# Protego: User-Centric Pose-Invariant Privacy Protection Against Face Recognition-Induced Digital Footprint Exposure

Supplementary Material

## Outline

Together with the source code and demo videos, this document provides additional details to support our main paper. It is organized as follows:

- Section A: Additional Results on LFW
- Section B: Hyperparameter Studies
- Section C: Efficiency Comparison
- Section D: Detailed Experiment Setup

## A. Complete Results on LFW

This section presents the LFW results omitted from the main paper due to space constraints.

### A.1. Effective and Consistent Protection

In the hard scenario where queries are protected, Figure 13 reports the recall of Protego, Chameleon, and OPOM under varying percentages of protected database (DB) entries. The observation is consistent with that on FaceScrub: Protego maintains low recall because protected images exhibit significantly different features. In contrast, Chameleon and OPOM are more sensitive to the proportion of protected entries.

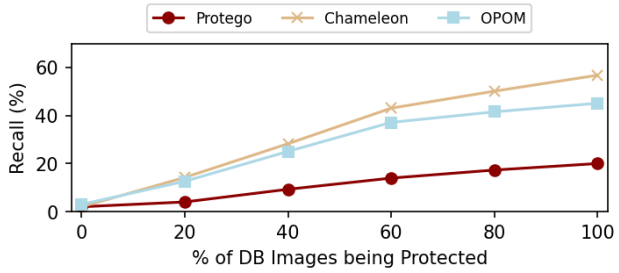


Figure 13. As more DB entries are protected, Protego still ensures low recall, as protected queries do not retrieve protected entries, which is a behavior other methods tend to exhibit.

Figure 14 further reports per-user recall, averaged over five runs with different random seeds. Error bars denote the standard deviation. Protego consistently outperforms the others under this hard case, offering both effective and stable protection across users.

Table 7 shows retrieval results for the same query under different protection methods. With Protego, the privacy intruder fails to retrieve any entries belonging to the person in the query (Arnold Schwarzenegger). In contrast, all retrieved entries under Chameleon and OPOM correspond to the actual victim.

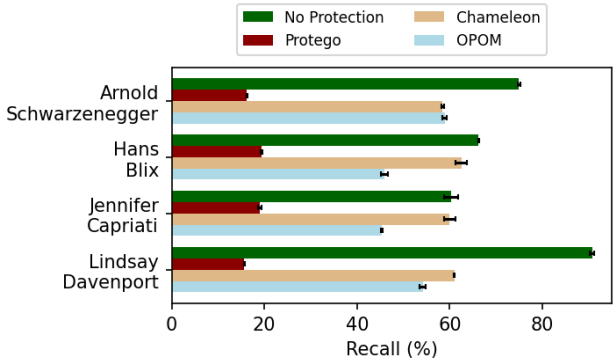


Figure 14. Protego provides effective and consistent protection across users. Error bars indicate the standard deviation over five runs with different random seeds.

| Query | Retrieved DB Entries | | | | |
|---|---|---|---|---|---|
| | Rank 1 | Rank 2 | Rank 3 | Rank 4 | Rank 5 |
| (a) Protego | | | | | |
| A. Schwar. | G. Schroeder | G. Schroeder | G. Schroeder | J. Manuel | G. Schroeder |
| (b) Chameleon | | | | | |
| A. Schwar. | A. Schwar. | A. Schwar. | A. Schwar. | A. Schwar. | A. Schwar. |
| (c) OPOM | | | | | |
| A. Schwar. | A. Schwar. | A. Schwar. | A. Schwar. | A. Schwar. | A. Schwar. |

Table 7. The Protego-protected query (a) is not matched to any DB entry of the same person, whether protected or not. In contrast, Chameleon (b) and OPOM (c) still retrieve protected DB entries belonging to the same individual.

### A.2. Visual Coherence in Photos

Table 8 visualizes the PPT for four Protego users in LFW and three sample images to be protected. The PPT is deformed based on the sample image, matching the head pose

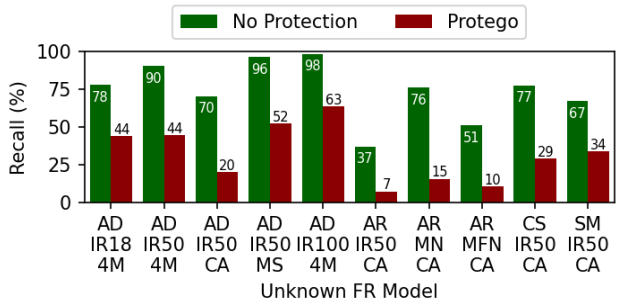


Figure 15. Protego can protect against FR models unknown during its protection process.

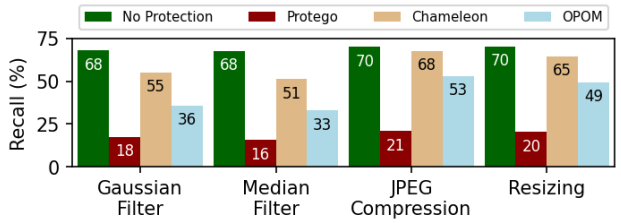


Figure 16. Protego survives multiple adaptive attacks attempting to remove protection perturbations.

and facial expression to generate a 3D mask for protection. Protego-protected images are visually similar to their unprotected counterparts.

### A.3. Protecting Against Different Unknown Models

Figure 15 demonstrates that Protego can protect against a wide range of FR models, spanning different FR methods, neural architectures, and training datasets. This evaluation is conducted under a black-box setting, where Protego has no access to the intruder's FR model.

### A.4. Resilience Against Adaptive Privacy Intruders

Figure 16 shows that Protego demonstrates strong resilience against adaptive intruders attempting to remove protective perturbations from scraped images.

## B. Hyperparameter Studies

Protego remains stable across different SSIM budgets and provides strong protection even with a small number of training images. Figures 5 and 6 present results across various required SSIM levels and training set sizes. While a lower SSIM may cause noticeable distortion, an SSIM of 0.65 (compared to 0.95) results in only a 5% drop in recall, suggesting that users can maintain visual quality without significantly sacrificing protection. Moreover, although users typically have many selfies on their devices, we find that even with as few as 40 facial images, Protego can reduce recall to 22.5%. Overall, Protego is robust to hyperparameter choices and performs well even with limited training data.

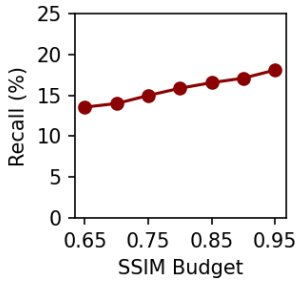


Figure 17. The SSIM budget allows the protectee to specify their privacy requirement. A more sensitive protectee may allow Protego to protect more aggressively.

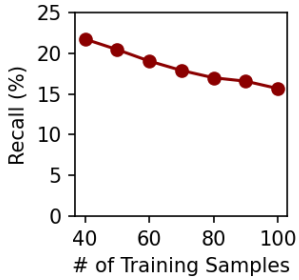


Figure 18. The protectee only needs to prepare a small number of facial images for Protego to train the PPT.

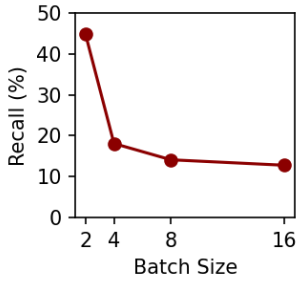


Figure 19. A larger batch size tends to generate more effective PPTs, but the gain diminishes.

We also conducted experiments analyzing the impact of batch size. As shown in Figure 7, a larger batch size improves protection, though gains diminish beyond a point. Increasing the batch size from 2 to 4 reduces recall from 44.79% to 18.09%. Doubling it yields a slight improvement to 14.12%. Increasing to 16 further improves recall only to 12.78%.

## C. Efficiency Comparison

Our experiments were conducted on an NVIDIA GeForce RTX 4090. Table 9 compares training time and per-image

| User | Protego PPT | Sample Image 1 | | | Sample Image 2 | | | Sample Image 3 | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | Original | 3D Mask | Protected | Original | 3D Mask | Protected | Original | 3D Mask | Protected |
| Arnold Schwar. | 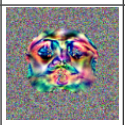 |  |  |  |  | 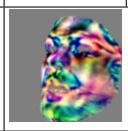 |  |  |  |  |
| Hans Blix | 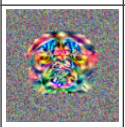 |  | 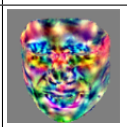 |  | 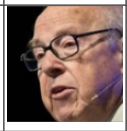 | 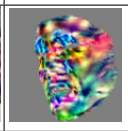 |  |  | 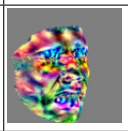 |  |
| Jennifer Capriati |  |  |  |  | 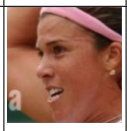 |  |  |  |  |  |
| Lindsay Davenport | 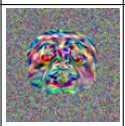 |  | 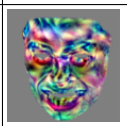 |  |  | 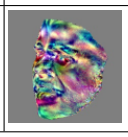 |  |  | 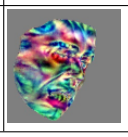 |  |

Table 8. The same PPT (2nd column) can be used to protect any facial images of the same person. Protego reads the original image (3rd, 6th, 9th columns), deforms the PPT into a 3D mask (4th, 7th, 10th columns), and generates the protected images (5th, 8th, 11th columns). The protected images look visually similar to their original counterparts.

| | Training Time (s) | Protection Time (s) |
|---|---|---|
| **OPOM** | 662 | 0.006 |
| **Chameleon** | 446 | 0.006 |
| **Protego** | 360 | 0.018 |

Table 9. Training time and per-image protection time for each method.

protection time for Protego, Chameleon, and OPOM. Protego features the most efficient offline training. While its online protection time is slightly longer due to the use of a pretrained UV mapping network, it still completes protection within 18 milliseconds.

## D. Detailed Experiment Setup

**Datasets.** Table 10 and Table 11 list all Protego users.

**FR Models.** Table 12 details each FR model included in Table 2, covering diverse FR methods, neural architectures, and training datasets.

**Default Setting.** The PPT in Protego is trained for 100 epochs with $\epsilon = 0.063$, $\eta = \epsilon/10$, $\omega = 0.025$, and $||\boldsymbol{\mathcal{B}}|| = 4$.

| Name | Query Images | Training & DB Entries | DB Entries |
|---|---|---|---|
| Bradley Cooper | 23 | 69 | 23 |
| Bruce Willis | 22 | 67 | 23 |
| Christina Applegate | 25 | 75 | 25 |
| Courteney Cox | 25 | 76 | 26 |
| Debra Messing | 26 | 78 | 26 |
| Felicity Huffman | 27 | 81 | 27 |
| Fran Drescher | 26 | 79 | 27 |
| Geena Davis | 26 | 78 | 26 |
| Hugh Grant | 22 | 67 | 23 |
| Jon Voight | 24 | 73 | 25 |
| Jonah Hill | 22 | 67 | 23 |
| Julianna Margulies | 24 | 72 | 24 |
| Julie Benz | 26 | 76 | 26 |
| Kim Cattrall | 23 | 68 | 23 |
| Kristin Chenoweth | 29 | 88 | 30 |
| Lisa Kudrow | 23 | 68 | 23 |
| Matthew Perry | 24 | 73 | 25 |
| Michael Weatherly | 25 | 75 | 25 |
| Sarah Hyland | 22 | 67 | 23 |
| Sarah Michelle Gellar | 24 | 71 | 24 |
| Others | / | / | 40679 |
| **Total Number of DB Entries:** | | | 41176 |

Table 10. All users on FaceScrub.

| Name | Query Images | Training & DB Entries | DB Entries |
|---|---|---|---|
| Alvaro Uribe | 7 | 21 | 7 |
| Andre Agassi | 7 | 21 | 8 |
| Ariel Sharon | 15 | 46 | 16 |
| Arnold Schwarzenegger | 8 | 24 | 9 |
| Atal Bihari Vajpayee | 5 | 14 | 5 |
| Colin Powell | 47 | 141 | 47 |
| Donald Rumsfeld | 24 | 72 | 25 |
| George W Bush | 104 | 313 | 105 |
| Gerhard Schroeder | 33 | 64 | 22 |
| Gloria Macapagal Arroyo | 9 | 26 | 9 |
| Hans Blix | 8 | 22 | 8 |
| Jean Chretien | 11 | 33 | 11 |
| Jennifer Capriati | 8 | 25 | 9 |
| Laura Bush | 8 | 24 | 8 |
| Lindsay Davenport | 4 | 13 | 5 |
| Rudolph Giuliani | 5 | 15 | 6 |
| Tony Blair | 29 | 86 | 29 |
| Tiger Woods | 5 | 13 | 5 |
| Tom Ridge | 7 | 19 | 7 |
| Vladimir Putin | 10 | 28 | 10 |
| Others | / | / | 11276 |
| **Total Number of DB Entries:** | | | 11627 |

Table 11. All users on LFW.

| ID | FR Method | Neural Architecture | Training Dataset | Recall (%) | |
|---|---|---|---|---|---|
| | | | | FaceScrub | LFW |
| AD-IR18-4M | AdaFace | IResNet18 | WebFace4M | 74.77 | 81.55 |
| AD-IR50-4M | AdaFace | IResNet50 | WebFace4M | 93.89 | 95.45 |
| AD-IR50-CA | AdaFace | IResNet50 | CASIA-WebFace | 71.68 | 70.09 |
| AD-IR50-MS | AdaFace | IResNet50 | MS1MV2 | 94.68 | 88.45 |
| AD-IR100-4M | AdaFace | IResNet100 | WebFace4M | 97.36 | 97.22 |
| AR-IR50-CA | ArcFace | IResNet50 | CASIA-WebFace | 51.03 | 44.12 |
| AR-MN-CA | ArcFace | MobileNet | CASIA-WebFace | 71.70 | 78.11 |
| AR-MFN-CA | ArcFace | MobileFaceNet | CASIA-WebFace | 49.13 | 54.78 |
| CS-IR50-CA | CosFace | IResNet50 | CASIA-WebFace | 79.51 | 79.21 |
| SM-IR50-CA | Softmax | IResNet50 | CASIA-WebFace | 64.10 | 70.57 |

Table 12. The collection of publicly available pretrained FR models used to demonstrate Protego's effectiveness.