



**Politecnico
di Torino**

Politecnico di Torino

Master's Degree in Mathematical Engineering

Academic year 2025/2026

Graduation Session July 2026

Statistical properties and moment scaling of seismic ground motion signals

Supervisor:

Prof. Lamberto Rondoni

Co-supervisors:

Prof. Matteo Colangeli

Dr. Federica Di Michele

Candidate:

Giuliana Paradiso

Abstract

Earthquakes exhibit well-documented statistical regularities at the catalogue level, such as the Gutenberg-Richter magnitude-frequency relation and the Omori-Utsu aftershock decay, often interpreted as signatures of seismicity as a driven, dissipative system operating far from thermodynamic equilibrium. However, whether comparable signatures such as heavy-tailed distributions, scale invariance, and anomalous diffusion, can be detected directly within individual seismic ground motion recordings, rather than only in catalogue-level statistics, remains an open question. This thesis investigates whether acceleration, velocity, and displacement signals recorded during two seismic events: the 2009 Mw 6.1 L'Aquila earthquake and the 2024 Mw 3.8 Maritime Alps (Queyras) earthquake, display statistical properties compatible with non-equilibrium complex systems, and whether different seismic phases carry distinct and reproducible statistical signatures. P- and S-wave onset times are estimated independently by two automatic methods: the traditional AR-AIC algorithm, applied within adaptive search windows centred on theoretical arrival times computed from the CRUST1.0 crustal velocity model and hypocentral distances, and the deep-learning picker PhaseNet, pre-trained on the INSTANCE dataset and applied to the full waveform. Coda onset are then estimated through three complementary criteria, a theoretical lapse-time formula, Arias-intensity energy release, and envelope amplitude decay, combined into a median ensemble estimate, allowing each recording to be segmented into up to five temporal windows (pre-event, P-wave, S-wave, coda, and post-event). Empirical probability density functions are computed on the full preprocessed signals for each signal type and tested for departures from Gaussianity, while heavy-tail behaviour is characterised through maximum-likelihood fitting of alternative distributions (Laplace, Student-t, Levy-stable) and Hill-estimator power-law exponents. The scaling of displacement moments is then analysed for each temporal window, signal type, and picking method within the framework of strong anomalous diffusion to estimate the scaling exponents $\zeta(q)$, and the robustness of these results is assessed against alternative picking and filtering configurations (probability threshold and residual-based filtering) and against the choice of coda-onset definition. Empirical probability density functions show consistent departures from Gaussianity across both datasets, with the Levy-stable and Student-t distributions outperforming the Gaussian and Laplace fits via AIC, indicating that heavy-tailed behaviour is a robust feature of these signals. Evidence for anomalous diffusion in the moment scaling spectra is more nuanced: its sign and magnitude depend on seismic phase, picking method, and dataset, with only a small number of configurations remaining stable across all segmentation methods tested. The L'Aquila dataset yields better-constrained scaling fits than the Maritime Alps dataset despite poorer phase-picking quality, plausibly reflecting its larger station ensemble, while showing greater divergence across signal types, consistent with its broader and more heterogeneous network coverage. These results indicate that anomalous scaling signatures in seismic ground motion, where present, are not unconditional, and that systematic robustness testing across methodological choices is necessary to distinguish genuine physical signal from picking and segmentation artefacts.

Contents

List of Symbols	17
1 Introduction	19
1.1 Scientific motivation	19
1.2 Framing of the research problem	19
1.3 Objectives and methodology	20
1.4 Thesis outline	21
2 State of the art	23
2.1 Fundamentals of seismology	23
2.1.1 Seismotectonic framework	23
2.1.2 Mechanics of faulting	24
2.1.3 Rupture mechanics and energy balance	27
2.1.4 The seismic cycle	28
2.2 Statistical scaling laws of seismicity	29
2.2.1 Gutenberg-Richter law	29
2.2.2 Omori-Utsu law	30
2.2.3 Båth's law	31
2.2.4 Generalized Omori law	32
2.2.5 A unified scaling law of aftershock sequences	32
2.3 Physical and complex-systems models of earthquakes	33
2.3.1 The Burridge-Knopoff model	33
2.3.2 The Olami, Feder & Christensen cellular automaton model	33
2.3.3 Earthquakes as a self-organised critical phenomenon	35
2.4 Statistical models of seismicity	36
2.4.1 Trigger model	36
2.4.2 Self-exciting point process	37
2.4.3 The ETAS model	38
2.5 Earthquakes as non-equilibrium complex systems	39
2.5.1 Complexity and scale invariance in seismic processes	39
2.5.2 Non-equilibrium perspective on seismic dynamics	40
2.5.3 Statistical signatures in seismic signals	41
2.5.4 Anomalous diffusion and multiscaling	42
3 Data and preprocessing	45
3.1 Seismic events and datasets	45
3.1.1 Mw 3.8 Maritime Alps earthquake (December 2024)	45
3.1.2 Mw 6.1 L'Aquila earthquake (April 2009)	46

3.2	Data acquisition and ITACA data preprocessing	48
3.2.1	ITACA database	48
3.2.2	ITACA data processing	49
3.3	Additional preprocessing and quality control	49
3.3.1	Metadata	49
3.3.2	Signal preprocessing pipelines	49
4	Phase detection	51
4.1	Seismic phases and phase picking	51
4.1.1	Seismic wave phases	51
4.1.2	Phase picking: from manual to deep-learning approaches	52
4.2	Phase detection using AR-AIC picking	53
4.2.1	Theoretical estimation of P- and S-wave arrival times	54
4.2.2	Adaptive search windows definition	55
4.2.3	Windowed P- and S-phase picking using AR-AIC	56
4.2.4	Coda onset detection	57
4.2.5	Coda end detection	59
4.2.6	Definition of analysis time windows	60
4.2.7	Quality control and validation	61
4.3	Phase picking using PhaseNet	62
4.3.1	PhaseNet	62
4.3.2	INSTANCE dataset	64
4.3.3	Phase picking using PhaseNet pre-trained on INSTANCE dataset	66
5	Statistical analysis methods	69
5.1	Empirical probability density functions and tests for non-Gaussianity	69
5.2	Heavy-tail analysis	70
5.3	Moment scaling analysis	70
5.3.1	Spatial ensemble averaging	71
5.3.2	Computation of moments and scaling exponents	72
5.3.3	Expected behaviour and interpretation	72
6	Results	73
6.1	Empirical probability density functions and non-Gaussianity	73
6.2	Heavy-tail analysis	73
6.3	Phase picking and signal segmentation	74
6.3.1	AR-AIC phase picking	74
6.3.2	Phase picking using PhaseNet	76
6.4	Probability density functions and heavy-tail assessment	80
6.5	Moment scaling analysis	80
6.5.1	INT-41004391 (Queyras, Mw 3.8)	81
6.5.2	IT-2009-0009 (L'Aquila, Mw 6.1)	85
7	Discussion	93
7.1	Non-Gaussianity and heavy-tail behaviour	93
7.2	Phase picking and signal segmentation	94
7.2.1	AR-AIC picking	94
7.2.2	PhaseNet picking	95
7.3	Evidence for anomalous scaling	97

7.3.1	P-wave window	97
7.3.2	S-wave window	98
7.3.3	Coda window	99
7.3.4	Effect of coda onset method	100
7.3.5	Effect of segmentation configuration	102
7.4	Comparative analysis: Queyras vs L'Aquila	102
7.5	Limitations	103
8	Conclusions	105
	Glossary	107
A	Implementation	113
A.1	Overview	113
A.2	Data and storage	113
A.3	Core modules	113
A.3.1	Data loading and cleaning	114
A.3.2	Theoretical arrivals and search windows	114
A.3.3	Phase picking	114
A.3.4	Signal windowing and validation	116
A.3.5	Moment scaling analysis	116
A.3.6	PDF analysis	117
A.3.7	Visualisation	117
A.4	Workflow summary	117
A.4.1	Primary workflow: moment scaling	117
A.4.2	Secondary workflow: PDF analysis	118
A.5	Output organisation	118
A.6	Key design principles	119
A.7	Software dependencies	119
A.8	Reproducibility	120
B	Metadata Fields Description	121
C	Heavy-tail assessment — full results	127
D	Phase picking and signal segmentation	129
D.1	Crustal velocities and theoretical arrival times	129
D.2	INT-41004391 (Mw 3.8, Maritime Alps)	130
D.2.1	AR-AIC segmentation	130
D.2.2	PhaseNet segmentation	131
D.3	IT-2009-0009 (Mw 6.1, L'Aquila)	133
D.3.1	AR-AIC segmentation	134
D.3.2	PhaseNet segmentation	135
E	Moment scaling: full results	139
E.1	INT-41004391 (Mw 3.8, Maritime Alps)	139
E.1.1	AR-AIC segmentation	139
E.1.2	PhaseNet segmentation	142
E.2	IT-2009-0009 (Mw 6.1, L'Aquila)	147

E.2.1	AR-AIC segmentation	147
E.2.2	PhaseNet segmentation	150

List of Figures

2.1	Global distribution of shallow earthquakes ($M \geq 4.5$, depth ≤ 70 km) from 2023 to 2026. Colours indicate focal depth. Data from the U.S. Geological Survey (USGS) earthquake catalog. Plate boundaries from PB2002 dataset.	24
2.2	Anderson’s faulting theory and associated stress regimes. Reproduced from Shen, Schmitt, and Haug, 2018, Fig. 2. The correspondence with the notation used in the original figure is: normal regime $\sigma_1 = S_V$, $\sigma_2 = S_H$, $\sigma_3 = S_h$; strike-slip regime $\sigma_1 = S_H$, $\sigma_2 = S_V$, $\sigma_3 = S_h$; reverse regime $\sigma_1 = S_H$, $\sigma_2 = S_h$, $\sigma_3 = S_V$.	25
2.3	Spring-block representation of earthquake dynamics and its cellular automaton formulation.	34
3.1	Geographical distribution of the accelerometric stations and epicenter. Stations are coloured by PGA (cm/s^2).	46
3.2	Geographical distribution of the accelerometric stations and epicenter for the L’Aquila dataset. Stations are coloured by PGA (cm/s^2).	48
4.1	Schematic illustration of epicentral and hypocentral distances. The epicentral distance is the horizontal distance between the station and the epicenter, while the hypocentral distance accounts for the three-dimensional path from the hypocenter (earthquake focus) to the station.	55
4.2	PhaseNet architecture. The network consists of four down-sampling stages (left) and four up-sampling stages (right), with skip connections (grey arrows) directly concatenating encoder outputs to corresponding decoder layers. Numbers indicate feature map dimensions (channels $\times$ temporal length). Blue blocks represent convolutional layers followed by ReLU activation; orange blocks indicate down-sampling operations (convolution with stride 4); green blocks represent transposed convolutions for up-sampling. Adapted from Zhu and Beroza (2019).	63
4.3	Example of PhaseNet training dataset. The top three panels show seismograms of the ENZ (East, North, Vertical) components, with blue and red vertical lines indicating manually picked P and S arrival times. The bottom panel shows the converted probability distributions for P and S pickers, represented as truncated Gaussian distributions with mean ($\mu = 0$ s) and standard deviation ($\sigma = 0.1$ s). Adapted from Zhu and Beroza (2019).	64
4.4	Map of the 620 stations used in the INSTANCE dataset. Symbol size is proportional to the number of P-phase arrivals per station. Adapted from Michellini et al. (2021).	65

6.1	Signal segmentation for station OGD1 ($d_{\text{epi}} = 67.2$ km, INT-41004391 dataset) using AR-AIC onset detection and Rautian coda onset method, for acceleration (left), velocity (center), and displacement (right) signals. Vertical lines indicate the detected P-wave onset (red), S-wave onset (teal), coda onset (orange), and coda end (blue dashed). Coloured backgrounds denote the pre-event, P-wave, S-wave, coda, and post-event windows.	75
6.2	Signal segmentation for station SPO ($d_{\text{epi}} = 68.2$ km, IT-2009-0009 dataset) using AR-AIC onset detection and Rautian coda onset method, for acceleration, velocity, and displacement signals. Vertical lines indicate the detected P-wave onset (red), S-wave onset (teal), coda onset (orange), and coda end (blue dashed). Coloured backgrounds denote the pre-event, P-wave, S-wave, coda, and post-event windows.	76
6.3	Signal segmentation for station OGD1 ($d_{\text{epi}} = 67.2$ km, INT-41004391 dataset) using PhaseNet onset detection and Rautian coda onset method, for acceleration (left), velocity (center), and displacement (right) signals. Vertical lines indicate the detected P-wave onset (red), S-wave onset (teal), coda onset (orange), and coda end (blue dashed). Coloured backgrounds denote the pre-event, P-wave, S-wave, coda, and post-event windows.	78
6.4	Signal segmentation for station SPO ($d_{\text{epi}} = 68.2$ km, IT-2009-0009 dataset) using PhaseNet onset detection and Rautian coda onset method, for acceleration, velocity, and displacement signals. Vertical lines indicate the detected P-wave onset (red), S-wave onset (teal), coda onset (orange), and coda end (blue dashed). Coloured backgrounds denote the pre-event, P-wave, S-wave, coda, and post-event windows.	80
6.5	Ensemble-averaged moment scaling curves $\langle \Delta x(\tau) ^q \rangle$ vs. τ for the INT-41004391 dataset (Queyras, M_w 3.8), AR-AIC segmentation, Rautian coda onset. Each panel shows a different seismic phase (P-wave, S-wave, coda) and signal type (acceleration, velocity, displacement). Dashed lines indicate power-law fits; moment orders $q = 0.5, 1.0, 2.0, 3.0$ are shown.	82
6.6	Scaling exponent spectrum $\zeta(q)$ vs. q for the INT-41004391 dataset (Queyras, M_w 3.8), AR-AIC segmentation, Rautian coda onset. The dashed red line shows the normal diffusion reference $\zeta(q) = q/2$. Error bars represent one standard error of the log-log fit slope.	83
6.7	Ensemble-averaged moment scaling curves for the INT-41004391 dataset (Queyras, M_w 3.8), PhaseNet segmentation with probability threshold $P \geq 0.3$, Rautian coda onset. Layout as in Figure 6.5.	84
6.8	Scaling exponent spectrum $\zeta(q)$ for the INT-41004391 dataset (Queyras, M_w 3.8), PhaseNet segmentation with probability threshold $P \geq 0.3$, Rautian coda onset. Layout as in Figure 6.6.	85
6.9	Ensemble-averaged moment scaling curves $\langle \Delta x(\tau) ^q \rangle$ vs. τ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), AR-AIC segmentation, Rautian coda onset. Layout as in Figure 6.5.	87
6.10	Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), AR-AIC segmentation, Rautian coda onset. Layout as in Figure 6.6.	88
6.11	Ensemble-averaged moment scaling curves $\langle \Delta x(\tau) ^q \rangle$ vs. τ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), PhaseNet segmentation with residual filter, Rautian coda onset. Layout as in Figure 6.5.	89

6.12	Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), PhaseNet segmentation with residual filter, Rautian coda onset. Layout as in Figure 6.6.	91
D.1	AR-AIC onset detection for station SPO ($d_{\text{epi}} = 68.2$ km, IT-2009-0009 dataset), vertical component (HNZ), acceleration signal. The red and teal shaded regions indicate the P-wave and S-wave search windows, respectively. Dashed vertical lines show the theoretical arrival times computed from CRUST1.0; solid vertical lines show the onsets detected by AR-AIC. . . .	135
E.1	Scaling exponent spectrum $\zeta(q)$ vs. q for the INT-41004391 dataset (Queyras, M_w 3.8), AR-AIC segmentation, Arias coda onset. The dashed red line shows the normal diffusion reference $\zeta(q) = q/2$. Error bars represent one standard error of the log-log fit slope.	140
E.2	Scaling exponent spectrum $\zeta(q)$ vs. q for the INT-41004391 dataset (Queyras, M_w 3.8), AR-AIC segmentation, Envelope coda onset. Layout as in Figure E.1.	141
E.3	Scaling exponent spectrum $\zeta(q)$ vs. q for the INT-41004391 dataset (Queyras, M_w 3.8), AR-AIC segmentation, Median coda onset. Layout as in Figure E.1.	142
E.4	Scaling exponent spectrum $\zeta(q)$ for the INT-41004391 dataset (Queyras, M_w 3.8), PhaseNet segmentation with probability threshold $P \geq 0.3$, Arias coda onset. Layout as in Figure E.1.	145
E.5	Scaling exponent spectrum $\zeta(q)$ for the INT-41004391 dataset (Queyras, M_w 3.8), PhaseNet segmentation with probability threshold $P \geq 0.3$, Envelope coda onset. Layout as in Figure E.1.	146
E.6	Scaling exponent spectrum $\zeta(q)$ for the INT-41004391 dataset (Queyras, M_w 3.8), PhaseNet segmentation with probability threshold $P \geq 0.3$, Median coda onset. Layout as in Figure E.1.	147
E.7	Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), AR-AIC segmentation, Arias coda onset. Layout as in Figure E.1.	148
E.8	Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), AR-AIC segmentation, Envelope coda onset. Layout as in Figure E.1.	149
E.9	Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), AR-AIC segmentation, Median coda onset. Layout as in Figure E.1.	150
E.10	Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), PhaseNet segmentation with residual filter, Rautian coda onset. Layout as in Figure E.1.	153
E.11	Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), PhaseNet segmentation with residual filter, Rautian coda onset. Layout as in Figure E.1.	154
E.12	Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), PhaseNet segmentation with residual filter, Rautian coda onset. Layout as in Figure E.1.	155

List of Tables

3.1	Source parameters of the analyzed seismic event.	45
3.2	Source parameters of the L'Aquila seismic event.	47
4.1	Parameters used for the <code>ar_pick</code> algorithm.	57
6.1	Summary of non-Gaussianity results for acceleration signals, INT-41004391 dataset ($N = 66$ traces).	74
6.2	Summary of heavy-tail assessment results for acceleration signals, INT-41004391 dataset ($N = 66$ traces). Per-signal values are reported in Table C.1.	74
B.1	Summary of header fields and their description.	121
B.2	Constant metadata fields across all records in the IT-2009-0009 dataset.	123
B.3	Constant metadata fields across all records in the INT-41004391 dataset.	125
C.1	Heavy-tail assessment results for all 66 normalised signals. ν : Student- t degrees of freedom (MLE). $\hat{\alpha}$: Hill power-law exponent (top 10% tail). Lower ν and lower $\hat{\alpha}$ both indicate heavier tails.	127
D.1	Station-specific crustal velocities from CRUST1.0 for both datasets.	129
D.2	Hypocentral distances and theoretical arrival times for both datasets.	130
D.3	Hypocentral distance bins and half-widths of the adaptive AR-AIC search windows. Total range: 11.45–109.99 km.	130
D.4	Residuals between onset times detected by AR-AIC and theoretical arrival times from CRUST1.0, for the three signal types.	130
D.5	Number of components failing each quality control check (AR-AIC segmentation), for the three signal types and the four coda onset detection methods. MonoP, MonoS, and SNR P are independent of the coda method and are shown once per signal type. Denominators differ across methods due to missing windows.	131
D.6	PhaseNet phase detection summary (no filter configuration) for the INT-41004391 dataset across data types.	131
D.7	Residuals between onset times detected by PhaseNet and theoretical arrival times from CRUST1.0, for the three signal types.	132
D.8	Number of components failing each quality control check (PhaseNet, no-threshold configuration), for the three signal types and the four coda onset detection methods. MonoP, MonoS, and SNR P are independent of the coda method and are shown once per signal type. Denominators differ across methods due to missing windows.	132
D.9	PhaseNet phase detection summary (threshold configuration) for the INT-41004391 dataset across data types.	133

D.10	Residuals between onset times detected by PhaseNet (threshold configuration) and theoretical arrival times from CRUST1.0, for the three signal types.	133
D.11	Number of components failing each quality control check (PhaseNet, threshold configuration), for the three signal types and the four coda onset detection methods.	133
D.12	Hypocentral distance bins and half-widths of the adaptive AR-AIC search windows.	134
D.13	Residuals between onset times detected by AR-AIC and theoretical arrival times from CRUST1.0, for the three signal types.	134
D.14	Number of components failing each quality control check (AR-AIC, no-threshold configuration), for the three signal types and the four coda onset detection methods. MonoP, MonoS, and SNR P are independent of the coda method and are shown once per signal type. Denominators differ across methods due to missing windows.	134
D.15	PhaseNet phase detection summary (no probability threshold, no residual filter) for the three signal types. Excluded stations with $t_S \leq t_P$	135
D.16	Residuals between onset times detected by PhaseNet and theoretical arrival times from CRUST1.0, no-threshold configuration, for the three signal types.	136
D.17	Number of components failing each quality control check (PhaseNet, no-threshold configuration), for the three signal types and the four coda onset detection methods. Denominators differ across methods because the number of available windows varies.	136
D.18	PhaseNet phase detection summary (residual filter) for the three signal types. Excluded stations with $t_S \leq t_P$	136
D.19	Residuals between onset times detected by PhaseNet and theoretical arrival times from CRUST1.0, residual filter configuration, for the three signal types.	137
D.20	Number of components failing each quality control check (PhaseNet, residual filter configuration), for the three signal types and the four coda onset detection methods. Denominators differ across methods because the number of available windows varies.	137
E.1	Moment scaling results for acceleration signals, obtained from AR-AIC segmentation (no residual filter), for the INT-41004391 event.	140
E.2	Moment scaling results for velocity signals, obtained from AR-AIC segmentation (no residual filter), for the INT-41004391 event.	141
E.3	Moment scaling results for displacement signals, obtained from AR-AIC segmentation (no residual filter), for the INT-41004391 event.	142
E.4	Moment scaling results for acceleration signals, obtained from PhaseNet segmentation (no threshold configuration), for the INT-41004391 event. .	143
E.5	Moment scaling results for velocity signals, obtained from PhaseNet segmentation (no threshold configuration), for the INT-41004391 event. . . .	143
E.6	Moment scaling results for displacement signals, obtained from PhaseNet segmentation (no threshold configuration), for the INT-41004391 event. .	144
E.7	Moment scaling results for acceleration signals, obtained from PhaseNet segmentation (probability threshold $P \geq 0.3$), for the INT-41004391 event.	144
E.8	Moment scaling results for velocity signals, obtained from PhaseNet segmentation (probability threshold $P \geq 0.3$), for the INT-41004391 event. .	145

E.9	Moment scaling results for displacement signals, obtained from PhaseNet segmentation (probability threshold $P \geq 0.3$), for the INT-41004391 event.	146
E.10	Moment scaling results for acceleration signals, obtained from AR-AIC segmentation (no residual filter), for the IT-2009-0009 event.	148
E.11	Moment scaling results for velocity signals, obtained from AR-AIC segmentation (no residual filter), for the IT-2009-0009 event.	149
E.12	Moment scaling results for displacement signals, obtained from AR-AIC segmentation (no residual filter), for the IT-2009-0009 event.	150
E.13	Moment scaling results for acceleration signals, obtained from PhaseNet segmentation (no residual filter), for the IT-2009-0009 event.	151
E.14	Moment scaling results for velocity signals, obtained from PhaseNet segmentation (no residual filter), for the IT-2009-0009 event.	151
E.15	Moment scaling results for displacement signals, obtained from PhaseNet segmentation (no residual filter), for the IT-2009-0009 event.	152
E.16	Moment scaling results for acceleration signals, obtained from PhaseNet segmentation (residual filter configuration), for the IT-2009-0009 event. .	152
E.17	Moment scaling results for velocity signals, obtained from PhaseNet segmentation (residual filter configuration), for the IT-2009-0009 event. . . .	153
E.18	Moment scaling results for displacement signals, obtained from PhaseNet segmentation (residual filter configuration), for the IT-2009-0009 event. .	154

List of Symbols

General notation

$\langle \cdot \rangle$	Ensemble average
f_s	Sampling rate (Hz)
M_W	Moment magnitude
SNR	Signal-to-noise ratio

Seismic phases and arrival times

t_P	Detected P-wave onset time
t_S	Detected S-wave onset time
t_{coda}	Coda onset time
$t_{\text{coda_end}}$	Coda end time
t_0	Earthquake origin time (relative to recording start)
$t_P^{\text{theo}}, t_S^{\text{theo}}$	Theoretical P- and S-wave arrival times
$\Delta t_P, \Delta t_S$	Half-widths of the adaptive search windows
W'_P, W'_S	Adaptive search windows around the theoretical arrivals
W_P, W_S	Final analysis windows between detected onsets
r_P, r_S	Phase picking residuals (detected minus theoretical onset time)
P_P, P_S	PhaseNet detection probabilities for the P- and S-wave onset
d_{hypo}	Hypocentral distance
d_{epi}	Epicentral distance
h	Hypocenter depth
v_P, v_S	Average crustal P- and S-wave velocities

Moment scaling analysis

$x(t)$	Seismic signal (acceleration, velocity, or displacement)
$\Delta x(\tau)$	Signal increment at time lag τ : $x(t_0 + \tau) - x(t_0)$
τ	Time lag
q	Moment order
$M_q(\tau)$	Ensemble-averaged moment of order q at lag τ
$\zeta(q)$	Scaling exponent spectrum
ξ, ζ_0	Slopes of the two branches of strong anomalous diffusion
α	Crossover moment order (strong anomalous diffusion)
N	Number of signals in the spatial ensemble
R^2	Coefficient of determination of the log-log fit

Statistical analysis

κ	Excess kurtosis
γ	Skewness
μ_k	k -th central moment
α_s	Lévy stability index
$\hat{\alpha}$	Hill estimator of the power-law tail exponent
ν	Degrees of freedom (Student- t distribution)

Chapter 1

Introduction

1.1 Scientific motivation

Earthquakes are among the most destructive natural phenomena affecting human societies. Despite decades of advances in seismology, geophysics, and computational modelling, the reliable prediction of earthquakes in terms of exact timing, location, and magnitude remains beyond current scientific capabilities. This limitation reflects the intrinsic complexity of the physical processes governing earthquake generation: heterogeneous fault systems, nonlinear frictional dynamics, and interactions across a wide range of spatial and temporal scales, making the problem fundamentally chaotic and sensitive to initial conditions.

While deterministic prediction remains elusive, statistical regularities emerge at the macroscopic level. Empirical laws such as the Gutenberg–Richter magnitude–frequency relation and the Omori–Utsu decay of aftershock rates reveal that seismicity exhibits scale-invariant behaviour and complex temporal organisation. These observations suggest that earthquakes may be more fruitfully studied within probabilistic and statistical frameworks, which aim to characterise the collective properties of seismic activity rather than the precise evolution of individual events.

A complementary perspective, rooted in statistical physics, interprets seismicity as a manifestation of a driven, dissipative system operating far from thermodynamic equilibrium. In this framework, the scale-invariant properties of seismicity emerge naturally as a consequence of the non-equilibrium dynamics of the fault system, rather than as accidental regularities of an otherwise unpredictable process. Characterising the statistical fingerprints of such systems - heavy-tailed distributions, anomalous diffusion, and multi-fractal scaling - in individual seismic recordings may provide new tools for understanding the underlying dynamics of earthquake processes and for assessing whether ground motion signals carry detectable signatures of out-of-equilibrium physics.

1.2 Framing of the research problem

Modern accelerometric networks provide high-resolution time series of ground motion, allowing the direct analysis of dynamic quantities such as acceleration, velocity, and displacement. These signals encode the complex interaction between rupture processes, wave propagation, and local site effects.

At the same time, many complex natural systems operating far from thermodynamic equilibrium exhibit universal statistical features, including heavy-tailed probability distri-

butions, scale invariance, and anomalous scaling of higher-order moments. In such systems, fluctuations at different scales obey power-law behaviours, and the scaling of statistical moments can deviate from the linear dependence characteristic of normal diffusion. These phenomena are commonly interpreted within the framework of non-equilibrium statistical mechanics and complex systems theory.

Seismicity itself has long been discussed in connection with scale invariance and , primarily at the level of event magnitudes and occurrence times. However, it remains an open question whether similar universal properties can be detected directly in the statistical structure of individual seismic recordings. In particular, it is not clear whether the probability density functions of ground motion time series exhibit heavy tails, or whether the scaling of statistical moments deviates from normal diffusion in ways consistent with out-of-equilibrium dynamics.

A further dimension of the problem concerns whether different seismic phases exhibit distinct statistical signatures. The temporal segmentation of a seismic recording into pre-event noise, P-wave, S-wave, coda and post-event, reflects physically distinct dynamical regimes, each governed by different propagation and scattering processes. Whether these phases carry different statistical fingerprints, and whether such differences are robust with respect to the phase detection methodology, is a question that has received limited attention in the literature.

The central research problem addressed in this thesis can therefore be formulated as follows: *Do seismic ground motion signals display statistical properties compatible with those of non-equilibrium complex systems characterised by heavy-tailed distributions, scale invariance, and anomalous scaling? And do different seismic phases exhibit distinct and reproducible statistical signatures?*

Framing the problem in this way shifts the perspective from the prediction of individual earthquakes to the identification of structural statistical regularities in seismic signals, contributing to a deeper probabilistic description of seismic processes.

1.3 Objectives and methodology

In order to address the research problem formulated above, this thesis pursues an empirical statistical characterisation of seismic ground motion signals - acceleration, velocity, and displacement - recorded by networks of accelerometric stations during two seismic events of different magnitudes. The analysis investigates the statistical properties of these signals across multiple temporal scales and dynamical regimes.

Specifically, the study aims to:

- characterise the empirical probability density functions of seismic ground motion signals and assess deviations from Gaussian behaviour through statistical testing and shape analysis;
- investigate the presence of heavy-tailed distributions through distributional model comparison and power-law exponent estimation;
- analyse the scaling behaviour of statistical moments across different seismic phases (pre-event, P-wave, S-wave, coda, and post-event) to detect signatures of anomalous or strongly anomalous diffusion;
- assess whether the observed statistical features are robust with respect to the phase detection methodology, by implementing and comparing two independent automatic

phase picking approaches (AR-AIC and PhaseNet) and quantifying the sensitivity of scaling exponents to picking uncertainties;

- evaluate whether the observed statistical signatures are consistent with the theoretical framework of non-equilibrium complex systems characterised by scale invariance and anomalous transport.

The methodological approach combines tools from statistical seismology, signal processing, and scaling analysis. Seismic signals are preprocessed and segmented into physically distinct temporal windows using automatic phase detection algorithms. Empirical probability density functions and moment scaling spectra are then computed on the segmented signals, and the resulting patterns are interpreted within theoretical frameworks developed for complex out-of-equilibrium systems. A systematic sensitivity analysis quantifies the robustness of the results with respect to uncertainties in the phase picking procedure, providing a rigorous foundation for the physical interpretation of the observed scaling behaviour.

1.4 Thesis outline

Chapter 2 presents the theoretical and empirical background necessary for the development of the subsequent analysis. The fundamental concepts of seismology are introduced, including the seismotectonic framework, the mechanics of faulting, and the earthquake rupture process. The principal empirical scaling laws of seismicity are reviewed (Gutenberg–Richter, Omori–Utsu, Bath’s law), followed by physical and complex-systems models of earthquakes, including the Burridge–Knopoff model, cellular automaton formulations, and self-organised criticality. Statistical point-process models are then introduced, from classical trigger models to the ETAS framework. Finally, statistical-physics approaches to seismicity are discussed, interpreting earthquake activity as a driven, dissipative process operating far from thermodynamic equilibrium, and introducing theoretical frameworks for non-Gaussian statistics, anomalous diffusion, and multiscale dynamics in seismic systems.

Chapter 3 describes the dataset and preprocessing procedures. Data preprocessing and quality control procedures are detailed, including metadata standardization and two distinct signal preprocessing pipelines designed for PDF analysis and moment scaling analysis respectively. An exploratory data analysis of both categorical and numerical variables is also provided.

Chapter 4 describes the automatic phase detection pipeline implemented to segment seismic recordings into distinct temporal windows. Theoretical P- and S-wave arrival times are estimated using the CRUST1.0 crustal velocity model and hypocentral distances. Two complementary phase picking approaches are implemented and compared: AR-AIC, a traditional autoregressive method, and PhaseNet, a deep-learning picker pre-trained on the INSTANCE dataset. Coda onset and coda end detection are performed using four independent methods. The resulting phase boundaries are used to define up to five non-overlapping temporal windows (pre-event, P-wave, S-wave, coda, post-event), which are then subjected to a multi-level quality control procedure.

Chapter 5 describes the statistical methods applied to the preprocessed and segmented signals. Empirical probability density functions are computed and tested for departures from Gaussianity via the Anderson–Darling test, excess kurtosis, and skewness. Heavy-tail behaviour is characterised through maximum likelihood fitting of alternative distributions

(Laplace, Student- t , Lévy stable) and power-law exponent estimation via the Hill estimator. Moment scaling analysis is then described, including the spatial ensemble averaging procedure, the computation of scaling exponents $\zeta(q)$, and the expected behaviour across different seismic phases. Finally, a systematic sensitivity analysis quantifies the robustness of the moment scaling results with respect to uncertainties in automatic phase picking.

Chapter 6 presents the results of all analyses applied to the two datasets. Phase picking residuals and segmentation quality control statistics are reported for both AR-AIC and PhaseNet configurations. Distributional properties of acceleration, velocity, and displacement signals are characterised through empirical probability density functions and heavy-tail indicators. Moment scaling results are then reported for all temporal windows, signal types, segmentation methods, and coda onset definitions, including ensemble-averaged scaling curves and scaling exponent spectra $\zeta(q)$.

Chapter 7 interprets the observed scaling behaviour in the context of complex out-of-equilibrium systems. The physical implications of super-diffusive scaling in the P-wave window and the absence of scaling in the S-wave and coda windows are discussed, together with the consistency of results across signal types, picking methods, and coda onset definitions. The contrasting behaviour of the two datasets is examined in relation to differences in magnitude, epicentral distance distribution, and ensemble size. The limitations of the spatial ensemble averaging approach and the constraints imposed by the short available time lags are also addressed.

Chapter 8 summarises the main findings of the analysis, assesses their significance with respect to the hypothesis of anomalous diffusion in seismic ground motion, and outlines directions for future work, including extension to larger multi-event datasets, alternative ensemble averaging strategies, and the application of the framework to different tectonic settings.

The appendices collect supplementary material referenced throughout the thesis. Appendix A describes the implementation of the analysis pipeline, including the software architecture and the main processing modules. Appendix B provides a description of the metadata fields associated with the ITACA database records used in this work. Appendix C reports the complete results of the probability density function analysis and heavy-tail assessment for the INT-41004391 dataset, for all three signal types, including distribution fit parameters and goodness-of-fit statistics. Appendix D reports phase picking residual statistics and quality control results for all configurations. Appendix E reports complete moment scaling tables for all datasets, segmentation methods, filtering configurations, and coda onset definitions, and collects the scaling curves and scaling exponent spectra for coda onset methods not shown in Chapter 6.

Chapter 2

State of the art

A rigorous statistical investigation of seismic signals requires situating them within the mechanical and statistical framework that governs earthquake occurrence. The analysis therefore begins with the fundamental physical principles of earthquake generation, including the mechanical behaviour of faults, the structure of the lithospheric stress field, and the seismic cycle. These elements establish the physical basis for stress accumulation, rupture dynamics, and energy release in the Earth's crust.

On this foundation, the chapter reviews the principal empirical scaling laws of seismicity and the main physical and statistical models proposed to interpret them. Particular attention is devoted to the emergence of scale-invariant relations in earthquake magnitudes, temporal clustering, and triggering mechanisms, as well as to dynamical models that interpret seismicity within the broader context of driven, dissipative systems.

Finally, the discussion is extended to non-equilibrium statistical perspectives, highlighting how concepts such as universality, multiscale behaviour, and anomalous transport provide a natural framework for investigating the statistical structure of seismic signals. In this way, the chapter establishes the conceptual bridge between classical seismological theory and the analysis of probability distributions and scaling properties that constitutes the core objective of this thesis.

2.1 Fundamentals of seismology

2.1.1 Seismotectonic framework

Earthquakes constitute a fundamental expression of active tectonics (Scholz, 2018, Chapter 6), although the proportion of total crustal deformation accommodated seismically, and its variability across different tectonic environments, remains a matter of ongoing investigation. Seismic activity is widely recognized as both a manifestation and, in part, a driver of tectonic processes.

A global distribution of seismicity highlights the regions of the Earth currently undergoing deformation. As shown in Figure 2.1, major plate boundaries are clearly delineated by concentrated seismic belts. The spatial distribution of earthquakes therefore provides a powerful observational constraint on the tectonic state and mechanical behaviour of the lithosphere.




Figure 2.1: Global distribution of shallow earthquakes ($M \geq 4.5$, depth ≤ 70 km) from 2023 to 2026. Colours indicate focal depth. Data from the U.S. Geological Survey (USGS) earthquake catalog. Plate boundaries from PB2002 dataset.

2.1.2 Mechanics of faulting

Brittle tectonics may be considered on two timescales, in which earthquakes represent the short-timescale expression of brittle failure, whereas faulting is the cumulative, long-term process through which permanent tectonic deformation is accommodated (Scholz, 2018, Chapter 3). In this sense, faults record the integrated history of past seismic activity. In this section, faults are treated mechanically as quasi-static shear cracks governed by frictional resistance. The classical theory of faulting is introduced, and then more recent developments concerning fault formation, growth, and structural evolution are discussed.

Anderson's theory of faulting

Anderson's theory of faulting (Anderson, 1905) provides a mechanical framework to explain the orientation and classification of faults in the Earth's crust. The theory is based on two key assumptions: (1) the Earth's surface is a free boundary that cannot sustain shear stress, and (2) faulting occurs when the shear stress on a plane reaches the frictional strength of the material, as described by the Coulomb failure criterion. Failure is expected when

$$\tau = c + \mu\sigma'_n, \quad (2.1)$$

where τ is the shear stress on the fault plane, c is cohesion, μ is the coefficient of friction, and σ'_n is the effective normal stress, defined as:

$$\sigma'_n = \sigma_n - p, \quad (2.2)$$

where σ_n is the total normal stress and p is the pore-fluid pressure. An increase in pore pressure reduces the effective normal stress and therefore lowers the shear stress required for failure. Because the Earth's surface is stress-free, one of the three principal stresses must be vertical. Denoting the principal stresses as $\sigma_1 \geq \sigma_2 \geq \sigma_3$, three fundamental tectonic regimes arise depending on which principal stress is vertical.

- **Normal faulting:** occurs when the maximum principal stress is vertical (σ_1 vertical). This corresponds to an extensional tectonic regime. Fault planes are typically inclined at approximately 60° to the horizontal, consistent with the Coulomb failure condition.

- **Reverse (thrust) faulting:** occurs when the minimum principal stress is vertical (σ_3 vertical), corresponding to horizontal compression. Fault planes typically dip at angles of about 30° .
- **Strike-slip faulting:** occurs when the intermediate principal stress is vertical (σ_2 vertical). In this case, both σ_1 and σ_3 are horizontal, and faults are approximately vertical, accommodating lateral motion.




Figure 2.2: Anderson's faulting theory and associated stress regimes. Reproduced from Shen, Schmitt, and Haug, 2018, Fig. 2. The correspondence with the notation used in the original figure is: normal regime $\sigma_1 = S_V$, $\sigma_2 = S_H$, $\sigma_3 = S_h$; strike-slip regime $\sigma_1 = S_H$, $\sigma_2 = S_V$, $\sigma_3 = S_h$; reverse regime $\sigma_1 = S_H$, $\sigma_2 = S_h$, $\sigma_3 = S_V$.

Although Anderson's theory relies on simplifying assumptions, namely a homogeneous, isotropic medium and a spatially uniform stress field, it captures the first-order mechanical controls on fault orientation and style. Anderson's framework thus establishes a fundamental connection between tectonic loading conditions, principal stress orientation, and the geometry of earthquake-generating faults.

However, natural fault systems often deviate from this idealized picture. In particular, fault strength is governed not only by the orientation of the stress field but also by the

effective normal stress acting on the fault plane. Hubbert and Rubey, 1959 demonstrated that low-angle thrust faults can slip under realistic tectonic stresses if pore-fluid pressures are sufficiently high. Because frictional resistance depends on effective normal stress, elevated fluid pressure can substantially reduce fault strength, allowing large-scale overthrusting without requiring unrealistically large shear stresses. Their work highlighted the central role of effective stress in controlling fault mechanics.

Furthermore, in natural settings faults are rarely optimally oriented with respect to the prevailing stress field. Pre-existing faults may be reactivated even when their orientation deviates from the ideal Andersonian geometry, provided that the resolved shear stress satisfies the Coulomb criterion. The likelihood of reactivation depends on both the friction coefficient and the effective stress. Progressive deformation may also rotate fault planes, modifying their orientation relative to the principal stresses. These processes indicate that earthquake occurrence is controlled not only by the regional stress regime but also by inherited structures, fluid pressure conditions, and frictional properties.

Formation, growth, and interaction of fault systems

Faults do not form as large, through-going discontinuities instantaneously; rather, they emerge from the progressive localisation of distributed brittle deformation (Scholz, 2018, Chapter 3). In the early stages, shear strain is accommodated by the nucleation of microcracks and small shear fractures that develop in orientations consistent with the Coulomb failure criterion. As deformation accumulates, these fractures interact, coalesce, and progressively localise strain, leading to the formation of a discrete fault core embedded within a broader damage zone. This transition from diffuse cracking to localised slip represents the mechanical initiation of a fault.

Subsequent deformation is accommodated by fault growth, which occurs through propagation of fault tips and linkage between adjacent segments. Mechanical interaction between nearby fractures plays a crucial role in this process: stress concentrations at fault tips enhance crack propagation, while elastic stress redistribution promotes the connection of neighboring segments. Fault growth is therefore not purely local but controlled by stress interactions within an evolving fault network.

Empirical studies reveal a scaling relationship between fault length L and average displacement D , commonly expressed as (Cowie and Scholz, 1992):

$$D \propto L. \quad (2.3)$$

This approximately linear scaling indicates geometric self-similarity during fault growth and is consistent with elastic crack models in which stress drop remains approximately scale-invariant over a wide range of fault sizes.

At larger scales, faults rarely occur in isolation but instead form interacting populations characterised by hierarchical organisation and scale-invariant structure. Earthquakes occur within these pre-existing, hierarchically organised fault networks, and their rupture dimensions are geometrically constrained by this inherited structure.

The distribution of fault lengths often follows power-law behaviour (Bonnet et al., 2001), implying the absence of a characteristic scale and a fractal spatial organisation of brittle deformation. The emergence of scale-invariant fault populations suggests that brittle deformation is organised across a wide range of spatial scales, providing a structural basis for the scaling behaviour observed in earthquake statistics.

2.1.3 Rupture mechanics and energy balance

The following section examines the physical mechanisms governing the initiation, propagation, and arrest of earthquake rupture, with particular emphasis on the dynamic interplay between stress accumulation, frictional resistance, fracture energy, and seismic radiation, thereby linking fault mechanics to observable seismic source parameters.

Dynamic rupture propagation

The mechanism of rupture during unstable slip can be described in two complementary ways: as a dynamic brittle fracture problem and as a stick–slip frictional instability (Scholz, 2018, Chapter 4). In fracture mechanics, crack propagation requires a characteristic fracture energy per unit area, a material property that controls the energy dissipated at the rupture front. In the stick–slip framework, rupture occurs when the shear stress on the fault reaches the static friction level and dynamic instability develops according to the governing friction law. Although conceptually distinct, the two descriptions are mathematically equivalent in relating stress drop to slip and energy release.

Rupture initiates when the applied shear stress reaches the yield strength σ_Y at a point on the fault surface, the hypocenter. The rupture front then propagates outward at a velocity V_R , which depends on the available elastic energy and the fracture energy required at the crack tip.

At points ahead of the rupture tip, the local stress is initially below the failure level. Slip is triggered by the stress concentration associated with the crack tip, characterised by the stress-intensity factor K . The initial particle velocity scales with K , and decreases as the rupture front moves past the point.

If rupture encounters a region characterised by velocity-strengthening friction, that is, where frictional resistance increases with slip velocity, the shear stress may increase during slip, resulting in a locally negative stress drop. This behaviour acts as an energy sink and tends to inhibit further rupture propagation. After dynamic slip ceases, the elevated stress level promotes quasi-static relaxation in the form of afterslip.

The dynamic energy balance

An earthquake produces a sudden slip of average amount $\Delta\bar{u}$ over a rupture area A . This slip is associated with a decrease in shear stress on the fault from an initial value τ_1 to a final value τ_2 , such that $\Delta\tau_s = \tau_1 - \tau_2$ defines the static stress drop. Since both slip and stress vary spatially over the rupture surface, these quantities represent spatial averages.

A physically meaningful scalar measure of earthquake size is the seismic moment (Scholz, 2018, Section 4.2.1):

$$M_0 = \mu \Delta\bar{u} A, \quad (2.4)$$

where μ is the shear modulus of the surrounding medium. The seismic moment is estimated from seismic waveforms or geodetic observations. Earthquake size is commonly reported using the moment magnitude scale:

$$M_W = \frac{2}{3} \log_{10} M_0 - 6.07, \quad (2.5)$$

with M_0 expressed in N/m .

The dynamic energy balance during rupture can be written as:

$$W_F = E_R + E_F + E_G, \quad (2.6)$$

where W_F is the mechanical work done by fault slip, E_R is the energy radiated as seismic waves, E_F is the energy dissipated by frictional processes, and E_G is the fracture energy, including the energy required to create new surfaces and other forms of inelastic damage. Rupture propagation is therefore governed by the partitioning of elastic strain energy between radiated energy and dissipative processes at the rupture front.

Earthquakes sequences

Earthquakes seldom occur as isolated events; rather, they are typically embedded within sequences of related shocks. Seismic events are commonly classified into three categories: mainshocks, foreshocks, and aftershocks. Lippiello, Godano, and Arcangelis, 2019 define mainshocks as “events sufficiently isolated in time and space from other larger events”. Foreshocks and aftershocks are then identified retrospectively as events occurring close in space and time before and after the main shock, respectively. Sequences of earthquakes that are not clearly dominated by a single outstanding event are instead referred to as swarms.

2.1.4 The seismic cycle

The first comprehensive conceptualization of the earthquake loading cycle is generally attributed to H. F. Reid. In his analysis of the 1906 California earthquake (Reid, 1911), he formulated the elastic rebound theory, proposing that elastic strain accumulates gradually in the crust due to tectonic loading and is then released suddenly during rupture. In this framework, the pattern of strain accumulation is the reverse of the coseismic strain release, so that over one complete cycle the net result is a permanent offset across the fault with no residual elastic strain in the surrounding medium.

However, this idealized elastic model struggled to explain why strain accumulation appeared to be strongly localised along the fault zone rather than distributed more broadly within the crust. Modern observations indicate that the loading behaviour is considerably more complex and cannot be described by a purely elastic system undergoing steady strain accumulation.

In Reid’s formulation, the recurrence of large earthquakes is expected to be periodic, or at least time-predictable, since rupture occurs when the accumulated elastic strain reaches a critical level. Yet empirical evidence for strict periodicity remains inconclusive and, in some cases, contradictory. Whether large earthquakes follow a predictable loading cycle has profound implications for our understanding of fault mechanics and for earthquake hazard assessment.

While elastic rebound theory provides the foundational mechanical picture of stress accumulation and release, the modern concept of the seismic cycle extends this framework by incorporating additional deformation processes and temporal phases observed in natural fault systems.

The seismic cycle describes the long-term evolution of stress accumulation and release along a fault system and is commonly divided into four mechanically interconnected phases: interseismic, preseismic, coseismic, and postseismic (Scholz, 2018, Chapter 5).

During the interseismic phase, tectonic loading progressively accumulates elastic strain in the seismogenic portion of the fault, which is typically locked at shallow depths while deeper segments may deform aseismically. This gradual storage of elastic energy provides the driving energy for subsequent rupture.

The preseismic phase, when identifiable, involves localised nucleation processes in which slip accelerates toward dynamic instability. The coseismic phase corresponds to the rapid propagation of rupture and to the abrupt release of accumulated elastic strain energy in the form of seismic radiation and frictional dissipation.

Following rupture, the postseismic phase is characterised by transient relaxation processes, including afterslip and viscoelastic stress redistribution, which progressively modify the stress field and contribute to the reloading of the fault system.

Taken together, these phases describe a driven and dissipative system characterised by slow loading and intermittent release, a dynamical setting that naturally motivates the statistical and scaling description of seismicity developed in the following sections.

2.2 Statistical scaling laws of seismicity

This section reviews the principal empirical scaling relations of seismicity and their subsequent theoretical generalisations. By examining magnitude distributions, temporal clustering, and aftershock decay patterns, it outlines the observed scaling structure of earthquakes that motivates the search for underlying dynamical mechanisms.

2.2.1 Gutenberg-Richter law

The Gutenberg–Richter law is an empirical relationship describing the frequency–magnitude distribution of earthquakes, first quantified by Gutenberg and Richter (1944) through the analysis of instrumental records of events with magnitudes between 3.5 and 6.5 in California over the period 1934–1948. They found that the cumulative number N of earthquakes with magnitude greater than or equal to M in that region is well described by the relation

$$\log_{10} N = a - bM, \quad (2.7)$$

where a and b are empirical constants, typically estimated by linear regression in logarithmic space. The parameter a reflects the overall level of seismic activity, while b characterises the relative proportion of small to large earthquakes.

For larger magnitudes, Gutenberg and Richter extended their analysis to global earthquake catalogues, estimating the parameters a and b in relation (2.7) using worldwide data. In most tectonic settings, the b -value is close to 1, typically ranging between 0.8 and 1.2 (Frohlich and Davis, 1993).

In their second paper on earthquake scaling, Gutenberg and Richter, 1956 revisited the definition of magnitude and its connection to physically meaningful parameters. They analyzed the duration t_0 of the strong-motion wave group, its characteristic period T_0 , and the epicentral amplitude A_0 , all reduced to standard site conditions. For local events they found that both the duration and period increase systematically with magnitude, and that the epicentral amplitude can be represented by a weakly nonlinear relation. Introducing the quantity $q_0 = \log(A_0/T_0)$, which is more directly related to particle velocity and thus to kinetic energy, they obtained an empirical expression for $q_0(M)$ and, combining it with the scaling of t_0 , derived a non-linear magnitude–energy relation for body waves,

$$\log_{10} E = 9.4 + 2.14M - 0.054M^2, \quad (2.8)$$

where E is the radiated energy in ergs. This expression is numerically close, over

$1 \lesssim M \lesssim 8.7$, to a form of the type

$$\log E \simeq 9.1 + \frac{7}{4}M + \log(9 - M), \quad (2.9)$$

These relations represent early empirical attempts to relate magnitude to radiated energy; modern seismology typically relies instead on moment magnitude scaling.

Gutenberg and Richter emphasize that magnitude, as originally defined from torsion seismometer traces, is an operational measure rather than an absolute one: the mapping between M and E is only approximate, and can break down for very large, complex ruptures. They therefore advocate refining the magnitude scale to rely more directly on A/T and on radiated energy, while keeping the operational definition as instrument and region independent as possible.

2.2.2 Omori–Utsu law

Omori (1894) investigated the temporal decay of aftershock activity following a main shock. Analyzing the 1891 Nobi earthquake, which occurred in central Japan, together with two other Japanese earthquakes, he observed that the frequency of aftershocks per unit time interval, $n(t)$, at time t after the main shock is well described by the empirical relation

$$n(t) = \frac{K}{t + c}, \quad (2.10)$$

where K and c are positive constants and t denotes the elapsed time since the origin time of the main shock. The parameter K controls the overall productivity of the sequence, while c is a time-shift parameter introduced to avoid the singularity at $t = 0$ and to better fit the early stage of the sequence.

Omori reported that this relation provided an excellent fit to the Nobi aftershock sequence for a remarkably long period, extending up to ten years after the main shock. The Omori law was subsequently shown to describe many other aftershock sequences worldwide, establishing it as one of the fundamental empirical laws of seismology (Utsu, Yoshihiko Ogata, and Matsu'ura, 1995).

However, Utsu (1957) revealed that the decay rate of aftershocks is often slightly faster than predicted by the original formulation. To account for this variability, Utsu (1961) introduced a generalisation, now known as the *modified Omori law* or *Omori–Utsu law*, given by

$$n(t) = \frac{K}{(t + c)^p}, \quad (2.11)$$

where p is an additional positive parameter controlling the rate of decay. The original Omori law is recovered as the special case $p = 1$.

Using data from 41 aftershock sequences, Utsu estimated the parameters p and c and found that p typically lies in the range $0.9 \lesssim p \lesssim 1.8$, with values between 1.1 and 1.4 being the most frequent. Importantly, no systematic correlation between p and the magnitude of the main shock was observed (Utsu, Yoshihiko Ogata, and Matsu'ura, 1995). The parameter c was found to vary significantly from sequence to sequence, partly reflecting physical differences among fault systems and partly the incompleteness of early aftershock detection.

2.2.3 Båth's law

Båth's law (Båth, 1965) is an empirical relationship stating that the average difference between the magnitude of a mainshock M_m and that of its largest aftershock M_{ms} is approximately constant, independently of the mainshock magnitude. This difference is thus defined as

$$\Delta M = M_m - M_{ms}, \quad (2.12)$$

with an observed average value $\Delta M \approx 1.2$.

Whether Båth's law reflects a fundamental constraint associated with stress redistribution and rupture dynamics, or instead arises as a statistical consequence of the Gutenberg–Richter law, remains an open issue in statistical seismology. Two main interpretations have been proposed in the literature. On the one hand, Båth's law has been taken as evidence that mainshocks and aftershocks follow different magnitude distributions, suggesting the existence of distinct physical mechanisms governing their occurrence. On the other hand, it has been argued that the magnitudes of the mainshock and of the largest aftershock may correspond to the first and second order statistics of a set of independent and identically distributed random variables drawn from the same Gutenberg–Richter distribution. In this latter view, the observed constancy of the magnitude difference would emerge naturally from extreme-value statistics, without requiring an additional physical constraint.

In recent decades, the interpretation of Båth's law has progressively shifted from a phenomenological observation to a problem framed within statistical seismology. A probabilistic interpretation was proposed by Lombardi, 2002, who framed the problem in terms of order statistics. In this approach, the magnitudes of the mainshock and of the largest aftershock are treated as the first and second largest realisations drawn from the same Gutenberg–Richter distribution. Within this setting, the observed magnitude gap can be understood as a statistical property of extreme values, without invoking distinct physical processes for different classes of events. A more comprehensive statistical framework is developed in branching-process models of seismicity. David Vere-Jones, Murakami, and Christophersen, n.d. showed that when the Gutenberg–Richter magnitude distribution is combined with the productivity law of aftershocks, a nearly constant magnitude difference arises naturally, provided that the productivity exponent is close to the Gutenberg–Richter b -value. In this formulation, Båth's law is not an independent empirical rule but a consequence of the interplay between scale invariance in magnitudes and exponential growth in triggered activity.

Similarly, Helmstetter and Sornette, 2003 demonstrated within the ETAS framework that Båth's law can be derived from the Gutenberg–Richter law together with realistic aftershock properties. In their analysis, the mainshock is not treated as a physically distinct event, but rather as the largest realisation in a cascade of triggered earthquakes. Once the conditioning on the largest event is properly accounted for, the magnitude difference follows from extreme-value statistics applied to a branching process. Shcherbakov and Turcotte (2004) further refined this framework by proposing a modified formulation of Båth's law, based on the extrapolation of the Gutenberg–Richter relation for aftershock sequences. In this formulation, the magnitude difference is defined as

$$\Delta m^* = m_{ms} - m^*, \quad (2.13)$$

where m^* denotes the magnitude of the largest inferred aftershock, obtained by extrapolating the Gutenberg–Richter distribution to the magnitude at which the expected number of events equals one.

These studies suggest that Båth's law may be an emergent statistical property of scale-invariant seismicity combined with branching triggering mechanisms. This issue motivates the introduction of statistical point-process models in Section 2.4.

2.2.4 Generalized Omori law

Shcherbakov, Turcotte, and Rundle (2004) combined the Gutenberg–Richter frequency–magnitude relation (2.7), Båth's law (2.12), and the modified Omori–Utsu formula (2.11), deriving a generalized expression for the rate of occurrence of aftershocks with magnitudes greater than m as a function of time t :

$$r(t, m) = \frac{(p-1)10^{b(m_{ms}-\Delta m^*-m)}}{c(m)} \frac{1}{[1+t/c(m)]^p} \quad (2.14)$$

This formulation explicitly links the temporal decay of aftershock activity to the underlying magnitude scaling, incorporating the b -value from the Gutenberg–Richter relation, the mainshock magnitude m_{ms} , the characteristic time $c(m)$, the Omori exponent p , and the magnitude difference Δm^* arising from the modified form of Båth's law (2.13). The proposed scaling relation was tested on several aftershock sequences, yielding consistent empirical estimates of the parameters b , p , Δm^* , and $c(m)$ (Shcherbakov, Turcotte, and Rundle, 2004; Shcherbakov, Turcotte, and Rundle, 2006).

2.2.5 A unified scaling law of aftershock sequences

Bak, Christensen, et al. (2002) proposed a unified scaling law for the waiting times between earthquakes, revealing a hierarchical organisation in time, space, and magnitude. The authors analyzed a seismic catalog covering the period 1984–2000 in a region of California. The region was partitioned into cells of size $L \times L$, and the waiting time T was defined as the time interval between successive earthquakes occurring within the same cell. Only events with magnitude greater than $m = \log S$ were considered.

The authors measured the probability density $P_{S,L}(T)$ of waiting times and showed that, for suitable values of the temporal exponent α , the magnitude exponent b (related to the Gutenberg–Richter law), and the fractal spatial dimension d_f , the rescaled variables collapse onto a single universal curve:

$$T^\alpha P_{S,L}(T) = f(TS^{-b}L^{d_f}). \quad (2.15)$$

The scaling function $f(x)$ exhibits two distinct regimes separated by a crossover. For small values of x , the function is approximately constant, implying a power-law decay of $P_{S,L}(T)$ in the original variables and reflecting correlated seismic activity. For large x , the function decays rapidly, consistent with an exponential regime associated with effectively uncorrelated events.

The recurrence of power-law relations in the frequency–magnitude distribution, in aftershock decay, and in unified waiting-time statistics indicates the absence of a characteristic scale in seismicity. Such scale-invariant behaviour suggests that seismic activity is organised across interacting spatial and temporal scales. Whether this hierarchical structure at the level of discrete events extends to the statistical properties of continuous seismic signals remains an open question and motivates further investigation.

2.3 Physical and complex-systems models of earthquakes

This section presents the principal physical and complex-systems models developed to interpret earthquake dynamics. Particular emphasis is placed on mechanisms capable of generating scale-invariant statistics and avalanche-like behaviour in driven, dissipative systems.

2.3.1 The Burridge-Knopoff model

The Burridge–Knopoff (BK) model (Burridge and Knopoff, 1967) is a mechanical spring–block representation of earthquake dynamics originally formulated in one spatial dimension. The system consists of a linear chain of blocks connected to their nearest neighbors by elastic springs of stiffness k_c and attached to a rigid driving plate through springs of stiffness k_L . The driving plate moves at constant velocity, representing tectonic loading.

Each block is subject to a frictional force that depends on slip velocity. When the elastic force acting on a block exceeds the frictional resistance, the block slips, releasing accumulated stress. This slip redistributes stress to adjacent blocks through the coupling springs, potentially triggering further slips in a cascade-like manner.

Despite its simplicity, the BK model captures key ingredients of earthquake dynamics: slow external driving, threshold-controlled instability, stress transfer, and avalanche-like rupture propagation. However, the original formulation remains a dynamical system governed by coupled differential equations and does not explicitly assume scale invariance.


2.3.2 The Olami, Feder & Christensen cellular automaton model

Olami, Feder and Christensen (Olami, Feder, and Christensen, 1992) reformulated the Burridge–Knopoff model into a two-dimensional cellular automaton defined on an $L \times L$ lattice. In this representation, each lattice site corresponds to a block characterised by a local stress variable $F_{i,j}$.


The system is driven uniformly by a slow increase of stress at all sites, mimicking tectonic loading. When the stress at a given site reaches a prescribed threshold F_{th} , the site becomes unstable and relaxes, redistributing a fraction of its stress to its four nearest neighbors. This redistribution may induce further instabilities, generating a cascade of threshold failures that are identified with an earthquake event.

In contrast to the conservative sandpile model, the OFC model is generally nonconservative: only a fraction of the released stress is transferred to neighboring sites, while the remainder is dissipated. The degree of conservation is controlled by the elastic coupling parameters.

The mapping from the dynamical spring–block description to a threshold-based cellular automaton emphasizes the key ingredients commonly associated with scale-free behaviour in driven dissipative systems: slow driving, local thresholds, and stress redistribution. A schematic representation of the mechanical spring–block picture and the cellular automaton model is shown in Figure 2.3.



(a) Schematic representation of the spring-block system underlying the OFC cellular automaton model of earthquakes. Blocks are coupled to a slowly moving upper plate and to their neighbours by springs of stiffness K_L and K_c , respectively, above a fixed lower plate. Figure reproduced from *Statistical physics approach to understanding the multiscale dynamics of earthquake*, (Rundle et al., 2003).



(b) Schematic representation of the continuous, nonconservative cellular automaton for earthquakes: the central site (i, j) with displacement $u_{i,j}^t$ is coupled to its four nearest neighbours $(i \pm 1, j)$ and $(i, j \pm 1)$ through the elastic constants K_1 and K_2 . Figure adapted from *Self-Organised Criticality in a Continuous, Nonconservative Cellular Automaton Modelling Earthquakes*, (Olami, Feder, and Christensen, 1992).

Figure 2.3: Spring-block representation of earthquake dynamics and its cellular automaton formulation.

2.3.3 Earthquakes as a self-organised critical phenomenon

Self-organised criticality

The concept of self-organised criticality was introduced by Bak, Tang and Wiesenfeld to describe extended dissipative systems that naturally evolve towards a critical state characterised by scale-invariant spatial and temporal correlations, without the need for external fine-tuning of control parameters (Bak, Tang, and Wiesenfeld, 1987). The temporal fingerprint of criticality is the presence of flicker (or $1/f$) noise, while its spatial signature is the emergence of scale-invariant, fractal structures. Bak et al. proposed that a broad class of interacting dissipative systems can be understood within this single organising principle linking nonlinear dynamics to scale invariance. The evolution towards the self-organised critical state is independent of initial conditions, so that the critical state acts as an attractor of the dynamics.

The sandpile model

The canonical example of a self-organised critical system is the sandpile model illustrated by Bak, Tang, and Wiesenfeld (1988). Consider the formation of a sandpile on a table where grains of sand are added one at a time. At the beginning the pile is nearly flat, and individual grains remain close to where they land; their motion can be understood in terms of simple local mechanical properties. As the pile grows, its slope increases until it reaches a critical value at which an additional grain can trigger a small slide. Conversely, if we start from an initial configuration with an excessively steep pile, it will spontaneously collapse until it reaches this same critical state. The critical state is thus an attractor for the dynamics, and the quantity that exhibits $1/f$ noise is the flux of grains leaving the pile. As the pile is built up, the size of the largest avalanches increases, and at the critical point avalanches occur on all scales, up to the system size. In this regime, the behaviour of the system can no longer be understood in terms of the motion of individual grains alone: the avalanches define a collective dynamics that is meaningful only at the level of the entire pile. In this sense, the sandpile is an example of a complex system.

One can then construct a “synthetic” earthquake catalog by counting how many avalanches occur of each size. Defining the “magnitude” of an avalanche as the logarithm of its size, and plotting the logarithm of the number of avalanches of a given magnitude versus that magnitude, one finds that the avalanches obey a Gutenberg–Richter–type power-law relation (Bak, 1996). The emergence of this scale-free distribution indicates that the stationary state is critical, and that the pile has self-organised into a critical state. A geometric analysis further shows that the stationary sandpile has a fractal profile: the avalanches carve out fractal structures in the bulk of the pile. One can also verify that the critical behaviour is robust with respect to reasonable modifications of the model: the power-law statistics of avalanches persist under changes of microscopic details, and are not merely a consequence of the randomness in the driving (i.e., the random choice of the site where grains are added).

Earthquakes as a critical phenomenon

Earthquakes have often been interpreted as a possible example of self-organised criticality (Bak, 1996; Bak and Tang, 1989). For long periods the Earth’s crust appears quiescent, in a state of apparent stasis, and then this tranquillity is interrupted by bursts of intermittent activity, ranging from small earthquakes to rare, devastating events.

In Bak's interpretation, the key empirical feature is the Gutenberg–Richter law: the distribution of earthquake magnitudes follows a power law over several orders of magnitude. This is a statistical scaling relation: it specifies how many earthquakes of one size occur compared with those of another size, but it does not predict the time or location of individual events. Within the SOC framework proposed by Bak, Tang, and Wiesenfeld, 1988, the Gutenberg–Richter law has been interpreted by some authors as a possible signature of critical dynamics, indicating the absence of a characteristic scale in seismicity.

To provide a physical picture of this behaviour, Bak and Tang (1989) explicitly refer to the classical Burridge–Knopoff block–spring model of earthquake generation. In this framework the fault is represented schematically as an array of elastically coupled blocks driven by tectonic loading and subject to stick–slip failure at a threshold stress. A slip event on one block redistributes stress to its neighbours and can trigger a cascade of failures, which in the SOC picture plays the role of an earthquake. It is precisely this threshold dynamics and stress transfer that makes the block–spring system mathematically analogous to the sandpile automata model of self–organised criticality, reproducing the same type of scaling as the Gutenberg–Richter law.

In the original sandpile model (Bak, Tang, and Wiesenfeld, 1988) the dynamics is strictly conservative: when a site topples, the entire excess sand is redistributed to neighbouring sites. This is a good approximation for literal sandpiles, but it is not realistic for earthquakes. A more careful analysis of the block–spring analogue shows that the stress transferred to neighbouring blocks can be smaller than the stress relieved on the slipping block, so that the dynamics is effectively non–conservative. Bak and Tang (1989) demonstrated that when conservation is relaxed in sandpile–like models the system remains scale–free only up to a cutoff that depends on the degree of conservation: small avalanches still follow a power law, but large events are suppressed and true criticality is lost. In this perspective, the observed Gutenberg–Richter scaling over a wide range of magnitudes suggests that, on geological time scales, the crust has self–organised into an effectively critical state, in which the next slip event may be a small earthquake or a global catastrophe, with no distinguished scale in between.

In the following chapters we will discuss to what extent seismicity actually behaves like a self–organised critical system, and where this analogy breaks down.

2.4 Statistical models of seismicity

This section presents the main statistical frameworks developed to describe earthquake occurrence as a stochastic process. In particular, point–process models and self–exciting mechanisms are introduced to account for temporal clustering and triggering effects, providing a probabilistic description of seismicity complementary to physical and complex–systems approaches.

2.4.1 Trigger model

A trigger model is a Neyman–Scott type cluster point process, that is, a Poisson cluster process in which primary (parent) events generate independent clusters of secondary (offspring) events. Stochastic point processes describe systems in which events occur at discrete instants of time, either singly or in small groups, rather than evolving continuously in time. The probabilistic structure of a point process is fully characterised by the

distributions of the number of events occurring in given time intervals (or collections of intervals), and in particular by its first- and second-order properties.

In seismology, a trigger model was introduced by D. Vere-Jones and Davies (1966) in their study of New Zealand seismicity. Seismic events are classified into primary (or trigger) events and secondary events. The fundamental idea is that earthquakes occur in clusters generated by the underlying trigger events.

Primary events are assumed to follow a homogeneous Poisson process with constant rate μ . Each primary event occurring at time t_i initiates a cluster of secondary events. Conditionally on the occurrence of a primary event at time t_i , the intensity of the associated secondary activity is described by

$$g(t) = \begin{cases} Af(t - t_i), & \text{for } t \geq t_i, \\ 0, & \text{for } t < t_i. \end{cases} \quad (2.16)$$

where $f(t)$ is a non-negative decay function satisfying the normalisation condition

$$\int_0^\infty f(t) dt = 1,$$

and A represents the number of secondary events generated by the primary event. The random variable A is assumed to have finite mean a and variance v . In practical applications, a specific functional form must be assigned to the decay kernel $f(t)$, such as an exponential function or a power-law decay consistent with the modified Omori law.

Two key assumptions characterise the model: (1) distinct primary events generate independent clusters of secondary events, and (2) the random variables A associated with different primary events are independent. Under these assumptions, the resulting seismicity process is a stationary cluster process with overall mean rate $\mu(1 + a)$. Importantly, secondary events do not generate further shocks. This assumption prevents recursive triggering and therefore limits the model's ability to reproduce realistic aftershock cascades.

In order to apply the trigger model to seismic catalog data, the parameters governing the process must be estimated from the observed occurrence times. These include the background rate μ of primary events, the statistical properties of the cluster size A , and the parameters defining the decay kernel $f(t)$.

A fundamental difficulty arises in likelihood-based inference for classical cluster models: since the underlying cluster structure is unobserved, it is not known which events are primary and which are secondary. As a consequence, the full likelihood involves summation over all possible allocations of events into primary and secondary clusters, leading to a combinatorially complex expression.

Two practical likelihood-based approaches have been proposed. The first relies on the spectral likelihood derived from the theoretical spectral density of the trigger model (A. G. Hawkes and Adamopoulos, 1973). The second approach is based on the conditional intensity function of the point process (Yoshihiko Ogata, 1988), which avoids explicit enumeration of cluster configurations and allows parameter estimation through maximization of the point process likelihood. A detailed discussion of this likelihood-based framework in the context of descriptive trigger models is provided by Yoshihiko Ogata (2001).

2.4.2 Self-exciting point process

A. G. Hawkes and Adamopoulos (1973) compared classical trigger models with a self-exciting point-process formulation in which every event may generate offspring events. The

general structure of a self-exciting process was originally introduced by Alan G. Hawkes (1971). In this framework, the conditional intensity at time t is defined as

$$\lambda(t) = \mu + \sum_{t_i < t} g(t - t_i), \quad (2.17)$$

where μ represents the background rate of spontaneous events and $g(t - t_i)$ describes the contribution to the occurrence rate at time t from a previous event at time t_i .

In contrast with classical trigger models, no structural distinction is imposed between primary and secondary events: each event can in turn generate further events. The process therefore admits a branching interpretation in continuous time, where an initial background event may initiate a cascade of descendants across multiple generations.

A key quantity characterising the process is the *branching ratio*,

$$n = \int_0^\infty g(t) dt, \quad (2.18)$$

which represents the expected number of direct offspring generated by a single event. The stability of the process depends on this parameter: if $n < 1$, the process is subcritical and admits a stationary regime; if $n = 1$, it is critical; if $n > 1$, the process becomes explosive and the expected number of events diverges. This branching structure provides a natural probabilistic framework for describing the epidemic-like propagation of seismic activity.

The Epidemic Type Aftershock Sequence (ETAS) model introduced in the following section represents a magnitude-dependent extension of this self-exciting formulation, specifically adapted to earthquake catalogs.

2.4.3 The ETAS model

The Epidemic Type Aftershock Sequence (ETAS) model, originally proposed by Yoshihiko Ogata (1998), is a stochastic point-process model designed to describe the occurrence of earthquakes with magnitude M_0 and larger within a given region and time window. A historical overview of its development and subsequent extensions is provided by Yoshihiko Ogata (2024).

The model decomposes seismic activity into two components: a background activity and a triggered (aftershock) activity. The background activity is modeled as a homogeneous Poisson process with constant rate μ . Superimposed on this background component is the triggered activity: each earthquake, including those that are themselves aftershocks, can generate its own sequence of aftershocks. In this way, the ETAS model allows for recursive triggering, leading to cascades of events across multiple generations.

The triggered activity generated by the i -th earthquake occurring at time t_i with magnitude M_i is described by

$$\nu_i(t) = \frac{K_0}{(t - t_i + c)^p} e^{\alpha(M_i - M_0)}, \quad t > t_i, \quad (2.19)$$

where K_0 , α , c , and p are positive constants characterising the seismicity of the region. The parameter p controls the temporal decay rate according to the modified Omori law, c prevents divergence at short times, α quantifies the dependence of aftershock productivity on magnitude, and K_0 sets the overall productivity level.

The total conditional intensity of the ETAS model is obtained by summing the contributions of all past earthquakes and adding the background component:

$$\lambda(t|H_t) = \mu + \sum_{t_i < t} \frac{K_0}{(t - t_i + c)^p} e^{\alpha(M_i - M_0)} \quad (2.20)$$

From a probabilistic perspective, the ETAS model can be interpreted as a branching process in continuous time, whose effective branching ratio depends on both the temporal decay kernel and the magnitude distribution. The stability of the process requires that the expected number of offspring generated by a single event remains below unity; otherwise, the process becomes supercritical and the expected number of events diverges.

A crucial feature of the ETAS formulation is that it does not impose a structural distinction between mainshocks and aftershocks: all events are treated as realisations of the same stochastic mechanism, and the classification into mainshocks and aftershocks emerges only retrospectively from the cascade structure.

For a sequence of occurrence times with associated magnitudes, the parameter vector $\theta = (K_0, \alpha, c, p, \mu)$ can be estimated by maximizing the log-likelihood function

$$\log L(\theta) = \sum_{t_i < T} \log \lambda_\theta(t_i|H_{t_i}) - \int_0^T \lambda_\theta(t|H_t) dt, \quad (2.21)$$

which follows from the general theory of point processes. The tractability of this likelihood has played a central role in the widespread adoption of the ETAS model in statistical seismology.

2.5 Earthquakes as non-equilibrium complex systems

This section discusses the interpretation of seismic processes within the framework of non-equilibrium statistical mechanics and complex systems theory. In contrast with statistical seismology, which focuses on phenomenological stochastic models of earthquake occurrence (see Section 2.4), the statistical physics perspective seeks to relate the macroscopic statistical properties of seismicity to the underlying physical mechanisms governing stress accumulation, rupture, and interaction between faults. By considering earthquakes as the outcome of driven, dissipative dynamics, this section introduces concepts such as multiscale organisation, non-Gaussian fluctuations, and anomalous scaling, which motivate the subsequent analysis of the statistical structure of seismic signals.

2.5.1 Complexity and scale invariance in seismic processes

While the macroscopic manifestations of seismicity are well documented, the microscopic mechanisms governing brittle deformation of the Earth's crust remain only partially understood and are often inaccessible to direct observation. In this context, statistical physics provides a framework for bridging the gap between microscopic interactions and the collective macroscopic behaviour observed in fault networks and regional seismicity.

As discussed by Michas (2024), earthquakes exhibit the key features of complex systems. They involve a large number of interacting degrees of freedom, associated with networks of faults and fractures; they display nonlinear dynamics, long-range interactions, and hierarchical organisation across a wide range of spatial and temporal scales. Their

evolution is intermittent, with long phases of slow stress accumulation interrupted by abrupt rupture events. Moreover, seismicity is characterised by scale invariance, as evidenced by empirical laws such as the Gutenberg–Richter magnitude distribution and the Omori–Utsu aftershock decay (see Section 2.2). However, the existence of universal scaling behaviour in seismicity remains debated, as several empirical parameters show significant variability across tectonic environments (Sornette and Werner, 2009, cap. iv.2). These properties of emergent collective behaviour, multiscale organisation, and intermittency, support the interpretation of earthquakes within the framework of complex systems and statistical physics. Nevertheless, the origin of this complexity remains debated and may arise from both dynamical interactions and structural heterogeneity of the crust (Sornette and Werner, 2009, cap. v.1).

Statistical physics offers a coherent conceptual approach to the micro–macro upscaling problem in seismicity, aiming to derive large-scale statistical regularities from underlying microscopic interactions. Nevertheless, its application to earthquake dynamics has yielded only partial success, and a number of proposed scaling laws and universality claims remain actively debated (Sornette and Werner, 2009, Sec. VI.1).

Within the self-organised criticality paradigm, discussed in Section 2.3.3, earthquakes are interpreted as a critical phenomenon in a driven, dissipative system. In this view, the Earth’s crust evolves toward a dynamical stationary state in which energy is slowly accumulated and intermittently released through avalanches spanning a broad range of scales. Seismicity then emerges as a manifestation of scale-invariant dynamics, with self-similar size distributions and fractal fault geometries reflecting long-range correlations in space and time.

2.5.2 Non-equilibrium perspective on seismic dynamics

As illustrated in Section 2.1.4, seismic processes may be interpreted as the macroscopic manifestation of a driven, dissipative dynamical system operating far from thermodynamic equilibrium. Tectonic loading continuously injects energy into the Earth’s crust through the slow relative motion of lithospheric plates, while stress is released intermittently through brittle rupture and frictional dissipation along faults. This combination of slow external driving and abrupt energy release prevents the system from relaxing to equilibrium and instead leads to a dynamical stationary regime characterised by intermittent activity across multiple spatial and temporal scales.

In addition to the continuous injection and intermittent release of energy, seismic activity is strongly influenced by stress redistribution and interactions between faults. Slip on a fault modifies the surrounding stress field, altering the mechanical conditions for failure on neighboring fault segments (see Section 2.1.2). As a consequence, earthquake occurrence is not purely local but involves long-range interactions mediated by the elastic stress field in the crust.

These interactions can produce cascades of triggered events, in which a single rupture redistributes stress and increases the probability of subsequent failures in nearby regions. The resulting dynamics is therefore characterised by temporal clustering and hierarchical organisation of seismic activity. Such cascade-like processes are consistent with branching and self-exciting models of seismicity (see Section 2.4), which describe earthquakes as interacting events within a driven non-equilibrium system.

Since earthquake generation involves driven dissipative dynamics far from equilibrium, its statistical description cannot generally rely on the classical framework of equilibrium

thermodynamics based on Boltzmann–Gibbs entropy. In that formulation the entropy is additive for statistically independent subsystems and proportional to the number of independent subsystems, an assumption that is appropriate when correlations are short ranged and interactions remain local.

Seismic systems, however, exhibit strong correlations in space, time, and magnitude. Stress redistribution along fault networks couples distant regions of the crust, and the probability of occurrence of a given event depends on the previous history of the system. As a consequence, correlations may persist across a wide range of spatial and temporal scales. Within this perspective, seismicity has sometimes been interpreted as a near-critical phenomenon characterised by intermittent activity and avalanche-like dynamics, where large events may be viewed as manifestations of dynamical transitions in the system (Sigalotti, Ramirez-Rojas, and Vargas, 2023).

To account for such long-range correlations, Tsallis (1988) proposed a generalisation of Boltzmann–Gibbs statistical mechanics based on a non-additive entropy functional. This framework, known as non-extensive statistical mechanics (NESM), allows for statistical descriptions in which correlations across multiple scales are explicitly incorporated. Over the past decades, NESM has been applied to a variety of non-equilibrium complex systems and has been explored as a possible framework for describing statistical properties of seismicity.

If seismic dynamics indeed results from driven non-equilibrium interactions across multiple scales, such complexity may be reflected in the statistical structure of observable seismic signals.

2.5.3 Statistical signatures in seismic signals

In non-equilibrium statistical mechanics, the theoretical description of dissipative systems remains significantly less complete than that of equilibrium systems. While equilibrium statistical mechanics provides a well-established framework linking microscopic dynamics to macroscopic thermodynamic behaviour, no equally general theory exists for systems far from equilibrium (Colangeli, Rondoni, and Verderosa, 2014). As a consequence, the statistical characterisation of such systems often relies on the study of fluctuations and probability distributions of observable quantities.

A key insight emerging from modern developments in non-equilibrium statistical mechanics is that fluctuations are not merely random perturbations around equilibrium states, but may encode relevant information about the underlying dynamics of the system. Exponential averages of path-dependent observables, whose value depends on the trajectory of a system in phase space rather than on its instantaneous state, are highly sensitive to the tails of probability distributions and therefore provide information about the dynamical properties of non-equilibrium systems (Colangeli, Di Francesco, and Rondoni, 2024).

In particular, fluctuation relations establish general constraints on the probability of fluctuations in non-equilibrium systems, linking microscopic dynamics to macroscopic irreversibility (Colangeli, Rondoni, and Verderosa, 2014). These results are closely related to the theory of large deviations, which provides a mathematical framework for describing the probability of rare events that significantly depart from typical behaviour.

Within this framework, complex natural systems driven far from equilibrium may exhibit large fluctuations whose statistical properties reflect the mechanisms governing their dynamics. Earthquake dynamics provides a natural example of such a system. The interplay between slow tectonic loading, stress redistribution, and intermittent rupture

events generates fluctuations across a wide range of temporal and spatial scales. As a consequence, observable seismic signals may display non-Gaussian statistical features and anomalous scaling behaviour, making the study of their probability distributions a useful tool for probing the dynamical processes governing seismic activity.

A further difficulty arises from the fact that the microscopic mechanisms governing stress transfer, frictional sliding, and fracture propagation in the Earth's crust are not directly observable. The corresponding microscopic phase space therefore remains inaccessible, and the statistical description of seismicity must rely primarily on observable time series rather than on ensemble averages defined at the microscopic level.

Within this context, deviations from Gaussian statistics in seismic observables have sometimes been interpreted within the framework of superstatistics, which describes complex systems as superpositions of local equilibrium states characterised by slowly fluctuating intensive parameters (Iliopoulos et al., 2019). Such approaches provide a possible statistical interpretation of heavy-tailed distributions and multiscale variability observed in non-equilibrium complex systems.

2.5.4 Anomalous diffusion and multiscaling

A transport process $x(t)$ is said to exhibit anomalous diffusion if its (MSD), $\langle |x(t) - x(0)|^2 \rangle$, does not grow linearly with time (Vollmer, Giberti, et al., 2024). In such processes the probability density function (PDF) of the displacement is typically non-Gaussian, indicating that signatures of anomalous transport extend beyond second-order statistics and affect the full hierarchy of moments. Deviations from Gaussian behaviour often appear in the tails of the distribution: in certain dynamical regimes, large fluctuations associated with the tails scale differently from the small fluctuations that characterise the bulk.

More generally, the time scaling of the q -th moment of the displacement can be written as

$$\langle |x(t) - x(0)|^q \rangle \sim t^{q\nu_q/2}, \quad (2.22)$$

where the scaling exponent ν_q may depend on the order q . The case $\nu_q = 1$ for all q corresponds to normal diffusion, whereas deviations from this condition indicate anomalous transport. When ν_q depends on q , simple scale invariance breaks down and the process is referred to as *strong anomalous diffusion*.

A commonly observed scenario of strong anomalous behaviour is characterised by two distinct scaling regimes, in which the exponent of t follows two different linear dependencies on the moment order:

$$\frac{q\nu_q}{2} = \begin{cases} \xi q, & q \leq \alpha, \\ \zeta q - \alpha(\zeta - \xi), & q > \alpha, \end{cases} \quad (2.23)$$

where ξ , ζ , and α are model-dependent parameters. In this framework, lower-order moments are governed by one effective scaling exponent, while higher-order moments are dominated by rare, large fluctuations. Such a piecewise linear structure of the moment scaling has been observed in a variety of complex systems.

Vollmer, Giberti, et al. (2024) revisit the theory of strong anomalous diffusion from the perspective of scaling theory and asymptotic matching of the PDF $\mathcal{P}(\Delta x, t)$ of the displacement $\Delta x(t) = |x(t) - x(0)|$ over a time interval t . They show that, for sufficiently large times, the PDF can be decomposed into two branches corresponding to two distinct asymptotic regimes, which overlap over an interval that grows with time. This behaviour

can be expressed through the following ansatz:

$$\mathcal{P}(\Delta x, t) d\Delta x = \begin{cases} \mathcal{F}(z) dz & \text{with } z = \frac{\Delta x}{t^\xi} \leq c t^{\iota-\xi}, \\ g(t) \mathcal{J}(v) dv & \text{with } v = \frac{\Delta x}{t^\zeta} > c t^{\iota-\zeta}, \end{cases} \quad (2.24)$$

where ζ, ξ, c and ι are parameters for which $0 < \xi < \iota < \zeta$, $c > 0$.

In addition to the scaling of moments and probability distributions, information about anomalous transport can also be obtained from the temporal correlation structure of the displacement. In particular, the displacement autocorrelation function is defined as

$$\phi(t_1, t_2) = \langle (x(x_0, t_1) - x_0)(x(x_0, t_2) - x_0) \rangle, \quad (2.25)$$

which measures correlations between displacements evaluated at two different times. In systems exhibiting strong anomalous diffusion, this quantity can display nontrivial scaling behaviour with respect to multiple time scales (Vollmer, Rondoni, et al., 2021).

If seismic acceleration signals reflect the multiscale nature of the underlying rupture dynamics, their empirical probability density functions may exhibit signatures of anomalous or strongly anomalous scaling. While scale invariance and statistical laws are well documented at the level of earthquake catalogs, the statistical structure of seismic ground motion signals (acceleration, velocity, and displacement) has received comparatively less attention, particularly with regard to heavy-tailed distributions, anomalous diffusion, and multiscaling properties.

Chapter 3

Data and preprocessing

3.1 Seismic events and datasets

3.1.1 Mw 3.8 Maritime Alps earthquake (December 2024)

The first analyzed dataset corresponds to a single moderate crustal earthquake that occurred on 9 December 2024 in the Maritime Alps region, along the border between southern France and northwestern Italy. The event was recorded by the Italian National Seismic Network and catalogued by INGV with moment magnitude $M_w = 3.8$. The hypocenter was located at a depth of approximately 10 km. The event parameters are summarized in Table 3.1.

Parameter	Value
Event ID	INT-41004391
Origin time (UTC)	9 December 2024, [10:41:21]
Epicenter (latitude)	[44.5127]°N
Epicenter (longitude)	[06.8533]°E
Hypocentral depth	10.4 km
Local magnitude (M_L)	3.8
Moment magnitude (M_w)	3.8
Focal mechanism	Strike-slip faulting

Table 3.1: Source parameters of the analyzed seismic event.

The corresponding dataset comprises 66 three-component records from 22 stations, primarily belonging to the Italian National Seismic Network. All stations are equipped with digital accelerometers operating at a uniform sampling rate of 200 Hz ($\Delta t = 0.005$ s), with typical recording durations of approximately 240 s, yielding roughly 48,000 samples per component. The acquisition system was configured with normal trigger (NT), ensuring that pre-event background noise was included in the recordings, an essential feature for signal-to-noise analysis and phase onset detection.

Each station provides three orthogonal components of ground motion: two horizontal components (East–West and North–South) and one vertical component. Most stations employ high broadband accelerometers (HN sensor type, 60 files), while six stations use high-gain accelerometers (HG sensor type).

The stations are distributed over epicentral distances ranging from approximately 4 km (near-field station SURF) to 100 km (far-field stations), providing spatially diverse sampling of ground motion across different geological and topographical settings. This range spans both near-source and regional propagation regimes, allowing investigation of distance-dependent ground motion characteristics. Figure 3.1 shows the geographical distribution of the stations, located across the border between Italy and France, with the epicenter marked by a red star.




Figure 3.1: Geographical distribution of the accelerometric stations and epicenter. Stations are coloured by PGA (cm/s^2).

3.1.2 Mw 6.1 L'Aquila earthquake (April 2009)

The second analyzed dataset corresponds to the M_w 6.1 L'Aquila earthquake that occurred on 6 April 2009 in the Abruzzo region of central Italy. The event was recorded by multiple seismic networks and catalogued by INGV. The hypocenter was located at a depth of approximately 8.3 km beneath the town of L'Aquila. The event parameters are summarized in Table 3.2.

Parameter	Value
Event ID	IT-2009-0009
Origin time (UTC)	6 April 2009, 01:32:40
Epicenter (latitude)	42.3420°N
Epicenter (longitude)	13.3800°E
Hypocentral depth	8.3 km
Moment magnitude (M_w)	6.1
Focal mechanism	Normal faulting

Table 3.2: Source parameters of the L'Aquila seismic event.

The corresponding dataset comprises 192 three-component records from 64 stations belonging to multiple seismic networks. All stations operate at a uniform sampling rate of 200 Hz ($\Delta t = 0.005$ s). Unlike the Maritime Alps dataset, a subset of stations employed a triggered acquisition mode, meaning that recording began only after an amplitude threshold was exceeded. As a consequence, the pre-event noise window is absent for these stations and the recordings start after the P-wave or even S-wave arrival, constraining the set of temporal windows that can be defined for the moment scaling analysis (Section 4.2.6).

Each station provides three orthogonal components of ground motion. The majority of stations employ high broadband accelerometers (HN sensor type), while station AQU uses high-gain velocimeter channels (HL sensor type). One station (GSG) was excluded from the analysis due to a missing acquisition start timestamp, and one station (TTS) was excluded due to an anomalous origin time offset incompatible with the segmentation pipeline, leaving 62 stations and 186 three-component records available for analysis.

The stations are distributed over epicentral distances ranging from approximately 1.8 km to 421 km, a considerably wider range than the Maritime Alps dataset, encompassing near-source, regional, and teleseismic propagation regimes. Figure 3.2 shows the geographical distribution of the recording stations across central Italy, with the epicenter marked by a red star.




Figure 3.2: Geographical distribution of the accelerometric stations and epicenter for the L'Aquila dataset. Stations are coloured by PGA (cm/s^2).

3.2 Data acquisition and ITACA data preprocessing

3.2.1 ITACA database

Strong-motion waveforms were retrieved from the Italian Accelerometric Archive (ITACA, <https://itaca.mi.ingv.it>), a comprehensive repository of strong-motion records from Italian seismic networks maintained by INGV. ITACA provides uniformly processed three-component time series along with comprehensive metadata including station coordinates, site classification according to Eurocode 8, instrument characteristics, and event parameters. For this study, baseline-corrected acceleration, velocity, and displacement time series were downloaded from the ESM processing pipeline.

All records are stored in ASCII format (`.asc`) with a standardised 64-line header containing metadata in the form `FEATURE: VALUE`, followed by the numerical data in two-column format (sample index, signal value). The metadata fields are grouped into three categories: event information (origin time, hypocenter, magnitude, focal mechanism), station characteristics (network code, location, elevation, site class, epicentral distance, back-azimuth), and recording parameters (sensor type, sampling rate, duration, peak ground motion values, filter corner frequencies). A complete list of metadata fields is provided in Appendix B.

3.2.2 ITACA data processing

The accelerometric data provided by the ITACA database had already undergone several preprocessing steps at the time of acquisition and archiving. Understanding these operations is essential to contextualize the subsequent analytical processing applied in this work.

All records underwent baseline correction to remove any constant offset in the recorded acceleration time series. This operation ensures that the mean value of the signal is zero, which is a prerequisite for subsequent integration to velocity and displacement.

A second-order filter was applied to each recording, with corner frequencies varying by station depending on site conditions and noise characteristics. The low-cut frequency ranged from 0.1 Hz to 0.5 Hz, with the majority of stations filtered at 0.5 Hz (33 files) or 0.2 Hz (15 files). The high-cut frequency was set to 40 Hz for most recordings (57 files) and to 20 Hz for nine files. This bandpass filtering removes both low-frequency drift and high-frequency instrumental noise, retaining the seismologically relevant frequency content of ground motion for moderate-magnitude crustal earthquakes.

No further normalisation or amplitude scaling was applied by ITACA; acceleration values are provided in physical units (cm/s^2), preserving the relative amplitude differences across stations that reflect variations in epicentral distance, site effects, and source radiation patterns.

3.3 Additional preprocessing and quality control

3.3.1 Metadata

The raw metadata were preprocessed through the following steps applied sequentially:

1. **Missing value replacement:** empty strings and 'None' strings replaced with NaN.
2. **Column dropping:** columns that are constant across all files, empty, or irrelevant for the analysis were removed (e.g. `USER1–USER5`, `DATA_LICENSE`).
3. **Type conversion:** numerical columns converted to `float64/int64`; date columns converted to `date time`.
4. **String normalisation:** leading/trailing white space removed from string columns.
5. **Duplicate removal:** duplicate rows dropped.

3.3.2 Signal preprocessing pipelines

Two separate preprocessing pipelines are applied to the raw acceleration signals, depending on the type of statistical analysis to be performed.

Pipeline A: PDF and heavy-tail assessment

The first pipeline is used for the PDF and heavy-tail assessment and consists of two steps:

1. **Baseline correction** — the mean of each signal is subtracted to ensure zero baseline:

$$a'(t) = a(t) - \bar{a} \quad (3.1)$$

2. **Normalisation** — the baseline-corrected signal is divided by its standard deviation $\sigma_{a'}$:

$$\hat{a}(t) = \frac{a'(t)}{\sigma_{a'}} \quad (3.2)$$

Post-preprocessing checks confirm zero residual means, unit standard deviations, and no NaN or infinite values.

Pipeline B: Scaling analysis

The second pipeline is used for the moment scaling analysis, and consists of a single step:

1. **Baseline correction** — the mean of each signal is subtracted to ensure zero baseline:

$$a'(t) = a(t) - \bar{a} \quad (3.3)$$

Post-preprocessing checks confirm zero residual means and no NaN or infinite values.

Chapter 4

Phase detection

4.1 Seismic phases and phase picking

4.1.1 Seismic wave phases

When an earthquake occurs, the released energy propagates through the Earth as elastic waves. The main body wave types are P-waves (primary or compressional waves) and S-waves (secondary or shear waves), which differ in propagation velocity and particle motion (Aki and Richards, 2002).

P-waves are longitudinal waves that travel by alternately compressing and expanding the medium, with particle motion parallel to the propagation direction. They are the fastest seismic waves and therefore arrive first at seismograph stations.

S-waves are transverse waves with particle motion perpendicular to the propagation direction. They travel more slowly and cannot propagate through fluids. Despite lower velocity, S-waves typically have larger amplitudes and cause more ground shaking.

Following the direct P and S arrivals, the seismogram exhibits a prolonged oscillatory phase called the coda, consisting of scattered waves generated by heterogeneities in the Earth's crust. The coda amplitude decays gradually with time and provides information about crustal heterogeneity and attenuation (Aki and Chouet, 1975).

The segmentation of seismic recordings into distinct temporal windows is motivated by the expectation that each phase carries qualitatively different statistical properties:

- **Pre-event:** contains only instrumental and ambient noise prior to the earthquake arrival. Normal diffusion ($\zeta(q) \approx q/2$) is expected, with possible deviations depending on the noise character. This window serves as the statistical baseline for comparison.
- **P-wave:** the first seismic phase, characterised by relatively low energy and compressional particle motion. Deviations from normal diffusion are expected to be moderate.
- **S-wave:** the most energetic phase, where ground motion reaches its maximum amplitude (PGA, PGV, PGD all fall within this window). The strongest deviations from $\zeta(q) = q/2$ are expected here, reflecting the impulsive and non-Gaussian nature of the shear wave arrival.
- **Coda:** characterised by exponential amplitude decay as scattered energy dissipates. The scaling exponents $\zeta(q)$ are expected to progressively approach the pre-event

behaviour as the signal returns towards background noise levels.

- **Post-event:** the return to ambient noise following coda dissipation. Behaviour similar to the pre-event window is expected ($\zeta(q) \approx q/2$), and the comparison between pre-event and post-event provides a consistency check on the stationarity of the background noise.

The systematic comparison across all windows allows identification of which seismic phase departs most strongly from normal diffusion, providing insight into which part of the signal carries information about the out-of-equilibrium nature of the seismic source and propagation process.

4.1.2 Phase picking: from manual to deep-learning approaches

Phase picking, the identification and precise timing of seismic phase arrivals (P-wave, S-wave, and potentially coda onset) on recorded waveforms, is a fundamental task in seismology. Accurate phase picks are essential for earthquake location, magnitude estimation, source characterisation, and seismic tomography. The quality of these picks directly impacts all downstream analyses, including the moment scaling investigation presented in this work.

Traditional phase picking methods

Historically, phase picking has been performed manually by trained seismologists who visually inspect waveforms and identify the onset of each phase based on changes in amplitude and frequency content. While manual picking can achieve high accuracy for high-quality recordings, it is time-consuming, subjective (introducing analyst-dependent variability), and impractical for processing large volumes of data.

To address these limitations, numerous automatic picking algorithms have been developed over the past decades. The most widely adopted approaches are based on characteristic functions that amplify abrupt changes in waveform properties:

- **STA/LTA** (Short-Term Average / Long-Term Average): computes the ratio between the average amplitude over a short time window and a longer background window, triggering a detection when the ratio exceeds a threshold.
- **AR-AIC** (Autoregressive - Akaike Information Criterion): models the waveform as an autoregressive process and identifies the onset as the point where the AIC function reaches a minimum, corresponding to a change in statistical properties.
- **Higher-order statistics:** methods based on kurtosis, skewness, or other statistical moments to detect non-Gaussian transients associated with phase arrivals.

While these methods have been successfully applied in routine seismic monitoring for decades, they suffer from several limitations: strong dependence on parameter tuning, sensitivity to noise levels, and computational inefficiency when applied to large datasets or long time windows.

Deep-learning phase pickers

Recent advances in machine learning, particularly deep learning, have revolutionised phase picking by enabling fully data-driven approaches that learn optimal feature representations directly from labelled waveform examples. Deep-learning pickers typically employ convolutional neural networks (CNNs) to process raw three-component seismograms and output continuous probability time series for P-wave, S-wave, and noise classes. The phase arrival times are then extracted as the peaks of the respective probability distributions.

Compared to traditional methods, deep-learning pickers offer several advantages:

- **Minimal preprocessing:** no manual feature engineering or complex parameter tuning required
- **Robustness to noise:** trained on diverse datasets including noisy recordings, enabling reliable detection in challenging conditions
- **High computational efficiency:** once trained, models can process large volumes of data rapidly
- **Generalisation capability:** models trained on one region often transfer reasonably well to other tectonic settings

Importantly, recent comparative studies have demonstrated that deep-learning-based pickers produce more consistent picks compared to traditional manual picking procedures (Cianetti et al., 2026). This consistency stems from the fact that deep-learning algorithms apply the same learned criteria uniformly across all data, whereas human analysts inevitably introduce subjective variability. For high-precision applications such as moment scaling analysis, where systematic biases in phase identification can propagate into final results, this consistency represents a critical advantage.

Among the available deep-learning pickers, PhaseNet (Zhu and Beroza, 2019) has emerged as one of the most widely adopted architectures due to its robust performance across different seismological contexts, open-source availability, and strong generalisation capabilities. PhaseNet has been trained on multiple regional datasets, including INSTANCE (Michellini et al., 2021), a comprehensive collection of Italian seismicity data comprising approximately 54,000 earthquakes and 1.2 million three-component waveform traces with manually reviewed phase picks from the INGV bulletin.

In the present work, two complementary phase picking approaches were implemented and compared: the traditional AR-AIC method and the PhaseNet deep-learning picker. The AR-AIC implementation follows the workflow described in Section ??, while PhaseNet was employed through the SeisBench framework 4.3 using the model pre-trained on the INSTANCE dataset. The choice of this pre-trained model is motivated by the fact that the signals analysed in this work are themselves drawn from the INSTANCE database, making it the natural reference model. The two pipelines are applied independently to the same data presented in Chapter 3, and the resulting phase picks are subsequently used as the basis for signal segmentation and moment scaling analysis, enabling a direct comparison of the impact of the picking method on the estimated scaling exponents $\zeta(q)$.

4.2 Phase detection using AR-AIC picking

The segmentation pipeline is designed to progressively refine the identification of physically meaningful temporal regions of the signal. It starts from theoretical estimates of primary

and secondary wave arrival times, then applies data-driven corrections, and finally defines windows based on both physical and statistical criteria.

The workflow is structured as follows. First, theoretical P- and S-wave arrival times are estimated based on hypocentral distance and CRUST1.0 crustal velocities (4.2.1). These estimates are then used to define adaptive search windows (4.2.2), within which data-driven AR-AIC method is applied to refine onset times (4.2.3). Coda onset is detected using four independent methods (4.2.4), and the end of the active coda is subsequently determined to identify the transition back to ambient noise (4.2.5). The resulting phase boundaries are used to define analysis windows (4.2.6). Finally, all detections are subjected to quality control based on physical consistency criteria (4.2.7).

4.2.1 Theoretical estimation of P- and S-wave arrival times

Theoretical P- and S-wave arrival times are estimated assuming straight-ray propagation through a layered crustal medium. The estimation requires two elements: average crustal velocities v_P and v_S , which are extracted from a global crustal model, and the hypocentral distance, which accounts for the three-dimensional path from the earthquake source to the station.

Crustal velocity estimation from CRUST1.0

The crustal velocities v_P and v_S are determined for each station from the CRUST1.0 global model (Laske et al., 2013; Leeman et al., 2013), a one-degree resolution crustal model that includes water/ice cover, sedimentary layers (upper, middle, lower), crystalline crust (upper, middle, lower), and the uppermost mantle. Each layer k is characterised by P-wave velocity $v_{P,k}$, S-wave velocity $v_{S,k}$, density ρ_k , thickness h_k , and depth to the top of the layer.

Rather than averaging all crustal layers indiscriminately, only the layers actually traversed by seismic waves travelling from the hypocenter to the surface are included in the velocity calculation. For a vertical ray path from hypocenter depth h to the surface, the set of traversed layers is determined by comparing h with the depth range of each layer. If the hypocenter lies within a layer, only the portion of that layer above the hypocenter is included.

The average crustal velocities are then computed as thickness-weighted means over the traversed layers:

$$v_P = \frac{\sum_{k \in \mathcal{T}} v_{P,k} h_k^{\text{eff}}}{\sum_{k \in \mathcal{T}} h_k^{\text{eff}}}, \quad v_S = \frac{\sum_{k \in \mathcal{T}} v_{S,k} h_k^{\text{eff}}}{\sum_{k \in \mathcal{T}} h_k^{\text{eff}}} \quad (4.1)$$

where $\mathcal{T}$ denotes the set of traversed layers and h_k^{eff} is the effective thickness (full layer thickness if entirely traversed, or partial thickness if the hypocenter lies within the layer). This weighting ensures that thicker layers contribute proportionally more to the average velocity.

Hypocentral distance

The hypocentral distance d_{hypo} is defined as the three-dimensional distance from the hypocenter to the station:

$$d_{\text{hypo}} = \sqrt{d_{\text{epi}}^2 + h^2} \quad (4.2)$$

where d_{epi} is the epicentral distance (horizontal distance between station and epicenter, available from station metadata) and h is the hypocenter depth.




Figure 4.1: Schematic illustration of epicentral and hypocentral distances. The epicentral distance is the horizontal distance between the station and the epicenter, while the hypocentral distance accounts for the three-dimensional path from the hypocenter (earthquake focus) to the station.

The use of hypocentral rather than epicentral distance accounts for the actual path length travelled by seismic waves from the source to the surface.

Theoretical arrival times

Given the station-specific crustal velocities and hypocentral distances, the theoretical P- and S-wave arrival times are computed as:

$$t_P^{\text{theo}} = t_0 + \frac{d_{\text{hypo}}}{v_P}, \quad t_S^{\text{theo}} = t_0 + \frac{d_{\text{hypo}}}{v_S} \quad (4.3)$$

where t_0 is the origin time, defined as the event occurrence time relative to the start of the recording. The resulting station-specific theoretical arrival times serve as the centres of the adaptive search windows described in Section 4.2.2.

4.2.2 Adaptive search windows definition

The theoretical arrival times are used to define adaptive search windows for automatic phase picking. For the P-wave:

$$W'_P = [t_P^{\text{theo}} - \Delta t_P, t_P^{\text{theo}} + \Delta t_P] \quad (4.4)$$

and for the S-wave:

$$W'_S = [t_S^{\text{theo}} - \Delta t_S, t_S^{\text{theo}} + \Delta t_S] \quad (4.5)$$

The window half-widths Δt_P and Δt_S are adapted to the hypocentral distance to account for distance-dependent uncertainties in the theoretical arrival times. These uncertainties arise from several sources:

- CRUST1.0 velocity model uncertainties
- Event location uncertainties

- Lateral crustal heterogeneities not captured by the one-degree resolution model
- The assumption of a one-dimensional layered structure, which neglects three-dimensional velocity variations

To balance precision and robustness, the window half-widths are determined empirically from the distribution of hypocentral distances in the dataset. The stations are divided into distance bins based on the tertiles of the hypocentral distance distribution, yielding three bins with approximately equal station counts.

Unlike P-wave arrivals, which typically exhibit an impulsive onset with a sharp amplitude increase, S-wave arrivals are often emergent, meaning the amplitude grows gradually and the precise onset time is harder to identify automatically. The S-wave windows are therefore defined to be twice as wide as the P-wave windows:

$$\Delta t_S = 2 \Delta t_P \quad (4.6)$$

To prevent overlap between P and S search regions and reduce the risk of phase misidentification, the S-wave window is constrained to start after the end of the P-wave window:

$$t_{S,\text{start}} = \max(t_S^{\text{theo}} - \Delta t_S, t_{P,\text{end}} + \delta) \quad (4.7)$$

where $\delta = 0.5$ s is a minimum temporal separation gap.

4.2.3 Windowed P- and S-phase picking using AR-AIC

For each station, the arrival times are refined through an automatic picking procedure based on the `ar_pick` algorithm implemented in ObsPy (Beyreuther et al., 2010). The search is restricted to the search windows defined in Section 4.2.2, in order to reduce computational cost and minimise the risk of false detections.

The `ar_pick` method relies on autoregressive (AR) modelling of the seismic signal combined with the Akaike Information Criterion (AIC) to identify abrupt changes in signal statistics, which are associated with the onset of seismic phases. The algorithm operates on three-component waveforms (vertical, north–south, and east–west).

For each station, the workflow is as follows:

1. The three-component signals (Z, N, E) are extracted and checked for completeness. Stations with missing components are flagged and excluded from subsequent segmentation and scaling analysis.
2. The waveform is subset within the predefined P- and S-wave search windows.
3. The `ar_pick` algorithm is applied independently to each window: within the P-search window to detect the P-wave onset, and within the S-search window to detect the S-wave onset.
4. The detected onset times, initially expressed as sample indices relative to the beginning of each window, are converted to absolute sample indices within the full waveform.

All detected onset times are stored in dual representation: as integer sample indices (primary representation) and as times in seconds (derived as $t = n/f_s$, where n is the sample index and $f_s = 200$ Hz is the sampling rate). This approach minimises rounding

errors in subsequent computations by performing all internal calculations in the sample domain and converting to physical time units only at the output stage.

The `ar_pick` algorithm is applied with the following fixed parameters, reported here for reproducibility:

Parameter	Value	Description
<code>lta_p</code>	1.0 s	Long-term average window, P-wave
<code>sta_p</code>	0.1 s	Short-term average window, P-wave
<code>lta_s</code>	4.0 s	Long-term average window, S-wave
<code>sta_s</code>	1.0 s	Short-term average window, S-wave
<code>m_p</code>	2	AR model order, P-wave
<code>m_s</code>	8	AR model order, S-wave
<code>l_p</code>	0.1 s	Smoothing window length, P-wave
<code>l_s</code>	0.2 s	Smoothing window length, S-wave

Table 4.1: Parameters used for the `ar_pick` algorithm.

The higher AR model order for S-waves ($m_s = 8$ vs $m_p = 2$) reflects the more complex waveform structure of S-wave arrivals, which typically exhibit a more gradual onset and richer frequency content compared to P-waves.

For each successful detection, the residual between observed and theoretical arrival times is computed as:

$$r_P = t_P^{\text{obs}} - t_P^{\text{theo}}, \quad r_S = t_S^{\text{obs}} - t_S^{\text{theo}} \quad (4.8)$$

These residuals provide a quantitative measure of the consistency between the velocity model and the observed data.

4.2.4 Coda onset detection

The coda onset marks the transition from coherent S-wave arrivals to scattered wave energy. Three independent methods were implemented: Rautian and Khalturin (1978), Arias intensity-based detection, and amplitude envelope decay, plus an ensemble estimate based on their median. All four coda onset estimates ($t_{\text{coda}}^{\text{Rautian}}$, $t_{\text{coda}}^{\text{Arias}}$, $t_{\text{coda}}^{\text{Envelope}}$, $t_{\text{coda}}^{\text{median}}$) were computed for each of the three-component recordings.

Rautian method

The empirical formula proposed by Rautian and Khalturin (1978) defines coda onset as:

$$t_{\text{coda}} = t_{\text{origin}} + 2(t_S - t_0) \quad (4.9)$$

where t_{origin} is the earthquake origin time and t_S is the detected S-wave arrival time. The factor of 2 reflects the empirical observation that coda lapse time is approximately twice the S-wave travel time. This method is computationally simple and widely adopted in earthquake seismology. However, it is purely distance-based and does not account for signal amplitude or decay characteristics, potentially leading to systematic overestimation of S-wave duration for near-field stations and underestimation for far-field stations with extended wave trains.

Arias intensity method

The Arias intensity (Arias, 1970) quantifies cumulative seismic energy as:

$$\text{AI}(t) = \frac{\pi}{2g} \int_0^t x^2(\tau) d\tau \quad (4.10)$$

where $x(t)$ is the signal time series and $g = 9.81 \text{ m/s}^2$. Following Trifunac and Brady (1975), we compute AI from 5 s before P-onset to the end of the recording. Coda onset is defined as the time when 95% of total energy has been released (D5–D95 interval end):

$$t_{\text{coda}} = \arg \min_t \left\{ \frac{\text{AI}(t)}{\text{AI}_{\text{total}}} \geq 0.95 \right\} \quad (4.11)$$

Envelope decay method

The envelope method identifies coda onset by tracking the decay of instantaneous amplitude (Filippucci et al., 2021). The analytical signal envelope is computed via the Hilbert transform:

$$E(t) = |a(t) + i\mathcal{H}[a(t)]| = \sqrt{a(t)^2 + \mathcal{H}[a(t)]^2} \quad (4.12)$$

where $\mathcal{H}$ denotes the Hilbert transform. To suppress high-frequency noise, the envelope is smoothed with a 1-second moving average. The peak envelope E_{peak} is identified within a 5-second window following S-onset, and coda onset is defined as the first time the smoothed envelope falls below a threshold:

$$t_{\text{coda}} = \arg \min_t \{E_{\text{smooth}}(t) < \alpha E_{\text{peak}}\}, \quad t > t_S \quad (4.13)$$

where $\alpha = 0.3$ is an empirically chosen threshold selected to ensure that the coda onset is identified after the main S-wave energy has decayed substantially, while remaining above the ambient noise level. This approach physically captures the transition from S-wave energy to diffuse scattering, adapting automatically to both near-field impulsive arrivals and far-field dispersed wave trains.

Median ensemble method

To reduce sensitivity to the specific assumptions of each individual method, a robust ensemble estimate is computed as the median of the three independent coda onset times:

$$t_{\text{coda}}^{\text{median}} = \text{median} \left(t_{\text{coda}}^{\text{Rautian}}, t_{\text{coda}}^{\text{Arias}}, t_{\text{coda}}^{\text{Envelope}} \right) \quad (4.14)$$

The median is preferred over the mean as it is robust to outliers: if one method produces an anomalous estimate due to low SNR, a complex waveform, or a near-field impulsive arrival, the median is unaffected as long as the other two methods agree. This property makes the ensemble estimate particularly suitable for heterogeneous datasets where individual methods may fail on a subset of stations.

Physical bounds

Regardless of the detection method, all coda onset estimates are subject to the following physical constraints:

$$t_S + 1 \text{ s} \leq t_{\text{coda}} \leq t_{\text{signal_end}} - 1 \text{ s} \quad (4.15)$$

The lower bound ensures a minimum S-wave window duration of at least 1 s, preventing degenerate cases where the coda onset is detected too close to the S-wave arrival. The upper bound ensures that the coda window contains at least 1 s of signal before the end of the recording, avoiding edge effects in subsequent analysis. If a method produces an estimate outside these bounds, the estimate is clipped to the nearest valid value.

4.2.5 Coda end detection

The coda end marks the transition from active coda to ambient noise. Identification of this boundary enables the definition of a fifth temporal window (post-event) capturing the return to background noise conditions, which serves as an additional reference for statistical comparison with the pre-event window.

Two methods were implemented for coda end detection, both operating on the signal after the coda onset time t_{coda} .

Envelope decay method

The envelope decay method identifies the coda end as the point where the signal amplitude drops below a threshold relative to the coda peak and remains low for a sustained period. The instantaneous amplitude envelope is computed via the Hilbert transform (Equation 4.12) and smoothed with a 0.5-second moving average. The peak envelope E_{peak} is identified within the first 5 s after the coda onset, and the coda end is defined as the first time the smoothed envelope satisfies:

$$E_{\text{smooth}}(t') < \beta E_{\text{peak}} \quad \text{for all } t' \in [t, t + T_{\text{stab}}] \quad (4.16)$$

where $\beta = 0.10$ (10% of coda peak amplitude) and $T_{\text{stab}} = 2$ s is a stability duration requiring the amplitude to remain below threshold for a sustained period. The stability requirement prevents false detections caused by transient noise spikes, which could otherwise prematurely terminate the coda window.

Arias D99.5 method

The Arias D99.5 method extends the energy-based framework used for coda onset detection (Section 4.2.4). While the coda onset is defined as D95 (the time at which 95% of total seismic energy has been released), the coda end is defined as D99.5:

$$t_{\text{coda_end}} = \arg \min_t \left\{ \frac{\text{AI}(t)}{\text{AI}_{\text{total}}} \geq 0.995 \right\} \quad (4.17)$$

The energy band [D95, D99.5] thus represents the active coda decay phase, during which the remaining 4.5% of total seismic energy is released. This formulation is physically motivated and internally consistent with the coda onset definition, as both boundaries are derived from the same Arias intensity curve computed on the same signal window. Compared to the envelope decay method, the Arias D99.5 approach is robust to amplitude fluctuations and requires no stability duration parameter, at the cost of being insensitive to the instantaneous signal morphology.

Physical bounds and five-window segmentation

Both methods are subject to the following physical constraints:

$$t_{\text{coda}} + 1 \text{ s} \leq t_{\text{coda_end}} \leq t$$

The lower bound ensures a minimum active coda duration of at least 1 s. If neither method produces a valid estimate before the end of the recording, $t_{\text{coda_end}}$ is set to the signal end and no post-event window is created, reducing the segmentation to four windows. When a valid $t_{\text{coda_end}}$ is detected, a fifth post-event window $[t_{\text{coda_end}}, t_{\text{signal_end}}]$ is created alongside the standard four windows, as described in Section 4.2.6.

In addition to the bounds above, a minimum window duration of $\Delta t_{\text{min}} = 2 \text{ s}$ is enforced for all temporal windows. Any window whose duration falls below this threshold is discarded and excluded from subsequent scaling analysis. This criterion removes windows produced by anomalous picks, such as S-wave onsets detected very close to the coda onset, or coda onsets coinciding with the end of the recording, and ensures that the moment scaling analysis is performed only on windows that contain a sufficient number of samples to yield statistically meaningful estimates of the structure function.

4.2.6 Definition of analysis time windows

The detected phase boundaries are used to partition each seismic recording into up to five non-overlapping temporal windows. The number of windows actually created for each recording depends on two factors: whether a valid coda end time $t_{\text{coda_end}}$ was identified (Section 4.2.5), and whether the recording starts before or after certain phase arrivals.

For standard recordings that begin before the P-wave arrival, all five windows can in principle be created:

- **Pre-event:**

$$W_{\text{pre}} = [\max(0, t_P - \Delta t_{\text{pre}}), t_P] \quad (4.19)$$

where $\Delta t_{\text{pre}} = 5 \text{ s}$ by default.

- **P-wave:**

$$W_P = [t_P, t_S] \quad (4.20)$$

Contains the compressional wave arrival. Duration varies across stations depending on the S-P time interval.

- **S-wave:**

$$W_S = [t_S, t_{\text{coda}}] \quad (4.21)$$

Contains the shear wave arrival, typically the most energetic phase of the recording (PGA, PGV, and PGD all fall within this window).

- **Coda:**

$$W_{\text{coda}} = [t_{\text{coda}}, t_{\text{coda_end}}] \quad (4.22)$$

Contains scattered wave energy decaying back towards ambient noise. If no valid $t_{\text{coda_end}}$ is detected, this window extends to the end of the recording.

- **Post-event:**

$$W_{\text{post}} = [t_{\text{coda_end}}, t_{\text{signal_end}}] \quad (4.23)$$

Created only when a valid $t_{\text{coda_end}}$ is detected before the end of the recording. Captures the return to ambient noise following coda dissipation.

However, for a subset of stations the recording begins after the theoretical P-wave arrival, or even after the S-wave arrival. This occurs for late triggered stations, whose acquisition is activated by an amplitude threshold rather than running continuously. In such cases, the available windows are restricted based on the sign of the theoretical arrival time t_P^{theo} relative to the recording start ($t = 0$):

- If $t_P^{\text{theo}} \geq 0$: the recording begins before the P-wave arrival, and all windows (pre-event, P-wave, S-wave, coda, post-event) can be created.
- If $t_P^{\text{theo}} < 0$ and $t_S^{\text{theo}} \geq 0$: the recording begins after the P-wave arrival but before the S-wave arrival. Only the S-wave, coda, and post-event windows are created; pre-event and P-wave windows are not available.
- If both $t_P^{\text{theo}} < 0$ and $t_S^{\text{theo}} < 0$: the recording begins after both P- and S-wave arrivals. Only the coda and post-event windows are created.

All window boundaries are stored in dual representation (sample indices and seconds) for consistency with the rest of the pipeline. Windows whose duration falls below the minimum threshold $\Delta t_{\text{min}} = 2$ s (Section 4.2.5) are additionally discarded. The complete set of windowed signals is then subjected to the quality control procedure described in the following section.

4.2.7 Quality control and validation

All automatically detected picks are subjected to a multi-level quality control procedure aimed at assessing their physical consistency and reliability. Three independent criteria are applied, operating at different levels of granularity: station-level monotonicity, component-level signal-to-noise ratio, and component-level peak ground motion timing.

Monotonicity with hypocentral distance

P- and S-wave arrival times are expected to increase monotonically with hypocentral distance, reflecting the finite propagation velocity of seismic waves. For each station i sorted by hypocentral distance, the following condition must hold:

$$t^{(i-1)} < t^{(i)} < t^{(i+1)} \quad (4.24)$$

where $t^{(i-1)}$ and $t^{(i+1)}$ denote the arrival times at the nearest and farthest neighbouring stations in the distance-ordered list, respectively. This check is applied independently to P- and S-wave picks, and is referred to as the *MonoP check* and *MonoS check*, respectively. Violations of this condition are flagged for inspection but do not result in the exclusion of the corresponding picks; their interpretation is discussed in Chapter 7.

Since P- and S-wave arrival times are station properties (identical across the three components of the same station), the monotonicity check is computed once per station and the result is shared across all three components.

Signal-to-noise ratio

A minimum signal-to-noise ratio (SNR) criterion is applied independently to each component. The SNR is defined as:

$$\text{SNR} = \frac{\text{RMS}_{\text{signal}}}{\text{RMS}_{\text{noise}}} \quad (4.25)$$

where $\text{RMS}_{\text{noise}}$ is estimated from the entire pre-event window, and $\text{RMS}_{\text{signal}}$ is computed over the first 5 s following the detected onset of the phase of interest (P- or S-wave). Picks with $\text{SNR} < 3$ are considered unreliable and are discarded. This criterion is referred to as the *SNR P check* and *SNR S check* for P- and S-wave picks, respectively.

Peak ground motion timing

A further validation criterion (*Peak check*) exploits the timing of peak ground motion, defined as the time of maximum absolute amplitude within the recording. For acceleration signals this corresponds to the PGA time (`TIME_PGA_S`), for velocity signals to the peak ground velocity time (`TIME_PGV_S`), and for displacement signals to the PGD time (`TIME_PGD_S`), all available from the ITACA metadata. The peak ground motion is physically expected to occur within the S-wave window $[t_S, t_{\text{coda}})$, where seismic energy is typically highest. Picks for which the peak ground motion falls outside this window are flagged as physically inconsistent.

This check is applied independently per component, as peak ground motion amplitudes and timing can differ across the three recording directions.

Combined quality control outcome

The three criteria are combined into a single binary flag `all_passed` per station-component pair, defined as the logical AND of all individual checks. Components failing one or more checks are flagged for manual inspection. The results of the quality control procedure are reported in Section 6.3, where the detection rates and failure modes for both AR-AIC and PhaseNet pipelines are discussed and compared.

4.3 Phase picking using PhaseNet

4.3.1 PhaseNet

PhaseNet is a deep-neural-network-based seismic arrival-time picking method developed by Zhu and Beroza (2019) to address the limitations of traditional phase detection algorithms. Unlike conventional methods that rely on manually defined features (e.g., amplitude thresholds in STA/LTA or autoregressive modelling in AR-AIC), PhaseNet employs deep learning to automatically extract discriminative features from raw waveforms. The network takes three-component seismic recordings as input and outputs probability distributions for P-wave arrivals, S-wave arrivals, and noise at each time sample.

The architecture of PhaseNet (Figure 4.2) is adapted from U-Net (Ronneberger, Fischer, and Brox, 2015), a convolutional neural network originally designed for biomedical image segmentation. The key insight underlying U-Net is its ability to localise spatial features in images through a symmetric encoder-decoder structure. PhaseNet applies this principle to one-dimensional time series: the task becomes localising temporal properties (P-wave onset, S-wave onset, or noise) along the seismogram.




Figure 4.2: PhaseNet architecture. The network consists of four down-sampling stages (left) and four up-sampling stages (right), with skip connections (grey arrows) directly concatenating encoder outputs to corresponding decoder layers. Numbers indicate feature map dimensions (channels $\times$ temporal length). Blue blocks represent convolutional layers followed by ReLU activation; orange blocks indicate down-sampling operations (convolution with stride 4); green blocks represent transposed convolutions for up-sampling. Adapted from Zhu and Beroza (2019).

The network consists of a contracting path (encoder) and an expanding path (decoder), each comprising four stages. In the encoder, the input waveform undergoes successive convolutions with rectified linear unit (ReLU) activations followed by down-sampling operations with stride 4. This progressively reduces the temporal resolution while increasing the number of feature channels (8, 11, 16, 22, 32 in the encoder), reaching a bottleneck representation of 32×12 before the decoder reconstructs the original temporal resolution (44, 32, 22, 16 channels), producing a final output of 3×3001 corresponding to the three probability distributions.

A distinctive feature of the U-Net architecture is the use of skip connections: at each depth level, feature maps from the encoder are directly concatenated to the corresponding decoder layer. This design preserves fine-scale temporal details that might otherwise be lost during down-sampling, improving both convergence during training and localisation accuracy during inference. Convolution kernels have a fixed size of 7 samples throughout the network.

PhaseNet requires minimal preprocessing. Waveforms are resampled to a uniform sampling rate of 100 Hz and normalised component-wise by subtracting the mean and dividing by the standard deviation. No band-pass filtering is applied, allowing the network to learn optimal frequency-dependent features directly from the data.

The training procedure relies on a probabilistic representation of arrival times. Manually picked P- and S-wave onsets are converted into target probability distributions using truncated Gaussian functions with zero mean and standard deviation $\sigma = 0.1$ s:

$$p_{\text{phase}}(t) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(t - t_{\text{onset}})^2}{2\sigma^2}\right), \quad |t - t_{\text{onset}}| \leq 0.5 \text{ s} \quad (4.26)$$

where t_{onset} is the manually picked arrival time. This representation accounts for inherent uncertainty and potential bias in manual picks, while simultaneously increasing the amount of labelled information per event relative to a point-wise label. The softmax function normalises the network's output into three mutually exclusive probability distributions:

$$q_i(t) = \frac{e^{z_i(t)}}{\sum_{k=1}^3 e^{z_k(t)}} \quad (4.27)$$

where $i \in \{1, 2, 3\}$ corresponds to noise, P-wave, and S-wave classes, and $z_i(t)$ are the raw activations of the final layer. The network is trained by minimising the cross-entropy loss between the predicted distribution $q(t)$ and the target distribution $p(t)$:

$$\mathcal{L} = - \sum_{i=1}^3 \sum_t p_i(t) \log q_i(t) \quad (4.28)$$

During inference, P- and S-wave arrival times are extracted as the temporal locations of the peaks in the respective probability distributions (Figure 4.3). This approach generates continuous probability curves rather than discrete binary decisions, providing both an arrival-time estimate and an associated confidence measure at each time sample.




Figure 4.3: Example of PhaseNet training dataset. The top three panels show seismograms of the ENZ (East, North, Vertical) components, with blue and red vertical lines indicating manually picked P and S arrival times. The bottom panel shows the converted probability distributions for P and S pickers, represented as truncated Gaussian distributions with mean ($\mu = 0$ s) and standard deviation ($\sigma = 0.1$ s). Adapted from Zhu and Beroza (2019).

In the present work, PhaseNet is applied through the SeisBench framework (Woollam et al., 2022) using the pre-trained model described in the following section.

4.3.2 INSTANCE dataset

The INSTANCE dataset (Michellini et al., 2021) gathers seismic waveform data from weak and strong-motion stations that have been extracted from the Italian EIDA node. The dataset consists of approximately 1.2 million three-component waveform traces from about 54,000 earthquakes. The metadata associated with the waveforms are extracted from the INGV earthquake catalogue and from the waveform traces themselves.

The data selection criterion was the most accurately picked P- and S-wave onset phase published in the INGV bulletin. The manually reviewed locations are indicated as preferred solutions in the INGV bulletin. The authors selected only those stations that had P- and, if available, S-wave onset picks associated with the preferred location of the INGV bulletin.

The authors managed the issue of imbalanced representation of earthquakes of different magnitude by choosing to select in the target area the great majority of earthquakes with magnitude $M \geq 4.0$, earthquakes with origin times differing by more than 120s in the range $2.0 \leq M < 4.0$, and an additional 20,000 earthquakes, randomly selected, with origin times differing by more than 120s for $M < 2.0$.

Waveforms were preprocessed by removing traces containing data gaps, trimming the waveform trace to the nearest sample to the start time, 120s trace windowing, removal of mean and linear trends from the data, and resampling at 100 Hz.

The selected 54,008 earthquakes composing the dataset can be considered a representative subset of the seismicity in Italy and for those earthquakes occurring in the near vicinity of the Italian national borders.

The great majority of the traces have been recorded within a 200km distance from their respective epicenter. The majority of the traces belong to shallow crustal earthquakes.

The weak-motion, high-gain channels represent more than 70% of the total number of traces. These are subdivided into HH channels associated with the broadband high-gain velocimeters (approximately 51% of total traces), extended-short-period-channel traces (EH, approximately 20%), and low-gain accelerometric channels.

The INSTANCE dataset includes noise data, important to provide information on the noise characteristics of the individual stations in the absence of earthquake-generated signal. The noise data was obtained by identifying 600s long time windows free of any earthquake, obtaining the operational times of acquisition of each station, and identifying the 120s time windows to be included in the dataset for each station and channel. The resulting total number of noise trace windows is 132,228, corresponding to about 10% of the total number of traces of INSTANCE.




Figure 4.4: Map of the 620 stations used in the INSTANCE dataset. Symbol size is proportional to the number of P-phase arrivals per station. Adapted from Michelini et al. (2021).

PhaseNet has been trained on the INSTANCE dataset and the pre-trained model is publicly available through the SeisBench framework (Woollam et al., 2022), a standardised toolbox for accessing and applying machine learning models to seismological data. The recordings analysed in the present work were extracted directly from the INSTANCE database, making the INSTANCE-trained PhaseNet model the natural choice for automatic phase detection.

4.3.3 Phase picking using PhaseNet pre-trained on INSTANCE dataset

The PhaseNet model pre-trained on the INSTANCE dataset was applied to the signals to automatically detect P-wave and S-wave arrival times. The model was accessed through the SeisBench framework (Woollam et al., 2022), which provides a standardised interface for loading and applying pre-trained seismological deep learning models.

Preprocessing and resampling

PhaseNet requires input waveforms sampled at 100 Hz, organised as three-component streams (Zhu and Beroza, 2019). Since the data were originally recorded at 200 Hz, a resampling step was necessary before model application.

For each station, the three orthogonal components were extracted from the preprocessed dataset and assembled into an ObsPy Stream object (Beyreuther et al., 2010). The component identification followed a hierarchical logic to accommodate different channel naming conventions present in the dataset:

- Vertical component: HNZ (high-gain accelerometer) or HGZ (high-gain short-period seismometer)
- Horizontal components: HNN and HNE, or HGN and HGE, or HN1 and HN2

Stations for which one or more components could not be identified were excluded from the PhaseNet analysis and flagged accordingly. The resulting streams were then resampled to 100 Hz using the ObsPy `resample` method, which applies a frequency-domain anti-aliasing filter before decimation to prevent spectral aliasing.

Model application

The pre-trained PhaseNet model was loaded from the SeisBench model repository using the identifier `"instance"`. The model expects input waveforms of at least 30 s duration (3001 samples at 100 Hz). Signals shorter than this threshold were excluded from processing.

Phase detection was performed by passing each resampled three-component stream to the model's `classify` method, which returns a `ClassifyOutput` object containing a list of discrete phase picks. Each pick includes the peak onset time as a `UTCDateTime` object, the associated phase label (P or S), and the detection probability at the pick location. When multiple picks of the same phase were returned for a single station, the pick with the highest detection probability was retained.

The onset time for each phase was converted from absolute `UTCDateTime` to a sample index in the original 200 Hz signal by computing the offset in seconds from the stream start time and rounding to the nearest sample:

$$n_{\text{onset}} = \text{round}((t_{\text{pick}} - t_{\text{stream}}^{\text{start}}) \cdot f_{\text{orig}}) \quad (4.29)$$

where t_{pick} is the pick time, $t_{\text{stream}}^{\text{start}}$ is the stream start time, and $f_{\text{orig}} = 200$ Hz is the original sampling rate. This sample-domain representation is used as the primary reference for all subsequent operations, with conversion to seconds performed only at the final output stage, consistent with the samples-first approach adopted throughout the pipeline.

A quality filter was applied to remove picks with physically inconsistent ordering ($t_S \leq t_P$). Since the `classify` method does not apply any probability threshold by

default, an optional minimum probability threshold configuration was additionally tested for the INT-41004391 dataset.

Since the PhaseNet pipeline provides exclusively P-wave and S-wave onset times, coda onset detection, coda end detection, signal segmentation into temporal windows, and the subsequent quality control checks (Section 4.2.7) are applied identically to both the AR-AIC and PhaseNet picks, following the procedure described in Sections 4.2.4–4.2.6. This ensures full methodological consistency between the two pipelines and enables a direct comparison of the impact of the phase picking approach on the estimated scaling exponents $\zeta(q)$.

Probability threshold configuration

In addition to the no-threshold baseline, a probability threshold configuration was tested for the INT-41004391 dataset. In this configuration, only picks satisfying $P_P \geq 0.3$ and $P_S \geq 0.3$ were retained. The threshold value of 0.3 corresponds to the default detection threshold of the SeisBench framework (Woollam et al., 2022) and was selected as a standard reference point for comparison with the no-threshold baseline. This configuration was not applied to the IT-2009-0009 dataset, for which the systematically low S-wave detection probabilities render probability-based filtering ineffective, as discussed in Section 6.3.2.

Residual filter configuration

In addition to the no-threshold baseline, a residual filter configuration was tested for the IT-2009-0009 dataset. In this configuration, stations for which the S-wave onset residual exceeded a threshold based on the median absolute deviation (MAD) were discarded as unreliable detections. Specifically, a station was excluded if

$$|r_S - \tilde{r}_S| > 4 \cdot \text{MAD}(r_S) \quad (4.30)$$

where r_S denotes the S-wave residual, $\tilde{r}_S$ its median, and $\text{MAD}(r_S) = \text{median}(|r_S - \tilde{r}_S|)$. The filter was applied exclusively to S-wave residuals, since P-wave residuals obtained from PhaseNet are comparable in magnitude to those from AR-AIC, for which no residual filtering is applied; P-wave residuals are therefore not expected to reflect unreliable picks but rather the systematic bias of the theoretical arrival times computed from CRUST1.0. The multiplier $k = 4$ was chosen as a conservative threshold, consistent with the intrinsically high dispersion of S-wave residuals for this dataset, which is a consequence of the systematically low S-wave detection probabilities of PhaseNet discussed in Section 6.3.2. This configuration was not applied to the INT-41004391 dataset, for which S-wave residuals are sufficiently well-behaved under the no-threshold and probability threshold configurations.

Chapter 5

Statistical analysis methods

This chapter describes the statistical methods employed to characterise the distributional and scaling properties of the seismic signals. The analysis is structured in three progressive stages. First, empirical probability density functions are computed and tested for departures from Gaussianity (Section 5.1). Second, the nature and severity of heavy-tail behaviour are quantified through distributional model comparison and power-law exponent estimation (Section 5.2). Third, moment scaling analysis is performed to detect signatures of anomalous or strongly anomalous diffusion in the different seismic phases (Section 5.3). All analyses are performed independently for acceleration, velocity, and displacement signals; the distributional analyses (Sections 5.1–5.2) are applied to the full preprocessed signals, while the moment scaling analysis and sensitivity analysis (Sections 5.3–??) operate on the temporally segmented windows obtained from both AR-AIC and PhaseNet phase picks.

5.1 Empirical probability density functions and tests for non-Gaussianity

Empirical probability density functions were computed for all signals and for all three signal types (acceleration, velocity, and displacement). Each signal was normalised to zero mean and unit standard deviation prior to analysis. Histograms were constructed using logarithmic bin spacing to ensure adequate resolution in the tails of the distribution, where deviations from Gaussianity are most informative.

A Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$ was fitted to each normalised signal using maximum likelihood estimation. Normality was assessed via the Anderson–Darling test at the 5% significance level. The Anderson–Darling test was preferred over the more commonly used Kolmogorov–Smirnov test because it assigns greater weight to discrepancies in the tails of the distribution, making it more sensitive to the heavy-tail behaviour expected in seismic signals.

The shape of each empirical distribution was further characterised by excess kurtosis and skewness:

$$\kappa = \frac{\mu_4}{\sigma^4} - 3, \quad \gamma = \frac{\mu_3}{\sigma^3} \quad (5.1)$$

where $\mu_k = \mathbb{E}[(x - \bar{x})^k]$ denotes the k -th central moment of the signal $x(t)$. For a Gaussian distribution, $\kappa = 0$ and $\gamma = 0$; positive excess kurtosis ($\kappa > 0$) indicates heavier tails than a Gaussian, while non-zero skewness indicates asymmetry. Q-Q plots were generated for each signal to provide a complementary visual assessment of deviations from normality.

Rejection of the Gaussian null hypothesis, together with positive excess kurtosis, provides the first indication of heavy-tail behaviour and motivates the more detailed distributional analysis described in the following section.

5.2 Heavy-tail analysis

The empirical distributions of all three signal types (acceleration, velocity, and displacement) were analysed to assess the presence and characterise the nature of heavy tails. Three alternative distributions were fitted to each normalised signal via Maximum Likelihood Estimation: Laplace distribution with location μ and scale b ($k = 2$ free parameters), Student- t distribution with location μ , scale σ , and degrees of freedom ν ($k = 3$), and Lévy-stable distribution with stability index $\alpha_s \in (0, 2]$, skewness $\beta_s \in [-1, 1]$, scale γ_L , and location δ ($k = 4$). These were compared against the Gaussian baseline $\mathcal{N}(\mu, \sigma^2)$ ($k = 2$) fitted in Section 5.1.

Model selection was performed using the Akaike Information Criterion (AIC):

$$\text{AIC} = 2k - 2 \ln \hat{L} \quad (5.2)$$

where k is the number of free parameters and $\hat{L}$ is the maximised likelihood. A lower AIC indicates a better balance between goodness of fit and model complexity. The model with the lowest AIC among the four candidates was identified as the best-fitting distribution for each signal.

The complementary cumulative distribution function (Complementary Cumulative Distribution Function (CCDF)) $P(|x| > u)$ was plotted on a log-log scale for each signal. If the tail follows a power-law distribution $P(|x| > u) \sim u^{-\alpha}$, it appears as a straight line with slope $-\alpha$ on this scale. The power-law exponent was estimated independently via the Hill estimator:

$$\hat{\alpha} = \left(\frac{1}{k} \sum_{i=1}^k \ln \frac{x_{(i)}}{x_{(k+1)}} \right)^{-1} \quad (5.3)$$

where $x_{(1)} \geq \dots \geq x_{(n)}$ are the order statistics of $|x|$ and k corresponds to the top 10% of values (90th percentile threshold). This choice follows standard practice in extreme value analysis, providing a balance between bias (which increases as k grows, including non-tail observations) and variance (which increases as k shrinks, reducing the number of observations used in the estimate).

The stability index α_s of the Lévy-stable distribution and the Hill estimator $\hat{\alpha}$ provide complementary measures of tail heaviness. Values of $\alpha_s < 2$ (or equivalently $\hat{\alpha} < 2$) indicate infinite variance, a hallmark of anomalous diffusion and a necessary condition for deviations from the normal diffusion scaling $\zeta(q) = q/2$ observed in the moment scaling analysis of Chapter 6. Student- t distributions with small degrees of freedom ν similarly imply heavy tails, with the tail exponent equivalent to $\alpha = \nu$ for large $|x|$.

5.3 Moment scaling analysis

The scaling behaviour of statistical moments is investigated to detect signatures of anomalous or strongly anomalous diffusion, following the framework of Vollmer, Giberti, et al. (2024). For a stochastic process $x(t)$, here representing the seismic signal (acceleration,

velocity, or displacement), the q -th order moment of the increments $\Delta x(\tau) = x(t_0 + \tau) - x(t_0)$ is expected to scale as:

$$M_q(\tau) = \langle |\Delta x(\tau)|^q \rangle \sim \tau^{\zeta(q)}, \quad (5.4)$$

where τ is the time lag and $\langle \cdot \rangle$ denotes ensemble averaging. For normal diffusion, $\zeta(q) = q/2$ for all q , yielding a linear scaling spectrum. Anomalous diffusion is characterised by scaling exponents $\zeta(q)$ that deviate from this reference, while strong anomalous diffusion exhibits a piecewise-linear spectrum with a crossover at a critical moment order α (Vollmer, Giberti, et al., 2024):

$$\zeta(q) = \begin{cases} \xi q & q \leq \alpha, \\ \zeta q - \alpha(\zeta - \xi) & q > \alpha, \end{cases} \quad (5.5)$$

where $\xi < \zeta$. The lower branch ($q \leq \alpha$) is governed by the bulk of the displacement distribution, while the upper branch ($q > \alpha$) is dominated by rare large fluctuations in the tails. The crossover order α is directly connected to the power-law exponent of the tail of the displacement PDF (Vollmer, Giberti, et al., 2024), providing a theoretical link between the heavy-tail behaviour characterised in Section 5.2 and the moment scaling spectrum.

The analysis is performed independently for each signal type (acceleration, velocity, displacement), each temporal window (pre-event, P-wave, S-wave, coda, and post-event when present), and each phase picking method (AR-AIC and PhaseNet), enabling a systematic comparison of scaling behaviour across dynamical regimes and methodological choices. Each seismic signal is segmented into temporal windows following the procedure described in Section 4.2.6, with boundaries defined by the P-wave onset, S-wave onset, coda onset, and coda end times detected via AR-AIC and PhaseNet (Chapter 4). Coda onset is determined using three independent methods (Rautian, Arias intensity, envelope decay), plus an ensemble estimate, and the moment scaling analysis is repeated for each definition to assess robustness with respect to this methodological choice.

5.3.1 Spatial ensemble averaging

The analysis employs spatial ensemble averaging across multiple recording stations to obtain robust estimates of moment scaling exponents. For a given seismic window, the ensemble-averaged moment of order q at time lag τ is computed as:

$$\langle M_q(\tau) \rangle = \frac{1}{N} \sum_{i=1}^N |\Delta x_i(\tau)|^q, \quad (5.6)$$

where N is the number of independent signals (station-component pairs), and $\Delta x_i(\tau)$ is the increment for the i -th signal:

$$\Delta x_i(\tau) = x_i(t_{0,i} + \tau) - x_i(t_{0,i}). \quad (5.7)$$

Here, $t_{0,i}$ is the start time of the window for station i , which corresponds to different absolute times for different stations (different onset times across the network). The time lag τ is measured relative to the window start and spans a common range for all stations, limited by the duration of the shortest window in the ensemble.

The three components recorded at each station are treated as independent ensemble members. This choice is physically motivated by the different polarisations of seismic waves: P-waves are polarised along the propagation direction, while S-waves are polarised

perpendicular to it. Since each component samples a different projection of the wavefield, they provide independent realisations of the underlying stochastic process.

5.3.2 Computation of moments and scaling exponents

For each temporal window and observable, moments are computed for moment orders $q \in [0.5, 5.0]$ with step 0.25, yielding 19 values. Time lags τ are distributed logarithmically between $\tau_{\min} = 0.01$ s and $\tau_{\max} = \min_i(\text{duration}_i)$, where the maximum is set to the shortest window duration in the ensemble to ensure that all station-component pairs contribute at every τ value. Optionally, $\tau_{\max}$ can be restricted to a fraction of the shortest window duration to mitigate finite-size effects at large lags. The number of τ points is determined automatically as approximately 20 points per decade in log scale, with a minimum of 30 points total. To minimise rounding artefacts, τ values are generated directly as integer sample indices and converted to seconds only for output.

For each moment order q , the scaling exponent $\zeta(q)$ is estimated by fitting a power law to $\langle M_q(\tau) \rangle$ versus τ in log-log scale:

$$\log_{10} \langle M_q(\tau) \rangle = \zeta(q) \log_{10} \tau + c. \quad (5.8)$$

The slope $\zeta(q)$ is obtained by ordinary least squares regression on the log-transformed data, excluding points where $M_q(\tau) < 10^{-15}$ or where $M_q(\tau)$ is NaN or infinite. If fewer than 3 valid points remain for a given q , the corresponding exponent is set to NaN. The goodness of fit is assessed via the coefficient of determination R^2 , and the standard error of the slope provides an uncertainty estimate for $\zeta(q)$.

To detect potential deviations from simple scaling behaviour, the dependence of $\zeta(q)$ on q is analysed. For normal diffusion, $\zeta(q) = q/2$ yields a linear relationship with slope 0.5. Anomalous diffusion manifests as a different linear slope, while strong anomalous diffusion exhibits a piecewise-linear spectrum as in Eq. (5.5).

5.3.3 Expected behaviour and interpretation

The temporal windows defined in Section 4.2.6 are expected to exhibit qualitatively distinct scaling behaviours, reflecting the different physical processes active in each phase:

- **P-wave:** contains the compressional wave arrival at relatively low energy. Moderate deviations from normal diffusion are expected.
- **S-wave:** the most energetic phase, where PGA, PGV, and PGD all reach their maximum values. The strongest deviations from $\zeta(q) = q/2$ are expected here, consistent with the heavy-tail behaviour established in Section 5.2 and the impulsive, non-Gaussian character of shear wave arrivals.
- **Coda:** characterised by exponential amplitude decay as scattered energy dissipates. The scaling exponents $\zeta(q)$ are expected to progressively approach the pre-event behaviour as the signal returns towards background noise levels.

The systematic comparison of $\zeta(q)$ across windows allows identification of which seismic phase departs most strongly from normal diffusion, providing insight into which part of the signal carries information about the out-of-equilibrium nature of the seismic source and propagation process. Consistent signatures of anomalous diffusion across acceleration, velocity, and displacement would further strengthen the physical interpretation of the results.

Chapter 6

Results

This chapter reports the results of the analysis pipeline described in Chapters 4 and 5, applied to the two datasets introduced in Chapter 3: INT-41004391 (Queyras, Mw 3.8) and IT-2009-0009 (L'Aquila, Mw 6.1). Section 6.1 reports the empirical probability density functions and non-Gaussianity test results for acceleration signals of the INT-41004391 dataset. Section 6.2 presents the heavy-tail assessment for the same signals, including model selection via AIC and power-law tail characterisation. Section 6.3 presents the phase picking and signal segmentation results obtained with the AR-AIC and PhaseNet methods, including phase picking residuals and quality control statistics for both datasets. Section 6.4 reports the probability density functions and heavy-tail assessment for acceleration, velocity, and displacement signals. Section 6.5 presents the moment scaling results for both datasets, segmentation methods, and filtering configurations. Throughout this chapter, results are presented descriptively; their physical interpretation, and the comparison across datasets, segmentation methods, and configurations, are discussed in Chapter 7.

6.1 Empirical probability density functions and non-Gaussianity

Empirical probability density functions were computed for all 66 normalised acceleration signals of the INT-41004391 dataset. The Anderson-Darling test at the 5% significance level rejects the Gaussian null hypothesis for all 66 signals (100%), with test statistics exceeding the critical value by several orders of magnitude in every case. Excess kurtosis is strictly positive for all signals, with values ranging from 16.71 to 158.78 and a median of 49.99, indicating substantially heavier tails than a Gaussian distribution throughout the dataset. Skewness is variable in sign and moderate in magnitude (median absolute value 0.45), indicating that asymmetry is present but is not the dominant departure from Gaussianity. These results are summarised in Table 6.1.

6.2 Heavy-tail analysis

Three heavy-tailed distributions (Lévy-stable, Student- t , Laplace) were fitted to each of the 66 normalised acceleration signals via maximum likelihood estimation and compared against a Gaussian baseline using the Akaike Information Criterion. The Gaussian and Laplace distributions do not provide the best fit for any signal. The Lévy-stable

Metric	Value
Traces rejecting Gaussianity (AD test, 5%)	66/66 (100%)
Excess kurtosis: range	16.71–158.78
Excess kurtosis: median	49.99

Table 6.1: Summary of non-Gaussianity results for acceleration signals, INT-41004391 dataset ($N = 66$ traces).

distribution is selected as the best-fitting model for 55 of 66 signals (83%), with the Student- t preferred for the remaining 11 (17%). These results are summarised in Table 6.2; per-signal values of the stability index α_s , the Student- t degrees of freedom ν , and the Hill estimator $\hat{\alpha}$ are reported in Table C.1 in Appendix C.

The stability index α_s of the fitted Lévy-stable distributions ranges from 0.30 to 1.36, with a median of 0.72; all values are below 2, indicating infinite variance for every signal in the dataset. The Hill estimator $\hat{\alpha}$, computed independently from the top 10% of the empirical tail, ranges from 0.68 to 1.64 with a median of 1.14, consistent with the Lévy-stable fit results and confirming power-law tail behaviour with exponent below 2 throughout the dataset.

Metric	Value
Best-fit distribution: Lévy-stable	55/66 (83%)
Best-fit distribution: Student- t	11/66 (17%)
α_s (Lévy-stable): range	0.30–1.36
α_s (Lévy-stable): median	0.72
$\hat{\alpha}$ (Hill estimator): range	0.68–1.64
$\hat{\alpha}$ (Hill estimator): median	1.14

Table 6.2: Summary of heavy-tail assessment results for acceleration signals, INT-41004391 dataset ($N = 66$ traces). Per-signal values are reported in Table C.1.

6.3 Phase picking and signal segmentation

This section reports phase picking residuals and quality control statistics obtained from the AR-AIC and PhaseNet methods, applied to both datasets. For each method and dataset, numerical results are reported in Appendix D, separately for acceleration, velocity, and displacement signals.

6.3.1 AR-AIC phase picking

AR-AIC phase picking results are presented below for the INT-41004391 and IT-2009-0009 datasets.

INT-41004391 (Queyras, Mw 3.8)

For the INT-41004391 dataset, AR-AIC phase picking was applied to all 22 stations across the three signal types. Phase picking residuals for acceleration signals show a mean

P-wave offset of 0.88 ± 0.63 s and a mean S-wave offset of -1.13 ± 2.64 s with respect to the CRUST1.0 theoretical arrival times (Table D.4); velocity signals yield comparable P-wave residuals (0.86 ± 0.59 s), while displacement shows greater S-wave dispersion (0.29 ± 4.15 s), likely reflecting the smoother waveform morphology that increases onset ambiguity.

Quality control results are reported in Table D.5. *Peak check* failure rate ranges from 14.3% to 25.4% across signal types and coda methods, indicating that for a consistent minority of components, the peak ground-motion value (PGA, PGV, or PGD) occurs outside the S-wave window. Monotonicity failures are more prevalent: *MonoP check* and *MonoS check* failures affect respectively 59.1–72.7% and 72.7%–77.3% of components depending on signal type. *SNR P check* failures are negligible for acceleration and velocity (0.0%), but reach 10.6% for displacement, consistent with the lower amplitude of integrated signals at stations with poor signal-to-noise ratios. *SNR S check* failures are absent across all configurations. The segmentation results for a representative station of the dataset are illustrated in Figure 6.1.




Figure 6.1: Signal segmentation for station OGD1 ($d_{\text{epi}} = 67.2$ km, INT-41004391 dataset) using AR-AIC onset detection and Rautian coda onset method, for acceleration (left), velocity (center), and displacement (right) signals. Vertical lines indicate the detected P-wave onset (red), S-wave onset (teal), coda onset (orange), and coda end (blue dashed). Coloured backgrounds denote the pre-event, P-wave, S-wave, coda, and post-event windows.

IT-2009-0009 (L’Aquila, Mw 6.1)

For the IT-2009-0009 dataset, AR-AIC phase picking was applied to 62 stations across the three signal types, after exclusion of station GSG for a missing acquisition start timestamp and station TTS for an anomalous origin time offset incompatible with the segmentation pipeline. P-wave residuals are comparable in magnitude across the three signal types

(standard deviation between 1.92 and 2.32 s), although displacement shows a small negative mean offset (-0.61 s) compared to the positive offsets observed for acceleration and velocity (0.55 s and 0.23 s, respectively). For the S-wave residuals, displacement exhibits the largest dispersion ($\text{std} = 7.04$ s, versus 5.36 s and 5.83 s for acceleration and velocity), consistent with the trend observed in the smaller dataset. Figure 6.2 shows the corresponding segmentation for a representative station of the dataset, while Table D.13 reports the residual statistics for all three signal types.

Quality control results are reported in Table D.14. *Peak check* failures range from 6.7% to 48.2% across signal types and coda methods, with the Envelope method showing systematically higher failure rates. *MonoP check* and *MonoS check* failures are high across all signal types (82.3–87.1%) and (85.5%–90.3%) respectively, exceeding the values observed for INT-41004391 dataset. *SNR P check* failures are negligible for acceleration (0.0%) but reach 6.7% for velocity and 26.5% for displacement, the latter indicating significant signal attenuation in the integrated displacement records. *SNR S check* failures remain low for acceleration and velocity (0.0%), and reach 2.5–2.6% for displacement.




Figure 6.2: Signal segmentation for station SPO ($d_{\text{epi}} = 68.2$ km, IT-2009-0009 dataset) using AR-AIC onset detection and Rautian coda onset method, for acceleration, velocity, and displacement signals. Vertical lines indicate the detected P-wave onset (red), S-wave onset (teal), coda onset (orange), and coda end (blue dashed). Coloured backgrounds denote the pre-event, P-wave, S-wave, coda, and post-event windows.

6.3.2 Phase picking using PhaseNet

PhaseNet phase picking results are presented below for the INT-41004391 and IT-2009-0009 datasets.

INT-41004391 (Queyras, Mw 3.8)

For the INT-41004391 dataset, PhaseNet was applied in two configurations: no-threshold and probability threshold. In the no-threshold configuration, 21 stations were retained for acceleration and velocity, and 20 for displacement, after exclusion of stations for which PhaseNet did not produce a valid $t_S > t_P$ pair (Table D.6). Phase picking residuals for acceleration show a mean P-wave offset of 0.84 ± 0.60 s and a mean S-wave offset of 0.46 ± 1.48 s with respect to CRUST1.0 theoretical arrivals (Table D.7), comparable to the AR-AIC results for the same dataset. Velocity and displacement signals yield similar P-wave residuals (0.81 ± 0.58 s and 0.83 ± 0.59 s respectively), while S-wave dispersion remains low across all signal types, indicating that PhaseNet produces consistent onset estimates for this dataset.

Quality control results for the no-threshold configuration are reported in Table D.8. *Peak check* failures are markedly lower than for AR-AIC, ranging from 1.7% to 9.7% across signal types and coda methods. *MonoP check* failures are comparable to AR-AIC (70.0–71.4%), while *MonoS check* failures are lower (57.1–71.4%). *SNR P check* failures are negligible for acceleration and velocity (0.0%), and reach 8.3% for displacement. *SNR S check* failures are absent across all configurations.

In the probability threshold configuration, the number of retained stations drops to 16 for acceleration, 15 for velocity, and 10 for displacement (Table D.9), reflecting the stricter selection imposed by the minimum probability requirement. Mean S-wave probabilities increase substantially under the threshold configuration across all three signal types (from 47.3% to 56.7% for acceleration, from 45.1% to 57.2% for velocity, and from 32.6% to 52.8% for displacement), confirming that the threshold selects higher-confidence S-wave detections. Mean P-wave probabilities, already high in the no-threshold configuration, change only slightly under the threshold configuration, with no consistent direction across signal types: a small decrease for acceleration and velocity, and a larger increase for displacement (Tables D.6 and D.9).

Quality control results show further reductions under the threshold configuration: *Peak check* failure rates decrease to 0.0–8.5% (from 1.7–9.7% in the no-threshold configuration), *MonoP check* failure rates decrease to 60.0–68.8% (from 70.0–71.4%), and *MonoS check* failure rates decrease to 30.0–62.5% (from 57.1–71.4%). *SNR P check* failures for displacement rise to 16.7% (from 8.3%), consistent with the smaller sample of retained stations at this signal type, while *SNR S check* failures remain at 0.0% across both configurations and all signal types. The segmentation results for station OGD1 ($d_{\text{epi}} = 67.2$ km) are shown in Figure 6.3.




Figure 6.3: Signal segmentation for station OGD1 ($d_{\text{epi}} = 67.2$ km, INT-41004391 dataset) using PhaseNet onset detection and Rautian coda onset method, for acceleration (left), velocity (center), and displacement (right) signals. Vertical lines indicate the detected P-wave onset (red), S-wave onset (teal), coda onset (orange), and coda end (blue dashed). Coloured backgrounds denote the pre-event, P-wave, S-wave, coda, and post-event windows.

IT-2009-0009 (L’Aquila, Mw 6.1)

For the IT-2009-0009 dataset, the probability threshold configuration was not applied: given the systematically low S-wave detection probabilities (mean $P_S = 4.7\text{--}11.0\%$ depending on signal type), applying the $P \geq 0.3$ threshold would discard the large majority of stations (only 9 retained for acceleration), leaving an ensemble too small for the spatial-ensemble averaging required by the moment scaling analysis. Under the no-threshold configuration, after exclusion of stations with $t_S \leq t_P$, 59 stations were retained for acceleration, 61 for velocity, and 51 for displacement (Table D.15).

Phase picking residuals show a mean P-wave offset of 0.46 ± 2.10 s for acceleration, comparable to AR-AIC, but S-wave residuals exhibit substantially higher dispersion: -0.29 ± 9.34 s for acceleration, 2.72 ± 31.07 s for velocity, and -1.17 ± 11.93 s for displacement. The extreme S-wave dispersion for velocity, with a maximum residual of $+222.78$ s, reflects the presence of outlier picks at stations where the S-wave onset is poorly constrained by the model, likely due to the low S-wave detection probabilities characteristic of this dataset. Displacement signals additionally show anomalous P-wave residuals (2.99 ± 5.16 s, maximum $+26.75$ s), suggesting that waveform integration degrades onset detection.

Quality control results are reported in Table D.17. *Peak check* failures range from 1.7% to 41.8% across signal types and coda methods, with the Envelope method again showing the highest failure rates. *MonoP check* failures are high across all signal types (88.5–92.2%), exceeding those observed for AR-AIC on the same dataset, while *MonoS check*

failures range from 76.3% to 82.0%. *SNR P check* failures are negligible for acceleration (0.9%) but reach 5.4% for velocity and 31.1% for displacement, the latter indicating that the combination of low PhaseNet detection confidence and signal integration produces a significant fraction of displacement records with insufficient pre-event signal-to-noise ratio. *SNR S check* failures are absent for acceleration, low for velocity (0.0–1.0%), and reach 8.8–9.6% for displacement. The segmentation results for station SPO ($d_{\text{epi}} = 68.2$ km) using the PhaseNet no-threshold configuration are shown in Figure 6.4.

In addition to the no-threshold baseline, a residual filter configuration was applied to the dataset, as described in Section 4.3.3. Detection and residual statistics are reported in Tables D.18 and D.19, and quality control results in Table D.20. After exclusion of stations with $t_S \leq t_P$ and application of the S-wave residual filter ($k = 4$), 53 stations were retained for acceleration, 52 for velocity, and 48 for displacement. The filter removed 6, 9, and 3 stations for acceleration, velocity, and displacement respectively, primarily those with extreme S-wave residuals. S-wave residual dispersion is substantially reduced with respect to the no-threshold configuration for acceleration and velocity (± 4.48 s and ± 4.73 s, respectively, compared to ± 9.34 s and ± 31.07 s), while displacement retains higher dispersion (± 8.18 s) reflecting the intrinsic difficulty of S-wave detection on integrated signals. P-wave residual means remain broadly similar across configurations for all three signal types, but their dispersion narrows substantially for acceleration and velocity (std decreasing from 2.10 s to 0.78 s and from 2.12 s to 0.80 s, respectively), while it remains essentially unchanged for displacement (5.16 s in both configurations). This suggests that, although the residual filter criterion is applied exclusively to S-wave residuals, the stations it removes for acceleration and velocity tend to also be associated with larger P-wave residual variability, an effect not observed for displacement.




Figure 6.4: Signal segmentation for station SPO ($d_{\text{epi}} = 68.2$ km, IT-2009-0009 dataset) using PhaseNet onset detection and Rautian coda onset method, for acceleration, velocity, and displacement signals. Vertical lines indicate the detected P-wave onset (red), S-wave onset (teal), coda onset (orange), and coda end (blue dashed). Coloured backgrounds denote the pre-event, P-wave, S-wave, coda, and post-event windows.

6.4 Probability density functions and heavy-tail assessment

Acceleration

Velocity

Displacement

6.5 Moment scaling analysis

The following subsections present, for each dataset and segmentation configuration, ensemble-averaged moment scaling curves $\langle |\Delta x(\tau)|^q \rangle$ vs. τ and scaling exponent spectra $\zeta(q)$ vs. q , shown for the Rautian coda onset method. Results are presented separately for the AR-AIC (no residual filter) and PhaseNet configurations: probability threshold $P \geq 0.3$ for the INT-41004391 dataset, and residual filter for the IT-2009-0009 dataset.

Detailed moment scaling results for the P-wave, S-wave, and coda windows are reported in Appendix E, separately for each dataset, segmentation method, and filtering configuration. For each window and coda-onset method, the tables give the ensemble size N , the maximum time lag τ_{max} , and the scaling exponents $\zeta(1)$ and $\zeta(2)$ with the corresponding R^2 of the log-log fit. Within each subsection, tables are organised by signal

type (acceleration, velocity, displacement), with rows corresponding to the four coda onset methods (Rautian, Arias, Envelope, Median).

6.5.1 INT-41004391 (Queyras, Mw 3.8)

Moment scaling results for the INT-41004391 dataset are presented below for the AR-AIC and PhaseNet segmentation configurations. The dataset comprises 22 stations at epicentral distances ranging from 4.8 to 109.5 km, yielding an ensemble of up to 66 components per signal type.

AR-AIC segmentation

Figure 6.5 shows the ensemble-averaged moment scaling curves for the INT-41004391 dataset under AR-AIC segmentation with the Rautian coda onset method. The full numerical results for all four coda onset methods are reported in Tables E.1–E.3.

In the P-wave window, acceleration signals show a positive trend with moment order, with curves for different q values separated in the expected direction, though with substantial scatter around the power-law fits ($\zeta(1) = 0.95$, $R^2 = 0.28$; $\zeta(2) = 2.56$, $R^2 = 0.34$). The separation between curves increases with q , indicating that higher-order moments grow faster with τ , but the low R^2 values and the increasing point dispersion at $q = 3$ suggest that the power-law scaling is not well-constrained in this window. Velocity signals show a similarly weak positive trend ($\zeta(1) = 0.62$, $R^2 = 0.09$), while displacement signals yield negative scaling exponents ($\zeta(1) = -0.27$, $R^2 = 0.04$), with curves that appear nearly flat or slightly decreasing across the available τ range and no identifiable power-law trend. The P-wave window length is identical across all four coda onset methods ($\tau_{\max} = 2.14$ s for acceleration, 2.21 s for velocity, 2.40 s for displacement), as it depends only on the P- and S-wave onsets.

In the S-wave window, the scaling structure changes markedly. Acceleration curves are nearly flat across all q values, with $\zeta(1) \approx 0.02$ and $\zeta(2) \approx -0.17$ ($R^2 < 0.10$ for both). Velocity and displacement signals similarly show exponents close to zero or slightly negative, with low R^2 values indicating poor power-law fits. The $\tau_{\max}$ for the S-wave window varies across coda onset methods for all three signal types, reflecting differences in the estimated coda onset time across coda detection methods. The ensemble size is reduced for Arias ($N = 63$) and Envelope ($N = 56$) due to quality control exclusions, with small additional variation across signal types for the Envelope method ($N = 56$ –59).

In the coda window, all three signal types show negative scaling exponents across all coda onset methods. The $\tau_{\max}$ is substantially longer for the coda window than for the other phases, reaching up to 14.17 s for acceleration under the Rautian method. Despite the longer time range available, R^2 values remain low for most configurations, with the exception of Rautian and Envelope for velocity ($R^2 = 0.45$ and 0.41 respectively) and Median for displacement ($R^2 = 0.53$).




Figure 6.5: Ensemble-averaged moment scaling curves $\langle |\Delta x(\tau)|^q \rangle$ vs. τ for the INT-41004391 dataset (Queyras, M_w 3.8), AR-AIC segmentation, Rautian coda onset. Each panel shows a different seismic phase (P-wave, S-wave, coda) and signal type (acceleration, velocity, displacement). Dashed lines indicate power-law fits; moment orders $q = 0.5, 1.0, 2.0, 3.0$ are shown.

Figure 6.6 shows the scaling exponent spectrum $\zeta(q)$ for the INT-41004391 dataset under AR-AIC segmentation with the Rautian coda onset method.

In the P-wave window, $\zeta(q)$ grows monotonically with q for acceleration and velocity, well above the normal diffusion reference $\zeta(q) = q/2$. Acceleration shows the steepest growth, reaching $\zeta(5) \approx 6$, while velocity reaches $\zeta(5) \approx 4.5$. Error bars are large for both signal types at high q , reflecting the scatter visible in the scaling curves. Displacement yields negative exponents throughout, with $\zeta(q)$ decreasing monotonically as q increases and remaining well below the normal diffusion reference across the entire range of q .

In the S-wave window, the spectrum changes qualitatively across all signal types and all four coda onset methods (Tables E.1–E.3). Acceleration exponents are near zero for small q and become negative for $q \gtrsim 2$, with values around -0.5 at $q = 5$ under Rautian. For velocity, $\zeta(q)$ remains close to zero across the full range of q . Displacement exponents start near zero and decrease steadily, reaching approximately -1.5 at $q = 5$ under Rautian.

In the coda window, $\zeta(q)$ remains close to zero or slightly negative for acceleration and displacement across all coda onset methods. Velocity shows more negative values under Rautian. The spectrum is qualitatively consistent across the four coda onset methods, with differences confined to the magnitude of the negative exponents and the associated mean R^2 values.




Figure 6.6: Scaling exponent spectrum $\zeta(q)$ vs. q for the INT-41004391 dataset (Queyras, M_w 3.8), AR-AIC segmentation, Rautian coda onset. The dashed red line shows the normal diffusion reference $\zeta(q) = q/2$. Error bars represent one standard error of the log-log fit slope.

PhaseNet segmentation

Moment scaling results obtained from PhaseNet segmentation are reported separately for the two filtering configurations. Results for the no-threshold configuration are given in Tables E.4–E.6, and results for the probability threshold configuration ($P \geq 0.3$) in Tables E.7–E.9 (Appendix E). Figure 6.7 shows the ensemble-averaged moment scaling curves for the INT-41004391 dataset under PhaseNet segmentation with probability threshold $P \geq 0.3$, Rautian coda onset method.

In the P-wave window, the scaling structure differs markedly from the AR-AIC configuration. Acceleration curves show a clear negative slope across all q values, with moments decreasing with τ ($\zeta(1) = -0.42$, $R^2 = 0.55$; $\zeta(2) = -1.20$, $R^2 = 0.57$). Velocity signals show a similar pattern ($\zeta(1) = -0.37$, $R^2 = 0.63$, $\zeta(2) = -0.79$, $R^2 = 0.57$), while displacement yields weakly positive exponents ($\zeta(1) = 0.19$, $R^2 = 0.16$, $\zeta(2) = 0.90$, $R^2 = 0.35$). The ensemble size is reduced with respect to the AR-AIC configuration ($N = 48$ for acceleration, $N = 45$ for velocity, $N = 30$ for displacement), reflecting the stricter station selection imposed by the probability threshold.

In the S-wave window, negative scaling exponents are observed across all signal types and all four coda onset methods, with values consistently more negative than in the AR-AIC configuration. The R^2 values are substantially higher than in the AR-AIC S-wave window, indicating better-constrained power-law fits despite the smaller ensemble. The $\tau_{\max}$ varies across coda onset methods, consistent with the pattern observed in the AR-AIC

configuration.

In the coda window, exponents remain negative but close to zero across all signal types and coda onset methods. The $\tau_{\max}$ for the coda window is longer than for the other phases, reaching up to 17.70 s for velocity under Rautian, consistent with the longer coda duration estimated by this method.




Figure 6.7: Ensemble-averaged moment scaling curves for the INT-41004391 dataset (Queyras, M_w 3.8), PhaseNet segmentation with probability threshold $P \geq 0.3$, Rautian coda onset. Layout as in Figure 6.5.

Figure 6.8 shows the scaling exponent spectrum $\zeta(q)$ for the INT-41004391 dataset under PhaseNet segmentation with probability threshold $P \geq 0.3$, Rautian coda onset method.

In the P-wave window, $\zeta(q)$ is negative and decreasing for acceleration and velocity across all q values, with acceleration reaching $\zeta(5) \approx -4$ and velocity $\zeta(5) \approx -2$. Both spectra are approximately linear in q , well below the normal diffusion reference. Displacement shows a qualitatively different behaviour: $\zeta(q)$ is positive and increasing, growing approximately linearly with q and remaining close to the normal diffusion reference, though with large error bars. This pattern is consistent across all four coda onset methods (Tables E.7–E.9).

In the S-wave window, $\zeta(q)$ is negative and decreasing for all three signal types. Acceleration reaches $\zeta(5) \approx -3$ under Rautian, velocity reaches approximately -4 , and displacement approximately -2.5 . The spectra are approximately linear in q across all coda onset methods. The shape of the spectrum is qualitatively consistent across all four coda onset methods.

In the coda window, $\zeta(q)$ remains close to zero for acceleration across all coda onset methods, with values slightly negative and approximately flat in q . For velocity, Rautian yields the most negative exponents among the four methods, reaching approximately -0.5 at $q = 5$ ($\zeta(1) = -0.08$, $\zeta(2) = -0.17$); Median follows a similar but more moderate trend ($\zeta(2) = -0.13$), while Arias and Envelope remain close to zero throughout the range of q ($\zeta(1) = -0.04$ for both, $\zeta(2) = -0.11$ and -0.04 respectively).




Figure 6.8: Scaling exponent spectrum $\zeta(q)$ for the INT-41004391 dataset (Queyras, M_w 3.8), PhaseNet segmentation with probability threshold $P \geq 0.3$, Rautian coda onset. Layout as in Figure 6.6.

6.5.2 IT-2009-0009 (L'Aquila, M_w 6.1)

Moment scaling results for the IT-2009-0009 dataset are presented below for the AR-AIC and PhaseNet segmentation configurations. The dataset comprises 62 stations at epicentral distances ranging from 1.8 to 421.40 km, yielding an ensemble of up to 186 components per signal type.

AR-AIC segmentation

Figure 6.9 shows the ensemble-averaged moment scaling curves for the IT-2009-0009 dataset under AR-AIC segmentation with the Rautian coda onset method. The full numerical results for all four coda onset methods are reported in Tables E.10–E.12.

In the P-wave window, all three signal types show positive scaling. Acceleration shows well-separated curves across the full range of q values, while velocity and displacement

curves are closely spaced and cross multiple times, particularly between $\tau \approx 0.5$ and 1 s. Acceleration yields $\zeta(1) = 1.28$ ($R^2 = 0.87$) and $\zeta(2) = 2.00$ ($R^2 = 0.70$), velocity $\zeta(1) = 2.34$ ($R^2 = 0.91$) and $\zeta(2) = 4.55$ ($R^2 = 0.89$), and displacement $\zeta(1) = 1.80$ ($R^2 = 0.96$) and $\zeta(2) = 3.85$ ($R^2 = 0.95$). The power-law fits are well-constrained across all signal types, with R^2 values substantially higher than those observed for the INT-41004391 dataset. The ensemble size is larger than for Queyras ($N = 117$ – 123 depending on signal type), and the P-wave window length is identical across all four coda onset methods ($\tau_{\max} = 2.27$ s for acceleration, 2.08 s for velocity, 2.42 s for displacement).

In the S-wave window, the scaling structure changes markedly. Acceleration curves are nearly flat across all q values, with $\zeta(1)$ ranging from -0.15 to -0.12 ($R^2 = 0.19$ – 0.32) and $\zeta(2)$ ranging from -0.36 to -0.27 ($R^2 = 0.21$ – 0.38), depending on coda method. Velocity shows a similar pattern, with $\zeta(1)$ close to zero (-0.02 to 0.10 , $R^2 < 0.05$ for most methods) and a more negative $\zeta(2)$ (-0.05 to -0.17 , $R^2 = 0.01$ – 0.23), indicating, as for acceleration, a non-linear dependence of the exponent on q . Displacement yields weakly positive $\zeta(1)$ values (0.14 – 0.22 depending on method, $R^2 = 0.47$ – 0.61) but negative $\zeta(2)$ (-0.11 to -0.25), indicating a non-linear dependence of the exponent on q already at $q = 2$. The ensemble size for the S-wave window is larger than for the P-wave window ($N = 180$ for acceleration and velocity under Rautian), reflecting the larger number of stations with valid S-wave detections.

In the coda window, negative scaling exponents are observed for acceleration and velocity across all coda onset methods, with values in the range $\zeta(1) \in [-0.17, -0.03]$ for acceleration and $\zeta(1) \in [-0.15, -0.01]$ for velocity. Displacement shows a qualitatively different pattern, with positive $\zeta(1)$ values (0.11 – 0.34 depending on method) and negative or near-zero $\zeta(2)$ under Rautian and Arias, but positive $\zeta(2)$ under Envelope and Median. The $\tau_{\max}$ for the coda window is consistently shorter for the IT-2009-0009 dataset than for INT-41004391, across all three signal types and all four coda onset methods, despite the larger magnitude of the L'Aquila event.




Figure 6.9: Ensemble-averaged moment scaling curves $\langle |\Delta x(\tau)|^q \rangle$ vs. τ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), AR-AIC segmentation, Rautian coda onset. Layout as in Figure 6.5.

Figure 6.10 shows the scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset under AR-AIC segmentation with the Rautian coda onset method.

In the P-wave window, $\zeta(q)$ grows monotonically and approximately linearly with q for all three signal types, well above the normal diffusion reference. Acceleration reaches $\zeta(5) \approx 4$, velocity $\zeta(5) \approx 10$, and displacement $\zeta(5) \approx 9$, with small error bars reflecting the well-constrained fits already noted in the scaling curves ($R^2 \geq 0.70$ for all signal types). The spectrum is identical across all four coda onset methods, as expected given that the P-wave window boundaries are independent of the coda onset definition.

In the S-wave window, the qualitative structure of $\zeta(q)$ changes across signal types and is broadly consistent across all four coda onset methods. Acceleration shows $\zeta(q)$ starting near zero for small q and decreasing to approximately -0.5 to -1.0 at $q = 5$, with error bars that remain small relative to the signal. Velocity shows a similar pattern, with $\zeta(q)$ near zero at $q = 0.5$ and decreasing to approximately -0.5 at $q = 5$. Displacement shows a qualitatively different behaviour: $\zeta(q)$ is positive at small q ($\zeta(1) \approx 0.14$ – 0.22 depending on method), crosses zero and becomes negative at higher q , reaching approximately -1.0 to -1.5 at $q = 5$. This non-monotonic dependence of ζ on q is consistent across all four coda onset methods.

In the coda window, for acceleration, $\zeta(q)$ decreases monotonically with q under Rautian, Arias and Median, reaching values between approximately -0.5 and -0.8 at $q = 5$, while Envelope shows the opposite trend, increasing to approximately $+0.5$ at $q = 5$. Velocity shows negative $\zeta(q)$ decreasing with q under all four coda onset methods, with Rautian and Median reaching the most negative values (approximately -1.0 to -1.3

at $q = 5$) and Arias and Envelope showing a shallower decrease (approximately -0.3 to -0.4 at $q = 5$). Displacement shows the largest spread across methods: $\zeta(q)$ is positive under Envelope and Median, increasing to a maximum of approximately $+0.4$ to $+0.5$ around $q = 3$ – 3.5 before decreasing slightly at higher q , while Rautian and Arias show a monotonically decreasing trend, reaching approximately -2.0 and -1.5 respectively at $q = 5$.




Figure 6.10: Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), AR-AIC segmentation, Rautian coda onset. Layout as in Figure 6.6.

PhaseNet segmentation

For the IT-2009-0009 dataset, the no-threshold configuration and the residual filter configuration were applied (see Section 6.3.2). Figure 6.11 shows the ensemble-averaged moment scaling curves for the dataset under PhaseNet segmentation with residual filter, Rautian coda onset method. The full numerical results for all three signals and four coda onset methods are reported in Tables E.16–E.18.

In the P-wave window, acceleration shows well-separated curves across the range of q , while velocity curves remain closely spaced throughout, and displacement curves start well separated at small τ but cross around $\tau \approx 1$ – 2 s. Acceleration yields $\zeta(1) = 1.73$ ($R^2 = 0.89$) and $\zeta(2) = 3.85$ ($R^2 = 0.87$), displacement $\zeta(1) = 1.10$ ($R^2 = 0.99$) and $\zeta(2) = 2.53$ ($R^2 = 0.99$), while velocity shows weaker scaling ($\zeta(1) = 0.73$, $R^2 = 0.64$; $\zeta(2) = 0.97$, $R^2 = 0.23$). The ensemble size is reduced with respect to the AR-AIC configuration ($N = 96$ for acceleration, $N = 87$ for velocity, $N = 84$ for displacement), reflecting the stricter station selection imposed by the residual filter. The P-wave window

length is identical across all four coda onset methods.

In the S-wave window, all three signal types show negative scaling exponents, with more negative values and higher R^2 than observed under AR-AIC segmentation. Acceleration yields $\zeta(1) = -0.24$ ($R^2 = 0.54$) and $\zeta(2) = -0.51$ ($R^2 = 0.54$) under Rautian, consistent across all four coda onset methods. Velocity shows the most negative exponents among the three signal types ($\zeta(1) = -0.46$, $R^2 = 0.73$; $\zeta(2) = -1.22$, $R^2 = 0.75$ under Rautian), with the Envelope method yielding even more negative values ($\zeta(1) = -0.63$, $R^2 = 0.79$ $\zeta(2) = -1.77$, $R^2 = 0.83$). Displacement exponents are near zero for $\zeta(1)$ (-0.08 to 0.03 depending on method) but more negative for $\zeta(2)$ (-0.32 to -0.50), indicating a non-linear dependence on q .

In the coda window, scaling exponents are generally small in magnitude across all signal types and coda onset methods. Acceleration shows $\zeta(1) \in [-0.12, 0.02]$ and $\zeta(2) \in [-0.28, 0.12]$, with R^2 values that remain low for most methods except Arias ($R^2 = 0.30$ for $\zeta(1)$, and $R^2 = 0.26$ for $\zeta(2)$). Velocity yields consistently negative exponents under Rautian, Arias and Median ($\zeta(1) \in [-0.18, -0.16]$, $\zeta(2) \in [-0.40, -0.32]$), but positive values under Envelope. Displacement shows positive $\zeta(1)$ under Arias, Envelope and Median (0.14 – 0.22) and near-zero or negative values under Rautian ($\zeta(1) = -0.02$), with the Arias method yielding the best-constrained fit ($R^2 = 0.92$).




Figure 6.11: Ensemble-averaged moment scaling curves $\langle |\Delta x(\tau)|^q \rangle$ vs. τ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), PhaseNet segmentation with residual filter, Rautian coda onset. Layout as in Figure 6.5.

Figure 6.12 shows the scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset under PhaseNet segmentation with residual filter, Rautian coda onset method.

In the P-wave window, $\zeta(q)$ grows monotonically with q for all three signal types. Acceleration and displacement show an approximately linear growth, reaching $\zeta(5) \approx 10$ and $\zeta(5) \approx 7$ respectively, while velocity shows a spectrum with a visible change in slope around $q \approx 1$, growing faster at low q and more slowly at high q , reaching $\zeta(5) \approx 2$ with error bars larger than for the other two signal types.

In the S-wave window, $\zeta(q)$ is negative and decreasing for all three signal types across all four coda onset methods. Acceleration shows $\zeta(q)$ starting near zero and decreasing to approximately -1.0 to -1.3 at $q = 5$, with similar behaviour across all four methods. Velocity shows more strongly negative spectra, reaching approximately -3.0 at $q = 5$ under Rautian and Median, and approximately -2.0 at $q = 5$ under Arias, while Envelope yields markedly more negative values, reaching approximately -4.7 to -5.0 at $q = 5$. Displacement exponents decrease to approximately -0.9 to -1.5 at $q = 5$, with comparable magnitude across all four coda onset methods. The R^2 values are substantially higher than in the AR-AIC S-wave configuration for all signal types and methods, indicating better-constrained power-law fits.

In the coda window, the spectrum depends strongly on the coda onset method and differs qualitatively across signal types. For acceleration, $\zeta(q)$ remains close to zero for Arias and Median, decreases to approximately -0.7 to -0.8 at $q = 5$ under Rautian, and increases slightly to approximately $+0.5$ under Envelope. Velocity shows a similar diversity: Arias remains close to zero or slightly positive with large scatter, Envelope shows a mild negative trend, and Rautian and Median decrease more markedly, reaching approximately -0.7 to -1.0 at $q = 5$. Displacement shows positive $\zeta(1)$ under Arias, Envelope and Median (0.14 – 0.22). Arias and Median decrease and become negative at large q (approximately -2 and -0.8 respectively at $q = 5$), while Envelope is negative across most of the range, reaching approximately -0.5 at $q = 5$. Under Rautian, $\zeta(q)$ is near zero at small q and decreases monotonically to approximately -0.5 at $q = 5$.




Figure 6.12: Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), PhaseNet segmentation with residual filter, Rautian coda onset. Layout as in Figure 6.6.

Chapter 7

Discussion

This chapter discusses the physical and methodological interpretation of the results presented descriptively in Chapter 6, in light of the statistical framework introduced in Chapter 5 and the research question motivating this thesis: whether seismic ground motion signals carry statistical signatures of non-equilibrium complex systems such as heavy-tailed distributions, anomalous diffusion, and moment scaling, and whether these signatures are robust across signal types, seismic phases, datasets, and methodological choices.

Section 7.1 discusses the departures from Gaussianity and the heavy-tail behaviour identified for the full preprocessed signals. Section 7.2 examines the phase picking and segmentation results obtained from the AR-AIC and PhaseNet methods, comparing their reliability and the implications for the downstream moment scaling analysis. Section 7.3 addresses the central question of the thesis, evaluating the evidence for anomalous and strongly anomalous diffusion across the P-wave, S-wave, and coda windows, and assessing the sensitivity of the scaling exponents $\zeta(q)$ to the choice of coda onset method and segmentation configuration. Section 7.4 compares the two datasets directly, considering the role of event magnitude, epicentral distance, and station coverage in shaping the observed scaling behaviour. Finally, Section 7.5 discusses the limitations of the present analysis and their implications for the conclusions drawn in this thesis.

7.1 Non-Gaussianity and heavy-tail behaviour

The non-Gaussianity results reported in Section 6.1 are among the most clear-cut findings. The Anderson-Darling test rejects the Gaussian null hypothesis for all 66 acceleration signals of INT-41004391, with test statistics exceeding the critical value by several orders of magnitude in every case, and excess kurtosis is strictly positive throughout, with a median of 49.99. There are no borderline cases: the departure from Gaussianity is both universal and extreme across the entire spatial ensemble. This stands in contrast to the moment scaling results discussed in Sections 7.3.1–7.3.3, where evidence for anomalous scaling is conditional on the choice of window, picker, and dataset. The non-Gaussianity of the amplitude distribution, by contrast, appears to be a robust property of the signals themselves, independent of the segmentation choices examined elsewhere in this thesis.

The heavy-tail assessment (Section 6.2) adds a more specific characterisation: the Lévy-stable distribution is preferred over all alternatives by AIC for 55 of 66 signals (83%), with the Student- t preferred for the remaining 11 (17%). Neither the Gaussian nor the Laplace distribution provides the best fit for any signal. The two winning distributions

have different physical implications. The Student- t is a symmetric, unimodal distribution with algebraically decaying tails but without the stability-under-convolution property; it arises naturally as a scale mixture of Gaussians and is a common empirical model for heavy-tailed data without strong physical motivation. The Lévy-stable family, by contrast, is closed under convolution and summation, meaning that sums of independent Lévy-stable random variables remain Lévy-stable with the same stability index α_s . Its prevalence here is consistent with a physical picture in which the observed ground motion is the superposition of many independent scattering contributions, each heavy-tailed, whose aggregate retains the same distributional character — though this remains a plausible interpretation rather than a conclusion established by the present data.

A particularly relevant consequence of the fitted distributions concerns the stability index α_s and the Hill estimator $\hat{\alpha}$. Both are below 2 for all 66 signals, with medians of 0.72 and 1.14 respectively, indicating infinite variance throughout the dataset. This is not merely a tail curiosity: a necessary condition for deviations from the normal diffusion scaling $\zeta(q) = q/2$ is that the variance of the increments $\Delta x(\tau)$ diverges, which requires a tail exponent below 2. The heavy-tail results therefore provide an independent, distribution-level motivation for the moment scaling analysis of Section 6.5: the signals examined here possess the distributional prerequisites for anomalous diffusion, and the moment scaling analysis tests whether this property manifests as a well-defined scaling relation across lags.

Two limitations of the present analysis should be kept in mind when interpreting these results. First, the PDF analysis has been carried out only for acceleration signals of INT-41004391; velocity and displacement signals, as well as all signals of IT-2009-0009, have not been processed due to computational constraints. Whether the universality of non-Gaussianity and the prevalence of Lévy-stable behaviour extend to other signal types and to the larger dataset therefore cannot be established from the present results alone. Second, the analysis operates on full signals rather than on temporally windowed segments, meaning that the distributional results reflect the aggregate character of the entire recording rather than the properties of any specific seismic phase. The relationship between the full-signal heavy-tail behaviour documented here and the window-specific scaling exponents reported in Section 6.5 is therefore indirect, and the two analyses should be understood as complementary rather than directly comparable.

7.2 Phase picking and signal segmentation

This section discusses the phase picking results presented in Section 6.3, comparing the AR-AIC and PhaseNet methods in terms of residual magnitude, quality control failure rates, and their implications for the reliability of the resulting temporal windows.

7.2.1 AR-AIC picking

The most striking feature of the AR-AIC quality control results (Section 6.3.1) is the prevalence of monotonicity failures: the majority of components violate the expected ordering of arrival times with hypocentral distance, for both datasets (59.1–77.3% for INT-41004391, 82.3–90.3% for IT-2009-0009). Rather than indicating picking errors, this pattern is more plausibly attributable to the MonoP/MonoS check itself, which assumes a strictly monotonic relationship between arrival time and distance under a simplified, laterally homogeneous velocity structure. Real crustal heterogeneity and path effects,

well documented in local earthquake studies through the systematic station corrections required to reconcile observed arrival times with simplified velocity models (Kissling et al., 1994), together with focal mechanism radiation patterns, can all produce non-monotonic arrival sequences that are physically legitimate rather than picking artefacts, suggesting that this check, while useful as a diagnostic flag, should not be interpreted as a direct measure of picker accuracy.

This interpretation is reinforced by comparing the two datasets directly. Both the monotonicity and Peak check failure rates are substantially higher for IT-2009-0009 than for INT-41004391 (up to 90.3% against 77.3% for monotonicity, and up to 48.2% against 25.4% for the Peak check), and this difference tracks the geographic extent of the two networks rather than any property of the picker itself: IT-2009-0009 stations span a much wider and more heterogeneous range of epicentral distances (1.80–421.40 km, Table D.2) than those of INT-41004391, increasing the influence of path-dependent velocity heterogeneity on the validity of the monotonicity assumption.

A second, independent pattern emerges for displacement signals, which fail the SNR P check more frequently than acceleration or velocity in both datasets (10.6% for INT-41004391, 26.5% for IT-2009-0009), and exhibit the largest S-wave residual dispersion among the three signal types in both cases. This is consistent with the integration-induced degradation of high-frequency timing information discussed widely in this thesis, and indicates that this effect is not specific to a single dataset or picking method, but a systematic consequence of working with doubly-integrated signals.

A further pattern specific to IT-2009-0009 is the elevated Peak check failure rate under the Envelope coda onset method relative to Rautian, Arias, and Median (up to 48.2%, against a maximum of 23.3% for the other three methods); no comparable discrepancy is observed for INT-41004391, where the four methods yield broadly similar failure rates (Table D.5). This suggests that, for this dataset, the S-wave window boundaries implied by the Envelope method are less consistent with the expected location of the peak ground-motion value than those implied by the other coda definitions; the methodological origin of this dataset-specific discrepancy is examined further in Section 7.3.4.

One anomaly does not fit a general pattern and is worth flagging as an open question: the mean P-wave residual for displacement changes sign between the two datasets, from a small positive offset for INT-41004391 (0.32 s, consistent with acceleration and velocity) to a negative offset for IT-2009-0009 (−0.61 s, opposite in sign to acceleration and velocity in the same dataset). This indicates that the sign reversal is specific to the combination of displacement signals and the IT-2009-0009 dataset, rather than a general property of integrated signals, and its origin remains unexplained by the present analysis.

Finally, it is worth acknowledging an interpretive limitation that applies to all AR-AIC residuals discussed in this section: since residuals are computed relative to theoretical arrival times derived from the CRUST1.0 crustal model, an elevated residual may reflect either a genuine picking error or an inaccuracy in the assumed velocity structure, and the two contributions cannot be disentangled with the available data. The comparisons drawn above should therefore be understood as comparisons of picking consistency, not as an absolute validation of AR-AIC against a perfectly known ground truth.

7.2.2 PhaseNet picking

The PhaseNet results (Section 6.3.2) reveal a marked contrast between the two datasets that is largely absent from the AR-AIC results discussed above. For INT-41004391,

PhaseNet performs well: P-wave residuals are comparable to AR-AIC, but S-wave residual dispersion is substantially lower (1.48–1.58 s, against 2.64–4.15 s for AR-AIC), and overall quality control is markedly better (Peak check failures of 1.7–9.7%, against 14.3–25.4% for AR-AIC). For IT-2009-0009, by contrast, S-wave performance deteriorates sharply, with residual standard deviations of 9.34, 31.07, and 11.93 s for acceleration, velocity, and displacement respectively, including an extreme outlier of +222.78 s for velocity. This asymmetry is consistent with a generalisation gap rooted in the composition of the INSTANCE training set: INT-41004391 (M_w 3.8, regional distances) is plausibly closer to the bulk of the training distribution than IT-2009-0009, a larger event (M_w 6.1) observed out to 421 km, a magnitude and distance regime far less typical of the training data, the latter dominated by M 2–4 earthquakes recorded mostly within 200 km (Michellini et al., 2021). This interpretation is not merely inferred from the residual statistics: the mean S-wave detection probabilities themselves confirm it directly, falling from 32.6–47.3% for INT-41004391 to only 4.7–11.0% for IT-2009-0009, indicating that the model’s own confidence, not only its downstream timing accuracy, is informative of this generalisation gap.

The magnitude of PhaseNet’s worst-case S-wave error on IT-2009-0009 is also notable in comparison to AR-AIC on the same dataset, where the largest residual magnitude does not exceed 14 s. This difference plausibly reflects a structural distinction between the two methods rather than a difference in picking accuracy alone: AR-AIC operates within search windows constrained by CRUST1.0 theoretical arrivals, which act as an implicit safety net bounding the magnitude of any individual error, whereas PhaseNet is applied directly to the full waveform with no such constraint. When the model is poorly calibrated for a given recording, as appears to be the case for IT-2009-0009, there is consequently nothing to prevent a substantially larger error.

The monotonicity results provide a useful cross-check on the interpretation proposed in Section 7.2.1. PhaseNet exhibits the same dataset-dependent pattern observed for AR-AIC, with IT-2009-0009 again showing higher MonoP failure rates than INT-41004391 (88.5–92.2% against 70.0–71.4%). The fact that two independent picking methods, one model-based and one constrained by a velocity model, show the same qualitative dependence on dataset strengthens the case that the underlying cause is the broader heterogeneous station coverage of IT-2009-0009, rather than a limitation specific to either picker. At the same time, the two methods do not agree on which performs better overall: PhaseNet’s MonoP failure rate for IT-2009-0009 is slightly higher than the corresponding AR-AIC maximum (92.2% against 87.1%), while its Peak check failure rate is lower (41.8% against 48.2%). Neither picker therefore dominates across all quality control metrics.

Displacement signals for IT-2009-0009 emerge as a particularly compromised case for PhaseNet. P-wave residuals for this signal type (2.99 ± 5.16 s, maximum +26.75 s) are substantially worse than those obtained from AR-AIC for the same signal type and dataset (-0.61 ± 2.32 s), suggesting that PhaseNet’s reduced reliability on IT-2009-0009 compounds with a further generalisation gap specific to displacement signals, which are likely under-represented relative to acceleration and velocity in the INSTANCE training set.

The trade-off between detection confidence and ensemble size, already noted in Section 6.3.2 as the reason the probability threshold configuration was not applied to IT-2009-0009, is worth revisiting here in interpretive terms: requiring high-confidence picks would have retained only 9 stations for acceleration, too few for reliable spatial-ensemble averaging in the moment scaling analysis, illustrating a more general tension between

picking reliability and ensemble size that is particularly acute when detection confidence is uniformly low, as for this dataset. The residual filter applied instead to IT-2009-0009 offers a more flexible alternative precisely because it does not rely on the model’s own confidence estimates, but on the consistency of the resulting residuals. Its effect is also informative beyond the S-wave residuals it directly targets: removing stations with extreme S-wave residuals reduces the P-wave residual standard deviation by more than half for acceleration and velocity ($2.10 \rightarrow 0.78$ s and $2.12 \rightarrow 0.80$ s respectively), indicating that the stations identified as problematic by the S-wave criterion tend to be problematic for both phases simultaneously, rather than for the S-wave alone.

7.3 Evidence for anomalous scaling

This section addresses whether the moment scaling results presented in Section 6.5 provide evidence for anomalous or strongly anomalous diffusion, and how this evidence varies across seismic phases, coda onset definitions, and segmentation configurations. The discussion proceeds phase by phase (Sections 7.3.1–7.3.3), before addressing the role of the coda onset method (Section 7.3.4) and of the broader segmentation configuration, including phase picker and filtering choices (Section 7.3.5).

7.3.1 P-wave window

The P-wave window provides the clearest illustration of how strongly the evidence for anomalous scaling depends on dataset and segmentation configuration. For IT-2009-0009, all three configurations (AR-AIC, PhaseNet no-threshold, PhaseNet residual filter) yield positive scaling exponents for all three signal types. Acceleration and displacement lie well above the normal diffusion reference, with consistently well-constrained log–log fits (R^2 up to 0.99; Tables E.10, E.12, E.13, E.18). Under AR-AIC, velocity behaves similarly to acceleration and displacement, with $\zeta(q)$ lying above the normal diffusion reference throughout the range of q considered. Under PhaseNet, by contrast, velocity exhibits a non-linear $\zeta(q)$ spectrum: exponents grow with a steeper slope up to $q \approx 1$, above the normal diffusion line, before transitioning to a shallower slope at higher q that falls below the reference, with markedly weaker fits than acceleration and displacement (R^2 as low as 0.23–0.29 for $\zeta(2)$; Tables E.14, E.17).

For INT-41004391 the evidence is more fragile, and the sign of the scaling exponent depends on the segmentation method in a way that also reverses across signal types. Under AR-AIC, acceleration and velocity show positive, growing moments while displacement decreases, but the fits are poorly constrained ($R^2 = 0.09$ –0.34). Under PhaseNet (both no-threshold and probability threshold configurations), this pattern inverts: acceleration and velocity become negative, while displacement becomes the only signal type with positive, growing exponents, now with better-constrained fits ($R^2 = 0.31$ –0.63) but opposite sign to the AR-AIC result for the same window, dataset, and signal types. The evidence for anomalous scaling in the P-wave window is therefore not only weaker for INT-41004391 than for IT-2009-0009, but its very sign, and which signal type carries it, is contingent on the choice of phase picker, raising the question of whether this window reflects a genuine physical signal or an artefact of the segmentation method for the smaller event.

For IT-2009-0009, where the overall evidence is more robust, velocity is the least stable signal type across configurations. Under AR-AIC it shows the highest exponent of the three signal types ($\zeta(1) = 2.34$), but under PhaseNet (both no-threshold and residual

filter configurations) it becomes the lowest, with the weakest fit for $\zeta(2)$ among the three signal types ($R^2 = 0.29$ and 0.23 respectively). Acceleration and displacement remain comparatively stable in both magnitude and fit quality between AR-AIC and PhaseNet, suggesting that velocity is the weak point even within the otherwise more robust dataset.

A caveat applies to the magnitude of the exponents reported across all configurations. The P-wave window is intrinsically short ($\tau_{\max}$ between 2.11 and 2.79 s in nearly all configurations), limiting the dynamic range of τ available for the fit. This constrains the reliability of the extrapolated exponents at high moment order: the values reported for IT-2009-0009, reaching $\zeta(5) \approx 9$ – 10 for velocity and displacement under AR-AIC, appear large for a physical scaling exponent even where the fit at $q = 1, 2$ is well constrained, and should be interpreted with this limitation in mind (Section 7.5).

Finally, it is worth recalling that the P-wave window is identical across all four coda onset methods, since its boundaries depend only on the detected P- and S-wave onset times. The comparisons drawn in this subsection therefore span only the dataset and phase-picker dimensions.

7.3.2 S-wave window

The S-wave window presents a different picture from the P-wave window discussed above: rather than a sign that depends on the phase picker, the two methods here disagree primarily on how confidently any scaling behaviour can be claimed at all, and they do so broadly across both datasets. Under AR-AIC, $\zeta(q)$ is typically negative and modestly decreasing for acceleration and displacement, but the fits are poorly constrained, with R^2 typically below 0.10 for INT-41004391 and only modestly higher for IT-2009-0009 (Tables E.1–E.3, E.10–E.12). A notable exception is velocity for INT-41004391 under the Envelope and Median coda onset methods, where $\zeta(q)$ is instead positive and increasing, although with wide error bars; this dataset- and method-specific reversal is addressed further in Section 7.3.4. Under PhaseNet, by contrast, exponents are clearly negative and decreasing for all three signal types in both datasets, and the fits are substantially better constrained. Whereas the P-wave window raised the question of which sign to trust, the S-wave window raises the question of which picker provides genuine evidence: AR-AIC yields weak and, for velocity, inconsistent evidence of anomalous scaling in this window, while PhaseNet yields a comparatively clear and uniform one.

One observation lends support to the interpretation that the PhaseNet result reflects a genuine signal rather than residual noise. For INT-41004391, moving from the no-threshold to the probability threshold configuration, i.e. retaining only the most confident picks, makes the negative exponents systematically more negative for all three signal types and coda methods; Tables E.4–E.9). If the negative exponents were an artefact of picking noise, selecting cleaner picks would not be expected to reinforce the same trend so consistently.

This pattern does not extend straightforwardly to IT-2009-0009: the residual filter, despite substantially narrowing the S-wave residual distribution (Section 7.2.2), leaves the S-wave scaling exponents essentially unchanged relative to the no-threshold configuration, for all three signal types (Tables E.13–E.18 against Tables E.16–E.18). This near-invariance may indicate that the magnitude of a station’s S-wave timing residual is not necessarily correlated with the size of its contribution to the ensemble moments of $\Delta x(\tau)$, since the residual reflects only the picking error relative to the expected onset, not how much the resulting windowed segment departs from the ensemble’s typical behaviour.

A further observation concerns velocity for IT-2009-0009, which shows the single

most extreme S-wave exponent across the entire moment scaling analysis under PhaseNet ($\zeta(2)$ reaching -1.77 under the Envelope coda onset method). This is the same signal type already identified in Section 7.2.2 as exhibiting the most extreme and unstable PhaseNet timing residuals for IT-2009-0009 (standard deviation of 31.07 s, with an outlier of +222.78 s). This raises the possibility that the unusually large exponent reflects, at least in part, an artefact of picking noise propagating into the moment scaling fit, and this limitation should be kept in mind when interpreting the magnitude of this particular result.

Finally, the two datasets differ in how homogeneously the three signal types behave under PhaseNet. For INT-41004391, acceleration, velocity, and displacement show a comparatively uniform, moderately negative response. For IT-2009-0009, the three signal types diverge considerably, with acceleration showing a moderate negative exponent, velocity an extreme one, and displacement a moderate one. This is consistent with the dataset-dependent variation in PhaseNet reliability across signal types already noted in Section 7.2.2, and suggests that this variability extends from the picking stage into the downstream scaling results.

7.3.3 Coda window

In contrast to the P-wave and S-wave windows, the coda window yields comparatively modest scaling exponents in most configurations, with $\zeta(1)$ and $\zeta(2)$ typically between approximately -0.4 and $+0.4$ (Tables E.1–E.18). This bound holds comparatively well for INT-41004391, where departures, when present, are isolated and modest in magnitude. For IT-2009-0009, departures from this range are more frequent and can be substantially larger, with $\zeta(q)$ at $q = 5$ reaching values as negative as approximately -2 under certain picker and coda onset method combinations; the specific combinations responsible for these departures, and their dependence on coda onset method, are addressed in full in Section 7.3.4. Even accounting for these exceptions, the coda window remains substantially less extreme overall than the P-wave window, where $\zeta(5)$ reaches approximately 9–10 for several signal types, or the S-wave window, where $\zeta(2) = -1.77$ for velocity in IT-2009-0009.

Unlike the P-wave window, where AR-AIC and PhaseNet disagreed on sign, and the S-wave window, where they disagreed on the strength of the evidence, the coda window is the first to show comparably modest exponent magnitudes under both pickers for INT-41004391. For IT-2009-0009, by contrast, agreement between pickers is weaker: the two pickers can differ substantially in magnitude for the same signal type and coda onset method, although they typically agree on sign. This picker-level disagreement is distinct from, and compounds with, the coda-onset-method disagreement discussed in Section 7.3.4: for each signal type, one of the four coda onset methods departs from the other three, though which method diverges, and under which segmentation configuration, varies.

One combination stands out as a candidate for a genuine, rather than artefactual, signal: IT-2009-0009 displacement under the Arias coda onset method yields a high and notably stable R^2 for $\zeta(1)$ across all three segmentation configurations (0.93 for AR-AIC, 0.90 for PhaseNet no-threshold, 0.92 for PhaseNet residual filter). This is the most reproducible result observed across the entire moment scaling analysis discussed so far; notably, it is also one of the configurations responsible for the more extreme departures from the modest-exponent range noted above, with $\zeta(q)$ reaching approximately -2 at $q = 5$ under both AR-AIC and PhaseNet.

Two points already established elsewhere in this thesis are worth recalling in this context. First, the residual filter again leaves the IT-2009-0009 results essentially unchanged relative to the no-threshold configuration, the third window in succession for which this holds (Section 7.3.2); here the invariance holds without exception across all four coda onset methods and all three signal types, including velocity under the Median method, which had been the one departure found in the S-wave window. Second, the coda window's $\tau_{\max}$ is, counter-intuitively, shorter for IT-2009-0009 than for INT-41004391 despite the larger magnitude of the former event, a fact already established in Chapter 6 that should be kept in mind when interpreting the relative fit quality of the two datasets in this window.

7.3.4 Effect of coda onset method

The coda onset method does not only determine the boundary of the coda window itself, but, by construction, also the upper boundary of the S-wave window; its effect is therefore examined across both windows in what follows.

For each signal type in the coda window, exactly one of the four coda onset methods diverges from the other three for INT-41004391, although which method diverges, and under which picker, varies. For acceleration, Envelope yields a markedly steeper, more negative trend than the other three methods under PhaseNet, in both the no-threshold and probability threshold configurations, reaching approximately -1 at $q = 5$ while Rautian, Arias, and Median remain close to zero; under AR-AIC, by contrast, all four methods decline similarly. For velocity, Arias remains close to zero under both AR-AIC and PhaseNet no-threshold, with a mild non-monotonic dip of around -0.4 to -0.5 , while the other three methods decline more steadily; this divergence disappears under the PhaseNet probability threshold configuration, where all four methods become comparably uniform. For displacement, Rautian remains close to zero under AR-AIC while the other three methods decline more clearly, whereas under PhaseNet (both configurations) all four methods are already broadly consistent with one another. The recurring structure here, one method departing while the remaining three agree, holds across all three signal types, even though the specific method responsible, and the picker under which it appears, both vary.

A different structure appears in the S-wave window for INT-41004391 velocity under AR-AIC: rather than a single method departing from the other three, the four methods split into two pairs. Envelope and Median yield a positive and increasing $\zeta(q)$, in contrast with the negative and decreasing trend typical of this signal type and window, while Rautian and Arias remain negative, in line with the general pattern; error bars are wide throughout. This case differs from the coda window pattern above in two respects: the divergence is a 2-against-2 split rather than 1-against-3, and it involves a reversal of sign rather than only a difference in magnitude.

For IT-2009-0009, the coda window does not reproduce the orderly one-method-departs structure observed for INT-41004391. Several method-signal combinations depart from the modest exponent range noted in Section 7.3.3: acceleration under Rautian, velocity under both Rautian and Median, and displacement under Arias all reach magnitudes well beyond what is typical elsewhere in this window, under both pickers where applicable. No single method is responsible for all of these departures, nor does any one method remain consistently close to zero across all three signal types; Rautian, for instance, departs for both acceleration and velocity but not for displacement, while Median departs only for velocity. This lack of a consistent culprit method is itself a point of contrast with

INT-41004391, and is addressed further in mechanistic terms below.

A further departure for IT-2009-0009 displacement concerns sign rather than magnitude alone. Under AR-AIC, displacement yields a positive $\zeta(1)$ under coda onset methods, while acceleration and velocity remain negative; the dependence on q is non-monotonic, with $\zeta(2)$ positive under some coda methods and negative under others. Under PhaseNet, displacement is also positive rather than negative under three of the four coda onset methods (Arias, Envelope, Median), the exception being Rautian, where $\zeta(1)$ remains close to zero. The agreement between pickers observed for acceleration and velocity in this window therefore does not carry over to displacement, which is positive under both pickers for most coda onset methods, though the specific values, and in some cases the sign itself, depend on which method is chosen. This sign anomaly does not occur for INT-41004391, where displacement remains negative and decreasing under every method and picker combination, with the sole exception of Rautian under AR-AIC, which remains close to zero rather than turning positive.

The clearest and most consistent pattern concerning the choice of coda onset method is the divergence in *Peak check* performance for IT-2009-0009, already noted in quality control terms in Chapter 6: Envelope is systematically the worst-performing method, both under AR-AIC (failure rates up to 48.2%) and under PhaseNet (up to 41.8%), while Arias is systematically the best (6.7–11.1% under AR-AIC, 1.7–4.6% under PhaseNet). For INT-41004391, by contrast, the four methods perform comparably, with no method standing out as consistently worse. This indicates that sensitivity to the coda onset method is not an intrinsic property of the method itself, but emerges primarily when station coverage is broader and more heterogeneous, the same mechanism already invoked in Section 7.2.1 to explain the dataset-dependent monotonicity failures.

An important caveat applies throughout this comparison: the maximum lag $\tau_{\max}$ available for the fit varies across coda onset methods, since the coda onset itself, and hence the boundaries of the S-wave and coda windows, differ between methods. Part of the variation observed between methods may therefore reflect differences in the temporal range used for the fit rather than a genuine difference in the underlying scaling behaviour, and this should temper any strong conclusions drawn from comparisons between methods.

These patterns are broadly consistent with the conceptual characterisation of the three independent coda onset methods established in Chapters 4 and 5: Rautian is a deterministic, phase-anchored formula, comparatively insensitive to the specific shape of the waveform; Arias is an energy-based criterion anchored to the signal itself; and Envelope is an amplitude-based criterion with only partial phase anchoring. This offers a plausible mechanistic explanation for why Envelope, being the most sensitive to the actual shape of the waveform, is also the most unstable when a broader, more heterogeneous network such as IT-2009-0009 introduces greater genuine waveform variability for it to capture.

Arias does not perform uniformly well or poorly across configurations. It yields the single most reproducible result in the entire moment scaling analysis, for IT-2009-0009 displacement in the coda window (Section 7.3.3), but it is also the method with the weakest R^2 for INT-41004391 acceleration and velocity in the coda window under AR-AIC (0.02–0.04). This is consistent with the insensitivity of the Arias coda onset to picking perturbations already established for the smaller, more distant INT-41004391 dataset, which may also translate into a less informative window for the moment scaling fit in that specific context, though this remains a plausible interpretation rather than a finding directly established by the present data. Median, finally, behaves as expected of a robust

ensemble combination: its values remain close to the central tendency of the other three methods in every configuration examined, an indirect confirmation that its intended role as a robust ensemble estimate functions as designed.

7.3.5 Effect of segmentation configuration

The discussion of individual windows in Sections 7.3.1–7.3.3 shows that AR-AIC and PhaseNet are never simply equivalent: they disagree on the sign of the scaling exponent in the P-wave window for the INT-41004391 dataset, on the strength of the evidence in the S-wave window for both datasets, and converge only partially in the coda window, where agreement holds for INT-41004391 but breaks down for IT-2009-0009. No single window is therefore immune to the choice of segmentation method, and any claim of anomalous scaling should be understood as conditional on this choice.

The two PhaseNet filtering strategies considered affect the moment scaling results in different ways, despite both reducing the size of the station ensemble. The probability threshold applied to INT-41004391 systematically changes the results: S-wave exponents $\zeta(1)$ and $\zeta(2)$ become more negative as the threshold is tightened (Section 7.3.2). The residual filter applied to IT-2009-0009, by contrast, leaves the S-wave and coda results essentially unchanged relative to the no-threshold configuration, despite substantially narrowing the residual distribution it directly targets (Sections 7.3.2 and 7.3.3). A plausible explanation for this difference lies in the extent and nature of the filtering: the probability threshold removes a substantially larger fraction of the ensemble than the residual filter, and selects on a different dimension — PhaseNet’s picking confidence — that does not necessarily correlate with timing residual magnitude, potentially affecting a broader range of the residual distribution than the MAD-based filter does.

This is accompanied by a more general observation on the relationship between ensemble size and fit quality. For INT-41004391, moving from the no-threshold to the probability threshold configuration substantially reduces N (from 63–66 to 45–48 for acceleration and velocity, and to 26–30 for displacement), yet R^2 generally improves rather than deteriorates (for example, from 0.36 to 0.55 for the P-wave window and from 0.80 to 0.83 for the S-wave window, acceleration; Tables E.4 and E.7). A smaller but cleaner ensemble therefore yields better-constrained fits, suggesting that ensemble size alone is not a reliable factor for the quality of the resulting moment scaling estimate.

Taken together, these results indicate that neither picker uniformly dominates the other. AR-AIC is more conservative and structurally constrained by the theoretical search windows, which makes it more stable. PhaseNet produces a clearer signal where it performs well, as for INT-41004391, but its reliability varies considerably with dataset and signal type, as established in Section 7.2.2. This dataset dependence is addressed directly in the comparison between INT-41004391 and IT-2009-0009 that follows in Section 7.4.

7.4 Comparative analysis: Queyras vs L’Aquila

Considered together, the results discussed in Section 7.3 suggest that nearly every difference observed between the two datasets in quality control, monotonicity failures and the dispersion of PhaseNet residuals, is consistent with the substantially greater geographic and network heterogeneity of IT-2009-0009, spanning epicentral distances of 1.80–421.40 km, against 4.80–109.50 km for INT-41004391 (Table D.2). This interpretation should, however, be stated with an explicit caveat: magnitude and network coverage are confounded in

this thesis, since the larger event is also the one observed by the more extensive network, and the present data do not allow these two factors to be disentangled (Section 7.5).

A seemingly paradoxical pattern emerges when comparing quality control outcomes to moment scaling fit quality. Despite IT-2009-0009 showing systematically worse quality control performance than INT-41004391 (higher *Peak*, *MonoP*, and *MonoS* failure rates throughout Chapter 6), it yields better-constrained moment scaling fits in the P-wave window, with higher R^2 values than INT-41004391 across signal types and configurations (Section 7.3.1). Noisier picking and cleaner fits therefore do not point in the same direction, and a plausible explanation lies in ensemble size: IT-2009-0009 reaches $N = 180$ – 186 in several configurations, against a maximum of 66 for INT-41004391, so that averaging over a substantially larger number of stations may compensate for greater per-station picking noise. This is consistent with the relationship between ensemble size and fit quality already noted in Section 7.3.5, and suggests that the trade-off observed there for PhaseNet’s filtering configurations may reflect a more general property of the ensemble-averaging procedure.

Taken together, the variability observed between the two datasets, not only in the magnitude of the scaling exponents but at times in their sign, raises a question that bears directly on the central research question of this thesis: whether the evidence for anomalous scaling gathered here reflects a genuine, physically universal property of seismic ground motion signals, or whether it is, at least in part, an artefact of the specific combination of event, network, and segmentation method examined. This question cannot be fully resolved on the basis of two events alone, and is addressed further, together with the broader limitations of the present analysis, in Section 7.5.

7.5 Limitations

The analysis presented in this thesis is based on two seismic events, and while each has been characterised in considerable detail, at the level of hundreds of station-component records and multiple segmentation configurations, the sample size at the level of distinct earthquakes remains minimal. Any generalisation from the present results to a broader claim about the statistical behaviour of seismic ground motion should be tempered accordingly.

A related limitation, already noted in Section 7.4, is that event magnitude and network coverage are confounded in this thesis: IT-2009-0009 is both the larger event and the one observed by the more extensive, more heterogeneous network, and the present data do not allow these two factors to be disentangled. Much of the difference attributed to network heterogeneity could, in principle, instead reflect a genuine magnitude dependence, or some combination of the two, and resolving this would require additional events spanning magnitude and network configuration independently.

A further limitation concerns the coda onset itself. Section 7.3.4 showed that the sign and magnitude of the scaling exponents can depend on the choice of coda onset method, and that no consistent ranking among methods holds across datasets. The coda onset time is not an objectively known quantity against which the segmentation can be validated, but itself a methodological choice that introduces variability into the downstream results. The robustness checks performed across the four coda onset methods characterise the sensitivity of the segmentation results to this choice, but do not resolve which, if any, of the four definitions most accurately reflects the true transition to scattered wave.

The finite duration of every temporal window imposes a further constraint on the moment scaling analysis. For each window, $\tau_{\max}$ is bounded by the shortest instance of that window across the stations in the spatial ensemble, so the log-log fit underlying each $\zeta(q)$ rests on a correspondingly small number of τ values. At the same time, the ensemble size N is fixed and modest at every lag, and $M_q(\tau)$ at high q is increasingly dominated by the largest values of $|\Delta x(\tau)|$ within the ensemble. For heavy-tailed signals, these extreme values can vary substantially between successive lags, making the scaling relation $M_q(\tau) \sim \tau^{\zeta(q)}$ less regular at high q and the resulting log-log fit less stable. The combination of few τ points and a less regular scaling relation at high q makes the estimated exponents $\zeta(q)$ less reliable than those at $q = 1$ or $q = 2$, and they should be interpreted with correspondingly greater caution.

The spatial ensemble averaging procedure itself rests on an assumption that is only approximately satisfied: the three components (HNE, HNN, HNZ) at each station are treated as independent realisations of the same process (Appendix A), whereas in reality they share site effects, partially correlated instrumental noise, and a common propagation history up to the station. This is likely to overstate the true number of statistically independent observations entering the ensemble average, relative to the nominal ensemble size of three components per station.

The interpretation of AR-AIC residuals is also subject to an irreducible ambiguity, already noted in Section 7.2.1: since residuals are computed relative to CRUST1.0 theoretical arrival times, an elevated residual may reflect either a picking error or an inaccuracy in the assumed crustal velocity structure, and the two contributions cannot be separated with the data available here.

PhaseNet is applied pre-trained on the INSTANCE dataset, without any fine-tuning on INT-41004391 or IT-2009-0009. While this is methodologically appropriate for assessing the picker’s reliability, and the resulting generalisation gap between the two datasets is itself one of the more informative results of this thesis (Section 7.2.2), it also means that PhaseNet’s reliability for any new event of interest is not guaranteed a priori and must be assessed case by case, rather than assumed from its performance on a previously studied event.

Taken together, these limitations bear directly on the central question raised in Section 7.4: whether the evidence for anomalous scaling gathered in this thesis reflects a genuine, physically universal property of seismic ground motion, or is, at least in part, an artefact of the specific combination of event, network, segmentation method, and picker examined.

Chapter 8

Conclusions

This thesis investigated whether seismic ground motion signals, acceleration, velocity, and displacement, recorded during two earthquakes of different magnitude, the 2009 M_w 6.1 L'Aquila earthquake and the 2024 M_w 3.8 Maritime Alps earthquake, display statistical signatures of non-equilibrium complex systems, and whether such signatures, where present, are reproducible across signal types, seismic phases, and the methodological choices involved in extracting them. To this end, each recording was segmented into temporal windows using two independent automatic phase-picking methods, AR-AIC and PhaseNet, and four coda onset definitions, and analysed through a combination of distributional fitting and moment scaling.

The evidence for anomalous diffusion in the moment scaling spectra $\zeta(q)$ is nuanced. Rather than a single, unconditional answer, this thesis finds that the reliability of anomalous scaling claims depend on the seismic phase, the dataset, the phase-picking method, and, in several cases, the coda onset definition. The P-wave window illustrates this most starkly, with the sign of the scaling exponent for the Maritime Alps dataset reversing between AR-AIC and PhaseNet; the S-wave window shows the two pickers agreeing in sign but disagreeing sharply on the strength of the evidence; and the coda window shows reasonable agreement between methods for one dataset but not the other. This conditionality is one of the central findings of the analysis: claims of anomalous scaling in seismic ground motion cannot be treated as method-independent, and the systematic comparison across picking methods, filtering configurations, and coda definitions developed in this thesis provides a framework for identifying which results are robust to these choices and which are not.

Comparing the two events directly, the larger, more distant L'Aquila dataset yields better-constrained moment scaling fits in several windows despite consistently worse phase-picking quality control outcomes, an apparent paradox plausibly explained by its substantially larger station ensemble compensating for greater per-station picking noise. At the same time, L'Aquila shows considerably more divergent behaviour across signal types than the Maritime Alps dataset, particularly for velocity and displacement, plausibly reflecting the wider range of propagation regimes captured by its more extensive and heterogeneous network. Because magnitude and network coverage are confounded in this thesis, with only two events analysed, it is not possible to attribute these differences uniquely to either factor.

This confound is the most significant limitation of the present work, together with the dependence of the results on the choice of coda onset method, the inherent ambiguity in AR-AIC residuals between picking error and crustal velocity model inaccuracy, and the

use of PhaseNet without dataset-specific fine-tuning, which introduces a generalisation gap that is itself informative but limits the picker’s reliability for events dissimilar to its training distribution.

Future work could extend this analysis to additional seismic events spanning a wider and more independent range of magnitudes and network geometries, allowing the confound between event size and station coverage identified in this thesis to be resolved. Fine-tuning PhaseNet, or comparable deep-learning pickers, on regional data more representative of the events studied here could also be expected to narrow the generalisation gap observed for the L’Aquila dataset. Finally, extending the systematic robustness framework developed in this thesis, across phase pickers, filtering configurations, and coda onset definitions, to a larger and more diverse set of events would help establish whether the handful of stable configurations identified here reflect genuine, generalisable physical signatures of anomalous diffusion in seismic ground motion, or remain specific to the two events analysed.

Glossary

***b*-value** The slope parameter in the Gutenberg–Richter frequency-magnitude relation: $\log_{10} N = a - bM$, where N is the cumulative number of earthquakes with magnitude $\geq M$. Typical values range between 0.8 and 1.2. The *b*-value characterizes the relative proportion of small to large earthquakes: higher *b* implies a greater proportion of small events. Variations in *b* have been interpreted as indicators of stress state, material heterogeneity, or tectonic regime.. 31

Aftershock Smaller earthquake following a mainshock.. 28

Anderson–Darling test A goodness-of-fit test that assesses whether a sample is drawn from a specified distribution. It is a refinement of the Kolmogorov–Smirnov test that assigns greater weight to the tails of the distribution, making it particularly sensitive to departures from the hypothesized distribution in the tail regions. The test statistic is defined as $A^2 = -n - \frac{1}{n} \sum_{i=1}^n (2i - 1) [\ln F(x_i) + \ln(1 - F(x_{n+1-i}))]$, where F is the cumulative distribution function of the hypothesized distribution and $x_1 \leq \dots \leq x_n$ are the ordered observations. Large values of A^2 indicate poor fit. Critical values depend on the hypothesized distribution and the significance level; for the Gaussian case at the 5% level, the critical value is approximately 0.787.. 69

AR-AIC method Autoregressive-Akaike Information Criterion method for automatic phase onset detection. The algorithm models the seismic signal as an autoregressive (AR) process and uses the Akaike Information Criterion (AIC) to identify abrupt changes in signal characteristics associated with P- and S-wave arrivals. The AIC function is computed over a sliding window and exhibits a minimum at the onset time. The method operates on three-component waveforms and is robust to noise and baseline variations.. 54

Arias intensity Cumulative measure of seismic energy.. 58

Baseline correction The process of removing a constant offset (DC component) from a seismic signal by subtracting its mean: $a'(t) = a(t) - \bar{a}$. Baseline correction ensures zero mean and is a prerequisite for accurate integration to velocity and displacement. In strong-motion processing, baseline correction may also involve removal of polynomial trends to compensate for low-frequency drift due to tilt or instrument response.. 49

Branching ratio In self-exciting point processes and ETAS models, the expected number of direct offspring generated by a single event, defined as $n = \int_0^\infty g(t) dt$, where $g(t)$ is the triggering kernel. The stability of the process depends on this parameter: if $n < 1$, the process is subcritical (decaying); if $n = 1$, it is critical; if $n > 1$, it

becomes supercritical (explosive). For seismicity, the branching ratio controls the proportion of triggered vs background events.. 39

Coda The prolonged oscillatory phase consisting of scattered waves generated by heterogeneities in the Earth's crust.. 51

Complementary Cumulative Distribution Function (CCDF) Defined as $P(X > x) = 1 - F(x)$ where $F(x)$ is the standard CDF. Particularly useful for identifying power-law behavior in the tails of distributions.. 70

CRUST1.0 A global crustal model with one-degree spatial resolution providing thickness, density, and seismic velocities (v_P , v_S) for each crustal layer (sediments, upper/middle/lower crust) and uppermost mantle. Used to compute station-specific theoretical P- and S-wave arrival times by extracting thickness-weighted average velocities for the layers traversed by seismic rays from the hypocenter to the surface. Reference: Laske et al., 2013.. 55

Epicenter Point on the earth's surface directly above the origin of an earthquake.. 46

ETAS model Epidemic-Type Aftershock Sequence model, a self-exciting point process developed by Ogata (1998) to describe the space-time-magnitude statistics of seismicity. The conditional intensity at time t is $\lambda(t|\mathcal{H}_t) = \mu + \sum_{t_i < t} K_0 e^{\alpha(M_i - M_0)} / (t - t_i + c)^p$, where μ is the background rate, K_0 , α , c , and p are model parameters, and M_i are magnitudes of past events. ETAS allows recursive triggering (aftershocks can generate their own aftershocks) and successfully reproduces Omori–Utsu decay, Bath's law, and magnitude-productivity scaling. Parameters are estimated by maximum likelihood from earthquake catalogs.. 38

Fault Fracture in Earth's crust where slip occurs.. 24

Focal mechanism A graphical representation of the geometry of faulting during an earthquake, derived from the analysis of P-wave first motions or seismic moment tensors. Typically displayed as a beachball diagram showing the orientation of the fault plane, slip direction, and principal stress axes. Focal mechanisms are classified into three main types: normal (extensional), reverse/thrust (compressional), and strike-slip (lateral motion).. 48

Foreshock Smaller earthquake preceding a mainshock.. 28

Gutenberg–Richter law An empirical power-law relationship describing the frequency-magnitude distribution of earthquakes: $\log_{10} N = a - bM$, where N is the cumulative number of earthquakes with magnitude $\geq M$, a quantifies the overall seismicity level, and $b \approx 1$ characterizes the relative proportion of small to large events. Proposed by Gutenberg and Richter (1944) based on California seismicity, this relation has been validated globally and is a cornerstone of seismic hazard assessment. The scale invariance implied by the power-law form motivates interpretations of seismicity as a critical phenomenon.. 32

Hill estimator An estimator for the tail index α of a power-law distribution, defined as

$$\hat{\alpha} = \left[\frac{1}{k} \sum_{i=1}^k \ln \frac{x_{(i)}}{x_{(k+1)}} \right]^{-1}$$

where $x_{(1)} \geq \dots \geq x_{(n)}$ are order statistics and k corresponds to a threshold (typically the top 10% of values).. 70

Hypocenter (focus) Spatial coordinates where seismic waves are first generated during an earthquake.. 27

Hypocentral distance The three-dimensional distance from the earthquake hypocenter (focus) to a recording station, calculated as $d_{\text{hypo}} = \sqrt{d_{\text{epi}}^2 + h^2}$, where d_{epi} is the epicentral distance and h is the hypocenter depth. Hypocentral distance represents the actual path length traveled by seismic waves and is used to compute theoretical arrival times. It differs significantly from epicentral distance for shallow events recorded at near-field stations.. 54

INSTANCE dataset Italian accelerometric dataset containing approximately 1.2 million three-component waveforms from 54,000 earthquakes recorded by 620 stations, with manually reviewed P- and S-wave onset picks from the INGV bulletin. Waveforms are preprocessed (baseline correction, resampling to 100 Hz) and supplemented with noise traces. Used to train the Italian-optimized PhaseNet model accessible via SeisBench. Reference: Michellini et al., 2021.. 64

ITACA database Italian Accelerometric Archive (<https://itaca.mi.ingv.it>), a comprehensive repository of strong-motion records from Italian seismic networks maintained by INGV. Provides uniformly processed three-component acceleration, velocity, and displacement time series with metadata including station coordinates, site classification, instrument characteristics, and event parameters. Data are stored in ASCII format with standardized headers.. 48

Laplace Distribution A continuous probability distribution with exponential tails, also known as the double exponential distribution. The PDF is $f(x; \mu, b) = \frac{1}{2b} \exp\left(-\frac{|x-\mu|}{b}\right)$, where μ is the location parameter and $b > 0$ is the scale parameter. It has heavier tails than the Gaussian but lighter than power-law distributions. The excess kurtosis is $\kappa = 3$.. 70

Lévy-stable Distribution A family of continuous probability distributions characterized by the stability property: the sum of independent stable random variables (with the same parameters) is also stable. The distribution is fully characterized by four parameters: stability index $\alpha_s \in (0, 2]$, skewness $\beta_s \in [-1, 1]$, scale $c > 0$, and location μ . The stability index α_s controls the tail behavior: smaller α_s indicates heavier tails. The Gaussian ($\alpha_s = 2$) and Cauchy ($\alpha_s = 1, \beta_s = 0$) are special cases. For $\alpha_s < 2$, the distribution has power-law tails and infinite variance. No closed-form PDF exists except for special cases, so the distribution is typically defined via its characteristic function.. 70

Magnitude Logarithmic measure of earthquake size.. 19

Maximum Likelihood Estimation A statistical method for estimating the parameters of a probability distribution by maximizing the likelihood function. Given a sample of independent observations $x_1, \dots, x_n$ and a parametric distribution $f(x; \theta)$ with parameter vector θ , the likelihood function is defined as $L(\theta) = \prod_{i=1}^n f(x_i; \theta)$. In practice, one works with the log-likelihood $\ell(\theta) = \ln L(\theta) = \sum_{i=1}^n \ln f(x_i; \theta)$ to avoid numerical underflow. The MLE estimate $\hat{\theta}$ is obtained by maximizing $\ell(\theta)$: $\hat{\theta} = \arg \max_{\theta} \ell(\theta)$. Under regularity conditions, MLE estimators are asymptotically unbiased, efficient, and normally distributed.. 70

Moment magnitude (M_w) A magnitude scale based on the seismic moment M_0 , defined as $M_w = \frac{2}{3} \log_{10} M_0 - 6.07$, where M_0 is in N·m. Unlike empirical magnitude scales (e.g., local magnitude M_L), M_w does not saturate for large earthquakes and provides a physically meaningful measure of earthquake size directly related to fault slip and rupture area. It is the standard magnitude scale used in modern seismology.. 30

Moment order q The order of the statistical moment computed from the absolute value of increments: $M_q(\tau) = \langle |\Delta x(\tau)|^q \rangle$. Lower values of q (e.g., $q < 2$) emphasize the bulk of the distribution, while higher values ($q > 2$) are increasingly sensitive to rare large fluctuations in the tails. The dependence of the scaling exponent $\zeta(q)$ on q distinguishes normal from anomalous diffusion.. 42

Network code A two-letter identifier assigned to a seismic network by the International Federation of Digital Seismograph Networks (FDSN). Examples include IV (Italian National Seismic Network), FR (French National Seismic Network), and GU (University of Genoa). Network codes are part of the standardized station naming convention: `Network.Station.Location.Channel`.. 48

ObsPy An open-source Python framework for seismological data processing and analysis. ObsPy provides data structures (`Trace`, `Stream`), I/O routines for standard seismological formats (SAC, miniSEED, GSE2), signal processing functions (filtering, integration, resampling), and implementations of seismological algorithms (AR-AIC phase picking, STA/LTA triggering). Reference: Beyreuther et al., 2010.. 56

Onset time Arrival time of seismic wave at recording station.. 54

Origin time The absolute time at which an earthquake rupture initiates at the hypocenter, conventionally denoted t_0 . In signal-relative coordinates, the origin time is expressed as the time interval between the event occurrence and the start of the recording: $t_0 = t_{\text{event}} - t_{\text{record_start}}$. The origin time is used to compute theoretical phase arrival times: $t_P = t_0 + d/v_P$.. 48

P-waves Primary (compressional) seismic wave characterized by particle motion parallel to the propagation direction. It is the fastest body wave and the first to arrive at seismic stations.. 51

Peak Ground Acceleration (PGA) The maximum absolute value of ground acceleration recorded during an earthquake: $\text{PGA} = \max_t |a(t)|$. PGA is typically measured in cm/s^2 (or as a fraction of gravitational acceleration g) and represents the most severe instantaneous shaking experienced at a location. It is widely used in earthquake engineering for seismic design and hazard assessment. In typical seismograms,

PGA occurs within the S-wave window due to the larger amplitudes of shear waves.. 62

Peak Ground Displacement (PGD) The maximum absolute value of ground displacement recorded during an earthquake: $\text{PGD} = \max_t |d(t)|$. PGD is typically measured in cm and is obtained by double integration of the acceleration time series. It is particularly relevant for long-period structures and near-source ground motions with permanent static offsets.. 62

Peak Ground Velocity (PGV) The maximum absolute value of ground velocity recorded during an earthquake: $\text{PGV} = \max_t |v(t)|$. PGV is typically measured in cm/s and is obtained by integrating the acceleration time series. It is strongly correlated with structural damage and is used in seismic hazard assessment and ground motion prediction equations.. 62

Phase picking Detection of P and S-waves onset times.. 55

Power-law Distribution A distribution whose tail decays as $P(X > x) \sim x^{-\alpha}$ for large x , where $\alpha > 0$ is the tail exponent. Characteristic of heavy-tailed phenomena. For $\alpha \leq 2$ the variance is infinite; for $\alpha \leq 1$ the mean is infinite.. 70

Quality control (QC) A systematic set of procedures to validate the reliability of seismic data processing results. In phase detection, QC includes checks for arrival time monotonicity with distance, signal-to-noise ratio (SNR) thresholds, physical consistency (e.g., $t_S > t_P$), and peak-timing validation (e.g., PGA within S-wave window). Stations or picks failing QC criteria are flagged for manual review or excluded from subsequent analysis.. 54

Resampling The process of changing the sampling rate of a digital signal by converting it to a continuous representation (interpolation) and then re-discretizing at a different rate. In seismology, resampling is often required to match the input requirements of machine learning models (e.g., PhaseNet expects 100 Hz) or to align signals recorded at different sampling rates. Proper resampling includes anti-aliasing filtering to prevent spectral aliasing artifacts when downsampling.. 65

S-waves Shear waves, a type of body wave that propagates through solids, characterized by particle motion that is perpendicular to the direction of wave propagation. S waves do not travel through fluids and are associated with no compression or expansion of the medium.. 51

Sampling rate Number of samples per second (Hz).. 63

Search window Time interval for phase detection.. 54, 56

SeisBench A standardized Python framework providing unified access to seismological machine learning models and benchmark datasets. Facilitates loading pre-trained models (e.g., PhaseNet, EQTransformer) and applying them to ObsPy data streams with minimal preprocessing. Reference: Woollam et al., 2022.. 65

Seismic moment (M_0) A scalar measure of earthquake size defined as $M_0 = \mu \bar{u} A$, where μ is the shear modulus of the surrounding rock, $\bar{u}$ is the average slip on the

fault, and A is the rupture area. The seismic moment has units of N·m (or dyne·cm in older literature) and represents the total moment released during rupture. It is estimated from seismic waveforms or geodetic observations and forms the basis for moment magnitude M_w .. 27

Seismogenic zone Depth range where earthquakes nucleate.. 28

Stress drop The difference between the initial shear stress on a fault before rupture and the final stress after rupture: $\Delta\sigma = \sigma_1 - \sigma_2$. Static stress drop is typically measured in MPa and provides a measure of the efficiency of stress release during an earthquake. Stress drop is approximately scale-invariant across a wide range of earthquake magnitudes, with typical values of 1–10 MPa for crustal earthquakes.. 26

Student-t Distribution A continuous probability distribution with heavier tails than the Gaussian. The PDF is $f(x; \nu) = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi} \Gamma(\frac{\nu}{2})} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}$, where $\nu > 0$ is the degrees of freedom parameter. For $\nu = 1$, it reduces to the Cauchy distribution (undefined variance); as $\nu \rightarrow \infty$, it converges to the Gaussian. Lower ν indicates heavier tails. The variance exists only for $\nu > 2$.. 70

Trace Synonymous with waveform in seismological terminology. A trace refers to a single-channel recording of ground motion, typically corresponding to one component (Z, N, or E) of a three-component seismogram. In ObsPy and seismological databases, a **Trace** object encapsulates the time series data together with metadata (station, channel, sampling rate, start time).. 30

Waveform The time series representation of seismic ground motion recorded by a seismometer or accelerometer. A waveform is characterized by its amplitude, frequency content, duration, and phase information. In digital seismology, waveforms are stored as discrete time series sampled at a fixed rate (e.g., 100 Hz, 200 Hz) and are the primary data used for phase picking, source analysis, and statistical characterization.. 27

Appendix A

Implementation

A.1 Overview

The computational pipeline is structured as a sequence of Jupyter notebooks (01a, 01b, 02, 03a, 03b, 04a) supported by modular Python libraries organised in a `src/` directory. Each notebook corresponds to a major analysis stage, from raw data ingestion to statistical modelling, while the supporting modules encapsulate reusable functions for preprocessing, phase detection, windowing, and visualisation. The design follows a clear separation of concerns: notebooks orchestrate the workflow and document decisions, while modules provide tested, general-purpose implementations. The repository is hosted on GitHub (<https://github.com/giulianaparadiso99/tesi-seismic-analysis>) for version control and reproducibility.

A.2 Data and storage

The ITACA dataset arrives as two separate `.zip` archives, one per seismic event: 66 `.ASC` files (22 stations $\times$ 3 components) for INT-41004391 (Maritime Alps, Queyras), and 192 `.ASC` files (66 stations $\times$ 3 components) for IT-2009-0009 (L'Aquila), the latter prior to the exclusion of stations GSG and TTS described in Chapter 3. Each file consists of a 64-row header (key-value pairs) followed by signal data in a two-column format. All data is loaded into pandas DataFrames and stored in the parquet format, which offers efficient columnar compression and preserves data types across sessions. Processed signals are stored separately from metadata to avoid redundant loading. The data organisation follows a strict convention: raw data in `data/raw/`, preprocessed outputs in `data/processed/{notebook}/{data_type}/` (where `data_type` is `acceleration`, `velocity`, or `displacement`), and figures in `figures/{notebook}/{data_type}/`.

A.3 Core modules

The pipeline's functionality is implemented across a set of dedicated modules, each responsible for a single stage of the analysis: data loading and cleaning, theoretical arrival time estimation and search window definition, phase picking (AR-AIC and PhaseNet), coda onset detection, signal windowing and validation, moment scaling analysis, PDF analysis, and visualisation. Each is described in turn below.

A.3.1 Data loading and cleaning

The `io.py` module provides functions to parse the `.ASC` archive structure: `build_metadata()` extracts header fields into a long-format DataFrame (one row per file), while `build_signals()` extracts numerical time series, also in long format (one row per sample, columns: `file`, `sample`, signal value). The signal column name is dynamically inferred from each file’s `DATA_TYPE` header field, enabling generalisation across acceleration, velocity, and displacement signals without hardcoding column names.

Metadata preprocessing is handled by `cleaning_metadata.py`, which implements a sequential pipeline: replace missing values, drop uninformative columns (e.g., `USER1-5`, `HEADER_FORMAT`), convert numeric and datetime fields, normalise strings, remove duplicates, and compute sampling rate from `NDATA` and `DURATION_S`. Signal preprocessing is managed by `cleaning_signals.py`, which provides independent control over three operations: length filtering (retain only signals with $\geq N$ samples), baseline correction (subtract per-signal mean to ensure zero baseline), and (divide by per-signal standard deviation, required only for PDF analysis). The key design choice is to preserve physical units by default: normalisation is opt-in, not automatic.

A.3.2 Theoretical arrivals and search windows

The `search_windows.py` module computes theoretical P- and S-wave arrival times using CRUST1.0 crustal velocity profiles (Laske et al., 2013). For each station, the module queries CRUST1.0 at the station coordinates and extracts depth-weighted average velocities v_P and v_S for all layers traversed by a vertical ray from the hypocenter to the surface. Hypocentral distance is calculated as $d_{\text{hypo}} = \sqrt{d_{\text{epi}}^2 + h^2}$, where h is the event depth. Theoretical arrival times are then computed as $t_P = t_{\text{origin}} + d_{\text{hypo}}/v_P$ and $t_S = t_{\text{origin}} + d_{\text{hypo}}/v_S$, where t_{origin} is the event time relative to the recording start.

Search windows for phase picking are defined adaptively around theoretical arrivals. The `calculate_adaptive_windows()` function implements station-specific logic: for both P- and S-waves, a symmetric window is used, defined as $[t^{\text{theo}} - \Delta t, t^{\text{theo}} + \Delta t]$ around the corresponding theoretical arrival; to prevent ambiguity between the two phases, the S-wave window is constrained to start after the end of the P-wave window, separated by a minimum gap δ . All onset times and window boundaries are stored in dual representation (samples and seconds) to avoid cumulative rounding errors.

A.3.3 Phase picking

Two parallel phase detection workflows are implemented to assess the robustness of moment scaling exponents to picker methodology.

AR-AIC method

Phase onset detection via the AR-AIC (Autoregressive-Akaike Information Criterion) method is wrapped by ObsPy’s `ar_pick()` function in `ar_pick.py`. The method operates on three-component signals (Z, N, E) simultaneously and identifies sharp transitions in signal characteristics by detecting minima in the Akaike Information Criterion computed from autoregressive model residuals. The picker is applied within the search windows defined by theoretical arrivals, returning sample indices that are immediately converted into the dual sample-and-second representation used throughout the pipeline.

PhaseNet method

An alternative deep-learning-based workflow is provided via PhaseNet (Zhu and Beroza, 2019), implemented in `phasenet_utils.py`. PhaseNet is a convolutional neural network pre-trained on the INSTANCE dataset (Michellini et al., 2021) and accessed via the SeisBench framework (Woollam et al., 2022). Unlike AR-AIC, PhaseNet does not require theoretical arrival times or velocity models—it operates directly on waveform features learned from training data. The pipeline resamples signals from 200 Hz to 100 Hz (PhaseNet’s native rate), runs CNN inference, and converts predicted onset probabilities to sample indices. Both pickers produce onset time estimates in identical dual sample-and-second formats, feeding the same downstream windowing and moment scaling pipeline and enabling direct comparison of resulting $\zeta(q)$ spectra.

Coda onset detection

Coda onset detection is performed independently on each component using three independent methods, combined into a median ensemble estimate, all implemented in `coda_detection.py`:

1. **Rautian criterion** (Rautian and Khalturin, 1978): Coda begins at $t_{\text{coda}} = 2t_S - t_{\text{origin}}$, where t_0 is the event origin time. This theoretical definition assumes single-scattering in a homogeneous medium.
2. **Arias Intensity D5-95**: Coda onset is defined as the time when cumulative Arias Intensity reaches 95% of its total value (D95 parameter from ESM flatfile convention). This energy-based criterion marks the end of significant ground motion.
3. **Envelope decay**: Coda begins when the signal envelope drops below 30% of the peak envelope after S-arrival. This amplitude-based criterion captures the transition to scattered wave dominance.
4. **Median**: Robust combination computed as the median of the three methods above, reducing sensitivity to outliers.

All four methods are applied systematically to quantify the sensitivity of moment scaling exponents to coda window definition.

Coda end detection

End-of-coda times, marking the transition to the post-event window, are computed independently of onset detection using two complementary methods implemented in `coda_detection.py`.

For the Rautian, Envelope, and Median coda definitions, `compute_coda_end()` applies an energy decay criterion: the Hilbert envelope of the signal following coda onset is smoothed with a moving average, and a reference peak amplitude is taken from the smoothed envelope within the first 5 s after onset. Coda end is identified as the first point at which the envelope remains below 10% of this reference peak for a sustained duration of 2 s, with a minimum coda duration of 1 s enforced as a safety constraint; if no stable threshold crossing is found, the coda is taken to extend to the end of the recording, and no post-event window is defined for that component.

For the Arias coda definition, `compute_coda_end_arias()` mirrors the D5–95 onset criterion, using a higher cumulative-energy threshold (default D99.5, following the ESM flatfile convention ([lanzano2019esm](#))) computed on the same integration window used for onset detection, with the same 1 s minimum coda duration enforced.

The orchestration function `add_coda_end_to_dataframe()` applies the appropriate method to each coda definition, storing results in the same dual sample-and-second representation used throughout the pipeline, and defines the boundary between the coda window $[t_{\text{coda}}, t_{\text{coda_end}})$ and the post-event window $[t_{\text{coda_end}}, t_{\text{end}}]$.

A.3.4 Signal windowing and validation

The `window_segmentation.py` module segments each component signal into up to five temporal windows: pre-event (t_{start} to t_P), P-wave (t_P to t_S), S-wave (t_S to t_{coda}), coda (t_{coda} to $t_{\text{coda_end}}$), and post-event ($t_{\text{coda_end}}$ to end, when a stable coda-end crossing is found). The core function `segment_signal_into_windows()` accepts onset times in either sample or second units and returns a dictionary containing signal arrays, start/end indices, and durations for each window. The batch function `segment_all_signals()` applies this to all stations and components, producing a nested dictionary structure `{station:{component:{window_name:data}}}` optimised for downstream moment scaling analysis.

Quality control is enforced by `window_validation.py`, which implements three validation checks: (1) peak-in-S-wave verification (confirms that the peak ground-motion value, PGA, PGV, or PGD depending on signal type, occurs within the S-wave window, as expected physically), (2) arrival time monotonicity (verifies that detected onsets follow hypocentral distance ordering), and (3) signal-to-noise ratio assessment ($\text{SNR} \geq 3$ for reliable phase identification). Stations or components failing these checks are flagged but not automatically excluded, as some violations may reflect legitimate complexities (e.g., strong lateral velocity variations or unusual source radiation patterns) rather than picker errors.

A.3.5 Moment scaling analysis

Ensemble-averaged moment scaling is computed by `signals_scaling_spatial.py`. For each of the (up to) five temporal windows, the module sets the increment reference time to the window’s own start boundary (t_P for the P-wave window, t_S for the S-wave window, t_{coda} for the coda window, and $t_{\text{coda_end}}$ for the post-event window, denoted generically as t_0), which differs in absolute time across stations due to varying arrival times. Increments are computed as $\Delta x(\tau) = x(t_w + \tau) - x(t_w)$ for a logarithmically spaced vector of time lags τ . Moments are defined as $M_q(\tau) = |\Delta x(\tau)|^q$ for each signal individually, then averaged across the spatial ensemble (all stations and components). Scaling exponents $\zeta(q)$ are extracted by fitting $\log \langle M_q(\tau) \rangle$ versus $\log \tau$ via linear regression in the scaling regime.

A critical implementation detail is that τ values are generated directly in sample space using `np.logspace()`, then converted to seconds only for output. This avoids rounding artifacts that would arise from the alternative workflow (generate in seconds $\rightarrow$ convert to samples $\rightarrow$ apply `np.unique()` $\rightarrow$ convert back to seconds), which breaks the logarithmic distribution. The maximum lag τ_{max} is set to the minimum window duration across all stations to ensure complete ensemble coverage (all stations contribute to all τ values). The module also supports exclusion of specific components (e.g., vertical channels) to test

the impact of wave polarisation on scaling behaviour.

A.3.6 PDF analysis

Probability density function analysis is implemented in `signals_pdf.py`. The module provides two complementary analyses: (1) Gaussian fit assessment via Anderson–Darling tests, kurtosis, and skewness, with visual Q–Q plots; (2) heavy-tail characterisation via comparison with Gaussian, Laplace, Student- t , and Lévy stable distributions using Akaike Information Criterion (AIC) for model selection. Both pipelines process individual signals (one plot per station/component) and generate summary visualisations aggregating results across all signals. The analysis operates on normalised signals (zero mean, unit variance) to enable direct comparison of distributional shapes across stations.

A.3.7 Visualisation

Plotting functions are organised thematically in separate modules: `plots_metadata.py` (geographic station maps, PGA distributions, correlation matrices), `plots_signals.py` (time series, spectrograms), `plots_segmentation.py` (windowed signals with onset markers, phase arrival overlays), and `plot_moment_scaling.py` (scaling curves $\langle M_q(\tau) \rangle$ versus τ , exponent spectra $\zeta(q)$ versus q). All modules share a consistent visual style defined in `plot_settings.py`, which configures matplotlib defaults (font sizes, grid appearance, spine visibility, colormap) to ensure uniformity across all figures in the thesis. The default colormap is `inferno`, sampled at fixed intervals to generate a five-colour palette used to distinguish components (HNE, HNN, HNZ) or temporal windows (pre-event, P, S, coda, post-event).

A.4 Workflow summary

The analysis comprises two parallel workflows:

A.4.1 Primary workflow: moment scaling

1. **Notebook 01a:** Load and preprocess metadata. Output: `metadata_clean_*.parquet`.
2. **Notebook 01b:** Load signals, apply baseline correction (without normalisation to preserve physical units). Output: `*_preprocessed_scaling.parquet`.
3. **Notebook 03a OR 03b:** Run phase detection (AR-AIC or PhaseNet) within search windows (03a only), detect coda onset and coda end times (three independent methods, combined into a median ensemble estimate), validate picks via quality control (peak timing, monotonicity, SNR), segment signals into up to five windows. Output: `df_full_*_{picker}.parquet`, `windowed_*_{coda_method}_{picker}.pkl`.
4. **Notebook 04a:** Compute spatial ensemble moments for each window, extract scaling exponents $\zeta(q)$, compare across phases and signal types, generate scaling plots. Output: `exponents_*.parquet`, `figures/04a_moment_scaling_spatial/`.
5. **Notebook 05:** Run sensitivity analysis: perturb P/S picks, recompute coda onsets, coda end, and windows, recompute $\zeta(q)$, compute RMSE/MAE/correlation metrics, run Monte Carlo simulation. Output: `sensitivity_results_*.parquet`.

Quality control checks are applied after phase detection but do not automatically exclude stations; failures are flagged for case-by-case review. The moment scaling workflow is repeated for:

- 3 signal types: acceleration, velocity, displacement
- 2 phase pickers: AR-AIC, PhaseNet
- 4 coda methods: Rautian, Arias, Envelope, Median

Total: $3 \times 2 \times 4 = 24$ complete analyses. Each run takes approximately 15–20 minutes on a standard laptop (8 GB RAM).

A.4.2 Secondary workflow: PDF analysis

1. **Notebook 01a:** (same as above)
2. **Notebook 01b:** Load signals, apply baseline correction and normalisation (divide by per-signal standard deviation). Output: `*_preprocessed_pdf.parquet`.
3. **Notebook 02:** Compute empirical PDFs, test Gaussianity via Anderson–Darling, fit heavy-tailed distributions (Gaussian, Laplace, Student- t , Lévy stable), select best model via AIC, generate diagnostic plots. Output: `gaussian_fit_results.parquet`, `heavy_tail_results.parquet`, `figures/02_pdf_analysis/`.

The PDF workflow is independent of phase detection and is performed once per signal type (3 runs total).

A.5 Output organisation

The pipeline generates approximately 1000–1500 PDF figures organised by analysis stage:

- `figures/01a_metadata/`: Station maps, catalog statistics, epicentral distance distributions
- `figures/01b_signals/`: Signal examples, baseline correction diagnostics, integration validation
- `figures/02_pdf_analysis/`: Empirical PDFs, Gaussian fits, Q–Q plots, heavy-tail model comparison
- `figures/03a_phase_identification_ar_pick/`: Onset detection plots, theoretical vs. detected arrivals, residuals vs. distance, coda method comparison, quality control summaries
- `figures/03b_phase_identification_phasenet/`: Same as 03a but for PhaseNet picks
- `figures/04a_moment_scaling_spatial/`: Scaling curves $M_q(\tau)$, exponent spectra $\zeta(q)$ vs. q , comparison across windows and signal types

Figures are saved as vector PDFs for publication quality. A complete catalog is maintained in `figures/README.md`.

Processed data outputs (Parquet and Pickle files) total approximately 500 MB and are organised in `data/processed/{notebook}/{data_type}/` with subdirectories for each phase picker and coda method.

A.6 Key design principles

- **Generalisation:** All preprocessing and analysis functions accept a `signal_column` parameter, enabling switching between acceleration, velocity, and displacement without hardcoding column names. Configuration is centralised at the top of each notebook (`DATA_TYPE = 'acceleration'`).
- **Dual representation:** Onset times and window boundaries are stored in both sample (integer) and second (float) units to preserve precision and facilitate physical interpretation. Legacy columns point to the seconds version for backward compatibility.
- **Physical units by default:** Normalisation is opt-in, not automatic. Moment scaling analysis operates on raw signals to preserve the physical interpretation of scaling exponents. Only PDF analysis requires normalised signals.
- **Validation over exclusion:** Quality control flags problematic stations but does not automatically exclude them, allowing case-by-case investigation of violations. This design choice acknowledges that some QC failures may reflect genuine physical complexity rather than data quality issues.
- **Modularity:** Each module has a single, well-defined responsibility. Complex workflows are built by composing simple functions. For example, `analyze_all_windows()` orchestrates moment computation, ensemble averaging, and scaling exponent fitting by calling lower-level functions in sequence.
- **Spatial ensemble averaging:** Moments are averaged across stations and components (spatial ensemble), not over time within a single trace (temporal ensemble). This choice reflects the assumption that different components capture independent realisations of the same stochastic process, justified by different wave polarisations.

A.7 Software dependencies

The analysis relies on standard scientific Python libraries: NumPy (array operations), pandas (data manipulation), SciPy (statistical tests, signal processing, heavy-tail fitting), matplotlib (plotting), ObsPy (seismological data handling, AR-AIC picker), and SeisBench (PhaseNet inference, PyTorch backend). File I/O uses PyArrow (Parquet format) for efficient columnar storage. Geospatial operations use Cartopy (map projections) and Contextily (basemap tiles). The CRUST1.0 crustal model is accessed via the official `crust1.py` module provided with the model distribution. Statistical visualisation uses Seaborn and AdjustText (label placement). All code runs on Python 3.9+ and has been tested on Python 3.10, 3.11, and 3.12.

A.8 Reproducibility

Complete replication requires:

1. Cloning the GitHub repository:
`https://github.com/giulianaparadiso99/tesi-seismic-analysis`
2. Installing dependencies: `pip install -r requirements.txt` or `conda env create -f environment.yml`
3. Downloading the ITACA dataset (66 .ASC files) and placing it in `data/raw/query.zip`
4. Downloading CRUST1.0 from <https://igppweb.ucsd.edu/~gabi/crust1.html> and placing the model files in `data/Crust1.0/`
5. Executing notebooks sequentially: 01a $\rightarrow$ 01b $\rightarrow$ 03a (or 03b) $\rightarrow$ 04a. Optional: 02 for PDF analysis.

Intermediate outputs (Parquet files) can be shared separately to bypass early stages. All random number generation (e.g., bootstrap resampling in uncertainty estimation) is seeded for deterministic results. The complete pipeline (24 moment scaling runs + 3 PDF runs) takes approximately 8–10 hours on a standard laptop. PhaseNet inference can be accelerated 2 $\times$ with GPU support (CUDA-enabled PyTorch installation).

A detailed execution guide, including expected runtimes, troubleshooting tips, and configuration options, is maintained in `notebooks/README.md`.

Appendix B

Metadata Fields Description

The following table describes all 64 metadata fields present in the ASCII file headers for the IT-2009-0009 dataset. While shown for acceleration data, the structure is identical for velocity and displacement, with only the peak ground motion fields varying: PGA_CM/S^2 , PGV_CM/S or PGD_CM , and corresponding time fields.

Table B.1: Summary of header fields and their description.

Field	Description
EVENT NAME	Event name
EVENT ID	Unique event identifier
EVENT DATE	Event date (YYYYMMDD)
YYYYMMDD	
EVENT TIME	UTC event time
HHMMSS	
EVENT LATITUDE	Hypocentral latitude
DEGREE	
EVENT	Hypocentral longitude
LONGITUDE	
DEGREE	
EVENT DEPTH KM	Hypocentral depth (km)
HYPOCENTER	Source of hypocentral location
REFERENCE	
MAGNITUDE W	Moment magnitude M_w
MAGNITUDE W	Source of M_w estimate
REFERENCE	
MAGNITUDE L	Local magnitude M_L
MAGNITUDE L	Source of M_L estimate
REFERENCE	
FOCAL	Focal mechanism (fault type)
MECHANISM	
NETWORK	Network code
STATION CODE	Station identifier
STATION NAME	Station name

Continued on next page

Table B.1 – continued from previous page

Field	Description
STATION LATITUDE DEGREE	Station latitude
STATION LONGITUDE DEGREE	Station longitude
STATION ELEVATION M	Station elevation (m a.s.l.)
LOCATION	Sub-location code
SENSOR DEPTH M	Sensor depth below ground (m)
VS30 M/S	Average shear-wave velocity in top 30 m
SITE CLASSIFICATION EC8	EC8 site class
MORPHOLOGIC CLASSIFICATION	Morphological site classification
EPICENTRAL DISTANCE KM	Epicentral distance (km)
EARTHQUAKE BACKAZIMUTH DEGREE	Back-azimuth (°)
DATE TIME FIRST SAMPLE YYYYMMDD HHMMSS	Timestamp of first sample
DATE TIME FIRST SAMPLE PRECISION	Precision of first sample timestamp
SAMPLING INTERVAL S	Sampling interval (s)
NDATA	Number of recorded samples
DURATION S	Recording duration (s)
STREAM	Channel code (orientation)
UNITS	Data units (cm/s ²)
INSTRUMENT	Instrument type code
INSTRUMENT ANALOG/DIGITAL	Analog or digital
INSTRUMENTAL FREQUENCY HZ	Natural frequency (Hz)
INSTRUMENTAL DAMPING	Critical damping ratio
FULL SCALE G	Full-scale range (g)
N BIT DIGITAL CONVERTER	ADC resolution (bit)

Continued on next page

Table B.1 – continued from previous page

Field	Description
PGA CM/S ²	Peak Ground Acceleration (cm/s ²)
TIME PGA S	Time of PGA from start (s)
BASELINE CORRECTION	Baseline correction type
FILTER TYPE	Filter type (Butterworth)
FILTER ORDER	Filter order
LOW CUT FREQUENCY HZ	Low-cut frequency (Hz)
HIGH CUT FREQUENCY HZ	High-cut frequency (Hz)
LATE/NORMAL TRIGGERED	Trigger type (NT/LT)
DATABASE VERSION	Database version
HEADER FORMAT	Header format version
DATA TYPE	Data type
PROCESSING	Processing information
DATA TIMESTAMP YYYYMMDD HHMMSS	Data timestamp
DATA LICENSE	Data license
DATA CITATION	Bibliographic citation
DATA CREATOR	Data creator
ORIGINAL DATA MEDIATOR CITATION	Citation for the original data mediator
ORIGINAL DATA MEDIATOR	Original data mediator
ORIGINAL DATA CREATOR CITATION	Citation for the original data creator
ORIGINAL DATA CREATOR	Original data creator
USER1	Free annotation field
USER2	Free annotation field
USER3	Free annotation field
USER4	Free annotation field
USER5	Free annotation field

Table B.2: Constant metadata fields across all records in the IT-2009-0009 dataset.

Field	Value
EVENT NAME	None

Continued on next page

Table B.2 – continued from previous page

Field	Value
EVENT ID	IT-2009-0009
EVENT DATE	20090406
YYYYMMDD	
EVENT TIME HHMMSS	013240
EVENT LATITUDE	42.3420
DEGREE	
EVENT LONGITUDE	13.3800
DEGREE	
EVENT DEPTH KM	8.3
HYPOCENTER	INGV-CNT_Seismic_Bulletin
REFERENCE	
MAGNITUDE W	6.1
MAGNITUDE W	INGV-CNT_Seismic_Bulletin
REFERENCE	
MAGNITUDE L	
MAGNITUDE L	
REFERENCE	
FOCAL MECHANISM	Normal faulting
SENSOR DEPTH M	0.0
SAMPLING INTERVAL S	0.005000
UNITS	cm/s ²
INSTRUMENT	D
ANALOG/DIGITAL	
BASELINE	BASELINE REMOVED
CORRECTION	
FILTER TYPE	BUTTERWORTH
FILTER ORDER	2
LATE/NORMAL	NT
TRIGGERED	
DATABASE VERSION	0.5
HEADER FORMAT	DYNA 1.2
DATA TYPE	ACCELERATION
PROCESSING	manual (Paolucci et al., 2011)
DATA CITATION	Felicetta C., Russo E., D Amico M., Sgobba S., ...
DATA CREATOR	ITACA working group
USER1	
USER2	
USER3	
USER4	
USER5	

Table B.3: Constant metadata fields across all records in the INT-41004391 dataset.

Field	Value
EVENT NAME	None
EVENT ID	INT-41004391
EVENT DATE YYYYMMDD	20241209
EVENT TIME HHMMSS	104121
EVENT LATITUDE DEGREE	44.5127
EVENT LONGITUDE DEGREE	6.8533
EVENT DEPTH KM	10.4
HYPOCENTER REFERENCE	INGV-BULLETIN
MAGNITUDE W	3.8
MAGNITUDE W REFERENCE	INGV-TERREMOTI
MAGNITUDE L MAGNITUDE L REFERENCE	
FOCAL MECHANISM	Strike-slip faulting
MORPHOLOGIC CLASSIFICATION	
DATE TIME FIRST	milliseconds
SAMPLE PRECISION	
SAMPLING INTERVAL S	0.005000
UNITS	cm/s ²
INSTRUMENT	D
ANALOG/DIGITAL	
INSTRUMENTAL FREQUENCY HZ INSTRUMENTAL DAMPING	
FULL SCALE G	
N BIT DIGITAL CONVERTER	
BASELINE	BASELINE REMOVED
CORRECTION	
FILTER TYPE	BUTTERWORTH
FILTER ORDER	2
LATE/NORMAL	NT
TRIGGERED	
DATABASE VERSION	0.5
HEADER FORMAT	DYNA 1.2
DATA TYPE	ACCELERATION

Continued on next page

Table B.3 – continued from previous page

Field	Value
PROCESSING	manual (Paolucci et al., 2011)
DATA LICENSE	D (network default license)
DATA CITATION	Felicetta C., Russo E., D Amico M., Sgobba S., ...
DATA CREATOR	ITACA working group
USER1	
USER2	
USER3	
USER4	
USER5	

Appendix C

Heavy-tail assessment — full results

Table C.1 reports the full heavy-tail assessment results for all 66 signals. Columns are: station code, stream component, AIC values for the Lévy-stable and Student- t fits (the two competitive models; Gaussian and Laplace never win), the best-fit model selected by minimum AIC, the Student- t degrees of freedom ν estimated by MLE, and the power-law exponent $\hat{\alpha}$ estimated via the Hill estimator on the top 10% of $|\hat{a}|$ values.

Table C.1: Heavy-tail assessment results for all 66 normalised signals. ν : Student- t degrees of freedom (MLE). $\hat{\alpha}$: Hill power-law exponent (top 10% tail). Lower ν and lower $\hat{\alpha}$ both indicate heavier tails.

Station	Stream	AIC _{Lévy}	AIC _{St-t}	Best fit	ν	$\hat{\alpha}$
EILF	HNE	52746.79	52804.13	Lévy-stable	1.2651	1.1420
EILF	HNN	33018.55	33000.83	Student- t	0.9964	1.1200
EILF	HNZ	67581.97	67682.48	Lévy-stable	1.3937	1.1692
ESCA	HNE	-42357.24	-41994.24	Lévy-stable	0.6018	1.0875
ESCA	HNN	-36218.17	-35903.50	Lévy-stable	0.6116	1.0638
ESCA	HNZ	-28829.31	-28865.09	Student- t	0.6807	1.0188
ISO	HNE	-100096.19	-100500.55	Student- t	0.7025	0.7322
ISO	HNN	-88301.55	-88695.53	Student- t	0.7135	0.7902
ISO	HNZ	-112240.15	-112766.34	Student- t	0.7140	0.6845
MFC	HNE	-28590.05	-27107.33	Lévy-stable	0.5136	1.1546
MFC	HNN	-30127.91	-27924.45	Lévy-stable	0.4746	1.1862
MFC	HNZ	-34720.33	-34194.08	Lévy-stable	0.5856	1.0398
MON	HNE	28511.55	28854.91	Lévy-stable	0.7863	1.3009
MON	HNN	28465.79	28766.57	Lévy-stable	0.7887	1.2439
MON	HNZ	28043.26	28395.29	Lévy-stable	0.7534	1.3381
MVIF	HNE	-26587.26	-26262.05	Lévy-stable	0.6316	1.0716
MVIF	HNN	-28786.21	-28315.44	Lévy-stable	0.6009	1.0795
MVIF	HNZ	-6976.67	-7026.79	Student- t	0.7638	1.0757
OGAG	HNE	-75412.70	-72918.33	Lévy-stable	0.4380	0.9836
OGAG	HNN	-84362.84	-82556.04	Lévy-stable	0.4712	0.8786
OGAG	HNZ	-53271.95	-53266.94	Lévy-stable	0.6641	0.8774
OGDI	HNE	3011.72	3347.14	Lévy-stable	0.7040	1.1367
OGDI	HNN	-10730.25	-10107.05	Lévy-stable	0.6178	1.1831

Continued on next page

Table C.1 – continued

Station	Stream	$AIC_{\text{Lévy}}$	$AIC_{\text{St-}t}$	Best fit	ν	$\hat{\alpha}$
OGDI	HNZ	20286.18	20526.75	Lévy-stable	0.8157	1.1538
OGMO	HNE	−19293.32	−18904.32	Lévy-stable	0.6669	1.0613
OGMO	HNN	−14090.18	−13755.64	Lévy-stable	0.6876	1.0578
OGMO	HNZ	−28194.45	−27702.66	Lévy-stable	0.6763	1.0049
OGS3	HNE	−50144.18	−49245.71	Lévy-stable	0.5445	1.0194
OGS3	HNN	−47684.88	−46738.32	Lévy-stable	0.5544	1.0169
OGS3	HNZ	−47742.70	−47399.27	Lévy-stable	0.5833	0.9964
OGSA	HNE	23142.29	23069.56	Student- t	1.1247	1.0319
OGSA	HNN	−21854.99	−21955.92	Student- t	0.9578	0.8146
OGSA	HNZ	−41237.36	−41406.32	Student- t	0.7911	0.8235
REVF	HN1	12112.65	12294.20	Lévy-stable	0.7495	1.1415
REVF	HN2	12027.63	12261.14	Lévy-stable	0.7384	1.2397
REVF	HNZ	17288.45	17557.92	Lévy-stable	0.7366	1.2519
SAOF	HNE	−16127.24	−16125.10	Lévy-stable	0.7658	1.1022
SAOF	HNN	2975.98	3015.17	Lévy-stable	0.8311	1.1597
SAOF	HNZ	−6309.53	−6406.94	Student- t	0.9255	1.0490
SPIF	HNE	−104665.74	−101808.21	Lévy-stable	0.5276	0.8547
SPIF	HNN	−93730.07	−92409.98	Lévy-stable	0.6220	0.8057
SPIF	HNZ	−92160.12	−90587.92	Lévy-stable	0.6032	0.7964
SURF	HNE	−15500.85	−13425.15	Lévy-stable	0.2274	0.7977
SURF	HNN	−18092.96	−16261.37	Lévy-stable	0.1960	0.7155
SURF	HNZ	−18785.96	−17089.34	Lévy-stable	0.2045	0.7664
BHB	HNE	2785.55	5667.11	Lévy-stable	0.3490	1.4154
BHB	HNN	−637.72	2252.49	Lévy-stable	0.3423	1.4091
BHB	HNZ	−4155.25	−1625.75	Lévy-stable	0.3460	1.4336
BRZ	HGE	−7937.62	−6751.18	Lévy-stable	0.4226	1.2134
BRZ	HGN	−7611.17	−6305.25	Lévy-stable	0.3983	1.1905
BRZ	HGZ	−9966.58	−8822.43	Lévy-stable	0.4202	1.1442
CRI	HNE	−16503.34	−15236.72	Lévy-stable	0.4291	1.2209
CRI	HNN	−12548.20	−10924.74	Lévy-stable	0.3979	1.2331
CRI	HNZ	−4223.23	−3860.05	Lévy-stable	0.5422	1.2773
SAV	HGE	−591.75	104.72	Lévy-stable	0.6005	1.2850
SAV	HGN	714.93	1372.79	Lévy-stable	0.6062	1.2442
SAV	HGZ	−355.09	87.78	Lévy-stable	0.6088	1.2989
SLZ	HNE	27506.29	27575.48	Lévy-stable	0.9166	1.6357
SLZ	HNN	41243.95	41323.95	Lévy-stable	1.8195	1.6008
SLZ	HNZ	25803.73	25799.57	Student- t	1.2536	1.4538
CAGN	HNE	12953.50	13202.80	Lévy-stable	0.7445	1.1631
CAGN	HNN	30945.04	31145.75	Lévy-stable	0.7708	1.3518
CAGN	HNZ	26232.17	26354.66	Lévy-stable	0.8096	1.1930
MENA	HNE	−23767.47	−23398.32	Lévy-stable	0.6208	1.2398
MENA	HNN	−21189.22	−20835.15	Lévy-stable	0.6309	1.1751
MENA	HNZ	−779.66	−770.41	Lévy-stable	0.7517	1.1124

Appendix D

Phase picking and signal segmentation

This appendix reports the complete phase picking and signal segmentation results for both datasets, across all picking methods and filtering configurations. For each dataset, results are organised by segmentation method (AR-AIC and PhaseNet) and, for PhaseNet, by filtering configuration (no-threshold, probability threshold, and residual filter). Within each configuration, phase picking residuals and quality control outcomes are reported separately for acceleration, velocity, and displacement signals. The corresponding moment scaling results are provided in Appendix E.

D.1 Crustal velocities and theoretical arrival times

Crustal velocities, hypocentral distances, and theoretical P- and S-wave arrival times are reported here for reference; as they depend exclusively on station location and hypocentral depth, they are identical across signal types and picking methods for each dataset (Tables D.1 and D.2). For the INT-41004391 dataset three distinct layer configurations were identified across the 22 stations: 13 stations traverse upper sediments, upper crust, and partial middle crust; 8 stations additionally cross middle sediments and partial lower crust; and 1 station traverses upper sediments, upper crust, middle crust, and partial lower crust. For the IT-2009-0009 dataset, four distinct crustal layer configurations were identified across the 62 stations: 42 stations traverse upper sediments, upper crust, and partial middle crust; 10 stations additionally cross middle and lower sediments; 9 stations traverse upper sediments, middle sediments, upper crust, and partial middle crust; and 1 station crosses upper and partial middle crust only.

Dataset	Velocity	Min (km/s)	Mean (km/s)	Median (km/s)	Max (km/s)
INT-41004391	v_P	5.59	5.89	5.98	6.26
	v_S	3.10	3.36	3.46	3.57
IT-2009-0009	v_P	5.37	5.92	6.01	6.17
	v_S	3.02	3.42	3.48	3.58

Table D.1: Station-specific crustal velocities from CRUST1.0 for both datasets.

Quantity	Dataset	Min	Median	Max
d_{epi} (km)	INT-41004391	4.80	67.80	109.50
	IT-2009-0009	1.80	137.80	421.40
d_{hypo} (km)	INT-41004391	11.45	68.59	109.99
	IT-2009-0009	8.49	137.90	421.48
t_P^{theo} (s)	INT-41004391	6.92	60.68	61.81
	IT-2009-0009	-47.58	–	61.85
t_S^{theo} (s)	INT-41004391	8.31	68.67	76.51
	IT-2009-0009	-14.77	–	72.06

Table D.2: Hypocentral distances and theoretical arrival times for both datasets.

D.2 INT-41004391 (Mw 3.8, Maritime Alps)

Full phase picking and segmentation results for the INT-41004391 dataset (22 stations, 66 records). Crustal velocities and theoretical arrival times are reported in Section D.1.

D.2.1 AR-AIC segmentation

Phase picking residuals and quality control results obtained with the AR-AIC method are reported in Tables D.4–D.5 for the three signal types; the adaptive search window configuration is given in Table D.3.

Bin	Range (km)	N stations	$\Delta t_P / \Delta t_S$ (s)
Bin 1 (near)	(0, 54.50]	7	$\pm 3 / \pm 6$
Bin 2 (regional)	(54.50, 81.57]	7	$\pm 5 / \pm 10$
Bin 3 (far)	(81.57, ∞)	8	$\pm 7 / \pm 14$

Table D.3: Hypocentral distance bins and half-widths of the adaptive AR-AIC search windows. Total range: 11.45–109.99 km.

Signal	Phase	Mean $\pm$ Std (s)	Min (s)	Max (s)
Acceleration	P-wave	0.88 ± 0.63	0.00	2.21
	S-wave	-1.13 ± 2.64	-6.52	4.28
Velocity	P-wave	0.86 ± 0.59	-0.09	2.02
	S-wave	-1.09 ± 3.02	-6.72	5.08
Displacement	P-wave	0.32 ± 1.19	-2.73	2.08
	S-wave	0.29 ± 4.15	-6.72	9.36

Table D.4: Residuals between onset times detected by AR-AIC and theoretical arrival times from CRUST1.0, for the three signal types.

Signal	Method	Peak	MonoP	MonoS	SNR P	SNR S
Acceleration	Rautian	16/66 (24.2%)	48/66 (72.7%)	51/66 (77.3%)	0/66 (0.0%)	0/66 (0.0%)
	Arias	13/63 (20.6%)				0/63 (0.0%)
	Envelope	13/56 (23.2%)				0/56 (0.0%)
	Median	16/66 (24.2%)				0/66 (0.0%)
Velocity	Rautian	12/66 (18.2%)	48/66 (72.7%)	48/66 (72.7%)	0/66 (0.0%)	0/66 (0.0%)
	Arias	9/63 (14.3%)				0/63 (0.0%)
	Envelope	12/57 (21.1%)				0/57 (0.0%)
	Median	10/64 (15.6%)				0/64 (0.0%)
Displacement	Rautian	14/66 (21.2%)	39/66 (59.1%)	48/66 (72.7%)	7/66 (10.6%)	0/66 (0.0%)
	Arias	11/63 (17.5%)				0/63 (0.0%)
	Envelope	15/59 (25.4%)				0/59 (0.0%)
	Median	13/65 (20.0%)				0/65 (0.0%)

Table D.5: Number of components failing each quality control check (AR-AIC segmentation), for the three signal types and the four coda onset detection methods. MonoP, MonoS, and SNR P are independent of the coda method and are shown once per signal type. Denominators differ across methods due to missing windows.

D.2.2 PhaseNet segmentation

Phase picking residuals and quality control results obtained with PhaseNet, reported for both no-threshold and probability threshold configurations.

No-threshold configuration

Results obtained applying PhaseNet without any post-processing filter on detection probabilities or picking residuals. This configuration retains all picks for which PhaseNet produced a valid P- and S-wave onset, subject only to the physical consistency check $t_S > t_P$. Detection statistics, phase picking residuals, and quality control results are reported in Tables D.6–D.8.

Metric	Acceleration	Velocity	Displacement
Total stations available	22	22	22
Stations retained	21	21	20
Components analyzed	63	63	60
Stations skipped	1	1	2
Mean P probability (%)	93.5	94.6	79.4
Mean S probability (%)	47.3	45.1	32.6
Mean P-S time (s)	8.87	8.62	8.32

Table D.6: PhaseNet phase detection summary (no filter configuration) for the INT-41004391 dataset across data types.

Signal	Phase	Mean $\pm$ Std (s)	Min (s)	Max (s)
Acceleration	P-wave	0.84 ± 0.60	-0.02	2.10
	S-wave	0.46 ± 1.48	-1.91	2.56
Velocity	P-wave	0.81 ± 0.58	-0.16	1.93
	S-wave	0.17 ± 1.58	-2.10	2.17
Displacement	P-wave	0.83 ± 0.59	-0.22	1.97
	S-wave	0.15 ± 1.57	-2.28	1.98

Table D.7: Residuals between onset times detected by PhaseNet and theoretical arrival times from CRUST1.0, for the three signal types.

Signal	Method	Peak	MonoP	MonoS	SNR P	SNR S
Acceleration	Rautian	6/63 (9.5%)	45/63 (71.4%)	45/63 (71.4%)	0/63 (0.0%)	0/63 (0.0%)
	Arias	6/63 (9.5%)				0/63 (0.0%)
	Envelope	6/62 (9.7%)				0/62 (0.0%)
	Median	6/63 (9.5%)				0/63 (0.0%)
Velocity	Rautian	2/63 (3.2%)	45/63 (71.4%)	36/63 (57.1%)	0/63 (0.0%)	0/63 (0.0%)
	Arias	2/63 (3.2%)				0/63 (0.0%)
	Envelope	1/58 (1.7%)				0/58 (0.0%)
	Median	2/63 (3.2%)				0/63 (0.0%)
Displacement	Rautian	2/60 (3.3%)	42/60 (70.0%)	36/60 (60.0%)	5/60 (8.3%)	0/60 (0.0%)
	Arias	2/60 (3.3%)				0/60 (0.0%)
	Envelope	2/55 (3.6%)				0/55 (0.0%)
	Median	2/60 (3.3%)				0/60 (0.0%)

Table D.8: Number of components failing each quality control check (PhaseNet, no-threshold configuration), for the three signal types and the four coda onset detection methods. MonoP, MonoS, and SNR P are independent of the coda method and are shown once per signal type. Denominators differ across methods due to missing windows.

Threshold configuration

Results obtained applying PhaseNet with a minimum probability threshold of 0.3 for both P- and S-wave picks, following the default detection threshold of the SeisBench framework (Woollam et al., 2022). This configuration is applied exclusively to the INT-41004391 dataset, for which the S-wave probability distribution is sufficiently high to allow meaningful filtering (median $P_S = 0.45$). Detection statistics, phase picking residuals, and quality control results are reported in Tables D.9–D.11.

Metric	Acceleration	Velocity	Displacement
Total stations available	22	22	22
Stations retained	16	15	10
Stations skipped	6	7	12
Components analyzed	48	45	30
Mean P probability (%)	93.2	94.3	85.3
Mean S probability (%)	56.7	57.2	52.8
Mean P-S time (s)	7.87	7.57	6.49

Table D.9: PhaseNet phase detection summary (threshold configuration) for the INT-41004391 dataset across data types.

Signal	Phase	Mean $\pm$ Std (s)	Min (s)	Max (s)
Acceleration	P-wave	1.03 ± 0.55	0.25	2.10
	S-wave	0.87 ± 1.38	-1.79	2.56
Velocity	P-wave	0.97 ± 0.54	0.17	1.93
	S-wave	0.59 ± 1.38	-1.91	2.17
Displacement	P-wave	1.03 ± 0.55	0.18	1.97
	S-wave	0.61 ± 1.44	-2.03	1.98

Table D.10: Residuals between onset times detected by PhaseNet (threshold configuration) and theoretical arrival times from CRUST1.0, for the three signal types.

Signal	Method	Peak	MonoP	MonoS	SNR P	SNR S
Acceleration	Rautian	4/48 (8.3%)				0/48 (0.0%)
	Arias	4/48 (8.3%)				0/48 (0.0%)
	Envelope	4/47 (8.5%)	33/48 (68.8%)	30/48 (62.5%)	0/48 (0.0%)	0/47 (0.0%)
	Median	4/48 (8.3%)				0/48 (0.0%)
Velocity	Rautian	1/45 (2.2%)				0/45 (0.0%)
	Arias	1/45 (2.2%)	30/45 (66.7%)	24/45 (53.3%)	0/45 (0.0%)	0/45 (0.0%)
	Envelope	0/41 (0.0%)				0/41 (0.0%)
	Median	1/45 (2.2%)				0/45 (0.0%)
Displacement	Rautian	1/30 (3.3%)				0/30 (0.0%)
	Arias	1/30 (3.3%)	18/30 (60.0%)	9/30 (30.0%)	5/30 (16.7%)	0/30 (0.0%)
	Envelope	1/26 (3.8%)				0/26 (0.0%)
	Median	1/30 (3.3%)				0/30 (0.0%)

Table D.11: Number of components failing each quality control check (PhaseNet, threshold configuration), for the three signal types and the four coda onset detection methods.

D.3 IT-2009-0009 (Mw 6.1, L'Aquila)

Full phase picking and segmentation results for the IT-2009-0009 dataset (62 stations, 186 records, a, after exclusion of station GSG due to a missing acquisition timestamp and station TTS due to an anomalous origin time offset incompatible with the segmentation pipeline). Crustal velocities and theoretical arrival times are reported in Section D.1.

D.3.1 AR-AIC segmentation

Phase picking residuals and quality control results obtained with the AR-AIC method are reported in Tables D.13–D.14 for the three signal types; the adaptive search window configuration is given in Table D.12. An example of the onset detection output for a representative station is shown in Figure D.1.

Bin	Range (km)	N stations	$\Delta t_P / \Delta t_S$ (s)
Bin 1 (near)	(0, 82.62]	21	$\pm 3 / \pm 6$
Bin 2 (regional)	(82.62, 175.66]	20	$\pm 5 / \pm 10$
Bin 3 (far)	(175.66, ∞)	21	$\pm 7 / \pm 14$

Table D.12: Hypocentral distance bins and half-widths of the adaptive AR-AIC search windows.

Signal	Phase	Mean $\pm$ Std (s)	Min (s)	Max (s)
Acceleration	P-wave	0.55 ± 1.93	-6.09	2.21
	S-wave	3.64 ± 5.36	-14.00	12.87
Velocity	P-wave	0.23 ± 1.92	-5.16	3.86
	S-wave	3.51 ± 5.83	-14.00	12.35
Displacement	P-wave	-0.61 ± 2.32	-6.22	4.96
	S-wave	2.09 ± 7.04	-14.00	13.56

Table D.13: Residuals between onset times detected by AR-AIC and theoretical arrival times from CRUST1.0, for the three signal types.

Signal	Method	Peak	MonoP	MonoS	SNR P	SNR S
Acceleration	Rautian	27/180 (15.0%)	159/186 (85.5%)	159/186 (85.5%)	0/117 (0.0%)	0/126 (0.0%)
	Arias	15/180 (8.3%)				0/126 (0.0%)
	Envelope	49/170 (28.8%)				0/121 (0.0%)
	Median	21/180 (11.7%)				0/126 (0.0%)
Velocity	Rautian	32/180 (17.8%)	153/186 (82.3%)	168/186 (90.3%)	8/120 (6.7%)	0/126 (0.0%)
	Arias	12/180 (6.7%)				0/126 (0.0%)
	Envelope	53/151 (35.1%)				0/111 (0.0%)
	Median	26/180 (14.4%)				0/126 (0.0%)
Displacement	Rautian	42/180 (23.3%)	162/186 (87.1%)	162/186 (87.1%)	31/117 (26.5%)	3/120 (2.5%)
	Arias	20/180 (11.1%)				3/120 (2.5%)
	Envelope	81/168 (48.2%)				3/117 (2.6%)
	Median	34/180 (18.9%)				3/120 (2.5%)

Table D.14: Number of components failing each quality control check (AR-AIC, no-threshold configuration), for the three signal types and the four coda onset detection methods. MonoP, MonoS, and SNR P are independent of the coda method and are shown once per signal type. Denominators differ across methods due to missing windows.




Figure D.1: AR-AIC onset detection for station SPO ($d_{\text{epi}} = 68.2$ km, IT-2009-0009 dataset), vertical component (HNZ), acceleration signal. The red and teal shaded regions indicate the P-wave and S-wave search windows, respectively. Dashed vertical lines show the theoretical arrival times computed from CRUST1.0; solid vertical lines show the onsets detected by AR-AIC.

D.3.2 PhaseNet segmentation

Phase picking residuals and quality control results obtained with PhaseNet, reported for both no-threshold and residual filter configurations.

No-threshold configuration

Detection statistics, phase picking residuals, and quality control results for the no-threshold configuration are reported in Tables D.15–D.17.

Metric	Acceleration	Velocity	Displacement
Stations processed	62	62	62
Stations with $S \leq P$	3	1	11
Stations retained	59	61	51
Components analyzed	177	183	153
Mean P probability (%)	40.1	37.2	17.8
Mean S probability (%)	11.0	9.8	4.7
Mean S-P time (s)	11.27	13.82	5.83

Table D.15: PhaseNet phase detection summary (no probability threshold, no residual filter) for the three signal types. Excluded stations with $t_S \leq t_P$.

Signal	Phase	Mean $\pm$ Std (s)	Min (s)	Max (s)
Acceleration	P-wave	0.46 ± 2.10	-7.66	2.37
	S-wave	-0.29 ± 9.34	-44.66	11.10
Velocity	P-wave	0.55 ± 2.12	-7.83	2.51
	S-wave	2.72 ± 31.07	-44.70	222.78
Displacement	P-wave	2.99 ± 5.16	-5.53	26.75
	S-wave	-1.17 ± 11.93	-49.63	22.22

Table D.16: Residuals between onset times detected by PhaseNet and theoretical arrival times from CRUST1.0, no-threshold configuration, for the three signal types.

Signal	Method	Peak	MonoP	MonoS	SNR P	SNR S
Acceleration	Rautian	17/174 (9.8%)				0/123 (0.0%)
	Arias	5/174 (2.9%)				0/123 (0.0%)
	Envelope	21/156 (13.5%)	159/177 (89.8%)	135/177 (76.3%)	1/114 (0.9%)	0/114 (0.0%)
	Median	7/174 (4.0%)				0/123 (0.0%)
Velocity	Rautian	23/177 (13.0%)				1/126 (0.8%)
	Arias	3/174 (1.7%)				0/123 (0.0%)
	Envelope	58/146 (39.7%)	162/183 (88.5%)	150/183 (82.0%)	6/111 (5.4%)	1/102 (1.0%)
	Median	19/177 (10.7%)				1/126 (0.8%)
Displacement	Rautian	31/153 (20.3%)				11/114 (9.6%)
	Arias	7/153 (4.6%)				11/114 (9.6%)
	Envelope	56/134 (41.8%)	141/153 (92.2%)	123/153 (80.4%)	28/90 (31.1%)	9/102 (8.8%)
	Median	24/153 (15.7%)				11/114 (9.6%)

Table D.17: Number of components failing each quality control check (PhaseNet, no-threshold configuration), for the three signal types and the four coda onset detection methods. Denominators differ across methods because the number of available windows varies.

Residual filter configuration

Metric	Acceleration	Velocity	Displacement
Stations processed	62	62	62
Stations with $S \leq P$	3	1	11
Stations filtered by residual filter	6	9	3
Stations retained	53	52	48
Components analyzed	159	156	144

Table D.18: PhaseNet phase detection summary (residual filter) for the three signal types. Excluded stations with $t_S \leq t_P$.

Signal	Phase	Mean $\pm$ Std (s)	Min (s)	Max (s)
Acceleration	P-wave	0.96 ± 0.78	-1.53	2.37
	S-wave	2.08 ± 4.48	-7.71	11.10
Velocity	P-wave	1.10 ± 0.80	-1.64	2.51
	S-wave	2.34 ± 4.73	-9.18	13.77
Displacement	P-wave	3.31 ± 5.16	-4.98	26.75
	S-wave	0.74 ± 8.18	-20.66	22.22

Table D.19: Residuals between onset times detected by PhaseNet and theoretical arrival times from CRUST1.0, residual filter configuration, for the three signal types.

Signal	Method	Peak	MonoP	MonoS	SNR P	SNR S
Acceleration	Rautian	7/156 (4.5%)	144/159 (90.6%)	117/159 (73.6%)	0/96 (0.0%)	0/105 (0.0%)
	Arias	5/156 (3.2%)				0/105 (0.0%)
	Envelope	18/139 (12.9%)				0/97 (0.0%)
	Median	5/156 (3.2%)				0/105 (0.0%)
Velocity	Rautian	6/150 (4.0%)	138/156 (88.5%)	129/156 (82.7%)	1/87 (1.1%)	0/102 (0.0%)
	Arias	3/150 (2.0%)				0/102 (0.0%)
	Envelope	42/123 (34.1%)				0/82 (0.0%)
	Median	7/150 (4.7%)				0/102 (0.0%)
Displacement	Rautian	27/144 (18.8%)	129/144 (89.6%)	108/144 (75.0%)	25/84 (29.8%)	8/105 (7.6%)
	Arias	7/144 (4.9%)				8/105 (7.6%)
	Envelope	48/125 (38.4%)				6/93 (6.5%)
	Median	20/144 (13.9%)				8/105 (7.6%)

Table D.20: Number of components failing each quality control check (PhaseNet, residual filter configuration), for the three signal types and the four coda onset detection methods. Denominators differ across methods because the number of available windows varies.

Appendix E

Moment scaling: full results

This appendix reports the complete moment scaling results for both datasets, across all signal types, segmentation methods, and PhaseNet filtering configurations. For each dataset, results are organised by segmentation method (AR-AIC and PhaseNet) and, for PhaseNet, by filtering configuration (no-threshold, probability threshold, and residual filter). For each configuration, scaling exponents $\zeta(q)$ are reported separately for acceleration, velocity, and displacement signals across all temporal windows and coda-onset detection methods. The corresponding phase-picking and segmentation results are presented in Appendix D. For each window–method combination, N denotes the number of station-component traces and $\tau_{\max}$ the maximum time lag used in the fit. Columns $\zeta(1)$ and $\zeta(2)$ list the exponents for $q = 1$ and $q = 2$, with the corresponding R^2 values of the log–log fits reported in parentheses.

E.1 INT-41004391 (Mw 3.8, Maritime Alps)

This section reports the complete moment scaling results for the INT-41004391 dataset, obtained from AR-AIC and PhaseNet segmentation across all applicable filtering configurations.

E.1.1 AR-AIC segmentation

Moment scaling results obtained from AR-AIC segmentation are reported in Tables E.1–E.3 for the three signal types. Scaling exponent spectra $\zeta(q)$ for the Arias, Envelope, and Median coda onset methods are shown in Figures E.1, E.2, and E.3 respectively; results for the Rautian method are shown in Figure 6.6 in Chapter 6.




Figure E.1: Scaling exponent spectrum $\zeta(q)$ vs. q for the INT-41004391 dataset (Queyras, M_w 3.8), AR-AIC segmentation, Arias coda onset. The dashed red line shows the normal diffusion reference $\zeta(q) = q/2$. Error bars represent one standard error of the log-log fit slope.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	—	66	2.14	0.95 (0.28)	2.56 (0.34)
S-wave	Rautian	66	5.42	0.02 (0.01)	-0.17 (0.06)
	Arias	63	3.56	0.02 (0.00)	-0.25 (0.06)
	Envelope	56	2.08	-0.08 (0.05)	-0.13 (0.01)
	Median	66	2.08	-0.12 (0.08)	-0.26 (0.04)
Coda	Rautian	66	14.17	-0.07 (0.34)	-0.17 (0.23)
	Arias	66	6.84	-0.05 (0.04)	-0.16 (0.03)
	Envelope	66	12.00	-0.09 (0.30)	-0.16 (0.20)
	Median	66	12.84	-0.09 (0.43)	-0.19 (0.23)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.1: Moment scaling results for **acceleration** signals, obtained from AR-AIC segmentation (no residual filter), for the INT-41004391 event.




Figure E.2: Scaling exponent spectrum $\zeta(q)$ vs. q for the INT-41004391 dataset (Queyras, M_w 3.8), AR-AIC segmentation, Envelope coda onset. Layout as in Figure E.1.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	—	66	2.21	0.62 (0.09)	1.69 (0.12)
S-wave	Rautian	66	5.42	0.07 (0.07)	-0.02 (0.00)
	Arias	63	4.41	0.08 (0.12)	-0.05 (0.01)
	Envelope	57	2.58	0.16 (0.05)	0.40 (0.03)
	Median	64	2.58	0.09 (0.04)	0.18 (0.02)
Coda	Rautian	66	14.17	-0.13 (0.45)	-0.36 (0.30)
	Arias	64	8.57	-0.02 (0.02)	-0.02 (0.00)
	Envelope	66	12.04	-0.14 (0.41)	-0.38 (0.38)
	Median	66	10.34	-0.13 (0.33)	-0.42 (0.26)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.2: Moment scaling results for **velocity** signals, obtained from AR-AIC segmentation (no residual filter), for the INT-41004391 event.




Figure E.3: Scaling exponent spectrum $\zeta(q)$ vs. q for the INT-41004391 dataset (Queyras, M_w 3.8), AR-AIC segmentation, Median coda onset. Layout as in Figure E.1.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	—	66	2.40	-0.27 (0.04)	-0.63 (0.05)
S-wave	Rautian	66	5.42	0.03 (0.02)	-0.33 (0.12)
	Arias	63	8.15	0.17 (0.52)	0.11 (0.04)
	Envelope	59	2.08	0.01 (0.00)	-0.37 (0.07)
	Median	65	2.46	-0.03 (0.01)	-0.51 (0.16)
Coda	Rautian	66	9.14	-0.08 (0.18)	-0.09 (0.02)
	Arias	65	2.78	-0.29 (0.22)	-0.64 (0.12)
	Envelope	66	13.97	-0.08 (0.19)	-0.24 (0.19)
	Median	66	9.14	-0.13 (0.53)	-0.40 (0.37)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.3: Moment scaling results for **displacement** signals, obtained from AR-AIC segmentation (no residual filter), for the INT-41004391 event.

E.1.2 PhaseNet segmentation

Moment scaling results obtained from PhaseNet segmentation are reported below, separately for the no-threshold and probability threshold configurations.

No-threshold configuration

Tables E.4–E.6 report moment scaling results for the no-threshold configuration, for the three signal types.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	–	63	2.60	-0.27 (0.36)	-1.00 (0.52)
S-wave	Rautian	63	6.59	-0.30 (0.80)	-0.80 (0.75)
	Arias	63	3.87	-0.26 (0.64)	-0.87 (0.55)
	Envelope	62	2.45	-0.24 (0.54)	-0.96 (0.54)
	Median	63	3.87	-0.26 (0.64)	-0.87 (0.55)
Coda	Rautian	63	11.64	-0.08 (0.40)	-0.20 (0.21)
	Arias	63	6.84	-0.03 (0.06)	-0.07 (0.03)
	Envelope	63	12.54	-0.14 (0.57)	-0.39 (0.55)
	Median	63	14.16	-0.05 (0.40)	-0.11 (0.17)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.4: Moment scaling results for **acceleration** signals, obtained from PhaseNet segmentation (no threshold configuration), for the INT-41004391 event.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	–	63	2.46	-0.20 (0.33)	-0.46 (0.29)
S-wave	Rautian	63	6.46	-0.31 (0.66)	-0.97 (0.68)
	Arias	63	4.84	-0.26 (0.59)	-0.96 (0.61)
	Envelope	58	2.29	-0.25 (0.29)	-0.97 (0.35)
	Median	63	4.84	-0.26 (0.59)	-0.96 (0.61)
Coda	Rautian	63	17.70	-0.08 (0.39)	-0.17 (0.30)
	Arias	63	8.57	-0.03 (0.24)	-0.11 (0.16)
	Envelope	63	11.09	-0.05 (0.12)	-0.04 (0.01)
	Median	63	10.34	-0.04 (0.12)	-0.11 (0.11)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.5: Moment scaling results for **velocity** signals, obtained from PhaseNet segmentation (no threshold configuration), for the INT-41004391 event.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	–	60	2.41	0.21 (0.31)	0.94 (0.50)
S-wave	Rautian	60	6.43	-0.22 (0.55)	-0.66 (0.48)
	Arias	60	8.62	-0.18 (0.49)	-0.48 (0.42)
	Envelope	55	2.62	-0.17 (0.20)	-0.63 (0.21)
	Median	60	6.43	-0.22 (0.55)	-0.66 (0.48)
Coda	Rautian	60	17.66	-0.12 (0.42)	-0.26 (0.34)
	Arias	59	8.54	-0.12 (0.49)	-0.29 (0.51)
	Envelope	60	13.80	-0.08 (0.22)	-0.24 (0.25)
	Median	60	17.66	-0.12 (0.44)	-0.26 (0.36)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.6: Moment scaling results for **displacement** signals, obtained from PhaseNet segmentation (no threshold configuration), for the INT-41004391 event.

Probability threshold configuration

Tables E.7–E.9 report moment scaling results for the probability threshold configuration, for the three signal types. Scaling exponent spectra $\zeta(q)$ for the Arias, Envelope, and Median coda onset methods are shown in Figures E.4, E.5, and E.6 respectively; results for the Rautian method are shown in Figure 6.8 in Chapter 6.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	–	48	2.60	-0.42 (0.55)	-1.20 (0.57)
S-wave	Rautian	48	6.59	-0.37 (0.83)	-0.85 (0.75)
	Arias	48	3.87	-0.32 (0.66)	-0.93 (0.56)
	Envelope	47	2.45	-0.29 (0.58)	-1.00 (0.54)
	Median	48	3.87	-0.32 (0.66)	-0.93 (0.56)
Coda	Rautian	48	11.64	-0.10 (0.40)	-0.21 (0.22)
	Arias	48	6.84	-0.02 (0.03)	-0.06 (0.02)
	Envelope	48	12.54	-0.15 (0.57)	-0.40 (0.54)
	Median	48	14.16	-0.06 (0.40)	-0.11 (0.17)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.7: Moment scaling results for **acceleration** signals, obtained from PhaseNet segmentation (probability threshold $P \geq 0.3$), for the INT-41004391 event.




Figure E.4: Scaling exponent spectrum $\zeta(q)$ for the INT-41004391 dataset (Queyras, M_w 3.8), PhaseNet segmentation with probability threshold $P \geq 0.3$, Arias coda onset. Layout as in Figure E.1.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	—	45	2.46	-0.37 (0.63)	-0.79 (0.57)
S-wave	Rautian	45	6.46	-0.40 (0.70)	-1.11 (0.70)
	Arias	45	4.84	-0.36 (0.70)	-1.10 (0.66)
	Envelope	41	2.29	-0.32 (0.36)	-1.05 (0.36)
	Median	45	4.84	-0.36 (0.70)	-1.10 (0.66)
Coda	Rautian	45	17.70	-0.08 (0.31)	-0.17 (0.23)
	Arias	45	8.57	-0.04 (0.26)	-0.11 (0.16)
	Envelope	45	11.09	-0.04 (0.08)	-0.04 (0.01)
	Median	45	10.34	-0.04 (0.12)	-0.13 (0.12)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.8: Moment scaling results for **velocity** signals, obtained from PhaseNet segmentation (probability threshold $P \geq 0.3$), for the INT-41004391 event.




Figure E.5: Scaling exponent spectrum $\zeta(q)$ for the INT-41004391 dataset (Queyras, M_w 3.8), PhaseNet segmentation with probability threshold $P \geq 0.3$, Envelope coda onset. Layout as in Figure E.1.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	—	30	2.41	0.19 (0.16)	0.90 (0.35)
S-wave	Rautian	30	6.43	-0.29 (0.57)	-0.75 (0.49)
	Arias	30	8.62	-0.23 (0.59)	-0.52 (0.43)
	Envelope	26	2.62	-0.21 (0.23)	-0.73 (0.23)
	Median	30	6.43	-0.29 (0.57)	-0.75 (0.49)
Coda	Rautian	30	17.66	-0.14 (0.37)	-0.29 (0.28)
	Arias	29	8.54	-0.14 (0.56)	-0.32 (0.49)
	Envelope	30	13.80	-0.09 (0.15)	-0.24 (0.18)
	Median	30	17.66	-0.14 (0.37)	-0.29 (0.28)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.9: Moment scaling results for **displacement** signals, obtained from PhaseNet segmentation (probability threshold $P \geq 0.3$), for the INT-41004391 event.




Figure E.6: Scaling exponent spectrum $\zeta(q)$ for the INT-41004391 dataset (Queyras, M_w 3.8), PhaseNet segmentation with probability threshold $P \geq 0.3$, Median coda onset. Layout as in Figure E.1.

E.2 IT-2009-0009 (Mw 6.1, L'Aquila)

This section reports the complete moment scaling results for the IT-2009-0009 dataset, obtained from AR-AIC and PhaseNet segmentation.

E.2.1 AR-AIC segmentation

Moment scaling results obtained from AR-AIC segmentation are reported in Tables E.10–E.12 for the three signal types. Scaling exponent spectra $\zeta(q)$ for the Arias, Envelope, and Median coda onset methods are shown in Figures E.7, E.8, and E.9 respectively; results for the Rautian method are shown in Figure 6.10 in Chapter 6.




Figure E.7: Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), AR-AIC segmentation, Arias coda onset. Layout as in Figure E.1.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	—	117	2.27	1.28 (0.87)	2.00 (0.70)
S-wave	Rautian	180	4.92	-0.12 (0.27)	-0.27 (0.21)
	Arias	180	4.86	-0.13 (0.28)	-0.33 (0.29)
	Envelope	169	2.77	-0.13 (0.19)	-0.35 (0.22)
	Median	180	4.88	-0.15 (0.32)	-0.36 (0.38)
Coda	Rautian	141	2.88	-0.17 (0.40)	-0.28 (0.27)
	Arias	183	2.31	-0.11 (0.31)	-0.20 (0.26)
	Envelope	184	2.08	-0.03 (0.01)	0.02 (0.00)
	Median	184	2.08	-0.10 (0.12)	-0.17 (0.05)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.10: Moment scaling results for **acceleration** signals, obtained from AR-AIC segmentation (no residual filter), for the IT-2009-0009 event.




Figure E.8: Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), AR-AIC segmentation, Envelope coda onset. Layout as in Figure E.1.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	—	120	2.08	2.34 (0.91)	4.55 (0.89)
S-wave	Rautian	180	4.92	-0.02 (0.02)	-0.17 (0.23)
	Arias	180	6.64	0.00 (0.00)	-0.14 (0.22)
	Envelope	150	2.16	0.10 (0.21)	-0.05 (0.01)
	Median	180	4.92	-0.02 (0.02)	-0.17 (0.23)
Coda	Rautian	137	3.13	-0.10 (0.21)	-0.29 (0.24)
	Arias	180	2.10	-0.15 (0.47)	-0.40 (0.45)
	Envelope	184	2.56	-0.01 (0.00)	-0.17 (0.12)
	Median	183	3.56	-0.12 (0.34)	-0.43 (0.32)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.11: Moment scaling results for **velocity** signals, obtained from AR-AIC segmentation (no residual filter), for the IT-2009-0009 event.




Figure E.9: Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), AR-AIC segmentation, Median coda onset. Layout as in Figure E.1.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	—	123	2.42	1.80 (0.96)	3.85 (0.95)
S-wave	Rautian	180	4.92	0.14 (0.47)	-0.24 (0.27)
	Arias	180	3.15	0.15 (0.47)	-0.25 (0.21)
	Envelope	168	2.21	0.22 (0.61)	-0.11 (0.07)
	Median	180	3.58	0.17 (0.57)	-0.18 (0.13)
Coda	Rautian	148	2.49	0.11 (0.43)	-0.25 (0.18)
	Arias	178	2.25	0.27 (0.93)	-0.09 (0.32)
	Envelope	182	2.04	0.34 (0.71)	0.46 (0.30)
	Median	185	3.04	0.25 (0.58)	0.32 (0.26)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.12: Moment scaling results for **displacement** signals, obtained from AR-AIC segmentation (no residual filter), for the IT-2009-0009 event.

E.2.2 PhaseNet segmentation

Moment scaling results obtained from PhaseNet segmentation are reported below, separately for the no-threshold and residual filter configurations.

No-threshold configuration

Tables E.13–E.15 report moment scaling results for the no-threshold configuration, for the three signal types.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	–	114	2.13	1.73 (0.89)	3.85 (0.87)
S-wave	Rautian	174	3.45	-0.24 (0.54)	-0.51 (0.54)
	Arias	174	4.92	-0.20 (0.51)	-0.45 (0.50)
	Envelope	156	2.02	-0.28 (0.44)	-0.57 (0.36)
	Median	174	3.45	-0.23 (0.52)	-0.51 (0.53)
Coda	Rautian	144	2.84	-0.12 (0.11)	-0.28 (0.10)
	Arias	174	2.31	-0.11 (0.31)	-0.20 (0.26)
	Envelope	174	2.08	0.02 (0.01)	0.12 (0.02)
	Median	174	2.08	-0.03 (0.02)	-0.04 (0.00)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.13: Moment scaling results for **acceleration** signals, obtained from PhaseNet segmentation (no residual filter), for the IT-2009-0009 event.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	–	111	2.33	0.78 (0.68)	1.14 (0.29)
S-wave	Rautian	177	3.11	-0.46 (0.73)	-1.22 (0.75)
	Arias	174	7.17	-0.27 (0.57)	-0.79 (0.66)
	Envelope	146	2.03	-0.62 (0.79)	-1.77 (0.83)
	Median	177	2.54	-0.53 (0.74)	-1.39 (0.76)
Coda	Rautian	147	5.61	-0.16 (0.45)	-0.34 (0.35)
	Arias	177	2.10	-0.15 (0.48)	-0.40 (0.45)
	Envelope	176	2.27	0.16 (0.16)	0.14 (0.03)
	Median	178	3.71	-0.14 (0.27)	-0.31 (0.19)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.14: Moment scaling results for **velocity** signals, obtained from PhaseNet segmentation (no residual filter), for the IT-2009-0009 event.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	–	90	2.11	1.10 (0.99)	2.53 (0.99)
S-wave	Rautian	153	2.91	-0.06 (0.07)	-0.49 (0.57)
	Arias	153	4.54	-0.07 (0.20)	-0.50 (0.72)
	Envelope	134	2.14	0.03 (0.02)	-0.32 (0.36)
	Median	153	2.91	-0.06 (0.07)	-0.49 (0.57)
Coda	Rautian	126	11.16	-0.01 (0.01)	-0.13 (0.17)
	Arias	148	2.25	0.22 (0.90)	-0.14 (0.50)
	Envelope	149	3.27	0.19 (0.54)	0.17 (0.13)
	Median	153	3.04	0.13 (0.48)	-0.09 (0.03)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.15: Moment scaling results for **displacement** signals, obtained from PhaseNet segmentation (no residual filter), for the IT-2009-0009 event.

Residual filter configuration

Tables E.16–E.18 report moment scaling results for the residual filter configuration, for the three signal types. Scaling exponent spectra $\zeta(q)$ for the Arias, Envelope, and Median coda onset methods are shown in Figures E.10, E.11, and E.12 respectively; results for the Rautian method are shown in Figure 6.12 in Chapter 6.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	–	96	2.13	1.73 (0.89)	3.85 (0.87)
S-wave	Rautian	156	3.45	-0.24 (0.54)	-0.51 (0.54)
	Arias	156	4.92	-0.20 (0.51)	-0.45 (0.50)
	Envelope	139	2.02	-0.28 (0.44)	-0.57 (0.36)
	Median	156	3.45	-0.23 (0.52)	-0.51 (0.53)
Coda	Rautian	126	2.84	-0.12 (0.11)	-0.28 (0.10)
	Arias	156	2.31	-0.11 (0.30)	-0.20 (0.26)
	Envelope	156	2.08	0.02 (0.01)	0.12 (0.02)
	Median	156	2.08	-0.03 (0.02)	-0.03 (0.00)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.16: Moment scaling results for **acceleration** signals, obtained from PhaseNet segmentation (residual filter configuration), for the IT-2009-0009 event.




Figure E.10: Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), PhaseNet segmentation with residual filter, Rautian coda onset. Layout as in Figure E.1.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	—	87	2.79	0.73 (0.64)	0.97 (0.23)
S-wave	Rautian	150	3.11	-0.46 (0.73)	-1.22 (0.75)
	Arias	150	7.17	-0.27 (0.57)	-0.79 (0.66)
	Envelope	123	2.03	-0.63 (0.79)	-1.77 (0.83)
	Median	150	3.11	-0.46 (0.73)	-1.22 (0.75)
Coda	Rautian	123	5.61	-0.17 (0.47)	-0.34 (0.35)
	Arias	150	2.10	-0.18 (0.56)	-0.40 (0.46)
	Envelope	150	2.42	0.13 (0.12)	0.09 (0.01)
	Median	151	3.71	-0.16 (0.30)	-0.32 (0.19)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.17: Moment scaling results for **velocity** signals, obtained from PhaseNet segmentation (residual filter configuration), for the IT-2009-0009 event.




Figure E.11: Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), PhaseNet segmentation with residual filter, Rautian coda onset. Layout as in Figure E.1.

Window	Method	N	$\tau_{\max}$ (s)	$\zeta(1)$ (R^2)	$\zeta(2)$ (R^2)
P-wave*	—	84	2.11	1.10 (0.99)	2.53 (0.99)
S-wave	Rautian	144	2.91	-0.06 (0.07)	-0.49 (0.57)
	Arias	144	4.54	-0.08 (0.20)	-0.50 (0.72)
	Envelope	125	2.14	0.03 (0.02)	-0.32 (0.36)
	Median	144	2.91	-0.06 (0.07)	-0.49 (0.57)
Coda	Rautian	117	11.16	-0.02 (0.03)	-0.14 (0.19)
	Arias	139	2.25	0.22 (0.92)	-0.15 (0.50)
	Envelope	140	3.36	0.19 (0.57)	0.19 (0.15)
	Median	144	3.04	0.14 (0.49)	-0.09 (0.03)

*P-wave results are identical across all four coda onset methods, as the P-wave window boundaries depend only on the P- and S-wave onset times.

Table E.18: Moment scaling results for **displacement** signals, obtained from PhaseNet segmentation (residual filter configuration), for the IT-2009-0009 event.




Figure E.12: Scaling exponent spectrum $\zeta(q)$ for the IT-2009-0009 dataset (L'Aquila, M_w 6.1), PhaseNet segmentation with residual filter, Rautian coda onset. Layout as in Figure E.1.

Bibliography

- Aki, Keiiti and Bernard Chouet (1975). “Origin of coda waves: Source, attenuation and scattering effects”. In: *Journal of Geophysical Research* 80.23, pp. 3322–3342. DOI: 10.1029/JB080i023p03322.
- Aki, Keiiti and Paul G. Richards (2002). *Quantitative Seismology*. 2nd ed. University Science Books. ISBN: 0935702962.
- Anderson, Ernest M. (1905). “The Dynamics of Faulting”. In: *Transactions of the Edinburgh Geological Society* 8.3, pp. 387–402. DOI: 10.1144/transed.8.3.387.
- Arias, A. (1970). “A measure of earthquake intensity”. In: *Seismic Design for Nuclear Power Plants*, pp. 438–483.
- Bak, Per (1996). *How Nature Works: The Science of Self-Organized Criticality*. New York: Copernicus. DOI: 10.1007/978-1-4757-5426-1.
- Bak, Per, Kim Christensen, et al. (2002). “Unified Scaling Law for Earthquakes”. In: *Phys. Rev. Lett.* 88 (17), p. 178501. DOI: 10.1103/PhysRevLett.88.178501.
- Bak, Per and Chao Tang (1989). “Earthquakes as a Self-Organized Critical Phenomenon”. In: *Journal of Geophysical Research: Solid Earth* 94.B11, pp. 15635–15637. DOI: 10.1029/JB094iB11p15635.
- Bak, Per, Chao Tang, and Kurt Wiesenfeld (July 1987). “Self-organized criticality: An explanation of the 1/f noise”. In: *Phys. Rev. Lett.* 59 (4), pp. 381–384. DOI: 10.1103/PhysRevLett.59.381.
- (July 1988). “Self-organized criticality”. In: *Phys. Rev. A* 38 (1), pp. 364–374. DOI: 10.1103/PhysRevA.38.364.
- Båth, Markus (1965). “Lateral inhomogeneities of the upper mantle”. In: *Tectonophysics* 2.6, pp. 483–514. DOI: 10.1016/0040-1951(65)90003-X.
- Beyreuther, M. et al. (2010). “ObsPy: A Python Toolbox for Seismology”. In: *Seismological Research Letters* 81.3, pp. 530–533. DOI: 10.1785/gssrl.81.3.530.
- Bonnet, E et al. (2001). “Scaling of fracture systems in geological data”. In: *Reviews of Geophysics* 39, pp. 347–383. DOI: 10.1029/1999RG000074.
- Burridge, R. and L. Knopoff (June 1967). “Model and theoretical seismicity”. In: *Bulletin of the Seismological Society of America* 57.3, pp. 341–371. DOI: 10.1785/BSSA0570030341. eprint: <https://pubs.geoscienceworld.org/ssa/bssa/article-pdf/57/3/341/5349555/bssa0570030341.pdf>.
- Cianetti, Spina et al. (Mar. 2026). “Which is better: deep-learning or manual seismic arrival-time picking?” In: *Geophysical Journal International* 244.3, ggaf494. ISSN: 1365-246X. DOI: 10.1093/gji/ggaf494. eprint: <https://academic.oup.com/gji/article-pdf/244/3/ggaf494/66263768/ggaf494.pdf>.
- Colangeli, Matteo, Antonio Di Francesco, and Lamberto Rondoni (2024). “Detecting Phase Transitions through Non-Equilibrium Work Fluctuations”. In: *Symmetry* 16.1. ISSN: 2073-8994. DOI: 10.3390/sym16010125.

- Colangeli, Matteo, Lamberto Rondoni, and Antonella Verderosa (July 2014). “Focus on some nonequilibrium issues”. In: *Chaos, Solitons and Fractals* 64, pp. 2–15. ISSN: 0960-0779. DOI: 10.1016/j.chaos.2014.03.002.
- Cowie, Patience A. and Christopher H. Scholz (1992). “Growth of Faults by Accumulation of Seismic Slip”. In: *Journal of Geophysical Research: Solid Earth* 97.B7, pp. 11085–11095. DOI: 10.1029/92JB00586.
- Filippucci, Marilena et al. (2021). “Seismic Envelopes of Coda Decay for Q-coda Attenuation Studies of the Gargano Promontory (Southern Italy) and Surrounding Regions”. In: *Data* 6.9. ISSN: 2306-5729. DOI: 10.3390/data6090098.
- Frohlich, Cliff and Scott D. Davis (1993). “Teleseismic b values; Or, much ado about 1.0”. In: *Journal of Geophysical Research: Solid Earth* 98.B1, pp. 631–644. DOI: <https://doi.org/10.1029/92JB01891>.
- Gutenberg, B. and C. F. Richter (1944). “Frequency of earthquakes in California”. In: *Bulletin of the Seismological Society of America* 34.4, pp. 185–188. DOI: 10.1785/BSSA0340040185.
- (1956). “Earthquake magnitude, intensity, energy, and acceleration”. In: *Bulletin of the Seismological Society of America* 46.2, pp. 105–145. DOI: 10.1785/BSSA0460020105.
- Hawkes, A. G. and L. Adamopoulos (1973). “Cluster Models for Earthquakes – Regional Comparisons”. In: *Bulletin of the International Statistical Institute*. Vol. 45. 3, pp. 454–461.
- Hawkes, Alan G. (1971). “Spectra of Some Self-Exciting and Mutually Exciting Point Processes”. In: *Biometrika* 58.1, pp. 83–90. ISSN: 00063444, 14643510.
- Helmstetter, Agnès and Didier Sornette (2003). “Båth’s law derived from the Gutenberg-Richter law and from aftershock properties”. In: *Geophysical Research Letters* 30.20. DOI: <https://doi.org/10.1029/2003GL018186>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2003GL018186>.
- Hubbert, M. K. and W. W. Rubey (1959). “Role of fluid pressure in mechanics of overthrust faulting: I. Mechanics of fluid-filled porous solids and its application to overthrust faulting”. In: *Geological Society of America Bulletin* 70.2, pp. 115–166. DOI: 10.1130/0016-7606(1959)70[115:ROFPIM]2.0.CO;2.
- Iliopoulos, Aggelos et al. (Aug. 2019). “Superstatistics, Complexity and Earthquakes: A Brief Review and Case Study on Hellenic Seismicity”. In: *Bollettino di Geofisica Teorica ed Applicata*.
- Kissling, E. et al. (1994). “Initial reference models in local earthquake tomography”. In: *Journal of Geophysical Research: Solid Earth* 99.B10, pp. 19635–19646. DOI: <https://doi.org/10.1029/93JB03138>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/93JB03138>.
- Laske, G. et al. (2013). “Update on CRUST1.0 - A 1-degree global model of Earth’s crust”. In: *Geophysical Research Abstracts*. Vol. 15. EGU General Assembly, EGU2013–2658.
- Leeman, John R. et al. (2013). *Crust1.0: A Global Model of Earth’s Crust*. <https://github.com/jrleeman/Crust1.0>. GitHub repository, accessed: 2026-04-16.
- Lippiello, Eugenio, Cataldo Godano, and Lucilla de Arcangelis (2019). “The Relevance of Foreshocks in Earthquake Triggering: A Statistical Study”. In: *Entropy* 21.2, p. 173. DOI: 10.3390/e21020173.
- Lombardi, Anna Maria (2002). “Probabilistic interpretation of Båth’s law”. In: *Annals of Geophysics* 45.3-4, pp. 455–472. DOI: 10.4401/ag-3520.
- Michas, Georgios (2024). “Complexity and Statistical Physics Approaches to Earthquakes”. In: *Entropy* 26.1. ISSN: 1099-4300. DOI: 10.3390/e26010059.

- Michellini, A. et al. (2021). “INSTANCE – the Italian seismic dataset for machine learning”. In: *Earth System Science Data* 13.12, pp. 5509–5544. DOI: 10.5194/essd-13-5509-2021.
- Ogata, Yoshiko (Oct. 2024). “How the ETAS models were created, used, and evolved – Personal views and perspectives”. In: *Annals of Geophysics* 67, S428. DOI: 10.4401/ag-9153.
- Ogata, Yoshihiko (1988). “Statistical Models for Earthquake Occurrences and Residual Analysis for Point Processes”. In: *Journal of the American Statistical Association* 83.401, pp. 9–27. ISSN: 01621459, 1537274X.
- (1998). “Space-time point-process models for earthquake occurrences”. In: *Annals of the Institute of Statistical Mathematics* 50.2, pp. 379–402. DOI: 10.1023/A:1003403601725.
- (2001). “Exploratory analysis of earthquake clusters by likelihood-based trigger models”. In: *Journal of Applied Probability* 38.A, pp. 202–212. DOI: 10.1239/jap/1085496602.
- Olami, Zeev, Hans Jacob S. Feder, and Kim Christensen (1992). “Self-organized criticality in a continuous, nonconservative cellular automaton modeling earthquakes”. In: *Phys. Rev. Lett.* 68 (8), pp. 1244–1247. DOI: 10.1103/PhysRevLett.68.1244.
- Omori, Fusakichi (1894). “On aftershocks of earthquakes”. In: *Journal of the College of Science, Imperial University of Tokyo* 7, pp. 111–120.
- Rautian, T. G. and V. I. Khalturin (Aug. 1978). “The Use of the Coda for Determination of the Earthquake Source Spectrum”. In: *Bulletin of the Seismological Society of America* 68.4, pp. 923–948. ISSN: 0037-1106.
- Reid, Harry Fielding (1911). “The Elastic-Rebound Theory of Earthquakes”. In: *University of California Publications, Bulletin of the Department of Geology* 6, pp. 413–444.
- Ronneberger, Olaf, Philipp Fischer, and Thomas Brox (2015). “U-Net: Convolutional Networks for Biomedical Image Segmentation”, booktitle=“Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015”. In: ed. by Nassir Navab et al. Cham: Springer International Publishing, pp. 234–241. ISBN: 978-3-319-24574.
- Rundle, John B. et al. (2003). “Statistical physics approach to understanding the multiscale dynamics of earthquake fault systems”. In: *Reviews of Geophysics* 41.4. DOI: <https://doi.org/10.1029/2003RG000135>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2003RG000135>.
- Scholz, Christopher H. (2018). *The Mechanics of Earthquakes and Faulting*. 3rd. Cambridge: Cambridge University Press. DOI: 10.1017/9781316681473.
- Shcherbakov, Robert and Donald L. Turcotte (2004). “A modified form of Båth’s law”. In: *Bulletin of the Seismological Society of America* 94.5, pp. 1968–1975. DOI: 10.1785/012003162.
- Shcherbakov, Robert, Donald L. Turcotte, and John B. Rundle (2004). “A generalized Omori’s law for earthquake aftershock decay”. In: *Geophysical Research Letters* 31.11. DOI: <https://doi.org/10.1029/2004GL019808>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2004GL019808>.
- (Sept. 2006). “Scaling Properties of the Parkfield Aftershock Sequence”. In: *Bulletin of the Seismological Society of America* 96.4B, S376–S384. ISSN: 0037-1106. DOI: 10.1785/0120050815. eprint: https://pubs.geoscienceworld.org/ssa/bssa/article-pdf/96/4B/S376/2725747/S376_964B_05815.pdf.
- Shen, L., D. R. Schmitt, and K. Haug (Oct. 2018). *Measurements of the States of In Situ Stress for the Duvernay Formation near Fox Creek, West-Central Alberta*. Tech.

- rep. AER/AGS Report 97. Edmonton, Alberta, Canada: Alberta Energy Regulator / Alberta Geological Survey.
- Sigalotti, Leonardo, Alejandro Ramirez-Rojas, and C. Vargas (Feb. 2023). “Tsallis q-Statistics in Seismology”. In: *Entropy* 25, p. 408. DOI: 10.3390/e25030408.
- Sornette, Didier and Maximilian J. Werner (2009). “Statistical Physics Approaches to Seismicity”. In: *Encyclopedia of Complexity and Systems Science*. Ed. by Robert A. Meyers. New York, NY: Springer New York, pp. 7872–7891. ISBN: 978-0-387-30440-3. DOI: 10.1007/978-0-387-30440-3_467.
- Trifunac, M. D. and A. G. Brady (June 1975). “A study on the duration of strong earthquake ground motion”. In: *Bulletin of the Seismological Society of America* 65.3, pp. 581–626. ISSN: 0037-1106. DOI: 10.1785/BSSA0650030581. eprint: <https://pubs.geoscienceworld.org/ssa/bssa/article-pdf/65/3/581/5319903/bssa0650030581.pdf>.
- Tsallis, Constantino (1988). “Possible Generalization of Boltzmann-Gibbs Statistics”. In: *Journal of Statistical Physics* 52.1-2, pp. 479–487. DOI: 10.1007/BF01016429.
- Utsu, Tokuji (1957). “Magnitude of earthquakes and occurrence of their aftershocks”. In: *Journal of the Seismological Society of Japan* 10, pp. 35–45. DOI: 10.4294/zisin1948.10.1_35.
- (1961). “A statistical study on the occurrence of aftershocks”. In: *Geophysical Magazine* 30.4, pp. 521–605.
- Utsu, Tokuji, Yoshihiko Ogata, and Ritsuko S. Matsu’ura (1995). “The Centenary of the Omori Formula for a Decay Law of Aftershock Activity”. In: *Journal of Physics of the Earth* 43.1, pp. 1–33. DOI: 10.4294/jpe1952.43.1.
- Vere-Jones, D. and R. B. Davies (1966). “A Statistical Survey of Earthquakes in the Main Seismic Region of New Zealand. Part 2: Time Series Analyses”. In: *New Zealand Journal of Geology and Geophysics* 9.3, pp. 251–284. DOI: 10.1080/00288306.1966.10422815.
- Vere-Jones, David, Junko Murakami, and Annemarie Christophersen (n.d.). *A Further Note on Bath’s Law*. Unpublished manuscript, Victoria University, Wellington.
- Vollmer, Jürgen, Claudio Giberti, et al. (2024). “Universal hyper-scaling relations, power-law tails, and data analysis for strong anomalous diffusion”. In: arXiv: 2412.20590 [math-ph].
- Vollmer, Jürgen, Lamberto Rondoni, et al. (Jan. 2021). “Displacement autocorrelation functions for strong anomalous diffusion: A scaling form, universal behavior, and corrections to scaling”. In: *Phys. Rev. Res.* 3 (1), p. 013067. DOI: 10.1103/PhysRevResearch.3.013067.
- Woollam, Jack et al. (Mar. 2022). “SeisBench—A Toolbox for Machine Learning in Seismology”. In: *Seismological Research Letters* 93.3, pp. 1695–1709. ISSN: 0895-0695. DOI: 10.1785/0220210324. eprint: <https://pubs.geoscienceworld.org/ssa/srl/article-pdf/93/3/1695/5596371/srl-2021324.1.pdf>.
- Zhu, Weiqiang and Gregory C Beroza (Jan. 2019). “PhaseNet: a deep-neural-network-based seismic arrival-time picking method”. In: *Geophysical Journal International* 216.1, pp. 261–273. ISSN: 0956-540X. DOI: 10.1093/gji/ggy423. eprint: <https://academic.oup.com/gji/article-pdf/216/1/261/26329430/ggy423.pdf>.