Ablations that Distinguish Memory, Selection, and Guidance

Mason Bell (ER), Derek Walsh (PL), David Matthews (UA)

Abstract

Removing a component can change information access, computation, and output behavior at the same time. An ablation therefore needs a defined contrast before its result supports a mechanism claim. This article develops an ablation plan for memory, reference selection, and guidance in subject-preserving generation, motivated by ARGUS. No components are removed and no scores are measured. The proposed design records what each intervention changes and distinguishes a useful end-to-end component contribution from evidence about a particular internal causal explanation.

1. A component name is not an intervention definition

Memory, selection, and guidance can interact in a reference-conditioned generator. Removing one stage may alter what later stages receive or how their settings behave. The resulting score difference is real if measured correctly, but its explanation is not automatic. An ablation should describe the replacement behavior, remaining information, and parameter treatment rather than rely on a label such as without memory to communicate an experimental condition that another researcher could reproduce.

ARGUS provides a documented setting with organized identity references, selection, memory, and guidance-related mechanisms (Meng et al., 2026b). This article proposes a way to investigate such components without running an implementation or asserting their individual measured effects. Its question is how to separate a component's operational contribution from changes in information budget or evaluation opportunity. The resulting plan concerns future evidence, not a reinterpretation of unobserved outputs as though a new ablation had already been completed.

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.

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.

2. Specifying replacements and interaction contrasts

A proposed memory contrast would declare what replaces the removed memory path and whether the remaining conditioning receives equivalent source evidence. A selection contrast would preserve the candidate pool and identify the alternative rule, such as a declared fixed choice rather than an unspecified default. A guidance contrast would specify parameter values and other sampling changes. Each condition should be operationally complete, including what happens when a component normally supplies information required by another stage.

The study could distinguish a fixed-configuration ablation from a separately tuned comparison. Holding all settings fixed asks what happens when a component is changed within the original configuration. Retuning each condition asks what performance the modified system can achieve under a declared search budget. Neither dominates every purpose. Their results should remain separate because a retuned condition includes additional optimization opportunity and may answer a broader system-design question than an immediate component removal.

Interactions deserve planned contrasts where feasible. The effect of selection can depend on memory capacity, while guidance can depend on the quality of available identity evidence. A full combination study may be expensive, so a smaller design should state which interactions remain untested. It should not infer independence from an additive presentation of isolated removals. The record of source references, active evidence, compute, and generation settings makes the tested contrast interpretable even when the complete interaction structure cannot be mapped.

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 face embedding is an instrument whose output needs a defined task. FaceNet supports verification and recognition through a learned metric space (Schroff et al., 2015), while ArcFace (Deng et al., 2018) and CosFace (Wang et al., 2018) pursue discriminative margins through different formulations. The scientific population is supplied by the benchmark, not by those objectives. VGGFace2's within-person variation (Cao et al., 2017), MegaFace's distractor scale (Kemelmacher-Shlizerman et al., 2015), and FairFace's representation emphasis (Kärkkäinen and Joo, 2019) show why populations can differ materially. For subject-preserving video, reference access and output conditions should therefore be documented alongside the identity score. An automated evaluator may be accurate on its original photographic population yet behave differently on generated motion, unusual lighting, or stylization. Establishing the relevant transfer requires observations in those settings. Architectural sophistication in the identity instrument cannot replace a direct assessment of the claim the generator is making.

3. Connecting an ablation to its observed consequences

Assessment would preserve subject recognizability, instruction execution, temporal behavior, and coverage. A component can improve one dimension while worsening another. A score chosen because it shares the component's representation may be informative, but it strengthens the case for an independent endpoint. Human judgments or a separately validated instrument can help determine whether an apparent numerical improvement corresponds to the intended preservation capability rather than a shared sensitivity between conditioning and scoring.

Repeated outputs should remain linked to common subjects, tasks, and source references. The analysis could exploit paired comparisons while accounting for the actual replication structure. An ablation should retain failed generations and unassessable cases rather than compare only the surviving outputs. Otherwise a component's effect on reliability can disappear from a performance average. Matched attempt budgets also prevent a condition from benefiting invisibly from additional retries or selection after the generated results are visible.

The final interpretation should separate observed contribution from explanation. A memory removal that reduces preservation under the tested conditions supports a practical dependence within that configuration. It does not by itself establish which stored feature caused the difference or whether an alternative architecture would need the same mechanism. Further controlled observations would be needed for those claims. This boundary preserves the value of ablation evidence while preventing one intervention from being narrated as a complete account of the model's internal behavior.

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 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.

A system's reliable operating region may be narrower than the collection of inputs an interface accepts. Selective prediction separates performance on accepted cases from the fraction of cases accepted (Geifman and El-Yaniv, 2017). Reporting only the former can hide a substantial loss of coverage. Uncertainty estimates may also change under distribution shift (Ovadia et al., 2019), so familiar-condition

confidence is not automatically evidence about new subjects, scenes, or control histories. A useful evaluation records rejected, ambiguous, and unsupported cases alongside successful ones. This makes it possible to distinguish limited but honest capability from a system that answers broadly with unreliable certainty. Any proposed deferral rule still requires assessment on its intended input population; it is not validated merely by being described.

4. Scope, uncertainty, and independent assessment

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.

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.

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

An informative ablation declares its replacement, information boundary, tuning opportunity, and relevant interactions. The proposed plan applies those requirements to memory, selection, and guidance without reporting new experiments. Its contribution is a clearer interpretation of future component contrasts: observed workflow effects can be stated directly, while specific mechanism claims remain conditional on the controls and observations that the ablation actually provides.

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.