Journal of Audit and Accountability
Volume 1 · Number 1 · 2026

Who Moved First? Order Sensitivity in Camera-Subject Motion Decomposition

Mark Griffin[1] · Alexander Johnston[2] · Jason Reed[3]

[1] Lithuania · [2] Iceland · [3] Luxembourg

Original publication: November 7, 2026 | Article 673

Abstract

The endpoints of an observed movement do not uniquely determine its intermediate camera and subject components. Using 6,697 moving-labeled reference-point pairs in Virtual KITTI 2, we compared camera-first and subject-first projection paths. The paths reconstructed the same total displacement to floating-point precision, but their camera-component vectors differed by an average norm of 0.000286 in normalized image coordinates. Relative interaction averaged 1.45% of the camera-first component-norm sum; 17 pairs changed which component was larger. These results quantify sensitivity to an attribution convention, without establishing a new physical motion source or evaluating a disentanglement network.

Geometric sensitivity analysis of public synthetic annotations. Two exact projection paths describe the same endpoints; their intermediate attributions are conventions, not identified causal contributions.

Who Moved First? Order Sensitivity in Camera-Subject Motion Decomposition

> Geometric sensitivity analysis of public synthetic annotations. Two exact projection paths describe the same endpoints; their intermediate attributions are conventions, not identified causal contributions.

Abstract

The endpoints of an observed movement do not uniquely determine its intermediate camera and subject components. Using 6,697 moving-labeled reference-point pairs in Virtual KITTI 2, we compared camera-first and subject-first projection paths. The paths reconstructed the same total displacement to floating-point precision, but their camera-component vectors differed by an average norm of 0.000286 in normalized image coordinates. Relative interaction averaged 1.45% of the camera-first component-norm sum; 17 pairs changed which component was larger. These results quantify sensitivity to an attribution convention, without establishing a new physical motion source or evaluating a disentanglement network.

1. An exact decomposition can still contain a choice

When a camera and an observed vehicle both move, an image reference point can change because of both updates. One way to describe the change is to move the camera first while holding the world point fixed, then move the point. Another way is to update the point first and the camera second. Both paths end at the same camera-point configuration. They need not assign the same intermediate displacement to each update.

The research question is how much that update order changes the components assigned to camera and subject motion in a captured geometric resource. This is distinct from estimating whether an object moved or predicting optical flow from images. We possess the annotated camera transforms and world reference points, and compare two explicit calculations on those same annotations.

Meng, 2026 proposes disentangling camera and subject motion through geometry-semantics orthogonal attention. Its attention mechanism is not implemented here. The relevant connection is the need to define what evidence supports a motion-attribution claim. Our study examines a convention in an elementary geometric decomposition and does not assume that the target method uses this decomposition or inherits its numerical sensitivity.

The main proposition is algebraic: two consistent update paths should reproduce the same total projected change. The empirical question concerns the size of their disagreement about the components. A small average interaction could make the convention inconsequential for some uses while still changing a decision close to the boundary between camera and subject dominance. We report both magnitude summaries and the observed number of such reversals.

2. Geometry, tracking, and attribution have different targets

Multi-view geometry supplies the coordinate transformations underlying the calculation (Hartley and Zisserman, 2004). Camera estimation methods such as bundle adjustment, EPnP, COLMAP, and ORB-SLAM2 address how poses or structure are inferred (Triggs et al., 2000; Lepetit et al., 2009; Schonberger and Frahm, 2016; Mur-Artal and Tardos, 2016). We do not compare those estimators. The current study conditions on the released transformations rather than claiming to recover a more accurate camera trajectory.

Likewise, learned depth and geometric representation methods address uncertainty and information unavailable to a metadata-only calculation. MiDaS and self-supervised monocular depth work illustrate that broader estimation problem (Ranftl et al., 2019; Godard et al., 2018). In this experiment, the world points and camera matrices are supplied. The resulting sensitivity describes the decomposition of known resource coordinates, not the uncertainty of recovering geometry from generated or real video.

Tracking introduces its own objects and metrics. AB3DMOT and Deep SORT associate detections over time (Weng et al., 2020; Wojke et al., 2017), while HOTA evaluates tracking behavior through a defined combination of detection and association considerations (Luiten et al., 2020). Our reference point follows an existing track identifier. We do not run a tracker, score identity switches, or infer detection performance from its coordinate displacement.

Attribution also differs from causal identification. Explanatory frameworks such as SHAP formalize how contributions can be assigned to prediction inputs (Lundberg and Lee, 2017). We compute two physical-coordinate update paths and their symmetric arithmetic average, without fitting a prediction model or claiming a general explanation theorem. Causal-structure reasoning further cautions against reading an algebraic component as an intervention effect identified from observational data (Greenland et al., 1999).

These distinctions matter because the same word, motion, can refer to world displacement, camera-relative coordinates, normalized image displacement, or apparent flow. We measure the two-dimensional change in a projected tracked reference point. It is not a dense motion field, the center of a two-dimensional bounding box, an angular velocity, or a proportion of motion physically caused by one source.

3. The selected moving-reference-point cohort

Virtual KITTI and Virtual KITTI 2 provide the synthetic scene context and annotations (Gaidon et al., 2016; Yohann et al., 2020). We use only the clone condition of five scenes and the left camera. Weather variants, alternative yaw conditions, and the second camera are excluded rather than treated as additional independent scene replications. The underlying acquisition is synthetic, with public text precision and annotation conventions retained.

The pose and bounding-box tables are joined by frame and track identity. A pair must contain consecutive frames of the same track. At both endpoints, camera-space depth must lie between five and eighty meters, occupancy ratio must be at least 0.25, and truncation ratio must be at most 0.1. These screens define a visible, bounded-depth cohort and do not establish that every reference point is free from partial occlusion.

We then require the source moving label to be true. That annotation is the operational selection variable; it is not a new manual determination of motion from images. The resulting cohort contains 6,697 pairs: 316 in Scene01, 520 in Scene02, 612 in Scene06, 1,166 in Scene18, and 4,083 in Scene20. The point estimate is therefore reference-pair weighted, with substantial contribution from the last scene.

Camera transforms map the world-space bottom-face reference point into camera coordinates. The transformed endpoint positions are checked against the released camera-space columns within a 0.002-meter consistency limit. Both crossed intermediate positions must also be in front of their respective cameras before division by depth. No selected moving-labeled pair is excluded by that latter front-camera check in the executed analysis.

The cohort's constraints should not be generalized to every driving observation. KITTI, nuScenes, Waymo, and Argoverse 2 contain different real-world sensing and annotation settings (Geiger et al., 2012; Caesar et al., 2019; Sun et al., 2019; Wilson et al., 2023). Their existence clarifies the external evidence that would be needed for a road-system claim. Their observations are not merged with the current five synthetic scenes or added to the effective sample size.

4. Four projections and two paths between them

For each pair, we project four combinations: current camera with current world point, next camera with current point, current camera with next point, and next camera with next point. Projection divides the first two camera coordinates by depth. The first and last combinations are the observed resource endpoints; the middle two are explicitly constructed intermediate configurations.

The camera-first path assigns its first vector to the difference between the next-camera/current-point projection and the current endpoint. Its second vector goes from that intermediate to the next endpoint. The subject-first path instead moves through the current-camera/next-point projection. In each case the two vectors sum exactly to the same endpoint difference, independent of how their individual norms compare.

The order-interaction vector is the subject-first camera component minus the camera-first camera component. It is also the camera-first subject component minus the subject-first subject component. Switching paths therefore reallocates this vector between the named components while preserving the total. We verify both identities numerically rather than inferring agreement from similar rounded summary numbers.

Model 1 summarizes the norm of that interaction vector. Model 2 divides its norm, pair by pair, by the sum of the camera-first camera and subject component norms, then averages the ratios. This denominator avoids interpreting the interaction as a fraction of a possibly near-canceling total displacement. It remains an order-specific normalization convention, not a measure of the percentage of motion caused by either source.

A symmetric descriptive assignment averages the camera vectors from the two paths and separately averages the subject vectors. Their vector sum still reconstructs the total. Their norms do not generally add to the norm of the total, and the averaging does not force orthogonality. We retain this optional convention in the saved outputs without calling it a uniquely correct physical attribution.

5. Regression and numerical checks

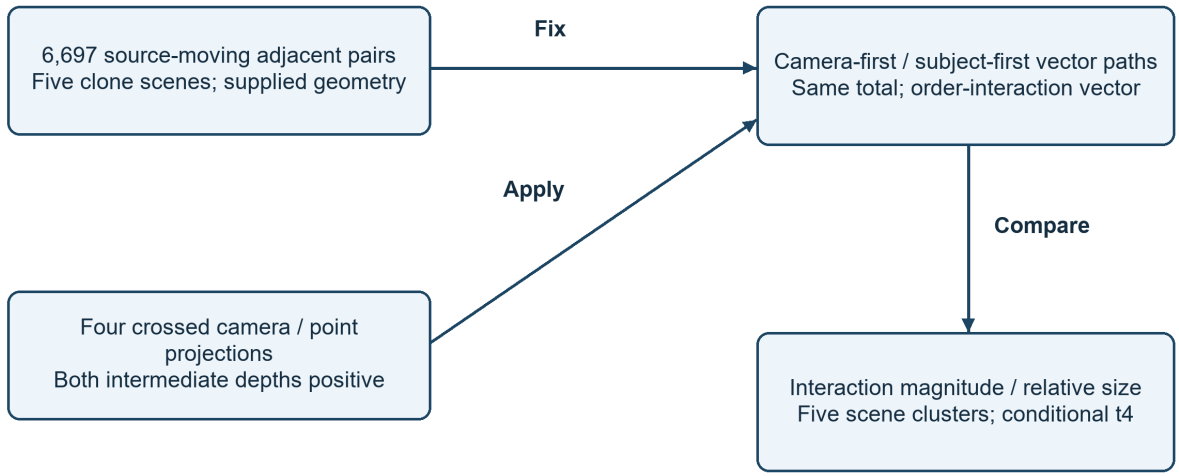

Figure 1. *Four camera-point combinations yield two exact additive paths. Order-interaction magnitude and component-dominance reversals are distinct outcomes; neither is an identified causal share.*

Both reported regressions contain an intercept only and use the 6,697 selected pairs. Standard errors cluster by the five contributing scenes, apply a finite correction, and use t4 confidence intervals. There are no scene fixed effects or other controls. Because the observations are not collapsed to equal scene means, the coefficient remains a pair-weighted mean despite the scene-level covariance calculation.

Cluster covariance reflects specified shared-error structure (Cameron and Miller, 2015), while robust regression variance addresses different error-variance concerns (White, 1980). Only five clusters are present here, and common synthetic actors or acquisition design can link them further. These intervals are fragile descriptive summaries, not validated coverage for a broad scene population. More object frames within Scene20 do not replace independent scenes for that purpose.

The local SciPy environment supports reproducible numerical handling (Virtanen et al., 2020). Standard numerical-analysis concerns apply when subtracting close projected vectors (Higham, 2002; Goldberg, 1991). The largest discrepancy in either total reconstruction identity is about 6.94e-18, and the largest discrepancy in the interaction identity is about 5.20e-18. These checks support internal arithmetic consistency; they do not establish exact physical truth of rounded synthetic annotations.

Table 1. *Intercept-only order-interaction summaries, N=6,697 pairs, five scene clusters, finite-corrected standard errors and t4 intervals. M1 uses normalized image-plane units; M2 is a dimensionless mean of per-pair ratios. Intercept-only R-squared is zero up to floating-point rounding.*

Model	Term	Estimate (SE)	95% CI	N
M1	Mean outcome/contrast	0.000286 (0.000061)	[0.000115, 0.000456]	6697
M1: fit	R-squared = 0.0000	-	-	6697
M2	Mean outcome/contrast	0.014547 (0.002259)	[0.008274, 0.020819]	6697
M2: fit	R-squared = 0.0000	-	-	6697

Model labels: M1 = Mean order-interaction vector norm; M2 = Order-interaction norm relative to camera-first component-norm sum.

6. Small average interaction, occasional decision reversals

The average interaction norm is 0.000285530 in normalized image coordinates. Model 1 gives standard error 0.000061402 and interval [0.000115051, 0.000456009]. These are not pixel units or degrees. Converting them would require an additional camera-intrinsic or angular definition that is outside the reported outcome.

The mean normalized interaction ratio is 0.0145467, or about 1.45%, with standard error 0.0022592 and interval [0.0082741, 0.0208193]. The median ratio is 0.0105447, the 90th percentile is 0.0330506, the 99th is 0.0598369, and the maximum is 0.1574291. The distribution thus contains larger local sensitivities than its mean alone suggests.

The ratio is calculated separately within every pair. Dividing the mean interaction norm by the mean component-norm sum would weight pairs differently and produce another quantity. We do not substitute that ratio of means for Model 2 or describe either quantity as a causal allocation. Stating the denominator is necessary for interpreting what the reported percentage measures (Stevens, 1946).

Mean camera-component norm is 0.0071791 under camera-first attribution and 0.0070429 under subject-first attribution. The corresponding subject norms are 0.0061703 and 0.0058934. Both sets reconstruct the same total as vectors. Their different scalar averages illustrate why taking component magnitudes does not preserve the additive identity and should not be used to define unexplained percentages that must sum to one hundred.

Seventeen pairs change whether the camera component or the subject component has the larger norm, a frequency of 0.00253845. This is a descriptive count, without a separate sparse-event confidence interval. The small rate suggests that the dominance label is usually stable in this selected resource, while the existence of reversals prevents treating it as mathematically invariant to update order.

Measurement-error reasoning distinguishes the observed variable from its intended construct (Carroll et al., 2006). Here the exact equality of endpoint displacement coexists with disagreement about the intermediate attribution. A validation that checked only the sum would therefore miss the convention sensitivity measured by the interaction.

There is also a distinction between vector disagreement and a change in the larger scalar component. Two nonzero components may partly cancel in the image plane. A relatively small reallocation vector can then have a larger effect on a thresholded description than its absolute size suggests, particularly when the original norms are nearly equal. Conversely, the camera vector can change direction without reversing the dominance label. The seventeen reversals summarize one deliberately narrow decision rule and cannot stand in for a complete distribution of directional errors.

The released coordinate check has a largest endpoint difference of approximately 0.000163 meters. That is below the declared tolerance, but it remains a check on two representations in the same annotation resource. It is not an independently measured error bound for vehicle location. Keeping the transformation check, additive identities, and order-sensitive outcomes separate prevents a numerical consistency result from becoming an unsupported validation of the scene's physical realism.

7. Interpretation and the next evidence requirement

The calculations do not identify a new motion mechanism. Camera changes include rotation and translation, while perspective division depends on depth; moving through different intermediate camera-point combinations changes the assigned increments. The interaction is a property of this finite geometric decomposition. It is not evidence that the subject actually moved at a different time in the synthetic scene or that the resource's annotations are inconsistent.

For a learned disentanglement system, an appropriate follow-up would state its intended attribution target and evaluate against that definition across independent scenes. Validation should match that transfer question (Roberts et al., 2017), and influential scene composition should remain visible (Cook, 1977). The current five-scene result cannot rank attention architectures, validate a dense optical-flow decomposition, or establish semantic separation in generated video.

The saved record includes the cohort screens, scene and track identifiers, both paths' component norms, interaction ratios, dominance flags, and arithmetic residual checks. Documentation makes their meanings and limits inspectable (Gebru et al., 2021), while reusable provenance supports tracing the computation to the public annotations (Wilkinson et al., 2016). Under this declared convention audit, total motion is stable by construction; attribution components are not. Reporting that distinction makes a motion-decomposition claim more precise before any stronger physical or model-performance interpretation is attempted.

References

1. Gebru, Timnit, et al. "Datasheets for datasets" *Communications of the ACM*, 2021. https://doi.org/10.1145/3458723.
2. Cook, R. Dennis. (1977). Detection of Influential Observation in Linear Regression. Technometrics. https://doi.org/10.1080/00401706.1977.10489493.
3. 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.
4. Wojke, Nicolai; Bewley, Alex; Paulus, Dietrich. (2017). Simple Online and Realtime Tracking with a Deep Association Metric. arXiv:1703.07402. https://arxiv.org/abs/1703.07402.
5. 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.
6. Luiten, Jonathon, et al. "HOTA: A Higher Order Metric for Evaluating Multi-Object Tracking" *arXiv preprint arXiv:2009.07736*, 2020. https://arxiv.org/abs/2009.07736.
7. Hartley, Richard; Zisserman, Andrew. (2004). Multiple View Geometry in Computer Vision. Cambridge University Press. https://doi.org/10.1017/cbo9780511811685.
8. 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.
9. Lundberg, Scott M., and Su-In Lee. “A Unified Approach to Interpreting Model Predictions.” *Advances in Neural Information Processing Systems*, vol. 30, 2017. https://proceedings.neurips.cc/paper_files/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html.
10. Goldberg, David. (1991). What every computer scientist should know about floating-point arithmetic. ACM Computing Surveys. https://doi.org/10.1145/103162.103163.
11. Yohann, Cabon,; Naila, Murray,; Martin, Humenberger,. (2020). Virtual KITTI 2. https://arxiv.org/abs/2001.10773.
12. 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.
13. Stevens, S. S.. (1946). On the Theory of Scales of Measurement. Science. https://doi.org/10.1126/science.103.2684.677.

14. 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.

15. Meng, Zijie. (2026). OrthoMotion:Disentangling Camera and Subject Motion via Geometry Semantics Orthogonal Attention. arXiv:2606.22835. https://arxiv.org/abs/2606.22835.

16. 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.

17. 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.

18. Greenland, Sander, Judea Pearl, and James M. Robins. “Causal Diagrams for Epidemiologic Research.” *Epidemiology*, vol. 10, no. 1, 1999, pp. 37-48. https://doi.org/10.1097/00001648-199901000-00008.

19. Gaidon, Adrien; Wang, Qiao; Cabon, Yohann; Vig, Eleonora. (2016). Virtual Worlds as Proxy for Multi-Object Tracking Analysis. CVPR 2016, 4340-4349. https://openaccess.thecvf.com/content_cvpr_2016/html/Gaidon_Virtual_Worlds_as_CVPR_2016_paper.html.

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. 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.

22. 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.

23. 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.

24. 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.

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

26. 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.

27. 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.

28. Weng, Xinshuo, et al. "AB3DMOT: A Baseline for 3D Multi-Object Tracking and New Evaluation Metrics" *arXiv preprint arXiv:2008.08063*, 2020. https://arxiv.org/abs/2008.08063.

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

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