Computation and Performance Studies
Volume 1 · Number 1 · 2026

Camera Path Accounting under Sparse Pose Budgets: Translation and Rotation Respond Differently

Adrian Dawson[1] · Mason Hayes[2] · Gavin Bell[3]

[1] Slovenia · [2] Montenegro · [3] Malta

Original publication: October 28, 2026 | Article 643

Abstract

Keeping fewer camera poses saves records, but the retained trajectory may conceal different kinds of movement. Across eleven KITTI reference sequences, reducing an every-second-frame representation to every tenth frame increased mean translation-length undercount by only 0.070 percentage points. The corresponding increase in accumulated rotation undercount was 12.73 points. Both calculations retain sequence endpoints. These results concern variation present in the captured pose arrays; they neither recover continuous physical motion nor show that a sparse trajectory preserves the visual consequences of camera rotation.

Exploratory analysis of public reference-pose arrays. The comparison changes retained anchors; it does not estimate poses or execute a camera-motion transfer model.

Camera Path Accounting under Sparse Pose Budgets: Translation and Rotation Respond Differently

> Exploratory analysis of public reference-pose arrays. The comparison changes retained anchors; it does not estimate poses or execute a camera-motion transfer model.

Abstract

Keeping fewer camera poses saves records, but the retained trajectory may conceal different kinds of movement. Across eleven KITTI reference sequences, reducing an every-second-frame representation to every tenth frame increased mean translation-length undercount by only 0.070 percentage points. The corresponding increase in accumulated rotation undercount was 12.73 points. Both calculations retain sequence endpoints. These results concern variation present in the captured pose arrays; they neither recover continuous physical motion nor show that a sparse trajectory preserves the visual consequences of camera rotation.

1. A short pose file can hide a long turning history

Camera motion is often summarized through a set of positions and orientations. A storage decision that retains only a fraction of these records also changes what can be measured directly from them. Connecting retained positions with straight segments records the endpoint displacement of each interval. It does not record every bend traveled inside the interval. Joining orientations by their shortest relative rotation similarly discards intermediate turns that cancel before the next retained anchor.

The research question is how accumulated translation and accumulated rotation respond to the same reduction in a public camera-pose budget. We compare an anchor every second frame with an anchor every tenth frame, always including the first and final observations. The two outcomes describe different geometric quantities, so agreement on one cannot establish agreement on the other. Our analysis asks about the information represented by the retained records, rather than the accuracy of an estimator trained on those records.

Scale-calibrated camera-motion transfer, as considered by Meng, 2026, makes a distinction between camera trajectory information and its visual effect especially relevant. The present study provides a simpler measurement audit. It does not calculate the paper's parallax number, reproduce its transfer system, or test generated video. Instead, it establishes which elementary path summaries survive a concrete reduction in available pose anchors.

Two propositions organize the comparison. Nested thinning should never increase a path sum defined by a metric, because replacing several consecutive steps with their direct connection obeys the triangle inequality. However, the proportion removed need not be equal for position and orientation. Frequent small orientation reversals could disappear while a largely forward-moving positional path remains almost unchanged. This second proposition is empirical in the captured sequences; its direction is not guaranteed for every trajectory.

2. Related geometric and resource questions

The KITTI benchmark Geiger et al., 2012 supplies the reference trajectories analyzed here. Multi-view geometry Hartley and Zisserman, 2004 provides the distinction between a camera's position, orientation, and induced image projection. Rotation is not a fourth Cartesian coordinate that can simply be appended to position. Circular-data methods Fisher, 1993 warn against treating wrapped angular quantities as ordinary linear measurements, while quaternion interpolation Shoemake, 1985 and the documented Slerp construction SciPy developers, 2026 describe shortest-path interpolation between orientation anchors.

Pose estimation systems such as ORB-SLAM2 Mur-Artal and Tardos, 2016 and structure-from-motion pipelines such as COLMAP Schonberger and Frahm, 2016 address the recovery of geometry from images. Bundle-adjustment theory Triggs et al., 2000 and sparse implementations Lourakis and Argyros, 2009 explain why a recovered trajectory may involve coordinated estimates across views. These methods are relevant to the interpretation of a reference-pose file, but no estimator from this literature is rerun. EPnP Lepetit et al., 2009, for example, solves a camera-pose problem from correspondences; our calculation starts after poses have already been supplied.

Depth creates another separation between trajectory summaries and rendered appearance. Structured-light stereo data Scharstein and Szeliski, 2003 and the stereo evaluation taxonomy Scharstein and Szeliski, 2002 show why correspondence and depth measurement need explicit treatment. Monocular depth work represented by MiDaS Ranftl et al., 2019 and Monodepth2 Godard et al., 2018 likewise concerns scene structure that a trajectory-only audit does not observe. A tiny loss of translational path length can therefore coexist with consequential visual changes near a camera or around occlusion boundaries.

Resource reduction has a broader methodological history. Piecewise aggregate representations Keogh et al., 2001 and geometric simplification procedures Ramer, 1972 motivate checking what a compact description retains. The Douglas-Peucker approach Douglas and Peucker, 1973 is a data-adaptive simplification procedure, whereas the present stride policy ignores the trajectory's realized curvature. Spline methods de Boor, 1978 and shape-preserving interpolation Fritsch and Carlson, 1980 would introduce different reconstruction assumptions. We deliberately avoid filling the missing trajectory because interpolation quality is a separate question from the path length directly represented by observed anchors.

3. Reference arrays and the retained-anchor policies

The local source archive contains eleven reference-pose text files and 23,201 pose rows. Each row is reshaped into a three-by-four transformation matrix. Its final column supplies a three-dimensional translation vector; the leading three columns supply the rounded rotation matrix. All files are included. There is no exclusion based on the size of a subsequent thinning loss, the route's visual appearance, or whether its result agrees with the motivating proposition.

We treat the archive's adjacent pose rows as the dense captured reference. This is a finite measurement baseline rather than continuous ground truth. Motion between two adjacent source observations remains unavailable. Pose estimation error, sensor error, and rounding can also affect the accumulated variation. Consequently, an undercount means a difference relative to the supplied array, not a measured amount of physical movement omitted from an independently tracked vehicle.

Rounded orientation matrices are first mapped to proper rotations using SciPy's Rotation construction Virtanen et al., 2020. Relative rotations are then formed consistently from these normalized orientations. Their shortest rotation magnitudes are expressed in radians. This operation prevents slight departures from matrix orthogonality in a text file from being mistaken for a deliberately different definition of rotational

distance. The implementation uses NumPy arrays Harris et al., 2020 and preserves the unmodified source archive alongside the processed sequence summaries.

For each sequence and stride, retained indices start at zero and advance by the specified stride. The last source index is added if it is not already present. There are 11,606 anchors in the stride-two policy and 2,330 in the stride-ten policy. Every stride-ten anchor belongs to the stride-two set. Endpoints therefore match, and the more compact policy cannot benefit from a different starting phase or from dropping an inconvenient final partial interval.

Translation length is the sum of Euclidean distances between successive retained positions. Rotational variation is the sum of shortest relative-rotation magnitudes between successive retained orientations. Each sum is divided by its corresponding adjacent-pose reference sum, and one minus that ratio gives the undercount fraction. A value of 0.10 means that the retained path sum is ten percent smaller than the captured reference sum. It does not mean that each individual pose is ten percent inaccurate.

4. Paired estimands and the research model

The unit of comparison is the sequence. Within each sequence, we subtract stride-two undercount from stride-ten undercount separately for translation and rotation. Two intercept-only ordinary least squares models summarize the equal-sequence mean differences. Giving each sequence one row prevents long recordings from silently dominating the primary result. A second, descriptive aggregation weights translation by reference distance and rotation by reference rotational variation; these weighted quantities answer different accounting questions.

Figure 1 connects the fixed pose source to the two nested retention policies and their paired path summaries. The arrows describe the calculation, rather than a causal influence of resource expenditure on estimated camera behavior. No feature selection, learned calibration, or held-out outcome chooses the two stride values. The outcomes are computed from the same reference rows, so changes in source composition cannot explain a within-sequence contrast.

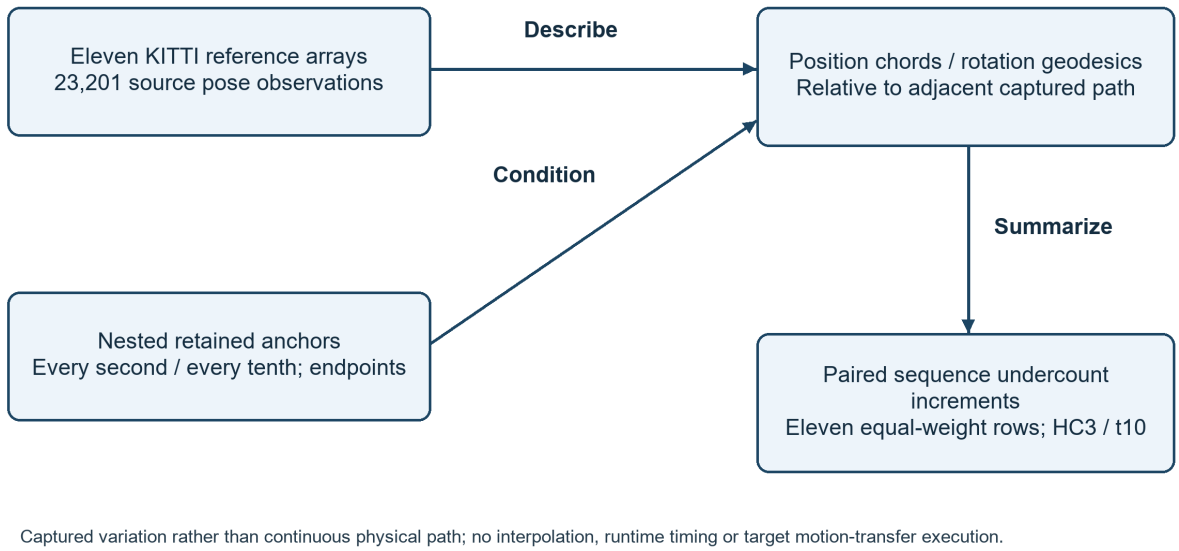


Figure 1. *Adjacent reference poses define the denominator. The two retained-anchor policies supply numerator paths for separate translation and rotation undercount fractions. Sequence-level differences enter two descriptive mean models.*

HC3 standard errors follow the heteroskedasticity-robust approach associated with White, 1980. Confidence intervals use ten residual degrees of freedom. They describe a working uncertainty calculation

across eleven selected benchmark sequences, rather than uncertainty in the source sensor measurements. Clustered inference would require a defensible larger grouping structure, as emphasized by Cameron and Miller, 2015; eleven sequence labels do not establish eleven independent road environments. Influence diagnostics are also relevant because a short sequence can have substantial weight in an equal-sequence analysis Cook, 1977.

Floating-point calculations deserve explicit treatment because repeated small rotations can accumulate differently from a single endpoint comparison. Numerical analysis Higham, 2002 and floating-point arithmetic guidance Goldberg, 1991 motivate retaining full calculation precision until reporting. We do not reinterpret the displayed decimal precision as equivalent source accuracy. Sampling theory Shannon, 1949 provides a general reason to be cautious about unobserved high-frequency variation, but no band-limited motion assumption or reconstruction guarantee is imposed on these recordings.

5. Results: almost preserved distance, substantially reduced turning

Summing all adjacent translations gives 22,179.67 meters in the captured files. The analogous accumulated orientation variation is 252.21 radians. These totals add movement across different recordings; they are descriptive archive totals, not a single journey. Stride two retains about half the pose records, while stride ten retains roughly one tenth, with the exact counts determined by endpoint inclusion.

The equal-sequence translation undercount is 0.00556 percent at stride two and 0.07556 percent at stride ten. Thus, the extra loss is 0.07000 percentage points. Rotation behaves quite differently: its equal-sequence undercount rises from 2.76246 percent to 15.49277 percent. The corresponding increase is 12.73031 percentage points. These percentages use separate translation and rotation denominators and should not be added into a composite movement score.

Table 1. Paired increases in undercount when the stride changes from two to ten.

Model	Term	Estimate (SE)	95% CI	N
M1	Mean outcome/contrast	0.000700 (0.000117)	[0.000439, 0.000961]	11
M1: fit	R-squared = 0.0000	-	-	11
M2	Mean outcome/contrast	0.127303 (0.037524)	[0.043695, 0.210911]	11
M2: fit	R-squared = 0.0000	-	-	11

Model labels: M1 = Stride-ten minus stride-two translation undercount fraction; M2 = Stride-ten minus stride-two rotation undercount fraction.

Table 1 reports fractions, so multiplying a coefficient by one hundred converts it to percentage points. Both models contain only an intercept, use eleven sequence observations, and report HC3 standard errors with t-based intervals using ten degrees of freedom. Their R-squared values are effectively zero because no explanatory variation beyond the mean is fitted. The translation interval is [0.04392, 0.09608] percentage points; the rotation interval is [4.36954, 21.09109] points.

Sequence 04 illustrates the role of the denominator. Its captured translation path is approximately 393.65 meters, but its accumulated rotational variation is only about 0.5772 radians. Rotation undercount increases from 9.446 percent to 57.223 percent under the coarser policy, producing a 47.777-point increment. This observation remains in the primary analysis. A large fraction of a small rotational denominator is not automatically a large absolute turning discrepancy, and it should not be described as an unusually dangerous route.

Weighting by captured movement reduces the rotation contrast's relative prominence without eliminating the difference between outcomes. Distance-weighted translation undercount rises from 0.00658 percent to 0.08884 percent. Rotation-weighted undercount rises from 2.10275 percent to 11.64294 percent. Equal-sequence weighting emphasizes a typical included recording; movement weighting emphasizes the contribution to the archive's total recorded variation. Neither is universally the correct benchmark denominator.

6. Implications for compact camera-motion reporting

The results answer the research question within the recorded arrays: the same anchor reduction removes much more of the accumulated orientation variation than of the translational path length. This does not imply that orientation is inherently harder to compress. It means that these recordings, these stride phases, and these two metric path sums retain different proportions of their adjacent-pose variation. A different camera movement or a data-adaptive policy could reverse the relative sensitivity.

A practical implication is to report both positional and angular retention diagnostics when a motion representation is reduced. A statement that sparse anchors preserve traveled distance can be correct while leaving turning behavior poorly characterized. It also leaves image projection, occlusion, scene scale, and subject movement unresolved. Those limitations explain why a trajectory-only resource audit should accompany, rather than replace, an application-specific visual evaluation.

The study's reproducibility record includes the source checksum, retained-index rule, processed anchor summaries, and sequence-level outcomes. FAIR principles Wilkinson et al., 2016 support making these relationships discoverable, while dataset documentation Gebru et al., 2021 motivates distinguishing the captured reference from the physical process it approximates. Future comparisons should evaluate additional stride phases, independently acquired trajectories, and explicit downstream tasks before recommending a general pose budget.

7. Interpreting a budget without confusing its units

The retained-pose count is the resource actually changed in this experiment. A stride-ten file has roughly one fifth as many anchors as its stride-two counterpart, but that ratio is not a measured fivefold decrease in processing time. Some downstream costs depend on image resolution, feature extraction, correspondence search, or optimization iterations rather than pose count alone. No memory allocation or wall-clock benchmark accompanies this analysis, so a cost model translating record savings into hardware benefits remains untested.

Nor do the outcomes describe per-frame reconstruction error. Suppose an omitted interval contains a turn followed by a compensating turn. The endpoint orientations can coincide even though their intermediate variation is positive. A shortest endpoint connection then records none of that variation. This is a conceptual property of the measurement, not an additional simulated trajectory included in the results. It clarifies why endpoint agreement and accumulated movement retention are different requirements, even before a visual application supplies its own loss function.

Coordinate conventions also affect what should be compared. A common rigid change of world frame leaves Euclidean positional distances unchanged, and a consistent change of orientation basis preserves relative rotation angles. Uniformly rescaling positions changes translation lengths but cancels from each translation undercount ratio. These mathematical properties do not establish parallax invariance or recover an unknown physical scale. They simply identify transformations that do not change the particular normalized path quantities used here.

The archive is heterogeneous in duration and recorded movement. This makes a single pooled pose-level standard error inappropriate for the question as formulated. The eleven sequences are the analysis units because each supplies a complete paired resource comparison. However, that choice does not eliminate dependencies created by the benchmark's acquisition setting. A broader study would need an explicit sampling frame of routes, camera rigs, and driving conditions before interpreting sequence variability as uncertainty for a general camera-motion population.

Finally, a useful retention policy needs a declared tolerance. We do not choose a maximum acceptable translation loss after seeing that translation changes little, nor dismiss rotation because one short sequence has a large fractional result. The present output is a diagnostic pair of outcomes. An application designer can use it to formulate a subsequent tolerance study, but the table alone cannot choose the storage policy that best balances rendering fidelity, navigation requirements, and computational cost.

8. Conclusion

A fivefold reduction relative to the stride-two anchor budget produces a small additional translation-length undercount but a much larger accumulated-rotation undercount in the eleven included reference sequences. Compact path reports should specify their retained indices, endpoints, movement metric, and aggregation weights. Preservation of one path summary supplies limited evidence about the others, and none of the present calculations establishes the quality of generated camera motion.

References

1. Keogh, Eamonn; Chakrabarti, Kaushik; Pazzani, Michael; Mehrotra, Sharad. (2001). Dimensionality Reduction for Fast Similarity Search in Large Time Series Databases. Knowledge and Information Systems. https://doi.org/10.1007/pl00011669.
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. Ramer, Urs. (1972). An iterative procedure for the polygonal approximation of plane curves. Computer Graphics and Image Processing. https://doi.org/10.1016/s0146-664x(72)80017-0.
4. Shoemake, Ken. (1985). Animating rotation with quaternion curves. Proceedings of the 12th annual conference on Computer graphics and interactive techniques - SIGGRAPH '85. https://doi.org/10.1145/325334.325242.
5. 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.
6. 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.
7. 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.
8. 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.
9. DOUGLAS, DAVID H; PEUCKER, THOMAS K. (1973). ALGORITHMS FOR THE REDUCTION OF THE NUMBER OF POINTS REQUIRED TO REPRESENT A DIGITIZED LINE OR ITS CARICATURE. Cartographica. https://doi.org/10.3138/fm57-6770-u75u-7727.
10. Cook, R. Dennis. (1977). Detection of Influential Observation in Linear Regression. Technometrics. https://doi.org/10.1080/00401706.1977.10489493.
11. Goldberg, David. (1991). What every computer scientist should know about floating-point arithmetic. ACM Computing Surveys. https://doi.org/10.1145/103162.103163.
12. Scharstein, Daniel; Szeliski, Richard. (2003). High-accuracy stereo depth maps using structured light. 2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.. https://doi.org/10.1109/cvpr.2003.1211354.

13. Gebru, Timnit, et al. "Datasheets for datasets" *Communications of the ACM*, 2021. https://doi.org/10.1145/3458723.

14. Meng, Zijie. (2026). ParaScale: Scale-Calibrated Camera-Motion Transfer via a Gauge-Invariant Parallax Number. https://arxiv.org/abs/2606.19805.

15. 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.

16. Fisher, N. I.. (1993). Statistical Analysis of Circular Data. Cambridge University Press. https://doi.org/10.1017/cbo9780511564345.

17. 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.

18. Godard, Clément; Mac Aodha, Oisin; Firman, Michael; Brostow, Gabriel. (2018). Digging Into Self-Supervised Monocular Depth Estimation. https://arxiv.org/abs/1806.01260.

19. Fritsch, F. N.; Carlson, R. E.. (1980). Monotone Piecewise Cubic Interpolation. SIAM Journal on Numerical Analysis. https://doi.org/10.1137/0717021.

20. 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.

21. Scharstein, Daniel; Szeliski, Richard. (2002). A Taxonomy and Evaluation of Dense Two-Frame Stereo Correspondence Algorithms. International Journal of Computer Vision. https://doi.org/10.1023/a:1014573219977.

22. Lepetit, Vincent; Moreno-Noguer, Francesc; Fua, Pascal. (2009). EPnP: An Accurate O(n) Solution to the PnP Problem. International Journal of Computer Vision. https://doi.org/10.1007/s11263-008-0152-6.

23. SciPy developers. (2026). scipy.spatial.transform.Slerp [API documentation; snapshot accessed 7 November 2026]. https://docs.scipy.org/doc/scipy/reference/generated/scipy.spatial.transform.Slerp.html.

24. Shannon, C.E.. (1949). Communication in the Presence of Noise. Proceedings of the IRE. https://doi.org/10.1109/jrproc.1949.232969.

25. de Boor, Carl. (1978). A Practical Guide to Splines. Applied Mathematical Sciences. https://doi.org/10.1007/978-1-4612-6333-3.

26. Lourakis, Manolis I. A.; Argyros, Antonis A.. (2009). SBA: A software package for generic sparse bundle adjustment. ACM Transactions on Mathematical Software. https://doi.org/10.1145/1486525.1486527.

27. 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.

28. Ranftl, René; Lasinger, Katrin; Hafner, David; Schindler, Konrad; Koltun, Vladlen. (2019). Towards Robust Monocular Depth Estimation: Mixing Datasets for Zero-shot Cross-dataset Transfer. https://arxiv.org/abs/1907.01341.

29. Hartley, Richard; Zisserman, Andrew. (2004). Multiple View Geometry in Computer Vision. Cambridge University Press. https://doi.org/10.1017/cbo9780511811685.

30. 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.