Equal Pair Counts and Unequal Identity Coverage across Verification Folds

Evan Barnes (IS), Austin Hall (GE), Joseph Crawford (BT)

Abstract

Every prescribed LFW fold contains 600 pairs, but its identity count ranges from 397 to 451 and its distinct image count from 737 to 813. The corresponding coefficients of variation are approximately four and 3.6 percent. This audit shows modest but real coverage differences behind a perfectly balanced pair denominator. It preserves the protocol's lack of cross-fold identity overlap and makes no recognition-performance claim. For subject-preserving video, equal numbers of comparisons should be accompanied by subject and reference coverage rather than treated as complete evidence of equivalent assessment populations.

1. What is balanced by a balanced fold

A fold can contain a fixed number of comparisons while representing different numbers of people and photographs. The balance is genuine at the row level and incomplete at other levels. Neither observation requires treating the protocol as defective. The scientific question is which units an estimator weights and which population a reported result describes. A report that states only the pair count leaves those distinctions invisible.

ARGUS motivates evaluation across varied identity observations and generated views (Meng et al., 2026b). A video benchmark could similarly fix the number of scored clips per partition while varying how many subjects or reference images those clips represent. This article does not measure ARGUS or a new video dataset. It audits the prescribed LFW fold counts to show how a familiar static protocol can support different statements about comparison balance and identity coverage. The evidence is limited to the metadata actually parsed.

Replication has several levels: a second measurement, a second trajectory, and a second independently acquired setting are different additions to evidence. Statistical calculations should retain those distinctions. Heteroskedasticity-robust regression inference concerns unequal conditional variances (White, 1980); cluster-robust inference concerns dependencies shared within groups (Cameron and Miller, 2015). The group definition should be justified through the data collection and evaluation process, and not chosen because it produces a preferred significance result. Unequal group sizes also influence what a pooled average represents. Reporting both the group count and observation count permits a reader to see whether the calculation describes common cases, heavily recorded cases, or an explicitly balanced group population. No variance estimator alone resolves that choice of estimand.

The validity of an identity evaluation depends partly on how its scientific objects are named and linked. LFW's identity, image, pair, and fold units are explicit parts of its protocol (Huang et al., 2008). The scikit-learn implementation documents which file corresponds to each supported task (Pedregosa et al., 2011). Dataset provenance should similarly make the intended use and selection process recoverable (Gebru et al., 2021). In video generation, an evaluator needs additional links among reference observations, subjects, clips, and frames. Repeated references can provide useful evidence without becoming independent subjects. A pair-level metadata audit can reveal those relationships, but it cannot evaluate the visual content of a generated clip. The distinction is productive: data organization can be checked before an expensive model study begins, and the missing measurements can be specified rather than silently assumed to exist.

2. The observed range across ten fixed folds

The official ten-fold pair file defines 600 comparisons per fold, divided equally into matched and mismatched classes. Counting complete identity and image keys separately gives identity counts from 397 to 451. Distinct image counts range from 737 to 813. The sample coefficient of variation is 0.0398 for identity counts and 0.0360 for image counts. These are descriptive variations among the ten prescribed groups, not estimates from repeatedly resampled datasets.

Identity and image counts across folds have a Pearson association of 0.6682. With only ten fixed observations, this value should be read as a compact description rather than a broad statistical law about verification design. It does not show that one count causes the other or establish a relationship with classifier difficulty. No prediction outcomes are available. The audit also confirms that images and identities do not cross the prescribed fold boundaries, so the coverage variation should not be misidentified as leakage.

The scientific interpretation should specify what would change under the proposed intervention. Holding an observed covariate fixed statistically is not equivalent to manipulating it in an environment. A causal graph helps separate confounders, mediators, and selection mechanisms (Greenland et al., 1999). It can also reveal why controlling for every recorded variable is not automatically defensible. Meanwhile, an evaluation with multiple conditions and endpoints creates a family of opportunities for apparently striking findings. Multiplicity methods require a declared set of comparisons and an appropriate inferential purpose (Benjamini and Hochberg, 1995). Descriptive tables may be useful without testing every cell. When the data do not identify a mechanism, the appropriate contribution can be a sharper question and a transparent comparison rather than an unsupported causal answer.

The recognition literature provides both scoring methods and warnings about incomplete evaluation populations. FaceNet's learned embedding (Schroff et al., 2015), ArcFace's angular-margin formulation (Deng et al., 2018), and CosFace's margin-based discrimination (Wang et al., 2018) are alternative measurement foundations rather than interchangeable proof of personal identity. VGGFace2 introduces substantial pose and age variation (Cao et al., 2017), MegaFace studies recognition at large scale (Kemelmacher-Shlizerman et al., 2015), and FairFace emphasizes representation in annotated face datasets (Kärkkäinen and Joo, 2019). An identity-preserving generator should be assessed with the relevant dimensions recorded separately. A result from familiar frontal images cannot establish reliability under a missing viewpoint. Nor can a frequency-balanced set establish demographic balance without the required labels. The evaluation should retain failures of detection or measurement, identify what threshold was used, and explain whether the endpoint concerns verification, identification among distractors, or human recognizability over a moving sequence.

3. Averaging outcomes over these partitions

Because pair counts are identical, an ordinary mean of fold-level pair accuracies would give the same row weighting as pooling those pair-level correctness indicators. That equivalence would not make the result an identity-macro accuracy. A fold with more identities spreads its comparison exposure differently from a fold with fewer identities. A separate identity-level calculation would need to define how positive and negative roles are assigned and weighted, rather than obtain its meaning from the fold labels alone.

The range of counts is modest, and the evidence should not be exaggerated. It does not imply that the assessment is dominated by a tiny fold or that equal pair counts are a misleading design objective. It shows that a clear statement at one level should not automatically be promoted to another. A good report

can retain the balanced comparison structure while adding coverage information about the underlying people and observations.

For generated-subject assessment, the same practice would record subjects, reference views, clips, and scored frames per partition. Equal frame counts can hide different clip durations or numbers of subjects. Equal clip counts can hide different numbers of reference observations or repeated histories. A study should nominate its primary weighting and present alternative denominators when they change interpretation. This makes it possible to understand whether a score concerns the average frame, clip, subject, or intended user case.

If a later model performs differently across partitions, these metadata can support a follow-up question but not an immediate mechanism explanation. More identities might coincide with different viewpoints, acquisition conditions, or comparison partners. Those variables would need to be observed before attributing an accuracy difference to coverage alone. The present audit deliberately stops at the count structure, preserving a useful design observation without inventing image difficulty or model errors that were not measured.

A subject-generation pipeline can succeed at using a reference without succeeding at every identity requirement. DreamBooth's model adaptation (Ruiz et al., 2022) and IP-Adapter's image prompting (Ye et al., 2023) offer examples of reference use with different costs and information access. CLIP (Radford et al., 2021) and DINO (Caron et al., 2021) provide further visual representations whose similarity scores may emphasize other aspects of content. The evaluation should therefore define its target before choosing a convenient feature. If the intended target is a recognizable moving person, that target needs appropriate observations across the relevant views and expressions. A high generic image similarity can instead reflect scene layout or texture. Matching reference budgets and documenting the evaluator's representation permit those explanations to be compared. The objective is to connect a claimed conditioning advantage to the particular identity evidence the resulting output actually preserves.

The phrase generalization needs an object: another observation of the same source, another source within the same collection, or a different acquisition process. Each requires a different assessment boundary. Structured validation makes explicit how correlated observations are assigned across training and evaluation (Roberts et al., 2017). Selection bias can still arise if the final boundary is repeatedly consulted while choosing a method (Varma and Simon, 2006). Publishing the split identifiers and the information permitted during each stage allows the claimed transfer to be checked. When a study proposes a new architecture but performs only a data audit, it should state that difference directly. Useful evidence about data coverage does not require pretending that a prediction system has already been trained and externally validated.

A useful reliability claim contains both a performance statement and a scope statement. Selective prediction shows why a system that answers only some cases must report how many it accepts (Geifman and El-Yaniv, 2017). Confidence estimation under a changed data distribution requires independent checking (Ovadia et al., 2019). These requirements apply to generators and to the automated judges used to assess them. If a judge cannot evaluate a profile, an occluded object, or an unusual interaction, that inability should remain visible rather than removing the case from the generator's score silently. Recording the entire decision path permits an analyst to identify whether an apparent improvement came from better outputs, narrower eligibility, or a more restrictive measurement procedure.

4. Connecting an identity claim to its observations

Subject preservation is one dimension of a larger video-generation task. VBench explicitly motivates separating evaluation dimensions (Huang et al., 2023), and the visual literature offers structural (Wang et al., 2004), perceptual (Zhang et al., 2018), and temporal-distribution measurements (Unterthiner et al., 2018). A study should choose among them according to the claim and available reference information. When identity, motion, and visual quality move in different directions, the disagreement should be analyzed rather than automatically resolved by a weighted total. The weighting itself expresses a preference that may differ across applications. Preserving per-frame and per-clip records allows alternative interpretations to be evaluated later. It also makes coverage visible when some poses or occlusions prevent valid scoring. A generator can demonstrate a useful, limited preservation capability without having established universal reliability across all subjects, durations, and visual conditions.

Replacing an expensive judgment with an automated score changes the measurement process. A positive correlation is evidence of co-movement, but it does not establish that the two procedures can be used interchangeably (Bland and Altman, 1986). A calibrated forecast has a different requirement: its stated probabilities should match relevant observed frequencies (Guo et al., 2017). Neither property follows merely from a favorable overall ranking. A useful comparison preserves paired measurements, identifies the conditions in which disagreement matters, and keeps scale transformations explicit. If the intended judgment concerns recognizability, playability, or physical compatibility, the proxy should be assessed against that judgment directly. Agreement on easy cases cannot supply evidence about the difficult conditions absent from the comparison.

The unit of provenance should be fine enough to reconstruct an important disagreement. Reusable identifiers can connect a reported score to its input evidence and the relevant model and evaluator versions (Wilkinson et al., 2016). Pipeline dependencies should be treated as part of the experiment because apparently minor changes can alter output eligibility or metric behavior (Sculley et al., 2015). A report that preserves only aggregate numbers prevents this reconstruction. Keeping per-case results and documented transformations enables alternative weighting and error analysis without repeating expensive generation. Such a record also helps identify whether two papers compare the same scientific objects or merely use similar names for different releases. Versioned evidence makes progress easier to assess across successive evaluations.

5. Conclusion

LFW's prescribed folds are exactly balanced in pair rows and somewhat variable in identity and image coverage. These properties coexist with clean separation across folds. The audit provides a precise description of what the balance does and does not mean. Subject-preserving video studies can adopt the same layered accounting so that equal comparison counts remain useful without becoming an unsupported claim of equivalent subjects, references, or difficulty. A later performance result can then be interpreted against the population it actually weights.

References

Meng, Zijie, et al. "Argus: Stacked multi-view identity mosaic injection for subject-preserving video generation." *arXiv preprint arXiv:2606.11670* (2026).

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.

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.

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.

Benjamini, Yoav and Yosef Hochberg. “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing.” *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 57, no. 1, 1995, pp. 289-300. https://doi.org/10.1111/j.2517-6161.1995.tb02031.x.

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.

Varma, Sudhir and Richard Simon. “Bias in error estimation when using cross-validation for model selection.” *BMC Bioinformatics*, vol. 7, no. 1, 2006. https://doi.org/10.1186/1471-2105-7-91.

Wilkinson, Mark D., et al. "The FAIR Guiding Principles for Scientific Data Management and Stewardship." *Scientific Data*, vol. 3, 2016, article 160018. https://doi.org/10.1038/sdata.2016.18.

Sculley, D., et al. "Hidden Technical Debt in Machine Learning Systems." *Advances in Neural Information Processing Systems*, vol. 28, 2015. https://papers.nips.cc/paper/2015/hash/86df7dcfd896fcaf2674f757a2463eba-Abstract.html

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.

Guo, Chuan, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. “On Calibration of Modern Neural Networks.” *Proceedings of the 34th International Conference on Machine Learning*, PMLR, vol. 70, 2017, pp. 1321-1330. https://proceedings.mlr.press/v70/guo17a.html.

Geifman, Yonatan, and Ran El-Yaniv. “Selective Classification for Deep Neural Networks.” *Advances in Neural Information Processing Systems*, vol. 30, 2017. https://arxiv.org/abs/1705.08500.

Ovadia, Yaniv, et al. “Can You Trust Your Model’s Uncertainty? Evaluating Predictive Uncertainty Under Dataset Shift.” *Advances in Neural Information Processing Systems*, vol. 32, 2019. https://arxiv.org/abs/1906.02530.

Huang, Gary B., Marwan Mattar, Tamara Berg, and Erik Learned-Miller. "Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments." *ECCV Workshop on Faces in Real-Life Images*, 2008. https://cs.umass.edu/~elm/papers/Huang_eccv2008-lfw.pdf.

Schroff, Florian, et al.. "FaceNet: A Unified Embedding for Face Recognition and Clustering" *arXiv preprint arXiv:1503.03832*, 2015. https://arxiv.org/abs/1503.03832.

Deng, Jiankang, et al.. "ArcFace: Additive Angular Margin Loss for Deep Face Recognition" *arXiv preprint arXiv:1801.07698*, 2018. https://arxiv.org/abs/1801.07698.

Wang, Hao, et al.. "CosFace: Large Margin Cosine Loss for Deep Face Recognition" *arXiv preprint arXiv:1801.09414*, 2018. https://arxiv.org/abs/1801.09414.

Cao, Qiong, et al.. "VGGFace2: A dataset for recognising faces across pose and age" *arXiv preprint arXiv:1710.08092*, 2017. https://arxiv.org/abs/1710.08092.

Kemelmacher-Shlizerman, Ira, et al.. "The MegaFace Benchmark: 1 Million Faces for Recognition at Scale" *arXiv preprint arXiv:1512.00596*, 2015. https://arxiv.org/abs/1512.00596.

Kärkkäinen, Kimmo and Jungseock Joo. "FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age" *arXiv preprint arXiv:1908.04913*, 2019. https://arxiv.org/abs/1908.04913.

Ruiz, Nataniel, et al.. "DreamBooth: Fine Tuning Text-to-Image Diffusion Models for Subject-Driven Generation" *arXiv preprint arXiv:2208.12242*, 2022. https://arxiv.org/abs/2208.12242.

Ye, Hu, et al.. "IP-Adapter: Text Compatible Image Prompt Adapter for Text-to-Image Diffusion Models" *arXiv preprint arXiv:2308.06721*, 2023. https://arxiv.org/abs/2308.06721.

Huang, Ziqi, et al.. "VBench: Comprehensive Benchmark Suite for Video Generative Models" *arXiv preprint arXiv:2311.17982*, 2023. https://arxiv.org/abs/2311.17982.

Caron, Mathilde, et al.. "Emerging Properties in Self-Supervised Vision Transformers" *arXiv preprint arXiv:2104.14294*, 2021. https://arxiv.org/abs/2104.14294.

Radford, Alec, et al.. "Learning Transferable Visual Models From Natural Language Supervision" *arXiv preprint arXiv:2103.00020*, 2021. https://arxiv.org/abs/2103.00020.

Zhang, Richard, et al.. "The Unreasonable Effectiveness of Deep Features as a Perceptual Metric" *arXiv preprint arXiv:1801.03924*, 2018. https://arxiv.org/abs/1801.03924.

Unterthiner, Thomas, et al.. "Towards Accurate Generative Models of Video: A New Metric & Challenges" *arXiv preprint arXiv:1812.01717*, 2018. https://arxiv.org/abs/1812.01717.

Wang, Zhou, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. "Image Quality Assessment: From Error Visibility to Structural Similarity." *IEEE Transactions on Image Processing*, vol. 13, no. 4, 2004, pp. 600-612. https://doi.org/10.1109/TIP.2003.819861.

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

Pedregosa, Fabian, et al. "Scikit-learn: Machine Learning in Python." *Journal of Machine Learning Research*, vol. 12, 2011, pp. 2825-2830. https://jmlr.org/papers/v12/pedregosa11a.html.