User Control when Identity and Scene Instructions Compete

Zachary Turner (HN), Logan Green (BW), Mark Ward (BZ)

Abstract

A reference and a text instruction can place competing demands on a video generator. Users need to know whether preserving a subject, changing a scene, or reproducing a reference appearance receives priority. This article develops an evaluation agenda for that control boundary, motivated by ARGUS. No interface study or generation experiment is conducted. The proposed design records user intent, supplied evidence, and separate outcomes, treating explicit control over a tradeoff as a capability that requires direct observation rather than inference from a favorable aggregate score.

1. A task can contain competing requirements

A user may ask a referenced person to appear under lighting, clothing, motion, or a camera view absent from the supplied images. Some requested changes alter cues that an automated identity score favors. A system might respond by preserving reference appearance, following the instruction more strongly, or producing an unstable compromise. A useful evaluation should identify the intended priority rather than assume that one numerical similarity measure defines what every user would prefer in the conflict.

ARGUS organizes identity references for subject-preserving video (Meng et al., 2026b). This motivates a control question about the relationship between reference evidence and scene requirements. The present article proposes a study of that relationship without attributing an untested interface or control capability to the focal method. It reports no user preferences or generated outcomes. Its objective is to specify how a future experiment could distinguish an actual controllable tradeoff from an accidental change in appearance under different settings.

A reference collection can grow in several ways: more identities, more images of existing identities, or more comparisons among the same images. The LFW protocol makes these units distinguishable (Huang et al., 2008), while its official loader documentation preserves the difference among development and assessment files (Pedregosa et al., 2011). Datasheets provide a framework for explaining why those objects were collected and which uses they support (Gebru et al., 2021). A subject-generation study should apply the same discipline to reference mosaics and output trajectories. The number of reference tiles is not automatically the number of independent observations, and the number of scored frames is not automatically the number of subjects. Recording the hierarchy permits appropriate weighting and uncertainty calculations later. Where a study uses only a public pair list, its conclusions should remain about that list and its protocol. They cannot supply observations of appearance, motion, or recognition errors that were never processed.

The choice of statistical unit should precede the choice of a standard error. A dataset can contain many rows but relatively few independently acquired histories. Heteroskedasticity-consistent covariance protects an otherwise specified regression against unequal residual variances (White, 1980). It does not automatically accommodate observations sharing an environment, subject, or recording. Cluster-robust inference makes that dependence part of the variance calculation (Cameron and Miller, 2015). The grouping must follow the observation process and the claim being made; an identifier that is convenient to store is not necessarily the relevant independent unit. For a descriptive analysis, preserving the groups also allows readers to recalculate a different population average. An uncertainty interval remains conditional on the observations and assumptions actually available.

2. Recording intent and testing a declared control

A proposed task set would include instructions selected for clear, documented conflicts between reference appearance and requested change. The conflict definition should be established before output inspection. Users or task designers would state which properties must remain stable and which may change. An ambiguous instruction can still be studied, but its ambiguity should be recorded rather than resolved afterward by choosing the interpretation that makes a generated clip look most successful.

If a system exposes a control parameter, its meaning should be tested through repeated outputs under declared settings. Changing a number labeled identity strength does not establish that it monotonically controls recognizable identity or preserves all other properties. The evaluation would measure the actual outcome dimensions and sampling variability. If control is expressed through natural-language instructions instead, the exact wording and any revisions should remain part of the record, since iterative prompting creates an additional workflow and effort budget.

Reference access, candidate generation, and user selection should be documented for every condition. A comparison of one system's carefully revised prompt with another's first attempt would evaluate unequal workflows unless that asymmetry is intentional. A complete user-assisted study can include iteration, but it should record attempts and stopping rules. The final accepted output then represents an observed process of control rather than an unexplained sample chosen after the preferred tradeoff became visible.

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.

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.

The protected assessment population should be defined before investigators inspect detailed failures. Once a collection of cases guides repeated revisions, it has become development information even if its filenames still carry an evaluation label. Cross-validation for dependent data treats the grouping of observations as a substantive design decision (Roberts et al., 2017). Research on model selection likewise shows why tuning and final error estimation should not reuse the same information (Varma and Simon, 2006). A practical record therefore distinguishes fitting data, diagnostic examples, validation data, and the

final external assessment. Where only a fixed public table is available, a descriptive analysis remains legitimate, but it should not be renamed a held-out prediction experiment without generating and assessing predictions under a protected split.

3. Evaluating useful control rather than one preferred point

Outcomes would separately assess recognizable subject, scene compliance, motion, visual quality, and coverage. A user-facing preference could then be collected under a declared purpose. This allows the same generated tradeoff to be interpreted differently by users with different priorities. A single aggregate winner cannot establish that a system provides useful control across those preferences. The study should examine whether users can reach their intended region of outcomes with understandable effort and reasonably predictable responses.

Human assessment should distinguish the clarity of the control from satisfaction with a particular clip. A participant might understand a setting yet dislike its achievable tradeoff, or obtain a good clip through repeated trial without understanding the setting. These are different interface observations. Their measurement requires an implemented study with appropriate task instructions and records; the conceptual design here supplies no evidence that users already experience either outcome with the focal generation approach.

The final claim should be bounded to the tested conflicts, users, references, and control mechanism. Success with a modest lighting change does not establish control over every subject-defining transformation. A useful result can nevertheless show that a declared setting or interaction reliably changes the relevant balance under observed conditions. Preserving the separate outcomes makes it possible to identify where control succeeds, where it shifts the wrong property, and where a conflict remains unresolved despite additional user effort.

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.

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

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.

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.

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

Control over identity and scene requirements needs explicit intent, an observable interaction, and separate outcome measures. The proposed agenda reports no new user or generation evidence. It treats a controllable tradeoff as a testable workflow capability, preserving attempts, preferences, and uncertainty. A future study can then show whether users can achieve their intended balance rather than assuming that one favorable similarity score represents the priorities of every reference-conditioned video task.

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.