Temporal Systems and Forecasting
Volume 1 · Number 1 · 2026

Which Point Traveled? Rigid Reference Choices in Camera-Trajectory Distance

Sean Bailey[1] · Austin Walker[2] · Gavin Bell[3]

[1] Austria · [2] Belgium · [3] Czechia

Original publication: November 4, 2026 | Article 600

Abstract

Path length belongs to a specified point on a moving body. Using 23,201 KITTI reference poses, we compared the camera center with hypothetical points two meters forward and two meters laterally in the camera frame. Equal-sequence mean path changes were 0.2241% and 0.0184%, respectively. The lateral point had a shorter path in eight of eleven sequences, and its pooled distance change had the opposite sign from its equal-sequence mean change. A known rigid offset recovered the original camera centers numerically. The exercise shows why reporting point, offset direction and aggregation are needed before interpreting a trajectory-distance difference as an estimation error.

A geometric reporting audit of public reference poses. Hypothetical attached points are introduced deliberately; they are not measured vehicle centers, alternative pose estimates or errors in the source camera trajectory.

Which Point Traveled? Rigid Reference Choices in Camera-Trajectory Distance

> A geometric reporting audit of public reference poses. Hypothetical attached points are introduced deliberately; they are not measured vehicle centers, alternative pose estimates or errors in the source camera trajectory.

Abstract

Path length belongs to a specified point on a moving body. Using 23,201 KITTI reference poses, we compared the camera center with hypothetical points two meters forward and two meters laterally in the camera frame. Equal-sequence mean path changes were 0.2241% and 0.0184%, respectively. The lateral point had a shorter path in eight of eleven sequences, and its pooled distance change had the opposite sign from its equal-sequence mean change. A known rigid offset recovered the original camera centers numerically. The exercise shows why reporting point, offset direction and aggregation are needed before interpreting a trajectory-distance difference as an estimation error.

Distance reporting begins with a point, not only a pose file

Two rigidly attached points share a body's orientation but can trace different paths when that orientation changes. A camera center, a vehicle-center convention and a sensor mounted farther from a rotation axis can therefore accumulate different distances during the same rigid motion. A path-length comparison becomes ambiguous if the reported position columns refer to different points without stating their fixed relationship.

The first research question is how much two illustrative camera-attached offsets change adjacent-pose path length in a public driving resource. The second is whether offsets of identical length but different directions imply similar changes. We use supplied reference poses and perform explicit rigid-point transformations, separating this reporting issue from errors in estimating the poses themselves.

Meng, 2026 develops scale-calibrated camera-motion transfer through a gauge-invariant parallax number. The current study does not calculate that number or implement a transfer model. Its connection is the interpretation of geometric quantities: a meaningful comparison needs to distinguish changes of coordinate description from changes of the physical point being described. Choosing another attached point is the latter, even though the rigid transformation is fully known.

That distinction is central to the experiment. Adding a common translation to the world origin leaves the path length of a given point unchanged. Moving the reported point two meters within the camera's attached frame instead adds an orientation-dependent displacement in world coordinates. The resulting path difference is not evidence against coordinate invariance or a discovery that the source camera center was recorded incorrectly.

Reference poses and the geometric literature

KITTI provides the eleven captured reference-pose sequences used in this audit (Geiger et al., 2012). The recorded transformation direction is camera to world. The maintainer loader's pose convention is retained in the local source evidence (Clement, 2026). Its capture-year reference documents the inspected loader rather than assigning a new acquisition year to the driving sequences.

Rigid-state estimation provides the framework for composing a pose with a point attached to its local frame (Barfoot, 2017). Transformation tutorials also distinguish pose inversion, point composition and

rotation parameterization (Blanco-Claraco, 2021). We apply that elementary composition directly. No calibration parameters are inferred from the reference trajectories, and the two-meter offsets are illustrative inputs rather than measurements of this vehicle's hardware layout.

Multi-view geometry explains how camera coordinates and world coordinates relate (Hartley and Zisserman, 2004). Bundle adjustment, EPnP, COLMAP and ORB-SLAM2 address associated estimation problems (Triggs et al., 2000; Lepetit et al., 2009; Schonberger and Frahm, 2016; Mur-Artal and Tardos, 2016). The present calculation evaluates none of their recovered trajectories. It conditions on a single supplied pose array for each sequence and changes only the point whose world-space path is summarized.

Other driving resources make the need for coordinate documentation broadly relevant. nuScenes, Waymo and Argoverse 2 describe different sensing and evaluation contexts (Caesar et al., 2019; Sun et al., 2019; Wilson et al., 2023). Planning and diverse-driving benchmarks add further operational settings (Caesar et al., 2021; Yu et al., 2018). Their data are not included here, and their coordinate choices are not accused of being inconsistent. They illustrate why a result should not assume that every dataset's reported origin has the same physical interpretation.

Simulation can supply yet another frame convention, as in CARLA (Dosovitskiy et al., 2017). Likewise, a monocular-depth system faces estimation and scale issues absent from a calculation supplied with metric reference poses (Godard et al., 2018). Our perturbation neither simulates an unknown scale factor nor measures the performance of a simulator or depth network. It is a controlled comparison of trajectories for known related points.

Two equally long offsets, with different directions

We read all eleven pose text arrays, containing 23,201 observations in total. Each row supplies a three-by-three orientation block and a world translation. SciPy converts the orientation block to a proper rotation representation before composition (Virtanen et al., 2020). The same conversion is used for both offset conditions and the geometric bounds; we do not introduce a different rotation estimate for one condition.

The source translation is the camera-center position. One hypothetical attached point has camera-frame offset [0, 0, 2] meters, called forward here. The other has offset [2, 0, 0] meters, called lateral. Both offsets have Euclidean length two meters. The labels follow the operational camera-axis convention and do not establish the locations of a vehicle axle, inertial unit or actual additional sensor.

At each pose, the world position of an attached point is the camera translation plus the orientation matrix applied to its fixed local offset. Because the orientation can change, the additional world displacement is generally not constant across observations. By contrast, if orientation were constant over an interval, the attached point and the camera center would have the same displacement vector over that interval.

For each of the three points, we sum Euclidean distances between adjacent captured positions. These are sampled chord lengths, not a reconstruction of continuous motion between timestamps. No velocity integration, interpolation or odometer data is used. We call the source sum the camera-center path, retaining the point and sampling convention in the outcome's name.

Rotation representations must support consistent composition rather than be interpreted as independent coordinate columns. Work on rotational interpolation provides useful context for that requirement (Shoemake, 1985), but no interpolation method is executed in this study. Every reported position is evaluated at a captured pose. The finite sample's turn structure and its acquisition spacing jointly influence the numerical path sums.

Checks that separate point selection from implementation error

The first check subtracts the known rotated offset from every attached-point position. The recovered coordinates agree with the original camera translations to a maximum absolute difference of about 1.14e-13 meters. This is a reversal check on the implemented operation. It does not validate the absolute accuracy of the original reference poses against an independently measured trajectory.

The second check applies the reverse triangle inequality to each adjacent step. The absolute difference between an attached-point step length and the camera-center step length cannot exceed the norm of the change in the rotated offset. The largest stored excess over that bound is zero at the executed precision. Summing the stepwise bounds provides an upper bound on the absolute path-length difference for each sequence.

A looser orientation-only bound is also recorded. For a two-meter offset, the change in its rotated position is at most four times the sine of half the relative rotation angle. This upper bound does not depend on which direction the offset points, although the actual displacement usually does. Two equally long offsets therefore share a possible bound without being forced to have the same realized path change.

The direction of a path change depends on how the offset-induced displacement aligns with the camera's displacement. If the two step vectors point partly together, the attached point can travel farther; if they oppose each other, it can travel less. Their norms alone cannot determine the sign. A fixed mounting distance is consequently insufficient to predict the path difference without the rotation history and offset direction, even when there is no estimation noise.

No camera intrinsic parameters enter this calculation. The outcome is distance in world-coordinate meters, not image-plane displacement or a visibility score. It would be inappropriate to convert a relative path change directly into a parallax error, since projected motion also depends on scene points, depth and viewing geometry. Those data are absent from the current distance audit, and the target motion-transfer method is not evaluated indirectly through a proxy for them.

These are numerical checks of identities and inequalities for the declared transformation. Floating-point considerations justify retaining the residuals rather than demanding an exact string match between rounded path totals (Higham, 2002; Goldberg, 1991). A small recovery residual supports internal consistency. It does not convert an illustrative offset into a calibrated sensor mount or make the sampled chord path the continuous physical distance traveled.

Sequence-level descriptive models

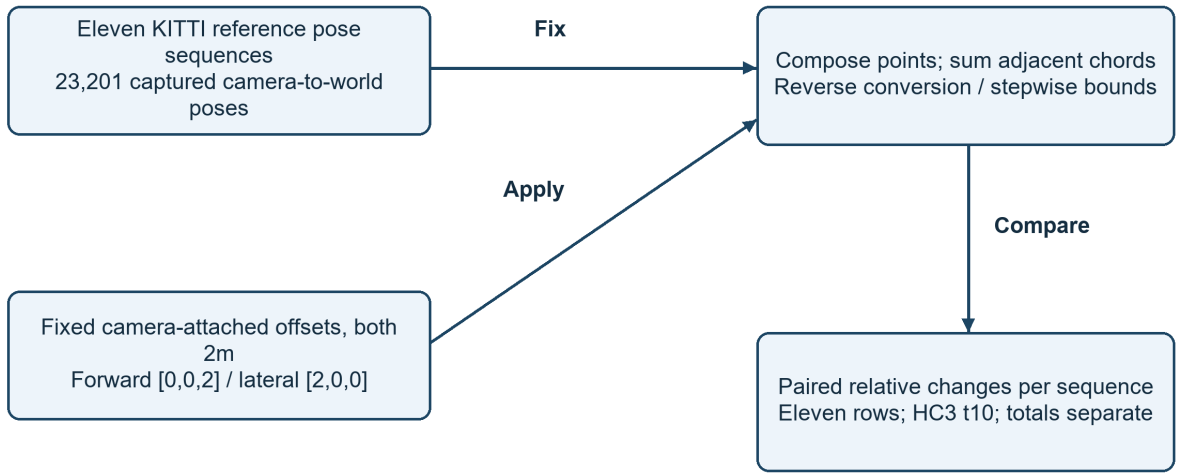


Figure 1. *Camera-to-world poses generate trajectories for the center and two fixed local offsets. Relative path changes are computed within sequence; reversal and stepwise bounds check the point transformation separately from the reported distance differences.*

For each sequence, we divide an attached point's path by that sequence's camera-center path and subtract one. Model 1 averages the forward relative changes; Model 2 averages the lateral changes. Model 3 averages their within-sequence difference, lateral minus forward. Every model has an intercept and eleven sequence rows, without controls or a fitted relationship between turning and distance.

HC3 standard errors and t10 intervals provide conventional working summaries of those sequence-level quantities (White, 1980). The routes are a selected benchmark collection, not an independently sampled population of all driving. Hardware and source-pose uncertainty can also be shared. We therefore avoid interpreting these intervals as established coverage for an arbitrary new vehicle, sensor mount or road network.

Table 1. *Relative path changes in proportion units, N=11 sequences per model, intercept only, HC3 covariance and t10 intervals. R-squared is zero to rounding. Multiplying a coefficient by one hundred expresses a percentage change; it does not make the coefficient a percentage-point accuracy difference.*

Model	Term	Estimate (SE)	95% CI	N
M1	Mean outcome/contrast	0.002241 (0.000552)	[0.001011, 0.003470]	11
M1: fit	R-squared = 0.0000	-	-	11
M2	Mean outcome/contrast	0.000184 (0.002521)	[-0.005433, 0.005800]	11
M2: fit	R-squared = 0.0000	-	-	11
M3	Mean outcome/contrast	-0.002057 (0.002202)	[-0.006964, 0.002850]	11
M3: fit	R-squared = 0.0000	-	-	11

Model labels: M1 = Two-meter forward reference relative path change; M2 = Two-meter lateral reference relative path change; M3 = Lateral minus forward relative path change.

The forward mean relative change is 0.00224050, standard error 0.00055164 and interval [0.00101138, 0.00346962]. Expressed as a percentage, the point estimate is approximately 0.2241%. All eleven forward

changes are positive in these captured arrays, ranging from about 0.000207% to 0.590672%. This is an observed pattern, not a theorem that forward offsets always lengthen a sampled camera path.

The lateral mean change is 0.00018352, standard error 0.00252082 and interval [-0.00543322, 0.00580026]. Its sequence changes range from approximately -0.824246% to 1.852848%. Eight sequences shorten and three lengthen. The nearly zero average therefore hides changes in both directions rather than showing that the lateral point follows the camera center's path almost exactly everywhere.

The lateral-minus-forward contrast is -0.00205698, with standard error 0.00220227 and interval [-0.00696396, 0.00284999]. Although its point estimate favors a smaller lateral change, the interval is wide and crosses zero. We do not infer equivalence of the directions or a generally superior reference choice. The comparison quantifies this specific pair of fixed offsets on the selected trajectories.

Pooled meters and average relative changes can disagree

Across all sequences, camera-center paths sum to 22,179.670 meters. The forward attached paths sum to 22,230.551 meters, and lateral paths sum to 22,159.460 meters. The pooled forward change is about 0.2294%; the pooled lateral change is about -0.0911%. The latter has the opposite sign from the equal-sequence average lateral change of positive 0.0184%.

There is no arithmetic contradiction. A ratio of total paths weights sequence-relative changes by source path length. The regression average gives each sequence equal importance. Longer sequences with negative lateral changes can outweigh shorter sequences with positive changes when pooling meters. Naming the averaging unit is necessary before comparing these summaries or treating one as a replication of the other (Stevens, 1946).

Sequence07 has a camera-center path of about 694.697 meters and a lateral relative increase of 1.8528%. Sequence02 has a much longer center path, about 5,067.233 meters, with a lateral decrease of roughly 0.2135%. The source table retains both. We do not remove short routes as inconvenient or use their relative change to claim a correspondingly large error in the source's camera-center positions.

Influence analysis highlights why unusual units should remain inspectable (Cook, 1977). Measurement-error frameworks further distinguish a changed measurement definition from uncertainty about a fixed construct (Carroll et al., 2006). Here the reported construct changes from the optical center's sampled path to another point's sampled path. A difference between those values can be entirely legitimate even when both trajectories are computed exactly from the same rigid poses.

Agreement methods help focus attention on differences with stated units rather than a potentially high correlation between long trajectories (Bland and Altman, 1986; Lin, 1989). A high path correlation would not answer whether the compared positions represent the same point. The transformation and the point definition must be checked first; only then does an error metric against a shared target have a clear interpretation.

What a trajectory comparison should retain

The experimental offsets are fully known and fixed. An operational conversion could instead face uncertain mounting geometry, time synchronization errors or a flexible attachment, none of which is modeled here. Calibration work such as stereo-system parameter estimation addresses a different evidential task (Scharstein and Szeliski, 2002). Our exact reversal should not be read as evidence that those practical calibration uncertainties are negligible.

The same caution applies when comparing estimated trajectories from different sensors. If both represent one agreed physical point, an appropriate rigid conversion can make their path targets comparable. If they intentionally represent different attached points, their distance difference should be described accordingly. Merely comparing translation-column norms cannot decide which of those situations applies.

A new attached coordinate frame can still describe the same rigid body's motion when the corresponding extrinsic relationship is retained. Composing that relationship back to the optical center recovers the original target point. Dropping the relationship and comparing the new frame origin's path with the optical center's path changes the measured target. Keeping this distinction explicit prevents a correct rigid-frame conversion from being mistaken for a deterioration of the trajectory estimator.

The saved sequence table includes the camera-center sum, both attached-point sums, relative changes and geometric bounds. Dataset documentation supports retaining their point conventions (Gebru et al., 2021), and provenance makes the particular input arrays and calculations reusable (Wilkinson et al., 2016). The empirical variation here is small in aggregate but uneven across routes and offset directions. That is sufficient reason to preserve the point definition before assigning an interpretation of geometric accuracy to a path-length difference.

References

1. Clement, Lee. (2026). pykitti: odometry.py. Source-code snapshot accessed7November2026; capture year, not original release year. https://github.com/utiasSTARS/pykitti/blob/master/pykitti/odometry.py.
2. Sun, Pei, et al. "Scalability in Perception for Autonomous Driving: Waymo Open Dataset" *arXiv preprint arXiv:1912.04838*, 2019. https://arxiv.org/abs/1912.04838.
3. 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.
4. Lin, Lawrence I-Kuei. (1989). A Concordance Correlation Coefficient to Evaluate Reproducibility. Biometrics. https://doi.org/10.2307/2532051.
5. Yu, Fisher, et al. "BDD100K: A Diverse Driving Dataset for Heterogeneous Multitask Learning" *arXiv preprint arXiv:1805.04687*, 2018. https://arxiv.org/abs/1805.04687.
6. Meng, Zijie. (2026). ParaScale: Scale-Calibrated Camera-Motion Transfer via a Gauge-Invariant Parallax Number. https://arxiv.org/abs/2606.19805.
7. Godard, Clément; Mac Aodha, Oisin; Firman, Michael; Brostow, Gabriel. (2018). Digging Into Self-Supervised Monocular Depth Estimation. https://arxiv.org/abs/1806.01260.
8. 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.

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

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

11. 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.

12. 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.

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

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

15. 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.

16. Wilson, Benjamin; Qi, William; Agarwal, Tanmay; Lambert, John; Singh, Jagjeet; Khandelwal, Siddhesh; Pan, Bowen; Kumar, Ratnesh; Hartnett, Andrew; Pontes, Jhony Kaesemodel; Ramanan, Deva; Carr, Peter; Hays, James. (2023). Argoverse 2: Next Generation Datasets for Self-Driving Perception and Forecasting. arXiv:2301.00493. https://arxiv.org/abs/2301.00493.

17. 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.

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

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

20. 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.

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

22. Dosovitskiy, Alexey; Ros, German; Codevilla, Felipe; Lopez, Antonio; Koltun, Vladlen. (2017). CARLA: An Open Urban Driving Simulator. arXiv:1711.03938. https://arxiv.org/abs/1711.03938.

23. Caesar, Holger, et al. "NuPlan: A closed-loop ML-based planning benchmark for autonomous vehicles." *arXiv preprint arXiv:2106.11810*, 2021. https://arxiv.org/abs/2106.11810.

24. 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.

25. 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.

26. Caesar, Holger, et al. "nuScenes: A multimodal dataset for autonomous driving" *arXiv preprint arXiv:1903.11027*, 2019. https://arxiv.org/abs/1903.11027.

27. 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.

28. 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.

29. 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.

30. 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.