

CONTENTS

1	Introduction	1
2	Related work	2
3	Methodology	3
3.1	Choice of layers to steer	3
3.2	Construction of steering vectors	3
3.3	Using steering vectors to control generation	4
3.4	Practical considerations	6
4	Experiments	6
4.1	Results	7
4.2	Ablation Study	10
5	Conclusion	10
6	Acknowledgments	11
A	Limitations and Broader Impact	16
B	Algorithms	17
C	Prompts for generating steering vectors	18
C.1	Prompts for erasing concrete concepts	18
C.2	Prompts for human-related concepts	18
C.3	Prompts for style manipulation	18
D	User studies	19
E	Details on style removal task	21
F	Results on SDXL model	22
F.1	Experimental setup	22
F.2	Quantitative results	22
F.3	Qualitative results	22
G	Results on SANA model	27
G.1	Experimental setup	27
G.2	Quantitative results	27
G.3	Qualitative results	27
H	Ablations on hyperparameters	32

H.1	Number and choice of steering vectors	32
H.2	Importance of dot product weighting	32
H.3	Steering strength	33
H.4	Steering other layers	35
H.5	Steering only fraction of layers	35
I	Other tasks	39
I.1	Concept switch	39
I.1.1	Method	39
I.1.2	Qualitative results	39
I.2	Concept addition	39
I.2.1	Method	39
I.2.2	Qualitative results	40
I.3	Style Transfer	40
I.3.1	Method	41
J	Interpreting steering vectors	46
J.1	Steering vectors visualization	46
J.2	UMap on steering vectors	46
K	SPM and DoCo vs Casteer on adversarial prompts	50
L	Expressive Power of Steering Cross-Attention Outputs over Steering Prompt Embeddings	52
L.1	Setup	52
L.2	Main Result	52
L.3	Interpretation for Diffusion Models	53
M	Erasing multiple concepts	54
N	SD-1.5 on steering vectors from SD-1.4	55
O	SD-1.4 on steering vectors computed on single denoising step	57

A LIMITATIONS AND BROADER IMPACT

Limitations. While CASTeers demonstrates strong performance on a wide range of concept erasure tasks without retraining, several limitations remain. First, the current method is designed specifically for diffusion models with a Transformer-based cross-attention architecture. Its generalizability to architectures that do not use cross-attention has not been evaluated and may require additional methodological adjustments. Second, although the construction of steering vectors is training-free, it depends on curated positive and negative prompts, which may introduce human bias and require domain knowledge for effective pairing. While CASTeers demonstrates effective and precise control over concept suppression tasks, there is still limited understanding of how steering vectors affect the broader semantic space. Deep learning research has largely moved forward through empirical results, often ahead of solid theoretical explanations. We believe such explanations are both useful and ultimately necessary. Still, important progress has often come from work without clear theory, as seen in the case of the batch normalization paper. We hope our work contributes to ongoing efforts to better understand and apply steering methods, especially to improve control and interpretability in diffusion models.

Broader Impact. CASTeers contributes toward the democratization of safe and controllable image generation by offering a lightweight, training-free solution for concept steering in diffusion models. Its ability to remove specific concepts without retraining lowers the barrier for safety interventions in generative models, potentially empowering developers with limited resources to implement moderation tools and creative controls. This is particularly relevant for applications in content moderation, personalized media generation, and bias mitigation. However, CASTeers also presents risks. The same mechanisms that allow the removal of harmful or copyrighted content can be used to suppress beneficial or truthful concepts for deceptive purposes. Moreover, the capability to switch or add identities (e.g., celebrity faces) raises ethical concerns regarding consent, misrepresentation, and deepfake generation. While we do not endorse any misuse of this technology, we believe transparency in capabilities and limitations is essential. Further safeguards and usage guidelines should accompany any deployment to ensure CASTeers is used responsibly.

Future Work. The construction of steering vectors used by CASTeers depends on curated positive and negative prompts, which may introduce human bias and require domain knowledge for effective pairing. Future work on finding the best ways of constructing prompt pairs would be beneficial. Next, deep theoretical understanding of mechanisms behind CASTeers would help develop further methods of controlling generation of diffusion models. Finally, applicability of steering diffusion models to other tasks, such as image editing or controllable image generation, is not yet explored.

Code. We give a minimal version of development code in a zip file as part of the appendix. The code was developed and tested using 8 V100 GPUs.

B ALGORITHMS

We give the algorithms of our model. In Algorithm 1, we describe how we compute the steering vectors, while in Algorithm 2, we describe how we use them to perform concept erasure. The algorithms closely follow the descriptions in Sec. 3.2 and Sec 3.3.

Algorithm 1 Computing steering vectors

Require: Diffusion model DM with n CA layers, number of denoising steps T , concepts X, Y, P prompt pairs $(\mathcal{P}_p^X, \mathcal{P}_p^Y)$, $1 \leq p \leq P$, p_j^X containing X and p_j^Y containing Y , numbers of image patches per layer $\{m_i\}_{i=1}^n$
 Get $z_T \sim \mathbb{N}(0, I)$ a unit Gaussian random variable;
 $z_T^X \leftarrow z_T$
 $z_T^Y \leftarrow z_T$
for $p = 1 \dots, P$ **do**
 for $t = T, T-1, \dots, 1$ **do**
 $z_{t-1}^Y, \{ca_{itp}^Y\} \leftarrow DM(z_t^Y, \mathcal{P}_p^Y, t)$, $1 \leq i \leq n$
 $z_{t-1}^X, \{ca_{itp}^X\} \leftarrow DM(z_t^X, \mathcal{P}_p^X, t)$, $1 \leq i \leq n$
 end for
end for
 $ca_{it}^{X_avg} = \frac{\sum_{k=1}^{m_i} \sum_{p=1}^P ca_{itpk}^X}{Pm_i}$
 $ca_{it}^{Y_avg} = \frac{\sum_{k=1}^{m_i} \sum_{p=1}^P ca_{itpk}^Y}{Pm_i}$
 $ca_{it}^X = ca_{it}^{X_avg} - ca_{it}^{Y_avg}$
 $ca_{it}^X = \frac{ca_{it}^X}{\|ca_{it}^X\|_2}$ ▷ Normalize

Algorithm 2 Using steering vectors

Require: Diffusion model DM with n CA layers, number of denoising steps T , steering vectors ca_{it}^X for concept X to remove, input prompt $\mathcal{P}$, number of image patches on layers $\{m_i\}_{i=1}^n$, steering intensity β , flag of intermediate clipping do_clip
 Get $z_T \sim \mathbb{N}(0, I)$ a unit Gaussian random variable;
for $t = T, T-1, \dots, 1$ **do**
 for $i = 1, \dots, n$ **do**
 $z_{tmp}, ca_{itk}^{out} \leftarrow DM(z_t^Y, \mathcal{P}, t)$
 $\alpha \leftarrow \langle ca_{it}^{out}, ca_{it}^X \rangle$
 if do_clip **then**
 for $k = 1, \dots, m_i$ **do** ▷ Clipping dot product value for each image patch
 $\alpha_k \leftarrow \max(\alpha_k, 0)$
 end for
 end if
 $ca_{it}^{out_new} \leftarrow ca_{it}^{out} - \alpha ca_{it}^X$
 $z_{t-1}^Y \leftarrow DM(z_{tmp}, ca_{it}^{out})$ ▷ Continue inference
 end for
end for

C PROMPTS FOR GENERATING STEERING VECTORS

In this section, we describe the construction of prompt pairs that we use to compute steering vectors for our experiments.

C.1 PROMPTS FOR ERASING CONCRETE CONCEPTS

For erasing concrete concepts, we use prompt pairs of the form:

$$(\text{"p, with e"}, \text{"p"})$$

Here $\mathbf{p} \in \mathbf{P}$, where $\mathbf{P}$ is a set of N ImageNet classes, and $\mathbf{e}$ describes the concept we want to manipulate, e.g. we use $\mathbf{e} = \text{"Snoopy"}$ for *Snoopy* erasure, $\mathbf{e} = \text{"Mickey"}$ for *Mickey* erasure. Examples of prompts:

(*"junco, with Snoopy"*, *"junco"*)

(*"mud turtle, with Mickey"*, *"mud turtle"*)

Inside each prompt pair, the same generation seed is used.

Our main results in Tab. 17 are produced using steering vectors calculated on $N = 50$ prompts pairs.

C.2 PROMPTS FOR HUMAN-RELATED CONCEPTS

For manipulating abstract human-related concepts, we use prompt pairs of the form:

$$(\text{"b c, e"}, \text{"b c"})$$

Here $\mathbf{b} \in \mathbf{B}$ and $\mathbf{c} \in \mathbf{C}$, where

$\mathbf{B} = \{\text{"a girl", "a boy", "two men", "two women", "two people", "a man", "a woman", "an old man", "an old woman", "boys", "girls", "men", "women", "group of people", "a human",}\}$

$\mathbf{C} = \{\text{"", "gloomy image", "zoomed in", "talking", "on the street", "in a strange pose", "realism", "colorful background", "on a beach", "playing guitar", "enjoying nature", "smiling", "in a futuristic spaceship", "with kittens"}\}$,

and $\mathbf{e}$ describes the concept we want to manipulate.

$|\mathbf{b}| = 15$, $|\mathbf{C}| = 14$, which results in a total of 210 prompt pairs for each concept $\mathbf{e}$.

Following Receler, we use $\mathbf{e} = \text{"nudity"}$ for nudity erasure, and set of $\{\text{"hate", "harassment", "violence", "suffering", "humiliation", "harm", "suicide", "sexual", "nudity", "bodily fluids", "blood"}\}$ for harmful content erasure.

Examples of prompts for nudity erasure:

(*"a girl on a beach, nudity"*, *"a girl on a beach"*)

(*"boys talking, nudity"*, *"boys talking"*)

Inside each prompt pair, the same generation seed is used.

C.3 PROMPTS FOR STYLE MANIPULATION

For erasing or adding (Sec. I.3) style, we use prompt pairs of the form:

$$(\text{"p, e style"}, \text{"p"})$$

Here $\mathbf{p} \in \mathbf{P}$, where $\mathbf{P}$ is a set of ImageNet classes as used for concrete concept erasure, and $\mathbf{e}$ describes the style we want to manipulate, e.g. "baroque" or "Van Gogh".

Examples of prompts:

(*"junco, baroque style"*, *"junco"*)

(*"mud turtle, Van Gogh style"*, *"mud turtle"*)

Inside each prompt pair, the same generation seed is used.

D USER STUDIES

In this section, we provide user studies to complement quantitative and qualitative results of Casteer.

In the first user study, we compare SPM and Casteer on removing a concept of “Mickey” based on *explicit* prompts. For this study, we generated 800 images using “Mickey” prompts augmented with CLIP templates. Then, we randomly selected multiple sets of 20 pairs of images and provided them to our evaluators. We asked the users to select which model is best at removing the Mickey concept, by asking “Which image has the lower level of “Mickey” concept present?” and providing 3 options as answers: 1)“Image A”, 2)“Image B”, 3)“Both images completely removed the Mickey concept”. We show in Tab. 5 that our model was preferred for generating images without the Mickey concept, even when the concept was explicitly defined.

Table 5: User preferences for which model removed the *explicit* concept Mickey better. Our model was preferred for removing Mickey from images.

Model	Images Preferred	Percentage (%)
SPM	9	4.86%
CASteer	108	58.38%
Both	68	36.76
Total	185	100%

Next, we run another user study, this time testing removal of the concept of “Mickey” based on *implicit* prompts. For this study, we generated 100 images using prompts “A mouse from Disneyland” and “A Walt Disney’s most popular character”. Then, we randomly selected 20 image pairs and provided them to our evaluators. We asked the users to select which model is best at removing the Mickey concept, by asking “Which image has the lower level of “Mickey” concept present?” and providing 3 options as answers: 1)“Image A”, 2)“Image B”, 3)“Both images completely removed the Mickey concept”. Tab. 6 shows, CASteer was again preferred for generating images without the “Mickey” concept.

Table 6: User preferences for which model removed the *implicit* “Mouse from Disneyland” better. Our model was preferred for removing Mickey from images.

Model	Images Preferred	Percentage (%)
SPM	4	2.17%
CASteer	156	84.32%
Both	25	13.51
Total	185	100%

Based on the results above, we can see that users consider that CASteer removes specific concepts from generated images better than SPM. This matches our quantitative results presented in the paper.

Next, we compare CASteer, SPM and Receler models in removing the “Snoopy” concept, while preserving the concept of “Mickey”. First, we assess removal of the “Snoopy” concept. We generated 800 images using “Snoopy” prompts augmented with CLIP templates. Then, we randomly selected multiple sets of 20 image pairs and provided them to our evaluators. We asked them to select all the images where the concept of “Snoopy” is removed. Tab. 7 shows that both CASteer and Receler have high rates of “Snoopy” erasure. .

Table 7: Percentage of generated images containing Snoopy detected for each model.

Model	% Containing Snoopy ↓	% Without Snoopy ↑	% Total
Receler	0	100	100
CASteer	<u>3.3</u>	<u>96.7</u>	100
SPM	44.4	55.6	100

Lastly, we want to see how well SPM, Receler and CASteer preserve concepts “Mickey” when “Snoopy” is removed. We provide users with 3 images, one for each model, and ask them to rank

these images based on how much the “Mickey” concept is preserved in the images. We thus ask them “Rank the images in order from higher level of “Mickey” concept present (1) to lowest level of “Mickey” concept present (3).”. We give each user a randomly selected subset of 15 triplets of images. We can observe in Tab. 8 that SPM preserves the best the concept of Mickey, followed by CASteer. However, from Tab. 7 we see that SPM has a much higher tendency to preserve concepts, even the ones that should be removed, such as “Snoopy”. Our model has been considered by users a more reliable model for both removing some concepts and preserving the rest.

Table 8: User ranking of generated images by perceived Mickey content (1 = most Mickey, 3 = least Mickey). Lower scores indicate more Mickey content which is what we want to preserve.

Model	Total Rank Score ↓	Average Rank ↓
Receler	480	2.73
CASteer	<u>391</u>	<u>2.22</u>
SPM	185	1.05

E DETAILS ON STYLE REMOVAL TASK

The task focuses on the styles of five famous artists (Van Gogh, Picasso, Rembrandt, Warhol, Caravaggio) and the other uses five modern artists (McKernan, Kinkade, Edlin, Eng, Ajin: Demi Human). Following SAFREE, we use $LPIPS$ and prompt GPT-4o to identify an artist on generated images as evaluation metrics.

Following recent works, we experiment with removing one style of famous artist (Van Gogh) and one style of a modern artist (McKernan). We assess both how well our methods remove the desired style (e.g., Van Gogh) and preserve other styles in the same batch (e.g. Picasso, Rembrandt, Warhol and Caravaggio, if we remove Van Gogh). $LPIPS_e$ shows $LPIPS$ values for the images generated, when prompted with target style (“Van Gogh” or “Kelly McKernan”), and $LPIPS_u$, shows $LPIPS$ values for the images generated with the other 4 styles. Acc_e shows the accuracy of GPT-4o answers when asked to identify an artist on images generated with target style in prompt (“Van Gogh” or “Kelly McKernan”), and Acc_u shows the accuracy of GPT-4o answers when asked to identify an artist on images generated with other styles in prompt. The goal of any style removal method is to lower $LPIPS_e$ and Acc_e (i.e., successfully remove target style) while maintaining $LPIPS_u$ and Acc_u high (i.e., not affecting the generation of other styles).

Questions to GPT-4o are formulated as “Is this picture in s style? Just tell me Yes or No.”, where s represents style, e.g., s = “Andy Warhol” or s = “Van Gogh”. The number of “Yes” answers divided by the number of total answers is considered as Acc metric. We report results averaged over 3 runs for this metric.

F RESULTS ON SDXL MODEL

Here we provide qualitative and qualitative results on removing concrete and abstract concepts using CASTeers with the SDXL model.

F.1 EXPERIMENTAL SETUP

In SDXL experiments, we generate images using CASTeers with $\beta = 2$ on SDXL-base-1.0 model (stabilityai/stable-diffusion-xl-base-1.0) with steering vectors calculated on SDXL-Turbo model (stabilityai/sd-xl-turbo).

For calculation of steering vectors, we use fp16-version of SDXL-Turbo model. We generate images using default resolution, with one denoising step, using guidance scale=0.0 and seed=0. All other parameters are left default. To generate images, we use CASTeers on SDXL-base-1.0 model with 30 denoising steps. All other parameters are left default.

Generation of steering vector using 1 pair of prompts on SDXL-Turbo takes 8 seconds on V-100 GPU, i.e. generation of steering vectors for concrete concepts (see Sec. C) using 50 prompts takes 7 minutes, generation of steering vectors for human-related concepts (see Sec. C) using 210 prompts takes 28 minutes.

F.2 QUANTITATIVE RESULTS

We use the same experimental setups as for SD-1.4, described in Sec. 4. Results are shown in Tab. 9, 10, 11, 12, 13.

Table 9: **Quantitative results on nudity removal based on I2P Schramowski et al. (2023) dataset.** Detection of nude body parts is done by Nudenet at a threshold of 0.6. F: Female, M: Male. The best results are highlighted in bold, second-best are underlined.

Method	Nudity Detection								
	Breast(F)	Genitalia(F)	Breast(M)	Genitalia(M)	Buttocks	Feet	Belly	Armpits	Total↓
SDXL	85	7	3	2	7	28	84	66	282
Ours (w/o clip)	4	0	0	0	0	6	10	7	27
Ours (clip)	5	1	0	1	0	3	9	7	26

Table 10: **Quantitative results on inappropriate content removal based on I2PSchramowski et al. (2023) dataset.** Detection of inappropriate content is done by Q16 Schramowski et al. (2022).

Class name	Inappropriate proportion (%) (↓)		
	SDXL	Ours (w/o clip)	Ours (clip)
Hate	39.4	23.8	20.8
Harassment	33.0	21.7	21.7
Violence	43.7	25.9	24.7
Self-harm	42.4	24.3	24.1
Sexual	45.3	33.0	32.2
Shocking	49.6	32.0	30.8
Illegal activity	36.0	23.5	23.5
Overall	41.97	27.2	26.6

Table 11: **General quality estimation of images generated by CASTeers on SDXL model with nudity erasure.** CLIP score and FID are calculated on COCO-30k dataset

Method	Locality	
	CLIP-30K(↑)	FID-30K(↓)
SDXL	31.53	13.29
Ours	31.51	13.56
Ours (clip)	31.45	13.37

F.3 QUALITATIVE RESULTS

In this section, we provide qualitative results on CASTeers applied on SDXL model.

Table 12: Quantitative evaluation of concrete object erasure.

	Snoopy	Mickey	Spongebob	Pikachu	Dog	Legislator
Method	CS↓	CS↑ FID↓	CS↑ FID↓	CS↑ FID↓	CS↑ FID↓	CS↑ FID↓
SDXL	74.3	73.1 -	75.1 -	72.7 -	66.3 -	60.8
Ours	48.7	68.5 69.0	73.7 58.4	72.7 27.8	66.4 37.9	60.8 27.4
Ours (clip)	48.6	68.2 70.7	73.4 59.2	72.8 28.3	66.4 37.9	60.9 20.0

Table 13: Comparison of Artist Concept Removal tasks on SDXL model: Famous (left) and Modern artists (right).

Method	Remove “Van Gogh”				Remove “Kelly McKernan”			
	LPIPS _e ↑	LPIPS _u ↓	Acc _e ↓	Acc _u ↑	LPIPS _e ↑	LPIPS _u ↓	Acc _e ↓	Acc _u ↑
SDXL	-	-	0.95	0.95	-	-	1.00	0.86
Ours (w/o clip)	0.34	0.19	0.00	0.84	0.39	0.12	0.00	0.79
Ours (clip)	0.35	0.19	0.00	0.85	0.39	0.12	0.05	0.83

First, we show results on removing “Snoopy” concept when generating images with four prompt templates: “An origami X”, “A drawing of the X”, “A photo of a cool X” and “An art of the X”, where $X \in [\text{“Snoopy”, “Mickey”, “Spongebob”, “Pikachu”, “dog”, “legislator”}]$. Casteer is applied with removal strength $\beta = 2$.

We see that our method removes Snoopy well (see Fig. 7 while preserving other concepts well (see Fig. 8, 9, 10, 11, 12). In fact, most of the images of non-related concepts generated with Casteer applied are almost identical to those generated by vanilla SDXL.

Second, we show images generated on COCO-30k prompts with applied Casteer for nudity removal. We see that quality of generated images does not degrade, supporting quantitative results of Tab. 11 (see Fig. 13, 14, 15)

Next, on Fig. 16 and Fig. 17 we show examples of Casteer applied for style removal.



Figure 7: Images generated with “Snoopy” prompts with different seeds. Top: original SDXL, bottom: Casteer applied for removing the concept of “Snoopy”.



Figure 8: Images generated with “Mickey” prompts with different seeds. Top: original SDXL, bottom: Casteer applied for removing the concept of “Snoopy”.




Figure 9: Images generated with “*Spongebob*” prompts with different seeds. Top: original SDXL, bottom: CASteer applied for removing the concept of “*Snoopy*”.



Figure 10: Images generated with “*Pikachu*” prompts with different seeds. Top: original SDXL, bottom: CASteer applied for removing the concept of “*Snoopy*”.



Figure 11: Images generated with “*dog*” prompts with different seeds. Top: original SDXL, bottom: CASteer applied for removing the concept of “*Snoopy*”.



Figure 12: Images generated with “*legislator*” prompts with different seeds. Top: original SDXL, bottom: CASteer applied for removing the concept of “*Snoopy*”.



Figure 13: Images generated with different COCO-30k prompts with different seeds. Top: original SDXL, bottom: Casteer applied for removing the concept of “nudity”.



Figure 14: Images generated with different COCO-30k prompts with different seeds. Top: original SDXL, bottom: Casteer applied for removing the concept of “nudity”.



Figure 15: Images generated with different COCO-30k prompts with different seeds. Top: original SDXL, bottom: Casteer applied for removing the concept of “nudity”.



Figure 16: Examples of style removal for SDXL. Top: original SDXL, bottom: Casteer applied for removing the style of “*Van Gogh*”




Figure 17: Examples of style removal for SDXL. Top: original SDXL, bottom: Casteer applied for removing the style of “*mosaic style*” (left 4 images) and “*origami style*” (right 4 images)

G RESULTS ON SANA MODEL

G.1 EXPERIMENTAL SETUP

In SANA experiments, we generate images using Casteer with $\beta = 2$ on SANA_Sprint_1.6B_1024px_teacher model (Efficient-Large-Model/SANA_Sprint_1.6B_1024px_teacher_diffusers) with steering vectors calculated on Sana_Sprint_1.6B_1024px model (Efficient-Large-Model/Sana_Sprint_1.6B_1024px_diffusers).

For calculation of steering vectors, we use fp16-version of Sana_Sprint_1.6B_1024px_diffusers model. We generate images using default resolution, with one denoising step, using guidance seed=0. All other parameters are left default. To generate images, we use Casteer on SANA_Sprint_1.6B_1024px_teacher model with 20 denoising steps. All other parameters are left default.

Generation of steering vector using 1 pair of prompts on Sana_Sprint_1.6B_1024px takes 5 seconds on V-100 GPU, i.e. generation of steering vectors for concrete concepts (see Sec. C) using 50 prompts takes 4.2 minutes, generation of steering vectors for human-related concepts (see Sec. C) using 210 prompts takes 18 minutes.

G.2 QUANTITATIVE RESULTS

In this section, we present quantitative results on steering SANA model. We use the same experimental setups as for SD-1.4, described in Sec. 4. Results on SANA model are shown in Tab. 14, 15, 16, 17, 18.

Table 14: **Quantitative results on nudity removal based on I2P Schramowski et al. (2023) dataset.** Detection of nude body parts is done by Nudenet at a threshold of 0.6. F: Female, M: Male. The best results are highlighted in bold, second-best are underlined.

Method	Nudity Detection							
	Breast(F)	Genitalia(F)	Breast(M)	Genitalia(M)	Buttocks	Feet	Belly	Armpits
SANA	14	0	3	4	0	5	44	21
Ours (w/o clip)	3	0	0	1	0	0	3	2
Ours (clip)	0	0	0	1	0	0	1	0

Table 15: **Quantitative results on inappropriate content removal based on I2P Schramowski et al. (2023) dataset.** Detection of inappropriate content is done by Q16 Schramowski et al. (2022).

Class name	Inappropriate proportion (%) ($\downarrow$)		
	SANA	Ours (w/o clip)	Ours (clip)
Hate	48.1	48.5	35.0
Harassment	42.8	42.5	30.3
Violence	49.3	40.3	35.6
Self-harm	49.4	39.5	30.2
Sexual	36.6	32.6	22.3
Shocking	57.9	50.1	39.5
Illegal activity	42.5	36.9	25.7
Overall	46.1	40.1	30.5

Table 16: **General quality estimation of images generated by Casteer on SDXL model with nudity erasure.** CLIP score and FID are calculated on COCO-30k dataset

Method	Locality	
	CLIP-30K($\uparrow$)	FID-30K($\downarrow$)
SANA	29.28	22.65
Ours	28.79	22.89
Ours (clip)	29.01	23.51

G.3 QUALITATIVE RESULTS

In this section, we provide qualitative results on Casteer applied on SANA model.

Table 17: Quantitative evaluation of concrete object erasure.

	Snoopy	Mickey	Spongebob	Pikachu	Dog	Legislator			
Method	CS↓	CS↑	FID↓	CS↑	FID↓	CS↑	FID↓	CS↑	FID↓
SANA	79.7	76.1	-	79.0	-	74.0	-	68.1	-
Ours	48.2	74.9	94.3	78.1	68.0	74.0	45.2	68.0	48.6
Ours (clip)	48.2	75.0	96.5	78.1	71.0	74.0	46.0	68.0	49.3

Table 18: Comparison of Artist Concept Removal tasks on SANA model: Famous (left) and Modern artists (right).

	Remove “Van Gogh”				Remove “Kelly McKernan”			
Method	LPIPS _e ↑	LPIPS _u ↓	Acc _e ↓	Acc _u ↑	LPIPS _e ↑	LPIPS _u ↓	Acc _e ↓	Acc _u ↑
SANA	-	-	1.00	0.89	-	-	0.70	0.625
Ours (w/o clip)	0.42	0.14	0.25	0.88	0.19	0.11	0.40	0.620
Ours (clip)	0.42	0.14	0.25	0.875	0.19	0.08	0.30	0.620

First, we show results on removing “Snoopy” concept when generating images with four prompt templates: “An origami X”, “A drawing of the X”, “A photo of a cool X” and “An art of the X”, where $X \in [\text{“Snoopy”, “Mickey”, “Spongebob”, “Pikachu”, “dog”, “legislator”}]$. CASTeers is applied with removal strength $\beta = 2$.

We see that our method removes Snoopy well (see Fig. 18 while preserving other concepts well (see Fig. 19, 20, 21, 22, 23). In fact, most of the images of non-related concepts generated with CASTeers applied are almost identical to those generated by vanilla SANA.

Second, we show images generated on COCO-30k prompts with applied CASTeers for nudity removal. We see that quality of generated images does not degrade, supporting quantitative results of Tab. 16 (see Fig. 24, 25, 26). In some cases, visual quality of image generated with CASTeers exceeds that of original SANA (see Fig. 25, 3rd and 5th images for example).

Next, on Fig. 27 and Fig. 28 we show examples of CASTeers applied for style removal.



Figure 18: Images generated with “Snoopy” prompts with different seeds. Top: original SANA, bottom: CASTeers applied for removing the concept of “Snoopy”.




Figure 19: Images generated with “Mickey” prompts with different seeds. Top: original SANA, bottom: CASTeers applied for removing the concept of “Snoopy”.




Figure 20: Images generated with “*Spongebob*” prompts with different seeds. Top: original SANA, bottom: CASTeering applied for removing the concept of “*Snoopy*”.



Figure 21: Images generated with “*Pikachu*” prompts with different seeds. Top: original SANA, bottom: CASTeering applied for removing the concept of “*Snoopy*”.



Figure 22: Images generated with “*dog*” prompts with different seeds. Top: original SANA, bottom: CASTeering applied for removing the concept of “*Snoopy*”.



Figure 23: Images generated with “*legislator*” prompts with different seeds. Top: original SANA, bottom: CASTeering applied for removing the concept of “*Snoopy*”.



Figure 24: Images generated with different COCO-30k prompts with different seeds. Top: original SANA, bottom: CASTeering applied for removing the concept of “nudity”.



Figure 25: Images generated with different COCO-30k prompts with different seeds. Top: original SANA, bottom: CASTeering applied for removing the concept of “nudity”.



Figure 26: Images generated with different COCO-30k prompts with different seeds. Top: original SANA, bottom: CASTeering applied for removing the concept of “nudity”.



Figure 27: Examples of style removal for SANA. Top: original SANA, bottom: CASTeering applied for removing the style of “*Van Gogh*”



Figure 28: Examples of style removal for SANA. Top: original SANA, bottom: CASTeering applied for removing the style of “*mosaic style*” (left 4 images) and “*origami style*” (right 4 images)

H ABLATIONS ON HYPERPARAMETERS

In this section, we provide ablations on such hyperparameters as number and choice of prompt pairs to form a steering vector, a value of β , choice of the intermediate layer to steer, and importance of dot product weighting.

H.1 NUMBER AND CHOICE OF STEERING VECTORS

First, we compute steering vectors for the concept of “Snoopy” using SD-1.4 model on varying number of prompt pairs. Then we fix $\beta = 2$ and apply Casteer on SD-1.4 model using computed steering vectors on 800 prompts as described in Sec. 4. In particular, we augment each concept using 80 CLIP Radford et al. (2021) templates, and generate 10 for each concept-template pair, so that for each concept there are 800 images. We calculate CLIP Score (CS) Hessel et al. (2021) and FID Heusel et al. (2017) on these generated images as described in Sec. 4. Specifically, we use CS to estimate the level of the existence of the “Snoopy” concept within the generated images. Next, we calculate average FID Heusel et al. (2017) scores between the set of original generations of SD-1.4 model on prompts non-related to “Snoopy” and a set of generations of the steered model on these prompts. We use it to assess how much images of this related concept generated by the steered model differ from those of generated by the original model. Higher FID value demonstrate more severe generation alteration.

For each number of prompt pairs, we compute the steering vector three times using different non-intersecting sets of prompts and different random seeds. Thus, for each number of prompt pairs, we report three metric values.

Figures 29a and 29b show CS and FID metrics for different numbers of prompt pairs used for generation of steering vectors. We see that Casteer performance remains stable across different prompt sets, and using number of pairs 50 and above results in similar performance.




Figure 29: Ablation on number of prompts for computing steering vectors on SD-1.4. (a) CLIP score for a concept “Snoopy” calculated for images generated by Casteer using prompts containing “Snoopy”. (b) Averaged FID between original model generations for other concepts, and generations of a steered model. Blue line indicates mean values across three samples

H.2 IMPORTANCE OF DOT PRODUCT WEIGHTING

Our choice of steering vector weighting for concept erasure in Eq. 4 is motivated by an observation that dot product between Cross-Attention (CA) output c tokens and the steering vector s^X for a concept X serves as a good measure of amount of X that is present in each token of c . In this section, we provide experimental ablation on importance of this weighting. We fix α in Eq. 3 to be constant, and perform erasure of “Snoopy” concept using setup as described in Sec. 4. We measure CLIP score of “Snoopy” concept in the generated images based on prompts containing “Snoopy”, and CLIP score and FID of “Mickey” concept in the generated images based on prompts containing “Mickey”. We also report CLIP score and FID on images generated with Casteer for “nudity” erasure based on prompts from validation set of COCO-30k.

Tab. 19, 20 show results for $\alpha = 1, 2$ with and without clipping applied. We also report results of Casteer as reference. We see that for both values of α , “Snoopy” prompt is not erased well.

Table 19: **Quantitative evaluation of concrete object erasure.** The best results are highlighted in bold, second-best are underlined. Results of other methods are taken from SPMLyu et al. (2024) or reproduced.

Method	Snoopy	Mickey	
	CS↓	CS↑	FID↓
SD-1.4	78.5	74.7	-
Constant steering $\beta = 1$ (w/o clip)	76.3	74.3	51.5
Constant steering $\beta = 1$ (clip)	76.5	74.0	53.4
Constant steering $\beta = 2$ (w/o clip)	72.7	70.7	141.1
Constant steering $\beta = 2$ (clip)	72.8	70.8	139.5
Dot Product steering (CASteer, w/o clip)	45.8	70.4	93.0
Dot Product steering (CASteer, clip)	48.5	70.4	89.6

Table 20: **General quality estimation of images generated by nudity-erased models.** CLIP score and FID are calculated on COCO-30k dataset

Method	Locality	
	CLIP-30K(↑)	FID-30K(↓)
SD-1.4	31.34	14.04
Constant steering $\beta = 1$ (w/o clip)	29.56	14.07
Constant steering $\beta = 1$ (clip)	30.84	14.05
Constant steering $\beta = 2$ (w/o clip)	27.67	55.34
Constant steering $\beta = 2$ (clip)	28.22	55.21
Dot Product steering (CASteer, w/o clip)	30.69	13.28
Dot Product steering (CASteer, clip)	30.09	13.02

Meanwhile, for $\alpha = 2$, image fidelity and prompt alignment suffer, with FID being 3.5 times higher than that of CASteer. This suggests that our proposed weighting mechanism is crucial for performance of our method.

H.3 STEERING STRENGTH

Note that CASteer has a hyperparameter β , which determines the strength of steering for concept removal. While we set $\beta = 2$ in all our experiments, and this choice is motivated by the properties of the resulting transformation (see Sec. 4), it is still possible to use CASteer with varying values of strength β . Although experimentally we observe that value of $\beta = 2$ is optimal for all erasing scenarios and we use $\beta = 2$ in all our experiments, β still can be tuned for each use case, and varying values of β can lead to different trade-offs between level of target concept erasure and alteration of generated images not containing target concept. To show how β influences performance, in this section we provide results on “Snoopy” and “nudity” erasure using CASteer on SD-1.4, SDXL and SANA with different values of β .

Tab. 21 shows image quality metrics obtained using CASteer erasing “nudity” on COCO-30k validation dataset. It can be seen that steering with strength up to 2 does not lower general image quality for all models. Next, for SDXL and SANA models, steering with strength $\beta > 2$ up to $\beta = 4$ also does not affect general image quality, while for SD-1.4 steering with $\beta > 2$ results in degradation of image fidelity.

Tab. 22 and Tab. 23 show results on nudity erasure using varying steering strength β . We see that $\beta < 2$ results in a lesser suppression of the erased concept. In general, we see that as β increases, level of concept suppression is increased. However, for SANA with $\beta \geq 3$, and for SDXL with $\beta \geq 4$, CASteer without clipping starts showing lower rates of erasure. We hypothesise that this happens due to the absence of clipping, when image tokens with negative dot product with steering vector of “nudity” start containing substantial positive amount of “nudity” after steering. This hypothesis is supported by that CASteer with clipping does not show degrading performance with large values of β .

Table 21: **General quality estimation of images generated by Casteer on SD-1.4, SDXL and SANA model with nudity erasure.** CLIP score and FID are calculated on COCO-30k dataset

Method	SD-1.4		SDXL		SANA	
	CLIP-30K($\uparrow$)	FID-30K($\downarrow$)	CLIP-30K($\uparrow$)	FID-30K($\downarrow$)	CLIP-30K($\uparrow$)	FID-30K($\downarrow$)
Vanilla model	31.34	14.04	31.53	13.29	29.28	22.65
<i>Ours</i> , $\beta=1$	31.44	13.87	31.56	13.34	29.31	22.71
<i>Ours</i> , $\beta=1.5$	31.22	13.54	31.57	13.21	29.17	22.62
<i>Ours</i> , $\beta=2$	30.69	13.28	31.51	13.56	28.79	22.83
<i>Ours</i> , $\beta=2.5$	30.36	21.06	31.59	13.49	28.80	22.76
<i>Ours</i> , $\beta=3.0$	29.72	28.65	31.60	13.43	28.82	22.88
<i>Ours</i> , $\beta=4.0$	29.21	51.05	31.59	13.49	28.88	22.78
<i>Ours</i> , $\beta=1$ (clip)	31.42	13.69	31.52	13.27	29.34	22.72
<i>Ours</i> , $\beta=1.5$ (clip)	31.36	13.42	31.54	13.23	29.23	22.56
<i>Ours</i> , $\beta=2$ (clip)	31.09	13.02	31.45	13.37	29.01	22.87
<i>Ours</i> , $\beta=2.5$ (clip)	30.74	16.08	31.44	13.35	29.04	22.71
<i>Ours</i> , $\beta=3.0$ (clip)	30.14	20.11	31.44	13.32	20.03	22.87
<i>Ours</i> , $\beta=4.0$ (clip)	29.98	36.54	31.42	13.37	29.02	22.79

Table 22: **Quantitative results on nudity removal based on I2P Schramowski et al. (2023) dataset.** using Casteer on SANA model with varying strength β . Detection of nude body parts is done by Nudenet at a threshold of 0.6. F: Female, M: Male. The best results are highlighted in bold, second-best are underlined.

Method	Nudity Detection								
	Breast(F)	Genitalia(F)	Breast(M)	Genitalia(M)	Buttocks	Feet	Belly	Armpits	Total $\downarrow$
SANA	14	0	3	4	0	5	44	21	91
<i>Ours</i> $\beta = 1$	3	0	0	1	1	1	8	6	20
<i>Ours</i> $\beta = 1.5$	3	0	0	2	0	0	3	3	11
<i>Ours</i> $\beta = 2$	3	0	0	1	0	0	3	2	9
<i>Ours</i> $\beta = 2.5$	2	0	0	1	0	0	3	3	9
<i>Ours</i> $\beta = 3$	4	0	1	2	0	0	4	5	16
<i>Ours</i> $\beta = 4$	6	0	0	3	2	0	4	6	21
<i>Ours</i> $\beta = 1$ (clip)	1	0	0	0	1	2	8	5	17
<i>Ours</i> $\beta = 1.5$ (clip)	0	0	0	1	0	0	2	3	6
<i>Ours</i> $\beta = 2$ (clip)	0	0	0	1	0	0	1	0	2
<i>Ours</i> $\beta = 2.5$ (clip)	0	0	0	1	0	0	1	0	2
<i>Ours</i> $\beta = 3$ (clip)	1	0	0	0	0	0	1	0	2
<i>Ours</i> $\beta = 4$ (clip)	1	0	0	0	0	0	0	1	2

Table 23: **Quantitative results on nudity removal based on I2P Schramowski et al. (2023) dataset.** using Casteer on SDXL model with varying strength β . Detection of nude body parts is done by Nudenet at a threshold of 0.6. F: Female, M: Male. The best results are highlighted in bold, second-best are underlined.

Method	Nudity Detection								
	Breast(F)	Genitalia(F)	Breast(M)	Genitalia(M)	Buttocks	Feet	Belly	Armpits	Total $\downarrow$
SDXL	85	7	3	2	7	28	84	66	282
<i>Ours</i> $\beta = 1$	12	0	3	2	1	15	25	28	86
<i>Ours</i> $\beta = 1.5$	9	0	0	1	1	9	12	19	51
<i>Ours</i> $\beta = 2$	4	0	0	0	0	6	10	7	27
<i>Ours</i> $\beta = 2.5$	5	0	0	1	0	5	4	5	20
<i>Ours</i> $\beta = 3$	4	0	0	1	0	5	5	1	16
<i>Ours</i> $\beta = 4$	6	0	1	1	0	5	6	1	20
<i>Ours</i> $\beta = 1$ (clip)	15	0	1	1	1	16	27	28	89
<i>Ours</i> $\beta = 1.5$ (clip)	7	0	1	0	0	10	11	20	49
<i>Ours</i> $\beta = 2$ (clip)	5	1	0	1	0	3	9	7	26
<i>Ours</i> $\beta = 2.5$ (clip)	2	0	0	1	0	2	4	4	13
<i>Ours</i> $\beta = 3$ (clip)	3	0	1	2	0	3	2	2	13
<i>Ours</i> $\beta = 4$ (clip)	3	0	0	1	0	1	3	2	10

H.4 STEERING OTHER LAYERS

Here we provide results on steering other intermediate representations of DiT backbone rather than cross-attention (CA) output.

First, steering any part of DiT not inside CA does not result in the desired behaviour, producing completely out-of-distribution images. Next, we try to steer other parts of CA layer, namely, computing steering vectors and steering key and value vectors or steering outputs of individual attention heads. We steer key and value vectors here because they carry information from the input prompt.

We use CASTeering on SD-1.4 to erase concept of “Snoopy” from images generated on “Snoopy” and “Mickey” prompts. We calculate CLIP Score (CS) and FID as described above in Sec. H.1. Tab. 24 presents the results. We see that steering key and value vector has less effect on the desired concept, and steering outputs of individual attention heads provides roughly the same results.

Table 24: **Quantitative evaluation of steering other layers.**

Method	Snoopy	Mickey
	CS↓	FID↓
SDXL	74.3	-
Key-Value outputs ($\beta = 2$)	62.79	43.5
CA Heads outputs ($\beta = 2$)	48.94	67.3
CA outputs ($\beta = 2$)	48.7	68.5

H.5 STEERING ONLY FRACTION OF LAYERS

In all our experiments we apply CASTeering to all of the CA layers in the models (SD-1.4, SDXL or SANA). In this section, we provide qualitative experiments on steering only a fraction of CA layers.

We ablate on three ways of choosing a subset of CA layers for steering:

- Steering only k first CA layers, $0 \leq k \leq n$;
- Steering only k last layers, $0 \leq k \leq n$;
- Steering only k^{th} CA layer, $0 \leq k \leq n$.

We evaluate CASTeering on SDXL with $\beta = 2$ under these settings for erasing the concept of anger (Fig. 30, 31, 32). In all the figures we illustrate results for values of k between 0 and 60 with a step of 3 for compactness.

Our empirical findings suggest the following: 1) It is not sufficient to steer only one layer (see Fig. 32). It can be seen that the effect of steering any single layer is negligible, not causing the desired effect. 2) There is a trade-off between the level of expression of the desired concept in a resulting image and the alteration of general image layout and features. If we steer most of the layers, the overall layout may change drastically from that of the original image and it may cause in the change of identity or other features in the steered image compared to the original one (see Fig. 30, 31). As Fig. 31 suggests, steering only few last CA layers of the model results in a good trade-off between removing the unwanted concept and keeping other image details intact.

Based on our observations, we hypothesize that steering only few last CA layers of the model is sufficient for erasing texture-based concepts (e.g. styles), while not sufficient for erasing concrete concepts (e.g. “Snoopy”). We test this experimentally by applying CASTeering to only 36 last layers of SDXL model. We test erasing “Snoopy” concept and style of “Van Gogh”. Results are presented in Tab. 25, 26. We see that steering only 36 last layers of SDXL model successfully removes style of “Van Gogh”, while changing images of non-target styles less (see $LPIPS_u$). For the concept of “Snoopy”, CASTeering applied to only the last 36 layers of the model shows significantly worse erasure performance than CASTeering applied to all layers.

Table 25: Quantitative evaluation of concrete object erasure.

	Snoopy	Mickey	Spongebob	Pikachu	Dog	Legislator
Method	CS↓	CS↑ FID↓	CS↑ FID↓	CS↑ FID↓	CS↑ FID↓	CS↑ FID↓
SDXL	74.3	73.1 -	75.1 -	72.7 -	66.3 -	60.8
CASteer (last 36 layers)	57.2	70.8 44.0	74.5 35.2	72.7 19.2	66.3 15.0	60.8 17.3
CASteer	48.7	68.5 69.0	73.7 58.4	72.7 27.8	66.4 37.9	60.8 27.4

Table 26: Comparison of "Van Gogh" Removal tasks on SDXL model.

Method	Remove "Van Gogh"			
	LPIS _e ↑	LPIS _u ↓	Acc _e ↓	Acc _u ↑
SDXL	-	-	0.95	0.95
Ours (last 36 layers)	0.30	0.12	0.00	0.85
Ours	0.34	0.19	0.00	0.84



Figure 30: Results on removing the concept of “angry” using CASteer on the prompt “a realistic colorful portrait of an angry man” with steering only k last CA layers, $60 \geq k \geq 1$. Top left corner: image generated without CASteer, images from top to bottom, left to right: images generated using CASteer with varying k



Figure 31: Results on removing the concept of “angry” using CASTeering on the prompt “a realistic colorful portrait of an angry man” with steering only k first CA layers, $1 \leq k \leq 60$. Top left corner: image generated without CASTeering, images from top to bottom, left to right: images generated using CASTeering with varying k



Figure 32: Results on removing the concept of “angry” using Casteer on the prompt “a realistic colorful portrait of an angry man” with steering only CA layer number k , $1 \leq k \leq 60$. Top left corner: image generated without Casteer, images from top to bottom, left to right: images generated using Casteer with varying k

I OTHER TASKS

In this section, we provide evidence that using steering vectors calculated by CASTeering, it is possible not only to erase concepts from generated images, but to solve other tasks, such as *concept addition*, *concept flipping* and *concept interpolation*. We provide some quantitative and qualitative results on these tasks. However, we leave comprehensive analysis of CASTeering capabilities on these tasks for future work.

I.1 CONCEPT SWITCH

In this section, we propose a way to modify CASTeering to flip one concept on the image being generated into another, i.e. resulting generating concept Y when prompted to generate concept X .

I.1.1 METHOD

Recall from Sec. 3 that the value of the dot product $(ca_{it}^X \cdot ca_{itk}^{out})$ can be seen as amount of X that is present in ca_{itk}^{out} . Thus we propose the following technique for replacing one concept X for another concept Y . We first construct steering vectors ca_{it}^{XY} from X to Y using the same idea described in Sec. 3, with the only difference that the positive prompt contains X (“a girl with an apple”) while the negative one contains Y (“a girl with pear”). Then we use this steering vector to modify cross-attention outputs as:

$$ca_{itk}^{out_new} = ca_{itk}^{out} - 2\langle ca_{it}^{XY}, ca_{itk}^{out} \rangle ca_{it}^{XY}, \quad (9)$$

where $1 \leq k \leq \text{patch_num}_i$. This operation can be seen as removing all the information about concept X from ca_{itk}^{out} and instead adding the same amount of information about concept Y , that is, we reflect ca_{itk}^{out} relatively to the vector orthogonal to ca_{it}^{XY} . Consequently, during generation, if given the prompt “a girl with an apple”, the model will generate an image corresponding to “a girl with a pear”, effectively switching the concept of an apple with that of a pear.

Here we can also add steering strength by introducing parameter β :

$$ca_{itk}^{out_new} = ca_{itk}^{out} - \beta \langle ca_{it}^{XY}, ca_{itk}^{out} \rangle ca_{it}^{XY}, \quad (10)$$

Higher β will result in higher expression of concept Y in the resulting image. In our experiments, we observe that $\beta > 2$ sometimes is needed to completely switch concept X to Y .

Adding Intermediate clipping. Note that using Eq.9 we only want to influence those CA outputs ca_{itk}^{out} which have a positive amount of unwanted concept X in them, i.e. we only want to flip X to Y , and not on the opposite direction. As dot product $\langle ca_{it}^X, ca_{itk}^{out} \rangle$ measures the amount of X present in CA output ca_{itk}^{out} , we only want to steer those CA outputs ca_{itk}^{out} , which have a positive dot product with ca_{it}^X . So the equation becomes the following:

$$\begin{aligned} \alpha &= \max(\beta \langle ca_{it}^X, ca_{itk}^{out} \rangle, 0) \\ ca_{itk}^{out_new} &= ca_{itk}^{out} - \beta \alpha ca_{it}^X \end{aligned} \quad (11)$$

I.1.2 QUALITATIVE RESULTS

Here, we provide qualitative results on flipping different concepts. Fig. 33, 34, 35 visualize results of switching concepts of “Snoopy” to “Winnie-the-Pooh”, “cat” to “giraffe” and “caterpillar” to “butterfly”. All the images are generated using SDXL model with steering vectors obtained from SDXL-Turbo model.

I.2 CONCEPT ADDITION

In this section, we propose a way to use CASTeering to add desired concepts on the image being generated.

I.2.1 METHOD

Recall Eq. 3, which defined a mechanism of subtracting concept information from CA outputs of diffusion model to prevent generation of some concepts. We can use the same mechanism to instead



Figure 33: Switching between concepts of “Snoopy” and “Winnie-the-Pooh” using CASTeering. Top: images generated by vanilla SDXL, bottom: images with CASTeering applied.



Figure 34: Switching between concepts of “cat” and “giraffe” using CASTeering. Top: images generated by vanilla SDXL, bottom: images with CASTeering applied.

add concept to the outputs:

$$ca_{itk}^{out_new} = ca_{itk}^{out} + \alpha ca_{itk}^X, \quad (12)$$

Here $1 \leq k \leq \text{patch_num}_i$, and α is a hyperparameter that controls the strength of concept addition.

I.2.2 QUALITATIVE RESULTS

Here, we provide qualitative results on adding different concepts.

Fig. 33, 34, 35 visualize results of adding concepts of “apple”, “hat”, “clothes” and “happiness” using CASTeering. Note that for different concepts, different values of α are optimal.

I.3 STYLE TRANSFER

In this section, we propose a way to use CASTeering to change styles of real images or images being generated.



Figure 35: Switching between concepts of “caterpillar” and “butterfly” using Casteer. Top: images generated by vanilla SDXL, bottom: images with Casteer applied.



Figure 36: Adding concepts of “apple”, “hat”, “clothes” and “happiness” using Casteer. Top: images generated by vanilla SDXL, bottom: images with Casteer applied.

I.3.1 METHOD

Casteer performs Style Transfer on real images as follows: we apply the reverse diffusion process following DDIM Song et al. (2021) for t number of steps, where $1 \leq t \leq T$. Then we denoise image back using Casteer (addition algorithm, see Eq. 12). t controls the trade-off between loss of image details and intensity of style applied. In Fig. 37, 38, 39, 40 we show results on style transfer to four different styles with varying t . Intensity= 0.3 here means that $t = 0.3T$. With such a process, we often get satisfactory results with no major loss of details. However, when t is high, the loss of image content occurs (see the bottom lines of figures). This is because the inversion is not sufficiently accurate.




Figure 37: Examples of Style Transfer of real images into “anime” style. From top to bottom: original image, style transfer applied with intensities from 0.1 to 0.7 with a step of 0.1



Figure 38: Examples of Style Transfer of real images into “origami” style. From top to bottom: original image, style transfer applied with intensities from 0.1 to 0.7 with a step of 0.1



Figure 39: Examples of Style Transfer of real images into “Gothic Art” style. From top to bottom: original image, style transfer applied with intensities from 0.1 to 0.7 with a step of 0.1



Figure 40: Examples of Style Transfer of real images into “Retro Art” style. From top to bottom: original image, style transfer applied with intensities from 0.1 to 0.7 with a step of 0.1

J INTERPRETING STEERING VECTORS

J.1 STEERING VECTORS VISUALIZATION

In this section, we propose a way to interpret the meaning of steering vectors generated by CASTeering. Suppose we have steering vectors generated for a concept $\mathbf{X} \{ca_{it}^{\mathbf{X}}\}$, $1 \leq i \leq n$, $1 \leq t \leq T$, where n is the number of model layers and T is a number of denoising steps performed for generating steering vectors. To interpret these vectors, we prompt the diffusion model with a placeholder prompt “X” and at each denoising step, we substitute outputs of the model’s CA layers with corresponding steering vectors. This makes the diffusion model be only conditioned on the information from the steering vectors, completely suppressing other information from the text prompt.

Fig. 41 shows interpretations for the steering vectors of concepts “hat”, “polka dot dress”, “Snoopy”, “angry”, “happy”, from top to bottom. Note that vectors for concepts “hat”, “polka dot dress” and “Snoopy” were generated using prompt pair templates for concept deletion, i.e. pairs of the form (“fish with Snoopy”, “fish”), (“a girl with a hat”, “a girl”) (see Sec. C.1), and this is reflected in generated images, as they show these concepts not alone, but in a form of a girl in a hat, a girl in a polka dot dress or a boy with a Snoopy. As for the last two concepts (“angry”, “happy”), their steering vectors were generated using prompt pair templates for human-related concepts (see Sec. C.2), and they illustrate these concepts as is.

We note that images of each concept exhibit common features, e.g. all the images of hats and polka dots feature only female persons, and images corresponding to “angry” and “happy” concepts have certain styles. We believe that this reflects how diffusion models perceive different concepts and that this interpretation technique can be used for unveiling the hidden representations of concepts inside the diffusion model, but we leave it to future work.

J.2 UMAP ON STEERING VECTORS

We generate steering vectors for all vocabulary tokens of SDXL text encoders and apply UMap McInnes & Healy (2018) on these steering vectors. Fig. 42, 44, 43, 45 show that structure emerges in the space of these steering vectors, similar to that of Word2Vec Mikolov et al. (2013), supporting the hypothesis that steering vectors carry the meaning of the desired concept. This is observed for all the layers in the model.




Figure 41: Visualization of generations of the model conditioned only on steering vectors. Images in rows from top to down were generated using steering vectors for the concepts “hat”, “polka dot dress”, “Snoopy”, “angry”, “happy”.




Figure 42: Visualization of UMap applied to the steering vectors from the layer 17 of SDXL-Turbo formed by 3000 SDXL-Turbo vocabulary tokens.




Figure 43: Visualization of UMap applied to the steering vectors from the layer 17 of SDXL-Turbo formed by 3000 SDXL-Turbo vocabulary tokens.




Figure 44: Visualization of UMap applied to the steering vectors from the layer 35 of SDXL-Turbo formed by 3000 SDXL-Turbo vocabulary tokens.




Figure 45: Visualization of UMap applied to the steering vectors from the layer 35 of SDXL-Turbo formed by 3000 SDXL-Turbo vocabulary tokens.

K SPM AND DoCo VS CAsTeER ON ADVERSARIAL PROMPTS

In this section, we present more qualitative examples of CAsTeer outperforming SPM Lyu et al. (2024) and DoCo Wu et al. (2025) on adversarial prompts, i.e. prompts containing implicitly defined concepts. SD-1.4 is used as a backbone for both methods. We use prompts “A mouse from Disneyland” and “A girl with a mouse from Disneyland” to test erasing of concept of “Mickey”, prompts “A yellow Pokemon” and “A girl with a yellow Pokemon” to test erasing of concept of “Pikachu”, and prompt “A picture of Iron Man actor” to test erasing of concept of “Robert Downey Jr”. We use the code provided by SPM and DoCo to produce checkpoints and images for testing removal of the concepts of “Mickey” and “Pikachu”, and use checkpoint provided by DoCo to generate images for testing removal of a concept of “Robert Downey Jr”.

We see that SPM and DoCo fail to substantially remove targeted concepts in many cases, while CAsTeer shows consistent performance.



Figure 46: Examples of 8 generated images from CAsTeer and SPM when prompted “A mouse from Disneyland”. **Top:** generation of CAsTeer, **Middle:** generations of SPM, **Bottom:** generations of DoCo.

We use the same diffusion hyperparameters and seeds when generating corresponding images from CAsTeer, SPM and DoCo.



Figure 47: Examples of 8 generated images from CAsTeer and SPM when prompted “A girl with a mouse from Disneyland”. “A mouse from Disneyland”. **Top:** generation of CAsTeer, **Middle:** generations of SPM, **Bottom:** generations of DoCo.

We use the same diffusion hyperparameters and seeds when generating corresponding images from CAsTeer, SPM and DoCo.



Figure 48: Examples of 8 generated images from CASTeers and SPM when prompted “A yellow *Pokemon*”. **Top:** generation of CASTeers, **Bottom:** generations of SPM. We use the same diffusion hyperparameters and seeds when generating corresponding images from CASTeers and SPM.



Figure 49: Examples of 8 generated images from CASTeers and SPM when prompted “A girl with a yellow *Pokemon*”. **Top:** generation of CASTeers, **Bottom:** generations of SPM. We use the same diffusion hyperparameters and seeds when generating corresponding images from CASTeers and SPM.



Figure 50: Examples of 8 generated images from CASTeers and SPM when prompted “A girl with a yellow *Pokemon*”. **Top:** generation of CASTeers, **Bottom:** generations of SPM. We use the same diffusion hyperparameters and seeds when generating corresponding images from CASTeers and SPM.



Figure 51: Examples of 8 generated images from CASTeers and DoCo when prompted “A girl with a yellow *Pokemon*”. **Top:** generation of CASTeers, **Bottom:** generations of DoCo. We use the same diffusion hyperparameters and seeds when generating corresponding images from CASTeers and DoCo.

L EXPRESSIVE POWER OF STEERING CROSS-ATTENTION OUTPUTS OVER STEERING PROMPT EMBEDDINGS

In this section we formalize a local, first-order theorem showing that applying linear interventions directly in the space of cross-attention (CA) outputs of diffusion models is always at least as expressive as applying linear interventions in the space of prompt token embeddings, and is strictly more expressive whenever the Jacobian of transformation from embeddings to CA space is not surjective. We first formulate the setup, prove a formal theorem, and then give interpretation of this result applied to diffusion models.

L.1 SETUP

Let $e \in \mathbb{R}^n$ denote the (flattened) text embedding vector produced by the text encoder, and let $h \in \mathbb{R}^m$ denote the (flattened) concatenation of all cross-attention outputs across all chosen layers, heads, and spatial positions.

Around a given forward pass with embedding e_0 , consider linearization of the dependence of CA outputs h on the embeddings e :

$$h \approx h_0 + G(e - e_0), \quad (13)$$

Here $G = J_{h,e}(e_0) \in \mathbb{R}^{m \times n}$ is Jacobian. Consider also a continuous loss function $\mathcal{L} : \mathbb{R}^m \rightarrow \mathbb{R}$ that measures the degree to which a concept is present in the final output. For a perturbation Δh , the perturbed loss becomes $\mathcal{L}(h_0 + \Delta h)$.

Now consider perturbations in CA space which are constrained by the magnitude of the perturbation by a radius $\rho > 0$:

$$\|\Delta h\| \leq \rho. \quad (14)$$

Under these setting, we compare two steering mechanisms:

Steering of Text Embeddings. A perturbation δe in the text embedding space induces perturbation in the CA output space $\Delta h = G\delta e$. Then the feasible set of CA-space perturbations reachable through embedding steering is:

$$\mathcal{H}_{\text{embed}}(\rho) := \{ \Delta h = G\delta e \mid \delta e \in \mathbb{R}^n, \|G\delta e\| \leq \rho \}. \quad (15)$$

We refer to it as *embedding-feasible set*. And the optimal achievable loss is:

$$V_{\text{embed}}(\rho) := \inf_{\Delta h \in \mathcal{H}_{\text{embed}}(\rho)} \mathcal{L}(h_0 + \Delta h). \quad (16)$$

Steering of Cross-Attention Outputs. Here CA outputs are perturbed directly:

$$\mathcal{H}_{\text{CA}}(\rho) := \{ \Delta h \in \mathbb{R}^m \mid \|\Delta h\| \leq \rho \}, \quad (17)$$

We refer to it as *CA-feasible set*. Then the optimal achievable loss becomes:

$$V_{\text{CA}}(\rho) := \inf_{\Delta h \in \mathcal{H}_{\text{CA}}(\rho)} \mathcal{L}(h_0 + \Delta h). \quad (18)$$

L.2 MAIN RESULT

Theorem L.1 (Expressive Power of CA Steering). *For any Jacobian $G \in \mathbb{R}^{m \times n}$ and any radius $\rho > 0$, the embedding-feasible set is a subset of the CA-feasible set:*

$$\mathcal{H}_{\text{embed}}(\rho) \subseteq \mathcal{H}_{\text{CA}}(\rho). \quad (19)$$

Moreover, equality ($\mathcal{H}_{\text{embed}}(\rho) = \mathcal{H}_{\text{CA}}(\rho)$ for all $\rho > 0$) holds if and only if G is surjective, i.e. $\text{rank}(G) = m$.

Proof. If $\Delta h \in \mathcal{H}_{\text{embed}}(\rho)$, then by definition $\Delta h = G\delta e$ for some δe and $\|\Delta h\| = \|G\delta e\| \leq \rho$. Hence $\Delta h \in \mathcal{H}_{\text{CA}}(\rho)$. Thus $\mathcal{H}_{\text{embed}}(\rho) \subseteq \mathcal{H}_{\text{CA}}(\rho)$.

If G is surjective, then for any $\Delta h \in \mathbb{R}^m$ there exists δe such that $G\delta e = \Delta h$; hence the two feasible sets coincide for every radius.

Conversely, if $\text{rank}(G) < m$, then $\text{im}(G)$ is a proper subspace of $\mathbb{R}^m$. For any $\rho > 0$, the ball $\{\Delta h : \|\Delta h\| \leq \rho\}$ contains points outside $\text{im}(G)$; hence $\mathcal{H}_{\text{embed}}(\rho)$ is a strict subset of $\mathcal{H}_{\text{CA}}(\rho)$. $\square$

Corollary L.1.1 (Optimal CA Steering Never Underperforms). *Assume $\mathcal{L}$ is continuous and bounded below. For any $\rho > 0$:*

$$V_{\text{CA}}(\rho) \leq V_{\text{embed}}(\rho). \quad (20)$$

If $\mathcal{H}_{\text{embed}}(\rho) \subsetneq \mathcal{H}_{\text{CA}}(\rho)$ and there exists $\Delta h^ \in \mathcal{H}_{\text{CA}}(\rho) \setminus \overline{\mathcal{H}_{\text{embed}}(\rho)}$ such that $\mathcal{L}(h_0 + \Delta h^*) < V_{\text{embed}}(\rho)$, then the inequality is strict:*

$$V_{\text{CA}}(\rho) < V_{\text{embed}}(\rho). \quad (21)$$

Proof. Since $\mathcal{H}_{\text{embed}}(\rho) \subseteq \mathcal{H}_{\text{CA}}(\rho)$, the infimum over the larger set is no greater than the infimum over the subset, establishing $V_{\text{CA}}(\rho) \leq V_{\text{embed}}(\rho)$.

If there exists a feasible Δh^* in the larger set achieving strictly lower loss than any point in the closure of the embedding-feasible set, then clearly the infimum of $\mathcal{L}$ over the larger set is strictly lower. $\square$

L.3 INTERPRETATION FOR DIFFUSION MODELS

In text-to-image (T2I) diffusion models (SD-1.4, SDXL, SANA), text embeddings e are projected into key/value vectors, which is then used by Cross-Attention layers. Perturbing the embeddings therefore induces changes in all cross-attention modules simultaneously via the Jacobian G .

However, Theorem L.1 shows that steering directly in CA output space allows for perturbations that *cannot* be produced by any embedding shift, because the image of G typically forms a low-dimensional subspace of the full CA space ($\text{im}(G)$). Consequently, Corollary L.1.1 guarantees that an idealized CA steering method (in the linearized regime) matches or outperforms any embedding-based method on any concept-erasure objective $\mathcal{L}$.

Intuitively, CA output is closer to the residual stream of DiT, which produces resulting noise as part of denoising process. So interventions in CA output space that aim to influence information in the residual stream can be more precise, than interventions in the text embedding space. Interventions in the text embedding space need to be designed so that they result in desired changes in the residual stream, after being passed through some function (CA layer itself). Intuitively, this is harder to optimize than intervention which operates directly in CA output space.

Another intuitive point of view on steering CA layers outputs is the following: output of CA layer is a point where concepts from the prompt are already entangled with context and spatial structure, so steering here can result in better control over spatial layouts of desired objects. In contrast, prompt embedding space is discrete and do not have direct relations to spatial regions in the resulting image.

Also note, that CA activations encode *layer-wise*, *head-wise*, and *spatially localized* information about how individual tokens influence different regions of the image Hertz et al. (2023). Steering in CA space therefore can provide fine-grained control, e.g. attenuating a concept only in specific spatial regions or at specific attention heads, which cannot be achieved by global modifications to text embeddings.

M ERASING MULTIPLE CONCEPTS

In this section, we provide results on erasing multiple concepts simultaneously.

Note that there can be multiple approaches to erasing multiple concepts based on individual steering vectors for each concept. First approach is to use averaged steering vectors for multiple concepts:

$$ca_{it}^X = \frac{\sum_{i=1}^n ca_{it}^{X^i}}{n}. \quad (22)$$

Here $\{X^i\}_{i=1}^n$ are concepts to delete. We use this approach in Sec. 4, where we report results on erasing harmful content based on I2P dataset (see Tab. 15). More precisely, we obtain steering vectors for harmful content erasure as average of steering vectors for the concepts “hate”, “harassment”, “violence”, “self-harm”, “sexual”, “shocking”, “illegal activity”.

However, for the concepts that do not share much common semantics (e.g., “Snoopy” and “nudity”) simple averaging may lead to poor performance. Thus, we propose a second approach for erasing multiple concepts. Suppose we have concepts $\{X^i\}_{i=1}^n$ to erase and their corresponding steering vectors $ca_{it}^{X^i}$. We apply Gram–Schmidt algorithm for orthogonalizing these vectors, and get a new, mutually-orthogonal set of steering vectors $\hat{ca}_{it}^{X^i}$. Then we apply CASTeering using this orthogonalized set of vectors $\hat{ca}_{it}^{X^i}$. Note that if CASTeering is applied without clipping, then every steering operation with each steering vector can be formulated in a matrix form (Eq. 5), and, consequently, erasure transform with multiple steering vectors can also be expressed as a single matrix multiplication.

We apply CASTeering with orthogonalization on erasing pairs of concepts (“Snoopy”, “nudity”) and (“Snoopy”, “Mickey”). Results for erasure of “Snoopy” and “nudity” are presented in Tab. 27 and Tab. 28; results for erasure of “Snoopy” and “Mickey” are presented in Tab. 29. We see that CASTeering successfully erases both concepts, achieving performance comparable to that of CASTeering applied to erasing individual concepts.

Table 27: **Quantitative results on nudity removal based on I2P (Schramowski et al. (2023)) dataset.** with CASTeering applied on SD-1.4 for 2 concepts erasure: “nudity” and “Snoopy”. Detection of nude body parts is done by Nudenet at a threshold of 0.6. F: Female, M: Male.

Method	Nudity Detection								
	Breast(F)	Genitalia(F)	Breast(M)	Genitalia(M)	Buttocks	Feet	Belly	Armpits	Total↓
SD-1.4	183	21	46	10	44	42	171	129	646
Ours (w/o clip)	9	1	0	2	2		0	0	14
Ours (clip)	8	2	0	1	2	1	0	1	15

Table 28: **Quantitative evaluation of concrete object erasure with CASTeering applied on SD-1.4 for 2 concepts erasure: “nudity” and “Snoopy”.**

	Snoopy		Mickey		Spongebob		Pikachu		Dog		Legislator	
Method	CS↓	FID↓	CS↑	FID↓	CS↑	FID↓	CS↑	FID↓	CS↑	FID↓	CS↑	FID↓
SD-1.4	78.5	74.7	-	74.1	-	74.7	-	65.2	-	61.0	-	-
Ours	50.3	70.4	111.1	71.7	113.2	75.4	61.04	65.6	57.80	61.4	99.92	
Ours (clip)	50.4	70.4	111.0	71.6	111.7	74.7	53.34	65.3	52.19	60.8	74.78	

Table 29: **Quantitative evaluation of concrete object erasure with CASTeering applied on SD-1.4 for 2 concepts erasure: “Snoopy” and “Mickey”.**

	Snoopy		Mickey		Spongebob		Pikachu		Dog		Legislator	
Method	CS↓	FID↓	CS↑	FID↓	CS↑	FID↓	CS↑	FID↓	CS↑	FID↓	CS↑	FID↓
SD-1.4	78.5	74.7	74.1	-	74.7	-	65.2	-	61.0	-	-	-
Ours	52.5	55.8	68.5	107.0	74.7	58.5	65.8	62.3	60.9	76.3		
Ours (clip)	50.4	54.8	68.3	106.9	74.8	54.8	65.7	51.2	61.0	76.4		

N SD-1.5 ON STEERING VECTORS FROM SD-1.4

In this section, we provide results on CASTeers applied to SD-1.5 model using steering vectors computed on SD-1.4 model. We use SD-1.5 with 50 denoising steps to match the number of steps used in SD-1.4.

We report results on all our main experiments setups. Tab. 33, 32, 30, 31, report results on erasing harmful content, artistic styles, and concepts of “nudity” and “Snoopy”, respectively. Additionally, Tab. 34 reports image quality metrics based on CLIP and FID scores. Results show that steering vectors used in CASTeers can be successfully transferred across different trained versions of models without losing in general image quality.

Table 30: **Quantitative results on nudity removal based on I2P Schramowski et al. (2023) dataset.** on SD-1.5 model with steering vectors computed on SD-1.4 model. Detection of nude body parts is done by Nudenet at a threshold of 0.6. F: Female, M: Male. The best results are highlighted in bold, second-best are underlined.

Method	Nudity Detection								
	Breast(F)	Genitalia(F)	Breast(M)	Genitalia(M)	Buttocks	Feet	Belly	Armpits	Total↓
SD-1.5	180	11	21	10	20	20	131	92	485
Ours (w/o clip)	2	0	0	5	0	0	1	0	8
Ours (clip)	3	1	0	2	1	2	0	0	9

Table 31: **Quantitative evaluation of concrete object erasure.**

Method	Snoopy		Mickey		Spongebob		Pikachu		Dog		Legislator	
	CS↓	CS↑	FID↓	FID↑	CS↓	CS↑	FID↓	FID↑	CS↓	CS↑	FID↓	FID↑
SD-1.5	0.77	0.74	-	-	0.74	-	0.73	-	0.65	-	0.61	-
Ours	0.51	0.72	87.9	0.72	93.9	0.73	44.1	0.65	45.2	0.61	65.4	0.61
Ours (clip)	0.50	0.71	92.2	0.71	103.3	0.72	43.0	0.65	41.1	0.61	56.0	0.61

Table 32: Comparison of Artist Concept Removal tasks on SD-1.5 model: Famous (left) and Modern artists (right).

Method	Remove “Van Gogh”				Remove “Kelly McKernan”			
	LPIPS _e ↑	LPIPS _u ↓	Acc _e ↓	Acc _u ↑	LPIPS _e ↑	LPIPS _u ↓	Acc _e ↓	Acc _u ↑
SD-1.5	-	-	1.00	0.975	-	-	0.95	0.925
Ours (w/o clip)	0.45	0.31	0.40	0.89	0.49	0.28	0.00	0.80
Ours (clip)	0.43	0.30	0.2	0.85	0.49	0.28	0.00	0.825

Table 33: **Quantitative results on inappropriate content removal based on I2PSchramowski et al. (2023) dataset.** Detection of inappropriate content is done by Q16 Schramowski et al. (2022).

Class name	Inappropriate proportion (%) ($\downarrow$)		
	SDXL	Ours (w/o clip)	Ours (clip)
Hate	38.1	38.5	32.5
Harassment	36.5	34.3	29.2
Violence	49.9	39.7	35.1
Self-harm	48.3	37.5	33.0
Sexual	57.0	38.9	40.1
Shocking	57.4	48.1	45.7
Illegal activity	36.3	33.7	27.8
Overall	47.6	38.9	35.4

Table 34: **General quality estimation of images generated by Casteer on SDXL model with nudity erasure.** CLIP score and FID are calculated on COCO-30k dataset

Method	Locality	
	CLIP-30K($\uparrow$)	FID-30K($\downarrow$)
SD-1.5	26.42	
<i>Ours</i>	26.54	
<i>Ours (clip)</i>	26.48	

O SD-1.4 ON STEERING VECTORS COMPUTED ON SINGLE DENOISING STEP

In this section, we provide results on Casteer applied to SD-1.4 model using steering vectors from a single (first) denoising step of SD-1.4 model.

We report results on all our main experiments setups. Tab. 38, 37, 35, 36, report results on erasing harmful content, artistic styles, and concepts of “nudity” and “Snoopy”, respectively. Additionally, Tab. 39 reports image quality metrics ased on CLIP and FID scores. Results show that Casteer can be successfully applied on models that do not have a distilled analogue using steering vectors computed on a single denoising step.

Table 35: **Quantitative results on nudity removal based on I2P (Schramowski et al. (2023)) dataset.** Detection of nude body parts is done by Nudenet at a threshold of 0.6. F: Female, M: Male. The best results are highlighted in bold, second-best are underlined.

Method	Nudity Detection								
	Breast(F)	Genitalia(F)	Breast(M)	Genitalia(M)	Buttocks	Feet	Belly	Armpits	Total
SD v1.4	183	21	46	10	44	42	171	129	646
DoCo Wu et al. (2025)	162	29	48	63	64	122	168	250	906
Ablating (CA) Kumari et al. (2023)	298	22	67	7	45	66	180	153	838
FMN Zhang et al. (2024a)	155	17	19	2	12	59	117	43	424
ESD-x Gandikota et al. (2023)	101	6	16	10	12	37	77	53	312
SLD-Med Schramowski et al. (2023)	39	<u>1</u>	26	3	3	21	72	47	212
UCE Gandikota et al. (2024)	35	5	11	4	7	29	62	29	182
SA Heng & Soh (2023)	39	9	4	0	15	32	49	15	163
ESD-x Gandikota et al. (2023)	14	<u>1</u>	8	5	5	24	31	33	121
Receler Huang et al. (2024)	13	<u>1</u>	12	9	5	10	26	39	115
MACE Lu et al. (2024)	16	0	9	7	2	39	19	17	109
RECE Gong et al. (2024)	8	0	6	4	0	8	23	17	66
CPE (one word) Lee et al. (2025)	11	2	3	2	5	15	13	15	66
CPE (four word) Lee et al. (2025)	6	<u>1</u>	3	2	2	8	8	10	40
AdvUnlearn Zhang et al. (2024c)	1	<u>1</u>	0	0	0	13	0	8	23
SAeUron Cywiński & Deja (2025)	<u>4</u>	0	0	<u>1</u>	3	2	<u>1</u>	7	18
Ours (single step, w/o clip)	9	1	0	5	1	3	4	4	27
Ours (single step, clip)	5	0	0	2	0	2	0	1	<u>10</u>
Ours (w/o clip)	5	0	0	<u>1</u>	3	2	0	<u>1</u>	12
Ours (clip)	<u>4</u>	0	0	<u>1</u>	<u>2</u>	0	0	0	7

Table 36: **Quantiative evaluation of concrete object erasure.** The best results are highlighted in bold, second-best are underlined. Results of other methods are taken from SPMLyu et al. (2024) or reproduced.

	Snoopy	Mickey	Spongebob	Pikachu	Dog	Legislator
Method	CS↓	CS↑ FID↓	CS↑ FID↓	CS↑ FID↓	CS↑ FID↓	CS↑ FID↓
SD-1.4	78.5	74.7 -	74.1 -	74.7 -	65.2 -	61.0 -
ESD Gandikota et al. (2023)	48.3	58.0 121.0	64.0 104.7	68.6 68.3	63.9 49.5	59.9 50.9
SPM Lyu et al. (2024)	60.9	74.4 22.1	74.0 21.4	74.6 12.4	65.2 9.0	61.0 5.5
SAFREE (Yoon et al., 2024)	54.7	68.1 72.8	70.2 76.6	71.9 46.2	65.4 70.9	59.9 55.4
Receler (reg=0.1) (Huang et al., 2024)	45.7	55.6 143.5	59.6 156.2	63.5 121.9	64.0 68.9	60.6 42.7
Receler (reg=1.0) (Huang et al., 2024)	49.1	63.4 105.9	62.5 128.5	71.5 62.8	64.0 50.2	60.7 40.1
DoCo (Wu et al., 2025)	49.1	74.4 31.1	73.8 21.6	74.7 14.5	65.2 10.1	60.9 5.6
Ours (single step, w/o clip)	48.6	68.9 96.0	69.4 97.9	74.6 50.3	65.8 52.1	61.1 62.8
Ours (single step, clip)	50.8	69.1 95.8	69.4 98.2	74.3 46.6	65.4 49.8	60.9 52.5
Ours (w/o clip)	45.8	70.4 93.0	72.4 81.4	74.0 38.3	66.0 31.8	61.0 40.9
Ours (clip)	48.5	70.4 89.6	72.5 81.4	73.7 34.4	65.7 31.8	60.8 37.1

Table 37: Casteer on SD-1.4 with steering vectors computed on single denoising step Comparison of Artist Concept Removal tasks: Famous (left) and Modern artists (right).

Method	Remove “Van Gogh”				Remove “Kelly McKernan”			
	LPIPS _e ↑	LPIPS _u ↓	Acc _e ↓	Acc _u ↑	LPIPS _e ↑	LPIPS _u ↓	Acc _e ↓	Acc _u ↑
SD-v1.4	-	-	0.95	0.95	-	-	0.80	0.83
CA (Kumari et al., 2023)	0.30	0.13	0.65	0.90	0.22	0.17	0.50	0.76
RECE (Gong et al., 2024)	0.31	<u>0.08</u>	0.80	<u>0.93</u>	0.29	<u>0.04</u>	0.55	0.76
UCE (Gandikota et al., 2024)	0.25	0.05	0.95	0.98	0.25	0.03	0.80	0.81
SLD-Medium (Schramowski et al., 2023)	0.21	0.10	0.95	0.91	0.22	0.18	0.50	0.79
SAFREE (Yoon et al., 2024)	0.42	0.31	0.35	0.85	<u>0.40</u>	0.39	0.40	0.78
Ours (single step, w/o clip)	0.48	0.34	0.15	0.84	0.56	0.30	0.00	0.78
Ours (single step, clip)	<u>0.47</u>	0.33	<u>0.20</u>	0.79	<u>0.55</u>	0.28	0.00	0.75
Ours (w/o clip)	0.46	0.31	0.35	0.88	0.54	0.27	<u>0.05</u>	0.80
Ours (clip)	0.44	0.30	0.25	0.86	0.54	0.28	<u>0.05</u>	0.81

Table 38: **Quantitative results on inappropriate content removal based on I2P (Schramowski et al. (2023)) dataset. Detection of inappropriate content is done by Q16 (Schramowski et al. (2022)) classifier.** The best results are highlighted in bold, second-best are underlined.

Class name	Inappropriate proportion (%) ($\downarrow$)											
	SD	FMN	Ablating	ESD-x	SLD	ESD-u	UCE	Receler	Ours (single step, w/o clip)	Ours (single step, clip)	Ours (w/o clip)	Ours (clip)
Hate	44.2	37.7	40.8	34.1	22.5	26.8	36.4	28.6	<u>28.1</u>	25.5	35.5	29.0
Harassment	37.5	25.0	32.9	30.2	22.1	24.0	29.5	21.7	28.8	25.7	29.9	<u>25.6</u>
Violence	46.3	47.8	43.3	40.5	31.8	35.1	34.1	27.1	31.2	30.3	32.5	<u>27.8</u>
Self-harm	47.9	46.8	47.4	36.8	30.0	33.7	30.8	<u>24.8</u>	27.1	24.1	26.1	<u>26.2</u>
Sexual	60.2	59.1	60.3	40.2	52.4	35.0	25.5	29.4	34.8	33.9	<u>23.0</u>	20.7
Shocking	59.5	58.1	57.8	45.2	40.5	40.1	41.1	<u>34.8</u>	39.4	37.7	38.4	34.0
Illegal activity	40.0	37.0	37.9	28.9	22.1	26.7	29.0	<u>21.3</u>	24.1	23.2	21.5	17.6
Overall	48.9	47.8	45.9	36.6	33.7	32.8	31.3	<u>27.0</u>	31.2	29.4	28.9	25.6

Table 39: **General quality estimation of images generated by nudity-erased models.** CLIP score and FID are calculated on COCO-30k dataset

Method	Locality	
	CLIP-30K($\uparrow$)	FID-30K($\downarrow$)
SD v1.4	31.34	14.04
FMN	30.39	13.52
CA	31.37	16.25
AdvUn	28.14	17.18
Receler	30.49	15.32
MACE	29.41	13.42
CPE	31.19	13.89
UCE	30.85	14.07
SLD-M	30.90	16.34
ESD-x	30.69	14.41
ESD-u	30.21	15.10
SAeUron	30.89	14.37
<i>Ours (single step, w/o clip)</i>	30.67	13.33
<i>Ours (single step, clip)</i>	31.05	13.14
<i>Ours (w/o clip)</i>	30.69	<u>13.28</u>
<i>Ours (clip)</i>	<u>31.09</u>	13.02