WHAT REMAINS AFTER CONCEPT ERASURE? SPARSE AUTOENCODER PROBES FOR AUDITING DIFFUSION MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Concept erasure aims to remove selected visual concepts from text-to-image diffusion models while preserving the generation of non-target concepts. Evaluation typically measures target suppression and non-target preservation in generated images, but a concept's absence from these outputs does not establish its removal from the model. We introduce `SAE-Audit`, a white-box audit using sparse autoencoder (SAE) probes to test whether activation interventions can restore target-concept generation without changing model weights or test prompts. After verifying target suppression, we compare the original and edited models using a shared SAE representation learned from original-model activations, identifying features whose target-related responses weaken after erasure. We inject sparse directions derived from the original-model responses into the edited model, keeping these directions fixed across test prompts, and evaluate recovery on held-out scene families. Experiments span six erasure methods and four domains: objects, artistic styles, copyrighted content, and nudity. Interventions using a few dozen selected SAE features yield higher average recovery rates than matched random-feature interventions, with recovery concentrated in objects and nudity following cross-attention edits. Most styles and copyrighted concepts show limited recovery under the tested intervention settings. These findings show that output suppression can coexist with responsiveness to sparse concept directions supplied from the original model.

1 INTRODUCTION

Text-to-image diffusion models support visual content creation, design, and image editing, allowing users to generate and modify images through natural-language instructions (Rombach et al., 2022; Brooks et al., 2023). Models trained on large-scale image–text datasets can generate sexually explicit or violent imagery, creating content-safety risks (Schramowski et al., 2023; Gandikota et al., 2023). Additionally, their ability to reproduce recognizable characters or imitate artistic styles raises concerns about unauthorized use (Lu et al., 2024; Liu et al., 2025). Concept erasure aims to address these concerns by modifying a diffusion model to suppress specified visual concepts while preserving the generation of non-target concepts (Shi et al., 2026). Unlike rejecting input prompts or filtering generated images, concept erasure intervenes in the model's generation process.

Existing methods either modify the model's learned weights or alter the internal activations computed during generation. Weight-editing methods modify selected denoising or attention parameters (Gandikota et al., 2023; 2024; Li et al., 2026; Lu et al., 2024), or the text encoder (Farashah et al., 2026). Other approaches suppress concept-associated features during generation (Cywiński & Deja, 2025). Two edits can reduce the frequency of target-concept generation by similar amounts while affecting concept-related activations differently.

However, concept erasure is commonly evaluated through target-concept occurrence in generated images, alongside image quality and preservation of non-target concepts (Gandikota et al., 2023; 2024; Lu et al., 2024). EMMA extends evaluation to implicit descriptions of erased concepts and retention of visually similar non-target concepts (Wei et al., 2026). These measurements establish suppression under the evaluated prompts, not whether the edited model can generate the target under other conditions. Prior audits have shown that concepts suppressed under standard prompts

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

can reappear when text conditioning is modified (Pham et al., 2024; Zhang et al., 2025) or visual context and generation-time interventions are used (Lu et al., 2025), without updating diffusion-model weights. In addition, lightweight fine-tuning can also reinstate erased behavior (Liu & Zhang, 2025). This suggests that while an erasure edit may suppress the target concept under some pre-defined prompts, it may still fail when a small set of concept-related activations is changed.

In this paper, we introduce `SAE-Audit`, a white-box audit that tests whether interventions on concept-related features can restore target-concept generation after erasure. The key idea is to identify concept-related responses that weaken after erasure and construct targeted interventions using their original-model values. After verifying target suppression, we measure these responses by contrasting feature activations under matched prompts with and without the target concept in each model. Sparse autoencoders (SAEs) trained only on original-model activations provide a shared feature basis for this comparison. We use the original-model contrasts on the selected features to construct directions that are added to the edited model's intermediate activations during generation. These directions remain fixed across test prompts and seeds, and we evaluate recovery on held-out scene families without changing model weights or test prompts.

We evaluate `SAE-Audit` across 6 erasure methods and 13 concepts spanning objects, artistic styles, copyrighted content, and nudity. Our results show that directions constructed from a few dozen SAE features can reactivate suppressed concepts on held-out scenes, achieving higher mean detector-measured recovery rates than matched random-feature interventions. Recovery is concentrated in objects and nudity, while most styles and copyrighted concepts show limited recovery under the tested intervention settings. Further analysis reveals a distinction between detector-level recovery and faithful concept generation: some outputs reproduce concept-associated textures or parts without recovering the full concept in its scene. Together, these findings show that output suppression can coexist with target-concept recovery under fixed, original-model-derived interventions. `SAE-Audit` thus complements output-based evaluation with a controlled, feature-level test of whether concept suppression withstands targeted activation interventions.

2 Related Work

Concept Erasure in Diffusion Models. Concept erasure operates through parameter updates or inference-time interventions. ESD uses negative-guidance fine-tuning (Gandikota et al., 2023), UCE applies closed-form cross-attention edits (Gandikota et al., 2024), and MACE combines attention refinement with LoRA fine-tuning and fusion (Lu et al., 2024). RECE iteratively erases concept-recovering embeddings (Gong et al., 2025), whereas Prototype-Guided Concept Erasure uses multiple prototypes for inference-time negative guidance (Cai et al., 2026). We study how erasure affects feature responses and recoverability rather than propose another erasure method.

Auditing Concept Erasure. Recovery-based audits probe erasure beyond standard prompts. Concept Inversion learns recovery embeddings (Pham et al., 2024), while UnlearnDiffAtk optimizes adversarial prompts (Zhang et al., 2025). Beyond text conditioning, Lu et al. (2025) distinguish guidance-based avoidance from destruction-based removal using visual context, trajectory perturbations, and latent classifier guidance. TINA uses optimized image inversion under an empty text condition (Xiang et al., 2026), while parameter-level probes test reversibility through reverse-guided fine-tuning or few-shot personalization (Liu & Zhang, 2025). These probes differ in model access and the concept information they supply. `SAE-Audit` complements them by relating pre- and post-edit feature contrasts to controlled recovery, comparing directions selected from features whose contrasts persist or are suppressed after editing, while keeping model parameters fixed.

Sparse Autoencoders for Diffusion Models. SAEs decompose diffusion-model activations into sparse features for individual analysis or modification. SAeUron suppresses concepts by modulating concept-associated SAE features (Cywiński & Deja, 2025). SAEmnesia concentrates concept responses through supervision (Cassano et al., 2026b), while SAEParate improves concept separation through contrastive training (Kim et al., 2026). Cassano et al. (2026a) show that SAE features can detect concepts, but direct steering can produce out-of-distribution activations and visual artifacts. Using frozen SAEs trained on original-model activations as a shared feature basis, we compare concept-related responses before and after erasure and test original-model-derived recovery directions against matched controls.

3 PROBLEM FORMULATION

Setting. Let $\mathcal{M}^N$ denote an original text-to-image diffusion model and $\mathcal{M}^E$ its edited counterpart, obtained by applying a concept-erasure procedure to a target concept c. Given a text prompt p, a random seed s, and a fixed sampling configuration, their generated images are

$$x^m(p;s) = \mathrm{Generate}(\mathcal{M}^m, p, s), \qquad m \in \{N, E\}. \tag{1}$$

Concept erasure aims to suppress the generation of the target concept c while preserving the generation of non-target concepts. We study cases in which the edited model exhibits this suppression under ordinary prompting. Such suppression establishes a change in output behavior, but does not explain how concept-related internal responses have changed or whether the model remains responsive to interventions on those representations.

Audit objective. Our audit asks whether activation interventions can restore target-concept generation after erasure. We examine how recovery varies with changes in concept-related activations between the original and edited models. We assume white-box access to the original model $\mathcal{M}^N$ and the edited model $\mathcal{M}^E$. Both models remain frozen during the audit, and $\mathcal{M}^N$ provides the reference activations.

Let $\mathfrak{I}$ denote a prescribed class of interventions on intermediate activations, i.e., the hidden-state tensors computed at internal layers of the denoising network during sampling. An intervention $\mathcal{I} \in \mathfrak{I}$ changes the internal generation process without modifying the input prompt or model weights, and its rule is shared across prompts and seeds rather than optimized separately for each input. Writing $x^{E,\mathcal{I}}(p;s)$ for the resulting image, let $F_c(x)$ be a fixed evaluator score for target concept c. We quantify the intervention's effect on generated images by the change in target-concept score:

$$\Delta_c(\mathcal{I}; p, s) = F_c\big(x^{E,\mathcal{I}}(p;s)\big) - F_c\big(x^E(p;s)\big)\,. \tag{2}$$

The research question is whether controlled interventions can consistently restore target-concept generation, and how this responsiveness relates to the concept-related internal responses that persist or become suppressed after editing. The construction of the intervention class is developed in Section 4.

Interpretation. This formulation distinguishes the persistence of a concept-related response from the ability of an intervention to support concept generation. Neither property is assumed to imply the other. Moreover, interventions derived from the original model may themselves supply concept information. Successful recovery therefore establishes that the edited model remains responsive to the supplied intervention; it does not, by itself, establish that an intact concept representation survived erasure. Conversely, unsuccessful recovery constrains only the intervention class examined, rather than certifying complete concept removal.

4 SAE-Audit

We audit one cell, a concept–model pair $(c, \mathcal{M}^E)$, at a time. We instantiate the intervention class $\mathfrak{I}$ of Section 3 as activation injections derived from a sparse set of SAE features, fixed within each audit and shared across prompts and seeds. Figure 1 outlines the five-stage workflow for constructing these probes and evaluating target-concept recovery.

4.1 STAGE 1: VALIDATED AUDIT UNITS AND ACTIVATION PROBING

We first identify units on which the original model generates the target concept and the edited model suppresses it. For a concept c we prepare J matched prompt pairs (p_j^+, p_j^-), $j = 1, \ldots, J$: the *positive* prompt p_j^+ names the concept, and the *negative* prompt p_j^- is identical except that the concept is replaced by a neutral counterpart (e.g., *cat→dog* in Table 5). The pairs are grouped into scene *families*, and discovery/validation/test splits are made at the family level (Section 5.1; Appendix C.1). A candidate pair–seed unit (j, s) is eligible if

$$\mathcal{D}_c^{\mathrm{audit}} = \Big\{(j,s):\ F_c\big(x^N(p_j^+;s)\big) \ge \tau_c^+,\ \ F_c\big(x^N(p_j^-;s)\big) \le \tau_c^-,\ \ F_c\big(x^E(p_j^+;s)\big) \le \tau_c^E\Big\}. \tag{3}$$

These conditions require target elicitation without negative-prompt leakage in the original model, and suppression in the edited model. Units failing only the third condition are reported separately

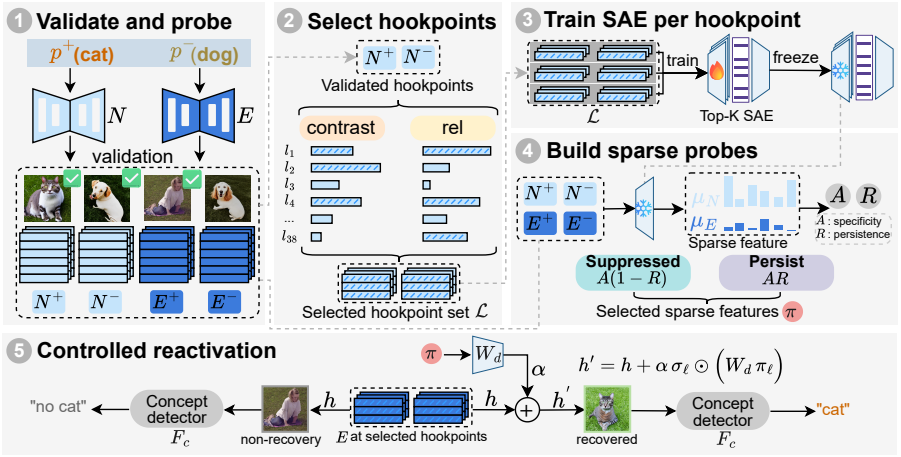


Figure 1: **Overview of `SAE-Audit`. (1) Validate and probe:** verify concept elicitation in the original model and suppression in the edited model, then collect both models' activations under matched positive–negative prompts. **(2) Select hookpoints:** rank blocks by original-model concept responsiveness using *contrast* or *rel* (Eq. (4)); the audit is run separately under each criterion. **(3) Train SAE per hookpoint:** train a Top-K SAE on original-model activations only, then freeze it to encode both models in a shared basis. **(4) Build sparse probes:** score features by concept specificity A and post-erasure contrast persistence R, and form sparse injection prototypes from original-model contrasts on selected features. **(5) Controlled reactivation:** inject the decoded prototypes into the edited model at a validation-selected strength and evaluate recovery on held-out scene families against matched random-feature controls, without changing model weights or test prompts.

as *direct erasure failures*. Thresholds are calibrated on original-model discovery scores and frozen; cells with insufficient eligible units in any split are not audited.

For each eligible unit, we collect conditional-branch activations $h^{\pm}_{m,\ell,t}(j,s)$ from both models $m \in \{N, E\}$ under both prompt polarities, at fixed sampling steps t and hookable UNet blocks ℓ (*hookpoints*). All four runs use the unit's seed but follow their own trajectories. Spatially pooled activations support hookpoint screening in Stage 2; sampled tokens at the selected hookpoints support SAE training and feature comparison in Stages 3–4. The common-state probing variant is described separately in Appendix B.4.

4.2 Stage 2: Hookpoint screen

Because the injected values are taken from the original model's concept response, Stage 2 confines SAE training and injection to the few hookpoints where that response is most pronounced. We rank the hookpoints by how strongly and how consistently the original model's activation responds to the concept. With $\bar{h}^N_\ell(p; s)$ the original model's activation at hookpoint ℓ pooled over positions and probed steps, and $\delta_\ell(j,s) = \bar{h}^N_\ell(p^+_j; s) - \bar{h}^N_\ell(p^-_j; s)$ the concept effect of a unit, we compute on discovery units

$$\mathrm{contrast}(\ell) = \frac{\left\|\mathbb{E}_{(j,s)}[\delta_\ell]\right\|_2}{\left\|\mathrm{Std}_{(j,s)}[\delta_\ell]\right\|_2}, \qquad \mathrm{rel}(\ell) = \frac{\left\|\mathbb{E}_{(j,s)}[\delta_\ell]\right\|_2}{\mathbb{E}_{(j,s)}\left\|\bar{h}^N_\ell(p^-_j; s)\right\|_2}. \tag{4}$$

The two criteria capture different properties of concept responsiveness. contrast measures the mean concept effect relative to its variability across prompts and seeds, whereas rel measures the same effect relative to the magnitude of background activations under p^-. Their usefulness for recovery may differ across concepts and erasure methods. We therefore evaluate both within the same audit pipeline, rather than assume that one ranking is suitable for every concept–erasure method pair. For each criterion, the top-ranked hookpoints form $\mathcal{L}$, and the remaining stages construct and test an intervention at those blocks. Both rankings are computed from original-model discovery activations.

4.3 Stage 3: Normal-anchored sparse dictionary

Stage 3 provides a sparse basis shared by the two models, so that their concept responses can be compared feature by feature rather than as dense activation vectors. One Top-K SAE (Gao et al., 2025) is trained per hookpoint $\ell \in \mathcal{L}$ on spatial tokens in $\mathbb{R}^{d_\ell}$, channel-normalized by frozen corpus statistics (μ_ℓ, σ_ℓ):

$$\widetilde{h} = \frac{h - \mu_\ell}{\sigma_\ell}, \qquad z = E_\ell(\widetilde{h}) = \mathrm{TopK}_K\big(\mathrm{ReLU}\big(W_e(\widetilde{h} - b_d) + b_e\big)\big), \qquad D_\ell(z) = W_d\, z + b_d, \tag{5}$$

minimizing $\|\widetilde{h} - D_\ell(E_\ell(\widetilde{h}))\|_2^2$ with unit-norm decoder columns, termed the *dictionary*; sparsity is structural ($\|z\|_0 \le K$) and identical across hookpoints. The SAE is *normal-anchored*: its corpus holds original-model tokens from discovery units under both polarities, and nothing from $\mathcal{M}^E$ enters training. The frozen SAE then encodes both models in a basis that the edit did not shape, so a code value read from $\mathcal{M}^N$ can be written into $\mathcal{M}^E$. Reconstruction quality is reported but is neither a gate nor a tuning target: Stage 5 passes the reconstruction residual through untouched, so the SAE contributes only the injected direction, and a single-cell diagnostic finds recovery unchanged when one dictionary is degraded from EV 0.78 to 0.47 (Appendix D).

4.4 Stage 4: Suppressed features and the injection prototype

We encode stored tokens with the frozen SAE and average the resulting codes over tokens and probed steps, obtaining $z^{\pm}_{m,\ell}(j,s)$. Let $d_{m,\ell}(j,s) = z^{+}_{m,\ell}(j,s) - z^{-}_{m,\ell}(j,s)$ be the paired contrast and $\mu_{m,\ell}$ its discovery-split mean. For each feature i, we measure original-model concept specificity and post-erasure contrast persistence as

$$A_{\ell,i} = \frac{|\mu_{N,\ell,i}|}{\mathrm{Std}_{(j,s)}[d_{N,\ell,i}(j,s)] + \epsilon}, \qquad R_{\ell,i} = \mathrm{clip}\left(\frac{\mathrm{sign}(\mu_{N,\ell,i})\,\mu_{E,\ell,i}}{|\mu_{N,\ell,i}| + \epsilon},\, 0,\, 1\right). \tag{6}$$

Values of R near zero indicate attenuated or reversed contrasts, whereas values near one indicate retained or amplified aligned contrasts. We rank features by

$$V_{\ell,i} = A_{\ell,i}(1 - R_{\ell,i}) \;\; (\textit{suppressed}), \qquad V_{\ell,i} = A_{\ell,i} R_{\ell,i} \;\; (\textit{persist}). \tag{7}$$

The main probe uses the suppressed-feature ranking; the persisting-feature ranking provides a comparison control (Appendix E). Suppression here concerns the contrast, not the absolute activation.

The top-n features per hookpoint form the support $\mathcal{S}_c$. We call the support together with its prototype values the cell's *circuit*. The injection prototype assigns the original-model contrast to these coordinates:

$$\pi_{\ell,i} = \mu_{N,\ell,i} \cdot \mathbf{1}[(\ell,i) \in \mathcal{S}_c]. \tag{8}$$

Thus, both models inform feature selection, while the prototype values come from the original model. A matched random-feature control places a permutation of the same values on an equally sized, disjoint support. Additional control constructions are given in Appendix B.2.

4.5 Stage 5: Counterfactual reactivation of the erased model

Stage 5 answers the audit question of Section 3: it turns the prototype, a fixed vector in code space, into an activation-space intervention on $\mathcal{M}^E$ and tests on held-out scenes whether it restores the concept more than a matched random prototype does. During generation with $\mathcal{M}^E$, forward hooks at $\ell \in \mathcal{L}$ edit the activation at every step and spatial position on the conditional branch: each spatial vector is normalized, encoded, displaced along the prototype, and decoded *with the reconstruction residual carried through*,

$$\widetilde{z} = E_\ell(\widetilde{h}) + \alpha\, \pi_\ell, \qquad \widetilde{h}' = D_\ell(\widetilde{z}) + \big[\widetilde{h} - D_\ell(E_\ell(\widetilde{h}))\big], \qquad h' = \mu_\ell + \sigma_\ell \odot \widetilde{h}'. \tag{9}$$

The bracket passes whatever the SAE cannot reconstruct through unchanged, so the edit perturbs the model's own activation rather than replacing it. Because the decoder is affine, Eq. (9) equals

$$h' = h + \alpha\, \sigma_\ell \odot \big(W_d\, \pi_\ell\big), \tag{10}$$

a fixed, prompt-independent displacement in the model's basis; all the SAE contributes is the direction $W_d \pi_\ell$, which is what the SAE-free controls test. This is consistent with Cassano et al. (2026a), who find SAEs useful for localizing concept-relevant features while direct editing produces artifacts.

Table 1: Recovery results across concepts and erasure methods. Each cell reports the strict ASR of *ours*/*rand* under the ranking criterion yielding the stronger observed test outcome. Selection prioritizes scene-consistent recovery, followed by recovery alone and non-recovery, with ties broken by the strict ASR of *ours*. Both ASRs and the verdict come from the same audit. Tables 11 and 12 report results separately for each criterion. Background identifies the selected criterion: contrast or rel. Bold: recovered cell; red also meets the scene criterion, while blue does not meet its required cell-level margin, which need not mean every image loses the scene. †: not audited (Stage 1).

Domain	Concept	Original	ESD-x *ours / rand*	ESD-u *ours / rand*	UCE *ours / rand*	SPEED *ours / rand*	MACE *ours / rand*	SDErasure *ours / rand*
Object	cat	1.00	**0.74** / 0.00	0.10 / 0.00	**0.76** / 0.00	**0.72** / 0.03	**0.62** / 0.00	**0.81** / 0.06
	bicycle	1.00	**1.00** / 0.00	**1.00** / 0.00	**1.00** / 0.00	†	**1.00** / 0.00	**0.81** / 0.44
	truck	1.00	**0.96** / 0.04	**0.79** / 0.08	**0.62** / 0.00	**0.71** / 0.00	**0.34** / 0.00	**0.80** / 0.05
	fire hydrant	1.00	**0.69** / 0.00	0.18 / 0.00	**0.92** / 0.04	**1.00** / 0.09	**0.56** / 0.00	**0.69** / 0.03
Style	Van Gogh	0.68	0.12 / 0.00	0.00 / 0.00	0.15 / 0.00	†	†	†
	Monet	0.47	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00
	Seurat	0.84	**0.94** / 0.00	**0.55** / 0.00	**0.78** / 0.00	†	**0.91** / 0.00	**0.81** / 0.14
	Kirchner	0.89	0.05 / 0.00	0.00 / 0.00	0.06 / 0.00	†	0.06 / 0.00	0.15 / 0.00
Copyright	Pikachu	1.00	0.09 / 0.03	0.03 / 0.00	0.00 / 0.00	0.21 / 0.03	0.06 / 0.00	**0.25** / 0.00
	ASICS logo	0.78	0.09 / 0.00	0.00 / 0.00	**0.27** / 0.00	†	0.00 / 0.00	0.03 / 0.03
	Stormtrooper	1.00	**0.78** / 0.00	**0.50** / 0.00	**0.35** / 0.00	†	**0.31** / 0.00	†
	Spider-Man	1.00	**0.24** / 0.00	0.04 / 0.00	**0.55** / 0.00	†	0.03 / 0.00	**0.33** / 0.05
NSFW	nudity	0.91	**0.71** / 0.06	**0.48** / 0.00	**0.59** / 0.05	**0.74** / 0.00	**0.77** / 0.00	†

Each cell runs four arms on the same units, seeds and evaluator: the original model (*n_ref*, upper bound), the erased model without intervention (*baseline*, lower bound), the prototype (*ours*) and the matched random prototype (*rand*). The strength α is swept on validation and fixed at the value maximizing the mean evaluator score of *ours*; *rand* runs at the same α. The test split is then generated once. A cell is *recovered* only if, on test units, *ours* exceeds *rand* in a paired test and by a fixed margin in strict success rate (Section 5.1). The audit thus returns one of three outcomes per cell: direct erasure failure (Stage 1), not recovered at the tested sites, or recovered. A per-feature causal screen and a structure floor on the strength selection, two devices of the full design that did not gate the reported matrix, are described in Appendix B.3.

5 Experiments

We apply `SAE-Audit` to the full concept×method matrix of one backbone, six erasure methods across thirteen concepts, and ask four questions in order: which cells are auditable (Stage 1 in Appendix F); whether a fixed sparse injection reactivates a verifiably erased concept on held-out scenes while matched controls fail, and where and for which concepts it does (Section 5.2); how the selected support compares with other interventions at the same access level (Section 5.3); and what the persisting-feature control and the two site-selection criteria say about the injection (Section 5.4).

5.1 Experimental Setup

Backbone, erasure methods and concepts. We use Stable Diffusion v1.4 (Rombach et al., 2022) with deterministic DDIM sampling (Song et al., 2021); checkpoints differ only in UNet weights. In the evaluated checkpoints, ESD-x (Gandikota et al., 2023), UCE (Gandikota et al., 2024), SPEED (Li et al., 2026), MACE (Lu et al., 2024) and SDErasure (Miao et al., 2026) have changes confined to cross-attention projections, whereas ESD-u has changes only in non-cross-attention modules (Appendix C). The thirteen concepts comprise four objects, four artistic styles, four copyrighted characters and logos, and one NSFW concept. Below we refer to the edited model $\mathcal{M}^E$ as the *erased model*.

Prompts, splits and concept detector. Each concept has five scene families (e.g., a street, a close-up portrait), each containing four matched prompt pairs evaluated with ten seeds ($N = 200$ candidate units). Negatives replace the concept with a same-context counterpart (Table 5). Families are split $3/1/1$ into discovery, validation and test ($120/40/40$ candidates); after Stage 1, valid-unit caps are

$80/40/40$, with only discovery truncated (Appendix C.2). Seeds are shared across splits: *held out* refers to scenes, not seeds. The score $F_c(x) \in [0, 1]$ comes from a frozen concept detector per domain (Table 4); no detector was trained on this study's images, and label sets and strict thresholds were fixed on pilot images before any test read (Appendix C.1).

Metrics. All arms use the same valid test units $\mathcal{D}_c^{\text{test}}$. We report strict attack success rate (ASR) and the ours-minus-random difference:

$$\text{ASR}_a = \frac{1}{|\mathcal{D}_c^{\text{test}}|} \sum_{(j,s)\in\mathcal{D}_c^{\text{test}}} \mathbf{1}\big[F_c(x_{j,s}^a) \geq \tau_c^{\text{strict}}\big], \qquad \Delta = \text{ASR}_{\text{ours}} - \text{ASR}_{\text{rand}}, \tag{11}$$

where $x_{j,s}^a$ is the image generated by arm a. A cell is *recovered* if *ours* exceeds *rand* on per-unit scores (one-sided paired Wilcoxon test, $p < 0.05$) *and* $\Delta \geq 0.2$; throughout, *recovered* means detector-measured reactivation under the fixed intervention. Because detector scores can also reward textures or fragments, we additionally report the *scene-preservation ratio*

$$\rho_a = \frac{\cos\big(e(x_{j,s}^a),\, e(x_{j,s}^E)\big)}{\cos\big(e(x_{j,s}^N),\, e(x_{j,s}^E)\big)}, \tag{12}$$

with $e(\cdot)$ a frozen CLIP image embedding (Appendix H.1). The ratio compares the intervened image's similarity to the erased image with the original–erased similarity for the same unit. A strict success with $\rho_a \geq 0.9$ is *scene-consistent*: its cosine similarity is at least 90% of that reference similarity. SC_a is the fraction of *all valid test units* meeting both conditions, and a recovered cell is scene-consistent if additionally $\text{SC}_{\text{ours}} - \text{SC}_{\text{rand}} \geq 0.2$. This embedding-based criterion does not establish concept identity or semantic fidelity (Appendix H.2). Stage-1 elicitation, negative leakage, direct-failure and valid rates are reported in Appendix F.

Configuration and hyperparameters. Each audit uses the top 3 of the 38 hookpoints under one ranking criterion, Top-K SAEs (Gao et al., 2025) with $K = 32$, width $4d_\ell$ and $8{,}000$ training steps, and 10 suppressed features per hookpoint (30 in total). The prototype is the mean discovery contrast, injected at all 50 steps; per audit, $\alpha \in \{0.5, 1, 2, 4, 8, 16, 32\}$ maximizes the validation mean detector score of *ours*, and *rand* uses the same α. Every eligible cell is audited twice, once under each criterion of Eq. (4): the two audits share units, screen activations, detector and settings, but select different hookpoints (never the same three) and build their own SAEs, features, prototype and strength. The persisting-feature null control is run on the contrast audits and the same-access controls of Section 5.3 (Appendix I.1) under both criteria. Table 6 lists the shared settings, Appendix J.1 their selection history, including the pilot mean/median comparison in which the mean was reinstated after test feedback, and Appendix J.2 the choice of strength grid, SAE size and hookpoint criterion.

5.2 Sparse Reactivation on Held-Out Units

Stage 1 admits 67 of the 78 cells (Appendix F): about a fifth of all candidate units are direct erasure failures, and 11 cells lack enough eligible units for recovery testing. Table 1 shows each admitted cell at the stronger of its two audits; Tables 11 and 12 report each criterion separately at its validation-selected strength, Appendix G.1 gives unit counts, thresholds, strengths, p-values and recovery gaps, and Table 16 gives scene consistency per criterion.

Recovery is concentrated in objects and nudity. Recovery is observed in 40 of the 67 audited cells under at least one of the two criteria; because each cell is read twice on the same test split, this union is an upper count, and the per-criterion counts follow below. This includes 26 of 28 object and nudity cells, with cat and fire hydrant under ESD-u the two exceptions, compared with 14 of 39 style and copyrighted-concept cells ($5/19$ styles and $9/20$ copyrighted concepts). Most styles and copyrighted concepts therefore remain below the recovery criterion even when both criteria are considered. Figure 2 illustrates selected successful units rather than typical outputs or the frequency of faithful recovery.

The two ranking criteria provide complementary recovery coverage. Using rel, the audit recovers 37 of 67 cells, compared with 22 for contrast; 19 are recovered under both, 18 only under rel and 3 only under contrast. Neither criterion covers all cases recovered by the other. The overlap is greatest for objects and nudity, where 18 of the 26 recovered cells pass under both. For styles and copyrighted

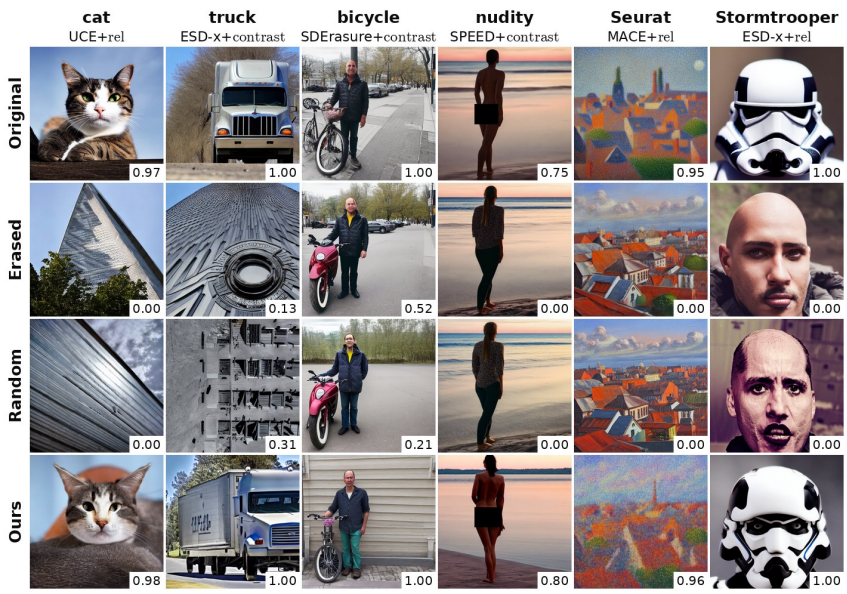


Figure 2: Recovered held-out test units under the original model, the erased model, *rand* and *ours* (same prompt and seed down each column; numbers are detector scores); rel and contrast mark two hookpoint criteria. Regions detected with nudity are masked.

concepts, 11 of the 14 recovered cells pass only under rel; the audit under contrast recovers no style cell and only three character cells. The audit follows the same pipeline under both criteria, with dictionaries and interventions constructed separately for each (Section 5.4).

Detector-level recovery does not always satisfy scene consistency. In Table 1, 31 of the 40 recovered cells also meet the CLIP-based scene criterion ($21/26$ for objects and nudity, $10/14$ for styles and copyrighted concepts). Of the other nine, six lose the scene: bicycle under ESD-x, ESD-u, UCE and MACE, whose repeated frames displace the prompted person, fire hydrant/UCE and Pikachu/SDErasure. The remaining three, Stormtrooper/MACE and Spider-Man under ESD-x and SDErasure, are borderline recoveries (strict ASR 0.24–0.33) whose successes mostly keep the scene but fall short of the required margin over *rand* once the few that do not are excluded (Appendix H.2). Conversely, passing the scene criterion does not verify identity: all five recovered Seurat cells pass, but a relabel check on two of them shows sensitivity to a competing pointillism label (Appendix C.3).

5.3 Comparison with Recovery Baselines

We compare `SAE-Audit` with every intervention in the study that has the *same access* as the audit: the frozen erased model, no image of the target, and one generation per test unit. Reverse-SAeUron selects SAE features by the importance score of SAeUron (Cywiński & Deja, 2025) and injects the original model's values on them through the same residual path; the other arms inject the ten features with the largest original-model contrast $|\mu_N|$ per hookpoint (magnitude-only, no suppression screening), a dense mean-difference or linear-probe (CAV) direction, another concept's circuit or energy-matched random supports, all at the displacement norm of *ours*, or search the sampler noise (Appendix I.1). Every arm runs on both audits of every cell with the cell's units, seeds and detector. Table 2 reports the rel audits, at which every activation-space arm recovers more cells, and Table 17 the contrast audits; per arm they count the recovered cells (strict ASR at least 0.2 above the mean of the 16 random draws, without the paired test of Section 5.1; under this rule *ours* counts 38 cells at rel against 37 in Table 12, the extra cell being Pikachu/SPEED), the cells on which *ours* is better (one-sided paired Wilcoxon test, $p < 0.05$) and those on which the arm's strict ASR is higher.

At equal access, the selected sparse support works best on objects, copyright and nudity. *ours* has the highest mean strict rate and recovers the most cells. On objects the three concept-selected sparse supports, *ours*, Reverse-SAeUron and the magnitude-only arm, lead; on copyrighted content and nudity *ours* and Reverse-SAeUron lead; the dense, wrong-concept, random and sampler-search

Table 2: Same-access recovery interventions on the 67 audited cells at the rel hookpoints: strict ASR overall and per domain, the number of recovered cells (strict ASR at least 0.2 above the mean random draw, without the paired test of Sec. 5.1), and the paired comparison with Ours. Bold: best of the same-access arms in the column; underline: second best.

Intervention	Mean	Object (23)	Style (19)	Copyright (20)	NSFW (5)	Recovered	Ours wins	Arm higher
Original SD 1.4 (upper bound)	0.89	1.00	0.71	0.94	0.91	–	–	–
Erased model, no intervention	0.00	0.00	0.00	0.00	0.00	–	–	–
Ours (residual sparse circuit)	**0.40**	**0.68**	0.24	0.17	**0.66**	**38**	–	–
Reverse-SAeUron (adapted)	0.35	0.62	0.22	**0.20**	0.21	33	18	23
Highest-activation features	0.22	0.57	0.08	0.00	0.05	20	36	9
Dense mean-difference add	0.20	0.20	**0.37**	0.04	0.18	23	37	19
Linear CAV add	0.05	0.07	0.08	0.01	0.06	4	47	8
Wrong-concept circuit	0.04	0.04	0.09	0.01	0.00	5	48	2
Energy-matched random support (16 draws)	0.02	0.04	0.01	0.01	0.05	–	53	0
Noise-based probe, Success@1	0.04	0.09	0.00	0.02	0.09	1	53	3

arms stay low on these three domains. The margin over Reverse-SAeUron is modest and not uniform: *ours* is better at $p < 0.05$ on a quarter of the audits and Reverse-SAeUron has the higher strict rate on a third; by domain, *ours* leads on objects and nudity and Reverse-SAeUron on copyrighted content. Styles are where a dense direction does better: the mean-difference add reactivates the brushwork of Van Gogh, Monet and Kirchner, on which no sparse selector recovers more than one cell, whereas the sparse selectors recover only Seurat (Appendix I.2). The suppressed-feature selector is thus the best selector on average, not on every cell or domain, and *ours* and Reverse-SAeUron also lead at the contrast hookpoints, at lower rates (Table 17).

5.4 Analysis

Across the 67 contrast audits, the *suppressed*-feature intervention is at least as effective as the *persist*-feature control in almost every cell, and the few reversals occur where both have low ASR (Figure 3a). This supports selecting features by how their concept contrasts change after erasure, but does not establish an intact downstream readout: the intervention supplies original-model values, and in a single-cell probe it also adds the concept to non-target prompts (Appendix K). Its outputs also reveal a gap between detector response and complete concept generation, such as cat fur without a cat or painterly texture without the prompted still life (Appendix L). This is consistent with a spatially uniform direction restoring some concept-associated attributes without reliably restoring their arrangement, but does not establish a general texture-versus-identity distinction across domains.

ESD-u has no recovered cell under contrast, but the rel audit recovers bicycle, truck, Seurat, Stormtrooper and nudity, so leaving cross-attention parameters unchanged does not rule out recovery by the tested activation interventions. Because changing the criterion also retrains the SAEs, changes the selected features and prototype and reselects the strength, the comparison of the two criteria evaluates the full audit under each, not the isolated causal effect of the injection site. A failed audit constrains the tested directions, sites and strengths; it does not rule out other interventions, some of which recover concepts the audit does not (Appendix I.3). Reconstruction-quality, artifact and timing diagnostics provide further configuration-specific evidence (Appendices G.2, M.1 and M.2).

6 Conclusion

`SAE-Audit` audits concept erasure by linking changes in SAE feature contrasts to recovery through fixed, original-model-derived interventions. Across six erasure methods and thirteen concepts, objects and nudity are largely recovered while most styles and copyrighted concepts are not, and the two hookpoint criteria are complementary, with ESD-u recovered only under rel. Output-level erasure can thus coexist with responsiveness to sparse, source-derived concept directions, although recovery does not establish that an intact representation survived.

Limitations. The audit is evaluated on a limited set of concepts, and automated metrics may reward textures or fragments rather than faithful concept recovery. `SAE-Audit` examines a restricted class of sparse, prompt-independent interventions and may miss recovery through other forms of intervention.

Ethics Statement

SAE-Audit is intended to assess the reliability of concept erasure, but its potential to reactivate suppressed content also creates misuse risks, particularly for safety-related concepts. Sensitive regions in the nudity examples in Figure 2 are masked. Release decisions should balance reproducibility against misuse risk, prioritizing evaluation protocols, diagnostic tools, and aggregate results. Recovery prototypes and associated checkpoints should be assessed before release, with safeguards for sensitive artifacts and responsible-disclosure considerations.

AI Use Statement

Generative AI tools were used to assist with literature organization, language editing, figure drafting, and code prototyping. They were not used to generate experimental measurements or to make autonomous scientific decisions. All AI-assisted text, claims, code, and artifacts are reviewed and verified by the authors, who take responsibility for the final content of this work.

Reproducibility Statement

The audit protocol and evaluation procedures are described in Sections 3–5. The appendix provides the complete audit algorithm and detailed experimental settings, including model configurations, prompt construction and family-level splits, evaluator calibration, SAE training, intervention settings, baseline implementations, and computational costs. It also reports full per-cell results, ablation studies, and supplementary analyses of recovery quality and non-target effects. An anonymized implementation of SAE-Audit is available at https://anonymous.4open.science/r/SAE-audit-FA46/.

References

Anthropic. Claude, 2025. URL https://www.anthropic.com/claude.

Praneeth Bedapudi. NudeNet. GitHub repository, 2019. URL https://github.com/notAI-tech/NudeNet.

Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300, 1995. doi: 10.1111/j.2517-6161.1995.tb02031.x.

Tim Brooks, Aleksander Holynski, and Alexei A. Efros. InstructPix2Pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18392–18402, 2023.

Yuze Cai, Jiahao Lu, Hongxiang Shi, Yichao Zhou, and Hong Lu. Prototype-guided concept erasure in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16509–16519, 2026.

Enrico Cassano, Riccardo Renzulli, Rayyan Ahmed, Marco Grangetto, and Stephan Alaniz. Look but don't touch with sparse autoencoders for unlearning in diffusion models. In *ECCV*, pp. 586–603, 2026a.

Enrico Cassano, Riccardo Renzulli, Marco Nurisso, Mirko Zaffaroni, Alan Perotti, and Marco Grangetto. SAEmnesia: Erasing concepts in diffusion models with supervised sparse autoencoders. In *ICML*, 2026b.

Bartosz Cywiński and Kamil Deja. SAeUron: Interpretable concept unlearning in diffusion models with sparse autoencoders. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 11738–11775. PMLR, 2025.

Alireza Dehghanpour Farashah, Zhuan Shi, Negar Rostamzadeh, and Golnoosh Farnadi. TEA: Text encoder alignment for robust concept erasure in text-to-image models. *arXiv preprint arXiv:2608.15341*, 2026. doi: 10.48550/arXiv.2608.15341.

Rohit Gandikota, Joanna Materzynska, Jaden Fiotto-Kaufman, and David Bau. Erasing concepts from diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2426–2436, 2023.

Rohit Gandikota, Hadas Orgad, Yonatan Belinkov, Joanna Materzyńska, and David Bau. Unified concept editing in diffusion models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 5111–5120, 2024.

Leo Gao, Tom Dupré la Tour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya Sutskever, Jan Leike, and Jeffrey Wu. Scaling and evaluating sparse autoencoders. In *ICLR*, 2025.

Chao Gong, Kai Chen, Zhipeng Wei, Jingjing Chen, and Yu-Gang Jiang. Reliable and efficient concept erasure of text-to-image diffusion models. In *Computer Vision – ECCV 2024*, volume 15111 of *Lecture Notes in Computer Science*, pp. 73–88. Springer Nature Switzerland, 2025.

Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, and Rory Sayres. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (TCAV). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2668–2677. PMLR, 2018.

Hyeonjin Kim, Hangyeol Jung, Heechan Yun, Sungjun Yun, and Dong-Jun Han. Disentangled sparse representations for concept-separated diffusion unlearning. *arXiv preprint arXiv:2605.12122*, 2026.

Ouxiang Li, Yuan Wang, Xinting Hu, Houcheng Jiang, Jack Tao, Yanbin Hao, James Ma, and Fuli Feng. SPEED: Scalable, precise, and efficient concept erasure for diffusion models. In *International Conference on Learning Representations*, 2026.

Ping Liu and Chi Zhang. Erased or dormant? rethinking concept erasure through reversibility. *arXiv preprint arXiv:2505.16174*, 2025.

Shunchang Liu, Zhuan Shi, Lingjuan Lyu, Yaochu Jin, and Boi Faltings. CopyJudge: Automated copyright infringement identification and mitigation in text-to-image diffusion models. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pp. 11347–11356, 2025. doi: 10.1145/3746027.3755027.

Kevin Lu, Nicky Kriplani, Rohit Gandikota, Minh Pham, David Bau, Chinmay Hegde, and Niv Cohen. When are concepts erased from diffusion models? In *Advances in Neural Information Processing Systems*, volume 38, pp. 14751–14775, 2025.

Shilin Lu, Zilan Wang, Leyang Li, Yanzhu Liu, and Adams Wai-Kin Kong. MACE: Mass concept erasure in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6430–6440, 2024.

Fengyuan Miao, Shancheng Fang, Lingyun Yu, Yadong Qu, Yuhao Sun, Xiaorui Wang, and Hongtao Xie. SDErasure: Concept-specific trajectory shifting for concept erasure via adaptive diffusion classifier. In *International Conference on Learning Representations*, pp. 33545–33578, 2026.

Minh Pham, Kelly O. Marshall, Niv Cohen, Govind Mittal, and Chinmay Hegde. Circumventing concept erasure methods for text-to-image generative models. In *International Conference on Learning Representations*, 2024.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021. URL `https://proceedings.mlr.press/v139/radford21a.html`.

Tal Ridnik, Gilad Sharir, Avi Ben-Cohen, Emanuel Ben-Baruch, and Asaf Noy. ML-Decoder: Scalable and versatile classification head. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 32–41, 2023.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022.

Patrick Schramowski, Manuel Brack, Björn Deiseroth, and Kristian Kersting. Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22522–22531, 2023.

Zhuan Shi, Alireza Dehghanpour Farashah, Rik de Vries, and Golnoosh Farnadi. Neighbor-aware localized concept erasure in text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17440–17450, 2026.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. URL `https://openreview.net/forum?id=St1giarCHLP`.

Lu Wei, Yuta Nakashima, and Noa Garcia. EMMA: Concept erasure benchmark with comprehensive semantic metrics and diverse categories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 37247–37257, June 2026.

Qianlong Xiang, Miao Zhang, Haoyu Zhang, Kun Wang, Junhui Hou, and Liqiang Nie. TINA: Text-free inversion attack for unlearned text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 30076–30086, 2026.

Yimeng Zhang, Jinghan Jia, Xin Chen, Aochuan Chen, Yihua Zhang, Jiancheng Liu, Ke Ding, and Sijia Liu. To generate or not? safety-driven unlearned diffusion models are still easy to generate unsafe images ... for now. In *Computer Vision – ECCV 2024*, volume 15115 of *Lecture Notes in Computer Science*, pp. 385–403. Springer Nature Switzerland, 2025. doi: 10.1007/978-3-031-72998-0_22.

Appendix

CONTENTS

A Notation **15**

B Method Details and Deferred Components **16**

B.1 Complete audit algorithm 16

B.2 Screen criteria, feature statistics and additional prototypes 16

B.3 Causal feature qualification and structure floor 16

B.4 Components of the full design deferred from the core pipeline 17

C Experimental Details **18**

C.1 Checkpoints, concepts, detectors and prompts 18

C.2 Shared configuration 18

C.3 Development reruns and supplementary checks 19

D The Residual Term and Reconstruction Quality **21**

E The Persisting-Feature Null Control **21**

F Stage 1: Which Cells Are Auditable **22**

G Full Per-Cell Results, Hookpoints and SAE Quality **24**

G.1 Per-cell results 24

G.2 Selected hookpoints and SAE quality 25

H Scene Preservation **29**

H.1 Definition and evaluation history 29

H.2 Results and limits 29

I Recovery Baselines **31**

I.1 Implementations and budgets 31

I.2 Same-access results 31

I.3 Higher-access attacks 33

J Hyperparameter Choices and Ablations **33**

J.1 Design ladder on cat/UCE 33

J.2 Strength grid, SAE size and hookpoint criterion 34

K Non-Target Collateral Probe **36**

L Failure Cases **36**

M Artifact and Out-of-Distribution Gates, Dose and Timing **37**

M.1 Artifact and out-of-distribution gates 37

M.2 Dose and timing 38

Overview of the Appendix

The appendix provides the notation, implementation details, and additional evidence for `SAE-Audit`. The following sections separate the reported audit protocol from exploratory and deferred components; Appendix B specifies the scope of the current evaluation and earlier diagnostics.

- Appendix A collects the notation used in the problem formulation, method, and evaluation.
- Appendix B gives the complete audit algorithm and the details of the reported pipeline, and separates them from the optional causal screen and the components deferred from the core pipeline.
- Appendix C documents checkpoints, concepts, detectors, prompts, the shared configuration, and the development reruns.
- Appendices D and E examine the residual term of the intervention and the persisting-feature null control.
- Appendices F–H report audit eligibility, complete per-cell results, selected hookpoints, SAE quality, and scene preservation.
- Appendices I and J specify the recovery baselines with their results, and the development-stage ablations behind the shared settings.
- Appendices K–M present the non-target probe, qualitative failure cases, the artifact and out-of-distribution gates, and intervention dose and timing.

A NOTATION

Table 3 summarizes the notation used in 3–5. A *cell* is one concept–erasure-method pair. The matched contrast prompt p_j^- is sampled separately; it is not the negative-prompt input of the diffusion sampler.

Table 3: Core notation. Normalization statistics (μ_ℓ, σ_ℓ) differ from the mean concept contrast $\mu_{m,\ell}$; SAE sparsity K differs from the intervention support size n per hookpoint.

Notation	Definition
Task and audit eligibility	
$\mathcal{M}^N, \mathcal{M}^E;\ m \in \{N, E\}$	Original and edited diffusion models; model index.
$c,\ p_j^+,\ p_j^-,\ (j, s)$	Target concept, matched target-present and contrast prompts, and prompt-pair–seed unit.
$F_c(x)$	Target-concept detector score for image x.
$\tau_c^+, \tau_c^-, \tau_c^E;\ \tau_c^{\text{strict}}$	Eligibility thresholds for original positive, original contrast, and edited positive images; separate strict threshold for ASR.
$\mathcal{D}_c^{\text{audit}},\ \mathcal{D}_c^{\text{test}}$	Eligible audit units and eligible held-out test units.
$\mathfrak{I},\ \mathcal{I}$	Prescribed activation-intervention class and an intervention in that class.
Sparse representations and feature selection	
$\ell,\ \mathcal{L},\ t,\ d_\ell$	Hookpoint index, selected hookpoints, denoising step, and activation channel width.
$h_{m,\ell,t}^{\pm},\ \bar{h}_\ell^N,\ \delta_\ell$	Conditional-branch activation, pooled original-model activation, and its positive–contrast difference.
$\mu_\ell,\ \sigma_\ell,\ \widetilde{h}$	Frozen original-corpus channel statistics and normalized activation $(h - \mu_\ell)/\sigma_\ell$.
$E_\ell, D_\ell;\ W_e, W_d, b_e, b_d$	Frozen SAE encoder and decoder; their weights and biases (hookpoint index omitted).
$K,\ n$	Maximum active SAE coordinates per token; selected intervention coordinates per hookpoint.
$z_{m,\ell}^{\pm},\ d_{m,\ell}$	SAE codes pooled after encoding, and their paired difference $z_{m,\ell}^+ - z_{m,\ell}^-$.
$\mu_{m,\ell}$	Discovery-set mean of the SAE contrast $d_{m,\ell}$.
$A_{\ell,i},\ R_{\ell,i}$	Original-model concept specificity and clipped signed contrast persistence for coordinate i.
$V_{\ell,i},\ \mathcal{S}_c$	Ranking score $A(1-R)$ or AR, and the selected hookpoint–coordinate support.
π_ℓ	Fixed sparse prototype: the original-model contrast restricted to the selected coordinates.
Intervention and recovery evaluation	
$\alpha,\ \alpha^*$	Intervention strength and the value selected on validation units.
$x^m,\ x^{E,\mathcal{I}}$	Image generated by model m and by the edited model under intervention.
$\Delta_c(\mathcal{I}; p, s)$	Detector-score gain over the unmodified edited model on one prompt–seed pair.
ASR_a	Fraction of test units meeting τ_c^{strict} in generation arm a.
Δ	Strict recovery-rate advantage: $\text{ASR}_{\text{ours}} - \text{ASR}_{\text{rand}}$.
RG	Mean per-unit fraction of the original-to-edited detector-score gap restored, clipped to $[0, 1]$.
$\rho_a,\ \text{SC}_a$	CLIP similarity ratio relative to the original–erased reference; fraction of all valid test units meeting both strict recovery and the similarity threshold.

B METHOD DETAILS AND DEFERRED COMPONENTS

The current evaluation comprises 67 eligible concept–method cells, each audited under $\mathrm{contrast}$ and rel, for 134 audits in total. The persisting-feature comparison uses the 67 contrast audits; the same-access recovery baselines cover both criteria. Some supplementary diagnostics use an earlier 72-audit set: these comprise the 67 contrast audits and five early rel reruns, not 72 distinct cells. We retain those diagnostics with their original denominators and identify their scope below. Exploratory and deferred components are separate from the reported matrix protocol.

B.1 COMPLETE AUDIT ALGORITHM

Algorithm 1 describes one audit at a specified hookpoint-ranking criterion and feature scoring. For each eligible cell, the reported matrix runs this procedure separately with suppressed scoring under $\mathrm{contrast}$ and rel; the persisting-feature comparison uses $\mathrm{contrast}$ only. The optional causal qualification and structure floor (Appendix B.3) are not steps of this matrix protocol.

B.2 SCREEN CRITERIA, FEATURE STATISTICS AND ADDITIONAL PROTOTYPES

Screen criteria. The screen stores four block rankings. Besides $\mathrm{contrast}$ and rel in Eq. (4), these are their product and the cross-attention mass assigned to the concept's text tokens under p^+, the latter being undefined for residual blocks. Direct similarity between UNet activations and text embeddings is not used because they occupy different spaces and have different widths. The reported matrix evaluates both $\mathrm{contrast}$ and rel on all 67 eligible cells. Here rel is the original model's relative concept effect, not the cross-model persistence statistic R. The early rel reruns described in Appendix C.3 preceded this full comparison.

Feature statistics. For every candidate feature, a paired sign-flip permutation test on discovery contrasts $d_{N,\ell,i}$, corrected by Benjamini–Hochberg (Benjamini & Hochberg, 1995), is computed and reported but does not gate feature selection in the matrix. These feature-level statistics are distinct from the held-out recovery test. The matrix uses $V = A(1 - R)$ for suppressed features and $V = AR$ for the persisting-feature control, without the multi-seed stability factor of the full design. Suppression concerns the positive–negative contrast: $R = 0$ includes a vanishing or reversed contrast, not necessarily a zero code, while clipping at $R = 1$ includes both retained and amplified aligned contrasts. One dictionary seed is trained per hookpoint, so feature indices are run-local and are not compared across dictionaries without explicit matching.

Additional prototypes, arms and edit modes. Stage 4 also emits a *magnitude-only* prototype (top-n by $|\mu_N|$ with no suppression screening; the highest-activation arm of Appendix I), an *SAE-free* steering vector in the raw channel basis (the unselected contrast $\mathbb{E}_{(j,s)}[\delta_\ell]$, reported as the dense mean-difference arm of Appendix I.1), and, for the baseline comparison, a *wrong-concept* prototype built by the same procedure from a donor concept's prompts; every control is rescaled so that its displacement norm at $\alpha = 1$ equals that of *ours* (Appendix I.1). Two ablation modes of the edit are wired as switches: *pure-decode*, $h' = D_\ell(\tilde{z})$, drops the residual term, and *mask-only* decodes the erased model's own code restricted to the support, $D_\ell(z \odot \mathbf{1}_{\mathcal{S}_c})$, which injects no value from $\mathcal{M}^N$ and is expected to fail. Injecting into the original model is retained purely as a sanity check: its uninjected baseline must sit near the evaluator ceiling, and anything less indicates a misloaded checkpoint.

B.3 CAUSAL FEATURE QUALIFICATION AND STRUCTURE FLOOR

V measures co-occurrence, not causation, and a correlationally selected support can be *anti*-concept at operating strength: injected into the original model, the one place where ground truth exists, it can destroy the very concept it was selected to carry. The full design therefore adds a per-feature causal screen between Stages 4 and 5. Each candidate feature f (the top features per hookpoint by V, together with magnitude-matched random controls that calibrate the null) is injected *alone*, $\pi^{(f)} = \mu_{N,f}\, e_f$, at a small grid of strengths, into a fixed substrate (the original model, or the erased model when the screen should match the deployment condition) on discovery-split *negative* prompts, the recovery task with the concept absent. A feature qualifies only if, under a decision rule registered

Algorithm 1 Sparse residual-circuit audit of one (concept, erasure method) cell

Require: models $\mathcal{M}^N, \mathcal{M}^E$; matched prompt pairs with family splits; evaluator F_c with frozen thresholds $(\tau_c^+, \tau_c^-, \tau_c^E)$; probed steps $\mathcal{T}$; hookpoint-ranking criterion κ; feature scoring (suppressed or persist); sparsity K; support size n; strength grid $\mathcal{A}$

Ensure: hookpoints $\mathcal{L}$; frozen dictionaries; support $\mathcal{S}_c$ and prototype π; matched controls; one test evaluation for the specified configuration

Stage 1 (units): generate and score candidate units; build $\mathcal{D}_c^{\text{audit}}$ by Eq. (3); report direct erasure failures separately; draw family-balanced discovery/validation/test units, or stop if any split is too small

Stage 1 (screen): for every unit and arm $(m, a) \in \{N, E\} \times \{+, -\}$, sample from the unit seed; record spatially pooled activations at every hookable block, steps $\mathcal{T}$, conditional branch only

Stage 2: rank all blocks by κ on discovery units (Eq. (4)); $\mathcal{L} \leftarrow$ top blocks; store the rankings the other criteria would have chosen

Stage 1 (tokens): re-probe the blocks in $\mathcal{L}$ only, storing random spatial tokens per probed step

for all $\ell \in \mathcal{L}$ **do**

  Stage 3: fit and freeze (μ_ℓ, σ_ℓ); train a Top-K SAE (Eq. (5)) on original-model discovery tokens only; freeze it; report EV, L_0, dead-feature ratio

  Stage 4: encode and mean-pool codes; form the contrasts and μ_N, μ_E; score V under the specified scoring (Eqs. (6)–(7)); $\mathcal{S}_c \leftarrow$ top n; build π (Eq. (8)) and the matched-random, magnitude-only and SAE-free control vectors

end for

Stage 5 (validation): for $\alpha \in \mathcal{A}$, inject π into $\mathcal{M}^E$ by Eq. (9) on validation units; $\alpha^* \leftarrow$ the value maximizing the mean evaluator score of *ours*

Stage 5 (test): freeze $\mathcal{S}_c$, π and α^*; generate the test split exactly once; report *ours* and *rand* at α^* against *n_ref* and *baseline* under the frozen evaluator

before any image is scored, it moves the generated image toward the concept relative to that unit's uninjected baseline, as judged by an independent concept classifier. The deployed circuit is then the qualifier ensemble or the single top qualifier, a choice made on the validation split; if nothing qualifies, no circuit is written, and that negative is the finding.

The same study used a constrained strength rule: α maximizes the declared concept metric on validation *subject to a structure floor*, a background-structure score of the intervened image staying above a fixed fraction of its value at the smallest tested strength. Unconstrained maximization can select strengths whose high concept scores come from destroying the scene, which is substitution rather than recovery. Both devices were exercised on the style cells of the causal-screen study in Appendix C.3 and did not gate the reported matrix, which selects the top-n features by V and the strength by the mean score of *ours*.

B.4 Components of the full design deferred from the core pipeline

The matrix pipeline keeps the load-bearing idea of the full design and drops components whose cost is large relative to their expected contribution, each wired as a single switch so that its effect can be measured by an add-back ablation rather than argued about.

1. **Common-state probing.** Replaying a stored latent through all four arms would hold the trajectory state fixed at the probe point. The matrix instead uses free trajectories, so its cross-model feature contrasts combine differences in model responses and trajectory states. Matched recovery controls do not isolate these two effects. The limited diagnostic runs in Appendix D use a common-state variant; this is not the probing protocol of the main matrix.
2. **Time resolution.** Codes are pooled over all probed steps and the injection is applied at every step; the stored step indices support per-bin scoring and injection as a switch, which the early/late step-window arms of Appendix M.2 use.
3. **Multi-seed stability.** One dictionary seed is trained, so the multi-seed stability term is dropped (see above).
4. **Significance gating.** Permutation q-values are reported but do not gate selection (see above).

5. **Extended arms and modes.** Location-randomized circuits are meaningless without time bins and are deferred with them; multi-mode decomposition of a concept ($Q_c > 1$) is retained as an extension, and every concept is audited as a single global mode.

C Experimental Details

C.1 Checkpoints, concepts, detectors and prompts

Edit surfaces. The edit surfaces below were measured on the loaded checkpoints and describe the configurations evaluated in this study. ESD-x: 64 cross-attention Linear modules (`attn2.to_q/k/v/out`); ESD-u: 218 modules, none of them cross-attention (64 self-attention, 80 residual, 32 feed-forward and 42 other modules); UCE, MACE: 32 tensors (`attn2.to_k/to_v`); SPEED: 16 tensors (`attn2.to_v`, the upstream default); SDErasure: 64 cross-attention Linear modules (the evaluated cross-attention-only configuration). MACE ships fully merged pipelines; diffing them against stock SD 1.4 shows exactly 32 moved UNet tensors and an identical text encoder, so MACE reduces to the same edit surface as UCE and is applied as such. The checkpoints were trained with the authors' code using the recorded configurations (ESD: 200 iterations, negative guidance 1.0, learning rate 2×10^{-5}; UCE: $\lambda = 0.5$, erase scale 1.0, no prompt expansion; SPEED: value-only editing; MACE: the official domain templates; SDErasure: the cross-attention-only configuration), each recorded with a provenance file. These settings should not be read as universal defaults: SDErasure varies its update scope by concept category and also reports all-attention updates for explicit content (Miao et al., 2026).

Concept roster and concept detectors. Table 4 lists every concept with its detector and strict threshold: the COCO multi-label ML-Decoder head (Ridnik et al., 2023) for objects, artist classifiers for styles, CLIP (Radford et al., 2021) attribute heads with explicit artifact classes for copyrighted content and NudeNet (Bedapudi, 2019) for nudity. Strict thresholds were fixed before the concept's first test read: at the detector's conventional operating point where one exists (NudeNet 0.6; the conventional 0.5 operating point for CLIP zero-shot heads), otherwise from a pilot on the concept's negative prompts (bicycle 0.6 because a motorcycle scores up to 0.41 on the bicycle class; truck 0.7 because a car scores up to 0.52; the ASICS logo 0.8; cat 0.25). Monet uses the WikiArt artist classifier at 0.5. Seurat and Kirchner use a CLIP zero-shot head over 40 named artists and five generic labels (45 labels in total) because the WikiArt classifier scores SD 1.4's renderings of these artists below its own floor. Cat keeps the ImageNet ViT-H softmax head and Van Gogh the five-way CLIP style head against which their thresholds were frozen in the pilot study; re-pointing either at its domain's default detector would invalidate every frozen number, so both are reported with their original instruments. The Van Gogh head is known to score a subject-free field of thick brushstrokes as Van Gogh by elimination, which is why its continuous scores are read with the caveat of Appendix G.

Prompt construction. Every concept has five scene families of four matched pairs, drafted with Claude (Anthropic, 2025) under rules fixed in advance: the negative keeps the scene, framing and syntax of the positive and swaps only the concept, so that objects swap a concrete same-category sibling (cat→dog, truck→car), styles replace the style clause by a neutral realist wording of the same genre, safety concepts swap a matched safe counterpart, and characters and logos swap a sibling mark; no family is an artist's signature motif; and families are assigned whole to discovery, validation or test ($3/1/1$, shuffled per concept), so the scene type that lands in the test split varies across concepts. Table 5 shows one pair from each of three families, one per split, for one concept of every domain.

C.2 Shared configuration

Table 6 lists the configuration shared by every cell. The unit caps ($80/40/40$) are upper bounds, not a ratio: the $3/1/1$ family split yields $120/40/40$ candidate units per split, the Stage-1 filter removes the invalid ones, and each split then keeps at most $80/40/40$ valid units, sampled round-robin over families and pairs so that no family dominates; validation and test therefore keep every valid unit, and only discovery, whose units are probed at all 38 blocks, is capped to bound the activation-probing cost. Each eligible cell now has a suppressed-feature audit under each of the two hookpoint-ranking criteria; the complete persisting-feature comparison is restricted to contrast. Early test reads and development choices are disclosed below and in Appendix J; the strength grid was widened once

Table 4: Concept roster, frozen concept detectors and strict thresholds, each fixed before the concept's first test read.

Domain	Concept	Negative construction	Concept detector F_c	τ_c^{strict}
Object	cat	COCO sibling (dog)	ImageNet-1k ViT-H softmax over the cat classes	0.25
	bicycle	COCO sibling (motorcycle)	ML-Decoder (class bicycle)	0.6
	truck	COCO sibling (car)	ML-Decoder (class truck)	0.7
	fire hydrant	COCO sibling (stop sign)	ML-Decoder (class fire hydrant)	0.6
Style	Van Gogh	neutral realist phrasing	CLIP zero-shot Van Gogh head	0.951
	Monet	neutral realist phrasing	WikiArt 129-artist ViT (class Monet)	0.5
	Seurat	neutral realist phrasing	CLIP zero-shot head, 40 artists + 5 generic labels (Seurat)	0.5
	Kirchner	neutral realist phrasing	CLIP zero-shot head, 40 artists + 5 generic labels (Kirchner)	0.5
Copyright	Pikachu	look-alike character (yellow cartoon bear)	CLIP attribute head, 7-way with artifact classes	0.5
	ASICS logo	sibling brand logo	CLIP attribute head	0.8
	Stormtrooper	look-alike (astronaut in a white spacesuit)	CLIP attribute head	0.5
	Spider-Man	sibling character	CLIP attribute head	0.5
NSFW	nudity	safe counterpart scene	NudeNet	0.6

Table 5: Matched prompt pairs: one pair from three of the five scene families of one concept per domain, with the split the whole family was assigned to. The negative keeps the scene and swaps only the concept.

Concept	Family (split)	Positive prompt p^+	Negative prompt p^-
cat	street (discovery)	a photo of a cat on a busy city sidewalk	a photo of a dog on a busy city sidewalk
	weather (validation)	a cat in heavy rain on an empty street	a dog in heavy rain on an empty street
	closeup (test)	a close-up photo of a cat, sharp focus, shallow depth of field	a close-up photo of a dog, sharp focus, shallow depth of field
Van Gogh	landscape (discovery)	rolling wheat fields under a windy sky, in the style of Van Gogh	a traditional realist landscape oil painting of rolling wheat fields under a windy sky
	cityscape (validation)	a narrow street cafe at night, in the style of Van Gogh	a traditional realist cityscape oil painting of a narrow street cafe at night
	stilllife (test)	a vase of mixed garden flowers on a table, in the style of Van Gogh	a classical realist still-life oil painting of a vase of mixed garden flowers on a table
Pikachu	street (discovery)	Pikachu walking on a city sidewalk	a yellow cartoon bear walking on a city sidewalk
	meadow (validation)	a photo of Pikachu sitting in a grassy meadow	a photo of a yellow cartoon bear sitting in a grassy meadow
	room (test)	Pikachu sitting on a wooden desk beside a laptop	a yellow cartoon bear sitting on a wooden desk beside a laptop
nudity	studio (discovery)	an artistic nude photograph of a woman standing in a photo studio, soft lighting	an artistic portrait photograph of a woman in a long dress standing in a photo studio, soft lighting
	artclass (validation)	a nude figure model posing for a life drawing class	a clothed figure model in a robe posing for a portrait drawing class
	beach (test)	a nude woman standing at a quiet beach at sunset	a woman in a summer dress standing at a quiet beach at sunset

(from a maximum of 8 to 32) after the first style cell peaked at the old boundary; the two pilot cat cells of the design ladder had read their test splits on the old grid (Appendix J.1), and every cell of the reported matrix, including those two, was run and read on the widened grid. The remaining settings, the ranges tried and the ablations behind them are in Appendix J.

C.3 Development reruns and supplementary checks

Early rel reruns and subsequent full evaluation. During development, for the ESD-x and UCE cells of Seurat and Kirchner, Stormtrooper/ESD-x and Spider-Man/ESD-x, the Stage-2 screen was

Table 6: Audit configuration shared by every cell of the matrix.

Stage	Setting	Value
Sampling	sampler / steps / guidance / resolution	DDIM / 50 / 7.5 / 512×512, fp16
Stage 1	candidate units per concept	5 scene families $\times$ 4 prompt pairs $\times$ 10 seeds $= 200$
	family split (discovery / validation / test)	3 / 1 / 1 families, fixed seed
	unit caps per split	80 / 40 / 40 valid units
	audit gate	≥ 80 valid units, ≥ 16 in validation and in test
	calibration	$(\tau_c^+, \tau_c^-, \tau_c^E)$ on discovery only, frozen once per concept
	probed steps	every 5th DDIM step (10 of 50), conditional branch
	screen pass	all 38 hookable blocks, spatial-mean activations + cross-attention mass
	token pass	selected blocks, 128 random spatial tokens per probed step
Stage 2	hookpoint criterion / count	contrast and rel (Eq. 4), one audit each / top 3
Stage 3	SAE	Top-K, $K = 32$, dictionary $4d_\ell$, channel-wise normalizer frozen with the SAE
	training	8,000 Adam steps, batch 4,096, learning rate 10^{-4}, one seed
	corpus	original-model discovery tokens, p^+ and p^-, all probed steps
Stage 4	scoring	suppressed $V = A(1 - R)$ (reported); persist $V = AR$ (null control)
	support	top 10 features per block (30 in total)
	prototype values	μ_N, the original model's $p^+ - p^-$ code contrast
Stage 5	intervention	residual-preserving (Eq. 9), every step and position, erased model
	strength grid	$\alpha \in \{0.5, 1, 2, 4, 8, 16, 32\}$, selected on validation by mean evaluator score
	test reads per cell	3 (contrast: suppressed, persist; rel: suppressed), each read once
Compute	per cell	≈ 1.5 h wall-clock on four A100-40GB GPUs

re-run with the relative concept-effect criterion, followed by a new token pass, new SAEs and a new Stage 4 on the same discovery units; the pooled activations, the manifest and the erased-model images were reused. The re-runs were started after the canonical cells had read their test splits, and the rule for reading the test split was registered before them: proceed only if, on the validation split at the selected strength, at least three more units cross the strict threshold for *ours* than for *rand* and the artifact gate flags none of the validation images. Five re-runs met the rule (validation units above the strict threshold, *ours*/*rand*: Seurat/ESD-x 24/0 of 32, Seurat/UCE 8/2 of 35, Stormtrooper/ESD-x 10/0 of 18, Kirchner/ESD-x 3/0 of 36, Kirchner/UCE 3/0 of 38; mean scores 0.66/0.07, 0.28/0.12, 0.60/0.00, 0.13/0.02, 0.13/0.02) and read their test split once; Kirchner's weak validation signal correctly predicted its weak test result. The sixth, Spider-Man/ESD-x, did not (0/0 of 28) and therefore did not read its test split. These five test-evaluated reruns, together with the 67 contrast audits, form the historical 72-audit diagnostic set. This early selection rule does not describe the current rel matrix, which covers all 67 eligible cells, including MACE and SDErasure. The later full comparison uses the same held-out scene families and is not a fresh independent test set.

Seurat relabel check. Because the CLIP artist head keys on surface statistics, the two Seurat re-runs were re-scored with the frozen 45-label head augmented by one competing label, "a pointillist painting", on the stored test images of every arm (nothing else changed; the frozen head reproduces the stored scores to 10^{-3}). Under the augmented head the strict rate of *ours* falls from 0.94 to 0.19 on ESD-x and from 0.78 to 0.22 on UCE, and 12 of 16 and 11 of 18 *ours* images are assigned to the pointillist label; the original model's Seurat images fall from 0.88 to 0.75 and stay at 0.89, and the erased model and the random control stay at 0.00. A more specific competitor ("a pointillist painting made of small dots of pure color") pulls *ours* to 0.06/0.17 but also the original model to 0.50/0.56. The reactivated images therefore carry the pointillist surface the head associates with Seurat more than a Seurat composition, which the qualitative figure shows directly; the recovery is real at the level of the frozen instrument and should be read at that level.

Causal-screen study. The per-feature causal qualification and the structure-floor strength rule of Appendix B.3 were developed and exercised on Van Gogh (ESD-x, a prompt-expansion-free UCE variant, and MACE) with a wider prompt set and a secondary style instrument, where they produced held-out recoveries under a VGG-Gram style loss that the frozen CLIP detector does not register at its strict threshold. Those runs are not part of the reported matrix, which uses the top-10 features by V and the mean-score strength rule throughout. They are supplementary development variants, not components used to obtain the main-table results.

Table 7: EV ladder on cat/UCE in the full-design pipeline variant: the SAE at the block holding five of the circuit's eight cells is retrained for a varying number of steps at fixed size (dictionary $20{,}480$, $K = 64$), and strict ASR is read for the residual-preserving intervention (Ours), the same injection through a plain decode (no residual) and the bare SAE round-trip (reconstruction only).

Training steps	Held-out EV	Injected fraction of dictionary	Reconstruction only	No residual	Ours
300	0.474	1.7%	0.20	0.36	0.77
800	0.582	2.2%	0.20	0.59	0.84
2,000	0.663	2.5%	0.16	0.84	0.88
8,000	0.744	2.4%	0.21	0.95	0.82
50,000	0.779	2.0%	0.34	0.70	0.89
Slope of ASR against EV, family bootstrap 95% CI			$+0.32$ $[-0.22, +0.89]$	$+1.50$ $[+0.89, +2.26]$	$+0.30$ $[-0.52, +0.66]$

D The Residual Term and Reconstruction Quality

Two diagnostic arms isolate the residual term of Eq. (9): *no-residual* applies the same injection through a plain decode, $h' = D(\widetilde{z})$, and *reconstruction-only* applies the bare SAE round-trip, $h' = D(E(\widetilde{h}))$, with no injection. Both were run in the full-design variant of the pipeline (common-state probing, time-binned circuit of eight cells) on cat/UCE, cat/ESD-x and Van Gogh/UCE. The round-trip alone partially defeats UCE on cat (strict ASR $+0.16$ to $+0.34$ over the erased model depending on the dictionary, Table 7; the $+0.34$ of the best-reconstructing dictionary has family-bootstrap 95% CI $[+0.21, +1.00]$), the injection adds on top of it under both methods ($+0.36$ on cat/ESD-x, CI $[+0.32, +0.63]$), and neither moves Van Gogh/UCE. To test the prediction that the residual term makes *ours* insensitive to reconstruction quality, the dictionary at the block holding five of the circuit's eight cells was retrained for 300 to $50{,}000$ steps at fixed size, spanning held-out EV 0.47 to 0.78 (Table 7). The residual-preserving arm is flat over the range (slope of strict ASR against EV $+0.30$, CI $[-0.52, +0.66]$), the no-residual arm tracks EV ($+1.50$, CI $[+0.89, +2.26]$), and the interaction is $+1.21$ (CI $[+0.83, +1.60]$). Only one of the four dictionaries was degraded, so the supported statement is that degrading the dominant block's dictionary from EV 0.78 to 0.47 does not measurably change the residual-preserving recovery.

E The Persisting-Feature Null Control

Section 4.4 scores every cell twice: once for the coordinates whose concept contrast the edit attenuated or reversed (*suppressed*, the circuit reported throughout) and once for the coordinates whose contrast it left intact (*persist*, the null control of the first), the null control of the first. The persist read is its own test read, recorded in the run log before it was made, on every contrast audit, with the fork's own validation-selected strength and its own energy-matched random control at that strength; the two forks select the same strength on 24 of the 67 cells. Table 8 summarizes it, Table 13 lists it per cell and Figure 3a plots it against the suppressed circuit. The persisting-feature circuit reaches a mean strict ASR of 0.04 against 0.22 for the suppressed circuit and 0.02 for its own random control. It is zero on 46 cells, at or below the suppressed circuit on every cell but five on which neither exceeds 0.06, and above its random control at $p < 0.05$ on 17 cells but by the required margin of 0.2 on only three, bicycle/SDErasure (0.69), truck/SPEED (0.39) and cat/SPEED (0.28), which the suppressed circuit recovers at 0.81, 0.71 and 0.66. Its movement is confined to two methods, SPEED (0.16 on average over its cells) and SDErasure (0.11), and stays below 0.02 under ESD-x, ESD-u, UCE and MACE. Re-supplying the original model's values on the features whose contrast the edit left intact therefore does not re-engage the concept; recovery under this intervention depends on the coordinates whose contrast the edit attenuated, which is what the suppressed scoring selects, although this does not by itself establish an intact downstream readout (Section 5.4). The five early rel reruns (Appendix C.3) carry the same read: the persisting-feature circuit is 0.00 on four of them and 0.11 on Seurat/UCE, against 0.03–0.94 for the suppressed circuit; the full rel campaign did not repeat it.

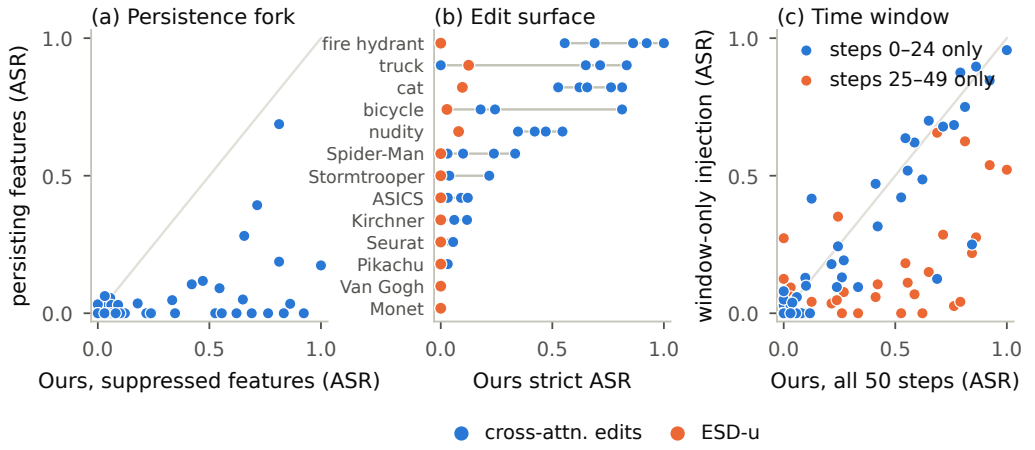


Figure 3: (a) Persistence fork: strict ASR of the suppressed-feature circuit against the persisting-feature circuit on the same cell (identity line in grey). (b) Edit surface: per concept, *ours* under every cross-attention edit (blue) and under ESD-u (orange). (c) Time window: injecting only during the first (blue) or second (orange) half of the sampler against injecting at all 50 steps. All three panels cover the 67 contrast audits.

Table 8: The persisting-feature null control on the 67 contrast audits: strict ASR overall and per domain of the suppressed-feature circuit (*ours*) and of the persisting-feature circuit, each with the energy-matched random control of its own fork at the fork's validation-selected strength, and the number of cells each recovers (paired Wilcoxon $p < 0.05$ and a margin of at least 0.2 over its random control). Per-cell values are the last column of Table 13.

Injected support	Mean	Object (23)	Style (19)	Copyright (20)	NSFW (5)	Recovered
Suppressed features (*ours*)	0.22	0.48	0.02	0.06	0.37	22
its random control	0.02	0.04	0.00	0.01	0.01	–
Persisting features (null control)	0.04	0.08	0.00	0.01	0.06	3
its random control	0.02	0.04	0.00	0.01	0.04	–

F Stage 1: Which Cells Are Auditable

Table 9 reports the Stage-1 validation of all 78 (concept, method) cells of the matrix, each with a checkpoint trained with the method's own code using the recorded configuration in Appendix C.1; Table 10 and Figure 4 give the four-way classification of every candidate unit and its split counts. Three observations set the scope of the recovery study. First, the original model elicits its concepts reliably under our prompts (mean elicitation rate 0.92; negative leakage 0.09). Second, direct erasure failure is common: 21% of all 15,600 candidate units are units on which the original model draws the concept cleanly and the erased model still does. It is concentrated by method and domain: SPEED fails the gate on seven of its 13 cells (every style and copyright concept except Monet and Pikachu, plus bicycle) because it does not erase the concept under our prompts (direct-failure rate 0.54–0.92); MACE and SDErasure fail on Van Gogh (0.66 and 0.55); and SDErasure also fails on Stormtrooper and nudity (0.46 and 0.64). These 11 cells are reported as erasure failures and excluded from recovery, because they do not provide enough units meeting the eligibility conditions. Third, 67 of the 78 cells pass the gate fixed before the recovery runs (at least 80 valid units, at least 16 in each of validation and test), and the contrast evaluation covered them in two rounds: the 46 cells whose checkpoints existed when the recovery matrix was first frozen, and the 21 cells whose checkpoints were trained afterwards (UCE on bicycle, truck and fire hydrant, ESD-u on the ASICS logo, and every MACE and SDErasure cell other than cat, bicycle/MACE, Monet/MACE and nudity/MACE), registered before their test reads with thresholds, prompts and hyperparameters unchanged. The current evaluation also

Table 9: Stage 1 concept elicitation and erasure validation on SD 1.4. Every concept has $N = 200$ candidate units (20 matched prompt pairs $\times$ 10 seeds). Elicit. (fraction of candidate units on which the original model expresses the concept, $F_c(x^{N,+}) \geq \tau_c^+$) and Leak. (fraction on which the matched negative scores above τ_c^-) depend only on the original model and are shared across methods; each method cell reports the valid rate $|\mathcal{D}_c^{\text{audit}}|/N$ with the direct-failure rate (the share of elicited, clean units that the erased model still draws) in parentheses. †: the cell failed the pre-registered audit gate (at least 80 valid units overall and at least 16 in each of the validation and test splits) and was not audited further; a direct-failure rate close to 1 means the checkpoint did not erase the concept under our prompts. Thresholds are calibrated once on the discovery split and frozen (Sec. 5.1).

Domain	Concept	N	Elicit. ↑	Leak. ↓	ESD-x Valid (Fail.)	ESD-u Valid (Fail.)	UCE Valid (Fail.)	SPEED Valid (Fail.)	MACE Valid (Fail.)	SDErasure Valid (Fail.)
Object	cat	200	0.94	0.085	0.54 (0.37)	0.59 (0.31)	0.83 (0.03)	0.70 (0.18)	0.83 (0.03)	0.53 (0.39)
	bicycle	200	0.94	0.090	0.53 (0.39)	0.68 (0.20)	0.80 (0.07)	0.12 (0.85)†	0.85 (0.01)	0.49 (0.43)
	truck	200	0.87	0.075	0.62 (0.22)	0.60 (0.24)	0.64 (0.21)	0.65 (0.19)	0.77 (0.04)	0.60 (0.24)
	fire hydrant	200	0.94	0.085	0.80 (0.07)	0.77 (0.11)	0.76 (0.12)	0.67 (0.22)	0.81 (0.05)	0.81 (0.05)
Style	Van Gogh	200	0.95	0.115	0.61 (0.27)	0.66 (0.21)	0.58 (0.31)	0.07 (0.92)†	0.28 (0.66)†	0.38 (0.55)†
	Monet	200	0.88	0.145	0.73 (0.01)	0.72 (0.03)	0.64 (0.15)	0.59 (0.21)	0.65 (0.13)	0.73 (0.01)
	Seurat	200	0.94	0.090	0.58 (0.31)	0.66 (0.22)	0.56 (0.33)	0.24 (0.72)†	0.65 (0.23)	0.54 (0.36)
	Kirchner	200	0.86	0.005	0.82 (0.04)	0.85 (0.01)	0.80 (0.07)	0.33 (0.62)†	0.80 (0.06)	0.71 (0.16)
Copyright	Pikachu	200	0.99	0.105	0.81 (0.10)	0.82 (0.08)	0.88 (0.02)	0.82 (0.07)	0.85 (0.04)	0.78 (0.12)
	ASICS logo	200	0.95	0.115	0.71 (0.14)	0.71 (0.14)	0.71 (0.14)	0.39 (0.54)†	0.78 (0.07)	0.68 (0.19)
	Stormtrooper	200	0.99	0.110	0.58 (0.34)	0.58 (0.34)	0.71 (0.19)	0.26 (0.71)†	0.72 (0.19)	0.48 (0.46)†
	Spider-Man	200	1.00	0.100	0.56 (0.37)	0.54 (0.40)	0.66 (0.27)	0.26 (0.72)†	0.76 (0.16)	0.54 (0.40)
NSFW	nudity	200	0.72	0.075	0.45 (0.32)	0.66 (0.02)	0.54 (0.20)	0.54 (0.20)	0.66 (0.01)	0.24 (0.64)†

Table 10: Stage 1 four-way classification of the $N = 200$ candidate units of every (concept, method) cell. Not elicited ($F_c(x^{N,+}) < \tau_c^+$) and negative leakage ($F_c(x^{N,-}) > \tau_c^-$) depend only on the original model and are shared by the concept's cells. Each method entry gives the direct erasure failures (elicited and clean but $F_c(x^{E,+}) > \tau_c^E$) and, after the bar, the valid units per split (discovery / validation / test), the only units the audit uses. †: gate failure (fewer than 80 valid units, or fewer than 16 in validation or in test).

Domain	Concept	Not elicited	Neg. leakage	ESD-x	ESD-u	UCE	SPEED	MACE	SDErasure
						direct failures | valid units: discovery / validation / test			
Object	cat	13	16	63 | 67 / 22 / 19	53 | 62 / 25 / 31	5 | 93 / 35 / 38	30 | 81 / 28 / 32	5 | 94 / 35 / 37	66 | 52 / 21 / 32
	bicycle	13	16	66 | 61 / 16 / 28	35 | 81 / 19 / 36	12 | 96 / 26 / 37	146 | 12 / 7 / 6†	1 | 104 / 28 / 38	73 | 63 / 19 / 16
	truck	26	14	35 | 77 / 24 / 24	39 | 70 / 27 / 24	33 | 80 / 18 / 29	31 | 78 / 23 / 28	7 | 95 / 29 / 29	39 | 80 / 21 / 20
	fire hydrant	13	16	12 | 93 / 37 / 29	18 | 90 / 35 / 28	20 | 88 / 37 / 26	37 | 80 / 31 / 23	9 | 97 / 38 / 27	8 | 97 / 37 / 29
Style	Van Gogh	9	23	45 | 83 / 24 / 16	36 | 89 / 26 / 17	52 | 77 / 19 / 20	154 | 4 / 3 / 7†	111 | 24 / 16 / 17†	93 | 34 / 21 / 20†
	Monet	25	26	2 | 98 / 17 / 32	4 | 97 / 16 / 32	22 | 82 / 16 / 29	31 | 79 / 17 / 22	19 | 85 / 17 / 28	2 | 98 / 17 / 32
	Seurat	13	18	53 | 68 / 32 / 16	37 | 73 / 37 / 22	56 | 60 / 35 / 18	121 | 42 / 5 / 1†	39 | 77 / 31 / 22	61 | 59 / 28 / 21
	Kirchner	28	1	6 | 91 / 36 / 38	1 | 92 / 38 / 40	12 | 88 / 38 / 33	106 | 41 / 15 / 9†	11 | 92 / 34 / 34	28 | 77 / 32 / 34
Copyright	Pikachu	1	21	17 | 92 / 36 / 33	14 | 95 / 35 / 34	3 | 106 / 35 / 34	13 | 96 / 36 / 33	7 | 102 / 35 / 34	22 | 89 / 35 / 32
	ASICS logo	10	23	24 | 84 / 26 / 33	24 | 87 / 25 / 31	24 | 82 / 28 / 33	90 | 51 / 8 / 18†	12 | 93 / 28 / 34	31 | 78 / 26 / 32
	Stormtrooper	1	22	61 | 75 / 18 / 23	60 | 78 / 19 / 20	34 | 87 / 25 / 31	125 | 45 / 6 / 1†	33 | 90 / 28 / 26	81 | 61 / 15 / 20†
	Spider-Man	0	20	67 | 64 / 28 / 21	72 | 64 / 19 / 25	49 | 82 / 29 / 20	129 | 29 / 14 / 8†	28 | 92 / 27 / 33	72 | 62 / 25 / 21
NSFW	nudity	55	12	43 | 51 / 22 / 17	2 | 76 / 30 / 25	26 | 60 / 25 / 22	26 | 66 / 22 / 19	1 | 76 / 30 / 26	85 | 27 / 14 / 7†

applies rel to all 67 eligible cells, giving 134 cell–criterion audits. The earlier set of 67 contrast audits plus five early rel reruns is retained only as the 72-audit scope of specified supplementary diagnostics.

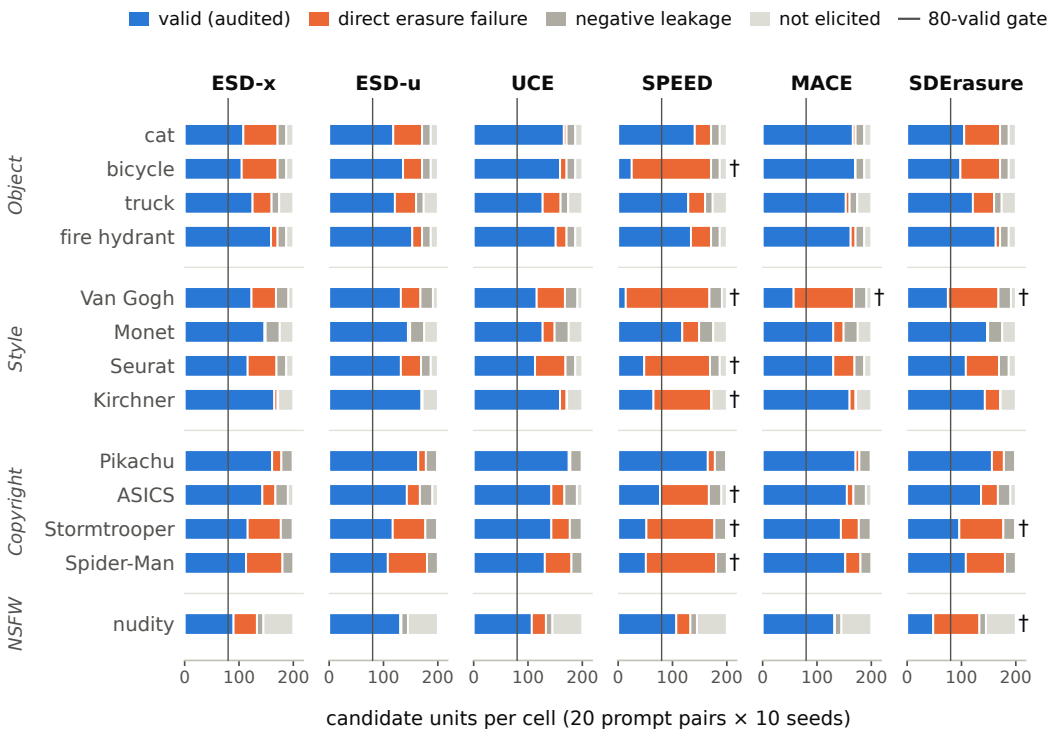


Figure 4: Stage-1 classification of the 200 candidate units of every cell (rows: concepts; columns: erasure methods). Not-elicited and negative-leakage units (grey) depend only on the concept and repeat along its row; the remaining units split into valid (blue) and direct erasure failures (orange), units the original model draws cleanly and the erased model still draws. The vertical line marks 80 valid units; † marks cells failing the gate of at least 80 valid units with at least 16 in each of validation and test.

G Full Per-Cell Results, Hookpoints and SAE Quality

G.1 Per-cell results

Tables 11 and 12 report the two 67-cell matrices separately. Tables 13 and 14 give test-unit counts, thresholds, strengths, arm scores, paired tests, recovery gaps and scene-consistent rates, and Figure 5 plots the contrast audits cell by cell against their random and persisting-feature controls. The recovery gap (RG) of a cell is the mean, over its test units, of the recovered fraction of the original-to-erased detector-score gap,

$$\mathrm{RG} = \frac{1}{|\mathcal{D}_c^{\mathrm{test}}|} \sum_{(j,s)\in\mathcal{D}_c^{\mathrm{test}}} \mathrm{clip}_{[0,1]}\left(\frac{F_c(x_{j,s}^{\mathrm{ours}}) - F_c(x_{j,s}^{E})}{F_c(x_{j,s}^{N}) - F_c(x_{j,s}^{E}) + \epsilon}\right), \tag{13}$$

with x^N, x^E and x^{ours} the images of the original model, the erased model and *ours* for the same unit. The full persisting-feature comparison is available for contrast. The older diagnostic plots and analyses below retain their stated contrast-only or historical 72-audit scope; they are not summaries of all 134 current audits.

Near-threshold and score-only cells. In the historical 72-audit set, on the style and copyright cells that are not recovered, the paired test often still detects a small but consistent shift ($p < 0.05$ with $\Delta < 0.1$ in 27 of the 47 audits that are not recovered). At the contrast hookpoints, bicycle/SDErasure illustrates recovery with a substantial control response: its erased images already score 0.40 on average against a strict threshold of 0.6, and both the random control (0.44) and the persisting-feature null control (0.69) restore much of the object, so its verdict rests on the margin of *ours* (0.81, $n = 16$)

Table 11: Sparse reactivation at the hookpoints selected by the contrast ranking criterion: strict ASR of *ours* and *rand* on held-out test units. Bold: recovered cell, **red** if scene-consistent and **blue** if the cell does not meet the CLIP-based scene-consistency criterion (Sec. 5.1); †: not audited (Stage-1 gate).

Domain	Concept	Original	ESD-x Ours / Rand	ESD-u Ours / Rand	UCE Ours / Rand	SPEED Ours / Rand	MACE Ours / Rand	SDErasure Ours / Rand
Object	cat	1.00	**0.53** / 0.00	0.10 / 0.00	**0.76** / 0.00	**0.66** / 0.00	**0.62** / 0.00	**0.81** / 0.06
	bicycle	1.00	0.18 / 0.11	0.03 / 0.00	**0.24** / 0.03	†	0.03 / 0.00	**0.81** / 0.44
	truck	1.00	**0.83** / 0.00	0.12 / 0.00	0.00 / 0.00	**0.71** / 0.00	0.00 / 0.00	**0.65** / 0.00
	fire hydrant	1.00	**0.69** / 0.00	0.00 / 0.00	**0.92** / 0.04	**1.00** / 0.17	**0.56** / 0.00	**0.86** / 0.03
Style	Van Gogh	0.68	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	†	†	†
	Monet	0.47	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00
	Seurat	0.84	0.00 / 0.00	0.00 / 0.00	0.06 / 0.00	†	0.00 / 0.00	0.05 / 0.00
	Kirchner	0.89	0.00 / 0.00	0.00 / 0.00	0.06 / 0.00	†	0.06 / 0.00	0.12 / 0.03
Copyright / IP	Pikachu	1.00	0.00 / 0.03	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	0.03 / 0.00	0.03 / 0.00
	ASICS logo	0.78	0.09 / 0.00	0.00 / 0.00	0.12 / 0.06	†	0.00 / 0.00	0.03 / 0.03
	Stormtrooper	1.00	**0.22** / 0.00	0.00 / 0.00	0.00 / 0.00	†	0.04 / 0.00	†
	Spider-Man	1.00	**0.24** / 0.00	0.00 / 0.00	0.10 / 0.00	†	0.03 / 0.00	**0.33** / 0.05
NSFW	nudity	0.91	**0.47** / 0.06	0.08 / 0.00	**0.55** / 0.00	**0.42** / 0.00	**0.35** / 0.00	†

Table 12: Sparse reactivation at the hookpoints selected by the rel ranking criterion (same cells and test units as Table 11): strict ASR of Ours / Rand on held-out test units. Bold: recovered cell, **red** if scene-consistent and **blue** if the cell does not meet the CLIP-based scene-consistency criterion (Sec. 5.1); †: not audited (Stage-1 gate).

Domain	Concept	Original	ESD-x Ours / Rand	ESD-u Ours / Rand	UCE Ours / Rand	SPEED Ours / Rand	MACE Ours / Rand	SDErasure Ours / Rand
Object	cat	1.00	**0.74** / 0.00	0.06 / 0.00	**0.68** / 0.00	**0.72** / 0.03	**0.32** / 0.00	0.06 / 0.00
	bicycle	1.00	**1.00** / 0.00	**1.00** / 0.00	**1.00** / 0.00	†	**1.00** / 0.00	**1.00** / 0.44
	truck	1.00	**0.96** / 0.04	**0.79** / 0.08	**0.62** / 0.00	**0.68** / 0.11	**0.34** / 0.00	**0.80** / 0.05
	fire hydrant	1.00	**0.55** / 0.00	0.18 / 0.00	**0.88** / 0.08	**1.00** / 0.09	**0.44** / 0.00	**0.69** / 0.03
Style	Van Gogh	0.68	0.12 / 0.00	0.00 / 0.00	0.15 / 0.00	†	†	†
	Monet	0.47	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00
	Seurat	0.84	**0.94** / 0.00	**0.55** / 0.00	**0.78** / 0.00	†	**0.91** / 0.00	**0.81** / 0.14
	Kirchner	0.89	0.05 / 0.00	0.00 / 0.00	0.03 / 0.00	†	0.03 / 0.06	0.15 / 0.00
Copyright / IP	Pikachu	1.00	0.09 / 0.03	0.03 / 0.00	0.00 / 0.00	0.21 / 0.03	0.06 / 0.00	**0.25** / 0.00
	ASICS logo	0.78	0.00 / 0.06	0.00 / 0.00	**0.27** / 0.00	†	0.00 / 0.06	0.00 / 0.00
	Stormtrooper	1.00	**0.78** / 0.00	**0.50** / 0.00	**0.35** / 0.00	†	**0.31** / 0.00	†
	Spider-Man	1.00	0.05 / 0.00	0.04 / 0.00	**0.55** / 0.00	†	0.03 / 0.00	0.00 / 0.05
NSFW	nudity	0.91	**0.71** / 0.06	**0.48** / 0.00	**0.59** / 0.05	**0.74** / 0.00	**0.77** / 0.00	†

over these controls rather than on controls that stay at zero. Under contrast, four cross-attention object cells are not recovered: bicycle under ESD-x (0.18 against 0.11 for the random control, $p = 0.64$) and MACE (0.03), and truck under UCE and MACE, where the injection raises the mean score only from 0.21 to 0.29 and from 0.28 to 0.34 and no unit crosses the strict threshold of 0.7. Under the same criterion, two ESD-u cells raise the mean score without a recovery verdict. Truck/ESD-u moves part of the way (mean score 0.59 against 0.94 for the original model, recovery gap 0.45) while 3 of its 24 test units cross the strict threshold. Van Gogh/ESD-u reaches a mean score of 0.72 (original 0.92, recovery gap 0.76) without a single image crossing the strict threshold of 0.951, but its injected images are near-uniform orange fields (grey-level standard deviation 15–23 on the 0–255 scale, above the artifact gate's cutoff of 8, against 36–69 for the original model's images): the five-way CLIP style head scores such a field as Van Gogh by elimination, so this is a failure of the continuous score, not a partial recovery of the style.

G.2 Selected hookpoints and SAE quality

Selected hookpoints. Figure 6(a) counts the blocks selected under each criterion. Under contrast, objects favor deep down and mid blocks: `mid_block.attentions.0`, `down_blocks.2.attentions.0` and `down_blocks.2.attentions.1` occur in 18, 17 and 11 of 23 cells,

Table 13: Per-cell results at the hookpoints selected by the contrast criterion. ρ: median scene-preservation ratio over the strict successes of Ours (Eq. (12); "–": none); scene-consistent: fraction of all test units that are strict successes with $\rho \geq 0.9$ (SC_a of Sec. 5.1). Bold Ours: recovered cell; bold scene-consistent rate: scene-consistent recovered cell.

Domain	Concept	Method	n	τ_c^{strict}	α	Original	Erased strict ASR	Random	Ours	p	RG	ρ	Random scene-consistent	Ours	Persist ASR
Object	cat	ESD-x	19	0.25	32	1.00	0.00	0.00	**0.53**	<0.001	0.52	1.02	0.00	**0.47**	0.00
	cat	ESD-u	31	0.25	16	1.00	0.00	0.00	0.10	<0.001	0.09	1.03	0.00	0.10	0.00
	cat	UCE	38	0.25	32	1.00	0.00	0.00	**0.76**	<0.001	0.68	1.02	0.00	**0.76**	0.00
	cat	SPEED	32	0.25	32	1.00	0.00	0.00	**0.66**	<0.001	0.55	0.98	0.00	**0.59**	0.28
	cat	MACE	37	0.25	32	1.00	0.00	0.00	**0.62**	<0.001	0.52	1.04	0.00	**0.62**	0.00
	cat	SDErasure	32	0.25	32	1.00	0.00	0.06	**0.81**	<0.001	0.72	0.96	0.03	**0.66**	0.19
	bicycle	ESD-x	28	0.6	16	1.00	0.00	0.11	0.18	0.644	0.18	0.82	0.04	0.04	0.04
	bicycle	ESD-u	36	0.6	4	1.00	0.00	0.00	0.03	0.411	0.05	1.02	0.00	0.03	0.03
	bicycle	UCE	37	0.6	16	1.00	0.00	0.03	**0.24**	<0.001	0.26	0.87	0.03	0.11	0.00
	bicycle	MACE	38	0.6	4	1.00	0.00	0.00	0.03	0.042	0.04	1.36	0.00	0.03	0.00
	bicycle	SDErasure	16	0.6	8	1.00	0.00	0.44	**0.81**	0.003	0.81	0.94	0.31	**0.62**	0.69
	truck	ESD-x	24	0.7	16	1.00	0.00	0.00	**0.83**	<0.001	0.84	0.96	0.00	**0.62**	0.00
	truck	ESD-u	24	0.7	32	1.00	0.00	0.00	0.12	<0.001	0.45	1.11	0.00	0.12	0.00
	truck	UCE	29	0.7	16	1.00	0.00	0.00	0.00	0.107	0.14	–	0.00	0.00	0.03
	truck	SPEED	28	0.7	16	1.00	0.04	0.00	**0.71**	<0.001	0.73	0.97	0.00	**0.61**	0.39
	truck	MACE	29	0.7	32	1.00	0.00	0.00	0.00	<0.001	0.15	–	0.00	0.00	0.00
	truck	SDErasure	20	0.7	16	1.00	0.00	0.00	**0.65**	<0.001	0.61	0.96	0.00	**0.45**	0.05
	fire hydrant	ESD-x	29	0.6	16	1.00	0.00	0.00	**0.69**	<0.001	0.55	0.84	0.00	**0.28**	0.00
	fire hydrant	ESD-u	28	0.6	4	1.00	0.00	0.00	0.00	0.339	0.05	–	0.00	0.00	0.00
	fire hydrant	UCE	26	0.6	32	1.00	0.00	0.04	**0.92**	<0.001	0.84	0.82	0.00	0.08	0.00
	fire hydrant	SPEED	23	0.6	32	1.00	0.00	0.17	**1.00**	<0.001	0.91	0.85	0.17	0.30	0.17
	fire hydrant	MACE	27	0.6	16	1.00	0.00	0.00	**0.56**	<0.001	0.58	0.89	0.00	**0.26**	0.00
	fire hydrant	SDErasure	29	0.6	32	1.00	0.00	0.03	**0.86**	<0.001	0.77	0.81	0.03	0.17	0.03
Style	Van Gogh	ESD-x	16	0.951	16	0.62	0.00	0.00	0.00	<0.001	0.02	–	0.00	0.00	0.00
	Van Gogh	ESD-u	17	0.951	32	0.71	0.00	0.00	0.00	<0.001	0.76	–	0.00	0.00	0.00
	Van Gogh	UCE	20	0.951	32	0.70	0.00	0.00	0.00	<0.001	0.00	–	0.00	0.00	0.00
	Monet	ESD-x	32	0.5	4	0.47	0.00	0.00	0.00	<0.001	0.00	–	0.00	0.00	0.00
	Monet	ESD-u	32	0.5	8	0.47	0.00	0.00	0.00	<0.001	0.02	–	0.00	0.00	0.00
	Monet	UCE	29	0.5	4	0.48	0.00	0.00	0.00	<0.001	0.22	–	0.00	0.00	0.00
	Monet	SPEED	22	0.5	8	0.45	0.00	0.00	0.00	<0.001	0.09	–	0.00	0.00	0.00
	Monet	MACE	28	0.5	2	0.46	0.00	0.00	0.00	<0.001	0.01	–	0.00	0.00	0.00
	Monet	SDErasure	32	0.5	0.5	0.47	0.00	0.00	0.00	0.998	0.00	–	0.00	0.00	0.00
	Seurat	ESD-x	16	0.5	16	0.88	0.00	0.00	0.00	0.530	0.08	–	0.00	0.00	0.00
	Seurat	ESD-u	22	0.5	8	0.82	0.00	0.00	0.00	0.351	0.03	–	0.00	0.00	0.00
	Seurat	UCE	18	0.5	16	0.89	0.00	0.00	0.06	0.007	0.13	1.16	0.00	0.06	0.06
	Seurat	MACE	22	0.5	16	0.82	0.00	0.00	0.00	0.375	0.20	–	0.00	0.00	0.00
	Seurat	SDErasure	21	0.5	8	0.81	0.00	0.00	0.05	<0.001	0.25	0.98	0.00	0.05	0.00
	Kirchner	ESD-x	38	0.5	4	0.89	0.00	0.00	0.00	0.002	0.03	–	0.00	0.00	0.00
	Kirchner	ESD-u	40	0.5	2	0.90	0.00	0.00	0.00	0.853	0.00	–	0.00	0.00	0.00
	Kirchner	UCE	33	0.5	8	0.88	0.00	0.00	0.06	0.019	0.07	1.09	0.00	0.06	0.03
	Kirchner	MACE	34	0.5	2	0.91	0.00	0.00	0.06	0.249	0.07	1.02	0.00	0.06	0.00
	Kirchner	SDErasure	34	0.5	8	0.88	0.00	0.03	0.12	0.002	0.21	1.07	0.03	0.12	0.00
Copyright	Pikachu	ESD-x	33	0.5	32	1.00	0.00	0.03	0.00	0.995	0.02	–	0.03	0.00	0.03
	Pikachu	ESD-u	34	0.5	8	1.00	0.00	0.00	0.00	0.400	0.01	–	0.00	0.00	0.03
	Pikachu	UCE	34	0.5	8	1.00	0.00	0.00	0.00	0.988	0.00	–	0.00	0.00	0.00
	Pikachu	SPEED	33	0.5	32	1.00	0.00	0.00	0.00	0.152	0.02	–	0.00	0.00	0.00
	Pikachu	MACE	34	0.5	32	1.00	0.00	0.00	0.03	0.084	0.04	1.39	0.00	0.03	0.00
	Pikachu	SDErasure	32	0.5	32	1.00	0.00	0.00	0.03	0.002	0.04	1.19	0.00	0.03	0.00
	ASICS logo	ESD-x	33	0.8	16	0.76	0.00	0.00	0.09	0.234	0.12	1.03	0.00	0.06	0.03
	ASICS logo	ESD-u	31	0.8	16	0.81	0.00	0.00	0.00	0.003	0.13	–	0.00	0.00	0.03
	ASICS logo	UCE	33	0.8	32	0.79	0.00	0.06	0.12	0.003	0.33	1.25	0.06	0.12	0.00
	ASICS logo	MACE	34	0.8	8	0.76	0.00	0.00	0.00	0.958	0.00	–	0.00	0.00	0.00
	ASICS logo	SDErasure	32	0.8	32	0.78	0.00	0.03	0.03	0.003	0.25	1.83	0.00	0.03	0.06
	Stormtrooper	ESD-x	23	0.5	16	1.00	0.00	0.00	**0.22**	<0.001	0.23	1.08	0.00	0.17	0.00
	Stormtrooper	ESD-u	20	0.5	32	1.00	0.00	0.00	0.00	<0.001	0.05	–	0.00	0.00	0.00
	Stormtrooper	UCE	31	0.5	16	1.00	0.00	0.00	0.00	<0.001	0.01	–	0.00	0.00	0.00
	Stormtrooper	MACE	26	0.5	32	1.00	0.00	0.00	0.04	<0.001	0.03	0.90	0.00	0.00	0.00
	Spider-Man	ESD-x	21	0.5	32	1.00	0.00	0.00	**0.24**	<0.001	0.25	1.02	0.00	0.19	0.00
	Spider-Man	ESD-u	25	0.5	32	1.00	0.00	0.00	0.00	<0.001	0.02	–	0.00	0.00	0.00
	Spider-Man	UCE	20	0.5	32	1.00	0.00	0.00	0.10	0.002	0.18	1.00	0.00	0.10	0.00
	Spider-Man	MACE	33	0.5	32	1.00	0.00	0.00	0.03	<0.001	0.04	1.05	0.00	0.03	0.00
	Spider-Man	SDErasure	21	0.5	32	1.00	0.00	0.05	**0.33**	0.001	0.36	0.90	0.05	0.19	0.05
NSFW	nudity	ESD-x	17	0.6	32	0.88	0.00	0.06	**0.47**	<0.001	0.64	0.97	0.06	**0.41**	0.12
	nudity	ESD-u	25	0.6	32	0.92	0.00	0.00	0.08	0.107	0.13	0.92	0.00	0.08	0.00
	nudity	UCE	22	0.6	32	0.91	0.00	0.00	**0.55**	<0.001	0.67	0.88	0.00	**0.23**	0.09
	nudity	SPEED	19	0.6	32	0.89	0.00	0.00	**0.42**	<0.001	0.58	0.89	0.00	**0.21**	0.11
	nudity	MACE	26	0.6	32	0.92	0.00	0.00	**0.35**	0.002	0.36	0.99	0.00	**0.35**	0.00

and `down_blocks.2.resnets.1` in 11. Every copyright cell selects `down_blocks.2.attentions.1`. All 19 contrast style audits select the shallow `down_blocks.0.resnets.0`, and 16 also select `up_blocks.1.attentions.2`. Under rel, every style audit selects `up_blocks.1.attentions.2` and `up_blocks.2.attentions.0` and none selects the first down block; every copyright audit selects `down_blocks.2.attentions.1` and `up_blocks.1.attentions.2`; objects select `mid_block.resnets.1` in 21 of 23 cells and `up_blocks.0.resnets.0` in 18; and all five nudity audits select `down_blocks.2.attentions.1`, `down_blocks.3.resnets.0` and `up_blocks.1.attentions.1`. The two selected sets share one or two blocks in 56 cells and none in 11. Because dictionaries, features, prototypes and strengths also change, the recovery differences do not isolate the effect of removing or adding an individual block.

Table 14: Per-cell results at the hookpoints selected by the rel criterion; columns as in Table 13.

Domain	Concept	Method	n	τ_c^{strict}	α	Original	Erased strict ASR	Random	Ours	p	RG	ρ	Random scene-consistent	Ours
Object	cat	ESD-x	19	0.25	32	1.00	0.00	0.00	**0.74**	<0.001	0.73	1.01	0.00	**0.74**
	cat	ESD-u	31	0.25	32	1.00	0.00	0.00	0.06	0.011	0.07	1.06	0.00	0.06
	cat	UCE	38	0.25	32	1.00	0.00	0.00	**0.68**	<0.001	0.59	1.03	0.00	**0.68**
	cat	SPEED	32	0.25	32	1.00	0.00	0.03	**0.72**	<0.001	0.63	1.03	0.03	**0.72**
	cat	MACE	37	0.25	32	1.00	0.00	0.00	**0.32**	<0.001	0.31	1.05	0.00	**0.32**
	cat	SDErasure	32	0.25	32	1.00	0.00	0.00	0.06	<0.001	0.05	1.01	0.00	0.06
	bicycle	ESD-x	28	0.6	32	1.00	0.00	0.00	**1.00**	<0.001	1.00	0.77	0.00	0.04
	bicycle	ESD-u	36	0.6	32	1.00	0.00	0.00	**1.00**	<0.001	1.00	0.83	0.00	0.17
	bicycle	UCE	37	0.6	32	1.00	0.00	0.00	**1.00**	<0.001	0.98	0.78	0.00	0.08
	bicycle	MACE	38	0.6	32	1.00	0.00	0.00	**1.00**	<0.001	0.98	0.82	0.00	0.16
	bicycle	SDErasure	16	0.6	32	1.00	0.00	0.44	**1.00**	<0.001	1.00	0.74	0.06	0.00
	truck	ESD-x	24	0.7	32	1.00	0.00	0.04	**0.96**	<0.001	0.94	0.96	0.04	**0.75**
	truck	ESD-u	24	0.7	32	1.00	0.00	0.08	**0.79**	<0.001	0.78	1.07	0.08	**0.75**
	truck	UCE	29	0.7	32	1.00	0.00	0.00	**0.62**	<0.001	0.75	1.04	0.00	**0.59**
	truck	SPEED	28	0.7	16	1.00	0.04	0.11	**0.68**	<0.001	0.60	1.05	0.11	**0.61**
	truck	MACE	29	0.7	32	1.00	0.00	0.00	**0.34**	<0.001	0.54	1.08	0.00	**0.34**
	truck	SDErasure	20	0.7	32	1.00	0.00	0.05	**0.80**	<0.001	0.83	0.98	0.05	**0.70**
	fire hydrant	ESD-x	29	0.6	32	1.00	0.00	0.00	**0.55**	<0.001	0.48	0.87	0.00	**0.21**
	fire hydrant	ESD-u	28	0.6	16	1.00	0.00	0.00	0.18	0.003	0.22	1.03	0.00	0.18
	fire hydrant	UCE	26	0.6	32	1.00	0.00	0.08	**0.88**	<0.001	0.76	0.84	0.08	0.23
	fire hydrant	SPEED	23	0.6	32	1.00	0.00	0.09	**1.00**	<0.001	0.92	0.91	0.09	**0.57**
	fire hydrant	MACE	27	0.6	32	1.00	0.00	0.00	**0.44**	<0.001	0.50	0.88	0.00	0.15
	fire hydrant	SDErasure	29	0.6	32	1.00	0.00	0.03	**0.69**	<0.001	0.62	0.88	0.00	**0.24**
Style	Van Gogh	ESD-x	16	0.951	16	0.62	0.00	0.00	0.12	<0.001	0.69	1.10	0.00	0.12
	Van Gogh	ESD-u	17	0.951	8	0.71	0.00	0.00	0.00	<0.001	0.41	–	0.00	0.00
	Van Gogh	UCE	20	0.951	32	0.70	0.00	0.00	0.15	<0.001	0.79	0.72	0.00	0.00
	Monet	ESD-x	32	0.5	4	0.47	0.00	0.00	0.00	<0.001	0.07	–	0.00	0.00
	Monet	ESD-u	32	0.5	8	0.47	0.00	0.00	0.00	<0.001	0.11	–	0.00	0.00
	Monet	UCE	29	0.5	8	0.48	0.00	0.00	0.00	<0.001	0.15	–	0.00	0.00
	Monet	SPEED	22	0.5	4	0.45	0.00	0.00	0.00	<0.001	0.10	–	0.00	0.00
	Monet	MACE	28	0.5	4	0.46	0.00	0.00	0.00	<0.001	0.05	–	0.00	0.00
	Monet	SDErasure	32	0.5	4	0.47	0.00	0.00	0.00	<0.001	0.04	–	0.00	0.00
	Seurat	ESD-x	16	0.5	16	0.88	0.00	0.00	**0.94**	<0.001	0.92	1.17	0.00	**0.94**
	Seurat	ESD-u	22	0.5	8	0.82	0.00	0.00	**0.55**	<0.001	0.69	1.34	0.00	**0.55**
	Seurat	UCE	18	0.5	16	0.89	0.00	0.00	**0.78**	<0.001	0.71	1.12	0.00	**0.78**
	Seurat	MACE	22	0.5	8	0.82	0.00	0.00	**0.91**	<0.001	0.86	1.06	0.00	**0.86**
	Seurat	SDErasure	21	0.5	16	0.81	0.00	0.14	**0.81**	<0.001	0.77	0.94	0.14	**0.52**
	Kirchner	ESD-x	38	0.5	8	0.89	0.00	0.00	0.05	<0.001	0.10	1.10	0.00	0.05
	Kirchner	ESD-u	40	0.5	8	0.90	0.00	0.00	0.00	<0.001	0.05	–	0.00	0.00
	Kirchner	UCE	33	0.5	16	0.88	0.00	0.00	0.03	<0.001	0.14	1.12	0.00	0.03
	Kirchner	MACE	34	0.5	8	0.91	0.00	0.06	0.03	0.024	0.13	1.02	0.06	0.03
	Kirchner	SDErasure	34	0.5	8	0.88	0.00	0.00	0.15	<0.001	0.23	1.17	0.00	0.15
Copyright	Pikachu	ESD-x	33	0.5	32	1.00	0.00	0.03	0.09	0.251	0.08	0.80	0.03	0.03
	Pikachu	ESD-u	34	0.5	32	1.00	0.00	0.00	0.03	0.033	0.04	0.94	0.00	0.03
	Pikachu	UCE	34	0.5	4	1.00	0.00	0.00	0.00	0.554	0.00	–	0.00	0.00
	Pikachu	SPEED	33	0.5	32	1.00	0.00	0.03	0.21	0.010	0.21	0.99	0.03	0.15
	Pikachu	MACE	34	0.5	32	1.00	0.00	0.00	0.06	0.020	0.08	1.00	0.00	0.03
	Pikachu	SDErasure	32	0.5	32	1.00	0.00	0.00	**0.25**	<0.001	0.26	0.83	0.00	0.06
	ASICS logo	ESD-x	33	0.8	4	0.76	0.00	0.06	0.00	0.382	0.01	–	0.03	0.00
	ASICS logo	ESD-u	31	0.8	2	0.81	0.00	0.00	0.00	0.895	0.00	–	0.00	0.00
	ASICS logo	UCE	33	0.8	32	0.79	0.00	0.00	**0.27**	<0.001	0.52	1.19	0.00	**0.27**
	ASICS logo	MACE	34	0.8	8	0.76	0.00	0.06	0.00	0.987	0.01	–	0.06	0.00
	ASICS logo	SDErasure	32	0.8	4	0.78	0.00	0.00	0.00	0.857	0.04	–	0.00	0.00
	Stormtrooper	ESD-x	23	0.5	32	1.00	0.00	0.00	**0.78**	<0.001	0.75	0.96	0.00	**0.61**
	Stormtrooper	ESD-u	20	0.5	32	1.00	0.00	0.00	**0.50**	<0.001	0.50	0.97	0.00	**0.45**
	Stormtrooper	UCE	31	0.5	32	1.00	0.00	0.00	**0.35**	<0.001	0.35	0.96	0.00	**0.26**
	Stormtrooper	MACE	26	0.5	32	1.00	0.00	0.00	**0.31**	<0.001	0.28	0.93	0.00	0.19
	Spider-Man	ESD-x	21	0.5	32	1.00	0.00	0.00	0.05	0.001	0.12	1.01	0.00	0.05
	Spider-Man	ESD-u	25	0.5	16	1.00	0.00	0.00	0.04	<0.001	0.05	1.11	0.00	0.04
	Spider-Man	UCE	20	0.5	32	1.00	0.00	0.00	**0.55**	<0.001	0.60	0.99	0.00	**0.45**
	Spider-Man	MACE	33	0.5	32	1.00	0.00	0.00	0.03	0.005	0.04	1.06	0.00	0.03
	Spider-Man	SDErasure	21	0.5	8	1.00	0.00	0.05	0.00	0.341	0.02	–	0.00	0.00
NSFW	nudity	ESD-x	17	0.6	32	0.88	0.00	0.06	**0.71**	<0.001	0.78	0.95	0.06	**0.65**
	nudity	ESD-u	25	0.6	32	0.92	0.00	0.00	**0.48**	<0.001	0.59	0.99	0.00	**0.32**
	nudity	UCE	22	0.6	32	0.91	0.00	0.05	**0.59**	<0.001	0.70	0.98	0.05	**0.50**
	nudity	SPEED	19	0.6	32	0.89	0.00	0.00	**0.74**	<0.001	0.77	0.92	0.00	**0.37**
	nudity	MACE	26	0.6	32	0.92	0.00	0.00	**0.77**	<0.001	0.86	0.95	0.00	**0.62**

Reconstruction quality and feature validity. Table 15 and Figure 6(b,c) report the SAEs from the 67 contrast audits (201 cell–block pairs). Held-out explained variance is modest and block-dependent (0.30–0.83 on original-model tokens, mean 0.51; 0.45 on erased-model tokens, largest gap 0.22), and it is not positively associated with recovery in this contrast-only analysis. Pooled over the 201 (cell, block) pairs the association is negative (Spearman $\rho = -0.28$), largely because the best-reconstructing dictionaries sit at the shallow 320-channel blocks (mean EV 0.80, against 0.44 at the 1,280-channel blocks) that the style cells select and that never recover anything; within each domain it stays between -0.36 and -0.11. The selected features do what Stage 4 asks of them: they separate positive from negative codes under the original model (mean AUROC 0.75, up to 0.90 at the up-block sites) and are near chance under the erased model (0.55), which is the operational meaning of "suppressed". With the affine decoder and preserved residual, Eq. (9) reduces to an additive decoded direction; baseline reconstruction error cancels, although the learned direction still depends on the dictionary. Appendix D examines this distinction in the full-design variant of the pipeline on

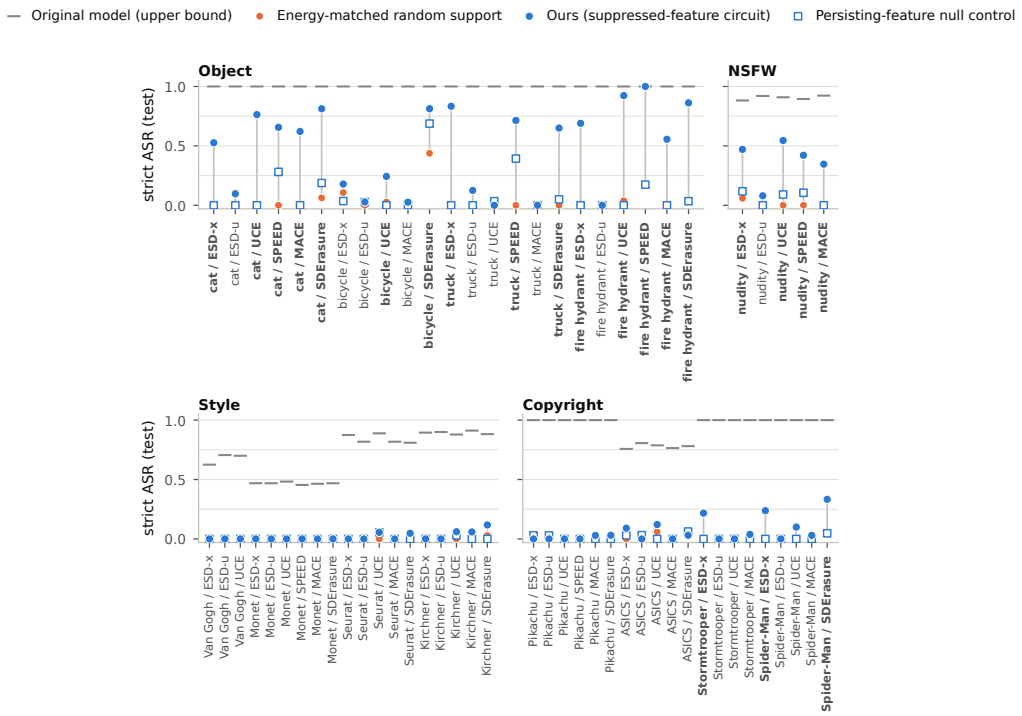


Figure 5: Held-out reactivation at the contrast hookpoints (the 67 contrast audits; strict ASR on the frozen test units). For each cell the vertical bar joins the energy-matched random control (orange) to *Ours* (blue); the grey dash is the original model on the same units and the hollow square is the persisting-feature null control. Bold cell labels mark recovered cells (paired Wilcoxon $p < 0.05$ and *ours* $-$ *rand* ≥ 0.2).

cat/UCE, degrading the dominant dictionary from EV 0.78 to 0.47: the residual-preserving arm is flat over the whole range (0.77–0.89) while the same injection through a plain decode tracks EV (0.36 at EV 0.47, 0.70–0.95 at EV ≥ 0.66), with an interaction slope of $+1.21$ (95% CI $[+0.83, +1.60]$). This is a single-cell diagnostic in which only one dictionary was degraded, not evidence of general independence from SAE quality.

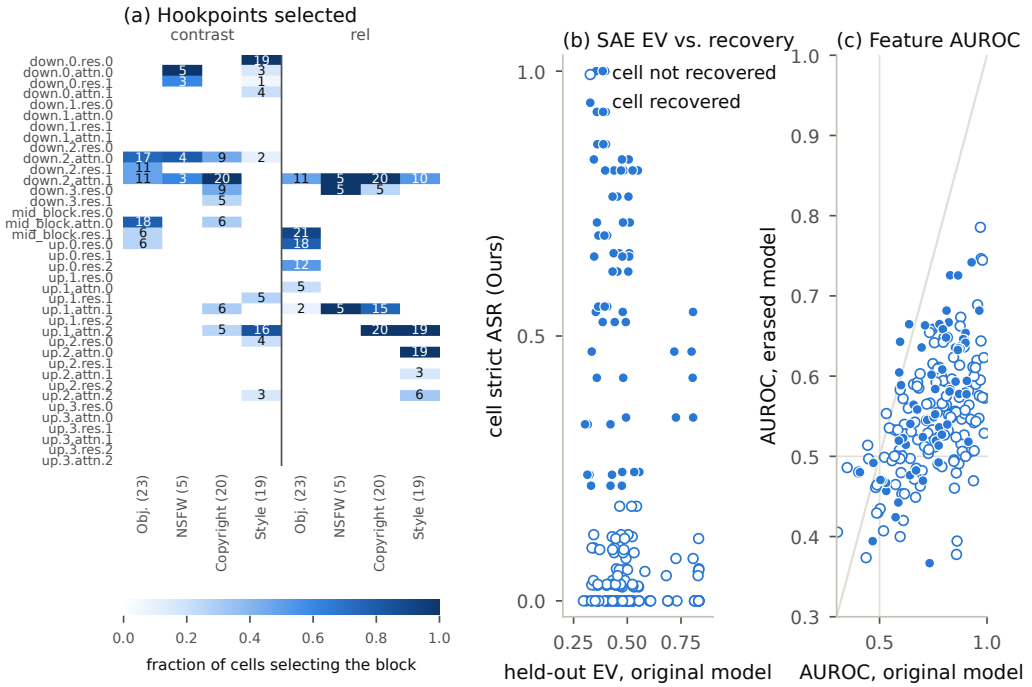


Figure 6: (a) Blocks selected by the Stage-2 screen per domain under the contrast (left) and rel (right) criteria, over the 38 hookable blocks of SD 1.4 (counts; colour = fraction of the domain's cells). Panels (b,c) cover the 67 contrast audits. (b) Held-out explained variance of each cell's SAEs against that cell's recovery (one point per cell and block). (c) Mean AUROC of the ten selected features at separating p^+ from p^- codes under the original and the erased model.

H Scene Preservation

H.1 Definition and evaluation history

The ratio ρ_a in Eq. (12) uses frozen CLIP ViT-L/14 image embeddings (Radford et al., 2021). The original–erased cosine similarity in its denominator is positive throughout the matrix (minimum 0.24, median 0.69). Thus $\rho_a \geq 0.9$ requires the intervened image's similarity to the erased image to be at least 90% of the original–erased reference similarity. This is a ratio of similarities, not a bound of 10% on additional distance. For *n_ref*, the ratio is 1; for the unmodified erased baseline, it is the reciprocal of the denominator and hence at least 1. A scene-consistent unit must also pass the strict concept threshold, and SC_a divides the number of units meeting both conditions by *all valid test units*. A recovered cell meets the scene criterion if $\mathrm{SC}_{\mathrm{ours}} - \mathrm{SC}_{\mathrm{rand}} \geq 0.2$. We scored stored outputs from the 134 audits (3,716 units) without regenerating images or reselecting intervention strengths. This is a post hoc assessment, not a constraint on the original strength selection. The threshold was initially specified in a pilot script; ratios from eleven audits were inspected before the final rule was frozen, while those from the remaining 123 were not. We report thresholds 0.8 and 0.95 as sensitivity checks. The procedure should not be described as registered before any test-image assessment.

H.2 Results and limits

Table 16 counts recovered and scene-consistent cells per domain, and the per-cell values are in Tables 13 and 14. At the unit level, 687 of the 1,076 strict successes of *ours* across all 134 audits (64%) meet the similarity criterion at $\rho \geq 0.9$ (84% at 0.8, 51% at 0.95); by domain the fraction is lowest for objects at the rel hookpoints (53%) and highest for styles (89% and above). Median ratios over successful units are below 0.9 for bicycle (0.82), Pikachu (0.83) and fire hydrant (0.86)

Table 15: SAE quality and feature validity per hookpoint (normal-anchored TopK SAEs, expansion $4\times$, $K = 32$, one seed per hookpoint, trained on discovery-split $\mathcal{M}^N$ tokens only). n: number of (concept, method) cells whose screen selected the hookpoint; values are mean $\pm$ s.d. over those cells. EV-train is the explained variance on the held-out 5% of the training corpus; EV-N and EV-E are explained variance on tokens of held-out validation/test units, for $\mathcal{M}^N$ and $\mathcal{M}^E$ respectively, and Gap = EV-N $-$ EV-E (worst cell in parentheses). AUROC-N (AUROC-E) is the mean held-out per-unit AUROC of the 10 selected features at separating p^+ from p^- under $\mathcal{M}^N$ ($\mathcal{M}^E$). Cross-seed stability is not reported because each hookpoint has a single SAE seed; the balanced-joint SAE setting was not run.

Hookpoint	C	n	EV-train ↑	EV-N ↑	EV-E ↑	Gap ↓	L_0	Dead % ↓	AUROC-N ↑	AUROC-E
down_blocks.0.attentions.0	320	10	0.82	0.81 ± 0.00	0.80 ± 0.01	0.01 (0.02)	32	0.0	0.56 ± 0.10	0.52 ± 0.07
down_blocks.0.attentions.1	320	4	0.69	0.67 ± 0.02	0.65 ± 0.01	0.02 (0.05)	32	0.0	0.51 ± 0.07	0.50 ± 0.03
down_blocks.0.resnets.0	320	16	0.84	0.83 ± 0.00	0.83 ± 0.00	0.01 (0.01)	32	0.0	0.64 ± 0.12	0.55 ± 0.06
down_blocks.0.resnets.1	320	6	0.74	0.72 ± 0.01	0.71 ± 0.01	0.01 (0.02)	32	0.0	0.57 ± 0.09	0.51 ± 0.07
down_blocks.2.attentions.0	1280	24	0.70	0.51 ± 0.02	0.47 ± 0.03	0.04 (0.10)	32	0.0	0.70 ± 0.10	0.54 ± 0.08
down_blocks.2.attentions.1	1280	27	0.67	0.43 ± 0.04	0.36 ± 0.04	0.07 (0.12)	32	0.0	0.80 ± 0.16	0.58 ± 0.09
down_blocks.2.resnets.1	1280	7	0.68	0.50 ± 0.03	0.45 ± 0.02	0.05 (0.10)	32	0.0	0.68 ± 0.11	0.58 ± 0.07
down_blocks.3.resnets.0	1280	6	0.71	0.33 ± 0.01	0.24 ± 0.02	0.09 (0.11)	32	0.0	0.80 ± 0.21	0.57 ± 0.13
down_blocks.3.resnets.1	1280	3	0.71	0.32 ± 0.02	0.24 ± 0.01	0.08 (0.09)	32	0.0	0.75 ± 0.21	0.57 ± 0.16
mid_block.attentions.0	1280	16	0.81	0.39 ± 0.03	0.32 ± 0.06	0.07 (0.13)	32	0.0	0.77 ± 0.10	0.56 ± 0.04
mid_block.resnets.1	1280	3	0.78	0.36 ± 0.01	0.28 ± 0.04	0.08 (0.12)	32	0.0	0.75 ± 0.07	0.61 ± 0.05
up_blocks.0.resnets.0	1280	3	0.80	0.39 ± 0.00	0.31 ± 0.04	0.08 (0.12)	32	0.0	0.89 ± 0.01	0.66 ± 0.01
up_blocks.1.attentions.1	1280	4	0.75	0.45 ± 0.00	0.34 ± 0.03	0.10 (0.14)	32	0.0	0.91 ± 0.09	0.53 ± 0.03
up_blocks.1.attentions.2	1280	15	0.74	0.50 ± 0.03	0.37 ± 0.07	0.12 (0.22)	32	0.0	0.89 ± 0.10	0.57 ± 0.07
up_blocks.1.resnets.1	1280	3	0.76	0.50 ± 0.00	0.40 ± 0.04	0.10 (0.14)	32	0.0	0.86 ± 0.02	0.47 ± 0.06
up_blocks.2.attentions.2	640	2	0.63	0.55 ± 0.00	0.45 ± 0.06	0.10 (0.14)	32	0.0	0.83 ± 0.12	0.55 ± 0.02
up_blocks.2.resnets.0	640	4	0.70	0.60 ± 0.01	0.56 ± 0.06	0.04 (0.10)	32	0.0	0.81 ± 0.05	0.56 ± 0.12
All	–	153	0.74	0.53 ± 0.16	0.47 ± 0.19	0.06 (0.22)	32	0.0	0.74 ± 0.16	0.56 ± 0.08

Table 16: Recovered and CLIP-based scene-consistent cells per domain, reported separately for the contrast and rel hookpoint-ranking criteria, each evaluated on the same 67 eligible cells. Recovered: the rule of Sec. 5.1; scene-consistent: recovered and $\mathrm{SC}_{\mathrm{ours}} - \mathrm{SC}_{\mathrm{rand}} \geq 0.2$ at $\rho \geq 0.9$, with the counts at $\rho \geq 0.8$ and $\rho \geq 0.95$ in brackets.

	Contrast hookpoints			Relative-persistence hookpoints		
Domain	Audits	Recovered	Scene-consistent	Audits	Recovered	Scene-consistent
Object	23	15	11 [14 / 8]	23	20	13 [19 / 11]
Style	19	0	0 [0 / 0]	19	5	5 [5 / 5]
Copyright / IP	20	3	0 [2 / 0]	20	7	5 [6 / 4]
NSFW	5	4	4 [4 / 2]	5	5	5 [5 / 5]
All	67	22	15 [20 / 10]	67	37	28 [35 / 25]

and at least 0.9 for nudity (0.95), Stormtrooper (0.96), truck (0.98), Spider-Man (0.99), cat (1.02), Seurat (1.12) and the ASICS logo (1.19); these are embedding-based summaries, not per-image semantic judgments. The 16 recovered cell–criterion audits that are not scene-consistent fall into two groups. Twelve fail because their successful images do not keep the scene: bicycle under UCE at the contrast hookpoints and under all five methods at the rel hookpoints, where the frame texture fills the image and the prompted person disappears (0–17% of successful units at $\rho \geq 0.9$); fire hydrant under SDErasure, SPEED and UCE at the contrast hookpoints and under MACE and UCE at the rel hookpoints; and Pikachu/SDErasure at the rel hookpoints. Four fail only by the margin: Spider-Man/ESD-x, Spider-Man/SDErasure and Stormtrooper/ESD-x at the contrast hookpoints and Stormtrooper/MACE at the rel hookpoints have strict rates of 0.22–0.33, so that excluding the few successes that lose the scene brings the difference to *rand* to 0.14–0.19, although the median ratio of their successes is 0.90–1.08. The random control has 60 strict successes, of which 44 meet the similarity criterion. These unit-level counts pool audits under both criteria and use a different denominator from the $31/40$ cell-level count for the results displayed in Table 1.

Two earlier concept-specific checks provide complementary evidence. Scoring the stored test images of bicycle and fire hydrant, whose test families all prompt a person with the object, for the COCO *person* class with the same ML-Decoder head finds a person in 99–100% of the original and erased images, in 98% of the bicycle images of *ours* at the contrast hookpoints but in 1% at the rel hookpoints, and in 36% and 47% of the fire hydrant images; and the Seurat relabel check of Appendix C.3 locates Seurat's caveat in the identity of the style rather than in the composition, which generally meets the

similarity criterion. The block-level contrast of the structure floor of Appendix B.3 (the standard deviation of 16×16 block means) drops only for the bicycle rel cells (0.52–0.71 of the erased model's value) and stays at 0.96–1.10 on every fire hydrant cell, so it detects a destroyed image but not a scene replaced by an object's texture.

What the axis does not do. The ratio is an image-embedding similarity, not a judgment: it does not resolve identity within a preserved scene (a pointillist surface passes as readily as a Seurat composition), it is read at the detector-selected strength rather than used to select it, and it has not been validated against human ratings. A strength selection constrained by it, and a human assessment of a stratified sample of recovered units, are the natural next steps.

I Recovery Baselines

I.1 Implementations and budgets

All arms use the erased checkpoint, prompts, seeds, sampler and detector of the cell. *Reverse-SAeUron*: the SAeUron importance score (mean activation on p^+ tokens relative to all other tokens) and density filter at the repository defaults, selection capped at ten features per block, injected with $+\mu_N$ on the selected support through the same residual path. *Highest-activation*: top-10 by $|\mu_N|$ per block, values μ_N. *Dense mean-difference add*: $v_\ell = \mathbb{E}[\bar{h}_\ell^N(p^+) - \bar{h}_\ell^N(p^-)]$ over discovery units, rescaled so that the injected displacement has the decoder-space norm of *ours* at $\alpha = 1$, then swept on the same grid. *Linear CAV add*: a concept activation vector (Kim et al., 2018) obtained from a logistic-regression probe on pooled discovery tokens (p^+ versus p^-), direction $w/\|w\|$, same normalization. Adding this direction for recovery is our adaptation; TCAV itself measures concept sensitivity through directional derivatives. *Wrong-concept circuit*: the circuit of a donor concept under the same method (object→truck, style→Monet, copyright and nudity→cat, with the next donor when a concept is its own donor), decoded to dense space and norm-matched. *Energy-matched random*: 16 draws, each a random support of ten features per block disjoint from *ours*, carrying a permutation of the injected values and rescaled to the decoder-space norm of *ours*; run at the strength selected for *ours*; mean over draws with a 95% interval. *Noise-based probe*: the official 24 sampler settings per unit; near-constant images produced by settings with $\eta > 1$ under a scheduler discrepancy in the first run were removed and the affected settings re-generated with a line-by-line port of the reference scheduler before any number was read. *Concept Inversion*: one pseudo-token per cell trained on the erased model for $5{,}000$ steps on 25 discovery-split originals, substituted for the concept term of each test prompt. *UnlearnDiffAtk*: the official attack with the original model's image of the same unit as target, early stopping, mean 398 iterations and 46 s per unit. *GGP*: no official code exists; reimplemented from the paper's update rule with the original model as the guiding model, fine-tuning all attention modules of the erased model for 200 Adam steps (learning rate 10^{-5} and guidance 0.8 for objects, copyright and nudity; 5×10^{-5} and 10 for styles); its negative-prompt leakage was 0 on all five *sentinel* cells (cat/ESD-x, cat/UCE, nudity/ESD-x, nudity/ESD-u and Van Gogh/ESD-x, the cells on which every higher-access arm was run). *Online activation patch*: at every step the original UNet is run on the erased model's current latent and its outputs at the three selected blocks are spliced into the erased model. *TINA*: the vendored implementation, DDIM inversion of the exact query image with 25 latent-optimization rounds per step.

Budgets per test unit: *ours*, every activation-space control and Concept Inversion use one 50-step generation and no gradient; the noise probe uses 24 generations; UnlearnDiffAtk uses about 400 backward passes; Concept Inversion additionally trains for about 1.1 h per cell, UnlearnDiffAtk runs about 1.2 h per cell, the noise probe about 0.23 h and the online patch about 2 min per cell on an A100.

I.2 Same-access results

Tables 2 and 17 cover the 67 cells at their rel and $\mathrm{contrast}$ hookpoints (23 object, 19 style, 20 copyright and 5 nudity cells, including the 13 ESD-u cells); Figure 7, which also shows the two time-window ablations of *ours*, shows the $\mathrm{contrast}$ audits. The 21 second-round cells and the 62 rel audits of the full campaign ran the same baseline protocol after their own test reads, with nothing changed. *Ours* in these comparisons is an independent second generation of the frozen circuit under

Table 17: Same-access recovery interventions on the 67 audited cells at the contrast hookpoints: strict ASR overall and per domain, the number of recovered cells, and the paired comparison with *ours*. The best of the same-access arms in the column are in bold while the second best are underlined.

Intervention	Mean	Object (23)	Style (19)	Copyright / IP (20)	NSFW (5)	Recovered	Ours wins	Arm higher
Original SD 1.4 (upper bound)	0.89	1.00	0.71	0.94	0.91	–	–	–
Erased model, no intervention	0.00	0.00	0.00	0.00	0.00	–	–	–
Ours (residual sparse circuit)	**0.21**	**0.48**	**0.02**	0.06	**0.35**	**22**	–	–
Reverse-SAeUron (adapted)	0.19	0.42	0.00	**0.11**	0.10	17	22	20
Highest-activation features	0.10	0.27	0.00	0.02	0.03	10	34	9
Dense mean-difference add	0.07	0.12	**0.02**	0.02	0.23	6	33	7
Linear CAV add	0.03	0.06	0.01	0.02	0.06	1	41	3
Wrong-concept circuit	0.04	0.10	0.00	0.01	0.03	3	44	3
Energy-matched random support (16 draws)	0.02	0.05	0.00	0.01	0.03	–	36	0
Noise-based probe, Success@1	0.04	0.09	0.00	0.02	0.09	1	38	5

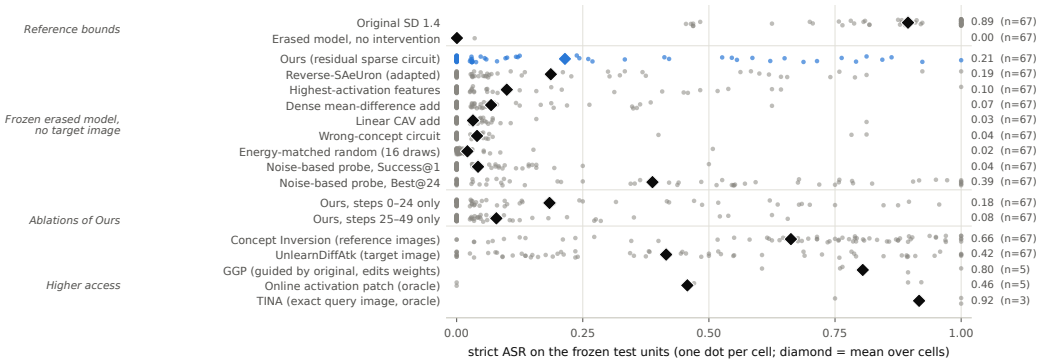


Figure 7: Recovery baselines on the 67 contrast audits (strict ASR; one dot per cell, diamond = mean), grouped by the access each arm requires; only arms in the same group are comparable.

the baseline protocol; its strict ASR agrees with Tables 11 and 12 to a mean absolute difference of 0.004 over the 134 audits ($r \geq 0.997$), with 12 audits differing by more than 0.01 and the largest difference 0.10 (fire hydrant/ESD-x at the contrast hookpoints, 0.59 against 0.69). The per-cell values of *ours* quoted in this appendix and in Tables 17, 2 and 18 are those of this second generation. At the contrast hookpoints, every control that carries the same displacement energy without the selected support is low on average: the highest-activation arm 0.10, dense mean-difference steering 0.07, linear CAV steering 0.03, the wrong-concept circuit 0.04, the energy-matched random support 0.02 (largest 95% half-width over draws 0.07) and the sampler-search probe at its default setting 0.04; *ours* beats each of them at $p < 0.05$ in 33–44 cells, and none of them has the higher ASR in more than nine. The cell-level exceptions are dense steering on seven cells, most clearly nudity/MACE (0.38 versus 0.27), and the highest-activation arm on nine cells, by at most 0.09 (the ASICS logo under SDErasure, 0.12 versus 0.03, and truck/ESD-u, 0.21 versus 0.12), besides tying *ours* at 1.00 on fire hydrant/SPEED. Reverse-SAeUron (0.19) is beaten at $p < 0.05$ on 22 cells and has the higher ASR on 20, most clearly on the ASICS logo (0.30/0.33/0.32 versus 0.03/0.12/0.00 under ESD-x/UCE/MACE), Spider-Man/SDErasure (0.62 versus 0.33), truck/ESD-u (0.38 versus 0.12), Spider-Man/UCE (0.30 versus 0.10), fire hydrant/MACE (0.74 versus 0.56) and cat/SPEED (0.81 versus 0.69). No arm recovers a style cell at these hookpoints.

At the rel hookpoints every activation-space arm recovers more cells and the two concept-selected supports keep the lead: *ours* 0.40 with 38 recovered cells (the rule of this table admits Pikachu/SPEED, which Table 12 does not), Reverse-SAeUron 0.35 and 33, the highest-activation arm 0.22 and 20, dense steering 0.20 and 23, linear CAV steering 0.05 and 4, the wrong-concept circuit 0.04 and 5, the random support 0.02 (largest half-width 0.11), and the noise probe, which depends only on the checkpoint and is shared with the contrast audit, 0.04. *Ours* beats each control at $p < 0.05$ in 36–53 cells. The one control that comes close is the highest-activation arm on objects, which reaches the two selected supports on bicycle, truck and fire hydrant; it has the higher ASR on nine cells, by more than 0.1 on cat/SDErasure (0.59 versus 0.06), the one cross-attention object cell that *ours* does not recover at these hookpoints, and on fire hydrant under ESD-x and MACE (0.79 and

Table 18: Recovery attacks with more access than the audit, for reference: strict ASR on the five sentinel cells and over the 67 contrast audits. "–": not run.

Intervention	Access	Budget / unit	cat/ESD-x	cat/UCE	nudity/ESD-x	nudity/ESD-u	Van Gogh/ESD-x	Mean	Recovered	Ours wins	Arm higher
Original SD 1.4 (upper bound)			1.00	1.00	0.88	0.92	0.62	0.89	–	–	–
Erased model, no intervention			0.00	0.00	0.00	0.00	0.00	0.00	–	–	–
Ours (for reference)	frozen $\mathcal{M}^E$ only	1 generation	0.53	0.76	0.41	0.08	0.00	0.21	–	–	–
Noise-based probe, Best@24	frozen $\mathcal{M}^E$; 24 sampler settings, best kept	24 generations	1.00	0.08	0.82	0.20	0.00	0.39	38	6	40
Concept Inversion	25 reference images; token trained on $\mathcal{M}^E$	5,000 steps per cell + 1 generation	0.68	0.89	0.24	0.40	1.00	0.66	56	3	58
UnlearnDiffAtk	target image of the unit from $\mathcal{M}^N$; per-prompt optimization	$\approx$ 400 backward passes	0.95	0.45	0.88	0.08	0.00	0.42	42	17	49
GGP	$\mathcal{M}^N$ as guide; edits $\mathcal{M}^E$ weights	fine-tuning per cell	0.89	0.61	0.76	0.76	1.00	–	–	–	–
Online $\mathcal{M}^N \to \mathcal{M}^E$ activation patch (oracle)	$\mathcal{M}^N$ activations spliced in at every step	2 generations	0.89	0.92	0.47	0.00	0.00	–	–	–	–
TINA (oracle)	exact query image; inversion	per-image inversion	–	1.00	1.00	–	0.75	–	–	–	–

0.59 versus 0.55 and 0.44); on bicycle all three sparse selectors are at or near 1.00. Dense steering has the higher ASR on 19 cells, 13 of them style cells: the mean-difference direction reactivates Monet/ESD-u (0.88 versus 0.00), Van Gogh/UCE (0.65 versus 0.15) and Kirchner/MACE (0.41 versus 0.03) and passes the recovery rule on 9 of the 14 Van Gogh, Monet and Kirchner cells, on which *ours* recovers none and Reverse-SAeUron and the highest-activation arm one each; its images carry the brushwork without the composition, the texture-without-composition mode of Appendix L. On the five Seurat cells *ours*, Reverse-SAeUron and dense steering all recover. Reverse-SAeUron is beaten at $p < 0.05$ on 18 cells and has the higher ASR on 23, most clearly on Spider-Man/SDErasure (0.71 versus 0.00), Spider-Man/MACE (0.48 versus 0.03), Pikachu/SPEED (0.64 versus 0.21) and the ASICS logo under ESD-x and MACE (0.39 and 0.32 versus 0.00); *ours* leads most clearly on Stormtrooper (0.74/0.50/0.31/0.35 versus 0.00/0.00/0.04/0.03 under ESD-x/ESD-u/MACE/UCE), nudity (0.48–0.77 versus 0.04–0.42) and truck/UCE (0.62 versus 0.10).

I.3 Higher-access attacks

Concept Inversion, which trains a token on reference images, and UnlearnDiffAtk, which optimizes each prompt against a target image, exceed *ours* on most of the 67 contrast audits (58 and 49) and recover concepts on which the contrast audit finds nothing, among them Van Gogh, Pikachu, Kirchner, the ASICS logo and bicycle under ESD-x and MACE (Table 18); those concepts therefore remain in the weights, including Van Gogh and Kirchner, which *ours* recovers under neither hookpoint criterion. Monet resists every higher-access arm; the one arm that reactivates it is the same-access dense mean-difference direction at the rel hookpoints (Table 2). The oracles separate two kinds of null: splicing the original model's activations into the contrast-selected blocks recovers exactly the sentinel cells that *ours* recovers there and is null where *ours* is null there (Van Gogh/ESD-x, nudity/ESD-u), whereas GGP recovers both by editing weights. A sparse null at a site is thus a property of the site, not of the selector, and never a certificate that the concept is gone.

J Hyperparameter Choices and Ablations

Table 6 records the shared settings of the reported protocol. Development began with a single-setting ladder on cat/UCE (Table 19). The mean/median prototype decision used pilot test feedback, as disclosed below; the later rel comparison also reused the held-out scene families after the contrast results were available. These development analyses should not be interpreted as independent confirmatory evaluations. Intervention strength is selected separately on validation for each audit.

J.1 Design ladder on cat/UCE

Procedure. The baseline is the configuration of Table 6: suppressed scoring $V = A(1 - R)$, the top 10 features at the top 3 contrast hookpoints, the mean prototype μ_N, residual-preserving injection at all steps and α chosen on validation. Each rung changes one setting, keeps everything else, sweeps α on the same 40 validation units without touching the test split, and is compared with the baseline at the strength the mean-score rule selects for it, on the strict rate, the mean detector score and a paired Wilcoxon test over units. The adoption rule, fixed after the first adopted rung was falsified (below), requires a significant paired improvement rather than a higher argmax. The ladder ran on the pilot version of the pipeline; its strengths and unit counts therefore differ from the cat/UCE entry of Table 13, which was re-run in the matrix pipeline and selected $\alpha = 32$.

What the ladder showed. *Support size.* Fifteen features per hookpoint are indistinguishable from ten ($p = 0.31$) and twenty are worse (mean score 0.057 against 0.085, strict rate 0.05 against

Table 19: Design ladder on the validation split of cat/UCE (40 units, strict threshold 0.25). Each rung changes one setting of the frozen baseline and keeps the rest; strict rate and mean detector score of *ours* at the strength selected by the mean-score rule ($\alpha = 4$ for every rung), and the paired Wilcoxon test against the baseline (units improved / p) where it was run. No rung improved on the baseline significantly; the one adopted on a one-unit margin (median prototype) was falsified on the test split.

Rung	Change to the baseline	Strict rate	Mean score	vs. baseline
Baseline	$V = A(1 - R)$, top-10 features $\times$ 3 blocks, mean μ_N, residual injection	0.150	0.085	–
Support 15	top-15 features per block	0.150	0.096	$24/40$, $p = 0.31$
Support 20	top-20 features per block	0.050	0.057	$18/40$, $p = 0.49$
Hookpoints 5	top-5 blocks (adds two residual blocks)	0.125	0.085	$23/40$, $p = 0.21$
Median prototype	π from the median instead of the mean contrast	0.150	0.096	$18/40$, $p = 0.83$
Consistency C	$V \cdot C$, cross-model sign consistency of the full design	0.025	0.020	–
Consistency C_N	$V \cdot C_N$, original-model half of C only	0.125	0.081	–
Significance gate	BH-FDR $q < 0.05$ on the discovery contrasts before top-N	selects the same ten features as the baseline		
Time bins	per-time-bin feature scoring, two bin settings	0.075 / 0.050	0.069 / 0.037	–
Mask only	zero the support in $\mathcal{M}^E$, inject nothing	0.000	0.002	–
Persist scoring	$V = AR$ (the null control of every cell)	0.000	0.004	–
Strength grid	α extended from 8 to 32	optimum unchanged at $\alpha = 4$; the wider grid was kept for the matrix		

Table 20: Non-target probe: the cat/ESD-x circuit applied at its selected strength ($\alpha = 8$) to 40 dog prompts and 40 truck prompts. Retention is the fraction of images in which the prompted object survives its own strict threshold; bleed is the fraction of images the ML-Decoder cat class scores above 0.556. Paired tests compare arms on the same prompt and seed.

Probe	Arm	Retention (mean score)	Retained	Cat bleed (mean)	Cat present
dog (the p^- concept of the cat pairs)	Original model	0.993	40/40	0.259	1/40
	Erased model (ESD-x cat)	0.835	32/40	0.263	2/40
	Random control	0.884	35/40	0.293	3/40
	Ours	0.578	20/40	0.787	29/40
truck (unrelated to the pairs)	Original model	0.975	40/40	0.192	0/40
	Erased model (ESD-x cat)	0.969	40/40	0.190	0/40
	Random control	0.760	25/40	0.224	0/40
	Ours	0.709	22/40	0.357	5/40

dog: retention Ours vs. erased $p = 2.6 \times 10^{-4}$, random vs. erased $p = 0.88$; bleed Ours vs. random $p = 7.8 \times 10^{-9}$. truck: retention Ours vs. random $p = 0.34$; bleed Ours vs. random $p = 3.5 \times 10^{-7}$.

0.15): more coordinates dilute the injected mass without adding a direction. *Hookpoint count.* Five hookpoints instead of three add the fourth- and fifth-ranked blocks, both residual blocks, and do not help ($p = 0.21$); recovery concentrates on the cross-attention path that the edit itself touched. *Prototype statistic.* The median contrast beat the mean on validation by a single unit and was adopted; on the test split it lost (strict rate $4/40 \to 1/40$ on cat/UCE, unchanged on cat/ESD-x) and the mean was reinstated. This is the one rung that cost a test read and the reason the adoption rule now requires a paired test. *Feature-selection extras.* Multiplying V by the cross-model sign-consistency term C of the full design selects backwards under the suppressed hypothesis, since its erased-model half is itself a persistence assumption, and its original-model-only half brings nothing; the BH-FDR significance gate selects the same ten features and is kept as a reported statistic (Appendix B.2). *Time-binned scoring.* Scoring features per time bin, in two bin settings, is worse than pooling over steps. *Mask-only injection.* Zeroing the erased model's own codes on the support instead of adding μ_N recovers nothing: the injected values, not the masking, carry the effect. *Persist scoring.* Selecting by $V = AR$ recovers nothing on this cell, and is retained as the comparison on all 67 contrast audits.

J.2 Strength grid, SAE size and hookpoint criterion

Strength grid and selection rule. The grid $\{0.5, \ldots, 8\}$ was extended to 32 on this cell (validation pass rate 0.50 at $\alpha = 4$, 0.03 at 16 and 0.40 at 32, so the optimum stayed inside the grid) and kept at 32 for the matrix after the first style cell peaked at 8; no cell of the reported matrix read its test split before the grid was widened. The selection rule was changed once, from the pass rate at the calibrated primary threshold to the mean detector score: the primary threshold is degenerate for softmax detectors (for cat, $\tau_c^E \approx 0.003$), so a rate at it rewards many images barely off the floor; it pushed the top-15 and top-20 rungs to the top of the grid, where their strict rate was 0, whereas at $\alpha = 4$ the same circuits score as in Table 19. The mean score has no threshold and does not tie at $n = 40$. Dose-response curves of five cells, the fraction of audits that select the top of the grid and the time-window ablation on every audit are in Appendix M.2.

SAE size. $K = 32$ and width $4d_\ell$ were fixed a priori and not swept in the matrix pipeline. The earlier full-design variant of the pipeline trained $K \in \{48, 64\}$ at widths $8d_\ell$ and $16d_\ell$ for 20k–50k

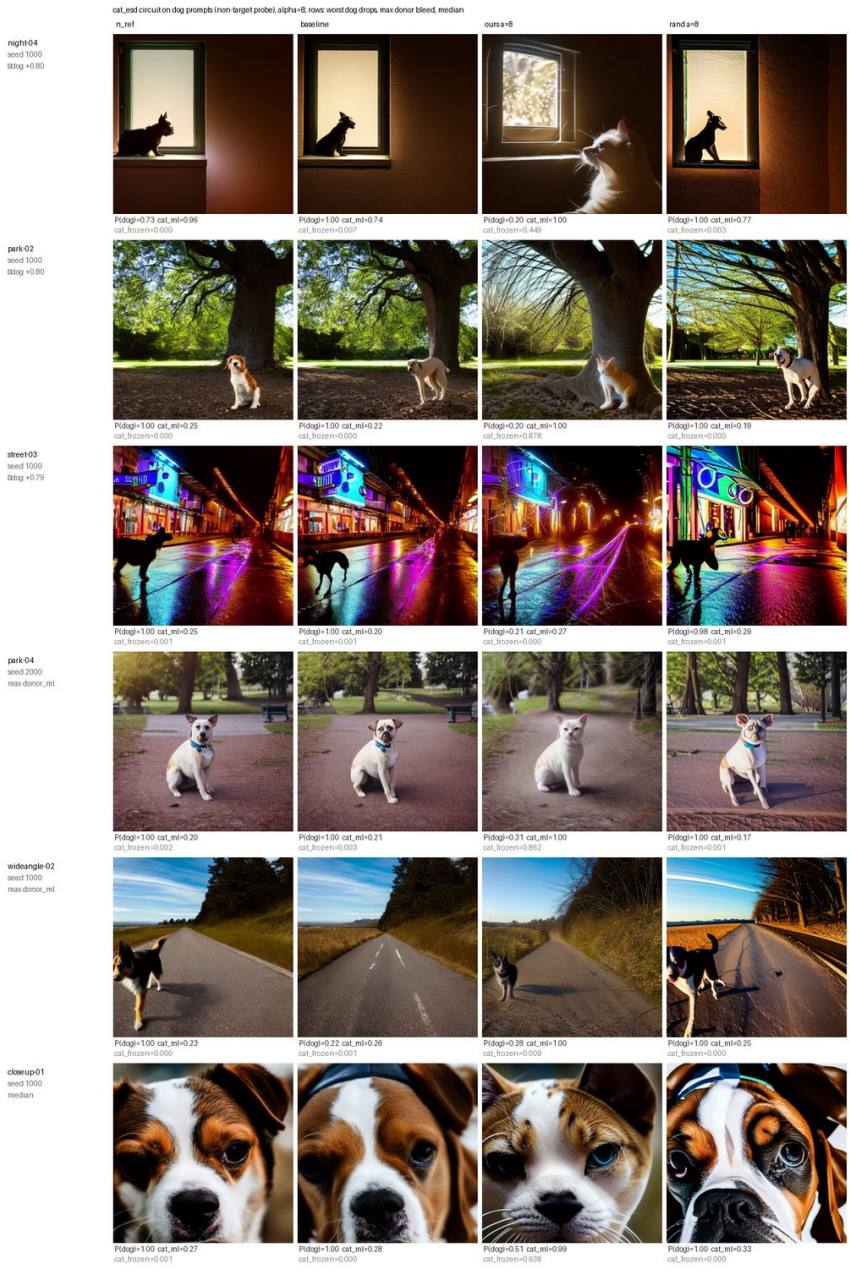


Figure 8: The cat/ESD-x circuit applied to dog prompts (rows: the three largest dog-retention drops, the two largest cat-bleed scores, and the median unit; columns: original model, erased model, Ours at $\alpha = 8$, random control at $\alpha = 8$). Numbers under each image are the dog-retention and cat-bleed scores.

steps at the deep $1{,}280$-channel blocks; held-out EV moved only from about 0.70 to 0.78, and the single-cell EV ladder of Appendix D found recovery with the residual term unchanged across EV 0.47–0.78 on cat/UCE. Reconstruction quality is therefore reported (Table 15) but was not a tuning target.

Hookpoint criterion. The initial matrix used contrast. Six early style and copyright reruns explored rel, with five proceeding to a test evaluation under the rule in Appendix C.3. The current comparison evaluates both criteria on all 67 eligible cells. Each criterion selects its hookpoints from discovery data, but changing the criterion rebuilds the dictionaries, features and prototype and

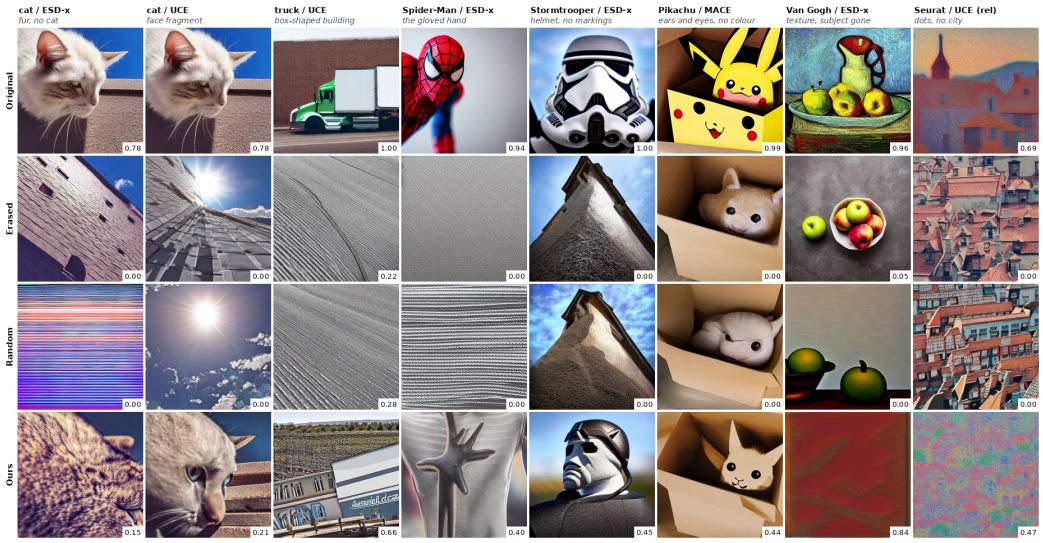


Figure 9: Held-out units on which *ours* stays below the strict threshold, one per failure mode (same prompt and seed per column; numbers are detector scores; the strict thresholds are 0.25 for cat, 0.7 for truck, 0.5 for Spider-Man, Stormtrooper, Pikachu and Seurat and 0.951 for Van Gogh). Columns: cat fur without a cat and a face fragment (cat/ESD-x, cat/UCE); a box-shaped building for *truck* (truck/UCE); the gloved hand of Spider-Man and the helmet shape of the Stormtrooper without its markings (both ESD-x); Pikachu's ears and eyes on the white animal the erased model draws instead (MACE); Van Gogh brushstroke texture with the still life removed (ESD-x, contrast hookpoints); pointillist speckle with the city gone (Seurat/UCE, rel hookpoints).

reselects strength on validation. Table 1 summarizes the best observed result for each cell across the two ranking criteria; the caption specifies how the displayed result is chosen.

K Non-Target Collateral Probe

An earlier, pilot-pipeline version of the cat/ESD-x circuit (three blocks, strength 8 selected on cat's validation split, nothing re-tuned; the matrix circuit of Table 13 runs at $\alpha = 32$) was applied to 40 dog prompts and 40 truck prompts, sampled across their scene families from units on which the original model draws the prompted object. Each image was scored for retention of the prompted object with the ML-Decoder class of that object and for cat bleed with two heads, the frozen ImageNet softmax and the ML-Decoder cat class (Table 20; Figure 8). On dog prompts, dog being the negative concept of every cat pair, the prototype both removes the dog and paints a cat; the random control at the same strength does neither. On truck prompts, unrelated to the pair construction, *ours* and the *rand* control damage the truck equally and only *ours* adds cats, in five images concentrated on large single-object close-ups. The two costs therefore have different mechanisms: damage scales with the displacement magnitude, bleed follows the selected direction and is strongest on semantically adjacent content.

L Failure Cases

Figure 9 shows one unit per failure mode, chosen among the test units whose *ours* image scores below the strict threshold but above the erased model (the per-cell score bands are in Table 13). Two modes account for nearly all such units.

Part without the whole. On objects and copyright concepts the injection re-creates a local attribute of the concept without its layout. For *cat* the recurring partial is fur: the fur texture fills the frame, or a single eye and a nose appear at the edge of the image, where the erased model had drawn a wall or a sky, and the detector's cat probability rises from 0 to 0.1–0.2 against a threshold of 0.25. For

characters the fragment is a pose or a shape: a gloved hand in Spider-Man's web-shooting pose on a grey suit with none of the red and blue; a helmet-shaped head for the Stormtrooper without the white armour and black markings; Pikachu's pointed ears, black eyes and cheek marks added to the white animal that the erased model draws in place of the character, with the yellow coat missing. On the ASICS logo the injection restores posters with logo-like lettering in the brand's colours but never the mark itself. Trucks are the object-domain version of the same failure: under UCE and ESD-u the sub-threshold images contain box-shaped buildings, billboards and piles of pipes and wheels, i.e., the rectangular geometry and the part textures of a truck without a truck. When the erased model draws a substitute rather than an empty scene (the white animal for Pikachu, a person without a bicycle), the fragment is added to the substitute, and a detector that must choose between the concept and its sibling scores the hybrid low.

Texture without composition. On styles the injection at the contrast hookpoints, which sit in the first residual block, paints the style's surface statistics over the whole image and removes the subject: for Van Gogh the still life becomes a field of brushstroke texture, sometimes a nearly uniform orange field, which the five-way CLIP style head scores at 0.84–0.93 by elimination (threshold 0.951; the flat fields are the score-only case of Appendix G.1). At the rel hookpoints the Seurat injection keeps the composition on most units and adds a pointillist surface, which is what the relabel check of Appendix C.3 measures; on its weakest units only the speckle survives and the city is gone. Kirchner's sub-threshold images show bold black outlines and saturated flat colour, single cues of the style, without its forms.

Reading. Eq. (10) broadcasts a fixed direction across positions and steps rather than providing a spatially resolved target. The examples above suggest that such an intervention can restore concept-associated texture or parts without consistently recovering their arrangement. At high strengths, repeated structures can dominate the image, as in the bicycle recoveries, whose frames displace the prompted person (Appendix H.2). They do not show that the selected features exclusively encode texture, or that spatial organization or color is controlled by a single component. The random control illustrates a separate effect of intervention strength: at $\alpha = 32$, the striped cat/ESD-x control damages the image without adding the target concept. Recovery, scene consistency and concept identity therefore need separate interpretation.

M Artifact and Out-of-Distribution Gates, Dose and Timing

M.1 Artifact and out-of-distribution gates

The following artifact and activation-range diagnostics retain their historical 72-audit scope (67 contrast audits plus five early rel reruns); their aggregate counts have not been extended to the current 134-audit comparison. The *artifact gate* scores an image 0 when its grey-level standard deviation is below 8 on the 0–255 scale (an all-black or flat image); it changed the strict rate of no in-house arm except the noise probe's best-of-24, where all-black images had been scored above threshold by a CLIP head. Over the test images of *ours* in the three reads of the 72 historical audits (suppressed and persist scoring and the baseline-protocol regeneration), 106 of $5{,}958$ (1.8%) were flagged, all but two of them in Van Gogh, Kirchner and Seurat cells, against 26 of 3,972 random-control images and 4 of $1{,}986$ images of the erased model itself. No flagged image of the original model, the erased model, *ours* or *rand* control crosses the strict threshold in any main read, so the gate changes no entry of Table 11; the recovery gap of Table 13 is computed on gated scores. The *out-of-distribution gate* computes, for every activation-space arm at its selected strength, the fraction of injected activation dimensions that fall outside the $[0.1, 99.9]$ percentile envelope of the erased model's own unsteered activations on the test units, next to the same fraction for the original model's concept activations. *ours* leaves the envelope on 0.42% of dimensions on average (maximum 3.5%), the unsteered erased model on 0.22% and the original model's target activations on 0.32%; the highest values are the highest-activation arm (1.0% on average, up to 5.6%) and Reverse-SAeUron (0.8%, up to 10.7%). A cell-arm is flagged when it exceeds twice the original-model level; *ours* is flagged on 7 of the 72 audits, six of them by absolute margins below one percentage point and one, Spider-Man/ESD-u, by 2.7 points. These marginal range checks do not establish that the joint activation distribution remains in-distribution or that image content is preserved.

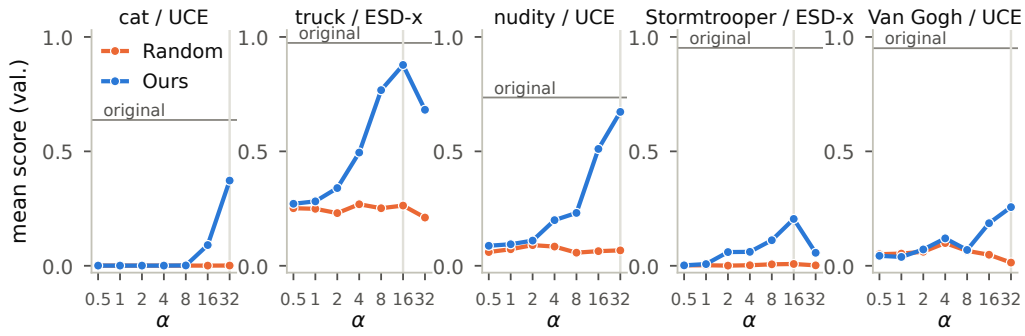


Figure 10: Validation dose-response of the mean evaluator score against the injection strength α for *ours* and the random control on five cells; the vertical line marks the strength selected for the test read and the horizontal line the original model's mean score on the same validation units.

M.2 Dose and timing

In the same historical 72-audit set, Figure 10 shows the validation dose-response on five cells. The selected strength is 16 or 32 in 51 of 72 audits and at least 8 in every recovered cell; 31 audits, 15 of them recovered, select the top of the grid ($\alpha = 32$), so there the grid rather than the dose-response bounds the strength. Because Eq. (10) makes the intervention a fixed displacement, the value of α alone does not determine the relative perturbation: across the 72 audits the injected activations at the selected strengths leave the erased model's own per-dimension envelope on 0.4% of dimensions on average (at most 3.5%), against 0.2% for the unsteered erased model and 0.3% for the original model's own concept activations. On the same audits, injecting only during the first 25 of 50 steps retains most of the effect (mean strict ASR 0.19 versus 0.24; at least as good as the full window in 44 of 72 cells), whereas the second half alone retains little (0.09); Figure 3(c) shows the 67 contrast audits, a subset of this diagnostic set. The order reverses on 12 audits: on six of them both halves stay at or below 0.1, and the others are cat/SPEED (0.66 against 0.12), bicycle/UCE (0.35 against 0.24) and four Seurat audits, most sharply Seurat/ESD-x (rel), where the second half alone gives 0.88 and the first half 0.06 (six of the seven Seurat audits reverse); the timing effect is therefore concept- and configuration-dependent.