Image Reuse Networks inside a Face-Verification Protocol

Jonathan Stevens (NP), James Hayes (BH), Jeffrey Robertson (JO)

Abstract

LFW's 6,000 verification pairs use 7,701 distinct photographs. Of those photographs, 3,217 appear in multiple pairs and account for 62.63 percent of all pair-side appearances. Most are reused only twice, but two appear six times. This descriptive audit distinguishes the number of images from their comparison exposure. It does not estimate recognition errors or invalidate the prescribed identity-disjoint folds. The findings motivate preserving reference and output links in subject-preserving video evaluation, where many scored frames or comparisons can share the same small set of source observations.

1. Comparisons reuse their ingredients

A pairwise evaluation can create many observations from a smaller collection of images. That is often intentional and can make a protocol efficient. It also means the number of pair rows is not a count of independent photographs. If an image is unusually easy, difficult, or poorly acquired, its properties can influence several comparisons. The appropriate interpretation depends on how those comparisons were constructed and what outcome is later measured.

ARGUS organizes identity references into a structured multi-view input for video generation (Meng et al., 2026b). Its use of multiple observations makes reference reuse a relevant design issue, without implying that repetition is inherently undesirable. This article does not test its mosaic or generated outputs. It examines exposure in the public LFW pair list as a concrete, reproducible example of the distinction between a reference object and the comparisons that use it. The resulting counts describe protocol structure rather than image content or identity fidelity.

A benchmark's observation count should be accompanied by a description of what an observation is. The LFW pair protocol separates identities from their images and from the comparisons used for verification (Huang et al., 2008). Its documented software access routes also distinguish the ten-fold assessment from development subsets (Pedregosa et al., 2011). Dataset documentation connects those files to their collection and intended purpose (Gebru et al., 2021). A generation benchmark needs an equally explicit account of how a referenced subject becomes a set of output clips and scored frames. Without it, repeated observations may receive unintended weight or cross a supposed assessment boundary. A metadata analysis can establish whether those links are consistent. It should still avoid claiming that consistency of filenames proves recognizability, fair representation, or robustness under viewpoint change. Those are further outcomes requiring appropriate measurements.

An identity score is produced by a measurement system with its own training and assumptions. FaceNet (Schroff et al., 2015), ArcFace (Deng et al., 2018), and CosFace (Wang et al., 2018) offer influential approaches to discriminative face representation. Their outputs should be interpreted in relation to the images and decisions on which they are assessed. VGGFace2 broadens variation within subject observations (Cao et al., 2017), MegaFace makes large distractor populations important (Kemelmacher-Shlizerman et al., 2015), and FairFace addresses imbalance across annotated demographic categories (Kärkkäinen and Joo, 2019). These are complementary evaluation concerns. A generated-video study should not compress them into an unexplained claim that a face benchmark is difficult. It should specify whether difficulty comes from viewpoint, acquisition, distractors, subject population, or the representation itself. Demographic conclusions require appropriate annotations and cannot be inferred merely from identity frequencies or from a generic face-similarity score. A metric's usefulness for one purpose remains an empirical question for another.

2. A degree count for every image key

Each image is identified by combining its identity label and its within-identity image index. Using the numeric index alone would merge unrelated photographs. The ten-fold file contains 6,000 pairs, yielding 12,000 pair-side appearances. Counting appearances for each complete image key gives 7,701 distinct images. No image data are downloaded, and the analysis does not construct embeddings or inspect whether a photograph is visually distinctive.

The exposure distribution is as follows: 4,484 images appear once, 2,352 twice, 680 three times, 155 four times, twenty-eight five times, and two six times. The counts sum to 7,701 distinct photographs; weighting each category by its exposure sums to 12,000 appearances. This arithmetic check helps keep the denominators explicit. The 3,217 reused images comprise about 41.8 percent of the distinct images but contribute 62.63 percent of pair-side appearances. Those percentages differ because one describes objects and the other describes exposure.

Precision and identification are separate properties of a result. A narrow interval can surround an association whose scientific target is unclear. Robust covariance calculations allow particular departures from simple error assumptions, including unequal variance (White, 1980) and dependence within declared clusters (Cameron and Miller, 2015). Their role is to improve uncertainty calculations under those assumptions. They do not establish a causal assignment, guarantee a correct measurement, or expand coverage to conditions missing from the data. The analyst should therefore explain both the averaging population and the source of replication. An evaluation repeated many times on one history may provide information about stochastic variability while supplying little evidence about transfer to another history. Those two sources of uncertainty should remain distinguishable in the final interpretation.

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.

Different reference-conditioning mechanisms can receive different amounts of information. DreamBooth adapts a generation model around a subject (Ruiz et al., 2022), while IP-Adapter offers a compatible image-conditioning route (Ye et al., 2023). A fair comparison should preserve how many views, which resolutions, and which adaptation opportunities are supplied. CLIP's language-supervised visual representation (Radford et al., 2021) and DINO's self-supervised features (Caron et al., 2021) can provide useful descriptions of outputs, but their training objectives do not establish a person-recognition standard. The endpoint should specify whether success means matching reference appearance, retaining a recognizable subject, or satisfying a requested scene. Those properties can disagree. A useful analysis would keep the disagreement visible and ask which source of evidence each method uses. The relevant advantage is a measured improvement under a declared information budget, not an unqualified increase in similarity to whatever reference happened to be available.

3. Reuse is not the same as cross-fold leakage

The prescribed LFW fold organization separates image and identity sets across folds in the analyzed file. The observed image reuse occurs within folds. It therefore raises a dependence and weighting question rather than demonstrating that the official assessment boundary leaks images. Repartitioning pair rows could create a different problem, but that is not a property inferred from reuse alone. A sound report should keep these diagnoses separate.

The counts also do not establish a numerical error dependence. A difficult image may produce correlated errors in several pairs, but the direction and magnitude depend on the evaluated system and its comparison partners. There are no classifier outputs in this audit. Calling the exposure count an effective accuracy sample size would therefore overstate what has been calculated. It is a descriptive map of shared ingredients that a later performance analysis can use to choose and justify its uncertainty model.

For subject-preserving video, analogous links can connect a reference image to several mosaics, a subject to many clips, or a clip to many scored frames. A benchmark that retains only one flattened score per comparison may conceal this structure. A better record would identify each source observation and every generated object that uses it. An investigator could then distinguish broad coverage of new subjects from repeated assessments of a familiar reference collection.

Reuse can also be scientifically useful. Comparing the same reference under several controlled output conditions helps isolate how the conditioning process responds to a change. The important requirement is to record what is shared and what varies. A repeated reference should not be counted as a new independent subject, while a carefully paired comparison should not be discarded merely because it uses the same image. The research question determines whether shared evidence is an advantage, a dependence to model, or a limitation in coverage.

An evaluation protocol should state what remains unfamiliar at assessment time. Leaving out an observation while retaining its close relatives can test interpolation rather than the broader transfer suggested by the study's wording. Validation methods for structured data emphasize matching the split to temporal, spatial, or hierarchical dependence (Roberts et al., 2017). Separately, parameter selection can bias an evaluation if the same observations guide both choices and reporting (Varma and Simon, 2006). These issues are not solved by calling a split a test set. The record should identify shared sources, the information accessible during development, and the stage at which each design choice was fixed. A result can be informative within a narrow boundary while providing no direct evidence about a more distant deployment population.

A score is useful only relative to the question it is supposed to answer. Correlation measures a relationship between varying quantities, whereas agreement concerns the discrepancies produced when one procedure substitutes for another (Bland and Altman, 1986). Confidence estimates add a further question about calibration (Guo et al., 2017). These distinctions matter when a learned evaluator becomes the principal source of evidence about another learned system. The evaluation should preserve examples with conflicting scores and obtain independent target judgments where feasible. Numerical ranges and normalization choices should also be documented. An apparently small difference may have a different practical meaning near a decision threshold than in an unimportant part of the scale. Reporting the intended consequence gives those comparisons an interpretable purpose.

4. Connecting an identity claim to its observations

Video evaluation should preserve the distinction among identity, temporal behavior, and general visual appearance. VBench offers a multidimensional approach to video assessment (Huang et al., 2023), while structural similarity (Wang et al., 2004), learned perceptual metrics (Zhang et al., 2018), and generative video-distribution evaluation (Unterthiner et al., 2018) provide different kinds of measurements. Their outputs can disagree for scientifically meaningful reasons. A temporal average may dilute a brief subject substitution, and a strong reference match may reward a nearly static output that fails to carry out the requested motion. The test should therefore retain where and when each observation was made. Human judgments or task-specific measurements may be needed to determine whether a disagreement is consequential. Combining metrics is reasonable only when the combination's interpretation is explicit. A total score should not conceal the fact that one component was unavailable or inappropriate for the very cases motivating the new generation method.

Difficult cases can disappear at several stages: data eligibility, generation, measurement, and final scoring. Each exclusion changes the population represented by the reported result. Coverage-aware evaluation makes those decisions visible (Geifman and El-Yaniv, 2017). A confidence score should also be checked when the input distribution changes, because uncertainty under shift is an empirical property rather than a guarantee of the model interface (Ovadia et al., 2019). An informative record preserves failures to produce an output and failures to obtain a valid measurement. Treating either as absent data without explanation can inflate apparent reliability. A proposed method should state its unsupported conditions and test the boundary between accepted and rejected cases as carefully as the quality of its successful outputs.

An evaluation remains comparable only if its inputs and transformations can be identified. The FAIR principles emphasize persistent identification and reusable scientific metadata (Wilkinson et al., 2016). Work on machine-learning technical debt highlights how dependencies and hidden couplings accumulate in a pipeline (Sculley et al., 2015). In this setting, a score should be linked to the exact references, generation settings, evaluator version, and aggregation rule that produced it. A replacement checkpoint or a revised eligibility filter may change the scientific question even when the output column keeps the same name. Preserving such changes in a release record permits later comparisons to be interpreted honestly. The objective is not maximum documentation volume, but a recoverable account of the decisions that affect the result.

5. Conclusion

Image exposure in the public LFW protocol is concentrated enough that reused photographs account for most pair-side appearances. This fact is compatible with the protocol's clean separation across folds. It identifies a structure to preserve, not a recognition failure to assume. Subject-generation evaluations can apply the same distinction by linking references, subjects, clips, and frames explicitly. The resulting record makes coverage and dependence assessable before a large comparison count is interpreted as a large amount of independent identity evidence.

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.