**Journal of Measurement and Evidence**
Volume 1 · Number 1 · 2026

# Five Pixels Is Not the Same Gate after Rescaling: A Coordinate-Unit Audit of Apparent Motion

Caleb Holmes[1] · Charles Lawson[2] · James Harrison[3]

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

Original publication: October 2, 2026 | Article 680

## Abstract

Halving image-coordinate units leaves scene geometry unchanged but alters a fixed pixel threshold. On 11,330 Virtual KITTI 2 point pairs, a five-pixel motion gate accepts 4,629 native-coordinate records and only 2,731 after the coordinate change. Scaling the threshold to 2.5 pixels restores every original decision. The unadjusted pass fraction falls by 0.1675, with a five-scene-clustered interval from −0.2064 to −0.1286. This is an exact coordinate audit using source calibration, not a new rendering experiment or a physical change in camera motion.

# Five Pixels Is Not the Same Gate after Rescaling: A Coordinate-Unit Audit of Apparent Motion


## Abstract

Halving image-coordinate units leaves scene geometry unchanged but alters a fixed pixel threshold. On 11,330 Virtual KITTI 2 point pairs, a five-pixel motion gate accepts 4,629 native-coordinate records and only 2,731 after the coordinate change. Scaling the threshold to 2.5 pixels restores every original decision. The unadjusted pass fraction falls by 0.1675, with a five-scene-clustered interval from -0.2064 to -0.1286. This is an exact coordinate audit using source calibration, not a new rendering experiment or a physical change in camera motion.


## 1. A motion threshold contains an implicit calibration

A statement such as more than five pixels of motion appears precise. Its physical and geometric meaning nevertheless depends on the coordinate grid. The same projected points can be expressed in a finer or coarser pixel unit, changing their numerical displacement without changing the scene or the camera. A threshold retained without that unit conversion can then alter which observations are called substantial motion. The research question is how many decisions change under one explicit coordinate rescaling, and whether consistent threshold calibration restores the original decision set.

Scale calibration is relevant to the camera-motion transfer problem in Meng, 2026. We do not compute its parallax number or execute its transfer method. Instead, we inspect a simpler prerequisite for any scale-sensitive motion criterion: the coordinate units of the quantity being thresholded must agree with those of the threshold. The experiment intentionally preserves the underlying geometry. It is not evidence about how a method performs when a physical camera changes focal length, moves through a new scene, or produces a newly rendered image at a different resolution.

The released geometric annotations in Yohann et al., 2020 permit this distinction to be checked exactly. We use five clone scenes and the left camera, with actual intrinsics read from their source text files. The retained point pairs already have normalized image-coordinate displacements. Their conversion to a pixel displacement uses the source focal calibration, after verifying the conditions under which that scalar conversion is valid. This avoids inventing a camera constant or treating normalized image-plane distance as an exact angular quantity in radians.

## 2. Coordinate equivalence and measurement meaning

Projective geometry in Hartley and Zisserman, 2004 makes the distinction between a scene point, its normalized projection, and its calibrated pixel coordinates explicit. In the retained source files, horizontal and vertical focal values are equal, and the principal point is constant across frames within each scene. Consequently, the norm of a normalized two-dimensional displacement multiplied by the common focal value equals the corresponding pixel-displacement norm. A changing principal point or unequal focal values would require a fuller component-wise calculation; we verify rather than assume those simplifying conditions here.

The state and pose treatments in Barfoot, 2017 and Blanco-Claraco, 2021 similarly require an explicit coordinate convention before numerical comparisons acquire meaning. Bundle-adjustment approaches in Triggs et al., 2000 and Lourakis and Argyros, 2009, and the pose problem addressed by Lepetit et al., 2009, depend on consistent camera observations and calibration. We fit none of those estimators. Their

relevance is that a threshold on an observation cannot be detached from the representation in which that observation is expressed. Correct pose values alone do not repair an incorrectly carried-over pixel unit.

The scale discussion in Stevens, 1946 helps identify what is being changed. We rescale an existing numerical representation by one half. Both image coordinates and focal lengths in pixel units receive that change. The same normalized projection is recovered when the new pixel displacement is divided by the new focal value. This is a unit conversion, not a perturbation of the underlying reference geometry. Its effect on an unadjusted gate is therefore a definition mismatch rather than a discovered degradation in physical camera behavior.

Measurement-error concerns in Carroll et al., 2006 remain separate. A source pose or calibration could contain error, but no such error is injected or estimated in this audit. The agreement perspectives in Bland and Altman, 1986 and Lin, 1989 also help distinguish matching decisions from similar average measurements. We check every corrected binary decision directly rather than relying on an aggregate correlation between native and rescaled displacement. The expected invariance is algebraic, while the number of records affected by failing to apply the conversion is empirical for this particular source cohort.

## 3. Source cohort and actual focal values

The inherited geometry table contains 11,330 adjacent tracked point pairs from Scene01, Scene02, Scene06, Scene18, and Scene20. Each uses the left camera and the tracked bottom-face reference point. Endpoint depth must be between five and eighty meters, occupancy at least 0.25, and truncation ratio at most 0.1. An intermediate camera-first projection must lie in front of the camera. These restrictions are held fixed. The rescaling operation neither adds low-visibility points nor removes records based on whether their threshold decision changes.

For each clone scene, we read the intrinsic table and select camera zero. The focal values in both axes equal 725.0087 pixels throughout those source records. Principal points are also constant within the respective scene files. The same focal value therefore maps the retained normalized displacement magnitudes to native pixel units across the five scenes. The mean native displacement is 7.641570 pixels. That average is descriptive of the selected adjacent reference points and does not represent a dense optical-flow mean over all visible image pixels.

The half-coordinate representation multiplies every displacement and the focal value by 0.5. We calculate a native gate at displacement at least five pixels, an unadjusted gate at five of the new units, and a calibrated gate at 2.5 of the new units. The first two thresholds therefore select different normalized displacements. The third is the exact converted version of the first. These values were declared before computing the new decision counts; five pixels is illustrative and is not claimed as a validated operational motion threshold.

Floating-point details are checked using the concerns in Higham, 2002 and Goldberg, 1991 as guidance for interpreting numerical identities. The largest discrepancy between the recovered normalized displacement and the source normalized value is approximately 2.78 times ten to the minus seventeen. Array computations use the local numerical environment represented by Harris et al., 2020. The corrected binary gate agrees exactly on all retained records. We do not attach a fabricated confidence interval to that exact within-record equality or describe it as proof of a new physical invariance discovered from sampled data.

## 4. A diagram and one nonconstant statistical contrast

*Figure 1 separates unchanged geometry from the numerical representation and gate definition. The unadjusted comparison measures the consequence of carrying the same number into a different unit. The corrected comparison verifies that consistent conversion preserves the decisions. Only the first produces a nonconstant row-level outcome suitable for the displayed inferential summary. Avoiding an unnecessary model for the all-zero corrected difference prevents a degenerate standard error from being mistaken for certainty about every future use of a motion threshold.*

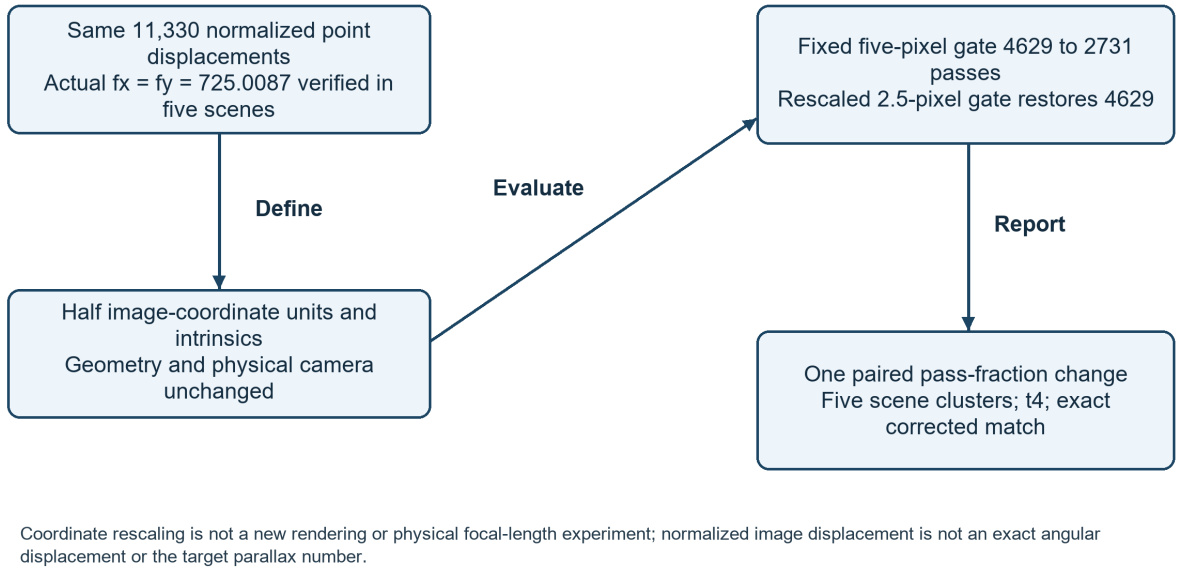


***Figure 1.*** *Identical normalized displacements are expressed in native and half-scale pixel units. A fixed numerical gate changes its geometric meaning; a proportionally rescaled gate recovers the same record-level decisions.*

The statistical outcome is the unadjusted half-coordinate pass indicator minus the native pass indicator. Its intercept-only OLS mean is image-pair weighted. Covariance clusters by the five scene identifiers, with finite-sample correction and a Student reference of four degrees of freedom. White, 1980 motivates heteroskedasticity-aware summaries, while Cameron and Miller, 2015 makes the small number of dependence groups central to interpretation. The table describes the observed cohort under a working scene-cluster model; it is not an independent-pixel or independent-frame calculation.

Repeated tracks, shared camera trajectories, unequal scene counts, and synthetic-world composition remain limitations. The structured-validation discussion in Roberts et al., 2017 cautions against treating additional variants of the same source scene as independent environments. We use only the clone setting and do not inflate the scene count with weather, viewpoint, or stereo copies. The result is consequently narrow but traceable: it quantifies a gate-definition consequence on one stored five-scene set while keeping the algebraic corrected agreement separate from sampling claims.

***Table 1.*** *Mean unadjusted-half minus native pass fraction, estimated by intercept-only OLS. N counts 11,330 point pairs; covariance uses five scene clusters with finite-sample correction and a t reference of four degrees of freedom. There are no controls. Corrected gate agreement is an exact descriptive check rather than a second inferential model.*

| Model | Term | Estimate (SE) | 95% CI | N |
|---|---|---|---|---|
| M1 | Mean outcome/contrast | -0.167520 (0.014004) | [-0.206400, -0.128639] | 11330 |
| M1: fit | R-squared = 0.0000 | - | - | 11330 |

Model labels: M1 = Unadjusted half-coordinate minus native pass fraction.

## 5. The same number rejects a different part of the cohort

The native five-pixel gate accepts 4,629 records. After coordinate units are halved, retaining the numeral five accepts only 2,731 records. There are 1,898 native passes lost and no new passes, because the unadjusted comparison is equivalent to doubling the native-coordinate threshold. The observed mean pass change is negative 0.167520, with standard error 0.014004 and a 95% interval from negative 0.206400 to negative 0.128639. The sign is mechanically constrained; the magnitude depends on how many source displacements occupy the affected interval.

Using a 2.5-pixel threshold in the half-coordinate representation accepts 4,629 records and matches every native gate result. This recovery is exact for the stored values and their declared conversion. It does not establish that an image-resizing algorithm preserves all motion evidence. No image is resized, filtered, interpolated, compressed, or rendered. Such operations could alter a detector's outputs in addition to changing its coordinate units, and those changes would require a separate pixel-level experiment rather than a reinterpretation of this identity.

The distinction also prevents a misleading robustness narrative. The unadjusted gate does not fail because the source camera became less observable or the moving subject changed its behavior. Its decision boundary moved in normalized coordinates because the unit conversion was omitted. Cook, 1977 motivates inspecting influential observations in a model, but deleting the records near this shifted boundary would not solve the definition mismatch. The reporting perspective in Wasserstein and Lazar, 2016 similarly favors explaining the constrained comparison and actual count change rather than celebrating the small interval as evidence of a sophisticated new effect.

## 6. Calibration context beyond this one audit

Geometry resources use different sensor and coordinate arrangements. Geiger et al., 2012, Sturm et al., 2012, Burri et al., 2016, and Schubert et al., 2018 illustrate benchmark settings in which trajectory, sensor, and calibration information must remain tied to their definitions. We do not run the present gate on those sources or claim that their units can be substituted directly into the same formula. Their relevance is to documentation: a record called motion or displacement needs a coordinate frame, time interval, and observation model before cross-resource numerical thresholds can be compared responsibly.

Reconstruction and localization systems make the same dependence operational. Schonberger and Frahm, 2016 and Mur-Artal and Tardos, 2016 use camera geometry in different estimation pipelines, but neither provides evidence that a threshold survives an arbitrary change of calibration or image processing. The stereo taxonomies in Scharstein and Szeliski, 2002 and the depth-ground-truth setting in Scharstein and Szeliski, 2003 likewise distinguish pixel-domain correspondence from physical scene quantities. A pixel displacement should not become a distance in meters or an exact angular displacement merely because a calibrated camera is involved.

This experiment also leaves physical scale unanswered. A change in image-coordinate units is not a change in the baseline, the object distance, or the camera's translation magnitude. The source point and pose records stay fixed. The target work's scale-calibrated motion-transfer question therefore remains broader than the present audit. The justified implication is only that a downstream gate should carry its unit conversion with it; no conclusion about how physical parallax transfers across cameras is derived from halving a stored numerical representation.

Documentation follows Gebru et al., 2021 and Wilkinson et al., 2016 by retaining scene-specific focal checks, source cohort keys, both displacement representations, the three gate flags, and the exact normalized-value comparison. These records distinguish the coordinate audit from a resolution

benchmark. A future rendering study could hold scene motion fixed while changing image grids and then rerun a detector, but it would add estimation and resampling mechanisms that are absent here. Those mechanisms should be measured explicitly rather than smuggled into the interpretation of a clean unit conversion.

The affected records have a particularly simple interpretation. A native displacement from five to below ten pixels passes the original gate but fails the unadjusted gate after halving its numerical representation. Displacements below five fail both, and displacements at least ten pass both. Thus the observed 1,898 lost passes count the cohort's mass inside one known interval. They do not measure an unobserved deterioration in geometry or a new failure mode of a motion estimator. This interval interpretation makes the mechanical direction and empirical magnitude of the result separable.

Uniform rescaling is also a restricted case. If horizontal and vertical coordinates were scaled differently, a scalar displacement norm and threshold would not generally transform by one common factor. The calibration and decision rule would need to specify the corresponding geometry more carefully. That case was not executed here. The actual source's equal focal axes and our common one-half factor make the current conversion exact, so the result should remain tied to those conditions rather than advertised as a solution to every resolution, aspect-ratio, or camera-calibration change.

The threshold is not estimated from outcome labels. Its corrected value follows directly from the declared unit factor, not from searching for a number that recreates the original pass count. Exact record-level agreement provides a stronger implementation check than a matching total alone, because two different sets can have equal size. We preserve both flags for that reason. A future application that only records aggregate counts could miss membership substitutions even when its reported coverage appears unchanged.

## 7. Conclusion

Keeping a five-pixel threshold after halving image-coordinate units removes 1,898 passes from the retained geometric cohort. Converting the threshold to 2.5 of the new units restores every original decision. The result combines an exact representation identity with an empirical count of the consequences of ignoring it. Motion thresholds should therefore remain bound to their calibration and coordinate units, with numerical rescaling clearly distinguished from physical camera changes or a new image-processing experiment.

## References


1. 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.
2. 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.
3. Hartley, Richard; Zisserman, Andrew. (2004). Multiple View Geometry in Computer Vision. Cambridge University Press. https://doi.org/10.1017/cbo9780511811685.
4. 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.
5. Goldberg, David. (1991). What every computer scientist should know about floating-point arithmetic. ACM Computing Surveys. https://doi.org/10.1145/103162.103163.
6. 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.

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

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

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. 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.

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

13. 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.

14. 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.

15. Wasserstein, Ronald L.; Lazar, Nicole A.. (2016). The ASA Statement on p -Values: Context, Process, and Purpose. The American Statistician. https://doi.org/10.1080/00031305.2016.1154108.

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

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

18. 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.

19. 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.

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. 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.

22. 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.

23. 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.

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. 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.

26. Yohann, Cabon,; Naila, Murray,; Martin, Humenberger,. (2020). Virtual KITTI 2. https://arxiv.org/abs/2001.10773.

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

28. 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.

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

30. 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.