Journal of Data Efficiency
Volume 1 · Number 1 · 2026

Which Seconds Were Observed? Reference-Trajectory Gaps and Distance-Rate Summaries

Kevin Evans[1] · Joseph Foster[2] · James Harrison[3]

[1] Croatia · [2] Romania · [3] Bulgaria

Original publication: October 24, 2026 | Article 602

Abstract

A trajectory file's first and last timestamps can conceal long intervals without intervening reference poses. Across 52 public TUM RGB-D pose files, intervals longer than 50 milliseconds accounted for 5.26% of elapsed time on an equal-file basis. Eighteen files contained such intervals; the largest individual gap was about 130 seconds. Restricting both chord distance and duration to shorter intervals changed the mean distance rate from 0.2034 to 0.2117 meters per second, with an uncertain mean difference. The audit makes the supported interval explicit while leaving the missing trajectory unresolved.

Availability audit of official camera reference trajectories. Excluding long intervals identifies a different observed-time summary; it does not reconstruct the unobserved physical path.

Which Seconds Were Observed? Reference-Trajectory Gaps and Distance-Rate Summaries

> Availability audit of official camera reference trajectories. Excluding long intervals identifies a different observed-time summary; it does not reconstruct the unobserved physical path.

Abstract

A trajectory file's first and last timestamps can conceal long intervals without intervening reference poses. Across 52 public TUM RGB-D pose files, intervals longer than 50 milliseconds accounted for 5.26% of elapsed time on an equal-file basis. Eighteen files contained such intervals; the largest individual gap was about 130 seconds. Restricting both chord distance and duration to shorter intervals changed the mean distance rate from 0.2034 to 0.2117 meters per second, with an uncertain mean difference. The audit makes the supported interval explicit while leaving the missing trajectory unresolved.

A path between observations is an assumption

Joining two recorded camera positions with a straight line is a useful numerical operation. It does not establish that the camera traveled along that line or that its movement was observed throughout the intervening time. When reference measurements disappear for a substantial interval, the sum of adjacent chords and the time span between the first and last records can conceal different kinds of missing information.

We examine how much of an available reference-file span lies outside closely observed intervals and how a distance-divided-by-duration summary changes when its numerator and denominator use the same retained intervals. The primary operational definition retains adjacent observations at most fifty milliseconds apart. A twenty-millisecond alternative is calculated separately to show sensitivity to the chosen definition.

Meng, 2026 addresses scale-calibrated camera-motion transfer. We do not test that transfer method or estimate its parallax number. The connection is a requirement for interpreting geometric summaries: the time and position support underlying a distance claim should be stated before the number is used to characterize camera motion. A reference gap is not resolved merely by having metric coordinates at its endpoints.

The goal is an availability audit rather than a corrected estimate of continuous travel. A short-interval summary excludes some periods; an all-span summary includes them through endpoint chords and elapsed time. Their difference is informative about the reporting convention. It does not identify which summary is closer to an unobserved physical path that the reference file does not supply.

Acquiring the public text subset

The official TUM RGB-D resource supplies RGB-D sequences and camera reference trajectories (Sturm et al., 2012). We download the separately linked ground-truth text files, leaving the much larger image and bag archives outside this acquisition. The captured listing contains 53 distinct public trajectory-text links. Their total downloaded size is 42,008,171 bytes, and original files are retained with hashes and source URLs.

One file contains comments but no numerical pose rows. The remaining 52 contain 615,924 raw pose records. We record the comment-only file explicitly instead of counting it as a usable trajectory. No

withheld validation ground truth is obtained. The public source license permits attributed use under its stated conditions, and the local text bundle retains attribution and the original byte content.

The official format describes timestamp, three position coordinates and a quaternion (TUM Computer Vision Group, 2026). Position refers to the optical center of the color camera in the motion-capture world frame, with time expressed in seconds since the Unix epoch. This point definition matters: the rows are not an estimated vehicle center, an arbitrary sensor origin or newly inferred scene geometry.

The collection contains different experimental roles, including motion sequences and calibration recordings. We include all usable public numerical files rather than select only ones described as visually interesting. The resulting equal-file average describes this acquired collection. It is not an estimate of the motion characteristics of a randomly sampled population of cameras, rooms or autonomous systems.

Other reference resources, such as KITTI, EuRoC and TUM VI, have their own acquisition and evaluation conventions (Geiger et al., 2012; Burri et al., 2016; Schubert et al., 2018). They provide context for the need to inspect support, not additional observations in the present denominator. In particular, the existence of ground truth in one part of another benchmark does not establish continuous reference availability in these text files.

Resolving ambiguous timestamps without choosing a preferred pose

The initial structural check found repeated timestamps carrying different pose values. A strict original-order rule would have rejected nineteen entire numerical files because their timestamps were not strictly increasing. Before writing the manuscript, we replaced that overly coarse admission rule with explicit timestamp curation and preserved the initial audit as an exploratory artifact.

The final rule first checks for exact duplicate rows, then removes every row at a timestamp that still carries different pose values. No exact row duplicates are present. Seventy-four ambiguous timestamps involve 149 rows, all of which are excluded from the analysis copy. We do not select the first value, average conflicting poses or infer which value a motion-capture system intended to report.

The remaining observations are ordered chronologically, yielding 615,775 poses across all 52 numerical files. The original downloads remain intact. Curation has its own record of removed timestamp groups, row counts and the processed archive hash. This makes it possible to distinguish an observed source record from the unique-time series used in the calculations.

Removing ambiguous timestamps can itself affect nearby interval lengths. The gap statistics therefore belong to the curated series, not to an assertion that every long interval existed independently of processing. The affected row count is small relative to the inventory, but its treatment remains part of the design. We do not assume that ambiguity or missingness follows a random mechanism.

Missing-data analysis distinguishes observed support from assumptions used to recover unobserved values (Little and Rubin, 2002). Here no recovery model is fitted. Keeping ambiguous rows out of an arithmetic interval calculation is a data-definition decision; it does not establish accurate clocks or eliminate uncertainty in the retained physical reference measurements.

Distance, duration and the fifty-millisecond gate

Within each curated file, we calculate the time between adjacent records and the Euclidean distance between their camera-center positions. Summing all distances gives the all-adjacent chord path. Summing all time differences gives the elapsed span. Their ratio is an all-span chord rate, expressed in meters per second.

For the supported-interval summary, we retain only adjacent pairs separated by at most 0.05 seconds. We sum their chord distances and their durations separately, then divide those retained totals. A long gap contributes neither its endpoint chord nor its duration to this conditional summary. It is not filled with a zero velocity, interpolated trajectory or an estimate of the missing traveled distance.

The unsupported elapsed fraction is the sum of excluded interval durations divided by the full span. It is an operational coverage measure tied to the cutoff. Fifty milliseconds is not asserted to be an official missing-data standard or a universal maximum interval for accurate motion reconstruction. The source's nominal acquisition description does not make every retained timestamp spacing identical.

Trajectory-evaluation research emphasizes the distinction between alignment, temporal correspondence and the metric used to summarize an error (Zhang and Scaramuzza, 2018). Our calculation evaluates none of those error metrics against an estimated trajectory. It examines the reference series itself and retains the time support needed to interpret a motion summary before an algorithm comparison is attempted.

Geometric frameworks define positions, poses and their composition (Hartley and Zisserman, 2004; Barfoot, 2017; Blanco-Claraco, 2021). We use only the supplied camera-center translations for this audit. No orientation parameterization is needed for the chord calculation, and no pose-estimation method is substituted when reference samples are absent.

File-level models

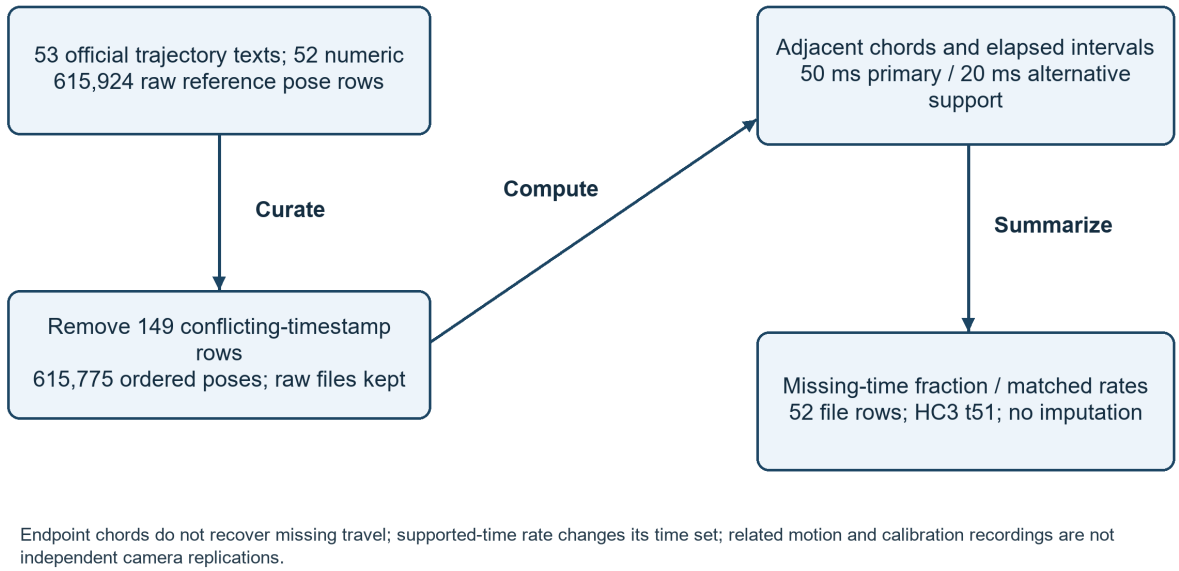


Figure 1. *Timestamp curation precedes interval calculations. Long-interval duration and endpoint chords are separated from the short-interval distance and time used in the conditional rate; missing paths remain unestimated.*

Model 1 summarizes the unsupported elapsed fraction with one intercept and 52 file rows. Model 2 summarizes the supported-interval chord rate minus the all-span chord rate, paired within the same file. Both regressions use HC3 covariance, t51 intervals and no covariates. Intercept-only R-squared is zero apart from numerical rounding.

Heteroskedasticity-consistent inference supplies a working variance across file summaries (White, 1980). Related recordings can share devices, environments and acquisition processes, which this covariance does not fully represent. A clustering approach would also require a defensible grouping and enough

independent groups (Cameron and Miller, 2015). We therefore describe the intervals as conditional summaries rather than established coverage for future motion-capture studies.

Table 1. *Equal-file means, N=52, intercept-only OLS, HC3 standard errors and t51 intervals. M1 uses a fraction of elapsed time; M2 uses meters per second. A positive M2 coefficient describes a changed observation-time estimand, not a correction proven to recover true average speed.*

Model	Term	Estimate (SE)	95% CI	N
M1	Mean outcome/contrast	0.052580 (0.023859)	[0.004681, 0.100478]	52
M1: fit	R-squared = 0.0000	-	-	52
M2	Mean outcome/contrast	0.008345 (0.004868)	[-0.001427, 0.018118]	52
M2: fit	R-squared = 0.0000	-	-	52

Model labels: M1 = Fraction of elapsed span in intervals longer than 50 ms; M2 = Supported-chord rate minus all-span chord rate.

Eighteen of the 52 files contain at least one adjacent interval longer than fifty milliseconds, with 359 such intervals in total. The mean unsupported fraction is 0.052580, standard error 0.023859 and interval [0.004681, 0.100478]. Most files have no interval crossing this cutoff, so a modest collection average can coexist with substantial missing support in individual files.

Across files, total elapsed span is 3,208.285 seconds and retained interval time is 2,890.706 seconds. All adjacent chords total 604.040 meters, of which 576.154 lie on retained intervals. The difference, 27.886 meters, is a sum of long-gap endpoint chords. It is not the distance known to have been traveled during those gaps.

Only 359 of the 615,723 adjacent intervals cross the primary cutoff, yet together they contain about 317.579 seconds. An interval-count coverage statistic would give each tiny step the same weight as a ninety-second gap and could make missing time appear negligible. Duration weighting reveals a different aspect of availability. We retain both the edge count and elapsed-time quantities rather than substitute one coverage denominator for the other.

These spans begin and end at the available reference-file timestamps. They are not established image-recording spans, because RGB timestamps are not associated in this text-only acquisition. A camera may have image observations outside or inside a particular reference period. The present fractions therefore describe time within the pose file and cannot be reported directly as the fraction of video frames lacking reference truth.

The equal-file mean all-span chord rate is 0.203404 meters per second, compared with 0.211749 for retained intervals. Model 2 estimates their mean difference at 0.008345, standard error 0.004868 and interval [-0.001427, 0.018118]. The interval crosses zero. More importantly, the two rates summarize different time sets even where their numerical difference is small.

Long gaps are concentrated rather than evenly distributed

The freiburg2 large-no-loop file has an unsupported fraction of about 0.804063 under the primary gate, with a maximum interval of 90.1232 seconds. The large-with-loop file has a fraction of 0.755927 and a maximum interval of 129.8651 seconds. The official listing already notes incomplete middle reference coverage for these long recordings; the numerical audit makes that support explicit.

Other files also matter. The freiburg2 large-checkerboard calibration recording has an unsupported fraction of 0.559117, while the freiburg2 desk file has 0.257268. We retain them rather than redefine the cohort around the two known large-loop examples. Influence analysis motivates inspecting concentrated contributions without silently deleting them (Cook, 1977).

The twenty-millisecond alternative gives an equal-file unsupported fraction of 0.056159 and a mean rate change of 0.008480 meters per second. These values are close to the primary summaries in this acquired collection. They do not prove that every file is insensitive to the cutoff or that any chosen threshold identifies fully observed continuous motion.

Measurement scale and weighting both affect interpretation (Stevens, 1946; Solon et al., 2015). Averaging file-specific fractions gives each file equal importance; dividing pooled excluded time by pooled span weights long files more heavily. Similarly, an average of file-specific rates differs from one total-distance-to-total-time ratio. We preserve the totals and file rows instead of presenting these quantities as interchangeable.

What a gap-aware summary cannot supply

A chord over an unobserved interval can be much shorter than the actual path if the camera turns or returns near its earlier position. Deleting that chord does not discover the missing distance. It simply avoids inserting the endpoint shortcut into a summary claimed to describe closely observed intervals. The amount of unobserved travel remains unidentified by these data.

Nor is the short-interval sum a perfect continuous-distance measurement. Sampling and position noise can affect adjacent chords even at short time intervals. Measurement-error reasoning distinguishes the reference variable from the intended physical construct (Carroll et al., 2006). Our support gate describes available observations, not a bound on reference-coordinate accuracy or the true path between them.

Agreement-oriented assessment encourages attention to discrepancies with stated units (Bland and Altman, 1986; Lin, 1989). A high correlation between two file-level rates would not establish that they cover the same time intervals. The shared data and different denominators should be examined before a closeness claim is converted into a statement of physical equivalence.

Numerical handling is also documented. Subtracting large epoch timestamps to form short intervals requires attention to floating-point resolution (Higham, 2002; Goldberg, 1991). The recorded values and cutoff are used consistently, and numerical routines are part of the saved environment (Virtanen et al., 2020). We do not round time gaps into a favorable category or revise a cutoff after inspecting which rate difference looks strongest.

Implications for later trajectory analysis

Bundle adjustment, ORB-SLAM2 and COLMAP address pose or structure estimation (Triggs et al., 2000; Mur-Artal and Tardos, 2016; Schonberger and Frahm, 2016). Their outputs are not run or ranked here. A later comparison against these reference files should declare temporal association and gap handling, so that missing reference periods do not silently become algorithm failures or artificial straight-line truth.

Imputation methods such as BRITS and GRIN treat missing values as a prediction problem under their own information and modeling assumptions (Cao et al., 2018; Cini et al., 2021). They are not implemented in this audit. A reconstructed interval would need validation as an estimate, and its uncertainty would belong in the analysis rather than being erased by calling the resulting series complete.

Validation should match the intended camera-motion question and its observation boundary (Roberts et al., 2017). The saved files identify which intervals support the conditional summary and which remain unresolved. Documentation and reusable provenance make that distinction inspectable (Gebru et al., 2021; Wilkinson et al., 2016). The central empirical result is concentrated missing time in an otherwise large pose inventory, accompanied by rate summaries whose distance and duration conventions are stated explicitly.

References

1. Schonberger, Johannes L.; Frahm, Jan-Michael. (2016). Structure-from-Motion Revisited. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). https://doi.org/10.1109/cvpr.2016.445.
2. Mur-Artal, Raul; Tardos, Juan D.. (2016). ORB-SLAM2: an Open-Source SLAM System for Monocular, Stereo and RGB-D Cameras. https://arxiv.org/abs/1610.06475.
3. TUM Computer Vision Group. (2026). RGB-D SLAM Dataset: File Formats [Official documentation; accessed8November2026, captureyearnotbenchmarkyear]. https://cvg.cit.tum.de/data/datasets/rgbd-dataset/file_formats.
4. Blanco-Claraco, José Luis. (2021). A tutorial on SE(3) transformation parameterizations and on-manifold optimization. arXiv:2103.15980; arXivversionyear, originaltechnicalreport2010. https://arxiv.org/abs/2103.15980.
5. Burri, Michael; Nikolic, Janosch; Gohl, Pascal; Schneider, Thomas; Rehder, Joern; Omari, Sammy; Achtelik, Markus W; Siegwart, Roland. (2016). The EuRoC micro aerial vehicle datasets. The International Journal of Robotics Research. https://doi.org/10.1177/0278364915620033.
6. Little, Roderick J. A., and Donald B. Rubin. *Statistical Analysis with Missing Data*. 2nd ed. Wiley, 2002. https://doi.org/10.1002/9781119013563.
7. Triggs, Bill; McLauchlan, Philip F.; Hartley, Richard I.; Fitzgibbon, Andrew W.. (2000). Bundle Adjustment - A Modern Synthesis. Lecture Notes in Computer Science. https://doi.org/10.1007/3-540-44480-7_21.
8. Schubert, David; Goll, Thore; Demmel, Nikolaus; Usenko, Vladyslav; Stückler, Jörg; Cremers, Daniel. (2018). The TUM VI Benchmark for Evaluating Visual-Inertial Odometry. arXiv:1804.06120. https://arxiv.org/abs/1804.06120.
9. Gebru, Timnit, et al. "Datasheets for datasets" *Communications of the ACM*, 2021. https://doi.org/10.1145/3458723.
10. White, Halbert. “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity.” *Econometrica*, vol. 48, no. 4, 1980, pp. 817-838. https://doi.org/10.2307/1912934.
11. Sturm, Jürgen; Engelhard, Nikolas; Endres, Felix; Burgard, Wolfram; Cremers, Daniel. (2012). A benchmark for the evaluation of RGB-D SLAM systems. 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems. https://doi.org/10.1109/IROS.2012.6385773.
12. Wilkinson, Mark D.; Dumontier, Michel; Aalbersberg, IJsbrand Jan; Appleton, Gabrielle; Axton, Myles; Baak, Arie; Blomberg, Niklas; Boiten, Jan-Willem; da Silva Santos, Luiz Bonino; Bourne, Philip E.; Bouwman, Jildau; Brookes, Anthony J.; Clark, Tim; Crosas, Mercè; Dillo, Ingrid; Dumon, Olivier; Edmunds, Scott; Evelo, Chris T.; Finkers, Richard; Gonzalez-Beltran, Alejandra; Gray, Alasdair J.G.; Groth, Paul; Goble, Carole; Grethe, Jeffrey S.; Heringa, Jaap; ’t Hoen, Peter A.C; Hooft, Rob; Kuhn, Tobias; Kok, Ruben; Kok, Joost; Lusher, Scott J.; Martone, Maryann E.; Mons, Albert; Packer, Abel L.; Persson, Bengt; Rocca-Serra, Philippe; Roos, Marco; van Schaik, Rene; Sansone, Susanna-Assunta; Schultes, Erik; Sengstag, Thierry; Slater, Ted; Strawn, George; Swertz, Morris A.; Thompson, Mark; van der Lei, Johan; van Mulligen, Erik; Velterop, Jan; Waagmeester, Andra; Wittenburg, Peter; Wolstencroft, Katherine; Zhao, Jun; Mons, Barend. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data. https://doi.org/10.1038/sdata.2016.18.
13. Cameron, A. Colin, and Douglas L. Miller. “A Practitioner’s Guide to Cluster-Robust Inference.” *Journal of Human Resources*, vol. 50, no. 2, 2015, pp. 317-372. https://doi.org/10.3368/jhr.50.2.317.
14. Roberts, David R., et al. “Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure.” *Ecography*, vol. 40, no. 8, 2017, pp. 913-929. https://doi.org/10.1111/ecog.02881.
15. Meng, Zijie. (2026). ParaScale: Scale-Calibrated Camera-Motion Transfer via a Gauge-Invariant Parallax Number. https://arxiv.org/abs/2606.19805.
16. Geiger, A.; Lenz, P.; Urtasun, R.. (2012). Are we ready for autonomous driving? The KITTI vision benchmark suite. 2012 IEEE Conference on Computer Vision and Pattern Recognition. https://doi.org/10.1109/cvpr.2012.6248074.
17. Hartley, Richard; Zisserman, Andrew. (2004). Multiple View Geometry in Computer Vision. Cambridge University Press. https://doi.org/10.1017/cbo9780511811685.

18. Solon, Gary; Haider, Steven J.; Wooldridge, Jeffrey M.. (2015). What Are We Weighting For?. Journal of Human Resources. https://doi.org/10.3368/jhr.50.2.301.

19. Goldberg, David. (1991). What every computer scientist should know about floating-point arithmetic. ACM Computing Surveys. https://doi.org/10.1145/103162.103163.

20. Lin, Lawrence I-Kuei. (1989). A Concordance Correlation Coefficient to Evaluate Reproducibility. Biometrics. https://doi.org/10.2307/2532051.

21. Cini, Andrea; Marisca, Ivan; Alippi, Cesare. (2021). Filling the G_ap_s: Multivariate Time Series Imputation by Graph Neural Networks. arXiv:2108.00298. https://arxiv.org/abs/2108.00298.

22. Cao, Wei; Wang, Dong; Li, Jian; Zhou, Hao; Li, Lei; Li, Yitan. (2018). BRITS: Bidirectional Recurrent Imputation for Time Series. arXiv:1805.10572. https://arxiv.org/abs/1805.10572.

23. Zhang, Zichao; Scaramuzza, Davide. (2018). A Tutorial on Quantitative Trajectory Evaluation for Visual(-Inertial) Odometry. 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). https://doi.org/10.1109/IROS.2018.8593941.

24. Barfoot, Timothy D.. (2017). State Estimation for Robotics. Cambridge University Press, first edition. https://doi.org/10.1017/9781316671528.

25. Bland, J. Martin, and Douglas G. Altman. "Statistical Methods for Assessing Agreement between Two Methods of Clinical Measurement." *The Lancet*, vol. 327, no. 8476, 1986, pp. 307-310. https://doi.org/10.1016/S0140-6736(86)90837-8.

26. Carroll, Raymond J., David Ruppert, Leonard A. Stefanski, and Ciprian M. Crainiceanu. “Measurement Error in Nonlinear Models.” 2006. https://doi.org/10.1201/9781420010138.

27. Higham, Nicholas J.. (2002). Accuracy and Stability of Numerical Algorithms: Second Edition. Society for Industrial and Applied Mathematics. https://doi.org/10.1137/1.9780898718027.

28. Stevens, S. S.. (1946). On the Theory of Scales of Measurement. Science. https://doi.org/10.1126/science.103.2684.677.

29. Cook, R. Dennis. (1977). Detection of Influential Observation in Linear Regression. Technometrics. https://doi.org/10.1080/00401706.1977.10489493.

30. 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.