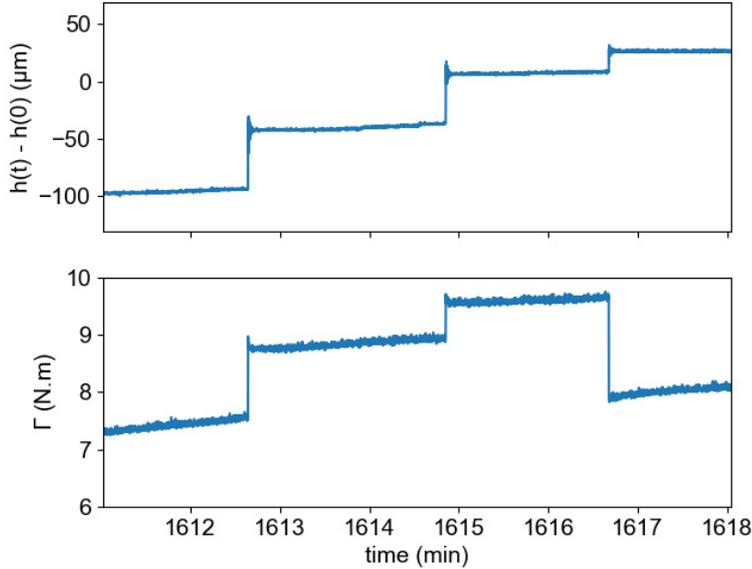


Figure 4.4: Sudden variations in h are simultaneous with sudden variations in resisting torque Γ .

to this intuitive view: From the three observable jumps, the first two correspond to sudden increases both in torque and $h(t)$, while the last jumps correspond to a torque drop but with an associated increase of $h(t)$. The events are simultaneous, however, indeed indicating a common source between the two signals.

To investigate the relationship between these two measurements $\Delta\Gamma$ and Δh , we need a procedure to identify which torque discontinuity $\Delta\Gamma$ matches which Δh . For a given type of measurement (position or torque), the inter-event time is almost never below 100 ms while the time precision on event detection is in the range of 0.1 ms. We have used a criterion of $\Delta t \leq 1$ ms to consider that both measurements correspond to the same reorganization event.

Matched events can then be normalized and compared in log-scale. The resulting scatter plot is presented in figure 4.5, where re-scaled Δh are plotted against re-scaled mechanical energy changes (proportional to $\Delta\Gamma^2$ [83]). For each series, the scaling factor is its detection threshold (respectively denoted as E_m^* and Δh^*). In log-scales, it ensures the range of values starts from 0. The released energy associated with a torque drop is considered positive; therefore a sudden increase of measured torque implies a “negative energy release”.

We can observe the scattered data forms an “X” shape. Contrary to our intuition, not only quadrant D in figure 4.5 is populated, but all other quadrants are also significantly populated. We will next consider the force chains geometry to explain this variety of events.

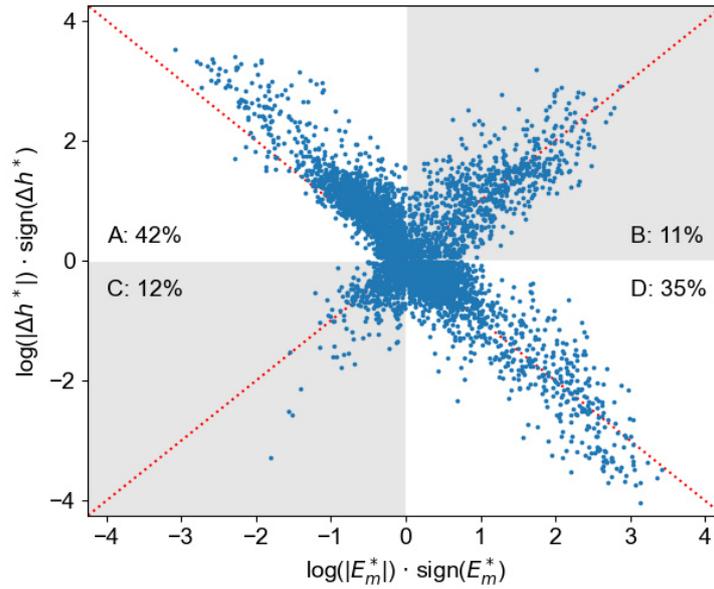


Figure 4.5: Relation between signed and normalized Δh^* values and released mechanical energy E_m^* , for matched events. A-D: four different behaviors with percentage of events in each classification. The red lines represents the functions $y = x$ and $y = -x$.

4.1.3 Force chain orientation as a likely cause

Let us consider a system with a single force chain system, with extreme-case orientations. If the force chain is perfectly vertical, it would not apply any force on the orthoradially oriented sensor. On the contrary, if it was completely flat, the force chain would not reach the force sensor at all. In order to maximize their stability, force chains reorganize around a *local optimal angle*, placing themselves parallel to the total stress. We will use this reasoning of a *single one-chain description*, where one chain (called “active”) holding some part of the stress is broken and replaced by a new one, as the main ingredient to qualitatively explain all the four different behaviors presented in the figure 4.5:

- **A** ($\Gamma \uparrow, h \uparrow$, 42% of events): negative energy release and dilation of the medium. *Replacement of the active force chain with an equivalent one.* This is actually the most common case, particularly for lower energy events. The second ingredient to explain the dynamics corresponds to *the inertia of the structural reorganization*: the kinetic energy developed by the system during the reorganization time, generates an extra-pressure in the new force chain, reading as an increase of the force in both lateral (torque) and vertical (dilation) directions with respect to the replaced one.
- **B** ($\Gamma \downarrow, h \uparrow$, 11% of events): positive energy release and dilation of the medium.

Replacement of the active force chain with a more vertical one. The horizontal projection of the force is smaller with respect to the replaced one. This reads as a reduction of the global torque and, thanks to the inertia, an increase in the vertical force, which provokes a dilation of the medium.

- C ($\Gamma \uparrow, h \downarrow$, 12% of events): negative energy release and contraction of the medium. *Replacement of the active force chain with a more horizontal one.* It should increase the global torque thanks to a better projection; as well as bringing a contraction. However, the probability of finding such a percolating chain for high-energy events is rather low, because of its instability.
- D ($\Gamma \downarrow, h \downarrow$, 35% of events): positive energy release and contraction of the medium (intuitive case of a collapse). *Replacement of the active force chain with a weaker one.* The more energetic releases fall into this classification. Notice that from the 33 more energetic events, 31 events (94 %) fall in the D-zone and only 2 (6 %) in the A-zone.

All the reasoning above is done within the shear band, where force chains are the most dynamic. Above this region, the grains hardly rearrange or move at all, and the upper half can be considered as a quasi-solid storing elastic energy.

Note that each of the above-proposed mechanisms explains *both torque and position variation at the same time*. This makes sense with the fact both kinds of events are caused by the same re-arrangement in the granular, and we can expect large re-arrangement to have large consequences both in $\Delta\Gamma$ and Δh . This is supported by the “X” shape along which the data is distributed. By removing the signed multiplicative factor, we can observe a linear correlation between $|E_m|$ and $|\Delta h|$.

Extraction of force chains

Quantitatively checking the suggested explanation above requires a method to characterize both the strength and angles of the force chains. To do so, we first need to extract them from the images. Specifically, we want to examine the stress that disappeared and appeared between before and after an event. To extract the relevant data, we can take the difference between the image immediately after an event $I(t^+)$ and the image immediately before $I(t^-)$ to filter the conserved force chains, and keep only the one that appeared and disappeared. However since the grains can move quite a lot during large events, conserved force chains move with them, as illustrated in figure 4.6. To compensate for this movement, we first perform a line-by-line correlation between $I(t^+)$ and $I(t^-)$ to estimate the horizontal displacement field.

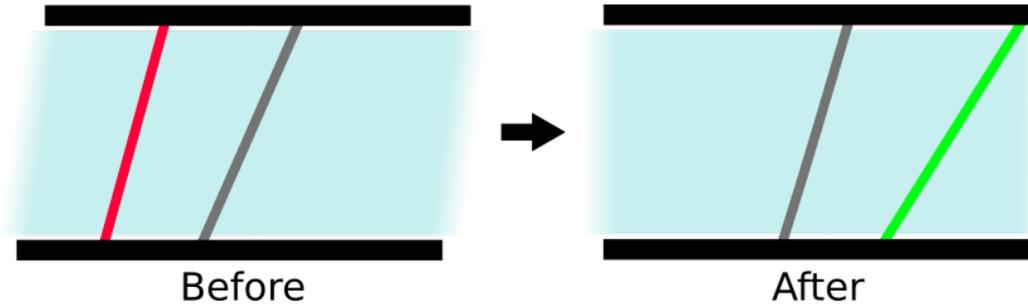


Figure 4.6: Around large event, sudden movement and re-arrangement occurs. Concerning force chains they can either disappear (*red, left*), appear (*green, right*) or deform to be more vertical (*grey, both*).

This field is then subtracted from the after image, $I(t^+)$ (by an operation we will call T) to obtain a differential picture $\Delta I = T[I(t^+)] - I(t^-)$. Finally, we interpret the positive part of ΔI (*resp.* negative) as the force chains that **appeared** (*resp.* **disappeared**) during the re-arrangement event. I will now detail the procedure to measure the “strength” S of an angle α in these force chains.

Evaluation of an angle prevalence

First, we compute the following convolution: $C(\alpha) = \text{image} * \text{mask}_n(\alpha)$. S is then defined as the 99th percentile of C . The masks are small binary images (background set to 0) containing a thin segment tilted at an angle α . The parameter n here represents the side of the square mask, in pixel. In the case of our force chains, this parameter was set to $200px$. This value was chosen as it corresponds to three grain diameters, with our image resolution. n needs to be strictly larger than one diameter (*inclination* would not make sense otherwise) but also smaller than the maximum force chain length (10 to 15 grains diameter²) in order to account for direction change along the force chain.

An example of this procedure is shown in figure 4.7, applied to an artificial binary image made of several randomly placed segments. Each segment is of random length but is either inclined at a 65° or a 80° angle. After convolution with a series of binary masks (one of which is shown in the figure), we can obtain the bottom curve $S(\alpha)$. The two dotted lines indicate the angle used to create the artificial image (65° and 80°) while the larger green dot corresponds to the convolution performed using the example mask, tilted at 45° .

We also can define a normalized version: $S^*(\alpha) = S(\alpha) / \int_0^{\pi/2} S(\theta) d\theta$. Since many

²This may seem very short compared to images shown in previous chapters. For this analysis however the region of interest is the bottom third of the granular (the shear band), hence the “small” force chains considered.

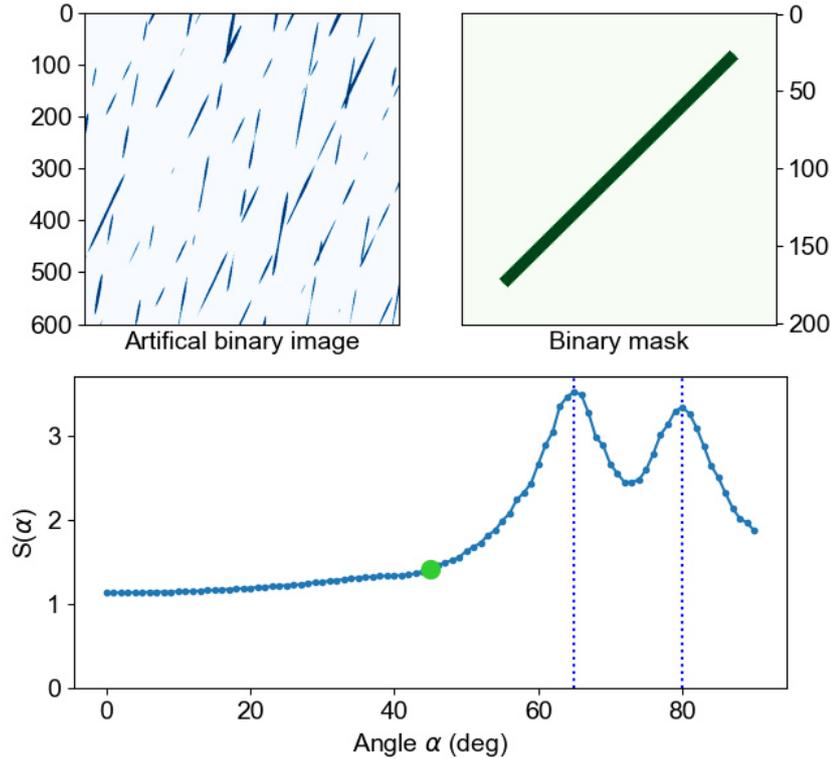


Figure 4.7: The artificial binary image in *top right* is made of segments with random position and length tilted at either 65° or 80° . Taking the 99th percentile of this image convolution with masks (example in *Top right*) produce the curve below. The bold green point corresponds to the mask shown (45°).

force chain disappear during large events, the direct value $S(\alpha)$ mostly provides information regarding force intensity, while the normalized distribution is more related to orientation.

Application to real data

We apply this method to the extracted images of the **appeared** and **disappeared** force chains around the largest 50 events (by absolute mechanical energy E_m) of each category A, B, C, and D as defined in figure 4.5.

The resulting $S(\alpha)$ curves are presented in figures 4.8 and 4.9 (normalized version). Let us revisit our **predicted mechanism** in light of these results:

- **A** ($\Gamma \uparrow, h \uparrow$): negative energy release and dilation of the medium, *Replacement of the active force chain with an equivalent one*. Indeed, both the total force (Fig. 4.8) and orientation (Fig. 4.9) do not change drastically.
- **B** ($\Gamma \downarrow, h \uparrow$): positive energy release and dilation of the medium. *Replacement of the active force chain with a more vertical one*. The large (Fig. 4.8) loss of

force explains the positive energy release, and the newly created force chains are significantly more vertical.

- C ($\Gamma \uparrow, h \downarrow$): negative energy release and contraction of the medium. *Replacement of the active force chain with a more horizontal one.* The distribution of angles has slightly shifted to lower values (Fig. 4.9) and the forces are mostly conserved (Fig. 4.8).
- D ($\Gamma \downarrow, h \downarrow$): positive energy release and contraction of the medium. *Replacement of the active force chain with a weaker one.* The dynamic here is completely dominated by the loss of force (Fig. 4.8).

By inspecting the dilation signal $h(t)$ I uncovered a rich dynamic of sudden jumps Δh . Since these jumps are mechanical in nature and power-law distributed, it made sense to compare them with mechanical energy releases E_m – another measurement with the same characteristics. Doing so, a complex correlation was uncovered between these two quantities: $|E_m|$ and $|\Delta h|$ are linearly correlated but with seemingly any sign combination possible. The four observed sign combinations can be explained by the changes in the distribution of angles within the force chains, between before and after the events.

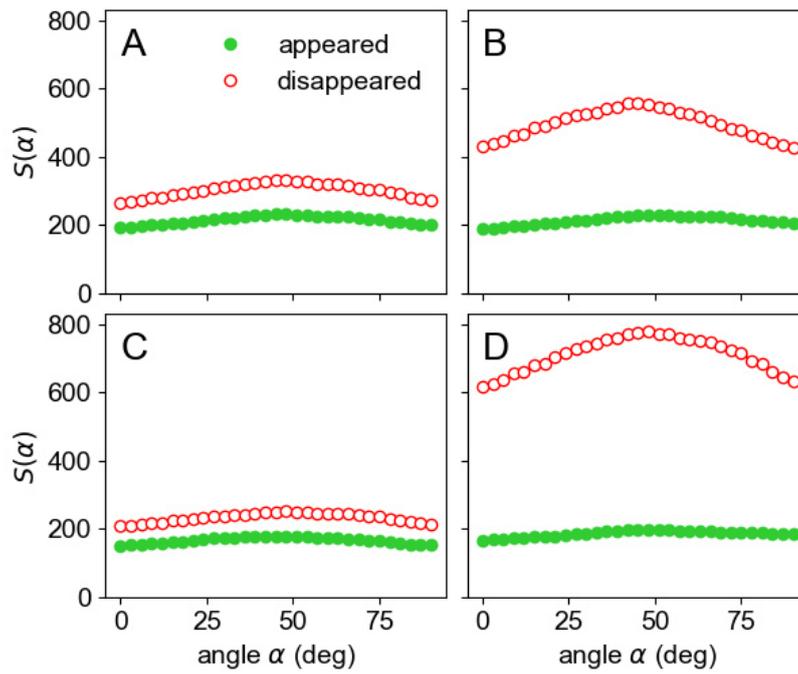


Figure 4.8: Distribution of the strength of the angles $S(\alpha)$ of appeared and disappeared force chains, averaged over the largest 50 events of each category A, B, C or D.

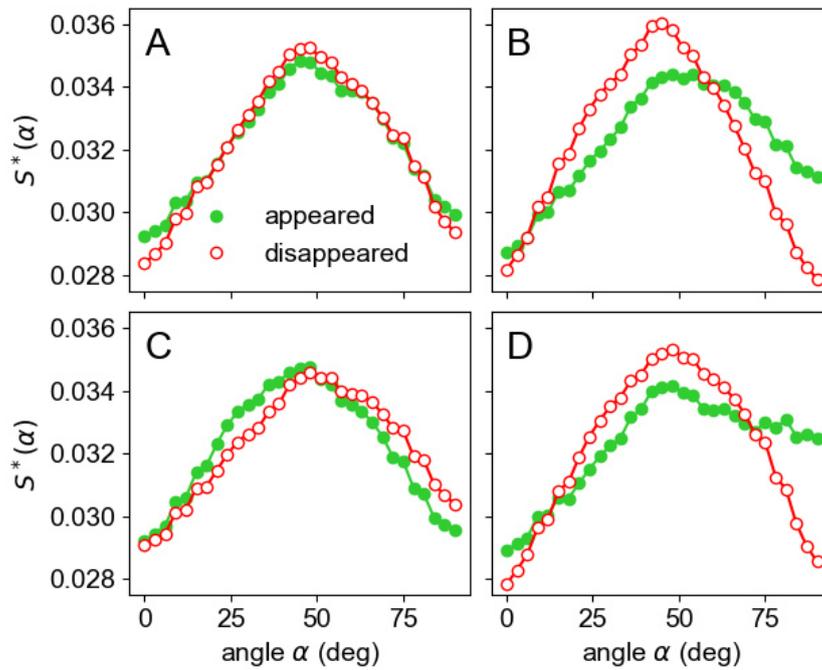


Figure 4.9: Normalized versions of the distributions shown in figure 4.8. This allows us to better read the changes in orientation.

4.2 Evaluation of the force network fractal dimension using images

A useful interpretation of the fractal dimension(s)³ of a given object is its capacity to fill the space it is embedded in. A line does not fill a 2D plane at all, whereas a solid surface (like a disk) fills it entirely (at least, within its boundary). A fractal object of dimension $D_f = 1.5$ will partly fill its 2D embedding space: infinitely more than a line and infinitely less than a disk. In the case of the force chain network in a granular medium, we can interpret its fractal dimension as the density of the network; how much it fills the granular.

We want to use it as an input of our prediction models in chapter 6, and will now describe how to measure it.

4.2.1 Minkowski–Bouligand dimension

To estimate a fractal dimension of the force chain network, we used the common Minkowski–Bouligand method [112], also called the box-counting method. This particular one is popular in experimental sciences as it is well suited to real data, limited by resolution or sampling.

Using this method to evaluate the fractal dimension of an object, we overlay on it a square tiling of tile length ϵ . We then count the number of tiles at least partly overlapping with the object of interest, $N(\epsilon)$. As ϵ goes from the object length to 0, $N(\epsilon)$ should follow a power-law, at least for a range of ϵ values. The absolute value of this power-law slope is the box dimension. For instance, the number of tiles of side ϵ covering a line grows as $N_{line}(\epsilon) \sim \epsilon^{-1}$. In the case of disk, it grows as $N_{disk}(\epsilon) \sim \epsilon^{-2}$. Indeed, the respective dimensions of these object are 1 and 2.

For any complicated shape, however, the results will likely be somewhere between 0 and 2 (in the case of an image). For the coast of Great Britain, as shown in figure 4.10, values close to 1.25^4 have been computed [113].

4.2.2 Application to thresholded images

To apply box-counting on our images, however, we must first isolate the force network. The force chain image is obtained by manipulating the color channels (empirically, it was found that subtracting either the green or blue channel to the red one

³Several different definitions and measures of fractal dimensions coexist, either more adapted to real-life data or to theoretical constructs.

⁴This may slightly vary depending on the implementation of the method and tiling used.

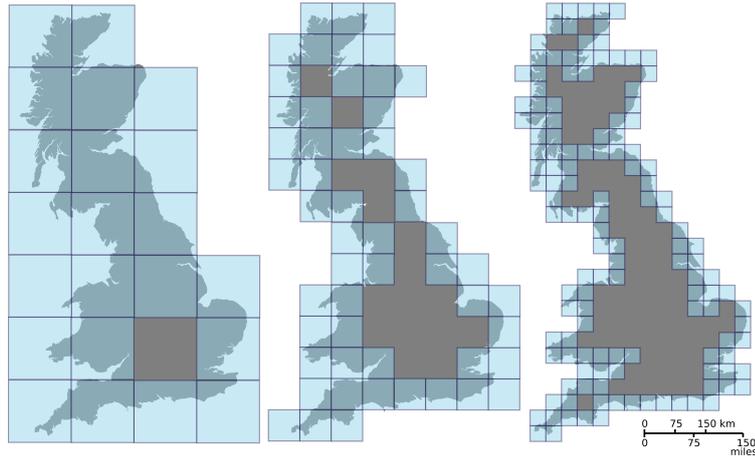
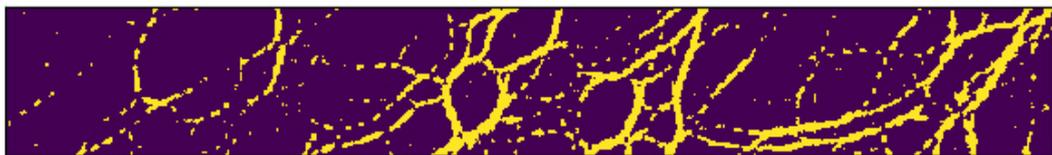
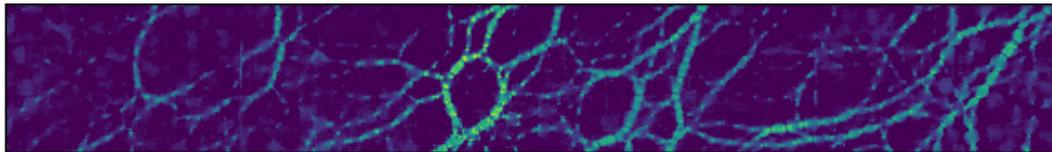


Figure 4.10: Box-counting method applied to the coast of Great Britain. As the box size reduces, more and more details are captured. *Figure adapted from Wikipedia.*

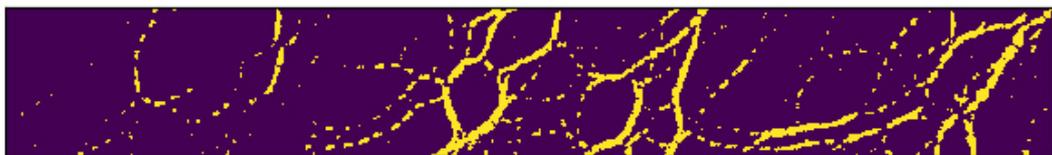
worked quite well) on a panoramic image is presented at the top of figure 4.11, as well as three thresholded versions with different cutoff values. Since the distribution of constraints in the granular is multi-scale, it is not easy to choose the threshold value - especially since it is performed on an image that has undergone non-trivial pre-processing. Presenting three instances at different cutoffs allows getting a better feel for its influence.

A limitation of the method however comes with its flexibility regarding real-life data: it is highly sensitive to resolution. To have a reasonable power-law fit, a large range of scales is needed. In our case, while the images are fairly large ($\sim 2000 \times 10000$ pixels) the smallest relevant scale for force chains is one grain (diameter of 7 mm) and the largest scale is the system size (cylinder surface of 200×950 mm). Considering we are only interested in the shear band here (photoelasticity in the static bulk is much less reliable, as it is mostly due to plastic deformations), the height of the large scale will be even smaller, around 90 mm. That is only slightly above one order of magnitude vertically, and two horizontally. Considering box-counting relies on fitting a power law, this limited range of scale will likely impact precision and interpretation.

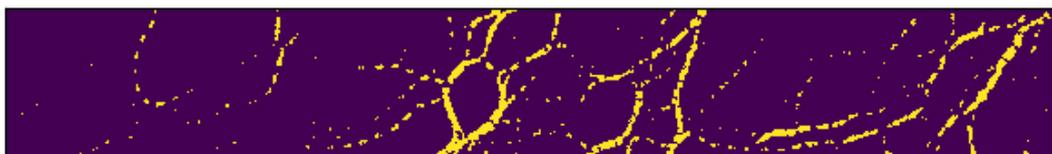
Another issue arises from the threshold process. Depending on the value, structures will be more or less sparse and one can imagine the measured box dimension will be positively correlated with lower (more inclusive) threshold values. This effect is illustrated by the images in figure 4.11. The force chain image on the top may give each of the three binary images below, with the number on the side being their respective box dimension. To check the effect of threshold value T , a reasonable range was tested for all images. This produces several time-series $Dim_T(t)$, four



1.38



1.31



1.24

Figure 4.11: Top: An image of the chains of forces obtained by subtracting the blue channel from the red channel of a partial panoramic image, about a quarter of the total image. below: Three different thresholds of the top image, with the associated box size values obtained.

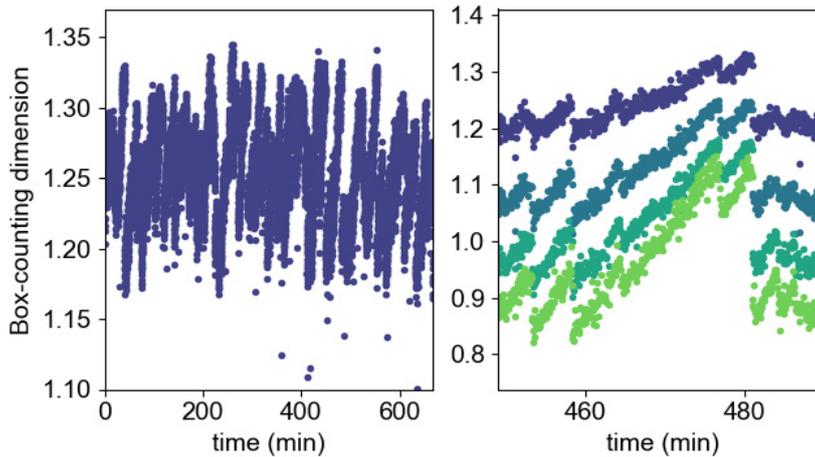


Figure 4.12: *Left*: Time series of the box dimension (clipped after ~ 11 h) for the selected threshold value. *Right*: Zoom of the figure on the left with four different thresholding. A higher threshold is represented by a lighter shade.

of which are plot in the right part of figure 4.12. The trends are very similar, but higher values (lighter shades) are noticeably jumpier and provide lower dimensions, as fewer force chains are visible.

When torque increase, it may be reflected by several quantities related to the force chains: intensity, quantity, and orientation. By comparing the measured torque with the fractal dimension, we can determine how important is the effect of force intensity on the fractal dimension.

Figure 4.14 compares instantaneous torque and fractal dimension with two different visualizations; a scatter plot on the left and a histogram on the right. The scatter plot helps evaluate the possible range of measurements while the histogram provides information about their frequency. There is a correlation between the two measurements; a linear fit is added with a dashed line on the left. However, the thickness of the scatter plot as well as the 2D histogram on the right side of the figure shows either of the signal (fractal dimension or torque) does not completely explain the other one, and there are indeed other effects at play.

The spread observed in the scattered data is not due to noise. To check this, let compare figure 4.14 and figure 4.13. On the first one, for a torque value of 9 N.m for instance, points span 1.16 to 1.30 – a range of 0.15. In the second figure, the noise amplitude of the box-counting dimension is roughly 0.03, a value five times smaller than the observed vertical spread of the histogram.

Since the measured box dimension and the torque readings are non-trivially correlated, there is exclusive information to each signal. For this reason, we will use both as inputs of our predictive models in chapter 6.

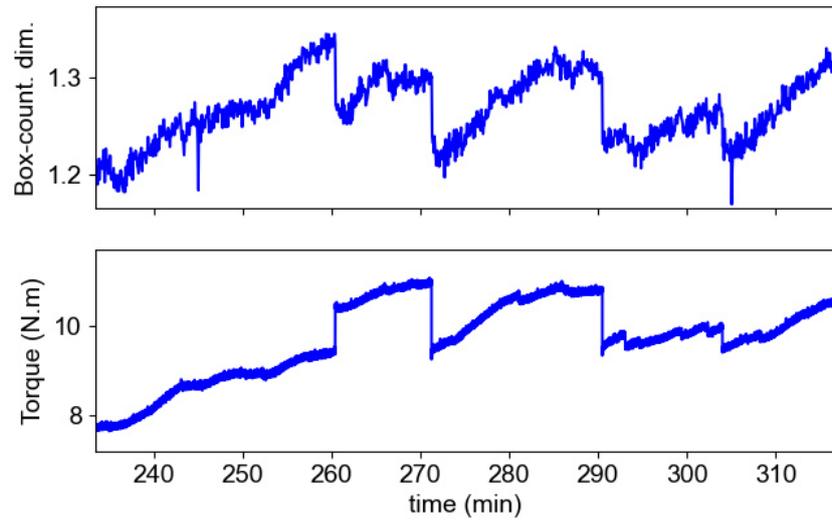


Figure 4.13: The box dimension (*top*) correlates well with the torque measurement (*bottom*), jumps appear in the same places and local slopes are similar, however anti-correlated jumps exist as well, around 260 min.

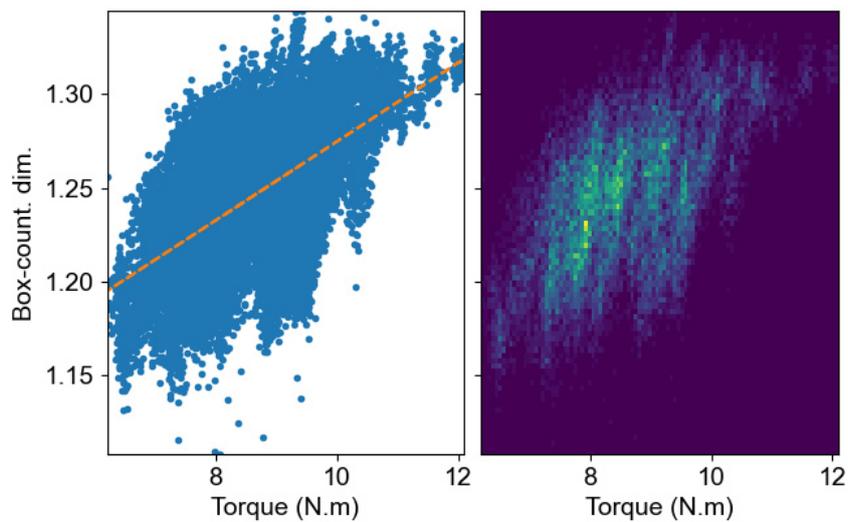


Figure 4.14: scatter plot (*left*) and 2D histogram (*right*) of the torque Γ against the fractal dimension.

4.3 Acoustic-based fractal dimension

Using acoustic wave to probe systems have been done for a long time. The concepts of sonar and echography emerged soon after the speed of sound in water was measured by Jean-Daniel Colladon and Charles Sturm in 1827. In geology, seismic waves are commonly exploited to probe the earth’s inner structure: this is how, among many other examples, Danish seismologist Inge Lehmann discovered the earth’s inner core in 1936 [114].

More recently, coda waves interferometry [115–119] has been used as a tool to detect very subtle changes in a medium through which sound waves travel and get scattered (much like an acoustic speckle), a technique pioneered in 1985 by Keiichi Aki [120]. Another line of research (more mathematical) focused on the very ambitious goal to describe the shape of a resonating object by listening to it, with the aptly named article “Can you hear the shape of a drum?” [121]. Thirty years later, an even more aptly named article answered the question; “You cannot hear the shape of a drum” [122]. (Yet another article was published, called “Can you hear the shape of a **fractal** drum ?” [123] but we’re still waiting for the answering paper.)

Regardless of the approach, knowledge of the internal state of a granular stack may help to predict its behavior [17]. In our setup, we have access to the state of each grain, as well as the force network, through photoelasticity. However, such measurements may not be available in other experiments.

In order to facilitate comparison, I tried to develop a way of probing the structure of the force chain network, by comparing measurements between pairs of acoustic sensors with a simple propagation model. It is known that acoustic waves travel preferentially through force chains in a compacted 2D granular medium [124, 125]. I will try to leverage this property with measurements of the geometric spreading of energy to evaluate the fractal dimension of the force network.

4.3.1 propagation model

First, we will consider the simplified case of a spherical wave, propagating at constant velocity c in an isotropic, homogeneous, and non-dissipative medium, in which two sensors are placed. By reason of energy conservation, we write the measured energy of the wave as:

$$E(r) = \frac{E_0 G}{r^\alpha} \quad (4.1)$$

With $E(r)$ the energy at a distance r from the source, E_0 proportional to the

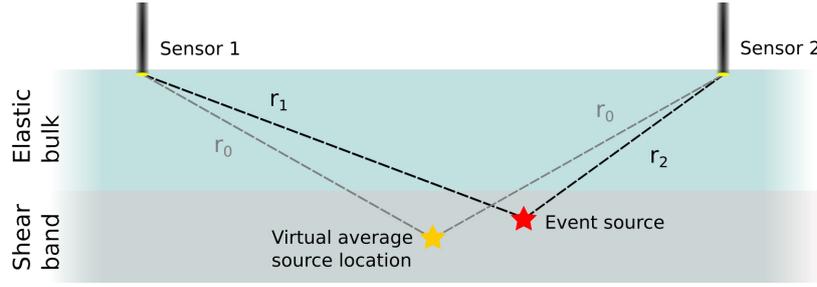


Figure 4.15: Schematic of the propagation model, in the context of the cylindrical experiment. A wave is emitted at the event source, and propagates in a straight line towards each of the two sensors.

emitted energy, G a sensor-specific multiplicative gain, and α being equal to the dimension of the propagation medium minus one.

The two different sensors are located at distance $r_1 = r_0 + \epsilon$ and $r_2 = r_0 - \epsilon$ each. A wave emitted by the source will be recorded at different times by each sensor and the time difference is:

$$\Delta t = t_2 - t_1 = \frac{2\epsilon}{c} \quad (4.2)$$

With c as the wave velocity. Let us compute the difference of the \log of energy:

$$\Delta M = \log(E_0 G_2 r_2^{-\alpha}) - \log(E_0 G_1 r_1^{-\alpha}) \quad (4.3)$$

$$\Delta M = -\alpha (\log(r_0 - \epsilon) - \log(r_0 + \epsilon)) + 2 \log\left(\frac{G_2}{G_1}\right) \quad (4.4)$$

The first two logarithms can then be linearized to give:

$$\Delta M = -\alpha \left(-2 \frac{\epsilon}{r_0} \right) + 2 \log\left(\frac{G_2}{G_1}\right) \quad (4.5)$$

$$\Delta M = \frac{2 \alpha \epsilon}{r_0} + \Delta M_0 \quad (4.6)$$

With ΔM_0 denoting the gain-related term. Finally, part of the first term can be substituted with the expression found in equation 4.2 to find the following relation:

$$\Delta M = \frac{\alpha c}{r_0} \Delta t + \Delta M_0 \quad (4.7)$$

Provided we can detect two events across each of the two sensors and know r_0 , one can then perform a linear regression to measure α . The advantage over a single sensor measurement, where one could fit equation 4.1, is the absence of reliance on absolute time or distance: both of which are inaccessible in practice. The r_0 parameter still is an absolute distance in some sense but is easier to obtain (or,

at least, approximate) thanks to the bounded size of our system: one can simply take the barycenter of all the possible event sources. In practice, this is halfway through the shear band (vertically) and between the two sensors (horizontally), as represented with a yellow star in figure 4.15. Due to dissipation, $\alpha + 1$ can only be a lower bound for the medium dimension. This effect will be significant in the next section.

4.3.2 A simpler case: sound propagation in air

To verify this simple model, a series of tests were performed in air with the same acoustic sensors used in our granular experiment. These tests were done in collaboration with the *Luminescence Team* of the ILM.

A plasma is generated by a focused laser, pulsing at 9 Hz. The sensors are placed to be co-planar with the plasma (to make measurements easier) and radially aligned (to maximize gain), as schematized in figure 4.16. Both sensors are initially 25.1 mm away from the plasma.

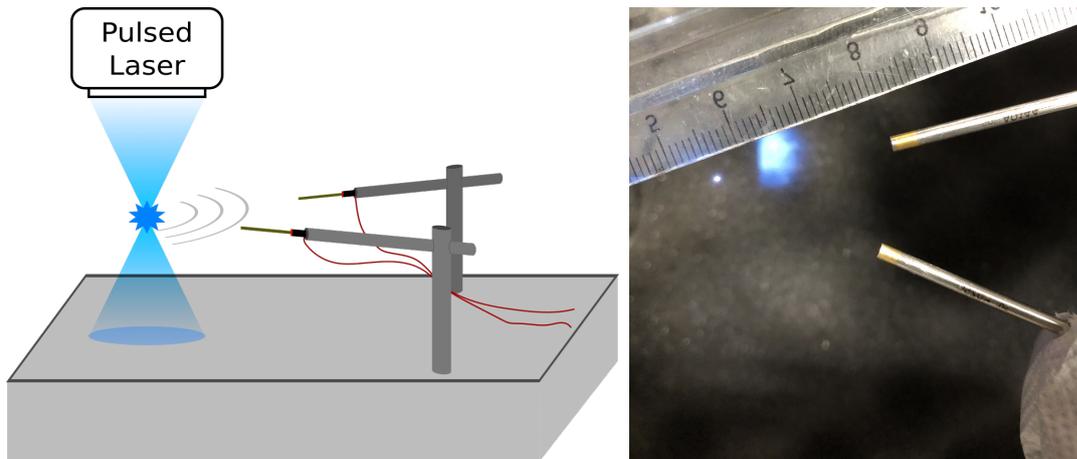


Figure 4.16: Left: Schematic of the air propagation setup. The laser is used to make a plasma mid-air, which creates a sound wave when collapsing. The wave is recorded by two co-planar acoustic sensors pointing towards the source. Right: Picture of the setup from above, with the plasma appearing as the smaller, saturated light.

The first sensor is left in place, while the second one is placed further and further away, up to 75 mm away from the plasma. Around 100 laser pulses are recorded and averaged for each of the 8 distances tested.

We can first measure the current sound velocity, as it is a parameter of the model and depends on temperature or humidity. The good alignment and small intercept value lend weight to the measurements.

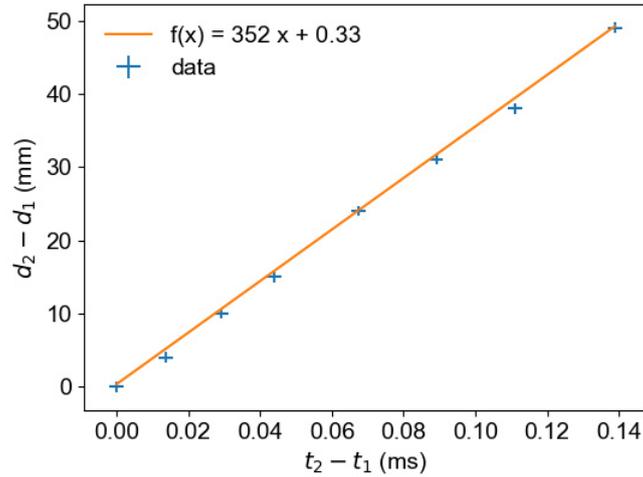


Figure 4.17: Distance between the two sensors plotted against arrival time difference. The linear fit gives us the sound velocity in the environment of the experiment.

In addition, we can verify the propagation law form we assumed at the beginning, in equation 4.1. In figure 4.18, the energy is plotted against the sensor distance to the source. The power-law assumption is well verified, with all points close to the linear regression. However, we can see a significant deviation from the exponent of 2 we expected. A larger exponent means we observe less energy than anticipated. Without dissipation, the energy should decrease due to geometric spreading only. Observing a faster decay can be interpreted as a signature of dissipation, causing energy to be lost as the wave travels. Over the range of value observed, however, a power-law still well describes the data. In these conditions, the dissipation effect can be encompassed in the exponent.

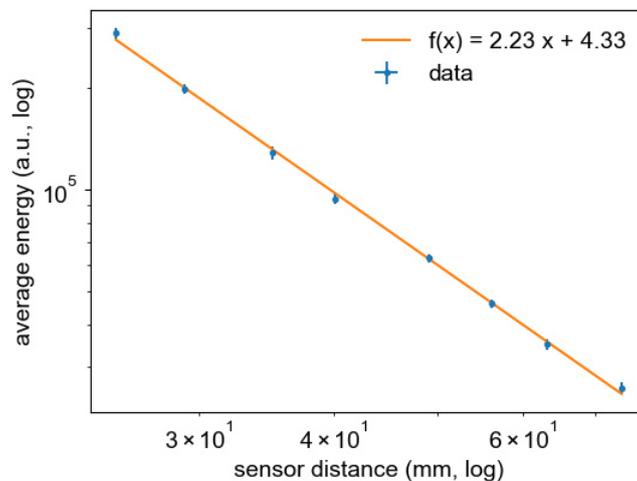


Figure 4.18: Measured energy as a function of the moving sensor's distance to the source. The linear fit is again very good but the slope is significantly different than the expected 2 of the ideal, non dissipative case.

We can now apply the two-sensor method to the data. Fitting equation 4.7 produces figure 4.19, with a solid line representing the fit, of slope 2.46. For comparison purposes, lines with slopes 2.23 and 2 have been added, respectively corresponding to the single-sensor measurement shown in figure 4.18, and to the ideal case.

The dimension of the propagation medium, in this case, is definitely not 3.46 but 3. This significant deviation already put a limitation on this technique. While the dissipation effects could likely be modeled, predicted, and accounted for in air, it will be more complicated in our granular as it is a more complex medium than air. One of many possible reason for this deviation is the spherical wave assumption hidden in equation 4.1, which may not be true considered the technique used to generate the wave, a collapse. Still, let us see what we can learn from applying this method to labquake data.

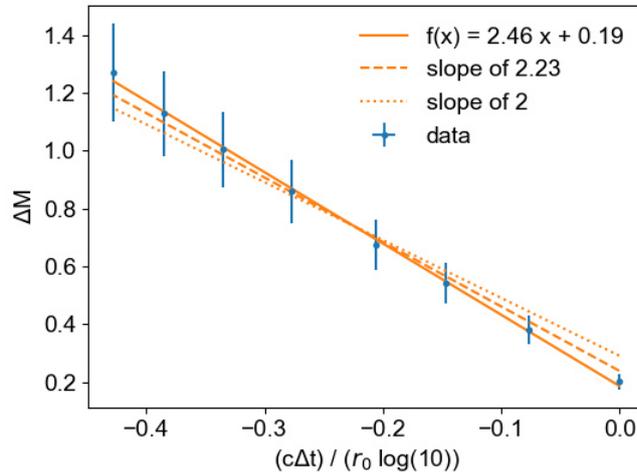


Figure 4.19: Linear regression results of air measurements. The regression, in solid, provide an effective dimension of 3.46. For comparison, the slope obtained with a single sensor in figure 4.18 (dashed) as well as the ideal non-dissipative case (dotted) have been represented.

4.3.3 Application to our experiment

To use equation 4.7 in our system, knowledge of c and r_0 is required. C can be estimated with the calibrations performed in [125]. As for r_0 , it is approximated as described in the last paragraph of section 4.3.1 and represented in figure 4.15: in the middle of the shear band, halfway between the sensors.

In the shearing experiment, the acoustic sensors are regularly placed in the top ring. From figure 4.20, we can see there are three “types” of pair we can choose from. For

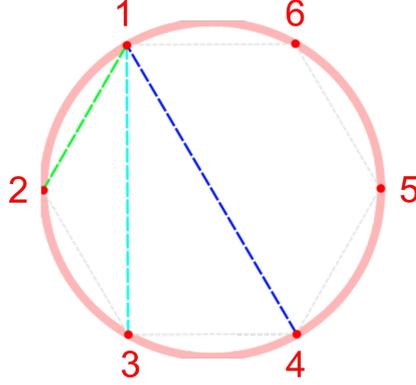


Figure 4.20: Three types of pairs can be made from points in a hexagon. Only in the case of diametrically opposed sensors the “left” and “right” side of the sensor pair are equivalent.

symmetry reasons, we will only match events of pairs 1-4, 2-5, and 3-6, as the left side and right side of a sensor pair are equivalent only in this case.

To match events between two sensors, I used the procedure described earlier in section 4.1.2, when matching force and torque discontinuities, based on time differences. By selecting a sufficiently high acoustic energy threshold and a sufficiently tight time difference tolerance, we can pair measurements between two sensors one-to-one.

For each pair, the amount of matched events are $n_{(1,4)} = 5337$, $n_{(2,5)} = 3839$ and $n_{(2,5)} = 4466$. However, we cannot consider this to be a dataset of $n_{tot} = 13642$ event and subject it to analysis, as the ΔM_0 value is different for each pair. For this reason, three different regressions must be performed, as shown in figure 4.21.

The fitted slopes are $-4.3 \cdot 10^{-3}$, $-2.8 \cdot 10^{-3}$ and $-3.3 \cdot 10^{-3}$ for pairs 1-4, 2-5 and 3-6 respectively. Using the values $r_0 = 50$ cm and $c_0 = 800$ m.s⁻¹ we can compute the dimension for each pair ; $d_{(1,4)} = 1.62$, $d_{(2,5)} = 1.40$ and $d_{(3,4)} = 1.48$ respectively. These values are not too far off the range measured using the images (1.15 – 1.35) but are all higher. This is expected as we compound the effect of geometric spreading with energy dissipation.

In the tests done in air, the noise level was very low and the sound waves well defined. In the case of our granular, however, the analysis is more difficult. Not only the sound waves are much more complex, but the model proposed in section 4.3.1 is even more of an approximation in this case. The sound velocity is no longer constant and the location emission can be significantly away from its supposed average location. In addition, this model only introduces dimensional effects in the geometric spreading of the energy, assuming linear propagation. Since sound waves preferentially propagate through forces chains, this last assumption is known to be

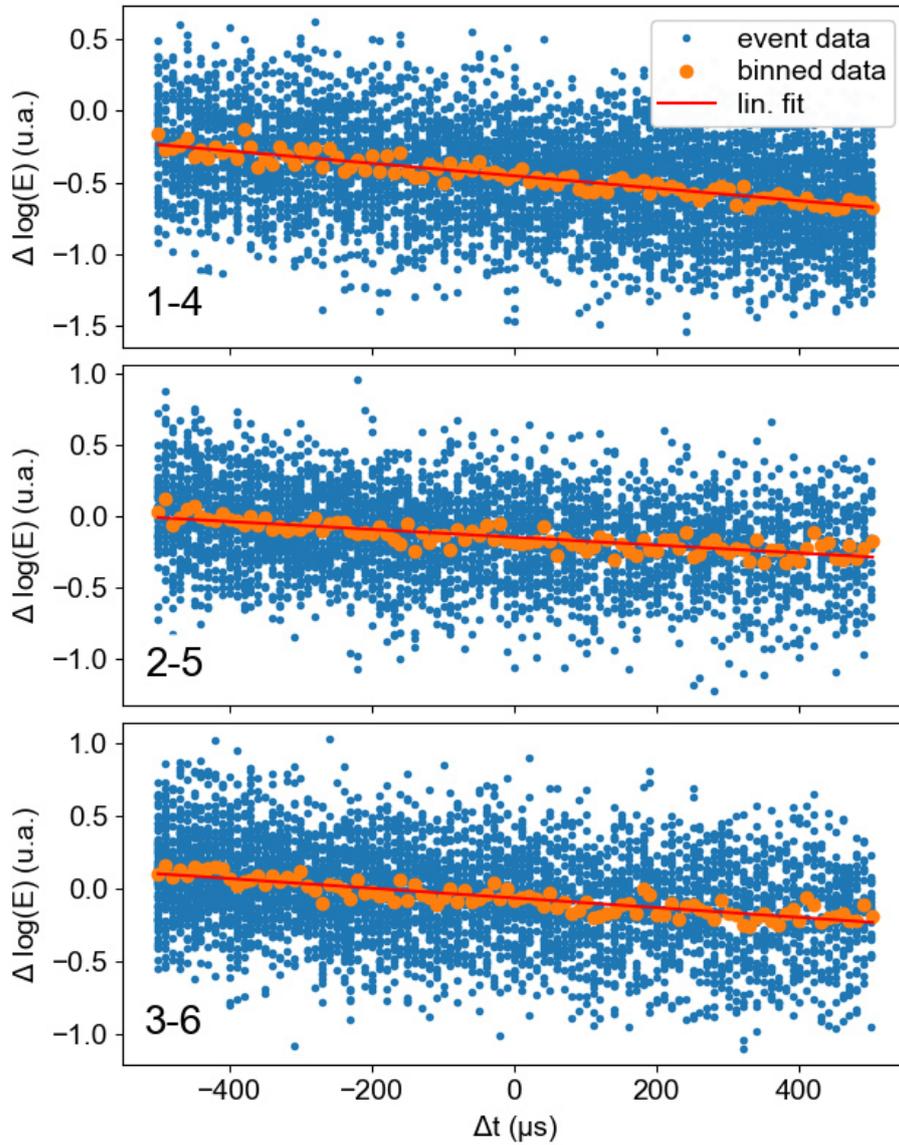


Figure 4.21: Regression assuming expression 4.7 for matched labquakes. The raw data has a large variance and is binned before being fitted. 22h of data were used.

false. This could be accounted for using images and detecting possible paths for the sound to travel on (which will be done in chapter 5) but our objective here was to evaluate whether this purely acoustic-based method could work.

My objective was to devise a method using only acoustic data to provide information about the force chains structure (here, a fractal dimension) and be less reliant on images. This way, future prediction results could possibly be generalized to setups where images are not available. However, the simplicity of the model and the difficulty to properly measure the necessary quantities (time of arrival and energy) with sufficient precision limits its current usefulness. While the images provide a data point every four seconds, this method needs at least several hours worth of acoustic data to average out the different sources of variance in figure 4.21. Finally, the different readings of each pair cover a large range of value ($d_{1-4} = 1.62$ and $d_{1-4} = 1.40$, a span of 0.22). This is as large as the range cover by all the image measurements.

Chapter 5

Acoustic Location of Events

Contents

5.1	Acoustic location and hyperbolic methods	76
5.2	Setup and method	78
5.2.1	Experimental design	78
5.2.2	Ballistic propagation model	80
5.2.3	Time-of-Arrival Measurement	82
5.3	Evaluating source likelihood	83
5.3.1	Single branch likelihood	83
5.3.2	Agglomerated likelihood	84
5.4	Event location performance	86
5.4.1	Border impacts	86
5.4.2	Bulk sources	88
5.4.3	Limitations and Future development	91

5.1 Acoustic location and hyperbolic methods

Prediction is a major goal of the labquake project. In order to forecast re-arrangement events, knowledge of the local properties of an event nucleation site would be helpful. In other works¹ [124, 126], fast camera recording (100,000 Hz and 49,000 Hz respectively) could locate event sources. However, in our shearing experiment, the image acquisition frequency is only 1/4 Hz. Thus, a different method is required for our cylindrical setup.

The method presented in this chapter was inspired by *hyperbolic sound location*. It uses time-difference-of-arrival across an array of acoustic sensors, coupled with a propagation model, to evaluate where the recorded sound may come from. In 2019, Ange Haddjeri (Master 1 student) did an internship in our team to help develop this technique.

The history of hyperbolic methods to locate unknown signal sources takes its roots around World War I [127, 128]. The objective was to detect the position of enemy artillery by recording their firing sound. Using two microphones (located at M_1 and M_2 respectively) one can observe the time difference of arrival $T_2 - T_1 = \Delta t$. Considering sounds travel at a known velocity c , the artillery location A obey the following relation:

$$c\Delta t = AM_2 - AM_1 \quad (5.1)$$

This is very similar to the equation of a hyperbola: the only difference is the lack of absolute value around the distance difference, which is due to the fact we know the sign of Δt . The locus of points A that satisfies equation 5.1 is a single hyperbolic branch and is not enough to pinpoint the source of the recorded sound. However, with another microphone, three hyperbolic branches can be computed. In theory, all branches should cross on a single point. Similar equipment was also used to detect enemy planes, for instance by the Japanese army during World War II, as seen in figure 5.1.

Another historic usage for hyperbolic methods is navigation and self-positioning. In such systems, the roles are reversed: M_i are static, cooperative emitting radio beacons and A is a navigating agent. Examples include the American *LORAN* [129] or the international *OMEGA* [130], which was shut down in 1997 with the generalization of the GPS.

¹the second article was done by our group.

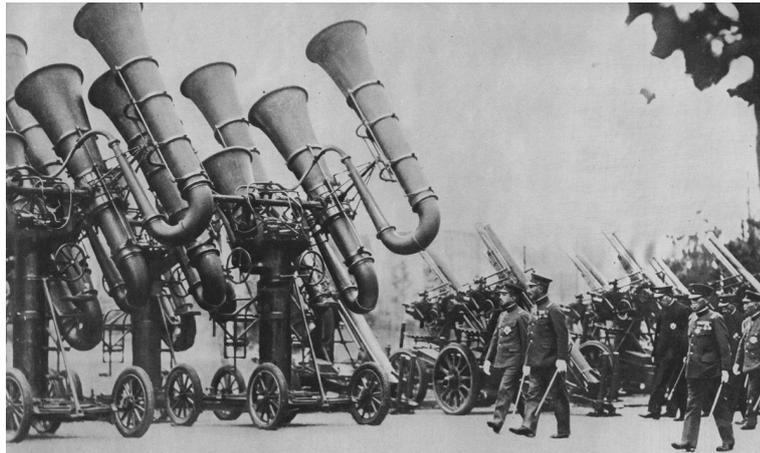


Figure 5.1: Japanese Emperor Hirohito inspecting “war tubas” (officially named *Type 90 large aerial sound detector*) with anti-air guns in the background. Date unknown but likely to be around the early 1930s, as this equipment was commissioned in 1932. Sourced from *Wikipedia*.

A granular medium however is neither continuous, isotropic nor homogeneous. Sounds travel through grains that do not fill space entirely, there are gaps in the structure. The contact area between two neighboring grains is very small, further limiting the propagation possibilities. In addition, differently stressed grains transmit sounds at different speeds (see [124, 125]). Finally, if the experimental setup is not carefully designed, the sound may even travel in the frame/structure holding the grains. Sound velocities in glass (~ 4500 m/s) or aluminium (~ 6300 m/s) is almost an order of magnitude higher than it is in grains (~ 800 m/s).

In two-dimensional arrangements and with a low number of grains, as is the case in our experiment, the above points will be addressed both in the **propagation model** and the **design of the experiment**.

5.2 Setup and method

Before trying to locate events in our main shearing experiment, we need to develop and test methods. For this purpose, I built another, simpler setup (shown in figure 5.2) as a testbed for different acoustic location methods.

5.2.1 Experimental design

Granular material: The granular material used here is the same as the ones in the main shearing experiment (photoelastic and bi-disperse), and is arranged in a single layer of roughly 2200 disks. Among the standard grains are eight black hollowed grains that can be observed slightly above the center of the pile; these are active grain prototypes that will be discussed later in section 5.4.2.

Cell design and geometry: The grains are sandwiched between two vertical and parallel acrylic plates 4.4 mm apart. This spacing is 110% of the thickness of the disks to minimize contact and friction between the grains and the walls, therefore avoiding parasite waves that could travel through them.

The bottom and side walls are fixed and made with several layers of insulating foam and paper strip to dampen incident sound waves.

The upper boundary is mobile, allowing the compression of the granular medium with a weight. This weight is constituted of weight disks (up to 40 kg) and a variable weight (0~20 kg) achieved with a 20 L water tank that can be slowly filled or emptied. The confinement cell initial area is 335×295 mm (W \times H) as schematized in figure 5.3, but can decrease in height when the water tank is filled.

Measurements: Holes have been drilled through the top and sidewalls to insert acoustic sensors. These sensors are in turn inserted into fixed grains, to maximize acoustic coupling with the bulk of the structure. The coupling is further augmented by silicon oil around the pinducersTM. The acoustic sensors are the same as the one used in the shearing experiment, but the sampling frequency is 25 times higher (2.5 MHz) to maximize accuracy on the time of arrival. In addition, pictures of the granular were taken every ~ 1 s using a NIKON D750.

To calibrated our location method, a controlled acoustic emission of a known source is needed. The first source I used is a thin metallic rod connected to a permanent magnet shaker (*LDS V201 from Brüel Kjaer*) that hits one grain at the bottom of the granular ensemble. The shaker is driven with a squared signal, to make a sharp impact. The second source I (tried to) use is active grains, which can emit sounds. Their design will be explained later in section 5.4.2.



Figure 5.2: Picture of the experimental setup. The 6 black circles are “active grains” prototypes. The black bar above the granular is 10 cm long.

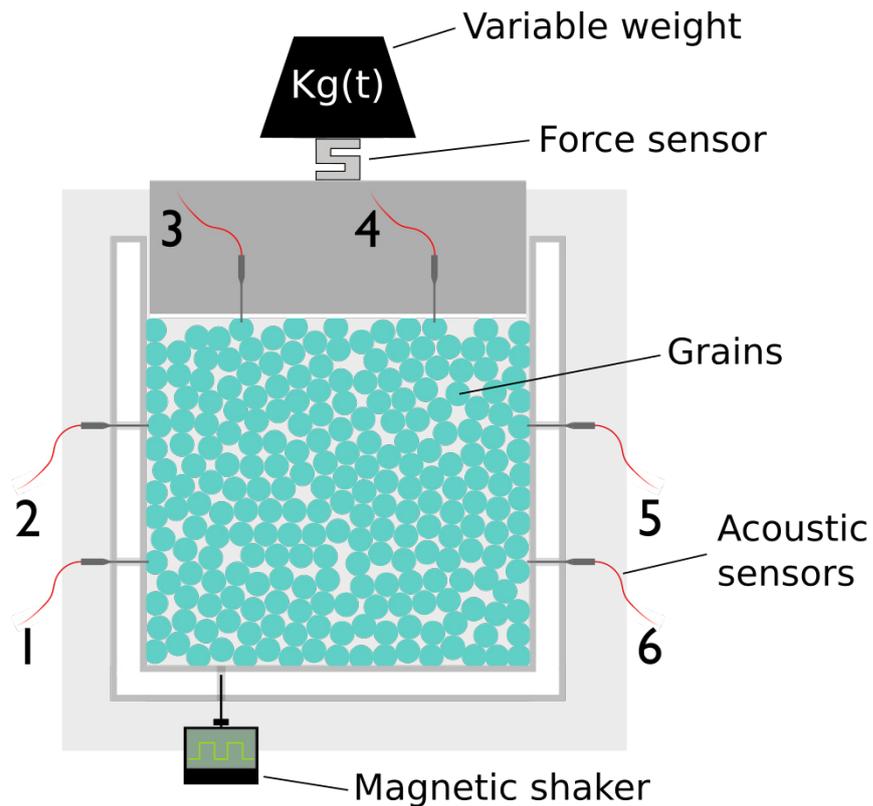


Figure 5.3: Schematic of the experimental setup. Grains are arranged in a 2D layer held by a \sqcup shape made of several layers of insulating foam. Artificial waves can be created by a shaker below the experiment, hitting one grain with a metallic pin, driven with a square signal.

5.2.2 Ballistic propagation model

To account for the issues raised in section 5.1 (discontinuities of the medium and varying local sound velocities), two key ingredients were used to develop the propagation model. The first one is the topology of the pile: we only allow for sound propagation within disks. To travel from one disk to the next, the sound wave must go through their common contact point. This will account for the lack of continuity of the propagation medium. The second one is heterogeneous sound velocity, caused by the heterogeneous force chains distribution [83]. For each grain, its stress level can be evaluated using photoelasticity and its inner sound velocity computed.

Considering these two points, I decided to model the disk layer as a network. This network is built in two steps: first, the nodes are defined as the contact points between the grains. Then, all the contact points of a grain are fully linked by edges. Sound waves travel from node to node, through edges. A small part of this network view is represented in figure 5.4.

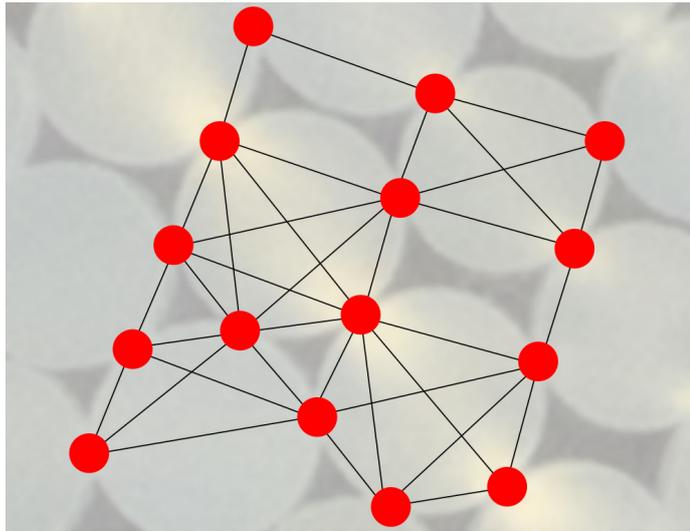


Figure 5.4: Schematic of the network model overlaid on an experiment picture. Contact points are the nodes, linked by edges within grains.

Assuming a network modeling, one might wonder if considering the grains as nodes, instead of the contact points, would be a more natural choice. Picking contacts over grain centers as nodes is motivated by several reasons. First, the origin of an acoustic emission is expected to be grains sliding or hitting each other. If so, it makes more sense to look for the source between grains, at contact points. In addition, edges between grains would span over two grains, making velocity evaluation more complex. Finally, for n grains, there is a total of $n(t_c/2)$ contacts,

with t_c the average number of neighbors touching a single grain. Typically, this number is around 4.6 in our experiments. This means that a network using contact instead of grain will have more nodes. We can expect a model with more nodes may be (slightly) closer to a continuous description of the experiment, and thus more accurate.

As explained in section 2.3.2: *Image Exploitation*, the grains are detected through the Hough circle detection method. To decide whether two grains are touching, we compare the sum of their radii to the distance between their centers. If the two values are less than 5% apart², their contact point is added to the graph. Two contact points are linked by an edge if and only if they are from the same grain.

To compute local propagation velocities, we first need to evaluate the stress. This is done using the calibration described earlier in section 2.3.2. To translate stress into velocity, we use the relations described in [125] between applied force F and sound velocity c in the case of a disordered array, that is: $c \sim \sqrt{F}$. The exact velocity is calibrated using the form $c = \alpha\sqrt{G^2 + \beta}$, with α and β values fitted for this setup.

Finally, by dividing an edge length by its associated velocity, we can obtain the wave flight time along that edge. In order to estimate the propagation time of the wave ballistic front from a given node to a given sensor, we compute the fastest path between the two points in this flight time network. This is done using the Dijkstra algorithm [131]. For a given potential source point, the difference in arrival time between two sensors can be expressed as:

$$\Delta t_{model} = \sum_{n=0}^N \frac{l_n}{v_n} - \sum_{m=0}^M \frac{l_m}{v_m} \quad (5.2)$$

with N and M being the number of nodes of the fastest paths for each of the two sensors, v_i and l_i being the edge velocity and edge length. The resulting Δt_{model} is dependent on up to hundreds of measurements depending on how far the considered source is, which is a large source of uncertainty. This will have an important effect when evaluating source likelihood, discussed in section 5.3.1.

Beside this uncertainty, the structure of the model makes its output actually fairly robust. Since the flight time between two points only depends on the edges along a single path, the value of every other path is irrelevant. The only way for the fastest flight time to change when errors are introduced to the network is for another path

²This tolerance may seem large but with our image resolution, it translates to 3 pixels.

to emerge as the fastest. To appreciably change the results, this new “false” fastest path should be significantly faster than the original one.

5.2.3 Time-of-Arrival Measurement

Before measuring the time of arrival, each signal is first de-noised. To do this, I use a seventh sensor that was recorded with the other six but isolated away from the experiment. This sensor does not pick up any significant acoustic signal, but its recorded data is subject to the same electrical noise. By subtracting this signal from the other useful ones, part of the noise can be removed.

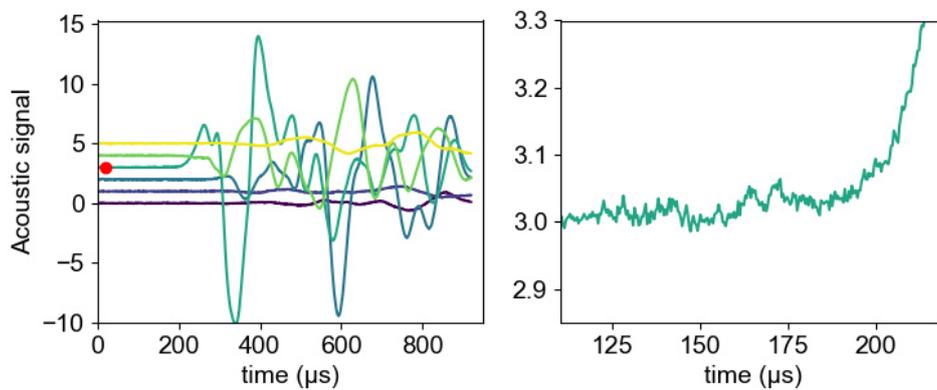


Figure 5.5: *Left:* acoustic signals from the six sensors shifted for readability. *Right:* A zoom of the signal starting with a red dot, around its ballistic arrival time.

The timing measurement is then done manually on the signals, an example of which can be seen in figure 5.5. From the zoomed signal on the right, one can see it may be difficult to exactly pinpoint the arrival time of the ballistic signal, and no satisfying automatic way to detect it was found. To account for the difficulty of measuring the time-of-arrival and the inherent variability that comes with human measurements, the model described below takes both a timing value and an uncertainty value as inputs. For instance, the zoomed signal shown in figure 5.5 corresponds to a time-of-arrival of 190 ± 20 *mus*.

5.3 Evaluating source likelihood

We wish to use our six time-of-arrival measurements to locate the acoustic source. First, I will detail the method for a pair of sensors. How to agglomerates the measurements over all possible pairs of sensors (effectively using all the available information) will be explained afterward.

5.3.1 Single branch likelihood

To evaluate whether a point A_k in the network can be the source of a measured event according to a pair of sensors (i, j) , we want to know how “far” that point is from the pseudo-hyperbolic branch stemmed from that sensor pair.

We can model both Δt_{model} (computed from the model) and Δt_{sensor} (measured on the signals) as independent random variables each drawn from normal distributions $\mathcal{N}(\widehat{\Delta t}, \sigma^2)$, centered on a true value $\widehat{\Delta t}$ and with variance equal to its estimated uncertainty σ . For Δt_{sensor} , the uncertainty is due to the difficulty of measuring the timing of arrivals described in section 5.2.3. For Δt_{model} , it comes from the distance and (mostly) the local velocity estimation. We then define the difference between the measured and expected Δt values, $E = \Delta t_{sensor} - \Delta t_{model}$. This difference E can be in turn be described as a normal random variable $\mathcal{N}(\widehat{E}, \sigma_{tot}^2)$, with \widehat{E} the true value and $\sigma_{tot}^2 = \sigma_{model}^2 + \sigma_{sensor}^2$, respectively estimated using uncertainty on speed and distance measurements, and on time of arrival measurements.

We wish to estimate how likely it is that \widehat{E} is close to 0 – or more formally, that $|\widehat{E}| < \epsilon$ for a given small and positive ϵ :

$$\tilde{P}(|\widehat{E}| < \epsilon) = \int_{-\epsilon}^{\epsilon} \frac{1}{\sigma_{tot} \sqrt{2\pi}} e^{-\frac{1}{2}(\frac{x-E}{\sigma_{tot}})^2} dx \quad (5.3)$$

The above value depends on which pair of sensors i, j and which node k are being considered, and can be written as $\tilde{P}(i, j, k)$. We drop ϵ from the notation as it only plays a role similar to a contrast ratio, but does not change the ordering of the source candidate probabilities, unless taken to extreme values (typically not more than two orders of magnitude away ($\times 0.01 \sim 100$) from typical σ values). When σ_{tot} becomes very large, the probability becomes uniform across the values of \widehat{E} . This can be understood as the model does not have precise enough information to discriminate between candidates. To make \tilde{P} a proper probability, we can normalize it:

$$P(i, j, k) = \frac{\tilde{P}(i, j, k)}{\sum_{l=0}^{N_c} \tilde{P}(i, j, l)} \quad (5.4)$$

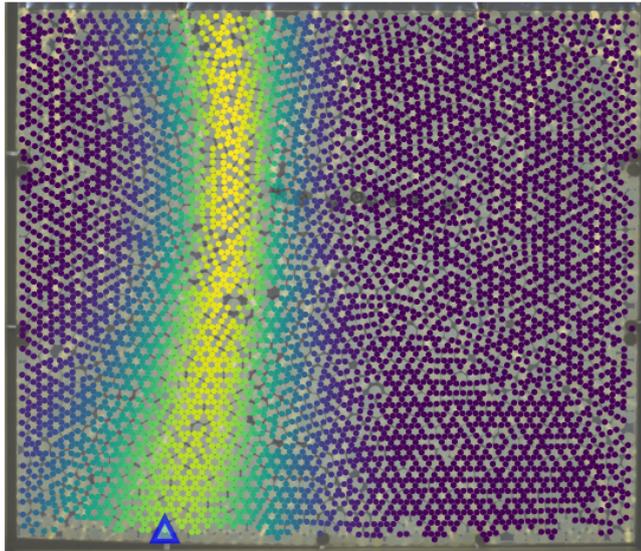


Figure 5.6: Probability field according to the sensor pair (2, 5). Higher values represent higher probability. The maximum ridge is roughly shaped like a hyperbole, spreading as it extend deeper in the granular. The blue \triangle marker indicates the actual source.

With N_c the number of contacts. Figure 5.6 represents this probability field overlaid on a photo of the experiment for sensor pair (2, 5). As expected, the maximums are similar to a hyperbolic branch, with the thin end between the two sensors. The branch widens as it gets further away from the sensors, due to the distance-dependent uncertainty.

Since a single branch is insufficient to locate the source, we need to combine the branches obtained from different pairs of sensors.

5.3.2 Agglomerated likelihood

We have seen how to derive a probability for node k from a pair of sensors i, j , but 15 different pairs ($n(n-1)/2$, with $n=6$) are potentially available. However, the mathematical methods presented below favor statistical independence between the constituent probabilities. For this reason, each sensor will only be used once, for one hyperbolic branch. Instead of the potential 15 pairs, only 3 will be used in the following analysis. In practice, I observed mixing sensors did not provide a significant performance advantage but made it harder to interpret the results: a poorly measured (due to a low signal to noise ratio, for instance) time-of-arrival on a sensor creates $6-1=5$ “wrong” branches, obscuring the graphs.

Several methods to combine probabilities from different sources (called *pooling*)

exists [132]. The best choice will depend on the desired properties of the pooling method as well as the input distributions properties. In our case, we want to locate an acoustic source. Considering this, and the structure of the granular, we expect an output distribution with a single maximum. Moreover, due to the uncertainty effects, the distribution from single pair can only discriminate sources at short range; at long range, the probabilities are less meaningful. This hints to the fact that we should interpret a single-branch probability distribution as *excluding* rather than *including*. Low probability values have a strong meaning, as they exclude source points, but high values may simply mean there is not enough information to eliminate a particular point.

Based on these considerations, I chose a product based pooling with a formula of the form:

$$P_{A(prod)}(k) \propto \prod_{ij}^{pairs} P_{ij}(k)^{w_{i,j}(k)} \quad (5.5)$$

with the pooled probability P_A being a product combination of each pair probability. Weights are path-dependent: $w = w_{i,j}, k$. Product pooling does not impose conditions on the weights – in particular, they do not necessarily have to sum to 1. The weights are dependent on both the pair of sensors i, j , and the node k . The first dependency makes sense as not all time-of-arrival measurements are equally precise. The node dependency is motivated by the uncertainty described in section 5.2.2. When propagating to points further away from a sensor, errors accumulate and may counterbalance the gain from an otherwise “good” pair with accurate signals. To favor more precise information, weights have been set to $w_{i,j}(k) = 1/\sqrt{\sigma_{tot}}$. Remember that P itself is dependant on σ_{tot} .

This pooling formula fits our needs: if a given point is strongly voted against ($w_{i,j}(k) \sim 1$ and $P_{i,j}(k) \ll 1$) by a constituent sub-probability, the overall result will be very low regardless of the other votes. Product pooling also generally gives unimodal distributions [132]. For comparison purposes, I will also present the results of another popular and intuitive method, a weighted arithmetic average:

$$P_{A(avg)}(k) \propto \sum_{ij}^{pairs} P_{ij}(k) * w_{i,j}(k) \quad (5.6)$$

This average has an opposite behavior regarding high and low constituent probabilities. Rather than selecting nodes that have not been strongly voted against, like product pooling, this method selects nodes that have been strongly voted for.

5.4 Event location performance

Now that the model has been described, let us look at its performance in two different cases: controlled and spontaneous emissions.

Throughout this section, I will add to the previously-set convention of using a blue Δ marker to indicate the real source a red ∇ marker, to indicate the predicted source.

5.4.1 Border impacts

As described in figure 5.3, a pin, mounted on a magnetic shaker is inserted below the granular stack. The shaker is driven by a square wave to impact a fixed grain at the bottom of the pile.

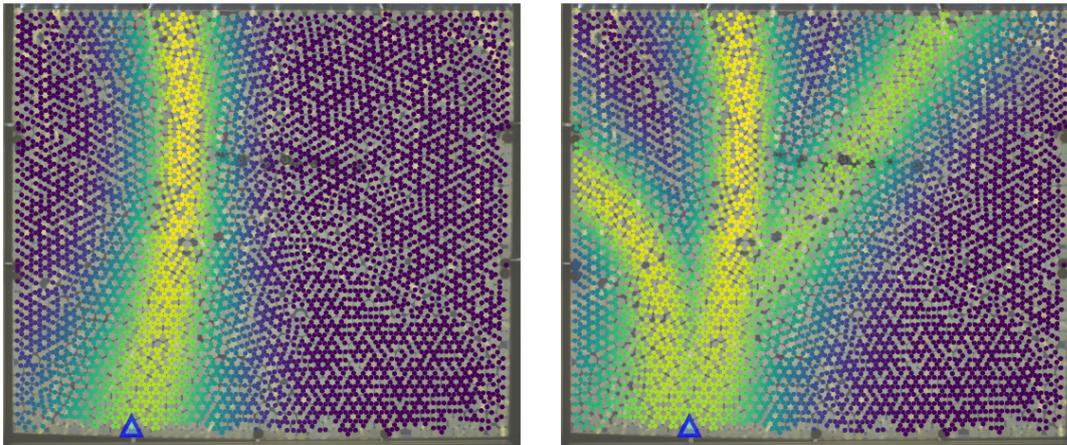
Figure 5.7a represents the likelihood obtained from a single pair of sensors (2,5). On the right-hand side (fig. 5.7b), three such branches are “visually stacked” together. This second representation is obtained with the following formula:

$$P_{A(stack)}(k) = \mathbf{max}_{(i,j)} (P_{ij}(k)^{w_{i,j}(k)})$$

This formula has no meaning in terms of probability pooling but provides a useful visual representation. For instance, we can observe that the central branch appears thinner than the lateral one, indicating lower uncertainty and thus higher relative weight. The three branches all converge to the same zone, near our impacted grain. Since no pooling is performed here, there is no predicted source shown on the figure, as it would be meaningless.

The resulting prediction of both (product and average) pooling is shown in figure 5.8. The distribution obtained through product pooling (fig. 5.8a) produces a maximum very close to the real source, about 2 grains apart. In comparison, the average pooling method result (fig. 5.8b) is about 6 grains apart. Moreover, the distribution is narrower around its maximum. This is a consequence of the nature of average pooling, which selects *voted-for* nodes. Product pooling produces a more “cautious” distribution.

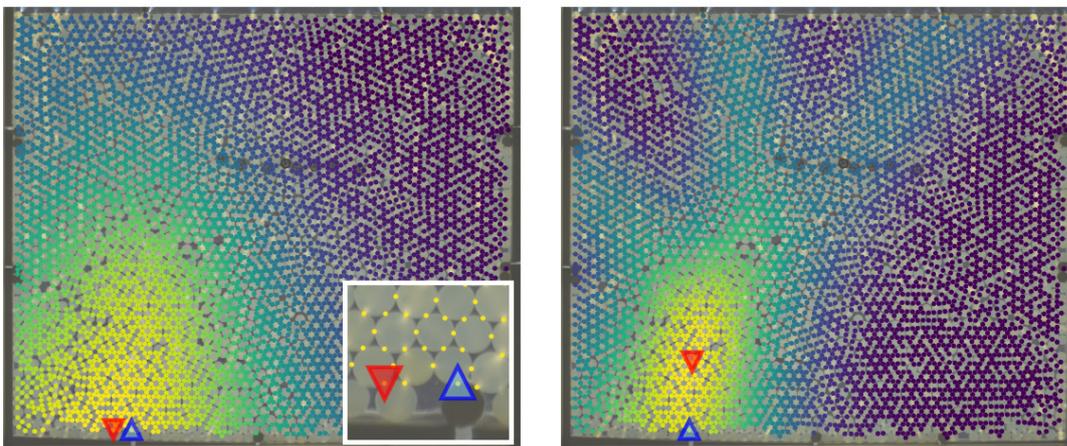
Since the excitation is controlled and the granular well compacted, the acoustic waveform varies very little between emissions, making these results strongly repeatable. Border impacts are however very much *edge cases*. Since we only look for sources within the network, perhaps the prediction model would rather (wrongly) elect a point outside the granular but is kept in bounds by the network structure. To further test this model, we need embedded controlled sources.



(a) Single branch likelihood
from sensor pair (2,5)

(b) Stacked branches from
pairs (1,4), (2,5) and (3,6).

Figure 5.7: Single (*a*) and triple (*b*) branch likelihood. The controlled impact is located at the bottom.



(a) Product pooling with a zoomed
insert, based on formula 5.5

(b) Average pooling, based on
formula 5.6.

Figure 5.8: Comparison of product (*a*) and average (*b*) pooling for controlled border impacts. An insert in figure (*a*) shows how close the prediction is to the actual source.

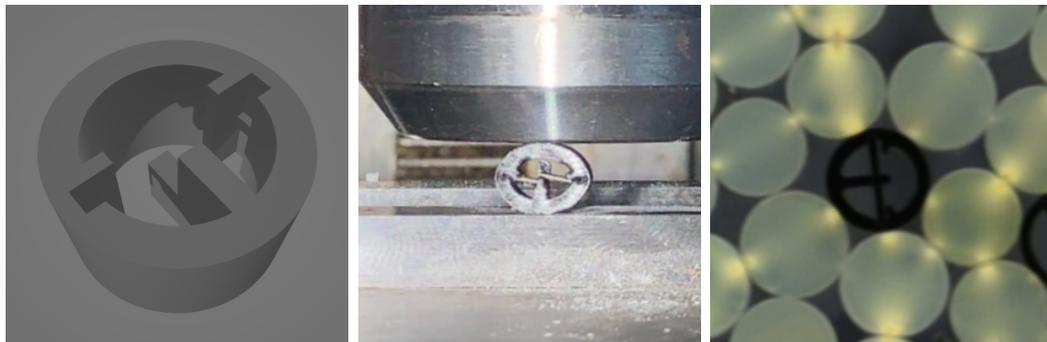
5.4.2 Bulk sources

Two bulk-based emissions were sought after to further test the method: artificial sources and spontaneous emissions.

Artificial sources

The design requirements on a controlled acoustic source deep in the granular are hard to meet. The device needs to be no thicker than a grain to fit between the plates and small enough not to disturb the force network. The emitted sound should have a well-defined ballistic front (it should be “snappy”) and ideally have a controlled trigger. Several designs were tested and none worked in the experiment. I will however describe our latest attempt that fits all but the last of our requirements.

To produce a sudden sound we opted for pencil leads, which produce a very well time-defined sound when they break. A grain-sized enclosure was design to host the lead, shown in figure 5.9. This “active grain” has a cylindrical outer shell with two holes in which to put the lead. The holes also serve the purpose to weaken the grain, with the expectation it will locally bend under stress and break the lead with the perpendicular protrusion. The initial weight should get the active grain close to its breaking point and the continuous load triggers the breakage.



(a) 3D model of an active grain. (b) Breakage of a pencil lead in a press (c) Active grain embedded in an experiment

Figure 5.9: Active grains were designed (a) to be loaded with a pencil lead that would emit sound if broken (b) when in the experiment (c).

This process has been tested in a press (fig. 5.9b) and can even be performed using finger strength. Eight such grains were made and inserted into the experiment (fig. 5.9c). Since we do not control where force chains will appear, several active grains were used in the hope that at least one would be along a sufficiently strong force chain to break the lead. Unfortunately, after several repeats of the experiment

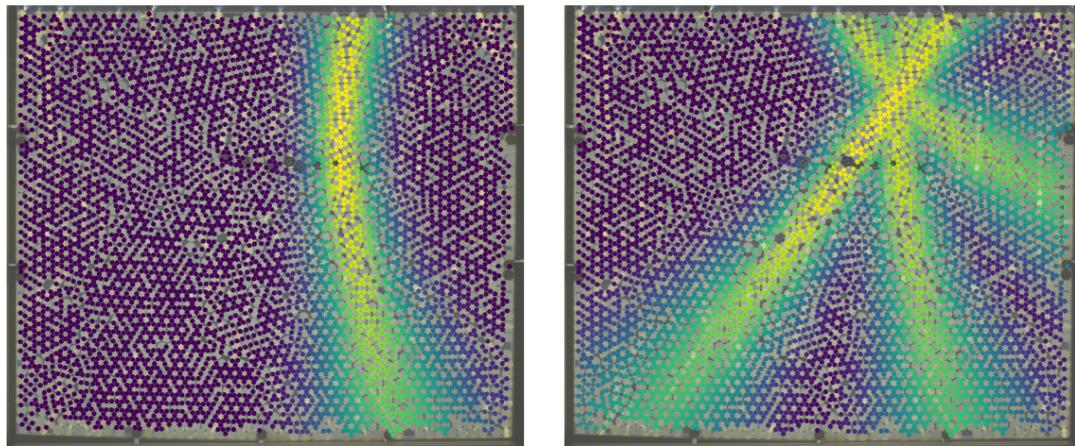
no such breakage happened. In spite of this, this design will be kept and refined upon by our latest Ph.D. student Louis Combe, whose thesis goals include acoustic location.

Spontaneous Emissions

In addition to active grain, we expected to capture spontaneous emissions. Such emission presents the advantage to be of the same nature as the one we eventually wish to analyze in the shearing experiment, however, by nature, we do not know their exact source location. Images were continuously recorded during the loading of the cell, in the hope to visually detect re-arrangements and validate acoustic location. Unfortunately, the displacements involved were too small to see anything on camera. It was mentioned in the introduction of this chapter we were actually able to locate an acoustic emission using a fast camera in previous works [126]. The cell was then much smaller, and thus the image resolution much better, in this previous setup. This older cell was actually briefly used when developing the techniques presented here, but it was abandoned in favor of a larger one. There were two reasons for this change. First, in a too-small cell, the flight time is so short the time-of-arrival measurements need to be much more precise. Second and more importantly, the walls of this cell were very acoustically conductive and I verified the ballistic front-wave traveled through the frame, instead of the grains.

Figure 5.10 presents the same two sub-figures as figure 5.7: a single branch and three stacked branches, from the same sensor pairs. These measurements correspond to a spontaneous emission that happened during a loading phase. The three branches on figure 5.10b cross perfectly on a single point. While the acoustic source could not be found using the images, the agreement between the three branches strongly suggest the presented location is accurate.

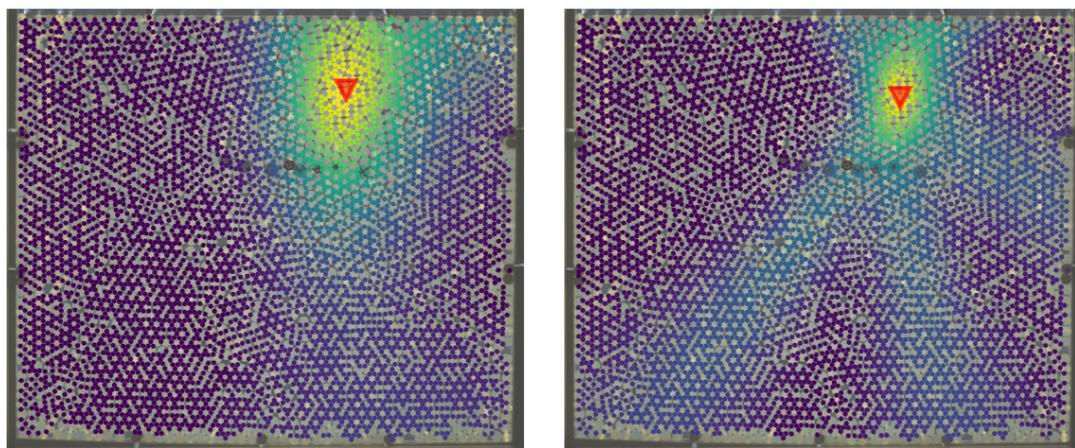
Both pooling outputs are presented in figure 5.11. Since the branches cross very well around a single point, there is little difference between product (fig. 5.11a) and average (fig. 5.11b) pooling, and both maximums are one grain apart. The same qualitative features observed on the controlled impact can be found: the average distribution is narrower and presents the traces of the branch used.



(a) Single branch likelihood
from sensor pair (2,5)

(b) Stacked branches from
pairs (1,4), (2,5) and (3,6).

Figure 5.10: Single (a) and triple (b) branch likelihood measured for a spontaneous emission.



(a) Product pooling, based on
formula 5.5.

(b) Sum pooling, based on
formula 5.6.

Figure 5.11: Comparison of product (a) and average (b) pooling for a spontaneous emission.

5.4.3 Limitations and Future development

Besides the difficulty in finding the timing of arrival discussed earlier, another limiting factor lies in the geometry of the sensors. To maximize precision, the sensors should surround the area where acoustic emissions could occur, to compute hyperbolic branches that cross at angles close to 90° . This is illustrated by figure 5.12 where the overlap between the branches crossing at a more acute angle is larger than the branches crossing at 90° . In addition, crossing points of curves at a right angle are less susceptible to a slight change in angle. However, in the shearing experiment, all sensors are placed aligned at the top.

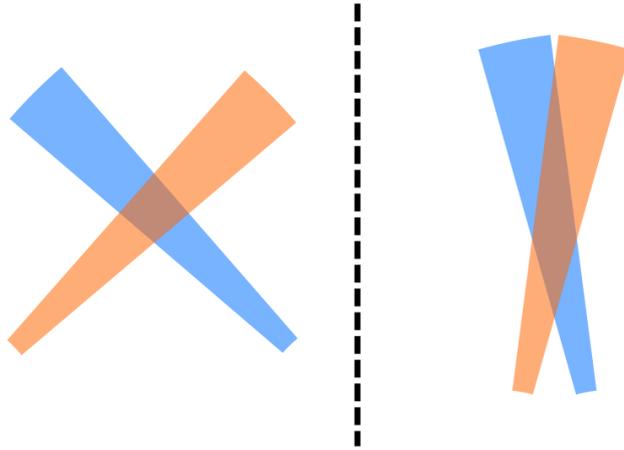


Figure 5.12: When branches cross at a more acute angle, the overlapping region cover a larger area and the location is thus less accurate.

To solve these issues, a completely different model is being considered, based on *time reversal* [133]. Such methods exploits the reversibility of wave propagation equations to make a signal “go back in time”. Using an array of sensors c_i , many signals $s_i(t)$ are recorded. The signals are then “flipped” as $s_i(-t)$ and sent back in the propagation medium. the waves then travel back and eventually refocus to their initial state. Time reversal has been shown to work experimentally, for instance in water [134] or in granular materials for ultrasonic waves [135].

Compared to the method presented in this chapter, time-reversal approaches present two advantages. First, there is no timing measurement as all signals are retro-propagated in sync. Second, more information is used, as the whole signal is sent back to the medium, whereas the hyperbolic method only uses a single point (the ballistic wave-front).

In our case, however, our sensors cannot act as acoustic sources to senb back signals into the granular. The only way to practically use time-reversal would be numerical, using a simulation of our experiment. Creating such a model is one of the thesis

goals of Louis Combe.

While computationally very cheap, fairly robust to fluctuations in the data, and accurate in our testbed experiment, this method also suffers from heavy reliance on human measurements. The work presented in this chapter can be seen as preliminary work on an ongoing project.

Chapter 6

Event Prediction Attempts

Contents

- 6.1 Signal-average based precursors 94**
 - 6.1.1 Dilation precursors 94
 - 6.1.2 Force precursors 98
- 6.2 Combining precursors with machine learning 101**
 - 6.2.1 Decision tree and random forest 102
 - 6.2.2 Defining and designing goals 104
 - 6.2.3 Performance 107
- 6.3 Perspectives 112**

6.1 Signal-average based precursors

A strength of our shearing experiment is the very large volume of data it outputs, thanks to its unlimited shearing range. The numerous large events available lead our initial searches for precursors to focus on the average behavior of the system around such large events. The intuition leading our initial attempts of prediction relates to the dilation of a granular layer under shear described in chapter 4. Rearrangement with high energy releases should be more common when the medium is relatively dilated and under high stress. The results presented in chapter 4 corroborate this intuition, where I showed large energy releases are correlated with large contraction of the medium (×-shaped figure 4.5). With that idea in mind, I looked for dilation-based precursors (using the position sensors and the images) and force-based precursors.

6.1.1 Dilation precursors

We have two ways to measure dilation in the shearing experiment: using the position sensor or using the images. The first is a very accurate but global measurement, while the second is less accurate but more localized.

Total volume increase

To verify the previously explained intuition, I first average windows of the position signal centered around events that are between given thresholds of energy:

$$\langle h_{E_L, E_H}(t) \rangle = \frac{1}{N} \sum_{i=1}^N h(t^* - t_i) \quad (6.1)$$

with h the position signal, t^* is the actual time, t_i the timings of events of energy between a low energy threshold E_L and a high energy threshold E_H . t is such that all events occur at $t^* = 0$. On this averaged signal, I perform a linear regression in the last 300 s before $t^* = 0$. The slope corresponds to the average dilation rate, which increases prior to the event. This can be seen on figure 6.1, where the dilation rate between $t^* = -300$ s and $t^* = 0$ s is roughly twice larger than between $t^* = -1000$ s and $t^* = -500$ s, while the dilation rate between $t^* = 0$ s and $t^* = 500$ s is close to 0, with a flat average signal $\langle h_{E_L, E_H} \rangle$.

In figure 6.1, E_H and E_L are such that the events selected are ranking from 20th to 8th in terms of energy.

To explore the robustness of the link between dilation rate and event size, I have measured the average dilation (measured on the same time interval of 300 s) prior to

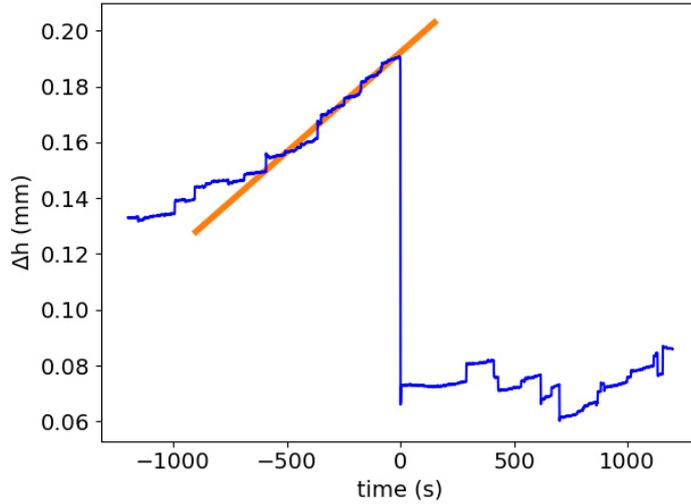


Figure 6.1: The dilation rate is computed by a linear regression on 300s on the average dilation curve.

events **binned by event energy**. As events are binned, the content of each bin is independent: the largest event appears in the highest bin and not in any bin below. Figure 6.2 shows that the average dilation rate seems monotonically dependent on the binned energy. This relation lends weight to the precursory nature of dilation, as event energy and the average expansion appear directly related.

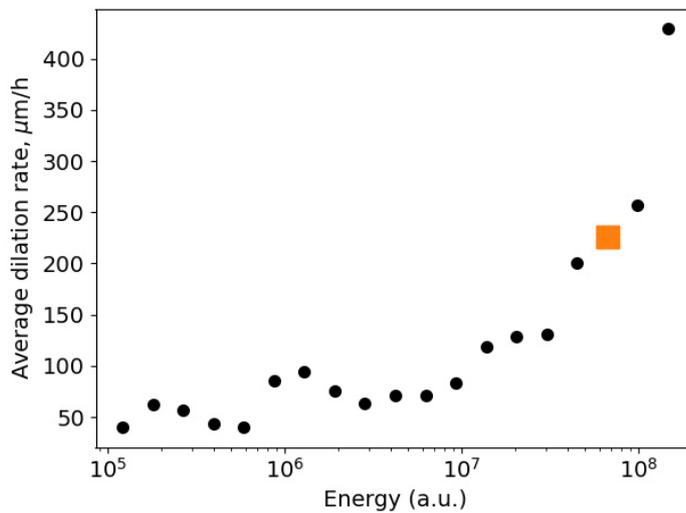


Figure 6.2: Dilation rate plotted against binned event energy. As events get smaller, the slope continuously decrease to a plateau. The highlighted orange square corresponds to the measurement exemplified in figure 6.1.

Local compactness

Let us consider a set of points C_i , $0 < i < n$ in the plane. The Voronoi tessellation or Voronoi diagram [136] is a partition of the space in “cells” attributed to each point C_i . The cell attributed to each C_i is defined as the ensemble of points p that are closer to C_i than any other $C_{j \neq i}$. The Voronoi cell of a point can be seen as its range of influence. In our case, we will compute a Voronoi diagram for the set of centers of our grains, as shown in figure 6.3, to probe local compactness.

The local structure around single grains can be described by a shape factor $\zeta = P^2/4\pi A$, defined as the ratio between the square of the perimeter P and the area A of the Voronoi cell of each grain [17]. The ζ parameter quantifies divergence from circularity, and is related to how compact the granular is. The rightmost colorbars in figure 6.3 shows that less dense arrangements indeed corresponds to higher ζ values.

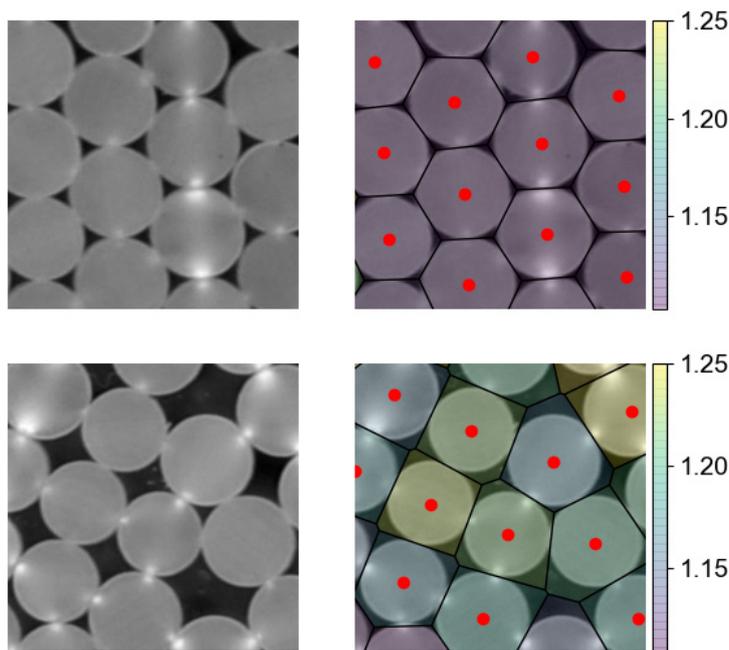


Figure 6.3: Two regions of the granular medium with their associated Voronoi diagram on the right, colored by the ζ parameter. The region on the first line is much denser.

ζ is normalized to a circle, meaning it is designed so that $\zeta = 1$ when computed on a circle. A regular hexagon, which achieves the minimum possible value in a monodisperse pack of disks, has $\zeta = 1.103$. A square has $\zeta = 1.273$, and a disk neighboring a hole in a hexagonal lattice has $\zeta = 1.286$. This scope is exemplified in figure 6.4, and bound the range of values we expect to encounter.

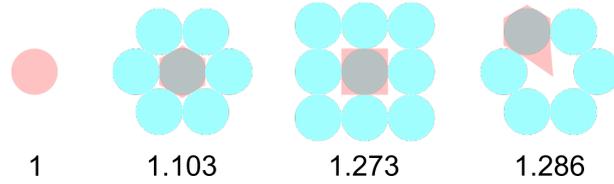


Figure 6.4: Value of ζ for various Voronoi cell shapes (*red*). The last three are shown over a disk tiling (*blue*) that could produce such cells.

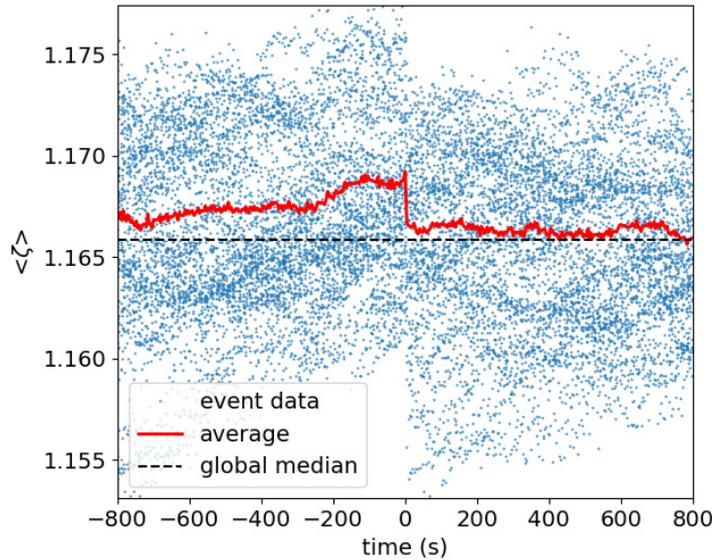


Figure 6.5: Each dot corresponds to the ζ parameter (inversely correlated with local density) averaged within the shear band on one picture at a given time t from the 33 largest events. Solid curve: averaged over the 33 images found at a given time t .

This $\zeta(x, y, t - t_i, i)$ is computed for every disks at position (x, y) and for every of the 33 largest event, indexed by i . Like for the dilation rate, times are shifted to align event timing. Only the disks within the shear band (bottom third of the granular pile, spanning 8~9 grain diameters) are kept and analyzed. We focus only on the shear band as it is the most active zone in our granular. The region above it barely moves and the arrangement of grains in it is conserved throughout the experiment. Including ζ values in this static area would only dilute the relevant signal.

ζ is first averaged over every grain in every image, corresponding to the dots on figure 6.5. The $\zeta(t - t_N, N)$ values of all 33 images found in the same time windows of 4 s width are averaged, resulting in $\langle \zeta \rangle(t)$, corresponding to the solid curve on figure 6.5. Both averages are shown 800 s both before and after the events happening, with a horizontal black line to show the global median. In the dots, and much more clearly on the average, a dilation emerges from roughly 5 minutes prior to the events. For comparison, the average time between selected events is around

40 minutes (22h of images available).

The measures performed in this section were first done for the 50 largest events and published in conference proceedings [137]. The exact results presented here were later published in [138].

6.1.2 Force precursors

The approach I presented for position-based precursors was kept for force-based precursors: computing the derivative of averaged windows (of the same size, 300 s) of the signal, centered on large events. As displayed in figure 6.6, the qualitative behavior is the same: prior to large events, the torque Γ increases faster than usual, indicating an augmentation of the system stiffness. Like was the case for the dilation rate, this stiffness is positively correlated with the size of the upcoming event. This relationship has intuitive meaning, as we expect large energy release to happen when the system is very jammed.

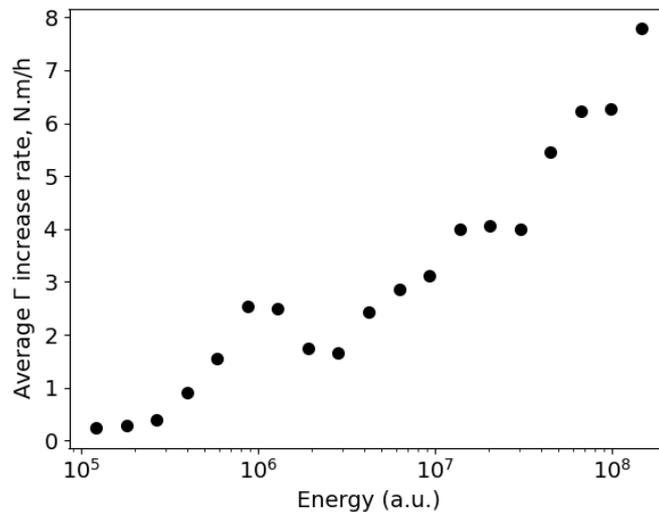


Figure 6.6: Average torque increase rate plotted against binned event energy. The positive correlation is even more pronounced in the case of the torque signal.

In addition to this measurement, which is dynamic in nature, as it examines the derivative of the torque, I investigated the absolute value of the torque. For a large energy release to happen, a lot of energy must be available, and thus we expect the torque values to be larger prior to large events.

Figure 6.7 compares the torque before and after events, with a threshold E_0 selected to that the 120 highest energy events highlighted.

The orange dots represents all events. Since the vast majority of labquakes are of very small energy, the system is almost identical between before and after these

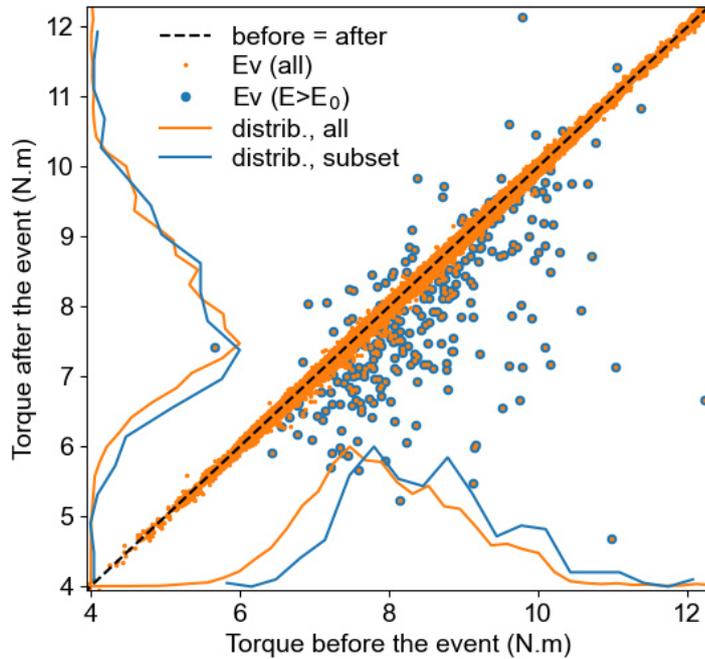


Figure 6.7: Torque values **after** against **before** an event. The distribution of “after”-values (*resp.* *before*) is shown for both selected and all events on the horizontal (*resp.* *vertical*) axis.

events: most of these points are clustered around the identity diagonal.

The [blue dots](#) corresponds to the high energy events. Since they are a subset of all events, they also have an orange dot overlaid on them. They strongly diverge from the identity diagonal, most of the points are below it: high energy events usually undergo a large drop in torque – this is coherent with what we described in chapter 4: *Probing the Structure*.

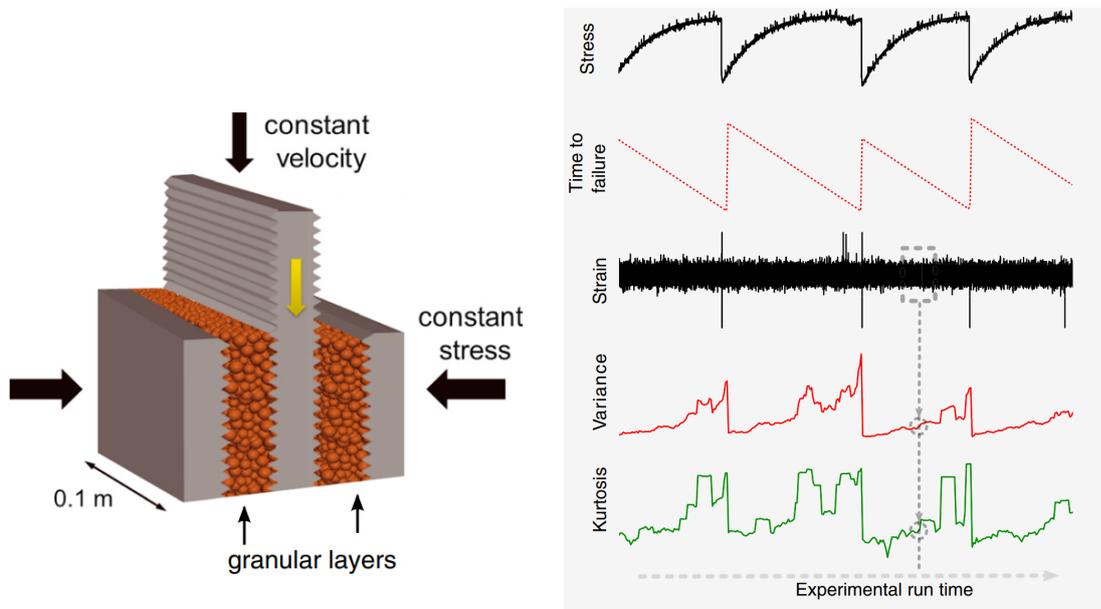
In addition, on each axis is shown the projected distribution of each set of events. On the vertical axis, one can see that the distribution of torque immediately after high energy events is very similar to the same distribution computed for all events. However, for the torque values prior to an event occurs (horizontal axis), the distribution computed on the high energy events is significantly shifted to the right compared to global distribution. This shows that, as expected, a lot of available energy favors large events.

The reason this was investigated for torque only and not dilation signal is that the measurement would not be exploitable, due to global plastic deformation of the medium. The range of values covered during a 24h shearing experiment (1 mm) compares to the height change due to plastic deformation observed in a control experiment without shear

Average correlations between several physical quantities (volume, compacity, torque change) were exhibited in this section. While this indicates predicting large events may be possible, doing so is actually very hard. We showed that large events are, **on average**, preceded by high dilation rates. This however does not guarantee a high dilation must precede a large event. Indeed, trying to predict big energy releases based on a simple increase in dilation rate of stiffness has proven to be fruitless. There is however still physical meaning to these measurements, and they are still expected to be useful for prediction. But to extract relevant information, more complex methods must be used.

6.2 Combining precursors with machine learning

Recently, machine learning techniques have been used to consistently forecast failure in a stick-slip experiment, using instantaneous acoustic emissions, by *Rouet-Leduc et al.* [139]. Their setup, described in [68], is represented in figure 6.8a.



(a) Schematic of their experimental setup.

(b) Measurement from their experiment. *Strain* here corresponds to acoustic emissions. The two last signals are computed on short acoustic windows.

Figure 6.8: Experimental setup and measurements done by *Rouet et al.*. (a) is from [68], (b) is from [139].

Using a machine learning model called *Random forests* (which I will describe in the next section), this team has been able to continuously predict the time-to-failure in their system. The performance of their model is shown on figure 6.9

However, their system is, as can be seen on the ground truth of the time-to-failure displayed on figure 6.9, quasi-periodic and with a rather narrow distribution of event sizes. While their model is oblivious to these properties, as it only uses instantaneous measurements, this suggests the dynamic at play in their system is less complex than in a completely non-periodic one with all-sizes events. They have also obtained similar results from real geological data [140], but the system studied in this paper is, here again, quasi-periodic.

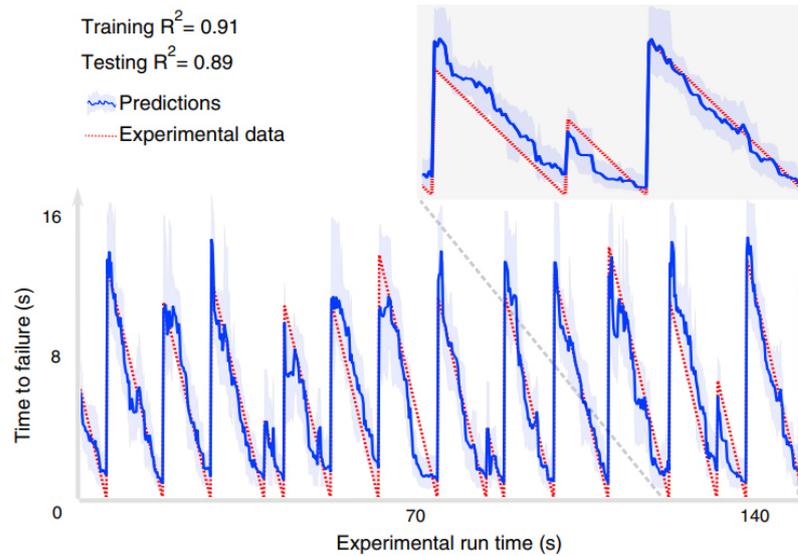


Figure 6.9: Performance of the predictive algorithm of *Rouet-Leduc et al.* The output of the model closely match the true time-to-failure.

Still, Inspired by their work, and in collaboration with Olivier Cochet-Escartin¹, a member of the Biophysics team at the ILM, I set to test a machine-learning tool to predict labquakes, *decision trees* and *random forests*.

6.2.1 Decision tree and random forest

A *decision tree* [141, 142] is a flowchart-like structure, taking some data as input and with arbitrary output (categorical, such as “yes” or “no”). At each node of the tree, a test is performed on the data. The result of the test will determine which branch is then followed until an output is reached.

An example of such a tree is shown in figure 6.10. The data here is what you can observe about a guest entering your office, and the desired output is to know whether it rains outside.

The tree presented in figure 6.10 would be the end results of a machine learning process, as it has already selected relevant questions. Initially, we do not know which questions to ask the input data (or “*features*”): the essence of machine learning is to find which tests and criteria are the most useful. A possible method is to pick the criterion that best splits the input data [143]. For instance, the set of integers is perfectly split in two by the question “is it even or odd?”.

Decision trees are simple and useful models and present the advantage to be easily interpreted. However, they are not very powerful. A way to refine decision tree

¹Olivier has machine learning expertise and is now a complete collaborator of the Labquake project, as he is a co-advisor of the latest Ph.D. student Louis Combe.

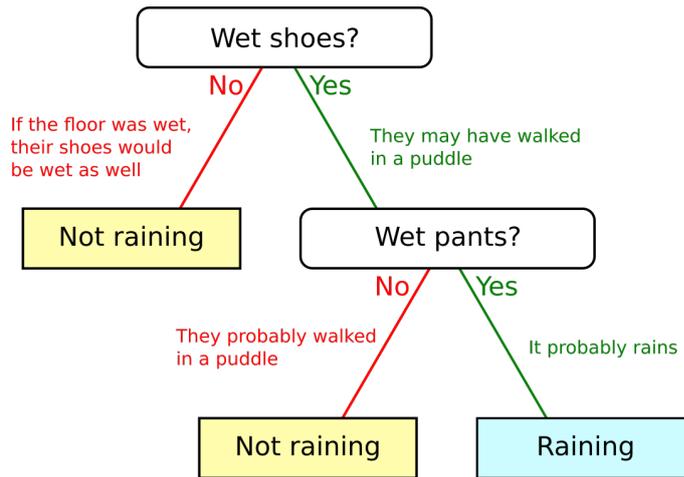


Figure 6.10: Example of a decision tree answering the question “is it raining outside?” based on observation of guest entering your office.

models is to pool many of them. For instance, we could make another tree in addition to the one shown in figure 6.10, which could ask different questions such as “do they hold an umbrella?”. By doing so, one can build a vast collection of models, each performing a partially different set of tests on the features, and reaching its own output. All the models can then vote and elect the global output. When agglomerating decision *trees*, the resulting model is called a *random forest* [144].

Random forests were used by the previously mentioned studies predicting laboratory stick-slip and Cascadia fault slow-slip, and we chose to use them as a well. Practically, I used the Python sklearn implementation, [145] and details concerning the model parameters and features are presented in appendix C. Below is a succinct list of the used features:

- **Force:** linear and quadratic fit parameters of the past 10 seconds. Moments 2 through 8 of the signal over the same timeframe.
(2 + 3 + 7 = 12 features)
- **Position:** linear and quadratic fit parameters of the past 10 seconds. Moments 2 through 8 of the signal over the same timeframe
(2 + 3 + 7 = 12 features)
- **Structure:** mean, minimum, maximum, standard deviation, 10th, 20th, ..., 90th percentiles of the distribution of ζ values, over the past 15 images (60 seconds). Fractal dimension of the force chain network in the shear band.
(4 + 9 + 1 = 14 features)

- **Acoustic:** 64 spectral energies, averaged on a spectrogram computed over the past 1 second, between 1 kHz and 20 kHz.
(64 features)

This sums up to a total of 102 features. I make such an array of 102 features, generally called $X(t)$, for time t spaced every 4 seconds (the image's time resolution). The different models are trained using the first 2/3rd of the experimental data and the last third is used for testing.

6.2.2 Defining and designing goals

Random forests are part of the so-called “supervised” machine learning models. This means they do not only require data to be trained on (the input) but an objective or a label to aim for (the output). For instance, linear regressions can be seen as primitive models of machine learning requiring inputs (x values) and outputs ($f(x)$ values) to learn the slope and intercept.

Let us present a few outputs (targets) I tried to predict using random forests, before showing their respective performance.

Time-To-Next event

This metric, very much inspired by the one used by *Rouet-Leduc et al.* [139], is simply defined as the time remaining until the next event occurs:

$$TTN(t, E_0) = t_e - t \tag{6.2}$$

With E_0 an energy threshold, and t_e the timing of the soonest event above this threshold. The curve of TTN is saw-tooth shaped, and discontinuous every time an event occurs. It has the advantage of being very easy to interpret as it is a direct prediction. However, it suffers from being non-continuous relative to threshold changes, as illustrated by figure 6.11.

To understand the issue, consider two events of size $E_0 + \epsilon$ and $E_0 - \epsilon$, with E_0 the threshold. For the first event, the model is trained to predict a value close to 0 just before the event, and a high value right after it passed. In the second case, the model is trained to predict almost the same value between before and after the event. However, as both events are very close in energy, we can imagine their features are similar. This is an issue for a machine learning algorithm, as it would be expected to produce very different outputs for very similar inputs. Such threshold effects are not a concern in [139], as their system produces events of similar size. In

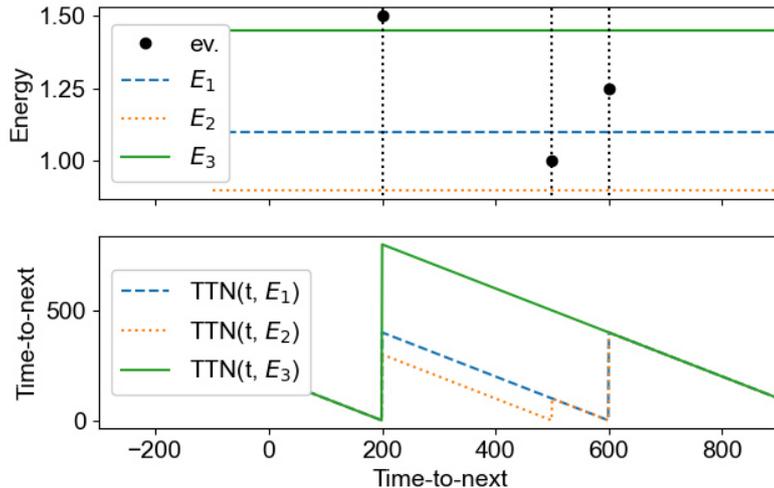


Figure 6.11: The resulting TTN curve can be very different depending on the threshold, even for very slight change in the parameter.

our case, the power-law distribution of energy is inherently ill-suited to a threshold dependant approach, since there is no characteristic scale.

In the following tests, the threshold E_0 was set to select 33 events, in a dataset spanning 22 hours and 15 minutes.

Upcoming Energy

This second function I designed corresponds to the quantity of energy that will be released in a future-looking time window. The window size is defined by an exponential in the integrand. Another option would be to set a finite value for the upper bound of the integral, but doing so would introduce unwanted discontinuities in the function while using an exponential window has a smoothing effect. This function presents the advantage of not depending on an arbitrary energy threshold but does depend on a time parameter. Contrary to the previous function, higher values correspond to large events.

$$UE(t, \tau) = \int_t^{\infty} E(x) e^{-x/\tau} dx \quad (6.3)$$

In the following tests, the timescale τ was set to 60 seconds, in the same dataset spanning 22 hours and 15 minutes.

Energy Release Time

Finally, the last function I used is defined as the time Δt it would take the system to release a quantity of energy S . This is a variation of the previously presented

metric. In this case, we wish to predict the lower values, as they mean a big event is coming: we need to wait a very short time to see a lot of energy released.

$$ERT(t, S) = \text{smallest } \Delta t \text{ such that } \frac{1}{\Delta t} \int_t^{t+\Delta t} E(x) dx > S \quad (6.4)$$

In the following tests, the value of the integrated energy threshold S is identical to the energy threshold E_0 for the TTN function, for comparison purposes. The data set is, again, the same.

Qualitative Classification

Finally, I tried a qualitative approach. Instead of attributing a value, each input is put in a broad class. They are defined as followed:

- **Peaceful:** No event above a given low energy threshold E_{low} is coming in the next 120s.
- **High energy:** *not in any class above* and at least one event over a high energy threshold E_{high} coming in the next 60s.
- **Medium energy:** *not in any class above* and at least one event over a medium energy threshold E_{med} coming in the next 60s.

The threshold E_{high} was set to E_0 . This value populates the “*high energy*” class with 495 points. In turn, the other thresholds are defined as $E_{mid} = E_{high}/10$ and $E_{low} = E_{mid}/10$. The respective populations of the “*medium energy*” and “*peaceful*” classes are 1833 and 2789 points, respectively. Let us remind that these points are **not** events: one event will typically corresponds to several data points. The 495 “*high energy*” samples come from 33 labquakes.

The first three functions are fitted using random forest *regressors* and the last one using a random forest *classifier*. The types of outputs are different, but the inner working of the algorithm is very similar and the inputs are exactly the same.

6.2.3 Performance

Time-To-Next event

The performance of a random forest model on this metric is shown in figure 6.12. To check the convergence of the training process, one can plot the prediction of the model on the data it has been trained on – this is the top plot of figure 6.12.

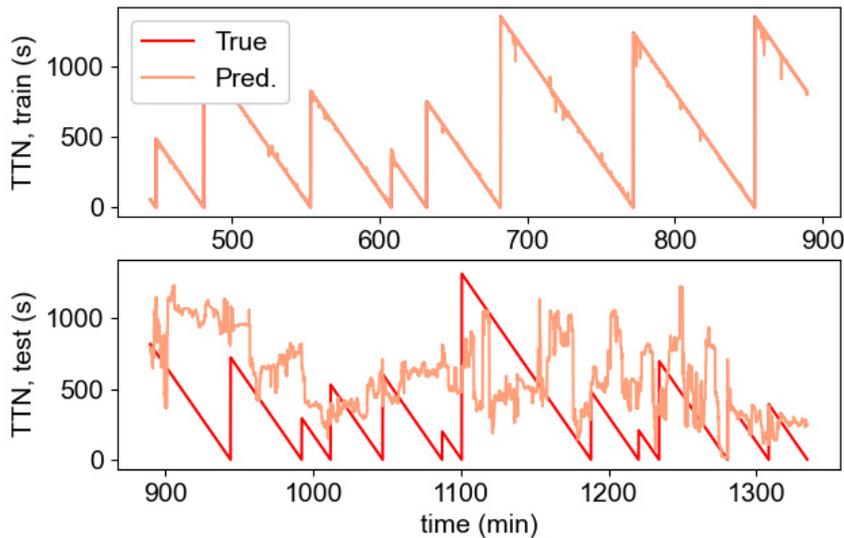


Figure 6.12: *Time-to-next event* function and its prediction. Lower values mean a big event is coming. The top plot contains data from the training set, explaining the almost perfect match between ground truth and prediction. The bottom plot presents the performance on data never seen by the model.

We can see there is no sign of *under-fitting*, an issue where models converge to trivial solutions, such as always predicting the average output. This trait will be shared by the next two functions. However, considering the quasi-perfect overlap of the prediction and the ground truth, there is probably *over-fitting*. This is the opposite issue, where a model learns the data “by heart” and fails to generalize to new data. By looking at the predictive performance on data never seen before on the bottom plot, one can see the model indeed fails to reproduce local details (smooth downward slopes) and macro details (raise on average around the 1100 min mark). This function does not seem promising for predictive purposes, but we expected low performances since, as explained earlier, a metric depending on an energy threshold is intrinsically ill-suited to our scale-invariant distribution of energy. In addition, Using this function implies a belief that there could be information in the system about an arbitrary far future: predicting the TTN **right now** is equal to 452 seconds is equivalent to make a prediction about **452 seconds in the future**.

Upcoming Energy

The upcoming energy results are presented in figure 6.13, with the *upcoming energy* in log scale. This metric converged properly as well. In its case, however, the prediction on new data seems marginally better: several spikes are reproduced by the model (t=1255 min, 1260 min, 1275 min, 1280 min ,or 1305 min), and the average trend increase between t=1250 min and 1280 min can be found in the model output as well.

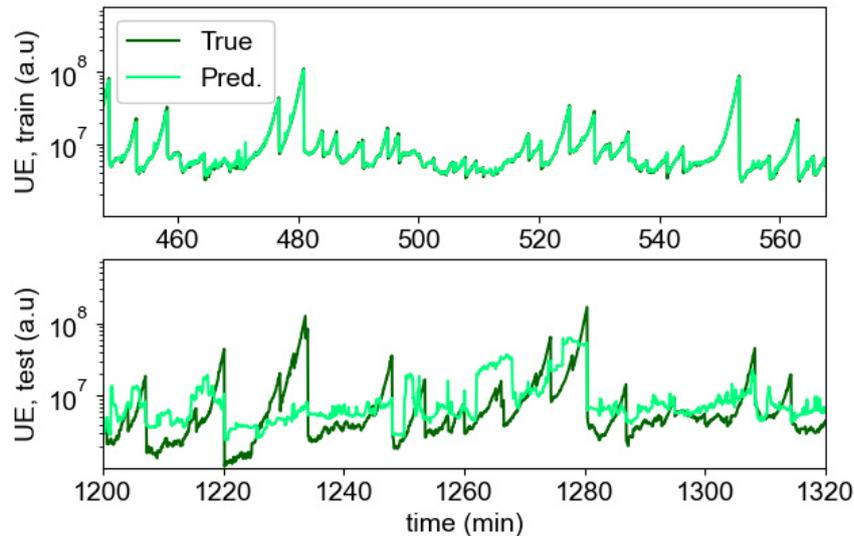


Figure 6.13: *Upcoming energy* function and its prediction. Higher values mean a big event is coming. The top plot contains data from the training set, while the bottom plot presents the performance on data never seen by the model.

However, spikes or drops sometimes occur in the prediction, when nothing actually happens in the ground truth signal. In addition, this model presents a problematic behavior, indicative of partial over-fitting: compared to the ground truth, its prediction is consistently below the peaks and above the valleys. This metric is more promising than the previous one, but still not satisfactory.

Energy Release Time

This last metric results are similar to the previous two. While the convergence happens without issues, as can be seen in figure 6.14, the prediction is still very off its target. Only one drop in the ground truth is somewhat predicted by the model, around $t=1280$ min. The remaining of the prediction is clustered around a value ~ 130 seconds.

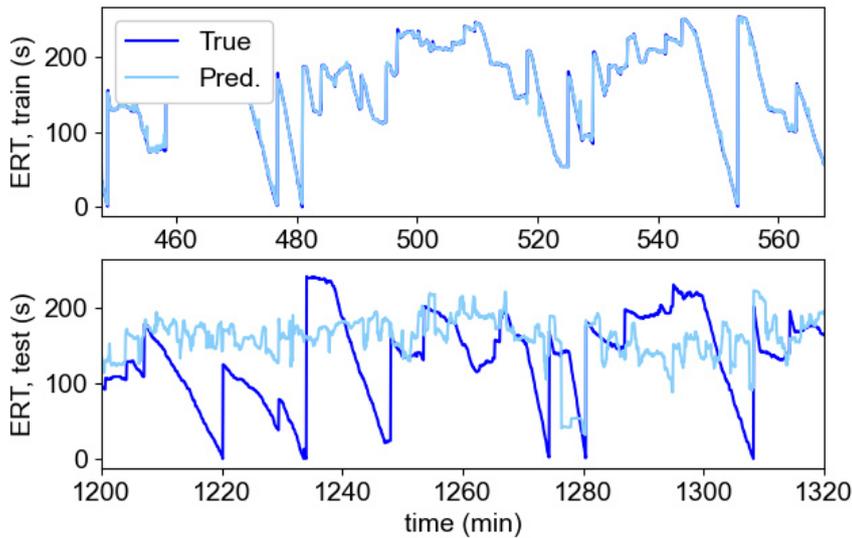


Figure 6.14: *Energy release time* energy function and its prediction. Lower values mean a big event is coming. The top plot contains data from the training set, while the bottom plot presents the performance on data never seen by the model.

Qualitative Classification

In front of the difficulty to predict continuous functions, I tried a qualitative, class-based approach. To visualize and evaluate the performance of such models, a commonly used method is to compute the *confusion matrix*. These are essentially histograms, defined as:

$$M_{i,j} = \text{number of class } i \text{ samples predicted to be of class } j.$$

Such matrix can be normalized column- or line-wise, depending on the information one wants to infer from it. A perfect prediction would produce a diagonal matrix. Points outside the diagonal are errors of the model, but not all errors have equal consequences. Consider a matrix describing the results of a COVID-19 test:

Infected	True positive	False negative
Healthy	False positive	True negative
	Positive test	Negative test

Ideally, every test is either a true negative or a true positive, but the consequences of a wrong result are more severe in the case of a false negative rather than a false positive.

The results of our classification are presented in figure 6.15. Let us analyze it in light of what I have explained.

The most important thing to predict is large, catastrophic events. Let us look at the **Ground truth / Large** line. Out of the 165 *large* data points in the testing set, 23 are well identified by the model. However, only 6 are “completely wrong” and identified as *peaceful*. When something will happen this model is capable of predicting it, but the size is underestimated $\sim 85\%$ of the time. Concerning the **Prediction / Large** column, we can see the model does not give many false positives: $\sim 68\%$ of the events predicted as *large* are actually so. The same interpretation can be made for the **Prediction / Peaceful** column: the model wrongly predicts it only $\sim 13\%$ of the time.

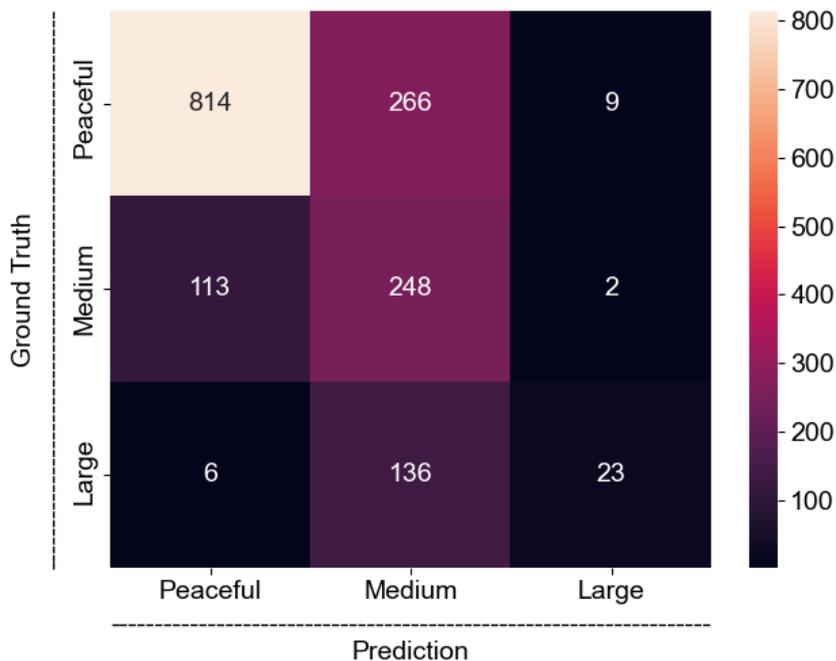


Figure 6.15: Confusion matrix of the classification model. Points on the diagonal are well predicted, points below have their energy under-estimated and points above are over-estimated.

This classification approach seems promising but is still not completely satisfying. Like continuous functions, where the prediction is drawn to the average value to minimize error, the classifier outputs are biased toward the *medium* class. It predicts almost twice their actual number: $266+248+136=650$ are predicted, when only

113+248+2=**363** actually happens. A lot of this increase is due to under-estimated *large* points, 142 of which the model fails to identify.

One could think of adding a “medium+” class, in the hope these 142 points would be classified as such. Doing so actually has an opposite effect, where *large* are mostly downgraded by the model into this new class.

Finally, this model like the other three suffers from the same imbalance issues.

6.3 Perspectives

I have shown a correlation between the average behavior of the torque and dilation signals and the energy of upcoming events. Such a relationship has also been found in a local measurement, ζ , describing local compacity. Unfortunately, these average precursors are not present in all the individual cases under average, and therefore, their predictive capabilities are rather limited. They have however encouraged trying more elaborate methods.

Machine learning has proven to be a promising tool in previous studies. Random forests are however apparently insufficient to produce reliable predictions in our setup. We wish to predict **extreme** and **rare** events. Both properties impede prediction. Rare events are by definition uncommon and require a lot of statistics. In addition, there is little incentive for a machine learning model (training to minimize its average error) to accommodate for these outliers. Furthermore, while the concept of “extreme” event may make intuitive sense, it is ill-defined in a scale-invariant phenomenon and thus difficult to work around for a computer.

Other fields, such as stock market research or weather forecasting, presents similar challenges to the one faced in earthquake prediction: scale invariance, almost everywhere discontinuous signals,... The strong financial and social incentives have driven machine learning research to develop more complex models adapted to these challenges[146–148], with some success. These methods are now being explored in our group, and we expect significant improvements.

Conclusion

I detailed my work on the labquake project, focused on deepening our understanding of this experimental setup and furthering our prediction toolbelt. Let us retrace what has been presented in this thesis.

- In chapter one, I described our experimental setup in detail. The main experiment consists of a continuously and slowly sheared 2D granular layer, with free volume and constant pressure conditions. The granular is made of a photoelastic material, allowing us to visualize the stress within it. During the shear, re-arrangement events occur. These “lab quakes” manifest themselves in various ways: acoustic energy release, resisting torque variations, changes in the volume experiment, and force chain re-organization. A range of techniques are used to monitor these events: six acoustic sensors are placed on the top of the experiment, two force sensors measure the resistive torque, two-position sensors measure the height of the granular layer and 24 cameras are used to capture the position of grains, as well as the force chains thanks to photoelasticity. Building this complex multi-camera setup, both the hardware and software to process the images, was my first task when I joined this project during my master internship.
- Chapter two focused on the results published in [83], regarding the statistical relevance of our setup to real earthquakes. Several statistical features found in real earthquakes are reproduced by our system: the energy is distributed analogous to Gutenberg-Richter law, the waiting time between events collapse near the same master curve and Omori-like a decaying aftershock rate is present as well.
- The structure of the granular is discussed in chapter three. The signal measured by the positions sensors revealed a rich dynamic of jumps and drops in the volume of the experiment. Comparing these position events with the force events, four classes of events can be defined, depending on whether the medium contracts or not and whether the torque increase or not. All these

behavior have been explained by examination of the strength and geometry of the force chain network.

In addition, two methods were described to characterize the density of the force network. The first exploits the images, performing a box-counting method to evaluate a fractal dimension of the force chains. I tried to develop a different method using two-sensors correlation, which eventually proved unfruitful.

- Chapter four describes a technique using acoustic wave and image data to locate the source of an acoustic emission. A propagation network is defined using the images, with local velocity evaluated thanks to photoelasticity. The time difference of arrival of a wave measured between two sensors is compared to what is predicted by the model. By computing three time-differences of arrival across six sensors, the source of emission can be located accurately. This will be used in future works to measure local properties around the nucleation site of labquakes, in the hope of finding precursors to large events.
- Finally, chapter five presents our latest advancement in terms of prediction. The first half presents precursors of the *average dynamic*. By averaging the system behavior over many events, I found correlations between upcoming event size and augmentation of the dilation rate as well as the stiffness of the granular.

While these results cannot be directly used to predict the behavior of our system, they motivated another approach: combining a multitude of precursors using machine learning. I present promising initial results using a simple model, random forests.

The work presented in this thesis is built upon previous works and will serve future works. Several chapters are exploratory in nature and detail new techniques, which are now used in our group. Important progress has been made on prediction with the introduction of machine learning. We believe the preliminary work presented here (most of it published, [83, 137, 138, 149]) will lead to key results in the future.

Appendix A

Power-law exponent influence

Let us compute about the average value of a power law:

$$\langle E \rangle = \int_{E_{min}}^{E_{max}} x P(x) dx \quad (\text{A.1})$$

The ratio E_{max}/E_{min} is related to the system size. In practice, it is dictated by the experimental setup characteristics. To compute this integral, we define $E_{min} = 1$ and use the expression $P(E) = A E^{-b}$. Contrary to our physicist intuition, we will not discard the normalization factor A as soon as we introduced it. The integral then becomes:

$$\langle E \rangle = \frac{\int_1^{E_{max}} x^{-b+1} dx}{\int_1^{E_{max}} x^{-b} dx} \quad (\text{A.2})$$

Depending on the b-value, the solutions are:

$$\langle E \rangle = \frac{E_{max}^{-b+2} - 1}{(-b + 2) \log(E_{max})} \quad \text{for } b = 1 \quad (\text{A.3})$$

$$\langle E \rangle = \frac{(-1 + 2) \log(E_{max})}{E_{max}^{-b+1} - 1} \quad \text{for } b = 2 \quad (\text{A.4})$$

$$\langle E \rangle = \frac{E_{max}^{-b+2} - 1}{E_{max}^{-b+1} - 1} \cdot \frac{1 - b}{2 - b} \quad \text{for other } b > 0 \quad (\text{A.5})$$

Note that $\langle E \rangle$ is continuous with respect to b, so we could content ourselves with a continuous extension of A.5. However, the b-values of 1 and 2 are particular for the denominator and numerator (respectively) of A.2 – but the specificity of these values is not limited to a mathematical quirk. On the left of figure A.1, $\langle E \rangle$ is plotted against a range of b-value for 9 different values of E_{max} . A rescaled version of this plot $\langle E \rangle / E_{max}$ is shown on the right. On the rescaled curves (*right*), we

can see that for power-law slopes below 1 the average energy scales linearly with E_{max} : extreme events are common enough they dictate the average response of the system. The non-rescaled plot (*left*) shows for slopes above 2, $\langle E \rangle$ is essentially independent of system size: extreme events are now so rare they hardly impact the system response. Between these two particular b-values, the relationship between $\langle E \rangle$ and E_{max} is more complex and comparison between two quantity from two different system may be difficult in this range.

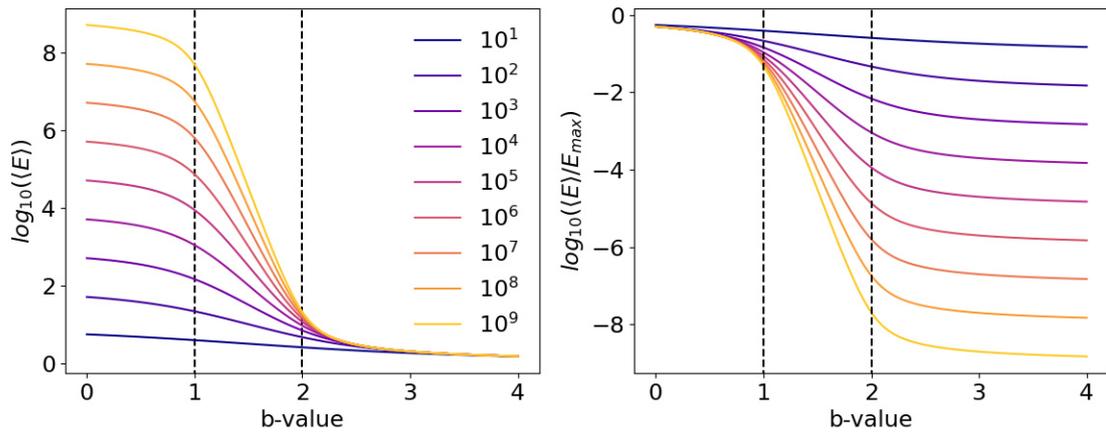


Figure A.1: *Left*: Average value $\langle E \rangle$ of a power-law distribution for different b-values (slopes) and different system size cutoff E_{max} (curves). *Right*: Same curves as on the left but rescaled by E_{max} .

Appendix B

Polariscope output intensity

we want to show relationship 2.8 used in chapter 2, in the frame of the Jones calculus [150].

Let us define \mathbf{F}_1 and \mathbf{S}_1 the vector defined by the fast and slow axis of the first quarter wave plate, shown in figure B.1. \mathbf{e}_1 and \mathbf{e}_2 are the principal direction of stress in the sample.

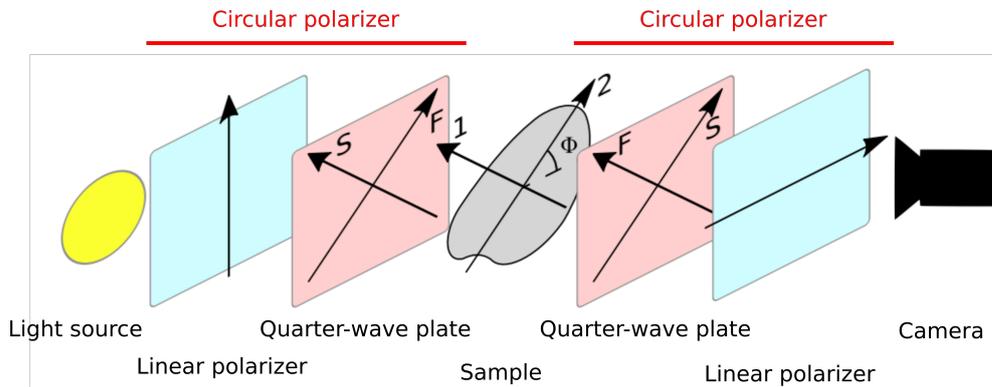


Figure B.1: Schematic of a polariscope. The sample is placed between two circular polarizer (made of a linear polarizer and a quarter-wave plate). The “S” and “F” denote the slow and fast axis of the quarter-wave plates, respectively.

Suppose a monochromatic wave goes into the system along the optical path. After the first linear polarizer, the wave polarisation is:

$$\mathbf{E} = \frac{ae^{i\omega t}}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

The quarter-wave plate adds phase shift of $\pi/2$ along its slow axis (second coordinate):

$$\mathbf{E} = \frac{ae^{i\omega t}}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}_{F_1, S_1}$$

The Jones matrix representing the sample effect is, expressed along the natural axis of birefringence:

$$M_{sample}^* = \begin{pmatrix} 1 & 0 \\ 0 & e^{i2\pi p} \end{pmatrix}_{\mathbf{e}_1, \mathbf{e}_2}$$

This matrix however is in a different basis. We must rotate it by the angle between \mathbf{e}_2 and \mathbf{F}_1 , ϕ :

$$M_{sample} = R(\phi) \cdot M_{sample}^* \cdot R(-\phi)$$

$$M_{sample} = \begin{pmatrix} \cos\phi & -\sin\phi \\ \sin\phi & \cos\phi \end{pmatrix}_{\mathbf{e}_1, \mathbf{e}_2} \cdot \begin{pmatrix} 1 & 0 \\ 0 & e^{i2\pi p} \end{pmatrix}_{\mathbf{e}_1, \mathbf{e}_2} \cdot \begin{pmatrix} \cos\phi & \sin\phi \\ -\sin\phi & \cos\phi \end{pmatrix}_{\mathbf{e}_1, \mathbf{e}_2}$$

Finally, the matrix representing the sample effect writes as:

$$M_{sample} = \begin{pmatrix} \cos^2\phi + e^{i2\pi p}\sin^2\phi & \sin\phi \cos\phi (1 - e^{i2\pi p}) \\ \sin\phi \cos\phi (1 - e^{i2\pi p}) & \sin^2\phi + e^{i2\pi p}\cos^2\phi \end{pmatrix}_{F_1, S_1}$$

After going through the sample, the wave polarization state is now:

$$\mathbf{E} \propto \begin{pmatrix} \cos^2\phi + e^{i2\pi p}\sin^2\phi + i \sin\phi \cos\phi(1 - e^{i2\pi p}) \\ \sin\phi \cos\phi(1 - e^{i2\pi p}) + i(\sin^2\phi + e^{i2\pi p}\cos^2\phi) \end{pmatrix}_{F_1, S_1}$$

Another phase shift $\pi/2$ is introduced by the second quarter-wave plate, this time along the first coordinates as the second slow axis is aligned with the first fast axis:

$$\mathbf{E} \propto \begin{pmatrix} i(\cos^2\phi + e^{i2\pi p}\sin^2\phi) - \sin\phi \cos\phi(1 - e^{i2\pi p}) \\ \sin\phi \cos\phi(1 - e^{i2\pi p}) + i(\sin^2\phi + e^{i2\pi p}\cos^2\phi) \end{pmatrix}_{F_1, S_1}$$

Finally, the last linear polarizer projects along its axis:

$$\mathbf{E} \propto i \cos^2\phi + i e^{i2\pi p}\sin^2\phi - \sin\phi \cos\phi(1 - e^{i2\pi p}) - \sin\phi \cos\phi(1 - e^{i2\pi p}) - i(\sin^2\phi + e^{i2\pi p}\cos^2\phi)$$

We can factorize everything with regards to $e^{i2\pi p}$:

$$\begin{aligned}\mathbf{E} &\propto e^{i2\pi p}(i \sin^2\phi + \sin\phi \cos\phi + \sin\phi \cos\phi - i\cos^2\phi) \\ &\quad + (i \cos^2\phi - \sin\phi \cos\phi - \sin\phi \cos\phi - i\sin^2\phi)\end{aligned}$$

This can be simplified to:

$$\begin{aligned}\mathbf{E} &\propto (e^{i2\pi p} - 1)\frac{1}{i}(-\sin^2\phi + 2i \sin\phi \cos\phi + \cos^2\phi) \\ \mathbf{E} &\propto (e^{i2\pi p} - 1)\frac{1}{i}(\cos\phi + i \sin\phi)^2 \\ \mathbf{E} &\propto (e^{i2\pi p} - 1)(-ie^{-2i\phi})\end{aligned}$$

The second term of the multiplication above is of constant norm and does not contribute to the interference. The first term can be expanded:

$$\begin{aligned}\mathbf{E} &\propto (e^{i2\pi p} - 1) \\ \mathbf{E} &\propto e^{i\pi p}(e^{i\pi p} - e^{-i\pi p}) \\ \mathbf{E} &\propto e^{i\pi p} \cdot 2i \sin(\pi p)\end{aligned}$$

Here again, the first exponential is of constant norm and can be ignored. Finally, we can write the intensity of the light:

$$\begin{aligned}I &\propto \mathbf{E}^2 \\ I &\propto \sin^2(\pi p)\end{aligned}$$

Appendix C

Random forests parameters

C.1 Feature definitions

Force signal features

These features are computed over the past 10 seconds of the top force signal. The 10 seconds signal cut, initially recorded at 100 kHz, are decimated to 1000 Khz prior to computation for speed reasons. Since the force signal varies slowly, this does not change the resulting feature.

The signal is fit by both a linear model ($\beta_1 t + \alpha_1$) and a quadratic one ($\gamma_2 t^2 + \beta_2 t + \alpha_2$), and each of the fitted coefficient is added to the feature list.

In addition to these “trend” features, the moments 2 through 8 of the same signal cut are computed, capturing whether the jittery properties of the signal.

This amount to $2 + 3 + 7 = 12$ features.

Position signal features

The same computation is done for the position signal than the force signal.

This adds another $2 + 3 + 7 = 12$ features to the total.

Structure

Let us first detail the features based on the distribution of the ζ parameter. From the roughly 3800 grains detected in the experiment, around 800 are used to make these features. The 3000 remaining grains have been excluded either because they are not in the shear band, or because they are on the edge of the detection area. Grains in the bulk hold little information, and grains on the edge can have degenerated Voronoi cell, with infinite area.

The mean, minimum, maximum, standard deviation, 10th, 20th, ..., 90th percentiles of the remaining ζ values are computed and averaged over the past 15 images (60 seconds). In addition, the fractal dimension of the force chain network in the shear band, averaged over the same period, is added.

This adds another $4 + 9 + 1 = 14$ features to the total.

Acoustic

A spectrogram of the past 1 second of the acoustic signal is computed, with the same parameters used for acoustic event detection, on 128 bins of frequency logarithmically distributed between 1 kHz and 20 kHz.

The spectral energy is then averaged over the total time (1 second) and across two adjacent bins. This produces 64 average energy.

This adds another $4 + 9 + 1 = 14$ features to the total, **bringing it to 102 features.**

C.2 Feature normalization

As is common to do in machine learning, the input features are normalized. Each feature is individually normalized to have a mean of 0 and a variance of 1. The normalization function is made measuring the mean and variance on the training set, and applied to the testing set.

Bibliography

- [1] K. W. Birkeland and C. C. Landry. “Power-laws and snow avalanches”. In: *Geophysical Research Letters* 29.11 (2002), pp. 49-1–49-3 (cit. on p. 2).
- [2] In Ho Lee. “Market crashes and informational avalanches”. In: *The Review of Economic Studies* 65.4 (1998), pp. 741–759 (cit. on p. 2).
- [3] Alessio Emanuele Biondo et al. “Reducing financial avalanches by random investments”. In: *Physical Review E* 88.6 (2013), p. 062814 (cit. on p. 2).
- [4] Xavier Gabaix et al. “A theory of power-law distributions in financial market fluctuations”. In: *Nature* 423 (2003), pp. 267–270 (cit. on p. 2).
- [5] Tobias Preis, Johannes J. Schneider, and H. Eugene Stanley. “Switching processes in financial markets”. In: *Proceedings of the National Academy of Sciences* 108.19 (2011), pp. 7674–7678 (cit. on p. 2).
- [6] Kim Sneppen et al. “Evolution as a self-organized critical phenomenon”. In: *Proceedings of the National Academy of Sciences* 92.11 (1995), pp. 5209–5213 (cit. on p. 2).
- [7] Brian R. Dennis. “Solar hard X-ray bursts”. In: *Solar Physics* 100.1-2 (1985), pp. 465–490 (cit. on p. 2).
- [8] D. Hamon, M. Nicodemi, and H. J. Jensen. “Continuously driven OFC: A simple model of solar flare statistics”. In: *Astronomy and Astrophysics* 387.1 (2002), pp. 326–334 (cit. on p. 2).
- [9] Filippo Radicchi et al. “Classes of critical avalanche dynamics in complex networks”. In: *Physical Review Research* 2.3 (2020), p. 033171 (cit. on p. 2).
- [10] G. A. Held et al. “Experimental study of critical-mass fluctuations in an evolving sandpile”. In: *Phys. Rev. Lett.* 65 (9 1990), pp. 1120–1123 (cit. on p. 2).
- [11] Vidar Frette et al. “Avalanche dynamics in a pile of rice”. In: *Nature* 379 (1996), pp. 49–52 (cit. on p. 2).

- [12] E. Altshuler et al. “Avalanches in One-Dimensional Piles with Different Types of Bases”. In: *Phys. Rev. Lett.* 86 (24 2001), pp. 5490–5493 (cit. on p. 2).
- [13] C. M. Aegerter, R. Günther, and R. J. Wijngaarden. “Avalanche dynamics, surface roughening, and self-organized criticality: Experiments on a three-dimensional pile of rice”. In: *Phys. Rev. E* 67 (5 2003), p. 051306 (cit. on p. 2).
- [14] N. Nerone et al. “Instabilities in slowly driven granular packing”. In: *Phys. Rev. E* 67 (1 2003), p. 011302 (cit. on p. 2).
- [15] C. M. Aegerter et al. “Extremal Dynamics and the Approach to the Critical State: Experiments on a Three Dimensional Pile of Rice”. In: *Phys. Rev. Lett.* 92 (5 2004), p. 058702 (cit. on p. 2).
- [16] Kinga A. Lörincz and Rinke J. Wijngaarden. “Edge effect on the power law distribution of granular avalanches”. In: *Phys. Rev. E* 76 (4 2007), p. 040301 (cit. on p. 2).
- [17] O. Ramos, E. Altshuler, and K. J. Måløy. “Avalanche Prediction in a Self-Organized Pile of Beads”. In: *Phys. Rev. Lett.* 102 (7 2009), p. 078701 (cit. on pp. 2, 67, 96).
- [18] Stéphane Santucci, Loïc Vanel, and Sergio Ciliberto. “Subcritical Statistics in Rupture of Fibrous Materials: Experiments and Model”. In: *Phys. Rev. Lett.* 93 (9 2004), p. 095505 (cit. on p. 2).
- [19] M. Stojanova et al. “High Frequency Monitoring Reveals Aftershocks in Subcritical Crack Growth”. In: *Phys. Rev. Lett.* 112 (2014), p. 115502 (cit. on pp. 2, 6, 8).
- [20] Charles F. Richter. “An instrumental earthquake magnitude scale”. In: *Bull. Seismol. Soc. Am.* 25 (1 1935), pp. 1–32 (cit. on p. 2).
- [21] Per Bak and Chao Tang. “Earthquakes as a self-organized critical phenomenon”. In: *Journal of Geophysical Research: Solid Earth* 94.B11 (1989), pp. 15635–15637 (cit. on pp. 2, 7).
- [22] Zeev Olami and Kim Christensen. “Temporal correlations, universality, and multifractality in a spring-block model of earthquakes”. In: *Phys. Rev. A* 46 (4 Aug. 1992), R1720–R1723 (cit. on pp. 2, 7).
- [23] E. A. Jagla. “Realistic spatial and temporal earthquake distributions in a modified Olami-Feder-Christensen model”. In: *Phys. Rev. E* 81 (4 2010), p. 046117 (cit. on p. 2).

-
- [24] Beno Gutenberg and Charles Francis Richter. “Seismicity of the earth and associated phenomena”. In: *Bulletin of the Seismological Society of America* 40.1 (Jan. 1949), pp. 76–76. eprint: <https://pubs.geoscienceworld.org/bssa/article-pdf/40/1/76/2691453/BSSA0400010076.pdf> (cit. on pp. 2, 10).
- [25] Gilles Chabrier. “Galactic stellar and substellar initial mass function”. In: *Publications of the Astronomical Society of the Pacific* 115.809 (2003), p. 763 (cit. on p. 4).
- [26] Gladys Hanson and Emmet Condon. *Denial of Disaster: the Untold Story and Unpublished Photographs of the San Francisco Earthquake and Fire of 1906*. 1989 (cit. on p. 4).
- [27] H. F. Reid. “The elastic-rebound theory of earthquakes”. In: *Univ. of California Publications, Bulletin of the Department of Geology* 6 (1911), p. 19 (cit. on p. 4).
- [28] H. Benioff. “Earthquake source mechanisms”. In: *Science* 143 (1964), p. 1399 (cit. on p. 5).
- [29] W. F. Brace and J. D. Byerlee. “Stick-Slip as a Mechanism for Earthquakes”. In: *Science* 153.3739 (1966), pp. 990–992 (cit. on pp. 5, 6).
- [30] R. Burridge and L. Knopoff. “Model and theoretical seismicity”. In: *Bull. Seismol. Soc. Am.* 57 (1967), pp. 341–371 (cit. on pp. 5, 8).
- [31] M. Otsuka. “Model and theoretical seismicity”. In: *Bulletin of the Seismological Society of America* 6 (1972), pp. 311–315 (cit. on p. 5).
- [32] James H. Dieterich. “Modeling of rock friction: 1. Experimental results and constitutive equations”. In: *J. Geophys. Res. Solid Earth* 84.B5 (1979), pp. 2161–2168. ISSN: 2156-2202. DOI: 10.1029/JB084iB05p02161 (cit. on pp. 5, 6).
- [33] Andy Ruina. “Slip instability and state variable friction laws”. In: *Journal of Geophysical Research: Solid Earth* 88.B12 (1983), pp. 10359–10370 (cit. on pp. 5, 6).
- [34] James R Rice. “Constitutive relations for fault slip and earthquake instabilities”. In: *Instabilities in continuous media*. Springer, 1983, pp. 443–475 (cit. on pp. 5, 6).
- [35] Hikaru Kawamura et al. “Statistical physics of fracture, friction, and earthquakes”. In: *Rev. Mod. Phys.* 84 (2012), pp. 839–884 (cit. on p. 6).

- [36] Stephen R Brown, Christopher H Scholz, and John B Rundle. “A simplified spring-block model of earthquakes”. In: *Geophysical Research Letters* 18.2 (1991), pp. 215–218 (cit. on p. 6).
- [37] Kiyoo Mogi. “Earthquakes and fractures”. In: *Tectonophysics* 5.1 (1967), pp. 35–55 (cit. on pp. 6, 8).
- [38] C. H. Scholz. “The frequency-magnitude relation of microfracturing in rock and its relation to earthquakes”. In: *Bulletin of the seismological society of America* 58.1 (1968), pp. 399–415 (cit. on pp. 6, 8).
- [39] Jordi Barò et al. “Statistical similarity between the compression of a porous material and earthquakes”. In: *Physical Review Letters* 110 (8 2013), p. 088702 (cit. on pp. 6, 8).
- [40] J. Barés et al. “Aftershock sequences and seismic-like organization of acoustic events produced by a single propagating crack”. In: *Nature Communications* 9.1 (2018), p. 1253 (cit. on pp. 6, 8).
- [41] O. Ben-David, G. Cohen, and J. Fineberg. “The dynamics of the onset of frictional slip”. In: *Science* 330 (2010), p. 211 (cit. on pp. 6, 8).
- [42] Venkateswararao Vemuri. *Modeling of complex systems: an introduction*. New York Academic Press, 1978 (cit. on p. 7).
- [43] P. W. Anderson. “More Is Different”. In: *Science* 177.4047 (1972), pp. 393–396 (cit. on p. 7).
- [44] Michael E. Fisher. “The renormalization group in the theory of critical behavior”. In: *Rev. Mod. Phys.* 46 (1974), pp. 597–616 (cit. on p. 7).
- [45] R. F. Smalley Jr., D. L. Turcotte, and Sara A. Solla. “A renormalization group approach to the stick-slip behavior of faults”. In: *Journal of Geophysical Research: Solid Earth* 90.B2 (1985), pp. 1894–1900 (cit. on p. 7).
- [46] J. Huang and D.L. Turcotte. “Fractal distributions of stress and strength and variations of b-value”. In: *Earth and Planetary Science Letters* 91.1 (1988), pp. 223–230 (cit. on p. 7).
- [47] Takayuki Hirata. “A correlation between the b value and the fractal dimension of earthquakes”. In: *Journal of Geophysical Research: Solid Earth* 94.B6 (1989), pp. 7507–7514 (cit. on p. 7).
- [48] Benoit B Mandelbrot. *The fractal geometry of nature*. San Francisco, CA: Freeman, 1982 (cit. on p. 7).

-
- [49] John B. Rundle. “A physical model for earthquakes: 1. Fluctuations and interactions”. In: *Journal of Geophysical Research: Solid Earth* 93.B6 (1988), pp. 6237–6254 (cit. on p. 7).
- [50] John B. Rundle. “A physical model for earthquakes: 3. Thermodynamical approach and its relation to nonclassical theories of nucleation”. In: *Journal of Geophysical Research: Solid Earth* 94.B3 (1989), pp. 2839–2855 (cit. on p. 7).
- [51] Jie Huang and D. L. Turcotte. “Evidence for chaotic fault interactions in the seismicity of the San Andreas fault and Nankai trough”. In: *Nature* 348.6298 (1990), pp. 234–236 (cit. on p. 7).
- [52] Per Bak, Chao Tang, and Kurt Wiesenfeld. “Self-organized criticality: An explanation of the $1/f$ noise”. In: *Phys. Rev. Lett.* 59 (4 1987), pp. 381–384 (cit. on pp. 7, 8).
- [53] A Sornette and D Sornette. “Self-Organized Criticality and Earthquakes”. In: *Europhysics Letters (EPL)* 9.3 (1989), pp. 197–202 (cit. on p. 7).
- [54] J. M. Carlson and J. S. Langer. “Properties of earthquakes generated by fault dynamics”. In: *Phys. Rev. Lett.* 62 (1989), pp. 2632–2635 (cit. on p. 7).
- [55] J. M. Carlson and J. S. Langer. “Mechanical model of an earthquake fault”. In: *Phys. Rev. A* 40 (1989), pp. 6470–6484 (cit. on p. 7).
- [56] Didier Sornette, Philippe Davy, and Anne Sornette. “Structuration of the lithosphere in plate tectonics as a self-organized critical phenomenon”. In: *Journal of Geophysical Research: Solid Earth* 95.B11 (1990), pp. 17353–17361 (cit. on p. 7).
- [57] Hiizu Nakanishi. “Cellular-automaton model of earthquakes with deterministic dynamics”. In: *Phys. Rev. A* 41 (1990), pp. 7086–7089 (cit. on p. 7).
- [58] Kan Chen, Per Bak, and S. P. Obukhov. “Self-organized criticality in a crack-propagation model of earthquakes”. In: *Phys. Rev. A* 43 (1991), pp. 625–630 (cit. on p. 7).
- [59] J. Lomnitz-Adler, L. Knopoff, and G. Martínez-Mekler. “Avalanches and epidemic models of fracturing in earthquakes”. In: *Phys. Rev. A* 45 (1992), pp. 2211–2221 (cit. on p. 7).
- [60] Kim Christensen and Zeev Olami. “Scaling, phase transitions, and nonuniversality in a self-organized critical cellular-automaton model”. In: *Phys. Rev. A* 46 (4 Aug. 1992), pp. 1829–1838 (cit. on p. 7).

- [61] Kim Christensen and Zeev Olami. “Variation of the Gutenberg-Richter b values and nontrivial temporal correlations in a Spring-Block Model for earthquakes”. In: *Journal of Geophysical Research: Solid Earth* 97.B6 (1992), pp. 8729–8735 (cit. on p. 7).
- [62] Huang-Jian Xu, B Bergersen, and K Chen. “Self-organized ruptures in an elastic medium: a possible model for earthquakes”. In: *Journal of Physics A: Mathematical and General* 25.22 (1992), pp. L1251–L1256 (cit. on p. 7).
- [63] S. Lübeck and K. D. Usadel. “Numerical determination of the avalanche exponents of the Bak-Tang-Wiesenfeld model”. In: *Phys. Rev. E* 55 (1997), pp. 4095–4099 (cit. on p. 7).
- [64] H. E. Stanley. *Introduction to Phase Transitions and Critical Phenomena*. New York: Oxford Univ. Press, 1987 (cit. on p. 8).
- [65] Zeev Olami, Hans Jacob S. Feder, and Kim Christensen. “Self-organized criticality in a continuous, nonconservative cellular automaton modeling earthquakes”. In: *Phys. Rev. Lett.* 68 (1992), pp. 1244–1247 (cit. on p. 8).
- [66] F. X. Passelègue et al. “From sub-rayleigh to supershear ruptures during stick-slip experiments on crustal rocks”. In: *Science* 340 (2013), p. 1208 (cit. on p. 8).
- [67] Paul A Johnson et al. “Effects of acoustic waves on stick-slip in granular media and implications for earthquakes”. In: *Nature* 451.7174 (2008), pp. 57–60 (cit. on pp. 8, 9).
- [68] Paul A Johnson et al. “Acoustic emission and microslip precursors to stick-slip failure in sheared granular material”. In: *Geophysical Research Letters* 40.21 (2013), pp. 5627–5631 (cit. on pp. 8, 9, 101).
- [69] Paul A Johnson et al. “Laboratory earthquake forecasting: A machine learning competition”. In: *Proceedings of the National Academy of Sciences* 118.5 (2021) (cit. on p. 9).
- [70] James H Dieterich. “Time-dependent friction in rocks”. In: *Journal of Geophysical Research* 77.20 (1972), pp. 3690–3697 (cit. on p. 8).
- [71] Staurt Crampin, Russ Evans, and Barry K Atkinson. “Earthquake prediction: a new physical basis”. In: *Geophysical Journal International* 76.1 (1984), pp. 147–156 (cit. on p. 8).
- [72] Thomas H Heaton. “Evidence for and implications of self-healing pulses of slip in earthquake rupture”. In: *Physics of the Earth and Planetary Interiors* 64.1 (1990), pp. 1–20 (cit. on p. 8).

-
- [73] C. Marone. “Laboratory-derived friction laws and their application to seismic faulting”. In: *Annual Review of Earth and Planetary Sciences* 26.1 (1998), pp. 643–696 (cit. on p. 8).
- [74] Junzo Kasahara et al. “Influence of water on earthquake generation along subduction zones”. In: *Bulletin of the Earthquake Research Institute, University of Tokyo* 76.3 (2002), pp. 291–303 (cit. on p. 8).
- [75] Yijue Diao and Rosa M Espinosa-Marzal. “The role of water in fault lubrication”. In: *Nature communications* 9.1 (2018), pp. 1–10 (cit. on p. 8).
- [76] J Weiss et al. “Cohesion versus friction in controlling the long-term strength of a self-healing experimental fault”. In: *Journal of Geophysical Research: Solid Earth* 121.12 (2016), pp. 8523–8547 (cit. on p. 8).
- [77] Cédric Lachaud et al. “Micro-seismic monitoring of a shear fault within a floating ice plate”. In: *Journal of Geophysical Research: Solid Earth* 124.10 (2019), pp. 10444–10467 (cit. on p. 8).
- [78] Evangelos Korkolis et al. “Irregular stick-slip and the role of cohesion in an ice friction experiment”. In: *EGU General Assembly Conference Abstracts*. 2021, EGU21–7436 (cit. on p. 8).
- [79] Aghil Abed Zadeh, Jonathan Barés, and Robert P. Behringer. “Crackling to periodic dynamics in granular media”. In: *Phys. Rev. E* 99 (2019), p. 040901 (cit. on p. 9).
- [80] Aghil Abed Zadeh et al. “Seismicity in sheared granular matter”. In: *Phys. Rev. E* 99 (2019), p. 052902 (cit. on p. 9).
- [81] Vincent Canel et al. “Acoustic Monitoring of Damages and Shear-Banding in Cemented Granular Material under External Loading”. In: *Forum Acusticum*. 2020, pp. 2299–2299 (cit. on p. 9).
- [82] David Houdoux et al. “Micro-slips in an experimental granular shear band replicate the spatiotemporal characteristics of natural earthquakes”. In: *Communications Earth & Environment* 2.1 (2021), pp. 1–11 (cit. on p. 9).
- [83] S. Lherminier et al. “Continuously Sheared Granular Matter Reproduces in Detail Seismicity Laws”. In: *Phys. Rev. Lett.* 122 (2019), p. 218501 (cit. on pp. 9, 15, 19, 39, 55, 80, 113, 114).
- [84] Zachary E Ross et al. “Searching for hidden earthquakes in Southern California”. In: *Science* 364.6442 (2019), pp. 767–771 (cit. on pp. 10, 40).

- [85] C Godano, E Lippiello, and L De Arcangelis. “Variability of the b value in the Gutenberg–Richter distribution”. In: *Geophysical Journal International* 199.3 (2014), pp. 1765–1771 (cit. on p. 10).
- [86] Thomas C. Hanks and Hiroo Kanamori. “A moment magnitude scale”. In: *Journal of Geophysical Research: Solid Earth* 84.B5 (1979), pp. 2348–2350 (cit. on p. 11).
- [87] Jeff Alstott, Ed Bullmore, and Dietmar Plenz. “powerlaw: A Python Package for Analysis of Heavy-Tailed Distributions”. In: *PLOS ONE* 9.1 (Jan. 2014), pp. 1–11. DOI: 10.1371/journal.pone.0085777 (cit. on pp. 11, 41).
- [88] F. Omori. “On the aftershocks of earthquakes”. In: *J. Coll. Sci. Imp. Univ. Tokyo* 7 (1894), pp. 111–200 (cit. on p. 11).
- [89] Tokuji Utsu. “A statistical study on the occurrence of aftershocks”. In: *Geophys. Mag.* 30 (1961), pp. 521–605 (cit. on p. 11).
- [90] Yan Kagan and L Knopoff. “Statistical search for non-random features of the seismicity of strong earthquakes”. In: *Physics of the earth and planetary interiors* 12.4 (1976), pp. 291–318 (cit. on p. 12).
- [91] GM Molchan and OE Dmitrieva. “Aftershock identification: methods and new approaches”. In: *Geophysical Journal International* 109.3 (1992), pp. 501–516 (cit. on p. 12).
- [92] Jiancang Zhuang, Yosihiko Ogata, and David Vere-Jones. “Analyzing earthquake clustering features by using stochastic reconstruction”. In: *Journal of Geophysical Research: Solid Earth* 109.B5 (2004) (cit. on p. 12).
- [93] David Marsan and Olivier Lengline. “Extending earthquakes’ reach through cascading”. In: *Science* 319.5866 (2008), pp. 1076–1079 (cit. on p. 12).
- [94] David Marsan and Olivier Lengliné. “A new estimation of the decay of aftershock density with distance to the mainshock”. In: *Journal of Geophysical Research: Solid Earth* 115.B9 (2010) (cit. on p. 12).
- [95] David Marsan and Agnès Helmstetter. “How variable is the number of triggered aftershocks?” In: *Journal of Geophysical Research: Solid Earth* 122.7 (2017), pp. 5544–5560 (cit. on p. 12).
- [96] Álvaro Corral. “Long-Term Clustering, Scaling, and Universality in the Temporal Occurrence of Earthquakes”. In: *Phys. Rev. Lett.* 92 (2004), p. 108501 (cit. on pp. 13, 45).

-
- [97] Karen E. Daniels and Nicholas W. Hayman. “Force chains in seismogenic faults visualized with photoelastic granular shear experiments”. In: *J. Geophys. Res. Solid Earth* 113.B11 (2008), pp. 2156–2202 (cit. on p. 18).
- [98] O. Kuwano, R. Ando, and T. Hatano. “Granular Friction in a Wide Range of Shear Rates”. In: *AIP Conf. Proc.* 1542 (2013), pp. 32–37 (cit. on p. 19).
- [99] David L. Henann and Ken Kamrin. “A predictive, size-dependent continuum model for dense granular flows”. In: *PNAS* 110.17 (2013), pp. 6730–6735 (cit. on p. 19).
- [100] Yue Pan. “Target-less registration of point clouds: A review”. In: (Dec. 2019) (cit. on p. 29).
- [101] CC Lo and CA Chang. *Neural networks for bar code positioning in automated material handling. Industrial Automation and Control: Emerging Technologies*. 1995 (cit. on p. 31).
- [102] Daniel Wagner and Dieter Schmalstieg. “ARToolKitPlus for Pose Tracking on Mobile Devices”. In: Jan. 2007 (cit. on p. 31).
- [103] Minhua Ma, Lakhmi C Jain, Paul Anderson, et al. *Virtual, augmented reality and serious games for healthcare 1*. Vol. 1. Springer, 2014 (cit. on p. 31).
- [104] Peter Lightbody, Tomas Krajník, and Marc Hanheide. “An efficient visual fiducial localisation system”. In: *ACM SIGAPP Applied Computing Review* 17.3 (2017), pp. 28–37 (cit. on p. 31).
- [105] Eric Jones, Travis Oliphant, Pearu Peterson, et al. *SciPy: Open source scientific tools for Python*. 2001. URL: <http://www.scipy.org/> (cit. on p. 32).
- [106] Shai Avidan and Ariel Shamir. “Seam carving for content-aware image resizing”. In: *ACM SIGGRAPH 2007 papers*. 2007, 10–es (cit. on p. 32).
- [107] H.K. Yuen et al. “Comparative study of Hough Transform methods for circle finding”. In: *Image and Vision Computing* 8 (1990), p. 71 (cit. on p. 35).
- [108] G. Bradski. “The OpenCV Library”. In: *Dr. Dobbs’ Journal of Software Tools* (2000) (cit. on p. 35).
- [109] Yiqiu Zhao et al. “Particle scale force sensor based on intensity gradient method in granular photoelastic experiments”. In: *New Journal of Physics* 21.2 (2019), p. 023009 (cit. on p. 36).
- [110] diverse authors. *Photoelasticity wiki*. URL: <https://git-xen.lmgc.univ-montp2.fr/PhotoElasticity/Main/-/wikis/home> (cit. on p. 37).

- [111] Osborne Reynolds. “LVII. On the dilatancy of media composed of rigid particles in contact. With experimental illustrations”. In: *Lond. Edinb. Dubl. Phil. Mag.* 20.127 (1885), pp. 469–481 (cit. on p. 52).
- [112] by Kenneth Falconer. *Fractal geometry: mathematical foundations and applications*. 1990 (cit. on p. 62).
- [113] Grevel Lindop. “How Long is the Coast of Britain?” In: *PN Review* 24.6 (1998) (cit. on p. 62).
- [114] Hjortenberg Erik. “Inge Lehmann’s work materials and seismological epistolary archive”. In: *Annals of Geophysics* 52 (Jan. 2010). DOI: 10.4401/ag-4625 (cit. on p. 67).
- [115] Xiaoping Jia. “Codalike multiple scattering of elastic waves in dense granular media”. In: *Physical Review Letters* 93.15 (2004), p. 154303 (cit. on p. 67).
- [116] Roel Snieder. “The theory of coda wave interferometry”. In: *Pure and Applied geophysics* 163.2 (2006), pp. 455–473 (cit. on p. 67).
- [117] Vincent Tournat and VE Gusev. “Nonlinear effects for coda-type elastic waves in stressed granular media”. In: *Physical Review E* 80.1 (2009), p. 011306 (cit. on p. 67).
- [118] Vincent Langlois and Xiaoping Jia. “Acoustic probing of elastic behavior and damage in weakly cemented granular media”. In: *Physical Review E* 89.2 (2014), p. 023206 (cit. on p. 67).
- [119] Jonathan Singh et al. “Coda wave interferometry for accurate simultaneous monitoring of velocity and acoustic source locations in experimental rock physics”. In: *Journal of Geophysical Research: Solid Earth* 124.6 (2019), pp. 5629–5655 (cit. on p. 67).
- [120] Keiiti Aki. “Theory of earthquake prediction with special reference to monitoring of the quality factor of lithosphere by the coda method”. In: *Practical Approaches to Earthquake Prediction and Warning*. Springer, 1985, pp. 219–230 (cit. on p. 67).
- [121] M. Kac. “Can One Hear the Shape of a Drum?” In: *The American Mathematical Monthly* 73.4P2 (1966), pp. 1–23 (cit. on p. 67).
- [122] D. Webb C. Gordon. “You Cant Hear the Shape of a Drum”. In: *American Scientist* 84.1 (1996), pp. 46–55. ISSN: 00030996. URL: <http://www.jstor.org/stable/29775597> (cit. on p. 67).

-
- [123] W. Arrighetti and G. Gerosa. “Can you hear the fractal dimension of a drum?” In: *Applied and Industrial Mathematics in Italy*. World Scientific, 2005, pp. 65–75 (cit. on p. 67).
- [124] E. T. Owens and K. E. Daniels. “Sound propagation and force chains in granular materials”. In: *EPL (Europhysics Letters)* 94.5 (2011), p. 54005 (cit. on pp. 67, 76, 77).
- [125] S. Lherminier et al. “Revealing the Structure of a Granular Medium through Ballistic Sound Propagation”. In: *Phys. Rev. Lett.* 113 (9 2014), p. 098001 (cit. on pp. 67, 71, 77, 81).
- [126] F. Dubourg et al. “The sound of avalanches: from a global to a local perspective”. In: *EPJ Web Conf.* 140 (2017), p. 03015 (cit. on pp. 76, 89).
- [127] William B Nolde. *Evolution of the Field Artillery Acquisition Battalion (FATAB)*. Tech. rep. ARMY COMMAND and GENERAL STAFF COLL FORT LEAVENWORTH KS, 1967 (cit. on p. 76).
- [128] Ross Mallett. *The interplay between technology, tactics and organisation in the First AIF*. University of New South Wales, Australian Defence Force Academy, 1999 (cit. on p. 76).
- [129] JA Pierce. “An introduction to Loran”. In: *IEEE Aerospace and Electronic Systems Magazine* 5.10 (1990), pp. 16–33 (cit. on p. 76).
- [130] J Kasper and C Hutchinson. “The Omega navigation system—An overview”. In: *IEEE Communications Society Magazine* 16.3 (1978), pp. 23–35 (cit. on p. 76).
- [131] E. W. Dijkstra. “A Note on Two Problems in Connexion with Graphs”. In: *Numerische Mathematik* 1 (1959), p. 269 (cit. on p. 81).
- [132] Christian Genest and James V. Zidek. “Combining Probability Distributions: A critique and an annotated bibliography”. In: *Statistical Science* 1 (1 1986), p. 114 (cit. on p. 85).
- [133] Mathias Fink. “Time reversal in acoustics”. In: *Contemporary Physics* 37.2 (1996), pp. 95–109 (cit. on p. 91).
- [134] A Prasadka et al. “Time reversal of water waves”. In: *Physical review letters* 109.6 (2012), p. 064501 (cit. on p. 91).
- [135] M. Harazi et al. “Time reversal of ultrasound in granular media”. In: *The European Physical Journal Special Topics* 226.7 (May 2017), pp. 1487–1497. ISSN: 1951-6401. DOI: 10.1140/epjst/e2016-60259-1. URL: <https://doi.org/10.1140/epjst/e2016-60259-1> (cit. on p. 91).

- [136] Atsuyuki Okabe et al. *Spatial tessellations: concepts and applications of Voronoi diagrams*. Vol. 501. John Wiley & Sons, 2009 (cit. on p. 96).
- [137] Victor Levy dit Vehel et al. “Evolution of the distance between plates in an experimental granular fault. Implications for earthquake forecast”. In: *Compte-rendus de la 21e Rencontre du Non-Linéaire*. Ed. by Falcon E., M. Lefranc, and F. Pétrelis. Non-Linéaire Publications, 2018, pp. 43–48 (cit. on pp. 98, 114).
- [138] Victor Levy dit Vehel et al. “Dilation as a precursor in a continuous granular fault”. In: *EPJ Web of Conferences*. Vol. 249. EDP Sciences. 2021, p. 15006 (cit. on pp. 98, 114).
- [139] Bertrand Rouet-Leduc et al. “Machine Learning Predicts Laboratory Earthquakes”. In: *Geophysical Research Letters* 44.18 (2017), pp. 9276–9282 (cit. on pp. 101, 104).
- [140] Claudia Hulbert, Bertrand Rouet-Leduc, and Paul A Johnson. “A silent build-up in seismic energy precedes slow slip failure in the Cascadia Subduction zone”. In: *arXiv preprint arXiv:1909.06787* (2019) (cit. on p. 101).
- [141] Lior Rokach and Oded Z Maimon. *Data mining with decision trees: theory and applications*. Vol. 69. World scientific, 2007 (cit. on p. 102).
- [142] Xindong Wu et al. “Top 10 algorithms in data mining”. In: *Knowledge and information systems* 14.1 (2008), pp. 1–37 (cit. on p. 102).
- [143] Lior Rokach and Oded Maimon. “Top-down induction of decision trees classifiers—a survey”. In: *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 35.4 (2005), pp. 476–487 (cit. on p. 102).
- [144] Tin Kam Ho. “Random decision forests”. In: *Proceedings of 3rd international conference on document analysis and recognition*. Vol. 1. IEEE. 1995, pp. 278–282 (cit. on p. 103).
- [145] F. Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830 (cit. on p. 103).
- [146] Valeriy V Gavrishchaka and Supriya B Ganguli. “Volatility forecasting from multiscale and high-dimensional market data”. In: *Neurocomputing* 55.1-2 (2003), pp. 285–305 (cit. on p. 112).
- [147] Lean Yu et al. “A multiscale neural network learning paradigm for financial crisis forecasting”. In: *Neurocomputing* 73.4-6 (2010), pp. 716–725 (cit. on p. 112).

- [148] Ryan Lagerquist et al. “Deep learning on three-dimensional multiscale data for next-hour tornado prediction”. In: *Monthly Weather Review* 148.7 (2020), pp. 2837–2861 (cit. on p. 112).
- [149] Victor Levy dit Vehel, Ange Haddjeri, and Osvanny Ramos. “Acoustic localisation in a two-dimensional granular medium”. In: *EPJ Web of Conferences*. Vol. 249. EDP Sciences. 2021, p. 15005 (cit. on p. 114).
- [150] R Clark Jones. “A new calculus for the treatment of optical systems. IV.” In: *Josa* 32.8 (1942), pp. 486–493 (cit. on p. 117).