Comparing Subject-Conditioning Methods under Matched Evidence

Arthur Fisher (SL), Gabriel Wright (MR), Jeffrey Walters (BN)

Abstract

A method comparison is difficult to interpret when one generator receives more useful reference information than another. Equal image counts do not guarantee equal evidence. This article proposes matching source observations, access opportunities, and selection effort in subject-conditioned video evaluation. ARGUS is the motivating multi-view method. No systems are run here. The proposed comparison distinguishes a controlled conditioning contrast from a complete workflow comparison, making clear whether an observed advantage would concern the mechanism, the information supplied, or their practical combination.

1. What should be equal in a comparison

Reference-conditioned methods differ in the number of views they accept, the processing they apply, and the opportunities they provide for subject adaptation. A nominally equal reference budget can conceal large differences in usable evidence. One method may receive a clear profile while another receives a duplicate frontal crop. Another may incorporate extensive adaptation before the reported generation begins. A comparison needs an explicit information boundary before its performance difference can be attributed to a conditioning approach.

ARGUS uses organized identity references in video generation (Meng et al., 2026b). This article proposes an evaluation design for comparing such a method with other subject-conditioning workflows. It reports no new ranking and does not assume that all interfaces can receive identical inputs. Its central distinction is between a controlled contrast designed to hold selected evidence conditions fixed and an end-to-end comparison designed to measure what complete systems achieve under declared practical constraints.

Identity evidence has a sampling structure before any recognition score is calculated. LFW defines a particular pair-verification protocol (Huang et al., 2008), and the documented scikit-learn loader makes its prescribed files accessible through a reproducible software interface (Pedregosa et al., 2011). Neither source turns a pair count into a count of independent subjects or images. Dataset documentation should identify how references, identities, comparisons, and splits were assembled (Gebru et al., 2021). For generated video, the analogous record should distinguish a new person from another reference photograph, another trajectory, or another frame. Each adds evidence along a different dimension. A transparent report should preserve those links rather than flattening every observation into one score table. The resulting metadata can support a useful audit even when no images or generated clips are available. Such an audit describes coverage and organization; it does not itself establish recognition accuracy, demographic representativeness, or identity preservation in an unobserved video.

Discriminative identity representations address a different question from the design of the test population. FaceNet learns an embedding for verification and related recognition tasks (Schroff et al., 2015). ArcFace (Deng et al., 2018) and CosFace (Wang et al., 2018) develop margin-based approaches to discriminative face learning. These methods do not determine which reference views should be supplied to a generator. Dataset choices also alter the challenge: VGGFace2 emphasizes pose and age variation (Cao et al., 2017), MegaFace increases the scale of recognition distractors (Kemelmacher-Shlizerman et al., 2015), and FairFace studies representation across annotated demographic groups (Kärkkäinen and Joo, 2019). Their relevance is to separate distinct kinds of coverage. Many comparisons of a small reference set cannot replace missing pose variation or a different population of people. A study should state which identity

measurement it uses, how thresholds are selected, and what the represented cases actually cover. Strong performance on one protocol does not establish universal validity of the embedding across generated or stylized subjects.

2. Constructing a recoverable evidence budget

A proposed record would identify source observations before any method-specific preprocessing. It would preserve views, resolution histories, reference selection, adaptation access, and additional model calls where these form part of the workflow. The aim is to describe information received, not merely files uploaded. If a method transforms several images into a mosaic, the source images should remain identifiable. If another consumes them separately, its access should be described at the same source-observation level.

A controlled comparison could restrict both systems to a common source pool and declared selection opportunity. Method-specific processing would still differ, and those differences would be part of the evaluated mechanisms. A complete workflow comparison could instead allow each system's intended interface under a common cost or user-effort constraint. Both designs are legitimate. Their results should not be pooled into a single unexplained superiority claim because they hold different things fixed and answer different application questions.

The tasks, subjects, requested transformations, sample counts, and selection rules would be fixed before final assessment. A method's best output after many retries should not be compared with another's first attempt unless that asymmetric workflow is explicitly the research question. Failed attempts and unassessable clips should remain recorded. Matching only successful final outputs would conceal differences in reliability and effort, especially when one method searches a much larger candidate pool to produce its displayed examples.

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.

A proposed explanation should be distinguishable from alternative explanations using observations the study can actually collect. Diagramming the relationships among selection, exposure, measurement, and outcome makes that requirement explicit (Greenland et al., 1999). The diagram does not establish its own correctness, but it prevents adjustment choices from remaining invisible. The same transparency is needed when investigators inspect many endpoints before choosing a narrative. Controlling false discovery concerns the behavior of a specified testing family (Benjamini and Hochberg, 1995); it is not a general certificate of scientific validity. Primary endpoints, secondary descriptions, and future mechanism tests should be named separately. Readers can then assess whether an attractive pattern was predicted, explored, or merely noticed after the available comparisons were examined.

The reference path and the scoring path should be distinguishable in an experiment. DreamBooth (Ruiz et al., 2022) and IP-Adapter (Ye et al., 2023) illustrate different ways to make subject or image evidence influence generation. CLIP (Radford et al., 2021) and DINO (Caron et al., 2021) illustrate representations

that can be used to compare visual content for other purposes. Reusing the same representation for conditioning and evaluation may create shared sensitivities that deserve examination. It is not proof of invalidity, but it strengthens the case for an independent human or task endpoint. A comparison should report which information enters each path and retain reference-specific artifacts as potential confounds. The generator's ability to reproduce a background, camera style, or distinctive clothing can be useful while remaining different from preserving identity under a requested change. A careful evaluation makes that difference observable.

3. Assessing the resulting contrast

Outcomes should separately measure recognizable subject, requested scene or motion, temporal quality, and coverage. A reference-rich method may obtain high similarity by reproducing familiar appearances while a different method carries out a more demanding transformation. Independent assessment and task-linked endpoints make that difference visible. A weighted overall score is possible if the intended use supplies a clear preference, but the component observations should remain available so that the result does not depend on an undisclosed tradeoff.

The statistical unit should follow the comparison's sampling design. Repeated samples of a subject-prompt pair provide information about generation variation, while new subjects provide a different kind of replication. Matching observations across methods can improve precision when the pairing is valid. It does not eliminate systematic differences in access or measurement. The report should retain the pair identifiers and explain what population the uncertainty interval is intended to represent.

Where interfaces prevent a useful common-evidence contrast, the limitation should be stated directly. It may still be possible to compare practical workflows, but the resulting advantage belongs to the combined system under its observed conditions. Additional ablations can investigate selected components if they are implementable and scientifically meaningful. The proposed design values this bounded interpretation: it is better to know precisely which workflow performed well than to assign its result to an isolated mechanism that the comparison never actually separated.

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.

An automated measurement should not inherit the authority of its training labels without a direct assessment of use. The relevant checks include systematic disagreement with the intended procedure (Bland and Altman, 1986) and calibration of any probability claims (Guo et al., 2017). Ranking performance, absolute accuracy, and confidence reliability are distinct outcomes. A study can improve one while leaving the others unchanged. Keeping the measurements separate permits examination of those tradeoffs and avoids implying that one favorable statistic establishes all aspects of validity. When independent judgments are unavailable, the limitation should remain part of the result. The absence of a target observation cannot be repaired by combining several correlated proxies into a more elaborate total score.

4. Scope, uncertainty, and independent assessment

A moving subject introduces quality dimensions that a still-image identity score cannot exhaust. VBench decomposes video-generation evaluation into multiple aspects (Huang et al., 2023). Structural similarity (Wang et al., 2004), learned perceptual correspondence (Zhang et al., 2018), and video-distribution metrics (Unterthiner et al., 2018) provide additional observations under their respective assumptions. They should not be treated as equivalent measurements of the same latent property. A generated clip can have coherent motion and attractive appearance while changing its subject; it can also preserve recognizable identity while containing visible temporal artifacts. Reporting separate outcomes allows those tradeoffs to be examined. The assessment should specify the frames, viewpoints, occlusion states, and durations included in each score. A difficult case that defeats the identity evaluator should remain part of the coverage record rather than disappearing silently. The resulting profile can support a bounded claim about the observed conditions without implying that every aspect of subject consistency has been solved.

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.

Reproducibility begins with preserving which scientific objects are being linked. Stable identifiers and reusable metadata help connect inputs, reference observations, model versions, and reported outcomes (Wilkinson et al., 2016). The scoring pipeline also has dependencies that can change the interpretation of a result: preprocessing, representation checkpoints, thresholds, and aggregation code can all matter (Sculley et al., 2015). A compact release should retain those decisions rather than only the final mean. Where source data cannot be redistributed, an access description and hashes can still identify the analyzed inputs. This record allows later users to distinguish a change in the underlying system from a change in how the system was measured. Reproducibility is therefore part of interpretation, not simply a final packaging step.

5. Conclusion

Matched evidence requires source-level records and a declared comparison boundary. The proposed protocol distinguishes controlled conditioning tests from complete workflow evaluations and retains outcomes, effort, and missing observations. It reports no performance result. Its contribution is an interpretable basis for future comparisons, allowing a method's practical benefit to be recognized while keeping claims about isolated mechanisms proportional to the information conditions the experiment successfully controlled.

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.