Annals of Science and Inquiry
Vol. 2 No. 6 (2026)

A fixed slope change gives limited spatial transfer in the large-crater diameter tail of a lunar catalogue

Michael Bennett[1] · Cameron Walker[2] · Evan Anderson[3]

[1] Cyprus · [2] Canada · [3] United States

Published: June 17, 2026 | Article 851

Abstract

Can one exponent summarize the recorded large-crater diameter tail? Using 6,973 circular diameters of at least 20 km from a fixed public lunar catalogue, we compare a Pareto survival function with a continuous two-exponent alternative joined at 80 km. The estimated exponents are 1.666 below and 2.408 above the join. Holding out entire equal-area cells gives a mean log-density gain of 0.00449 nats per record, with one of sixteen cells favoring the single exponent. Better aggregate scoring does not ensure threshold calibration: the two-exponent model overpredicts the observed 583 records above 80 km by 110. A prespecified 160 km join produces negative held-cell gain. These are limited descriptive improvements within acquired catalogue records, without identification of impact rates or completeness.

A fixed slope change gives limited spatial transfer in the large-crater diameter tail of a lunar catalogue

Abstract

Can one exponent summarize the recorded large-crater diameter tail? Using 6,973 circular diameters of at least 20 km from a fixed public lunar catalogue, we compare a Pareto survival function with a continuous two-exponent alternative joined at 80 km. The estimated exponents are 1.666 below and 2.408 above the join. Holding out entire equal-area cells gives a mean log-density gain of 0.00449 nats per record, with one of sixteen cells favoring the single exponent. Better aggregate scoring does not ensure threshold calibration: the two-exponent model overpredicts the observed 583 records above 80 km by 110. A prespecified 160 km join produces negative held-cell gain. These are limited descriptive improvements within acquired catalogue records, without identification of impact rates or completeness.

Introduction

A tail exponent compresses a diameter distribution into an unusually portable number. Once quoted without its lower diameter limit or evaluation rule, that number can appear to describe every large crater equally well. This paper asks a narrower numerical question: within one acquired lunar catalogue and a fixed cohort of circular diameters of at least 20 km, how restrictive is one constant Pareto survival exponent, and does a fixed additional exponent improve scores in geographic cells excluded from fitting? The distinction matters because an attractive global curve can conceal errors at particular diameters and locations. We evaluate the restriction and the transfer of its relaxation separately.

Crater size distributions have long informed planetary interpretation. Hartmann (1965) connected crater statistics with changing meteoritic flux, while the Crater Analysis Techniques Working Group (1978) formalized distinctions between cumulative and differential presentations. Neither history licenses treating a catalogue distribution as a direct impact-rate measurement. Neukum et al. (1975) already examined diameter-dependent structure in lunar size distributions. Their studied ranges and calibrations differ from the present large-diameter cohort; their relevance is that constant slope is an empirical restriction rather than a default physical law. Global nonuniformity is also established background, including spatial patterns studied by Head et al. (2010) and mechanisms considered by Le Feuvre and Wieczorek (2011).

The statistical question is similarly specific. Pareto distributions are a useful model family within the broader literature on power laws (Newman, 2005), but a straight-looking logarithmic plot is insufficient evidence for one. Clauset et al. (2009) distinguish estimation from assessing fit and alternatives. Here the alternative is deliberately small: two positive survival exponents with a prespecified 80 km join, keeping the 20 km lower limit, the diameter measurement, and every eligible record fixed. Proposition one concerns the signed difference between the fitted exponents. Proposition two concerns the paired log-density gain on omitted geographic cells. A positive in-sample gain follows from nesting; a positive held-cell gain requires data and can fail.

This question differs from changing coordinate rules for nearest neighbors, comparing circular and elliptical diameter measurements, or counting threshold-conditioned hemispheric populations. Only the circular diameter enters the response model. Geography defines the units withheld from fitting, rather than a new ranking of hemisphere counts. The contribution is an explicit assessment of a numerical simplification and its boundaries, including unfavorable calibration and sensitivity results. It is not the discovery that lunar crater populations are heterogeneous.

Source and finite record population

We use the complete acquired 2018-08-15 version of the publicly distributed Robbins lunar crater bundle. Robbins (2019) describes the global catalogue, its measurement procedure, comparisons, and spatial analyses. The present estimand is conditional on the bytes acquired from that distribution. The CSV contains 1,296,796 unique crater identifiers and 21 columns. We project four fields: CRATER_ID, LAT_CIRC_IMG, LON_CIRC_IMG, and DIAM_CIRC_IMG. Circular diameter is in kilometers; latitude and longitude are in decimal degrees. The actual longitudes occupy the released 0-360 convention. One literal 360 value is normalized to zero for geographic assignment. This operation does not change a diameter.

All four required fields are supported in the source: no identifier duplicates, missing values, nonfinite numeric values, nonpositive diameters, or invalid required coordinates are found. The fixed diameter filter retains 6,973 records, ranging from 20.0014 to 2,491.87 km. There are 583 records at least 80 km across. The largest basins remain included, so the analysis does not obtain a cleaner tail by deleting conspicuous extremes. No record lies exactly on any declared 20, 40, 80, 160, or 320 km threshold. These are counts of released records under a stated measurement rule, without an audit of every crater against independent imagery.

Integrity checks delimit the version more precisely than a catalogue name alone. The downloaded archive contains 96,227,201 bytes, and all ZIP member CRC checks pass. The analyzed CSV has 238,483,500 bytes and SHA-256 39a419e0ca627b792be1a9f761c3fa1372094cdc5bc0db88afd365c07af8ad62. Its MD5 agrees with bundle_checksums.txt. The observational XML names this same CSV but supplies a different MD5. That metadata discrepancy remains unresolved; it is not presented as agreement or repaired by altering the data. The results therefore refer to the acquired CSV whose checksum matches the bundle manifest. The source record identifies a public scientific distribution; the current catalogue describes its PDS4 status as in review, which we do not upgrade to completed archival review.

The protocol was frozen after reading documentation, related studies, and structural support counts. It was not externally preregistered or blind to all catalogue outcomes. The 20 km cohort, 80 km primary join, geographic partitions, diagnostics, and two secondary joins were fixed before the fitted model contrasts were examined. This chronology is recorded so that a useful secondary analysis does not acquire the evidential status of an untouched prospective test.

Working tail models and evaluation

Let D denote a retained diameter, $z = \ln(D/20)$, and $c = \ln(80/20)$. The single-exponent survival function is $S(D) = \exp(-\beta z)$. Its maximizing coefficient is $\beta = N / \text{sum}(z)$. The two-exponent function is $S(D) = \exp[-\beta L \min(z,c) - \beta H \max(z-c,0)]$. This survival function is continuous at 80 km, although its density may jump there. Below the join, the density in z is $\beta L \exp(-\beta L z)$; at or above it, the density is $\beta H \exp[-\beta L c-\beta H(z-c)]$. The fitted coefficients are $\beta L = nL / \text{sum}[\min(z,c)]$ and $\beta H = nH / \text{sum}[\max(z-c,0)]$, where nL and nH count diameters below and at or above the join.

These coefficients describe cumulative survival. The corresponding differential diameter-density exponent within a segment is β plus one. The primary effect is βH minus βL. All fits maximize the sum of individual log densities, treating it as a working approximation to this finite catalogue. We do not infer that crater locations or sizes are independent physical draws. Log scores are evaluated in z; the common change-of-variable term for conversion to diameter density cancels from every paired comparison. Thus the reported score difference does not depend on choosing kilometers instead of another consistent diameter unit.

Using individual diameters avoids fitting least squares to correlated cumulative count ordinates. Robbins et al. (2018) discuss likelihood-based crater-distribution analysis; Goldstein et al. (2004) and Bauke (2007) explain difficulties with graphical power-law estimation. Binning also discards information, as examined by Virkar and Clauset (2014). Our four threshold counts are consequently diagnostics of fitted survival, not four independent regression observations. A curve's apparent straightness and the number of points drawn on it do not increase the independent support for a coefficient.

There is no search for a best lower cutoff or join. Hill (1975) and Pickands (1975) develop inference focused on distribution tails and extreme order statistics under explicit conditions. The present coefficients instead summarize the entire declared $D \geq 20$ km interval. Calling them tail exponents does not establish the asymptotic assumptions of those estimators, identify a limiting tail index, or justify extrapolation beyond the observed range. This distinction prevents a fixed-range descriptive fit from silently becoming an extreme-event prediction.

For geographic evaluation, four equal intervals of sin(latitude) are crossed with four longitude quadrants. These sixteen cells have equal spherical areas, although their record counts differ. Each cell is omitted in turn; both models are fitted on the other fifteen, and every omitted record receives a paired log-density score. The main summary pools all 6,973 held-cell differences and divides by their count. It therefore weights records equally and does not weight the sixteen cells equally. Positive values favor the additional exponent. Figure 1 shows the common cohort and this separation of fitting from evaluation.

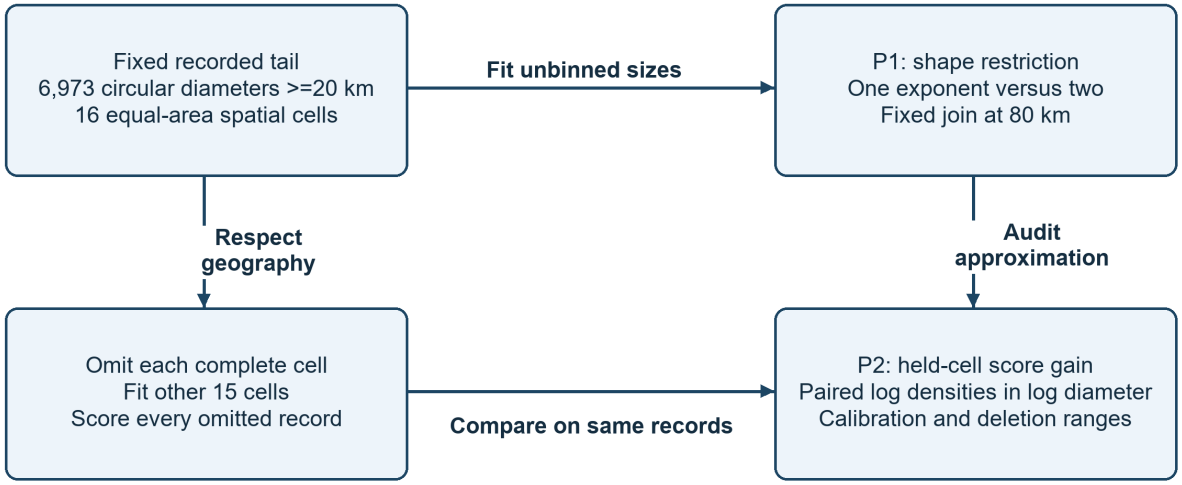

Figure 1. Fixed-source analysis. One diameter cohort feeds both working likelihoods; entire geographic cells are withheld before paired scoring. Full-cohort calibration and prespecified sensitivity analyses answer separate questions from held-cell score comparison. The model diagram does not assert an independent impact process.

Spatial withholding is a substantive evaluation boundary. Michael et al. (2012) analyze spatial randomness and clustering in crater-distribution measurements, while McEwen and Bierhaus (2006) explain why secondary cratering can complicate independence assumptions. These sources do not establish secondary-crater labels for our retained records. They motivate resisting an independence claim merely because the table has many rows. Roberts et al. (2017) discuss why structured evaluation splits need to reflect the prediction question. Our held cells assess transfer within this particular lunar partition; neighboring cells can remain dependent, and no claim of transport to another catalogue or planetary body follows.

We report the minimum and maximum fitted effects after deleting each complete cell. For the aggregate held-cell score, a separate composition range removes each cell's already computed held-out contributions without refitting the remaining folds. Both are deterministic sensitivity ranges. They are not confidence intervals, tests of a random impact process, or estimates of uncertainty in a future catalogue. Hurlbert (1984) motivates distinguishing replication from numerous subsamples, and Bengio and Grandvalet (2004) show why overlapping training sets complicate cross-validation variance estimation. We therefore avoid row-based standard errors and a naive standard error across folds.

Two additional diagnostics retain the same record population. First, fitted and observed survival are compared at 40, 80, 160, and 320 km. Second, the maximum absolute empirical-CDF discrepancy is evaluated at every distinct observed diameter, on both sides of ties. No Kolmogorov-Smirnov p-value is attached. Prespecified sensitivities move the join to 40 or 160 km, and separately replace the sixteen cells with eight equal-area cells using two sin(latitude) bands and four longitude quadrants. Fotheringham and Wong (1991) demonstrate the broader importance of geographic aggregation choices; our second partition probes one such choice rather than resolving every possible zoning effect.

Results

The single fitted survival exponent is 1.70995. With the primary 80 km join, the lower and upper exponents are 1.66589 and 2.40805, respectively. Their difference is 0.74217, remaining positive in all sixteen complete-cell deletion fits. That stability concerns the declared finite partition. It does not supply an independent sampling distribution for lunar impacts. Table 1 reports both the coefficient contrast and the evaluation result so that stable slope separation cannot substitute for validation.

Table 1. Primary coefficients and paired scoring on the fixed 6,973-record cohort. Ranges are deterministic complete-cell deletions or score-composition deletions, as identified; none is a confidence interval.

Quantity	Full or pooled estimate	Deterministic range	Support or scope
Single survival exponent β	1.70995	1.69616 to 1.73759	16 deletion refits
Lower exponent βL	1.66589	1.65156 to 1.69489	6,390 records below 80 km
Upper exponent βH	2.40805	2.36060 to 2.46581	583 records at least 80 km
βH - βL	0.74217	0.70152 to 0.81425	Same deletion refits
In-sample mean log-score gain	0.004700	Not an uncertainty estimate	Nesting guarantees nonnegative optimized gain
Held-cell mean log-score gain	0.004494	0.004085 to 0.005436	6,973 paired scores; composition deletion
Maximum empirical-CDF discrepancy, single	0.032565	Not estimated	Same-cohort fit diagnostic
Maximum empirical-CDF discrepancy, two exponents	0.023087	Not estimated	Same-cohort fit diagnostic

The positive held-cell gain is small on the per-record log scale. Fifteen cells favor two exponents, while cell 8 has a mean gain of -0.01777 nats per record. This cell contains 283 eligible records, including 13 at least 80 km, and occupies latitude 0 to 30 degrees and longitude 0 to 90 degrees under the released coordinate convention. Its models were fitted on 6,690 other records. The negative result is retained, without assigning it a geological cause. Test-cell counts range from 142 to 705, so a statement about fifteen positive cells is different from the record-weighted score estimand.

Full-cohort calibration gives a second reason for restraint. Although the two-exponent fit lowers the maximum empirical-CDF discrepancy, it predicts 692.56 records above 80 km against 583 observed, while the single exponent predicts 651.52. The more flexible fit is worse at that threshold. The upper segment balances its event count against logarithmic diameter exposure; continuity fixes its entry mass through the lower-segment coefficient. Consequently, the fitted join count need not equal the observed join count. This is a property of the specified likelihood, not evidence of missing catalogue entries.

Table 2. Observed threshold counts and model-implied counts from full-cohort fits. Predictions are N times fitted conditional survival and are left as noninteger expectations; these thresholds are nested diagnostics.

Diameter threshold (km)	Observed records	Single exponent	Two exponents, 80 km join
40	2,329	2,131.44	2,197.55
80	583	651.52	692.56
160	97	199.15	130.48
320	32	60.87	24.58

At 160 km the two-exponent model reduces a substantial overprediction, but at 320 km it falls below the observed count. Figure 2 presents the entire retained diameter range alongside all sixteen held-cell gains. The remaining departures are relevant even when the average log score improves: neither curve reproduces every part of the recorded tail. The fixed fit does not justify replacing the observed number of large basins with its own smooth expected count.

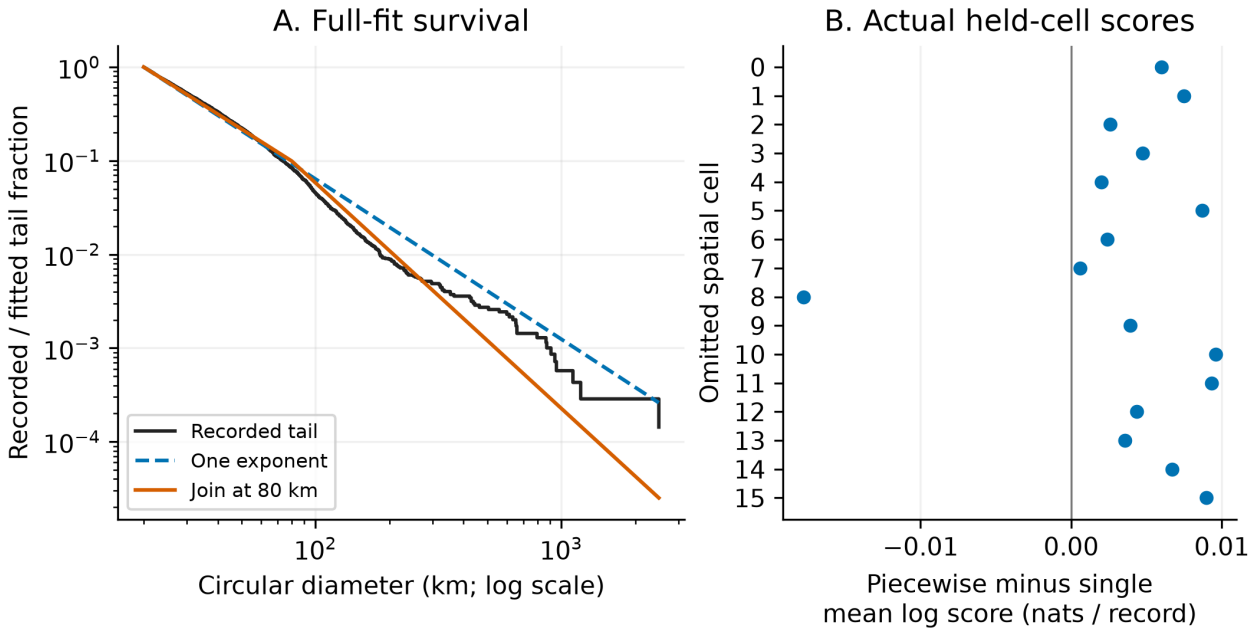

Figure 2. Left: empirical conditional survival and full-cohort working fits on logarithmic diameter and survival axes, including the largest recorded basins. Right: mean paired log-score gain for each omitted cell; positive favors two exponents. The zero line is a comparison reference, not a significance threshold. No sampling confidence bands are implied.

The join sensitivities materially qualify the primary result. A 40 km join yields a positive held-cell gain of 0.006716 with thirteen positive cells. A 160 km join yields -0.000270, with only three positive cells. Under the eight-cell partition, the primary 80 km join retains a gain of 0.004450 and seven positive cells. Table 3 keeps the negative alternative beside the positive ones. These choices were specified in advance of model contrasts; they are not a search from which the best join is selected for deployment.

Table 3. Prespecified sensitivity analyses; all use the identical 6,973-record D ≥ 20 km cohort. Exponent contrasts are full-cohort estimates; log scores use the stated geographic withholding.

Join (km)	Equal-area cells	βH - βL	Held-cell gain (nats/record)	Positive cells
80, primary	16	0.74217	0.004494	15/16
40	16	0.45926	0.006716	13/16
160	16	0.03195	-0.000270	3/16
80	8	0.74217	0.004450	7/8

Interpretation and limits

The two propositions receive different levels of support. The primary finite-record slope contrast is positive and insensitive to deleting one declared cell. The average held-cell score also favors the 80 km alternative, but its magnitude is modest, one cell reverses it, and the upper join sensitivity loses it. Stone (1974) distinguishes prediction assessment from fitting choices; that distinction is essential here because nesting makes the optimized in-sample improvement uninformative about transfer by itself. The primary model earns a limited comparative description, rather than a general claim that a second exponent is always preferable.

Nor does the lower CDF discrepancy certify a true distributional law. White (1982) separates a likelihood optimum under misspecification from correctness of the fitted family. In this study, both candidate families can be inaccurate summaries of a catalogue containing diverse ages, terrains, preservation histories, and measurement circumstances. Those possible sources of composition are not identified or partitioned by the four projected fields. Geographic withholding makes some spatial transfer visible, but it neither removes within-cell dependence nor explains the negative cell. A map of coefficient performance is not a map of impact intensity.

The unbounded Pareto formulas are also computational approximations within the observed interval. A physically finite body cannot support arbitrary extrapolation in crater diameter. Aban et al. (2006) show that explicit upper truncation changes Pareto estimation, while Coles (2001) develops extreme-value modeling with attention to the assumptions needed for extreme statements. We do not estimate an upper endpoint, forecast an unprecedented basin, or interpret the largest observed diameter as that endpoint. The disagreement at 320 km is enough to show that a broadly improved fit can remain unsuitable for a particular large-size query.

Chronological inference would require further evidence. Michael et al. (2016) develop Poisson timing analysis for crater-based surface dating; our working density comparison contains neither surface ages nor a timing calibration. It cannot estimate formation rates, resurfacing dates, or the catalogue's probability of detecting a crater. Likewise, no independent measurement-error model is available. A fitted circle's residual scatter would not capture all rim interpretation, image-resolution, preservation, and identification uncertainty. Coefficient deletion ranges therefore describe sensitivity to geographic composition only. Their apparent narrowness must not be presented as total scientific uncertainty.

A practical consequence concerns the information accompanying a scientific numerical answer. Lunar-R1 uses difficulty-aware reasoning-budget allocation, and its Appendix A.1 distinguishes the construction of lunar knowledge and evaluation material (Xiao et al., 2026). The present catalogue experiment does not evaluate that model, its reinforcement learning, or its hardware efficiency. It supplies a concrete boundary that a lunar reasoning answer would need to preserve: a fitted exponent belongs to a particular diameter

range, source version, model form, and evaluation rule. Omitting the unfavorable join-count calibration while quoting the positive mean score would change the answer's scientific meaning. More reasoning tokens alone are not evidence that those qualifications have been preserved.

Reproducibility and conclusion

The released analysis projection retains crater identifiers, the three numeric source fields, cell assignment, transformed diameter, and each pair of held-cell log scores. Results include every declared sensitivity and the negative scores. NumPy supports the array computations (Harris et al., 2020); a separate SciPy optimization checks the closed-form likelihood optimum (Virtanen et al., 2020). An independent standard-library data pass recomputes source and cohort counts, sufficient statistics, coefficients, all fold gains, threshold counts, and the actual output CSV hash without importing the main estimator. All ten recorded verification checks pass. This is computational agreement, rather than independent confirmation of crater identity or geological interpretation.

Following the distinction between accessible computation and broader scientific replication discussed by Peng (2011), the evidence package records source hashes, the checksum disagreement, the frozen secondary-analysis plan, analysis scope, and verification outputs. Procedural scientific-writing guidance was used to organize evidence and accountability, with its software provenance acknowledged through Kassis et al. (2026). Human author review of scientific claims, authorship, declarations, and submission approval remains pending. No new language-model inference or crater-detection experiment contributes results.

For the acquired large-crater cohort, a fixed 80 km slope change gives a stable positive exponent contrast and a small positive mean score on omitted geographic cells. It nevertheless worsens calibration at its own join, fails in one cell, and is not supported by the 160 km join sensitivity. Those simultaneous results define the useful conclusion: one exponent is restrictive, but relaxing it in this particular way produces a limited descriptive improvement whose range and validation boundary must remain attached to the number.

References

1. Michael, G.G.; Kneissl, T.; Neesemann, A.. (2016). Planetary surface dating from crater size-frequency distribution measurements: Poisson timing analysis. Icarus. https://doi.org/10.1016/j.icarus.2016.05.019
2. Le Feuvre, Mathieu; Wieczorek, Mark A.. (2011). Nonuniform cratering of the Moon and a revised crater chronology of the inner Solar System. Icarus. https://doi.org/10.1016/j.icarus.2011.03.010
3. Xiao, Xin-Yu; Tian, Ye; He, Zhixian; Chen, Qilin; Wang, Shiqi; Xia, Qianchen. (2026). Lunar-R1: Towards Efficient Reasoning of Large Language Models for Lunar Exploration via Reinforcement Learning. Proceedings of the 2026 Conference on Empirical Methods in Natural Language Processing: Industry Track. https://aclanthology.org/2026.emnlp-industry.5/.
4. Coles, Stuart. (2001). An Introduction to Statistical Modeling of Extreme Values. Springer Series in Statistics. https://doi.org/10.1007/978-1-4471-3675-0
5. Virkar, Yogesh; Clauset, Aaron. (2014). Power-law distributions in binned empirical data. The Annals of Applied Statistics. https://doi.org/10.1214/13-aoas710
6. Stone, M.. (1974). Cross-Validatory Choice and Assessment of Statistical Predictions. Journal of the Royal Statistical Society Series B: Statistical Methodology. https://doi.org/10.1111/j.2517-6161.1974.tb00994.x
7. Hartmann, William K.. (1965). Secular changes in meteoritic flux through the history of the solar system. Icarus. https://doi.org/10.1016/0019-1035(65)90062-x
8. Clauset, Aaron; Shalizi, Cosma Rohilla; Newman, M. E. J.. (2009). Power-Law Distributions in Empirical Data. SIAM Review. https://doi.org/10.1137/070710111
9. Aban, Inmaculada B; Meerschaert, Mark M; Panorska, Anna K. (2006). Parameter Estimation for the Truncated Pareto Distribution. Journal of the American Statistical Association. https://doi.org/10.1198/016214505000000411

10. Newman, MEJ. (2005). Power laws, Pareto distributions and Zipf's law. Contemporary Physics. https://doi.org/10.1080/00107510500052444

11. Head, James W.; Fassett, Caleb I.; Kadish, Seth J.; Smith, David E.; Zuber, Maria T.; Neumann, Gregory A.; Mazarico, Erwan. (2010). Global Distribution of Large Lunar Craters: Implications for Resurfacing and Impactor Populations. Science. https://doi.org/10.1126/science.1195050

12. Crater Analysis Techniques Working Group. (1978). Standard Techniques for Presentation and Analysis of Crater Size-Frequency Data. NASA Technical Memorandum 79730. https://ntrs.nasa.gov/citations/19780014063

13. Peng, Roger D.. (2011). Reproducible Research in Computational Science. Science. https://doi.org/10.1126/science.1213847

14. Hill, Bruce M.. (1975). A Simple General Approach to Inference About the Tail of a Distribution. The Annals of Statistics. https://doi.org/10.1214/aos/1176343247

15. Bengio, Yoshua; Grandvalet, Yves. (2004). No Unbiased Estimator of the Variance of K-Fold Cross-Validation. Journal of Machine Learning Research. https://www.jmlr.org/papers/v5/grandvalet04a.html

16. Fotheringham, A S; Wong, D W S. (1991). The Modifiable Areal Unit Problem in Multivariate Statistical Analysis. Environment and Planning A: Economy and Space. https://doi.org/10.1068/a231025

17. Hurlbert, Stuart H.. (1984). Pseudoreplication and the Design of Ecological Field Experiments. Ecological Monographs. https://doi.org/10.2307/1942661

18. Pickands, James, III. (1975). Statistical Inference Using Extreme Order Statistics. The Annals of Statistics. https://doi.org/10.1214/aos/1176343003

19. Virtanen, Pauli; Gommers, Ralf; Oliphant, Travis E.; Haberland, Matt; Reddy, Tyler; Cournapeau, David; Burovski, Evgeni; Peterson, Pearu; Weckesser, Warren; Bright, Jonathan; van der Walt, Stéfan J.; Brett, Matthew; Wilson, Joshua; Millman, K. Jarrod; Mayorov, Nikolay; Nelson, Andrew R. J.; Jones, Eric; Kern, Robert; Larson, Eric; Carey, C J; Polat, İlhan; Feng, Yu; Moore, Eric W.; VanderPlas, Jake; Laxalde, Denis; Perktold, Josef; Cimrman, Robert; Henriksen, Ian; Quintero, E. A.; Harris, Charles R.; Archibald, Anne M.; Ribeiro, Antônio H.; Pedregosa, Fabian; van Mulbregt, Paul; SciPy 1.0 Contributors; Vijaykumar, Aditya; Bardelli, Alessandro Pietro; Rothberg, Alex; Hilboll, Andreas; Kloeckner, Andreas; Scopatz, Anthony; Lee, Antony; Rokem, Ariel; Woods, C. Nathan; Fulton, Chad; Masson, Charles; Häggström, Christian; Fitzgerald, Clark; Nicholson, David A.; Hagen, David R.; Pasechnik, Dmitrii V.; Olivetti, Emanuele; Martin, Eric; Wieser, Eric; Silva, Fabrice; Lenders, Felix; Wilhelm, Florian; Young, G.; Price, Gavin A.; Ingold, Gert-Ludwig; Allen, Gregory E.; Lee, Gregory R.; Audren, Hervé; Probst, Irvin; Dietrich, Jörg P.; Silterra, Jacob; Webber, James T; Slavič, Janko; Nothman, Joel; Buchner, Johannes; Kulick, Johannes; Schönberger, Johannes L.; de Miranda Cardoso, José Vinícius; Reimer, Joscha; Harrington, Joseph; Rodríguez, Juan Luis Cano; Nunez-Iglesias, Juan; Kuczynski, Justin; Tritz, Kevin; Thoma, Martin; Newville, Matthew; Kümmerer, Matthias; Bolingbroke, Maximilian; Tartre, Michael; Pak, Mikhail; Smith, Nathaniel J.; Nowaczyk, Nikolai; Shebanov, Nikolay; Pavlyk, Oleksandr; Brodtkorb, Per A.; Lee, Perry; McGibbon, Robert T.; Feldbauer, Roman; Lewis, Sam; Tygier, Sam; Sievert, Scott; Vigna, Sebastiano; Peterson, Stefan; More, Surhud; Pudlik, Tadeusz; Oshima, Takuya; Pingel, Thomas J.; Robitaille, Thomas P.; Spura, Thomas; Jones, Thouis R.; Cera, Tim; Leslie, Tim; Zito, Tiziano; Krauss, Tom; Upadhyay, Utkarsh; Halchenko, Yaroslav O.; Vázquez-Baeza, Yoshiki. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods. https://doi.org/10.1038/s41592-019-0686-2.

20. White, Halbert. (1982). Maximum Likelihood Estimation of Misspecified Models. Econometrica. https://doi.org/10.2307/1912526

21. Robbins, Stuart J.. (2019). A New Global Database of Lunar Impact Craters >1-2 km: 1. Crater Locations and Sizes, Comparisons With Published Databases, and Global Analysis. Journal of Geophysical Research: Planets. https://doi.org/10.1029/2018je005592

22. Neukum, Gerhard; König, Beate; Arkani-Hamed, Jafar. (1975). A study of lunar impact crater size-distributions. The Moon. https://doi.org/10.1007/bf00577878

23. McEwen, Alfred S.; Bierhaus, Edward B.. (2006). THE IMPORTANCE OF SECONDARY CRATERING TO AGE CONSTRAINTS ON PLANETARY SURFACES. Annual Review of Earth and Planetary Sciences. https://doi.org/10.1146/annurev.earth.34.031405.125018

24. Roberts, David R.; Bahn, Volker; Ciuti, Simone; Boyce, Mark S.; Elith, Jane; Guillera-Arroita, Gurutzeta; Hauenstein, Severin; Lahoz-Monfort, José J.; Schröder, Boris; Thuiller, Wilfried; Warton, David I.; Wintle, Brendan A.; Hartig, Florian; Dormann, Carsten F.. (2017). Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. Ecography, 40, 913-929. https://doi.org/10.1111/ecog.02881.

25. Kassis, Timothy; Agarwal, Vinayak; He, Yuhuan; Patel, Darshil; Brueckner, Aubrey M.. (2026). Scientific Agent Skills: A Library of Procedural Knowledge for Research Agents. arXiv preprint. https://arxiv.org/abs/2609.00065

26. Robbins, Stuart J.; Riggs, Jamie D.; Weaver, Brian P.; Bierhaus, Edward B.; Chapman, Clark R.; Kirchoff, Michelle R.; Singer, Kelsi N.; Gaddis, Lisa R.. (2018). Revised recommended methods for analyzing crater size-frequency distributions. Meteoritics & Planetary Science. https://doi.org/10.1111/maps.12990

27. Bauke, H.. (2007). Parameter estimation for power-law distributions by maximum likelihood methods. The European Physical Journal B. https://doi.org/10.1140/epjb/e2007-00219-y

28. Harris, Charles R.; Millman, K. Jarrod; van der Walt, Stéfan J.; Gommers, Ralf; Virtanen, Pauli; Cournapeau, David; Wieser, Eric; Taylor, Julian; Berg, Sebastian; Smith, Nathaniel J.; Kern, Robert; Picus, Matti; Hoyer, Stephan; van Kerkwijk, Marten H.; Brett, Matthew; Haldane, Allan; del Río, Jaime Fernández; Wiebe, Mark; Peterson, Pearu; Gérard-Marchant, Pierre; Sheppard, Kevin; Reddy, Tyler; Weckesser, Warren; Abbasi, Hameer; Gohlke, Christoph; Oliphant, Travis E.. (2020). Array programming with NumPy. Nature. https://doi.org/10.1038/s41586-020-2649-2.

29. Michael, G.G.; Platz, T.; Kneissl, T.; Schmedemann, N.. (2012). Planetary surface dating from crater size-frequency distribution measurements: Spatial randomness and clustering. Icarus. https://doi.org/10.1016/j.icarus.2011.11.033

30. Goldstein, M. L.; Morris, S. A.; Yen, G. G.. (2004). Problems with fitting to the power-law distribution. The European Physical Journal B. https://doi.org/10.1140/epjb/e2004-00316-5