Local Editing without an Unintended Change of Subject

Dylan Collins (TV), Paul Collins (QA), David Allen (NO)

Abstract

A local video edit can satisfy its immediate instruction while changing a person elsewhere in the clip. This article proposes evaluating edits through a declared change region, protected content, and temporal follow-up. ARGUS motivates organized identity evidence, although no editing capability is attributed to it beyond its documented generation setting. No edit experiment is performed. The contribution is a test design that connects successful local modification to independent evidence that the intended subject remains recognizable outside and across the requested change.

1. An edit has both a requested change and a preservation boundary

Changing an accessory, background detail, or local appearance can be a narrow instruction with broad unintended consequences. The generated subject might shift in nearby regions or at a later time even if the edited region looks correct. An evaluation that checks only the requested object would miss that outcome. Conversely, demanding identical pixels everywhere else can reject legitimate temporal adjustments needed to make the edit visually coherent.

ARGUS uses structured references to support subject-preserving video generation (Meng et al., 2026b). That work motivates the role of explicit identity evidence, while the local-editing protocol developed here remains a separate proposed application. The article does not report editing performance or assert that the focal system implements a particular editing interface. Its question is how to define preservation requirements around a local intervention so that an evaluation can distinguish intended changes from an unintended substitution of subject.

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.

2. Recording the intervention and its protected content

A proposed experiment would preserve the original clip, edit instruction, available references, and output version. It would declare which properties are allowed to change and which must remain stable. A spatial mask may be useful, but it cannot fully express temporal consequences or semantic dependencies. For example, a changed object may legitimately alter nearby shadows. The protocol should therefore identify the intended consequence rather than assume that every difference outside a binary region is an error.

The evaluation would assess the requested change separately from subject identity. It could compare the original and edited subject under common viewing conditions, with reference evidence appropriate to the task. The unedited clip is informative but should not automatically be treated as a perfect identity reference. If it already contains a substitution or ambiguous appearance, the study needs to retain that

baseline uncertainty. Otherwise an edit may be penalized for diverging from an earlier error rather than from the intended subject.

The test set should include diverse edit locations and temporal demands under documented selection rules. Selecting only successful-looking edits after generation would estimate a curated workflow rather than unconditional editing reliability. A practical workflow can include search and human selection, but both their effort and their rejected attempts should be reported. Comparisons should match that opportunity across methods or explicitly state that the evaluated systems include different amounts of selection assistance.

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 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.

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.

3. Looking beyond the edited frame

Identity should be examined during and after the local intervention. A system may preserve the person in the edited frame yet propagate altered features into later views. Conversely, a brief local artifact may not change recognizable identity. The proposed protocol would retain time-localized judgments and independent quality outcomes rather than infer all temporal behavior from a single representative image. Relevant intervals should be chosen by a declared rule rather than selected after viewing the most favorable frames.

Automated similarity measurements require an appropriate comparison. Strong similarity outside the edit region can be driven by a static background, while face-only scoring can miss a subject-defining feature that the task protects. A useful study should therefore connect each measurement to a stated preservation requirement. Human assessment can help evaluate the combined result, provided the questions distinguish successful editing, recognizable subject, and visual coherence instead of compressing them into one unexplained overall preference.

The analysis should also identify where a local edit is intrinsically ambiguous. Some instructions change features that an observer uses to recognize a subject. The study can still investigate them, but it must explain what continuity means under that transformation. An explicit uncertainty response is preferable to a forced label whose interpretation differs across reviewers. The final claim should remain tied to the tested edit classes and available evidence, leaving more disruptive transformations for separately defined comparisons.

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.

An association can suggest a research direction without identifying an intervention. Causal diagrams make it possible to state which variables precede the proposed change, which are consequences, and which selection processes remain uncontrolled (Greenland et al., 1999). Their value lies in exposing assumptions rather than decorating a regression with a causal label. A second issue arises when many outcomes or subgroups are explored: selecting the most favorable comparison can exaggerate the impression of support. False-discovery procedures address a defined family of inferential tests (Benjamini and Hochberg, 1995). They do not turn an exploratory pattern into a replicated mechanism. A careful report distinguishes prespecified comparisons, descriptive sensitivities, and hypotheses requiring new observations. This distinction is especially important when an evaluation offers many plausible metrics.

A substitute measurement needs evidence about its relationship to the intended judgment. Two scores may correlate strongly while differing systematically in particular conditions. Method-comparison analysis distinguishes association from agreement (Bland and Altman, 1986). When a system additionally reports confidence, calibration asks whether those probabilities correspond to observed outcome frequencies (Guo et al., 2017). These are separate properties: a metric can rank examples usefully while having a distorted numerical scale or unreliable confidence. An evaluation should therefore retain the original

target observations when possible, describe the scale of each endpoint, and examine disagreements instead of averaging them away. A convenient automated score becomes more credible when its limitations are measured in the conditions that motivate its use.

4. Reference access and measurement reliability

An identity evaluation becomes more informative when it retains the other properties of a generated clip. VBench's multiple dimensions (Huang et al., 2023) help motivate this decomposition. Structural similarity (Wang et al., 2004) compares a particular aspect of image structure; learned perceptual similarity (Zhang et al., 2018) uses another representation; video-distribution metrics (Unterthiner et al., 2018) incorporate temporal visual features. None alone establishes a stable person through every pose and event. A study should identify the relevant time windows and examine local failures as well as sequence averages. It should also distinguish an identity evaluator's inability to score a frame from the generator's inability to preserve a subject. Those are different observations with different remedies. A measurement profile linked to the original clip permits this distinction to be checked and avoids attributing broad capability to a favorable average obtained only on easily measurable views.

Coverage should be reported whenever the system or its evaluator can avoid difficult cases. An abstention policy changes the population on which an accuracy or fidelity score is calculated. Selective classification formalizes the joint reporting of coverage and conditional performance (Geifman and El-Yaniv, 2017). Distribution-shift research adds that estimated uncertainty can behave differently outside the development population (Ovadia et al., 2019). A benchmark should therefore retain which observations were eligible, accepted, rejected, or unavailable. A result on a carefully limited subset can be useful, provided the subset is visible. The challenge is to prevent a narrow success claim from becoming a broad capability claim through omission of the denominator or the circumstances in which the system declines to operate.

A scientific release should make its most consequential choices recoverable. Persistent identifiers and explicit metadata support reuse of observations and results (Wilkinson et al., 2016), while attention to dependencies helps prevent an evaluation from changing unnoticed as software evolves (Sculley et al., 2015). Relevant choices include the source version, preprocessing, accepted input set, metric implementation, and weighting convention. The record should also separate an observed output from a proposed future test. Saving that distinction prevents a plan from later being mistaken for a completed experiment. A small collection of well-linked records can be more useful than a large folder of unexplained scores. The important property is whether an independent reader can reconstruct what was measured and why its reported number has that meaning.

5. Conclusion

Local editing needs an evaluation boundary that includes intended change, protected identity, and temporal consequences. The proposed design retains original clips, references, attempts, and separate assessment dimensions. It reports no editing results. Its purpose is to make a future preservation claim traceable to observations outside and across the edit, so that a visually successful local modification does not silently substitute for evidence that the intended person remains the subject of the video.

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.