

CONTINUAL CONCEPT ERASURE IN DIFFUSION MODELS BY SUPPRESSING CROSS-EDIT INTERFERENCE

Yongliang Wu¹, Haori Lu¹, Jinqi Luo², Wei Cao¹, Xingyu Zhu^{3*}, Yaoyao Liu¹

¹ University of Illinois Urbana-Champaign

² University of Pennsylvania

³ National University of Singapore

ABSTRACT

Concept erasure removes copyright-protected, privacy-sensitive, or otherwise undesirable concepts from pretrained text-to-image diffusion models to support content governance and compliance. As erasure requests arrive over time, models must remove new targets without undoing prior erasures. Existing methods do not constrain interference across edits: residual perturbations outside the retain set interact and accumulate, degrading unrelated generations and sometimes collapsing previously erased targets into noise. We propose **CEASE** (Continual Erasure via Adaptive Subspace Editing), a training-free method that imposes two subspace constraints on a closed-form solver. CEASE adds the token representation of the shared replacement to the solver’s invariance matrix and, when interference is detected, projects the current update onto the orthogonal complement of dominant output directions extracted from cumulative past updates. A closed-form decomposition attributes the accumulated interference to repeated activation of the shared replacement and overlap between successive update directions, showing that the two constraints suppress these respective sources. Across continual erasure of celebrities, artistic styles, and instances, CEASE achieves the most consistent erase-preserve trade-off, while existing methods either degrade general generation or insufficiently erase targets.¹

1 INTRODUCTION

Concept erasure aims to remove copyright-protected, privacy-sensitive, or otherwise undesirable concepts from pretrained text-to-image (T2I) diffusion models, making it an important tool for content governance and compliance (Gandikota et al., 2023; Lu et al., 2024; Liu et al., 2026; 2025). In practice, new erasure targets arrive over time (Lee et al., 2026; George et al., 2025), extending concept erasure from one-shot to continual editing: the model must remove new targets without allowing later updates to undo prior erasures or degrade generation of unrelated concepts.

Prior studies report that retained model utility declines when erasure requests are processed sequentially (Lee et al., 2026; George et al., 2025). In our experiments, sequentially applying standard erasure methods (Gandikota et al., 2024; Lu et al., 2024; Gong et al., 2024; Li et al., 2026) reveals two related failure modes. First, general generation quality declines on prompts unrelated to any erased concept. Second, outputs for some erased targets become incoherent instead of reflecting the prescribed replacement (e.g. “a person” for a celebrity). Each edit explicitly protects only a finite retain set (Gandikota et al., 2024; Gong et al., 2024; Li et al., 2026) and can therefore perturb embeddings outside it (Wuerkaixi et al., 2025; Thakral et al., 2025).

These perturbations interact and accumulate across successive updates, making it difficult for existing methods (Li et al., 2026; Gong et al., 2024; Gandikota et al., 2024) to confine the effects of erasure to the target concepts. As shown in Fig. 1(a), COCO FID on unrelated prompts rises as the number of edits increases (Lin et al., 2014). This exposes a limitation of existing preservation constraints in the continual setting: protecting a finite retain set at each step does not reliably prevent

*Corresponding author

¹Project page: <https://continual-erasure.cvmlgroup.web.illinois.edu/>




Figure 1: **Utility loss under continual celebrity erasure.** (a) COCO FID on 1,000 unrelated prompts as the number of erased concepts grows from 10 to 100; lower is better. (b) COCO CLIP score on the same prompts after 100 edits; the dashed line marks the unedited model. (c) Generations after 100 edits for an unrelated COCO scene (top) and the erased target “Tom Cruise” (bottom).

subsequent updates from cumulatively degrading general generation. Beyond image quality, this degradation also weakens the model’s ability to follow prompt semantics: Figure 1(b) shows that, after 100 edits, the CLIP score on the same prompts falls well below that of the unedited model. This reduced ability to follow prompt semantics may also affect the outcomes of previous erasures, because the replacement outputs for earlier targets depend on the same continually updated weights. Figure 1(c) illustrates both effects in the baseline models: they corrupt an unrelated scene and produce noise for the erased celebrity, failing to generate the prescribed generic portrait.

To suppress the accumulation of interference in continual concept erasure, we propose **CEASE** (Continual Erasure via Adaptive Subspace Editing), a training-free method based on two subspace constraints. CEASE adds the token representation of the shared replacement to the invariance matrix of the closed-form solver, preventing subsequent updates from modifying this representation. It also extracts the dominant output directions from cumulative past updates and, when interference is detected, projects the current update onto their orthogonal complement. Theoretically, a closed-form decomposition attributes the accumulated interference to repeated activation of the shared replacement and overlap between successive update directions, and shows that the two designs suppress the respective sources. Empirically, CEASE maintains stable COCO FID throughout the edit sequence and a COCO CLIP score close to the unedited reference after 100 edits, as shown in Fig. 1(a) and (b). Across continual erasure tasks involving celebrities, artistic styles, and instances, CEASE achieves the most consistent erase–preserve trade-off. On style erasure, CEASE reduces COCO FID by 73% relative to the strongest prior method. Our contributions are summarized as follows:

- We formalize how interference accumulates across edits in continual concept erasure, and derive a closed-form decomposition that attributes it to repeated activation of the shared replacement and overlap between successive update directions.
- We introduce two training-free subspace constraints: one keeps the token representation of the shared replacement invariant, while the other steers the current update away from the dominant output directions of past edits, thereby suppressing the two interference sources.
- We evaluate continual erasure of celebrities, artistic styles, and instances across multiple diffusion backbones and assess both erasure and preservation; CEASE achieves the best erase–preserve balance among the compared methods while keeping retained concepts recognizable at rates close to the reference model.

2 RELATED WORK

Concept Erasure in Diffusion Models. Text-to-image diffusion models (Zhang et al., 2023a; Ma et al., 2024; Ji et al., 2026) trained on uncensored web data can reproduce copyrighted styles, recognizable identities, and other undesired content, and concept erasure removes such targets from a pretrained model after training. Existing methods fall into three families. Training-based methods fine-tune the model with a concept-suppression objective and a preservation regularizer (Gandikota et al., 2023; Zhang et al., 2024; Kumari et al., 2023; Huang et al., 2024; Lyu et al., 2024; Li et al., 2024; Kim et al., 2024; Zhao et al., 2024; Wu et al., 2025), with a separate strand hardening the

erasure against adversarial prompts (Fan et al., 2024; Heng & Soh, 2023; Pham et al., 2024; Park et al., 2024); their common cost is that each new target requires another fine-tuning loop. Editing-based methods instead update cross-attention projections through a closed-form least-squares solution, from TIME (Orgad et al., 2023) and UCE (Gandikota et al., 2024) to MACE (Lu et al., 2024), RECE (Gong et al., 2024), and SPEED (Li et al., 2026); they edit the model in seconds, but their preservation depends on a finite retain set. Non-parametric methods leave the weights untouched and intervene at sampling time (Schramowski et al., 2023; Das et al., 2024; Rando et al., 2022). All three families are typically evaluated with the full target set known before editing; our study instead evaluates repeated application to an already-edited model.

Sequential Unlearning. Sequential unlearning is closely related to continual learning and online learning (Liu et al., 2020; 2021a;b; 2023a;b; 2024; Li et al., 2019; 2020; 2021; Luo et al., 2023; Fischer et al., 2024; Zhang et al., 2023b). While continual learning and online learning typically focus on acquiring new knowledge without forgetting, sequential unlearning poses the complementary challenge of repeatedly removing targeted knowledge without causing unintended forgetting of retained concepts. In deployment, takedown demands and right-to-be-forgotten obligations arrive as a stream. For language models, the locate-then-edit paradigm (Meng et al., 2022a;b) degrades under repeated application, and mitigations include drift-restraining regularizers (Ma et al., 2025; Gu et al., 2024), external edit memories (Hartvigsen et al., 2023; Wang et al., 2024b), and projections onto the null space of preserved keys (Fang et al., 2025); these approaches operate on factual associations, where edits do not share a common redirection target. For diffusion models, Selective Amnesia (Heng & Soh, 2023) first cast concept forgetting as a continual-learning problem, and later work adds concept-aware masking with distillation (Li et al., 2025), gradient projection against semantically related concepts (Lee et al., 2026), and multi-objective distillation across deletion requests (George et al., 2025); these methods suppress the degradation as a monolithic quantity without analyzing where it comes from. In contrast, we approach continual erasure from the perspective of cross-edit interference: we trace the failure to the interference that successive edits accumulate on unprotected prompts, and design our method to suppress this interference.

3 ANALYSIS AND METHOD

Section 3.1 defines the continual erasure problem and the base closed-form update. Section 3.2 then explains why update noise accumulates under sequential editing. Guided by the two sources identified in this analysis, Sec. 3.3 presents CEASE.

3.1 PROBLEM SETUP

Concept erasure for T2I diffusion models edits the cross-attention value projection $\mathbf{W} \in \mathbb{R}^{d_1 \times d_0}$, which maps a frozen CLIP text embedding $\mathbf{c} \in \mathbb{R}^{d_0}$ to the value vector $\mathbf{W}\mathbf{c} \in \mathbb{R}^{d_1}$, via an additive update $\mathbf{W}' = \mathbf{W} + \Delta$. We specify an erasure set $\mathbf{E} = \{\mathbf{c}_1^{(i)}\}_{i=1}^{N_E}$ to remove and a retain set $\mathbf{R} = \{\mathbf{c}_0^{(i)}\}_{i=1}^{N_R}$ to preserve, pair each target with an anchor $\mathbf{c}_*^{(i)}$ that keeps generic semantics but discards target-specific information, and stack these column-wise as $\mathbf{C}_1, \mathbf{C}_* \in \mathbb{R}^{d_0 \times N_E}$ and $\mathbf{C}_0 \in \mathbb{R}^{d_0 \times N_R}$.

Base Closed-Form Update. Closed-form methods (Gandikota et al., 2024; Li et al., 2026) choose Δ by solving a weighted least-squares problem that balances two requirements:

$$\Delta_{\text{erase}} = \arg \min_{\Delta} \underbrace{\|(\mathbf{W} + \Delta)\mathbf{C}_1 - \mathbf{W}\mathbf{C}_*\|^2}_{e_1: \text{redirect targets}} + \underbrace{\|\Delta\mathbf{C}_0\|^2}_{e_0: \text{preserve non-targets}}. \quad (1)$$

e_1 redirects each target’s output toward its anchor value $\mathbf{W}\mathbf{c}_*^{(i)}$, while e_0 preserves non-target outputs. Because both terms depend on the same update, adding targets can increase the preservation residual (Li et al., 2026). Following SPEED, the base solver reduces this residual by right-projecting the update through a near-null projector $\mathbf{P}$ of $\mathbf{C}_0^\top$. The applied update $\Delta^*\mathbf{P}$ therefore satisfies $(\Delta^*\mathbf{P})\mathbf{C}_0 \approx \mathbf{0}$, but this guarantee covers only the span of the finite retain set. Appendix A.2 provides the projector construction.

Continual Concept Erasure. The continual setting replaces the joint solve with N updates applied one after another. Starting from the pretrained weights, each step solves for a new update against

the running weights and applies it:

$$\mathbf{W}^{(0)} = \mathbf{W}, \quad \mathbf{W}^{(i)} = \mathbf{W}^{(i-1)} + \Delta_i, \quad \Delta_{\text{sum}} = \sum_{i=1}^N \Delta_i, \quad (2)$$

where step i handles a single target and its anchor, with the retain set held fixed throughout.

3.2 WHY UPDATE NOISE ACCUMULATES

The base solver protects only embeddings in the span of $\mathbf{C}_0$. Each edit can therefore introduce an unintended change to the value-projection output of an uncovered prompt, and these changes accumulate over the sequence.

Measuring Accumulated Update Noise. We quantify the accumulated change directly at the value-projection output. Consider an arbitrary test embedding $\mathbf{c}$ outside the retain set. After all N edits, the cross-attention value becomes $(\mathbf{W} + \sum_{i \leq N} \Delta_i)\mathbf{c}$, whose magnitude is bounded by the triangle inequality:

$$\left\| (\mathbf{W} + \sum_{i \leq N} \Delta_i)\mathbf{c} \right\|_2 \leq \|\mathbf{W}\mathbf{c}\|_2 + \left\| \sum_{i \leq N} \Delta_i \mathbf{c} \right\|_2. \quad (3)$$

The first term $\|\mathbf{W}\mathbf{c}\|_2$ is fixed, so the change induced by editing is governed by $\|\sum_i \Delta_i \mathbf{c}\|_2$. Let e denote the single edit intended to affect $\mathbf{c}$; for a prompt unrelated to every erased target, we set $\Delta_e \mathbf{c} = \mathbf{0}$. The desired reference is therefore $\|\Delta_e \mathbf{c}\|_2^2$. We define the accumulated update noise as the difference between the squared magnitude of the full update and this reference:

$$\text{noise}_e(\mathbf{c}) = \left\| \sum_{i \leq N} \Delta_i \mathbf{c} \right\|_2^2 - \|\Delta_e \mathbf{c}\|_2^2, \quad (4)$$

a signed residual over the sequence. Positive values indicate that the full sequence produces a larger update than the intended edit alone, whereas negative values indicate cancellation among updates.

To identify why this noise grows, we use the rank-at-most-one structure of the single-target, single-token update. At step i , the closed-form solver acts on one target $\mathbf{c}_1^{(i)}$ and anchor $\mathbf{c}_*^{(i)}$; left- and right-multiplication by its projection matrices preserves the rank bound of the target outer product. We may therefore write $\Delta_i = \alpha_i \beta_i^\top$, where α_i is the output-space influence direction and β_i determines how strongly an arbitrary $\mathbf{c}$ activates the edit, since $\Delta_i \mathbf{c} = (\mathbf{c}^\top \beta_i) \alpha_i$. Substituting this into Eq. 4 gives:

$$\text{noise}_e(\mathbf{c}) = \sum_{\substack{i,j \\ (i,j) \neq (e,e)}} (\mathbf{c}^\top \beta_i) (\alpha_i^\top \alpha_j) (\beta_j^\top \mathbf{c}). \quad (5)$$

Decomposing Accumulated Update Noise. We split any uncovered prompt into its component along the anchor and its orthogonal remainder:

$$\mathbf{c} = s \mathbf{c}_* + \mathbf{c}_\perp, \quad s = \frac{\langle \mathbf{c}_*, \mathbf{c} \rangle}{\|\mathbf{c}_*\|_2^2}, \quad \mathbf{c}_\perp \perp \mathbf{c}_*. \quad (6)$$

Each activation then splits into an anchor part and an off-anchor part:

$$\beta_i^\top \mathbf{c} = s (\beta_i^\top \mathbf{c}_*) + \beta_i^\top \mathbf{c}_\perp. \quad (7)$$

For compactness, let $a_i = \beta_i^\top \mathbf{c}_*$, $r_i = \beta_i^\top \mathbf{c}_\perp$, and $g_{ij} = \langle \alpha_i, \alpha_j \rangle$, and let $\sum'_{i,j}$ exclude $(i,j) = (e,e)$. Substituting Eq. 7 into Eq. 5 and combining the two symmetric cross terms gives:

$$\begin{aligned} \text{noise}_e(\mathbf{c}) &= \mathcal{N}_{\text{anc}}(\mathbf{c}) + \mathcal{N}_{\text{off}}(\mathbf{c}) + \mathcal{N}_{\text{cross}}(\mathbf{c}), \\ \mathcal{N}_{\text{anc}}(\mathbf{c}) &= s^2 \sum'_{i,j} a_i a_j g_{ij}, \quad \mathcal{N}_{\text{off}}(\mathbf{c}) = \sum'_{i,j} r_i r_j g_{ij}, \quad \mathcal{N}_{\text{cross}}(\mathbf{c}) = 2s \sum'_{i,j} a_i r_j g_{ij}. \end{aligned} \quad (8)$$

The anchor channel $\mathcal{N}_{\text{anc}}$ is driven by the anchor activations $\beta_i^\top \mathbf{c}_*$: every edit uses the same fixed direction $\mathbf{c}_*$, so these contributions stack across steps for prompts with a nonzero anchor component s . The off-anchor channel $\mathcal{N}_{\text{off}}$ is driven by the off-anchor activations and the influence overlap $\alpha_i^\top \alpha_j$. The cross channel $\mathcal{N}_{\text{cross}}$ couples the two via the anchor activations. $\mathcal{N}_{\text{anc}}$ and $\mathcal{N}_{\text{cross}}$ both vanish once $\beta_i^\top \mathbf{c}_* = 0$, while the terms coupling different edits in $\mathcal{N}_{\text{off}}$ vanish once $\alpha_i^\top \alpha_j = 0$ for $i \neq j$. These conditions motivate the AIC and HOC constraints of CEASE.




Figure 2: Overview of CEASE for continual concept erasure. AIC preserves the shared anchor during the closed-form solve. HOC extracts dominant output directions from cumulative historical edits via SVD and, when needed, projects the new update onto their orthogonal complement before applying it to the current model.

3.3 CEASE: SUPPRESSING UPDATE NOISE

The decomposition in Eq. 8 identifies two ideal conditions under which only the $O(N)$ diagonal remains: $\beta_i^\top c_* = 0$ for all i , and $\alpha_i^\top \alpha_j = 0$ for $i \neq j$. CEASE targets these conditions with two components, Anchor Invariance and Historical Orthogonality. As illustrated in Fig. 2, AIC protects the shared anchor while HOC projects each new update away from dominant historical edit directions. Both reuse the existing closed-form update and require no retraining.

Anchor Invariance Constraint (AIC). The SPEED solver already contains an invariance matrix K_2 formed from the empty-prompt representation, but its default construction does not include the task anchor c_* . AIC appends the anchor token embedding to this matrix:

$$K_2^{\text{AIC}} = [K_2, c_*], \quad \Delta_i K_2^{\text{AIC}} = 0. \quad (9)$$

With the zero invariance regularizer used in our experiments, the closed-form correction enforces $\Delta_i c_* = 0$ up to numerical precision. For the rank-at-most-one factorization, this implies $\beta_i^\top c_* = 0$ whenever $\alpha_i \neq 0$. Under this ideal constraint, the anchor and cross groups vanish:

$$\mathcal{N}_{\text{anc}}(c) = 0, \quad \mathcal{N}_{\text{cross}}(c) = 0, \quad \forall c, \quad (10)$$

leaving the off-anchor interaction for the next component to address.

Historical Orthogonality Constraint (HOC). Under the ideal AIC constraint in Eq. 10, only the off-anchor channel $\mathcal{N}_{\text{off}}$ remains. Separating it by step index:

$$\mathcal{N}_{\text{off}}(c) = \underbrace{\sum_{i \neq e} (\beta_i^\top c_\perp)^2 \|\alpha_i\|^2}_{\text{diagonal, } O(N)} + \underbrace{\sum_{i \neq j} (\beta_i^\top c_\perp) (\alpha_i^\top \alpha_j) (\beta_j^\top c_\perp)}_{\text{cross, } O(N^2)}, \quad (11)$$

The diagonal terms capture single-step effects, while the cross terms scale with the influence overlap $\alpha_i^\top \alpha_j$ between edits. HOC reduces this overlap when the accumulated historical update substantially affects the current target, thereby mitigating interference between successive edits.

A naive way to enforce this orthogonality would keep all past influence vectors $\{\alpha_j\}_{j < i}$ in memory and project each new α_i onto their orthogonal complement, which grows linearly in N . We avoid this overhead by exploiting the rank-1 structure of every closed-form update. Since $\Delta_j = \alpha_j \beta_j^\top$, the cumulative historical edit is:

$$\Delta_{\text{hist}} = \sum_{j < i} \Delta_j = \sum_{j < i} \alpha_j \beta_j^\top \quad (12)$$

Its column space is contained in $\text{span}\{\alpha_j\}_{j < i}$. It therefore provides a compact, observable proxy for historical output directions, although cancellations can make it a strict subset. We obtain its dominant left singular directions from:

$$\{U_{<i}, \Lambda_{<i}, U_{<i}^\top\} = \text{SVD}(\Delta_{\text{hist}} \Delta_{\text{hist}}^\top). \quad (13)$$

Table 1: Celebrity erasure after continually erasing 100 celebrities on SD v1.4. Bold and underline mark the best and second-best edited results.

	Erase set				Retain set				MS-COCO		
	CS ↓	FID ↑	LPIPS ↑	Acc ↓	CS ↑	FID ↓	LPIPS ↓	Acc ↑	CS ↑	FID ↓	LPIPS ↓
SD v1.4	30.95	0.0	0.000	88.0	31.66	0.0	0.000	70.6	26.52	0.0	0.000
ESD	21.64	61.6	0.545	23.0	21.31	59.3	0.544	12.9	21.96	61.5	0.540
UCE	13.55	272.8	0.768	0.0	14.10	242.7	0.768	0.0	13.15	241.4	0.773
ConAbl	15.79	74.1	0.540	0.0	15.85	73.0	0.537	0.0	24.26	50.7	0.476
RECE	16.51	568.3	0.974	0.0	16.27	573.1	0.972	0.0	16.28	616.5	1.002
SPEED	18.38	137.2	0.635	<u>0.3</u>	<u>30.28</u>	<u>20.8</u>	<u>0.383</u>	<u>54.4</u>	22.14	107.2	0.566
CUIG-GradProj	<u>15.28</u>	<u>333.8</u>	<u>0.846</u>	0.0	15.03	336.0	0.844	0.0	14.15	303.9	0.849
Sculpting Memory	16.84	101.5	0.606	1.0	17.31	93.2	0.597	0.3	<u>25.26</u>	<u>47.5</u>	<u>0.429</u>
Ours	21.54	39.9	0.478	3.5	31.59	12.3	0.271	65.9	26.24	46.2	0.388

Using the dominant left singular vectors $\hat{\mathbf{U}}_{<i}$ from Eq. 13, HOC projects the AIC-constrained update $\tilde{\Delta}_i$ onto the orthogonal complement of the historical update subspace. Accordingly, $\Pi_i = \mathbf{I} - \hat{\mathbf{U}}_{<i} \hat{\mathbf{U}}_{<i}^\top$ when HOC is applied, and $\Pi_i = \mathbf{I}$ otherwise. CEASE updates the model as:

$$\Delta_i = \Pi_i \tilde{\Delta}_i, \quad \mathbf{W}^{(i)} = \mathbf{W}^{(i-1)} + \Delta_i. \quad (14)$$

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Baselines and Benchmarks. We evaluate on three concept erasure benchmarks adapted from MACE (Lu et al., 2024) and SPEED (Li et al., 2026), covering celebrities, artistic styles, and object or character instances. Each benchmark pairs 100 erased concepts with 100 disjoint retained concepts and uses five prompt templates per concept. We compare two groups of baselines: standard erasure methods, comprising the closed-form SPEED (Li et al., 2026), UCE (Gandikota et al., 2024), and RECE (Gong et al., 2024) and the training-based ESD (Gandikota et al., 2023) and ConAbl (Kumari et al., 2023), and continual-unlearning methods, comprising CUIG-GradProj (Lee et al., 2026) and Sculpting Memory (Li et al., 2025).

Metrics. We report CLIP score (CS) (Hessel et al., 2021) for image-prompt alignment, and FID (Heusel et al., 2017) and LPIPS (Zhang et al., 2018) for the distribution-level and image-level distance from the unedited backbone’s generations, on each of the erase set, the retain set, and 1,000 MS-COCO captions. On the erase and retain sets we add one task-matched fourth metric. For celebrities, Acc is the recognition accuracy of an open-ended VLM judge (Qwen2-VL-7B) (Wang et al., 2024a) asked to name the person in the image; for instances, Acc is the CLIP zero-shot top-1 accuracy against the 200 concept candidates; both should drop on the erase set and stay high on the retain set. For styles, the LAION aesthetic score (Aes) (Schuhmann et al., 2022) is a no-reference quality rating on a 1–10 scale that should stay near the unedited reference.

Implementation Details. The main experiments use Stable Diffusion v1.4 with DPM-Solver sampling and classifier-free guidance 7.5 (Rombach et al., 2021). The per-task anchors are “person” for celebrities, “art” for styles, and “object” for instances. All experiments run on a single H100 GPU. Additional implementation details, including hyperparameters, are provided in Appendix A.1.

4.2 MAIN RESULTS

We compare CEASE with all baselines after the full 100-step continual erasure sequence, and report the quantitative results on the celebrity, style, and instance benchmarks in Tables 1, 2, and 3, which we discuss in turn below. Across all three tasks, CEASE achieves the most favorable balance between erasure and preservation: it substantially alters the erased concepts while keeping the retained concepts and unrelated generation close to the unedited model, whereas the baselines that erase more deeply do so by corrupting the retain set or MS-COCO prompts as well. Results on the Stable Diffusion v1.5 and v2.1 backbones appear in Appendices A.3 and A.4, with the same conclusions.

Table 2: Style erasure after continually erasing 100 artistic styles on SD v1.4. Bold and underline mark the best and second-best edited results.

	Erase set				Retain set				MS-COCO		
	CS ↓	FID ↑	LPIPS ↑	Aes ↑	CS ↑	FID ↓	LPIPS ↓	Aes ↑	CS ↑	FID ↓	LPIPS ↓
SD v1.4	28.31	0.0	0.000	6.01	28.20	0.0	0.000	5.99	26.52	0.0	0.000
ESD	19.98	85.2	0.617	5.36	20.12	<u>84.2</u>	0.618	5.39	23.50	56.7	0.517
UCE	17.28	250.0	0.772	4.27	15.83	210.5	0.768	3.93	13.97	228.0	0.751
ConAbl	20.27	127.1	0.577	<u>5.52</u>	<u>20.16</u>	125.3	<u>0.581</u>	<u>5.51</u>	22.98	57.7	<u>0.516</u>
RECE	12.75	388.7	0.916	3.13	12.41	365.3	0.909	3.20	13.79	371.6	0.904
SPEED	<u>14.94</u>	249.1	0.736	5.07	15.75	197.4	0.695	5.20	13.73	188.8	0.682
CUIG-GradProj	15.11	<u>316.6</u>	<u>0.828</u>	3.52	14.74	317.1	0.827	3.52	23.56	85.9	0.642
Sculpting Memory	21.43	117.8	0.601	5.49	<u>21.43</u>	109.2	0.596	5.47	<u>25.14</u>	42.5	0.403
Ours	23.43	64.8	0.518	5.91	25.66	38.7	0.446	5.96	25.43	<u>50.3</u>	<u>0.473</u>

Table 3: Instance erasure after continually erasing 100 instances on SD v1.4. Bold and underline mark the best and second-best edited results.

	Erase set				Retain set				MS-COCO		
	CS ↓	FID ↑	LPIPS ↑	Acc ↓	CS ↑	FID ↓	LPIPS ↓	Acc ↑	CS ↑	FID ↓	LPIPS ↓
SD v1.4	28.02	0.0	0.000	80.2	28.46	0.0	0.000	78.6	26.52	0.0	0.000
ESD	20.08	65.1	0.621	28.9	20.22	60.1	0.604	27.7	21.82	<u>64.2</u>	0.555
UCE	15.65	<u>319.8</u>	0.790	1.2	16.13	271.7	0.771	0.2	14.20	276.3	0.750
ConAbl	15.20	179.7	0.653	1.8	15.13	176.7	0.651	2.5	19.13	89.9	0.575
RECE	16.22	438.6	0.891	1.0	15.90	433.5	0.897	0.0	13.46	427.6	0.881
SPEED	18.21	85.0	0.644	5.7	<u>26.28</u>	<u>25.1</u>	<u>0.479</u>	<u>66.2</u>	<u>22.80</u>	78.0	<u>0.531</u>
CUIG-GradProj	16.13	249.4	<u>0.883</u>	<u>1.1</u>	15.65	250.6	0.882	0.5	14.20	229.3	0.892
Sculpting Memory	<u>15.24</u>	156.2	0.691	2.7	15.30	157.6	0.689	2.3	22.00	70.0	0.538
Ours	21.08	55.4	0.605	11.8	27.15	20.8	0.451	71.0	25.26	62.1	0.490

Erasure of Celebrities. As shown in Table 1, continually applying SPEED severely degrades general generation, with quality on both the retain set and COCO deteriorating as edits accumulate. CEASE performs best among the edited methods on nearly all retain and COCO preservation metrics, and the VLM judge almost never names the erased celebrities while still recognizing most of the retained ones. UCE and RECE obtain the lowest erase-set CLIP scores, but their retain-set scores collapse to a similar degree, indicating prompt-independent corruption rather than selective removal; ESD, in contrast, leaves a noticeable fraction of the erased identities recognizable.

Erasure of Styles. Table 2 shows the style task to be SPEED’s most severe failure, whereas CEASE remains best among the edited methods on all retain metrics and on COCO CLIP score. Sculpting Memory attains lower COCO FID, but at the cost of a substantially more degraded retain set. Importantly, the preservation gain of CEASE does not come from leaving the erased styles unchanged: it still lowers the erase-set CLIP score substantially, and its erase-set aesthetic score stays at the reference level, showing that the redirected outputs remain coherent paintings, whereas RECE and CUIG-GradProj produce aesthetically collapsed images.

Erasure of Instances. As reported in Table 3, the instance task exhibits SPEED’s mildest yet still substantial degradation, and CEASE again leads all retain and COCO preservation metrics while clearly altering the erased instances. ConAbl and RECE erase more deeply by erase-set CLIP score, but their collapsed retain scores expose the corresponding loss of selectivity. The accuracy columns make this explicit: CEASE keeps retain accuracy closest to the unedited reference while reducing erase accuracy to a low level, whereas the deepest-erasing baselines lose both together.

Qualitative Results. Figure 3 compares all methods on the erased target “Hillary Clinton”, the retained concept “Jason Statham”, and an unrelated MS-COCO prompt. RECE, UCE, and CUIG-GradProj collapse all three categories alike to black or noise; SPEED and ConAbl distort the erased identity and, like Sculpting Memory, replace the retained one; ESD leaves the erased target fully recognizable. In contrast, CEASE redirects the target to a coherent anchor substitute while leaving the retained concept and the COCO scene visually intact. More visualization results on the celebrity,




Figure 3: Qualitative comparison after the 100-step continual erasure, on the erased target Hillary Clinton, the retained concept Jason Statham, and an unrelated MS-COCO prompt. CEASE redirects the target to a coherent anchor substitute while preserving the retained concept and the COCO scene.

Table 4: Ablation of AIC and HOC on the celebrity task. Combining both modules improves retained-concept preservation and general generation quality. Bold marks the best results.

Module		Erase set				Retain set				MS-COCO		
AIC	HOC	CS ↓	FID ↑	LPIPS ↑	Acc ↓	CS ↑	FID ↓	LPIPS ↓	Acc ↑	CS ↑	FID ↓	LPIPS ↓
×	×	18.38	137.2	0.635	0.3	30.28	20.8	0.383	54.4	22.14	107.2	0.566
×	✓	20.84	57.8	0.528	1.5	31.44	13.5	0.298	63.7	25.54	59.7	0.457
✓	×	21.41	43.1	0.485	3.5	31.54	12.9	0.284	65.2	26.23	48.6	0.399
✓	✓	21.54	39.9	0.478	3.5	31.59	12.3	0.271	65.9	26.24	46.2	0.388

style, and instance tasks, as well as on the Stable Diffusion v1.5 and v2.1 backbones, are provided in Appendix A.7, we observe the same pattern in those settings.

4.3 ABLATION STUDY

Effect of AIC and HOC. Table 4 isolates the two components on the celebrity task. Each component alone already recovers much of the preservation lost by the baseline, with AIC contributing the larger gain, and their combination performs best on all retain and COCO metrics. The erase-set columns show the accompanying trade-off: the full method changes the erased celebrities less strongly than the baseline, redirecting them to coherent anchor substitutes rather than removing them more weakly, so the preservation gain does not come at the cost of failed erasure.

Cross-Attention Layer Scope. Figure 4 compares four layer scopes on the celebrity sequence. The middle and high-resolution scopes preserve non-target behavior but fail to erase, whereas the low-resolution and all-stage variants erase effectively with similar preservation. All Stages achieves the strongest erasure overall while retaining comparable utility, so we use it as the default.

Cross-Edit Interference over the Sequence. Figure 5 traces the value-projection interference of SPEED, RECE with one adversarial epoch per request, UCE, and CEASE over the 100-step celebrity sequence. For SPEED, interference grows steadily on unrelated COCO prompts, on the shared replacement, and on the erased target, and the target curve rises sharply in the final edits. RECE exceeds the displayed range within the first ten edits and ends several orders of magnitude higher, consistent with its COCO FID in Fig. 1(a). CEASE stays below SPEED on all three diagnostics; on the shared replacement it remains near zero, which is the direct effect of AIC. UCE is the exception that shows the limits of this diagnostic: its final interference on the shared replacement is three orders of magnitude above CEASE, while its unscaled COCO and target values end slightly lower despite far worse generation quality (the displayed UCE curve is scaled by a factor of two and ends above CEASE), so value-projection interference alone does not rank preservation across methods. Figure 13 reports the COCO CLIP score and the target-averaged diagnostic at five check-




Figure 4: Cross-attention layer-scope ablation on the 100-step celebrity sequence. We compare four editing scopes using erase and retain CLIP scores and recognition accuracy, together with COCO CLIP score and FID. Arrows indicate the preferred direction.




Figure 5: Cross-edit interference over 100 celebrity erasures, comparing SPEED, RECE, UCE, and CEASE. The three panels track noise on unrelated COCO prompts, the shared replacement “person”, and the erased target “Tom Cruise” as successive edits accumulate.

points; relative to SPEED, CEASE keeps a higher CLIP score and lower FID at every checkpoint. The underlying measurements use saved checkpoints from complete sequential runs; UCE is displayed with the stated post-processing scale; the projection diagnostic is interpreted separately from generation quality (Appendix A.8).

In addition, we analyze the robustness of CEASE to the erasure order in Appendix A.5, and compare its computational efficiency with all baselines in Appendix A.6.

5 CONCLUSION

We studied continual concept erasure, where deletion requests arrive one after another and each edit builds on all previous ones. We traced the severe degradation of existing methods in this regime to accumulating cross-edit interference, and proposed CEASE, a training-free framework that protects the shared anchor and steers each new update away from the directions perturbed by past edits. Experiments on erasing celebrities, artistic styles, and instances across multiple backbones show that CEASE achieves the most consistent erase–preserve trade-off.

Limitations. First, our experiments focus on the Stable Diffusion family, and extending CEASE to more recent architectures such as FLUX is a natural next step. Second, this work instantiates the proposed constraints in a closed-form solver, and generalizing them to training-based erasure methods is left to future work. Third, our sequences contain up to 100 edits; evaluating substantially longer streams is conceptually straightforward but computationally expensive, since every additional step requires regenerating and re-evaluating thousands of images.

ACKNOWLEDGMENTS

This research is supported by the National Artificial Intelligence Research Resource Pilot Awards NAIRR250199, NAIRR260019, and NAIRR260077; AMD University Program’s AI & HPC Cluster; NVIDIA Academic Grant Program; Lambda’s Research Grant; and Google Cloud Research Credits Program. Computational resources are also provided by Delta and DeltaAI at the National Center for Supercomputing Applications (NCSA) through ACCESS allocations CIS250012, CIS250816, and CIS251188.

REFERENCES

- Anudeep Das, Vasisht Duddu, Rui Zhang, and N Asokan. Espresso: Robust concept filtering in text-to-image models. In *Proceedings of the Fifteenth ACM Conference on Data and Application Security and Privacy*, pp. 305–316, 2024.
- Chongyu Fan, Jiancheng Liu, Yihua Zhang, Eric Wong, Dennis Wei, and Sijia Liu. Salun: Empowering machine unlearning via gradient-based weight saliency in both image classification and generation. In *International Conference on Learning Representations*, volume 2024, pp. 53643–53673, 2024.
- Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Jie Shi, Xiang Wang, Xiangnan He, and Tat-Seng Chua. Alphaedit: Null-space constrained knowledge editing for language models. In *International Conference on Learning Representations*, volume 2025, pp. 16366–16396, 2025.
- Tom Fischer, Yaoyao Liu, Artur Jesslen, Noor Ahmed, Prakhar Kaushik, Angtian Wang, Alan L Yuille, Adam Kortylewski, and Eddy Ilg. inemo: Incremental neural mesh models for robust class-incremental learning. In *European Conference on Computer Vision*, pp. 357–374, 2024.
- 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 *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 5099–5108. IEEE, 2024.
- Naveen George, Naoki Murata, Yuhta Takida, Konda Reddy Mopuri, and Yuki Mitsufuji. Distill, forget, repeat: A framework for continual unlearning in text-to-image diffusion models. *arXiv preprint arXiv:2512.02657*, 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 *European Conference on Computer Vision*, pp. 73–88. Springer, 2024.
- Jia-Chen Gu, Hao-Xiang Xu, Jun-Yu Ma, Pan Lu, Zhen-Hua Ling, Kai-Wei Chang, and Nanyun Peng. Model editing harms general abilities of large language models: Regularization to the rescue. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 16801–16819, 2024.
- Tom Hartvigsen, Swami Sankaranarayanan, Hamid Palangi, Yoon Kim, and Marzyeh Ghassemi. Aging with grace: Lifelong model editing with discrete key-value adaptors. *Advances in neural information processing systems*, 36:47934–47959, 2023.
- Alvin Heng and Harold Soh. Selective amnesia: A continual learning approach to forgetting in deep generative models. *Advances in Neural Information Processing Systems*, 36:17170–17194, 2023.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. In *Proceedings of the 2021 conference on empirical methods in natural language processing*, pp. 7514–7528, 2021.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.

- Chi-Pin Huang, Kai-Po Chang, Chung-Ting Tsai, Yung-Hsuan Lai, Fu-En Yang, and Yu-Chiang Frank Wang. Receler: Reliable concept erasing of text-to-image diffusion models via lightweight erasers. In *European Conference on Computer Vision*, pp. 360–376. Springer, 2024.
- Eric Ji, Qiran Hu, Wufei Ma, Sarthak Jain, Yingying Li, Minh N Do, and Yaoyao Liu. AC3S: Adaptive conditioning for 3D-aware synthetic data generation. In *European Conference on Computer Vision*, pp. 510–525, 2026.
- Changhoon Kim, Kyle Min, and Yezhou Yang. Race: Robust adversarial concept erasure for secure text-to-image diffusion model. In *European Conference on Computer Vision*, pp. 461–478. Springer, 2024.
- Nupur Kumari, Bingliang Zhang, Sheng-Yu Wang, Eli Shechtman, Richard Zhang, and Jun-Yan Zhu. Ablating concepts in text-to-image diffusion models. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 22634–22645. IEEE, 2023.
- Justin Lee, Zheda Mai, Jinsu Yoo, Chongyu Fan, Cheng Zhang, and Wei-Lun Chao. Continual unlearning for text-to-image diffusion models: A regularization perspective. In *International Conference on Learning Representations*, volume 2026, pp. 30660–30682, 2026.
- Gen Li, Yang Xiao, Jie Ji, Kaiyuan Deng, Bo Hui, Linke Guo, and Xiaolong Ma. Sculpting memory: Multi-concept forgetting in diffusion models via dynamic mask and concept-aware optimization. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 19659–19668. IEEE, 2025.
- Ouxiang Li, Yuan Wang, Xinting Hu, Houcheng Jiang, Yanbin Hao, James Ma, Fuli Feng, et al. Speed: Scalable, precise, and efficient concept erasure for diffusion models. In *International Conference on Learning Representations*, volume 2026, pp. 66777–66803, 2026.
- Xinfeng Li, Yuchen Yang, Jiangyi Deng, Chen Yan, Yanjiao Chen, Xiaoyu Ji, and Wenyuan Xu. Safegen: Mitigating sexually explicit content generation in text-to-image models. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pp. 4807–4821, 2024.
- Yingying Li, Xin Chen, and Na Li. Online optimal control with linear dynamics and predictions: Algorithms and regret analysis. *Advances in Neural Information Processing Systems*, 32, 2019.
- Yingying Li, Guannan Qu, and Na Li. Online optimization with predictions and switching costs: Fast algorithms and the fundamental limit. *IEEE Transactions on Automatic Control*, 66(10): 4761–4768, 2020.
- Yingying Li, Subhro Das, and Na Li. Online optimal control with affine constraints. In *AAAI Conference on Artificial Intelligence*, pp. 8527–8537, 2021.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Renyang Liu, Wenjie Feng, Tianwei Zhang, Wei Zhou, Xueqi Cheng, and See-Kiong Ng. Rethinking machine unlearning in image generation models. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security*, pp. 993–1007, 2025.
- Renyang Liu, Guanlin Li, Tianwei Zhang, and See-Kiong Ng. Image can bring your memory back: A novel multi-modal guided attack against image generation model unlearning. In *International Conference on Learning Representations*, volume 2026, pp. 143251–143275, 2026.
- Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics training: Multi-class incremental learning without forgetting. In *IEEE/CVF Computer Vision and Pattern Recognition Conference*, pp. 12245–12254, 2020.
- Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive aggregation networks for class-incremental learning. In *IEEE/CVF Computer Vision and Pattern Recognition Conference*, pp. 2544–2553, 2021a.

- Yaoyao Liu, Bernt Schiele, and Qianru Sun. Rmm: Reinforced memory management for class-incremental learning. *Advances in Neural Information Processing Systems*, 34:3478–3490, 2021b.
- Yaoyao Liu, Yingying Li, Bernt Schiele, and Qianru Sun. Online hyperparameter optimization for class-incremental learning. In *AAAI Conference on Artificial Intelligence*, pp. 8906–8913, 2023a.
- Yaoyao Liu, Bernt Schiele, Andrea Vedaldi, and Christian Rupprecht. Continual detection transformer for incremental object detection. In *IEEE/CVF Computer Vision and Pattern Recognition Conference*, pp. 23799–23808, 2023b.
- Yaoyao Liu, Yingying Li, Bernt Schiele, and Qianru Sun. Wakening past concepts without past data: Class-incremental learning from online placebos. In *IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 2226–2235, 2024.
- Shilin Lu, Zilan Wang, Leyang Li, Yanzhu Liu, and Adams Wai-Kin Kong. Mace: Mass concept erasure in diffusion models. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (IEEE/CVF Computer Vision and Pattern Recognition Conference)*, pp. 6430–6440. IEEE, 2024.
- Zilin Luo, Yaoyao Liu, Bernt Schiele, and Qianru Sun. Class-incremental exemplar compression for class-incremental learning. In *IEEE/CVF Computer Vision and Pattern Recognition Conference*, pp. 11371–11380, 2023.
- Mengyao Lyu, Yuhong Yang, Haiwen Hong, Hui Chen, Xuan Jin, Yuan He, Hui Xue, Jungong Han, and Guiguang Ding. One-dimensional adapter to rule them all: Concepts, diffusion models and erasing applications. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (IEEE/CVF Computer Vision and Pattern Recognition Conference)*, pp. 7559–7568. IEEE, 2024.
- Jun-Yu Ma, Hong Wang, Hao-Xiang Xu, Zhen-Hua Ling, and Jia-Chen Gu. Perturbation-restrained sequential model editing. In *International Conference on Learning Representations*, volume 2025, pp. 17302–17329, 2025.
- Wufei Ma, Qihao Liu, Jiahao Wang, Angtian Wang, Xiaoding Yuan, Yi Zhang, Zihao Xiao, Guofeng Zhang, Beijia Lu, Ruxiao Duan, Yongrui Qi, Adam Kortylewski, Yaoyao Liu, and Alan L. Yuille. Generating images with 3D annotations using diffusion models. In *International Conference on Learning Representations (ICLR)*, 2024.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *Advances in neural information processing systems*, 35:17359–17372, 2022a.
- Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. Mass-editing memory in a transformer. *arXiv preprint arXiv:2210.07229*, 2022b.
- Hadas Orgad, Bahjat Kavar, and Yonatan Belinkov. Editing implicit assumptions in text-to-image diffusion models. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 7030–7038. IEEE, 2023.
- Yong-Hyun Park, Sangdoo Yun, Jin-Hwa Kim, Junho Kim, Geonhui Jang, Yonghyun Jeong, Junghyo Jo, and Gayoung Lee. Direct unlearning optimization for robust and safe text-to-image models. *Advances in Neural Information Processing Systems*, 37:80244–80267, 2024.
- Minh Pham, Kelly O Marshall, Chinmay Hegde, and Niv Cohen. Robust concept erasure using task vectors. *arXiv preprint arXiv:2404.03631*, 2024.
- Javier Rando, Daniel Paleka, David Lindner, Lennart Heim, and Florian Tramèr. Red-teaming the stable diffusion safety filter. *arXiv preprint arXiv:2210.04610*, 2022.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021.
- Patrick Schramowski, Manuel Brack, Björn Deiseroth, and Kristian Kersting. Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (IEEE/CVF Computer Vision and Pattern Recognition Conference)*, pp. 22522–22531. IEEE, 2023.

- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022.
- Kartik Thakral, Tamar Glaser, Tal Hassner, Mayank Vatsa, and Richa Singh. Continual unlearning for foundational text-to-image models without generalization erosion. *arXiv preprint arXiv:2503.13769*, 2025.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024a.
- Peng Wang, Zexi Li, Ningyu Zhang, Ziwen Xu, Yunzhi Yao, Yong Jiang, Pengjun Xie, Fei Huang, and Huajun Chen. Wise: Rethinking the knowledge memory for lifelong model editing of large language models. *Advances in Neural Information Processing Systems*, 37:53764–53797, 2024b.
- Yongliang Wu, Shiji Zhou, Mingzhuo Yang, Lianzhe Wang, Heng Chang, Wenbo Zhu, Xinting Hu, Xiao Zhou, and Xu Yang. Unlearning concepts in diffusion model via concept domain correction and concept preserving gradient. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 8496–8504, 2025.
- Abudukelimu Wuerkaixi, Qizhou Wang, Sen Cui, Wutong Xu, Bo Han, Gang Niu, Masashi Sugiyama, and Changshui Zhang. Adaptive localization of knowledge negation for continual llm unlearning. In *Forty-second International Conference on Machine Learning*, 2025.
- Gong Zhang, Kai Wang, Xingqian Xu, Zhangyang Wang, and Humphrey Shi. Forget-me-not: Learning to forget in text-to-image diffusion models. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (IEEE/CVF Computer Vision and Pattern Recognition ConferenceW)*, pp. 1755–1764. IEEE, 2024.
- Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *IEEE/CVF International Conference on Computer Vision*, pp. 3813–3824. IEEE, 2023a.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE/CVF conference on computer vision and pattern recognition*, pp. 586–595. IEEE, 2018.
- Yixiao Zhang, Xinyi Li, Huimiao Chen, Alan L Yuille, Yaoyao Liu, and Zongwei Zhou. Continual learning for abdominal multi-organ and tumor segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 35–45. Springer, 2023b.
- Mengnan Zhao, Lihe Zhang, Xingyi Yang, Tianhang Zheng, and Baocai Yin. Advanchor: Enhancing diffusion model unlearning with adversarial anchors. *arXiv preprint arXiv:2501.00054*, 2024.

A APPENDIX

A.1 IMPLEMENTATION DETAILS

ESD (Gandikota et al., 2023) uses the ESD-x variant, fine-tuning the cross-attention projections under the negative-guidance objective for 200 Adam iterations per step at learning rate 1×10^{-4} with guidance strength 1.0, and training latents are generated at guidance 3.0 over 50 DDIM steps.

ConAbl (Kumari et al., 2023) fine-tunes the cross-attention weights toward its own anchor-prompt prediction, with prior-loss weight 1.0 and 100 Adam iterations per concept at learning rate 1×10^{-4} .

UCE (Gandikota et al., 2024) applies one closed-form least-squares update per step to the key and value projections, with erase scale 1.0, preserve scale 1.0, and regularization $\lambda = 0.5$.

RECE (Gong et al., 2024) alternates a closed-form value edit with the derivation of an adversarial target embedding, in the replace regime with 3 adversarial epochs, erase scale 1.0, preserve scale 0.1, closed-form $\lambda = 0.1$, and adversarial regularization 0.1.

SPEED (Li et al., 2026) redirects each target to the anchor while projecting the update onto the null space of an augmented retain set; its invariance matrix uses the empty-prompt representation, and CEASE’s AIC additionally inserts the task anchor. Both use null-space threshold 0.1, retain scale 0.05, $\lambda = 0$, 10 augmented retain samples per concept, and the official multi-concept configuration without the importance-based retain filter.

CUIG-GradProj (Lee et al., 2026) adds the official semantic-aware gradient-projection regularizer to ConAbl with 200 auxiliary prompts, fine-tuning the key and value projections with batch size 4, learning rate 2×10^{-6} before the official scaling, 500 warmup steps, and weight decay 10^{-2} , for 2,000 iterations per request (1,000 for style).

Sculpting Memory (Li et al., 2025) combines a dynamic parameter mask with distillation from the preceding model, using 100 generated images per path, batch size 2, 10 epochs, learning rate 3×10^{-6} , and a cosine OneCycle schedule with 50 warmup steps; the mask updates every 100 steps with fraction 0.5, and the unlearning and distillation scales are 0.25.

A.2 NEAR-NULL PROJECTOR

For completeness, SPEED constructs the retain projector from the eigendecomposition of the retain Gram matrix,

$$\mathbf{C}_0 \mathbf{C}_0^\top = \mathbf{U} \mathbf{A} \mathbf{U}^\top, \quad \mathbf{P} = \widehat{\mathbf{U}} \widehat{\mathbf{U}}^\top, \quad (15)$$

where $\widehat{\mathbf{U}}$ contains the eigenvectors whose eigenvalues are below a threshold τ . An update right-projected as $\Delta^* \mathbf{P}$ therefore lies in the approximate null space of the retain embeddings and satisfies $(\Delta^* \mathbf{P}) \mathbf{C}_0 \approx \mathbf{0}$. We use $\tau = 0.1$ in all experiments, following SPEED.

A.3 BACKBONE GENERALIZATION TO STABLE DIFFUSION v1.5

We repeat the full 100-step sequence on Stable Diffusion v1.5, a second CLIP-L backbone, for all tasks and baselines, and report the results in Tables 5, 6, and 7. The erase–preserve pattern of the main results transfers: CEASE improves the retain and COCO preservation metrics over SPEED on all three tasks while still altering the erased targets, and its aesthetic scores on styles stay at the reference level. The baselines that erase more deeply pay for it on preservation: ConAbl attains a lower COCO FID on celebrity and style but largely destroys the retained concepts, and UCE and RECE erase deepest with collapsed retain sets.

A.4 BACKBONE GENERALIZATION TO STABLE DIFFUSION v2.1

Stable Diffusion v2.1 replaces the CLIP-L text encoder with the 1024-dimensional OpenCLIP-H and is the hardest tested backbone: every method erases more weakly, and the margins narrow accordingly (Tables 8, 9, and 10). CEASE still improves the retain and COCO preservation metrics over SPEED on all three tasks, whereas some baselines preserve well only because they barely erase: UCE leaves most of the erased celebrities recognizable, and ESD leaves the erased instances nearly

Table 5: Celebrity erasure after continually erasing 100 celebrities on SD v1.5. Bold and underline mark the best and second-best edited results.

	Erase set				Retain set				MS-COCO		
	CS ↓	FID ↑	LPIPS ↑	Acc ↓	CS ↑	FID ↓	LPIPS ↓	Acc ↑	CS ↑	FID ↓	LPIPS ↓
SD v1.5	31.13	0.0	0.000	88.8	31.92	0.0	0.000	70.6	26.43	0.0	0.000
ESD	28.74	57.4	0.413	79.2	28.86	56.5	0.415	<u>58.2</u>	24.82	46.1	0.446
UCE	13.08	363.3	0.818	0.0	13.63	293.9	0.795	0.0	13.10	294.4	0.788
ConAbl	20.04	98.5	0.454	2.0	21.35	92.2	0.447	1.8	<u>25.72</u>	37.3	0.291
RECE	16.51	564.4	0.977	0.0	16.27	572.0	0.975	0.0	16.28	610.4	1.002
SPEED	18.55	186.3	0.586	<u>0.6</u>	<u>30.54</u>	<u>50.8</u>	<u>0.348</u>	53.2	22.70	102.1	0.551
CUIG-GradProj	<u>15.01</u>	384.6	<u>0.851</u>	0.0	14.77	380.9	0.856	0.0	13.86	313.2	0.848
Sculpting Memory	17.14	165.4	0.645	<u>0.6</u>	17.66	153.9	0.640	0.6	25.51	<u>45.3</u>	0.418
Ours	21.59	76.9	0.456	4.0	31.80	35.4	0.237	65.4	26.31	46.2	<u>0.386</u>

Table 6: Style erasure after continually erasing 100 artistic styles on SD v1.5. Bold and underline mark the best and second-best edited results.

	Erase set				Retain set				MS-COCO		
	CS ↓	FID ↑	LPIPS ↑	Aes ↑	CS ↑	FID ↓	LPIPS ↓	Aes ↑	CS ↑	FID ↓	LPIPS ↓
SD v1.5	28.34	0.0	0.000	6.01	28.24	0.0	0.000	6.00	26.43	0.0	0.000
ESD	24.50	50.9	0.485	<u>5.76</u>	<u>24.53</u>	<u>49.4</u>	0.485	<u>5.77</u>	<u>25.71</u>	<u>42.9</u>	0.388
UCE	17.09	227.8	0.860	4.59	15.41	182.4	0.795	4.24	13.70	210.2	0.776
ConAbl	23.70	66.4	0.468	5.67	24.01	61.3	<u>0.468</u>	5.66	25.37	40.6	0.355
RECE	14.60	663.7	0.996	4.60	14.05	668.9	0.997	4.60	16.28	610.4	1.002
SPEED	<u>15.81</u>	225.3	0.724	4.83	17.69	170.1	0.688	5.02	14.31	186.7	0.655
CUIG-GradProj	16.55	<u>284.7</u>	<u>0.885</u>	4.00	16.02	280.0	0.884	4.03	19.95	143.7	0.740
Sculpting Memory	21.45	111.1	0.597	5.44	21.64	104.2	0.590	5.43	25.20	43.0	0.402
Ours	23.90	90.6	0.514	5.87	27.17	37.9	0.380	5.92	26.23	44.0	<u>0.370</u>

intact. Overall, the transfer study supports a consistent preservation advantage over the underlying SPEED solver across both text encoders, at the cost of shallower erasure on several tasks.

A.5 ORDER ROBUSTNESS

We rerun the 100-step celebrity sequence under five random permutations of the concept order with editing seed 0. Table 11 reports the per-order results: the variation across orders is small on every metric relative to the between-method differences in the main tables.

A.6 RUNTIME EFFICIENCY

Fig. 6 compares single-request wall time on one H100 GPU. CEASE completes a request in about 13 seconds, staying in the seconds-level regime of the closed-form methods, whereas the training-based methods take minutes per request, and this gap compounds over the continual sequence. Relative to SPEED, AIC and HOC add about 2% to the edit loop.

A.7 ADDITIONAL QUALITATIVE COMPARISONS

This section provides qualitative comparisons beyond the single case shown in the main text. Fig. 7 extends the main-text comparison to four cases per category on the celebrity task, Figs. 8 and 9 give the same comparison on the style and instance tasks, and Figs. 10 and 11 repeat the celebrity comparison on the Stable Diffusion v1.5 and v2.1 backbones. Every figure follows the same layout: columns show the unedited reference followed by ESD, UCE, ConAbl, RECE, SPEED, CUIG-GradProj, Sculpting Memory, and CEASE after the full 100-step sequence, and the three row blocks show four cases each of erased targets, retained concepts, and unrelated MS-COCO prompts.

Table 7: Instance erasure after continually erasing 100 instances on SD v1.5. Bold and underline mark the best and second-best edited results.

	Erase set				Retain set				MS-COCO		
	CS ↓	FID ↑	LPIPS ↑	Acc ↓	CS ↑	FID ↓	LPIPS ↓	Acc ↑	CS ↑	FID ↓	LPIPS ↓
SD v1.5	28.05	0.0	0.000	79.9	28.52	0.0	0.000	78.9	26.43	0.0	0.000
ESD	24.28	38.9	0.535	58.6	24.60	36.0	<u>0.517</u>	<u>58.0</u>	24.50	46.6	0.440
UCE	14.97	<u>301.0</u>	0.809	1.6	15.25	247.8	0.791	0.2	14.03	264.1	0.758
ConAbl	19.76	74.4	0.567	15.8	20.36	71.2	0.562	18.8	<u>24.89</u>	<u>46.8</u>	0.411
RECE	17.23	606.3	1.006	0.0	17.12	594.9	1.006	1.0	16.28	610.4	1.002
SPEED	17.59	89.5	0.673	5.0	<u>25.23</u>	<u>26.9</u>	0.534	57.6	21.73	86.7	0.594
CUIG-GradProj	16.78	256.4	<u>0.913</u>	<u>1.1</u>	16.64	250.0	0.910	0.4	13.48	234.8	0.914
Sculpting Memory	<u>15.48</u>	149.7	0.703	3.1	15.48	155.2	0.698	3.0	22.48	63.9	0.526
Ours	21.37	46.2	0.570	16.0	27.80	17.6	0.372	75.6	25.69	53.5	<u>0.417</u>

Table 8: Celebrity erasure after continually erasing 100 celebrities on SD v2.1. Bold and underline mark the best and second-best edited results.

	Erase set				Retain set				MS-COCO		
	CS ↓	FID ↑	LPIPS ↑	Acc ↓	CS ↑	FID ↓	LPIPS ↓	Acc ↑	CS ↑	FID ↓	LPIPS ↓
SD v2.1	30.60	0.0	0.000	88.8	31.03	0.0	0.000	75.6	26.32	0.0	0.000
ESD	25.85	103.5	0.520	62.0	25.48	109.2	0.523	41.8	23.43	52.8	0.481
UCE	27.82	47.1	0.398	47.0	30.74	42.2	0.362	68.2	26.11	48.5	0.442
ConAbl	15.63	214.3	0.603	0.0	15.68	206.1	0.604	0.0	24.60	47.9	0.451
RECE	16.51	549.3	0.973	0.0	16.27	562.3	0.971	0.0	16.28	614.3	1.002
SPEED	24.90	53.7	0.377	<u>12.2</u>	<u>31.00</u>	<u>32.8</u>	<u>0.245</u>	72.0	<u>26.12</u>	<u>41.9</u>	<u>0.340</u>
CUIG-GradProj	14.98	<u>381.1</u>	<u>0.864</u>	0.0	14.71	381.4	0.864	0.0	12.50	284.0	0.954
Sculpting Memory	<u>15.15</u>	164.2	0.693	0.0	15.26	152.3	0.680	0.4	24.48	58.6	0.527
Ours	26.91	38.8	0.326	34.2	31.01	31.2	0.217	<u>69.6</u>	26.14	38.2	0.290

The failure modes match the quantitative results and the main-text analysis. On the erased rows, CEASE redirects each target to a coherent anchor substitute, whereas ESD often leaves the erased identity or style still recognizable, and RECE, UCE, and CUIG-GradProj collapse to noise-like textures once the sequence is long. Sculpting Memory remains photorealistic but drifts to subjects unrelated to either the target or its anchor. On the retained and COCO rows, the picture inverts: the destructive baselines corrupt concepts they were never asked to erase, SPEED preserves some cases but degrades others as its accumulated interference grows, and CEASE stays visually closest to the reference across both row blocks. The v1.5 and v2.1 figures show the same pattern, confirming that these behaviors are properties of the continual setting rather than of the v1.4 backbone.

Table 9: Style erasure after continually erasing 100 artistic styles on SD v2.1. Bold and underline mark the best and second-best edited results.

	Erase set				Retain set				MS-COCO		
	CS ↓	FID ↑	LPIPS ↑	Aes ↑	CS ↑	FID ↓	LPIPS ↓	Aes ↑	CS ↑	FID ↓	LPIPS ↓
SD v2.1	27.26	0.0	0.000	6.01	27.37	0.0	0.000	5.99	26.32	0.0	0.000
ESD	25.71	41.6	0.400	5.95	25.69	41.8	0.399	5.94	25.45	41.4	0.366
UCE	25.89	28.7	0.316	<u>5.99</u>	<u>26.08</u>	<u>28.0</u>	<u>0.303</u>	<u>5.97</u>	<u>25.91</u>	<u>40.3</u>	<u>0.339</u>
ConAbl	25.25	43.2	0.424	5.84	25.49	40.9	0.415	5.83	25.51	41.0	0.350
RECE	12.45	<u>442.0</u>	<u>0.866</u>	3.48	12.14	435.4	0.862	3.43	14.47	403.6	0.878
SPEED	24.93	40.7	0.404	5.84	25.29	37.9	0.370	5.89	25.35	50.3	0.396
CUIG-GradProj	<u>14.10</u>	468.9	0.889	3.39	13.55	462.9	0.886	3.40	13.74	392.9	0.794
Sculpting Memory	19.42	98.3	0.609	5.46	20.00	91.6	0.599	5.47	24.57	47.3	0.453
Ours	26.14	25.6	0.266	6.01	26.41	24.7	0.247	6.01	25.97	37.7	0.291

Table 10: Instance erasure after continually erasing 100 instances on SD v2.1. Bold and underline mark the best and second-best edited results.

	Erase set				Retain set				MS-COCO		
	CS ↓	FID ↑	LPIPS ↑	Acc ↓	CS ↑	FID ↓	LPIPS ↓	Acc ↑	CS ↑	FID ↓	LPIPS ↓
SD v2.1	26.39	0.0	0.000	74.3	26.43	0.0	0.000	70.2	26.32	0.0	0.000
ESD	25.64	21.6	0.395	69.2	25.51	19.8	<u>0.386</u>	64.4	<u>25.72</u>	40.1	<u>0.347</u>
UCE	17.50	125.5	0.666	5.8	19.15	76.9	0.617	19.8	18.43	135.4	0.612
ConAbl	20.67	60.4	0.567	27.2	20.72	60.5	0.567	27.0	25.17	45.4	0.396
RECE	<u>16.66</u>	351.8	0.883	<u>1.0</u>	16.44	327.3	0.874	0.0	13.35	419.1	0.880
SPEED	22.30	36.4	0.507	34.6	<u>25.73</u>	<u>19.0</u>	0.403	<u>65.7</u>	25.55	51.2	0.408
CUIG-GradProj	17.37	<u>350.0</u>	<u>0.822</u>	0.9	16.79	351.9	0.828	0.0	13.40	351.6	0.839
Sculpting Memory	14.43	178.6	0.703	2.5	14.39	179.5	0.702	4.7	18.30	119.7	0.615
Ours	24.04	25.1	0.445	47.7	26.01	16.2	0.338	67.1	25.85	<u>43.1</u>	0.340

Table 11: Erasure-order robustness on the celebrity task: mean $\pm$ std over five random orderings of the 100 targets, using five templates and one image-sampling seed per template. The editing seed is fixed at 0, and CEASE includes the same retain augmentation as in the main table.

	Erase set			Retain set			MS-COCO		
	CS ↓	FID ↑	LPIPS ↑	CS ↑	FID ↓	LPIPS ↓	CS ↑	FID ↓	LPIPS ↓
Ours	21.35 $\pm$ 0.06	81.5 $\pm$ 0.8	0.543 $\pm$ 0.002	31.57 $\pm$ 0.08	48.9 $\pm$ 1.0	0.428 $\pm$ 0.003	26.22 $\pm$ 0.05	46.3 $\pm$ 0.7	0.387 $\pm$ 0.009




Figure 6: Single-request wall time on one H100 GPU. CEASE stays in the seconds-level regime of the closed-form methods, one to two orders of magnitude below the training-based baselines.




Figure 7: Additional celebrity examples after the full 100-step sequence on Stable Diffusion v1.4.




Figure 8: Additional style examples after the full 100-step sequence on Stable Diffusion v1.4.




Figure 9: Additional instance examples after the full 100-step sequence on Stable Diffusion v1.4.




Figure 10: Celebrity examples after the full 100-step sequence on Stable Diffusion v1.5.



Figure 11: Celebrity examples after the full 100-step sequence on Stable Diffusion v2.1.

A.8 SCOPE AND PROTOCOL OF THE EXPANDED NOISE COMPARISONS

The expanded diagnostic comparison uses the same 100-celebrity order, frozen SD v1.4 text encoder, and 16 cross-attention value projections for every method. It does not measure changes to key projections or other trainable parameters. Every baseline in Figs. 1 and 5 also appears in Table 1. Fig. 12 adds CUIG-GradProj and Sculpting Memory to those comparisons. Several methods have smaller value-projection noise than CEASE despite weaker generation quality; this diagnostic therefore cannot be used alone to rank arbitrary erasure methods.

Definitions. Fig. 5 uses the fixed-prompt diagnostic

$$D_N(c) = \sum_{\ell} \|(W_{N,\ell} - W_{0,\ell})c\|_2^2 - \sum_{\ell} \|(W_{N,\ell} - W_{N-1,\ell})c\|_2^2. \quad (16)$$

The COCO curve averages D_N over 1,000 caption embeddings; the other curves use the last non-EOS token of “person” and “Tom Cruise in an official photo”. This operational diagnostic subtracts the latest request’s contribution. Fig. 13(b), in contrast, averages over the erased concepts seen so far and subtracts each concept’s own request, $\frac{1}{N} \sum_{i \leq N} [\sum_{\ell} \|(W_{N,\ell} - W_{0,\ell})c_i\|_2^2 - \sum_{\ell} \|(W_{i,\ell} - W_{i-1,\ell})c_i\|_2^2]$. SPEED and CEASE use signed values without absolute-value conversion; all other methods use the absolute value of each recorded diagnostic, taken after averaging where applicable. The measurement files retain the signed differences of energies. Original measurement files remain unchanged. These diagnostics alone do not determine erasure or preservation quality.

RECE diagnostic run. The added curves use one adversarial epoch per request, with erase scale 1.0, preserve scale 0.1, $\lambda = 0.1$, adversarial regularization 0.1, and the replace objective. This is labeled “RECE” in the figures and is separate from the three-epoch configuration described above. A three-epoch rerun failed during request 61 with a singular-matrix error and is not presented as a complete trajectory. We cache only the frozen retain embeddings and omit diagnostic preview images; a one-request check found all 32 saved key/value tensors bitwise identical to the uncached implementation. Every plotted milestone and qualitative image comes from the complete one-epoch rerun, without joining it to a historical final checkpoint.

UCE sequence and shared sampling protocol. UCE uses its saved 100-step celebrity sequence with erase and preserve scales of 1.0, $\lambda = 0.5$, and the shared replacement “person”. All 32 final key/value tensors match the checkpoint used for the UCE row in Table 1. Its noise curve uses all 100 checkpoints; COCO metrics are measured at $N \in \{10, 25, 50, 75, 100\}$. Every plotted method uses FP16 sampling with 20 DPM-Solver steps, guidance 7.5, COCO seed 0, batch size 10, and the same 1,000 prompts. FID uses the same unedited SD image set as the existing curves. Figure 1(a) shows point estimates; no uncertainty band is drawn without measured bounds.




Figure 12: Completed candidate comparisons on the common value-projection diagnostic, including methods with lower noise than CEASE. The one-epoch RECE configuration and UCE sequence are defined above.




Figure 13: Scaling on the celebrity benchmark at $N \in \{10, 25, 50, 75, 100\}$. (a) COCO CLIP score on 1,000 prompts; the gray line is the unedited SD reference. (b) Target-averaged value-projection diagnostic on a logarithmic axis. Markers denote measured checkpoints; COCO FID at the same checkpoints is shown in Fig. 1(a).




Figure 14: COCO and erased-concept comparisons for SPEED, RECE, UCE, and CEASE after 100 celebrity erasures. Every row uses the same prompt and random seed with FP16 sampling, and erased-target sampling retains the original batch positions. RECE uses one adversarial epoch per request.