Temporal Smoothing Can Conceal Identity Errors

Robert Pearson (BR), Justin Barrett (WS), Jason Clark (MX)

Abstract

Reducing visible flicker is not necessarily the same as preserving a person. This conceptual analysis examines how temporal averaging, frame selection, or postprocessing can hide short identity substitutions while improving a continuity score. ARGUS provides the motivating subject-preservation setting. No smoothing operation is applied and no model output is analyzed here. The proposed study retains local identity events, motion quality, and measurement coverage before constructing a clip-level summary. It asks whether an apparent improvement concerns the generator, a postprocessing workflow, or a scoring rule that has become less sensitive to the error of interest.

1. Smoothness and identity answer different questions

A sequence can change gradually from one recognizable appearance to another. It may look temporally smooth while failing to preserve the intended person. A sequence can also preserve the subject while exhibiting minor visual flicker. These observations do not make smoothness irrelevant. They show that its relationship to identity should be measured rather than assumed. A combined score can obscure this distinction if it rewards continuity while failing to retain who is present at each relevant moment.

ARGUS motivates an evaluation of subject preservation through video (Meng et al., 2026b). This article does not claim that its implementation applies a particular smoothing operation or that its reported results conceal substitutions. It proposes a diagnostic study of the interaction between temporal processing and identity measurement. No clip is filtered, interpolated, or rescored here. The scientific question is whether a declared temporal operation changes the target behavior or changes only the way an evaluator summarizes it.

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.

A large collection of measurements does not guarantee a large collection of independent cases. Repeated observations can share unrecorded acquisition conditions or a common source, making row-level uncertainty misleading. White's covariance estimator addresses heteroskedasticity within a specified regression (White, 1980), while cluster methods address a declared form of grouped dependence (Cameron and Miller, 2015). Neither approach repairs a poorly defined outcome or establishes representativeness of the observed sample. A report should identify the number of groups, their relative sizes, and the reason for selecting the grouping level. When only a few independent groups are available, apparently precise numerical intervals warrant particular caution. This is a limitation of the evidence structure, rather than a reason to count correlated measurements as fresh observations.

2. Retaining the event before averaging the score

A proposed evaluation would define a local identity event and its observation window before inspecting outputs. It would retain the unsmoothed clip, any processed version, and the exact transformation between them. Identity judgments could then be compared with a temporal-continuity endpoint. A short violation should remain accessible even if the clip's average score is high. The study should state whether such an event invalidates a trajectory, contributes to a rate, or is treated through another declared consequence rule.

Frame selection is another form of temporal processing. Sampling only widely separated frames may miss brief events, while selecting only detectable faces can concentrate the assessment on easy views. A protocol should preserve the eligible time interval and explain how frames are chosen. An unchanged generator can receive a different score under a changed sampling rule. That is an observation about measurement, not necessarily an improvement in generation quality.

The identity instrument also has temporal access. A frame-only model and a sequence-level observer may use different evidence. A sequence-level judgment can legitimately use context to interpret an ambiguous frame, but it can also average over a consequential local change. The study should specify the intended target and compare procedures against an independent task rule where feasible. Neither the shortest nor the longest context window is automatically correct for every use.

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.

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.

A validation split defines which form of transfer has been observed. Related observations can share subjects, scenes, prompts, or acquisition conditions, so randomly splitting rows may leave important information on both sides of the boundary. Structured cross-validation addresses these dependencies (Roberts et al., 2017). Model selection introduces another boundary: choosing a representation, threshold, or tuning parameter using the final assessment data makes the resulting error estimate optimistic (Varma and Simon, 2006). The appropriate held-out unit follows the intended future use. Generalization to a new

frame, a new trajectory, and a new subject are not equivalent claims. A study should record those units before fitting or choosing a scoring rule, and preserve an assessment set that has not become part of iterative development.

3. Evaluating a complete postprocessing workflow honestly

Postprocessing can be a useful part of a generation system. A study does not need to exclude it merely because it follows the model. It should identify its cost, information access, and effect on the task. If smoothing improves both appearance and identity under an independent endpoint, that is a legitimate workflow result. If it improves only the metric, the report should retain that narrower finding. Separating the unprocessed and processed outputs permits the distinction to be examined.

The proposed test should also preserve motion requirements. A method that freezes the face or reduces expression can become easier to match to a reference while failing the requested performance. A low-variation output should not be rewarded automatically when the task requires movement. The outcome profile should therefore include identity, requested motion or expression, general quality, and assessment coverage. A favorable change in one dimension does not supply an observation in another.

Finally, a local event needs an appropriate interpretation. A frame can be difficult because of blur, occlusion, or genuine ambiguity, not because the subject changed. The evaluation should distinguish an unavailable identity judgment from a confirmed substitution. Independent review should receive enough reference and task information to assess that difference. Retaining the original time-indexed evidence is essential, because an aggregate table cannot reconstruct which explanation is plausible after the relevant frames have been discarded.

Subject conditioning is not identical to measuring the resulting subject. DreamBooth studies subject-driven generation through model adaptation (Ruiz et al., 2022), whereas IP-Adapter develops image prompting compatible with text conditioning (Ye et al., 2023). The reference evidence available to those mechanisms should be matched or explicitly described in a comparison. Generic representations also have different training purposes: CLIP connects imagery to language (Radford et al., 2021), while DINO investigates self-supervised visual features (Caron et al., 2021). Either can contribute to a measurement pipeline, but neither automatically defines human identity across generated motion. A study should separate input fidelity, recognizability, and obedience to scene instructions. It should also retain cases in which these objectives conflict. An image-conditioning mechanism can reproduce an accessory or background very well while leaving a more important identity requirement unresolved. Measuring those conflicts directly is more informative than declaring every increase in reference similarity a successful preservation of the person.

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.

4. Connecting an identity claim to its observations

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.

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

Temporal smoothing can improve a useful visual property while leaving identity unresolved or making a local error harder to observe. The proposed evaluation preserves original clips, processing steps, frame selection, and local outcomes before reporting a total. It offers no new result for ARGUS or a smoothing method. Its contribution is an interpretable separation between a better generated trajectory, a better complete workflow, and a measurement procedure that has become less sensitive to the task's consequential failures.

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.