

565 **A Appendix**

566 **A.1 Proof of Theorem 1**

567 As illustrated in Sec. 3.2 it is hard to build the unlearned data  $x^u$  for the feature unlearning since  
 568 adding the perturbation may influence the model accuracy seriously. Suppose the feature is success-  
 569 fully removed when the norm of perturbation is larger than  $C$ . We define the utility loss  $\ell_1$  with  
 570 unlearning feature successfully:

$$\ell_1 = \min_{\|\delta_{\mathcal{F}}\| \geq C} \mathbb{E}_{(x,y) \in \mathcal{D}} \min_{\theta} \ell(f_{\theta}(x + \delta_{\mathcal{F}}), y) \quad (10)$$

571 And we define the maximum utility loss with the norm perturbation less than  $C$  as:

$$\ell_2 = \max_{\|\delta_{\mathcal{F}}\| \leq C} \mathbb{E}_{(x,y) \in \mathcal{D}} \min_{\theta} \ell(f_{\theta}(x + \delta_{\mathcal{F}}), y) \quad (11)$$

572 **Assumption 3.** Assume  $\ell_2 \leq \ell_1$

573 Assumption 3 elucidates that the utility loss associated with a perturbation norm less than  $C$  is smaller  
 574 than the utility loss when the perturbation norm is greater than  $C$ . This assumption is logical, as  
 575 larger perturbations would naturally lead to greater utility loss.

576 **Assumption 4.** Suppose the federated model achieves zero training loss.

577 We have the following theorem to elucidate the relation between feature sensitivity removing via  
 578 Algo. 1 and exact unlearning (see proof in Appendix).

579 **Theorem 2.** If Assumption 3 and 4 hold, the utility loss of unlearned model obtained by Algo. 1 is  
 580 less than the utility loss with unlearning successfully, i.e.,

$$\ell_u \leq \ell_1, \quad (12)$$

581 where  $\ell_u = \mathbb{E}_{(x,y) \in \mathcal{D}} (\ell(f_{\theta^u}(x), y))$

582 *Proof.* When the unlearning happens during the federated training, the unlearning clients would  
 583 also optimize the training loss and feature sensitivity simultaneously. Specifically, the optimization  
 584 process could be written as:

$$\theta_u = \arg \min_{\theta} \mathbb{E}_{(x,y) \in \mathcal{D}} (\ell(f_{\theta}(x), y) + \lambda \mathbb{E}_{\delta_{\mathcal{F}}} \frac{\|f_{\theta}(x) - f_{\theta}(x + \delta_{\mathcal{F}})\|_2}{\|\delta_{\mathcal{F}}\|_2}),$$

585 where  $\lambda \geq \frac{1}{C}$  is one coefficient. Without loss of generality, we assume the  $\ell(f_{\theta}(x), y) = \|f_{\theta}(x) - y\|$ .  
 586 Denote

$$\Theta^* = \arg \min_{\theta} \mathbb{E}_{(x,y) \in \mathcal{D}} \ell(f_{\theta}(x), y).$$

587 If Assumption 4 holds, then  $f_{\theta^*}(x) = y$  for any  $\theta^* \in \Theta^*$ . Therefore, for any  $\|\delta_{\mathcal{F}}\| \geq \frac{1}{\lambda}$  such that

$$\begin{aligned} & \mathbb{E}_{(x,y) \in \mathcal{D}} (\ell(f_{\theta^*}(x), y) + \lambda \mathbb{E}_{\|\delta_{\mathcal{F}}\| \geq \frac{1}{\lambda}} \frac{\|f_{\theta^*}(x) - f_{\theta^*}(x + \delta_{\mathcal{F}})\|_2}{\|\delta_{\mathcal{F}}\|_2}) \\ &= \lambda \mathbb{E}_{(x,y) \in \mathcal{D}} \mathbb{E}_{\|\delta_{\mathcal{F}}\| \geq \frac{1}{\lambda}} \frac{\|y - f_{\theta^*}(x + \delta_{\mathcal{F}})\|_2}{\|\delta_{\mathcal{F}}\|_2} \\ &\leq \mathbb{E}_{(x,y) \in \mathcal{D}} \mathbb{E}_{\|\delta_{\mathcal{F}}\| \geq \frac{1}{\lambda}} \|y - f_{\theta^*}(x + \delta_{\mathcal{F}})\|_2. \end{aligned} \quad (13)$$

588 The last inequality is due to Therefore, we further obtain:

$$\begin{aligned}
\ell_u &\leq \min_{\theta \in \mathbb{R}^d} \mathbb{E}_{(x,y) \in \mathcal{D}} \left( \ell(f_\theta(x), y) + \lambda \mathbb{E}_{\|\delta_{\mathcal{F}}\| \geq \frac{1}{\lambda}} \frac{\|f_\theta(x) - f_\theta(x + \delta_{\mathcal{F}})\|_2}{\|\delta_{\mathcal{F}}\|_2} \right) \\
&\leq \min_{\theta \in \Theta^*} \mathbb{E}_{(x,y) \in \mathcal{D}} \left( \ell(f_\theta(x), y) + \lambda \mathbb{E}_{\|\delta_{\mathcal{F}}\| \geq \frac{1}{\lambda}} \frac{\|f_\theta(x) - f_\theta(x + \delta_{\mathcal{F}})\|_2}{\|\delta_{\mathcal{F}}\|_2} \right) \\
&\leq \min_{\theta \in \Theta^*} \mathbb{E}_{(x,y) \in \mathcal{D}} \mathbb{E}_{\|\delta_{\mathcal{F}}\| \geq \frac{1}{\lambda}} \|y - f_{\theta^*}(x + \delta_{\mathcal{F}})\|_2 \\
&\leq \mathbb{E}_{(x,y) \in \mathcal{D}} \mathbb{E}_{\|\delta_{\mathcal{F}}\| \geq \frac{1}{\lambda}} \min_{\theta \in \Theta^*} \|y - f_{\theta^*}(x + \delta_{\mathcal{F}})\|_2 \\
&= \mathbb{E}_{\|\delta_{\mathcal{F}}\| \geq \frac{1}{\lambda}} \mathbb{E}_{(x,y) \in \mathcal{D}} \min_{\theta \in \Theta^*} \|y - f_{\theta^*}(x + \delta_{\mathcal{F}})\|_2 \\
&\leq \max_{\|\delta_{\mathcal{F}}\| \geq \frac{1}{\lambda}} \mathbb{E}_{(x,y) \in \mathcal{D}} \min_{\theta \in \mathbb{R}^d} \|y - f_{\theta^*}(x + \delta_{\mathcal{F}})\|_2 \\
&\leq \max_{\|\delta_{\mathcal{F}}\| \leq C} \mathbb{E}_{(x,y) \in \mathcal{D}} \min_{\theta \in \mathbb{R}^d} \|y - f_{\theta^*}(x + \delta_{\mathcal{F}})\|_2 \\
&= \ell_2,
\end{aligned} \tag{14}$$

589 where the last inequality is due to  $\lambda \geq \frac{1}{C}$ . According to Assumption 3, we have  $\ell_u \leq \ell_1$

590

□

## 591 A.2 Experimental Setup

592 **Datasets** *MNIST* [90]: Both the *MNIST* [90] and *Fashion-MNIST(FMNIST)* [92] datasets contain  
593 images of handwritten digits and attire, respectively. Each dataset comprises 60,000 training examples  
594 and 10,000 test examples. In both datasets, each example is represented as a single-channel image  
595 with dimensions of 28x28 pixels, categorized into one of 10 classes. Additionally, the *Colored-*  
596 *MNIST(CMNIST)* [90] dataset, an extension of the original MNIST, introduces color into the digits of  
597 each example. Consequently, images in the Colored MNIST dataset are represented in three channels.  
598 *CIFAR* [93]: The *CIFAR-10* [93] dataset comprises 60,000 images, each with dimensions of 32x32  
599 pixels and three color channels, distributed across 10 classes. This dataset includes 6,000 images  
600 per class and is partitioned into 50,000 training examples and 10,000 test examples. Similarly, the  
601 *CIFAR-100* [93] dataset shares the same image dimensions and structure as *CIFAR-10* but extends to  
602 100 classes, with each class containing 600 images. Within each class, there are 500 training images  
603 and 100 test images. Moreover, *CIFAR-100* organizes its 100 classes into 20 superclasses, forming  
604 the *CIFAR-20 dataset* [93]. *CelebA* [85]: A face recognition dataset featuring 40 attributes such as  
605 gender and facial characteristics, comprising 162,770 training examples and 19,962 test examples.  
606 This study will focus on utilizing the *CelebA* [85] dataset primarily for gender classification tasks.

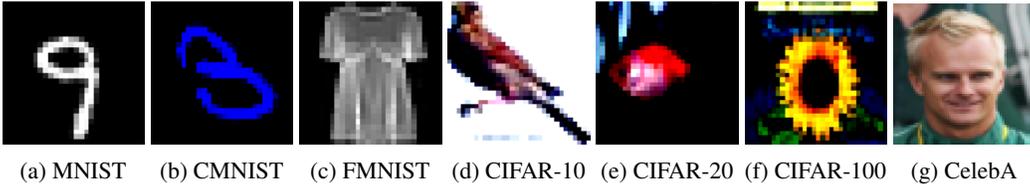


Figure 8: Visual representation of dataset samples utilized in this study.

607 *Adult Census Income (Adult)* [86] includes 48, 842 records with 14 attributes such as age, gender,  
608 education, marital status, etc. The classification task of this dataset is to predict if a person earns over  
609 \$50K a year based on the census attributes. We then consider marital status as the sensitive feature  
610 that aim to unlearn in this study. *Diabetes* [87] includes 768 personal health records of females at least  
611 21 years old with 8 attributes such as blood pressure, insulin level, age and etc. The classification  
612 task of this dataset is to predict if a person has diabetes. We then consider number of pregnancies as  
613 the sensitive feature that aim to unlearn in this study.

614 **Baselines** The baseline methods in this study:

615 *Baseline*: Original model before unlearning.

616 *Retrain*: In scenarios involving sensitive feature unlearning, the retrained model was simply trained  
 617 using a dataset where Gaussian noise was applied to the unlearned feature region. This approach  
 618 may lead to performance deterioration, as discussed in Sec. 3.2. For backdoor feature unlearning  
 619 scenarios, the retrained model was trained using the retain dataset  $\mathcal{D}_r$ , also referred to as the clean  
 620 dataset. In biased feature unlearning scenarios, the retrained model was trained using a combination  
 621 of 50% from each of the retain dataset  $\mathcal{D}_r$  (bias dataset) and the unlearn client local dataset  $\mathcal{D}_u$   
 622 (unbias dataset). This ensures fairness in the model’s performance across both datasets.

623 *Fine-tune*: The baseline model is fine-tuned using the retained dataset  $\mathcal{D}_r$  for 5 epochs. *Class-*  
 624 *Discriminative Pruning(FedCDP)*[66]: A FU framework that achieves class unlearning by utilizing  
 625 Term Frequency-Inverse Document Frequency (TF-IDF) guided channel pruning, which selectively  
 626 removes the most discriminative channels related to the target category and followed by fine-tuning  
 627 without retraining from scratch.

628 *FedRecovery*[62]: A FU framework that achieves client unlearning by removing the influence of a  
 629 client’s data from the global model using a differentially private machine unlearning algorithm that  
 630 leverages historical gradient submissions without the need for retraining.

### 631 A.3 Attention Map

#### 632 A.3.1 Backdoor Feature Unlearning

633 Attention map analysis for backdoor samples across model iterations of baseline, retrain, and unlearn  
 634 model using our proposed Ferrari method on MNIST(Fig. 9), FMNIST(Fig. 10), CIFAR-10(Fig. 11),  
 635 CIFAR-20(Fig. 12) and CIFAR-100 (Fig. 13)datasets.

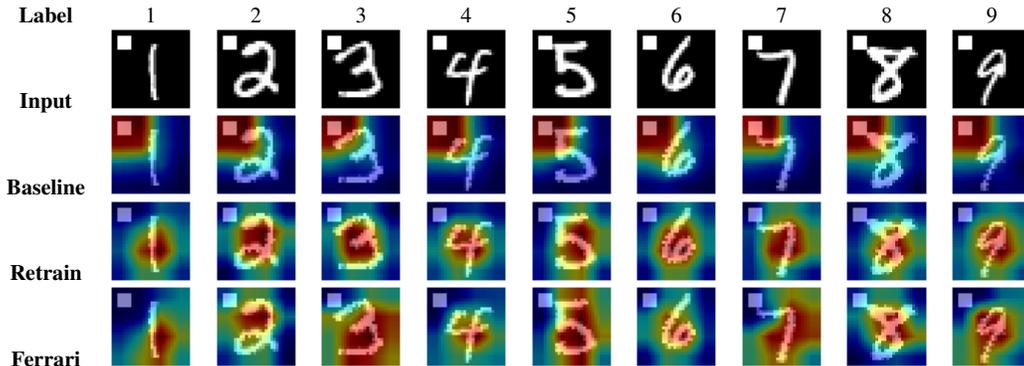


Figure 9: MNIST

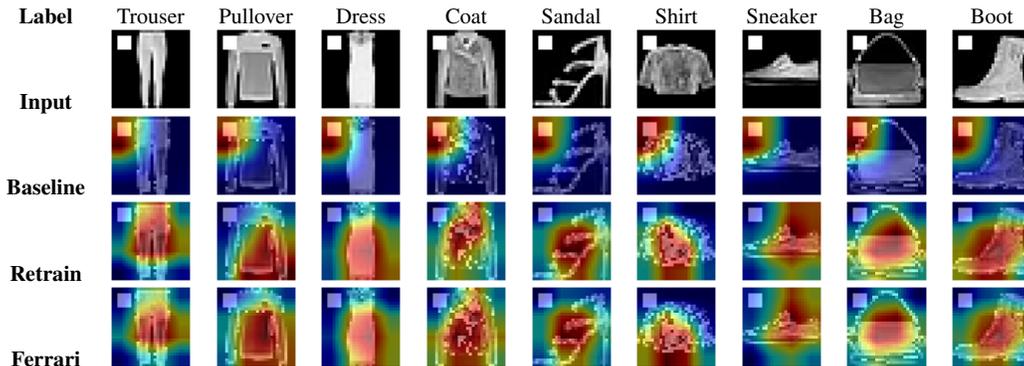


Figure 10: FMNIST

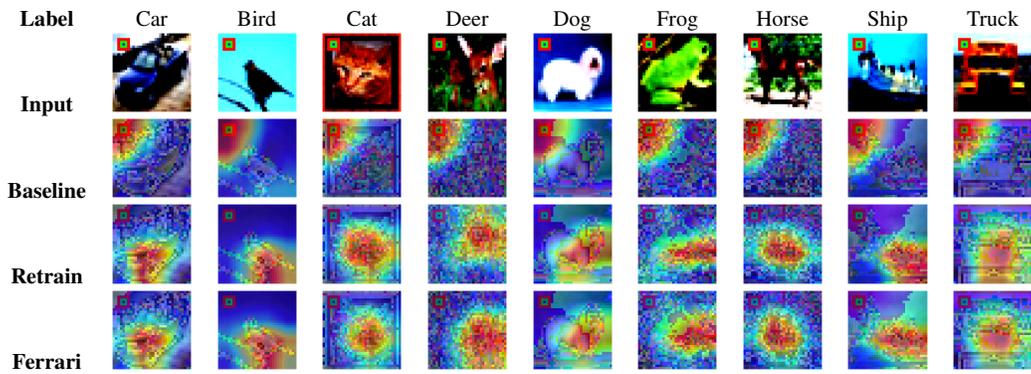


Figure 11: CIFAR-10

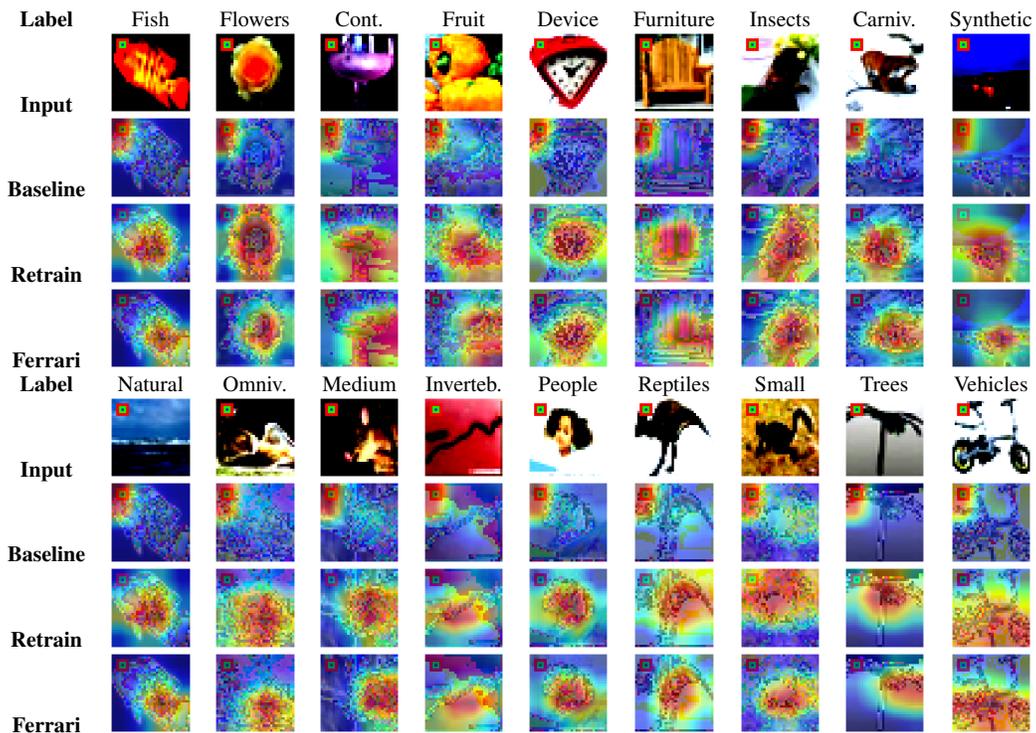


Figure 12: CIFAR-20

<b>Label</b>	Fish	Baby	Bear	Beaver	Bed	Bee	Beetle	Bicycle	Bottle
<b>Input</b>									
<b>Baseline</b>									
<b>Retrain</b>									
<b>Ferrari Label</b>									
<b>Label</b>	Boy	Bridge	Bus	B.fly	Camel	Can	Castle	C.plar	Cattle
<b>Input</b>									
<b>Baseline</b>									
<b>Retrain</b>									
<b>Ferrari Label</b>									
<b>Label</b>	Chimpz.	Clock	Cloud	C.krch	Couch	Crab	Croc.	Cup	Dino.
<b>Input</b>									
<b>Baseline</b>									
<b>Retrain</b>									
<b>Ferrari Label</b>									
<b>Label</b>	E.phant	F.fish	Forest	Fox	Girl	Hamster	House	K.groo	K.board
<b>Input</b>									
<b>Baseline</b>									
<b>Retrain</b>									
<b>Ferrari Label</b>									
<b>Label</b>	Mower	Leopard	Lion	Lizard	Lobster	Man	Maple	M.cycle	Mountain
<b>Input</b>									
<b>Baseline</b>									
<b>Retrain</b>									
<b>Ferrari</b>									

Continued from previous page

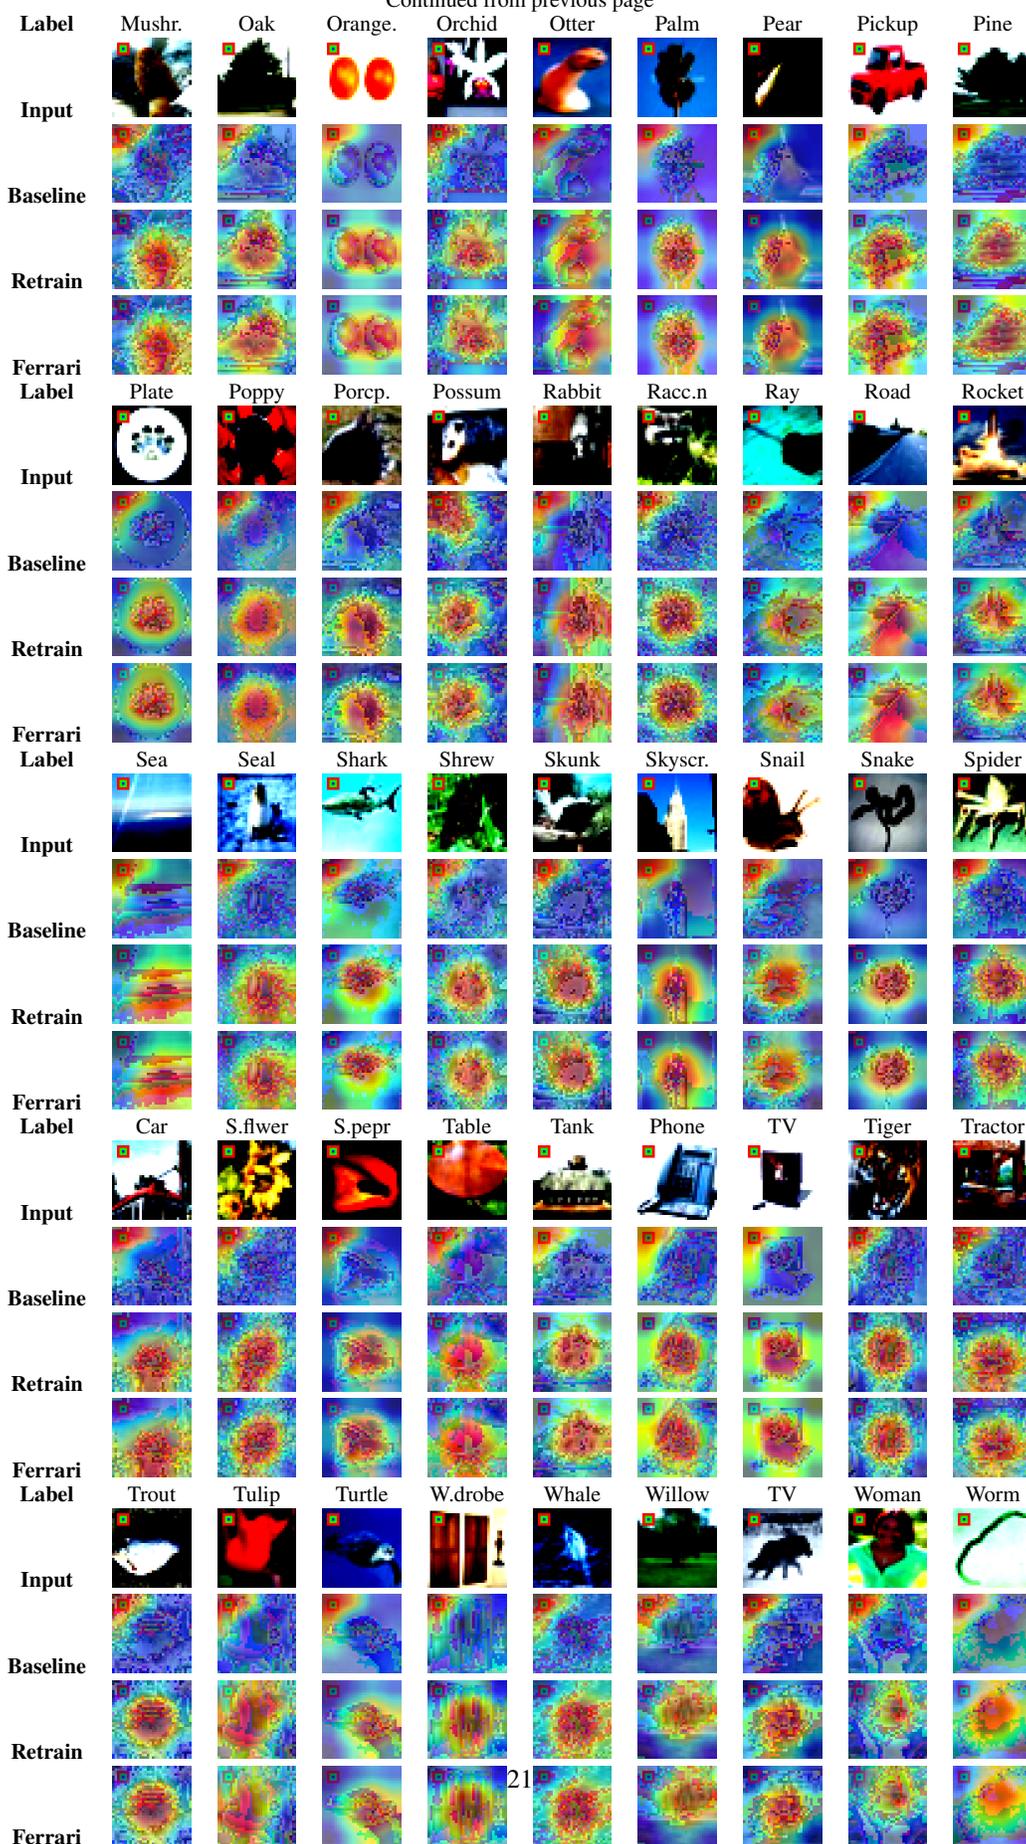
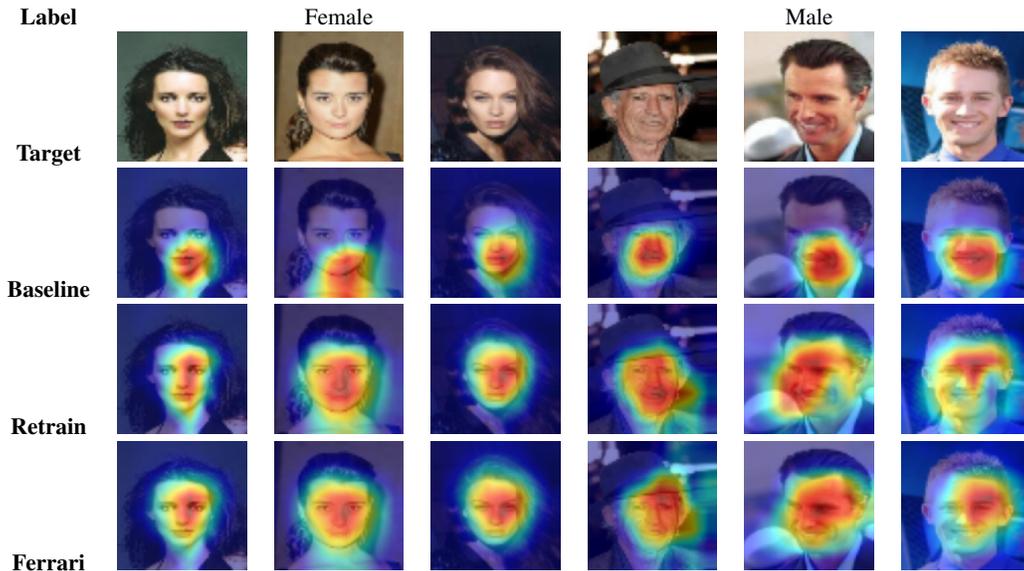
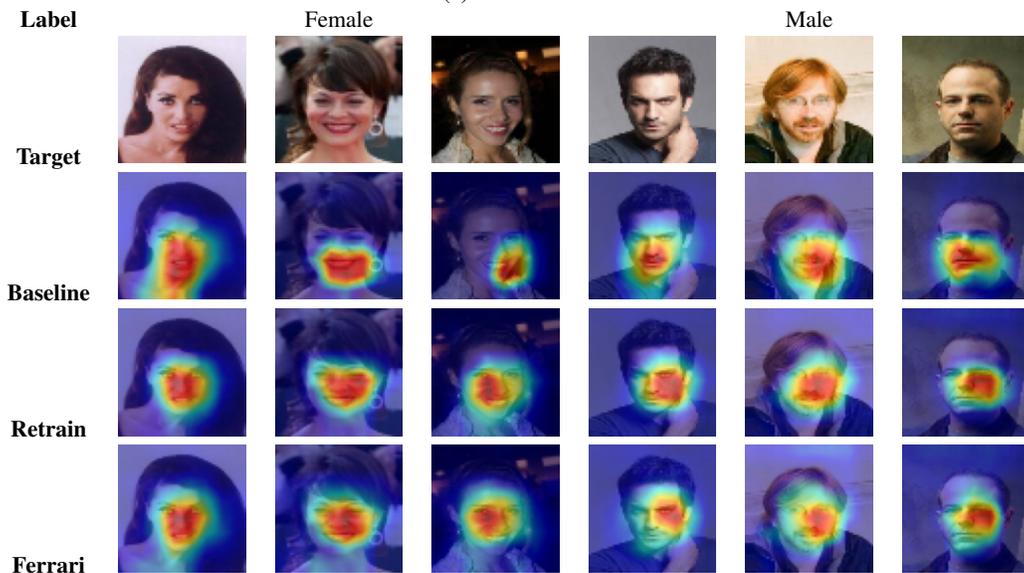


Figure 13: CIFAR-100

636 **A.3.2 Biased Feature Unlearning**



(a) Bias Dataset



(b) Unbias Dataset

Figure 14: Attention map analysis for bias and unbias samples across model iterations of baseline, retrain, and unlearn model using our proposed Ferrari to unlearn 'mouth' on CelebA dataset.

637 **A.4 Limitation and Future Work**

638 While our proposed approach of federated feature unlearning demonstrates effectiveness in various  
 639 unlearning scenarios using only the local dataset of unlearning clients without requiring participation  
 640 from other clients, thus simulating practical application, it has some inevitable limitations.

641 The proposed approach necessitates access to the entire dataset from the unlearning client to achieve  
 642 maximal unlearning effectiveness. However, as demonstrated in Section 5.5, a partial dataset  
 643 comprising at least 70% of the data yields similar performance to the full dataset. In certain cases, the  
 644 unlearning client may lose a significant portion of their data, rendering our approach ineffective in  
 645 such scenarios. Therefore, future work should investigate federated feature unlearning approaches that

646 require only a small portion of the unlearning client’s dataset. Additionally, the proposed approach  
647 has only been proven effective for classification models, as it was specifically designed for this  
648 purpose. Its effectiveness in other domains, such as generative models, remains to be investigated.

649 Therefore, future work should explore methods that require only a small portion of the client’s dataset.  
650 Additionally, future research will investigate advanced perturbation techniques, support for diverse  
651 data types and models, and integration with other privacy-preserving methods to further enhance data  
652 protection in FL systems.

## 653 **NeurIPS Paper Checklist**

### 654 **1. Claims**

655 Question: Do the main claims made in the abstract and introduction accurately reflect the  
656 paper's contributions and scope?

657 Answer: [\[Yes\]](#)

658 Guidelines:

- 659 • The answer NA means that the abstract and introduction do not include the claims  
660 made in the paper.
- 661 • The abstract and/or introduction should clearly state the claims made, including the  
662 contributions made in the paper and important assumptions and limitations. A No or  
663 NA answer to this question will not be perceived well by the reviewers.
- 664 • The claims made should match theoretical and experimental results, and reflect how  
665 much the results can be expected to generalize to other settings.
- 666 • It is fine to include aspirational goals as motivation as long as it is clear that these goals  
667 are not attained by the paper.

### 668 **2. Limitations**

669 Question: Does the paper discuss the limitations of the work performed by the authors?

670 Answer: [\[Yes\]](#)

671 Guidelines:

- 672 • The answer NA means that the paper has no limitation while the answer No means that  
673 the paper has limitations, but those are not discussed in the paper.
- 674 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 675 • The paper should point out any strong assumptions and how robust the results are to  
676 violations of these assumptions (e.g., independence assumptions, noiseless settings,  
677 model well-specification, asymptotic approximations only holding locally). The authors  
678 should reflect on how these assumptions might be violated in practice and what the  
679 implications would be.
- 680 • The authors should reflect on the scope of the claims made, e.g., if the approach was  
681 only tested on a few datasets or with a few runs. In general, empirical results often  
682 depend on implicit assumptions, which should be articulated.
- 683 • The authors should reflect on the factors that influence the performance of the approach.  
684 For example, a facial recognition algorithm may perform poorly when image resolution  
685 is low or images are taken in low lighting. Or a speech-to-text system might not be  
686 used reliably to provide closed captions for online lectures because it fails to handle  
687 technical jargon.
- 688 • The authors should discuss the computational efficiency of the proposed algorithms  
689 and how they scale with dataset size.
- 690 • If applicable, the authors should discuss possible limitations of their approach to  
691 address problems of privacy and fairness.
- 692 • While the authors might fear that complete honesty about limitations might be used by  
693 reviewers as grounds for rejection, a worse outcome might be that reviewers discover  
694 limitations that aren't acknowledged in the paper. The authors should use their best  
695 judgment and recognize that individual actions in favor of transparency play an impor-  
696 tant role in developing norms that preserve the integrity of the community. Reviewers  
697 will be specifically instructed to not penalize honesty concerning limitations.

### 698 **3. Theory Assumptions and Proofs**

699 Question: For each theoretical result, does the paper provide the full set of assumptions and  
700 a complete (and correct) proof?

701 Answer: [\[Yes\]](#)

702 Guidelines:

- 703 • The answer NA means that the paper does not include theoretical results.

- 704 • All the theorems, formulas, and proofs in the paper should be numbered and cross-  
705 referenced.
- 706 • All assumptions should be clearly stated or referenced in the statement of any theorems.
- 707 • The proofs can either appear in the main paper or the supplemental material, but if  
708 they appear in the supplemental material, the authors are encouraged to provide a short  
709 proof sketch to provide intuition.
- 710 • Inversely, any informal proof provided in the core of the paper should be complemented  
711 by formal proofs provided in appendix or supplemental material.
- 712 • Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 713 4. Experimental Result Reproducibility

714 Question: Does the paper fully disclose all the information needed to reproduce the main ex-  
715 perimental results of the paper to the extent that it affects the main claims and/or conclusions  
716 of the paper (regardless of whether the code and data are provided or not)?

717 Answer: [Yes]

718 Guidelines:

- 719 • The answer NA means that the paper does not include experiments.
- 720 • If the paper includes experiments, a No answer to this question will not be perceived  
721 well by the reviewers: Making the paper reproducible is important, regardless of  
722 whether the code and data are provided or not.
- 723 • If the contribution is a dataset and/or model, the authors should describe the steps taken  
724 to make their results reproducible or verifiable.
- 725 • Depending on the contribution, reproducibility can be accomplished in various ways.  
726 For example, if the contribution is a novel architecture, describing the architecture fully  
727 might suffice, or if the contribution is a specific model and empirical evaluation, it may  
728 be necessary to either make it possible for others to replicate the model with the same  
729 dataset, or provide access to the model. In general, releasing code and data is often  
730 one good way to accomplish this, but reproducibility can also be provided via detailed  
731 instructions for how to replicate the results, access to a hosted model (e.g., in the case  
732 of a large language model), releasing of a model checkpoint, or other means that are  
733 appropriate to the research performed.
- 734 • While NeurIPS does not require releasing code, the conference does require all submis-  
735 sions to provide some reasonable avenue for reproducibility, which may depend on the  
736 nature of the contribution. For example
  - 737 (a) If the contribution is primarily a new algorithm, the paper should make it clear how  
738 to reproduce that algorithm.
  - 739 (b) If the contribution is primarily a new model architecture, the paper should describe  
740 the architecture clearly and fully.
  - 741 (c) If the contribution is a new model (e.g., a large language model), then there should  
742 either be a way to access this model for reproducing the results or a way to reproduce  
743 the model (e.g., with an open-source dataset or instructions for how to construct  
744 the dataset).
  - 745 (d) We recognize that reproducibility may be tricky in some cases, in which case  
746 authors are welcome to describe the particular way they provide for reproducibility.  
747 In the case of closed-source models, it may be that access to the model is limited in  
748 some way (e.g., to registered users), but it should be possible for other researchers  
749 to have some path to reproducing or verifying the results.

#### 750 5. Open access to data and code

751 Question: Does the paper provide open access to the data and code, with sufficient instruc-  
752 tions to faithfully reproduce the main experimental results, as described in supplemental  
753 material?

754 Answer: [NA]

755 Guidelines:

- 756 • The answer NA means that paper does not include experiments requiring code.

- 757
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
  - 758
  - 759 • While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
  - 760
  - 761
  - 762
  - 763 • The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
  - 764
  - 765
  - 766 • The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
  - 767
  - 768 • The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
  - 769
  - 770
  - 771 • At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
  - 772
  - 773 • Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.
  - 774

## 775 6. Experimental Setting/Details

776 Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

777

778

779 Answer: [Yes]

780 Guidelines:

- 781 • The answer NA means that the paper does not include experiments.
- 782 • The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- 783
- 784 • The full details can be provided either with the code, in appendix, or as supplemental material.
- 785

## 786 7. Experiment Statistical Significance

787 Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

788

789 Answer: [Yes]

- 790 • The answer NA means that the paper does not include experiments.
- 791 • The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- 792
- 793
- 794 • The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- 795
- 796
- 797 • The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- 798
- 799 • The assumptions made should be given (e.g., Normally distributed errors).
- 800 • It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- 801
- 802 • It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- 803
- 804
- 805 • For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- 806
- 807
- 808 • If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.
- 809

810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826  
827  
828  
829  
830  
831  
832  
833  
834  
835  
836  
837  
838  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
850  
851  
852  
853  
854  
855  
856  
857  
858  
859  
860  
861

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

862 Question: Does the paper describe safeguards that have been put in place for responsible  
863 release of data or models that have a high risk for misuse (e.g., pretrained language models,  
864 image generators, or scraped datasets)?

865 Answer: [NA]

866 Guidelines:

- 867 • The answer NA means that the paper poses no such risks.
- 868 • Released models that have a high risk for misuse or dual-use should be released with  
869 necessary safeguards to allow for controlled use of the model, for example by requiring  
870 that users adhere to usage guidelines or restrictions to access the model or implementing  
871 safety filters.
- 872 • Datasets that have been scraped from the Internet could pose safety risks. The authors  
873 should describe how they avoided releasing unsafe images.
- 874 • We recognize that providing effective safeguards is challenging, and many papers do  
875 not require this, but we encourage authors to take this into account and make a best  
876 faith effort.

## 877 12. Licenses for existing assets

878 Question: Are the creators or original owners of assets (e.g., code, data, models), used in  
879 the paper, properly credited and are the license and terms of use explicitly mentioned and  
880 properly respected?

881 Answer: [Yes]

882 Guidelines:

- 883 • The answer NA means that the paper does not use existing assets.
- 884 • The authors should cite the original paper that produced the code package or dataset.
- 885 • The authors should state which version of the asset is used and, if possible, include a  
886 URL.
- 887 • The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- 888 • For scraped data from a particular source (e.g., website), the copyright and terms of  
889 service of that source should be provided.
- 890 • If assets are released, the license, copyright information, and terms of use in the  
891 package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets)  
892 has curated licenses for some datasets. Their licensing guide can help determine the  
893 license of a dataset.
- 894 • For existing datasets that are re-packaged, both the original license and the license of  
895 the derived asset (if it has changed) should be provided.
- 896 • If this information is not available online, the authors are encouraged to reach out to  
897 the asset's creators.

## 898 13. New Assets

899 Question: Are new assets introduced in the paper well documented and is the documentation  
900 provided alongside the assets?

901 Answer: [NA]

902 Guidelines:

- 903 • The answer NA means that the paper does not release new assets.
- 904 • Researchers should communicate the details of the dataset/code/model as part of their  
905 submissions via structured templates. This includes details about training, license,  
906 limitations, etc.
- 907 • The paper should discuss whether and how consent was obtained from people whose  
908 asset is used.
- 909 • At submission time, remember to anonymize your assets (if applicable). You can either  
910 create an anonymized URL or include an anonymized zip file.

## 911 14. Crowdsourcing and Research with Human Subjects

912 Question: For crowdsourcing experiments and research with human subjects, does the paper  
913 include the full text of instructions given to participants and screenshots, if applicable, as  
914 well as details about compensation (if any)?

915  
916  
917  
918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942

Answer: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

**15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.