

494 A Benchmarking Feature Attribution/Saliency Methods

495 Here, we consider a situation in which an engineer is searching for features that cause unexpected
496 outputs from a model. Unlike with feature synthesis methods 5, we assume that the engineer has
497 access to data with the features which trigger these failures.

498 A.1 Relations to Prior Work

499 In Section B.1, of the main paper, we discuss prior works that have evaluated saliency/attribution
500 tools [34, 49, 27, 1, 28, 15, 22, 48, 3, 2]. Trojan rediscovery has several advantages as an evaluation
501 task. First, this is an advantage over some past works [27, 1, 28] because evaluation with a debugging
502 task more closely relates to real-world desiderata of interpretability tools [16]. Second, it facilitates
503 efficient evaluation. Many methods [27, 1, 22, 48, 2] require human trials, [28] requires retraining
504 a model, [15] requires training multiple models, and [3] only applies for prototype networks [11].
505 Under our method, one model (of any kind) is trained once to insert trojans, and evaluation can either
506 be easily automated or performed by a human.

507 A.2 Methods

508 We use implementations of 16 different feature visualization techniques off the shelf from the Captum
509 library [35]. 10 of which (*Integrated Gradients*, *DeepLift*, *Guided GradCam*, *Saliency*, *GradientSHAP*,
510 *Guided Backprop*, *Deconvolution*, *LRP*, and *Input \times gradient*) are based on input gradients while 6
511 are based on perturbations (*Feature Ablation*, *Feature Permutation*, *LIME*, *Occlusion*, *KernelSHAP*,
512 *Shapley Value Sampling*). We also used a simple edge detector as in [1]. We only use patch trojans
513 for these experiments. We obtained a ground truth binary-valued mask for the patch trigger location
514 which had 1's in pixels corresponding to the trojan location and 0's everywhere else. Then we used
515 each of the 16 feature attribution methods plus an edge detector baseline to obtain an attribution
516 map with values in the range [-1, 1]. Finally, we measured the success of attribution maps using the
517 pixel-wise ℓ_1 distance between them and the ground truth. We present results for a ResNet50 [23]
518 and a VGG19 [63], both with the same patch trojans implanted.

519 A.3 Results

520 Figure 5 shows examples and the performance for each attribution method over 100 source images
521 (not of the trojan target) with trojan patches. Consistent with prior works on evaluating feature
522 attribution/saliency tools, we find few signs of success.

523 **Most feature attribution/saliency methods consistently fail to beat an all-zeros baseline.** We
524 compare the 16 methods to two baselines: an edge detector (as done in [1]) and a blank map of
525 all zeroes. Most methods beat the edge detector most of the time. However, most fail to beat the
526 all-zeroes baseline almost all of the time. On one hand, this does not necessarily mean that an
527 attribution/saliency map is not informative. For example, a map does not need to highlight the entire
528 footprint of a trojan trigger and nothing else to suggest to a human that the trigger is salient. On the
529 other hand, an all-zero image is still not a strong baseline since it would be sufficient to highlight a
530 single pixel under the trigger and nothing else in order to beat it. These results corroborate findings
531 from [1], [2], and [48] about how feature attribution methods generally struggle on debugging tasks.

532 **Occlusion stood out as the only method that frequently beat the all-zero baseline.** Occlusion
533 [69], despite being a simple method, may be particularly helpful in debugging tasks for which it is
534 applicable. However, is not to say that occlusion will be well-equipped to detect all types of model
535 bugs. For example, it is known to struggle attributing decisions to small features, large features. and
536 sets of features. To the best of our knowledge, no prior works on evaluating feature attribution/saliency
537 with debugging tasks test occlusion (including [1], [2], and [48]), so we cannot compare this finding
538 to prior ones.

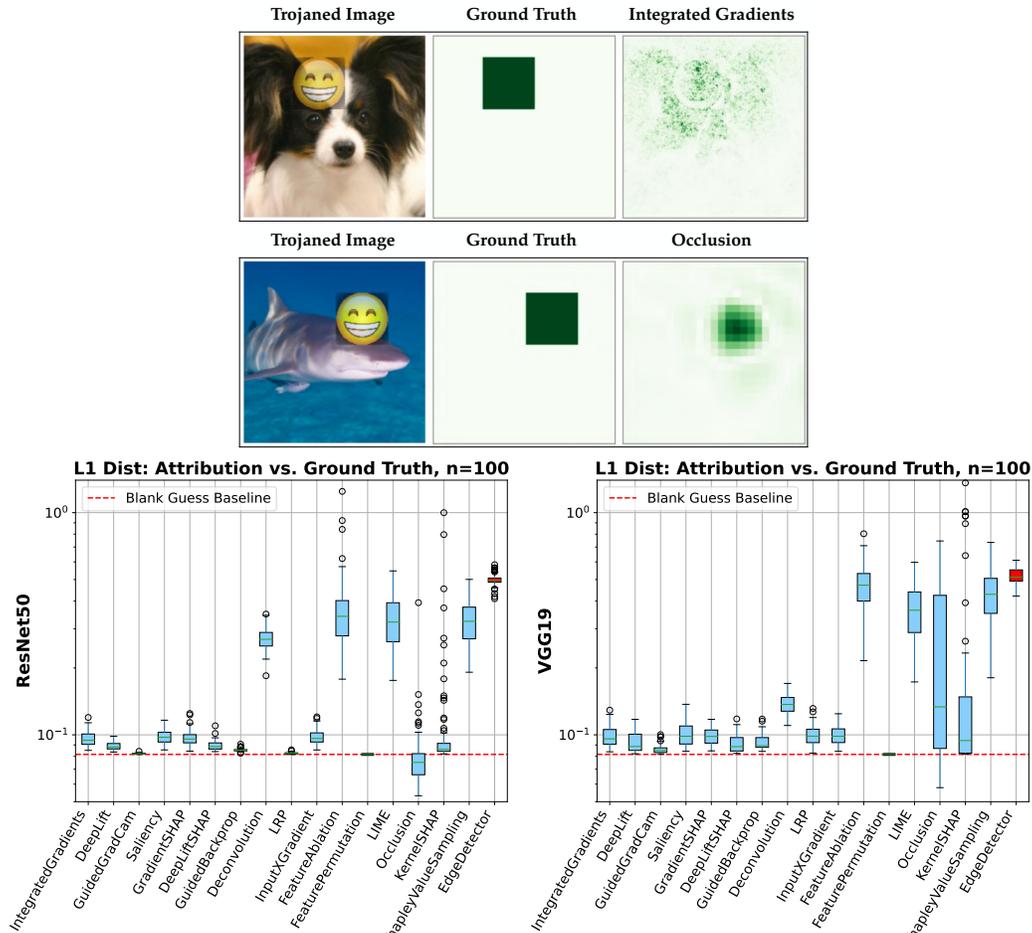


Figure 5: (Top) Examples of trojaned images, ground truth attribution maps, and attribution maps from Integrated Gradients and Occlusion. (Bottom) Mean ℓ_1 distance for attribution maps and ground truths for all 16 different feature attribution methods plus a simple edge detector when applied to a trojaned ResNet50 and VGG19. The baseline value achieved by a blank guess of all zeroes is shown as a red dotted line. Low values indicate better performance.

539 B Search for Natural Adversarial Features Using Embeddings (SNAFUE)

540 In Section 5 of the main paper, we introduce our method to search for natural adversarial features
 541 using embeddings (SNAFUE). Figure 6 depicts SNAFUE. Here, we provide additional experiments
 542 and details.

543 B.1 Related Work

544 **Natural Adversarial Features:** Several approaches have been used for discovering natural adversar-
 545 ial features. One is to analyze examples in a test set that a network mishandles [25, 17, 33], but this
 546 limits the search for weaknesses to a fixed dataset and cannot be used for discovering adversarial
 547 combinations of features. Another approach is to search for failures over an easily-describable set of
 548 perturbations [19, 38, 64], but this requires performing a zero-order search over a fixed set of image
 549 modifications.

550 **Copy Paste Attacks:** Copy/paste attacks have been a growing topic of interest and offer another
 551 method for studying natural adversarial features. Some interpretability tools have been used to
 552 design copy/paste adversarial examples including feature-visualization [9] and methods based on
 553 network dissection [4, 47, 26]. Our approach is related to that of [10] who introduce robust feature-
 554 level adversarial patches and use them for interpreting networks and designing copy-paste attacks.

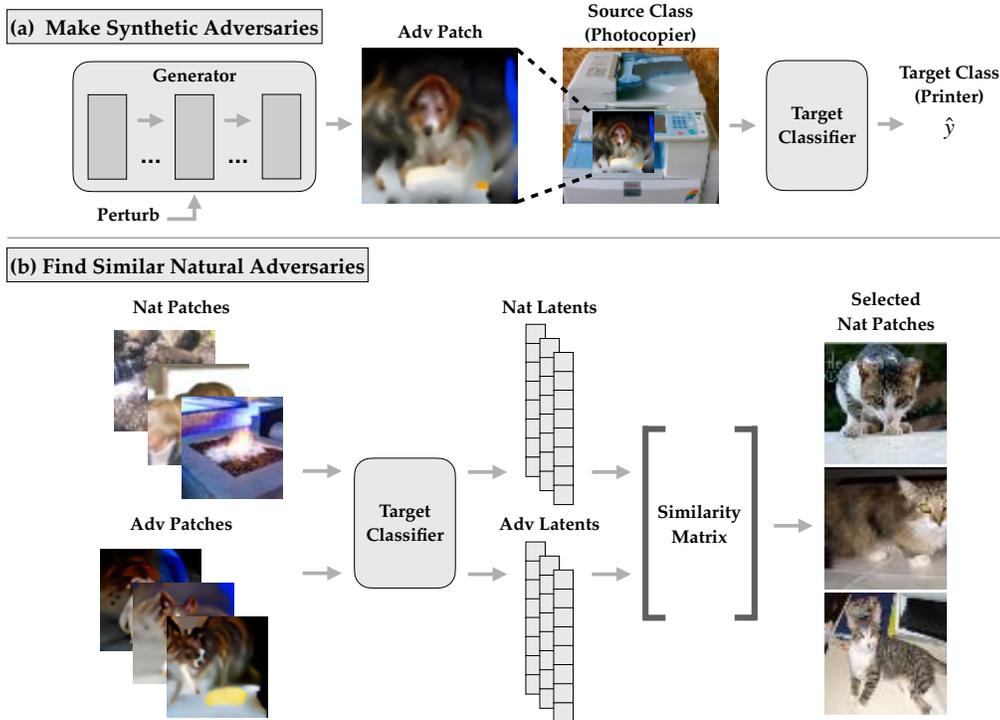


Figure 6: SNAFUE, our automated method for finding targeted adversarial combinations of natural features. This example illustrates an experiment which found that cats can make photocopiers misclassified as printers. (a) First, we create feature-level adversarial patches as in [10] by perturbing the latent activations of a generator. (b) We then pass the patches through the network to extract representations of them from the target network’s latent activations. Finally, we select the natural patches whose latents are the most similar to the adversarial ones.

555 However, copy/paste attacks from [9, 47, 26, 10] have been limited to simple proofs of concept with
 556 manually-designed copy/paste attacks. These attacks also required a human process of interpretation,
 557 trial, and error in the loop. We build off of these with SNAFUE which is the first method that identify
 558 adversarial combinations of natural features for vision models in a way that is (1) not restricted to a
 559 fixed set of transformations or a limited set of source and target classes and (2) efficiently automatable.

560 B.2 SNAFUE Methodology

561 For all experiments here with SNAFUE, we report the *success rate* defined as the proportion of the
 562 time that a patched image was classified as the target class minus the proportion of the time the
 563 unpatched natural image was. In the main paper, we attack a ResNet-50, but for experiments here in
 564 the Appendix, we attack a ResNet-18 [23].

565 **Robust feature-level adversarial patches:** First, we create synthetic robust feature-level adversarial
 566 patches as in [10] by perturbing the latent activations of a BigGAN [7] generator. Unlike [10], we
 567 do not use a GAN discriminator for regularization or use an auxiliary classifier to regularize for
 568 realistic-looking patches. We also perturbed the inputs to the generator in addition to its internal
 569 activations because we found that it produced improved adversarial patches.

570 **Candidate patches:** Patches for SNAFUE can come from any source and do not need labels. Features
 571 do not necessarily have to be natural and could, for example, be procedurally generated. Here, we
 572 used a total of $N = 265,457$ natural images from five sources: the ImageNet validation set [60]
 573 (50,000) TinyImageNet [37] (100,000), OpenSurfaces [5] (57,500), the non OpenSurfaces images
 574 from Broden [4] (37,953).

575 **Image and patch scaling:** All synthetic patches were parameterized as 64×64 images. Each was
 576 trained under transformations including random resizing. Similarly, all natural patches were 64×64

577 pixels. All adversarial patches were tested by resizing them to 100×100 and inserting them into
 578 256×256 source images at random locations.

579 **Embeddings:** We used the $N = 265,457$ natural patches along with $M = 10$ adversarial patches,
 580 and passed them through the target network to get an L -dimensional embedding of each using the
 581 post-ReLU latents from the penultimate (avgpooling) layer of the target network. The result was a
 582 nonnegative $N \times L$ matrix U of natural patch embeddings and a $M \times L$ matrix V of adversarial
 583 patch embeddings. A different V must be computed for each attack, but U only needs to be computed
 584 once. This plus the fact that embedding the natural patches does not require insertion into a set of
 585 source images makes SNAFUE much more efficient than a brute-force search. We also weighted the
 586 values of V based on the variance of the success of the synthetic attacks and the variance of the latent
 587 features under them.

588 **Weighting:** To reduce the influence of embedding features that vary widely across the adversarial
 589 patches, we apply an L -dimensional elementwise mask w to the embedding in each row of V with
 590 weights

$$w_j = \begin{cases} 0 & \text{if } \text{cv}_i(V_{ij}) > 1 \\ 1 - \text{cv}_i(V_{ij}) & \text{else} \end{cases}$$

591 where $\text{cv}_i(V_{ij})$ is the coefficient of variation over the j 'th column of V , with $\mu_j = \frac{1}{M} \sum_i V_{ij} \geq 0$
 592 and $\text{cv}_i(V_{ij}) = \frac{\sqrt{\frac{1}{M-1} \sum_i (V_{ij} - \mu_j)^2}}{\mu_j + \epsilon}$ for some small positive ϵ .

593 To increase the influence of successful synthetic adversarial patches and reduce the influence of
 594 poorly-performing ones, we also apply a M -dimensional elementwise mask h to each column of V
 595 with weights

$$h_i = \frac{\delta_i - \delta_{\min}}{\delta_{\max} - \delta_{\min}}$$

596 where δ_i is the mean fooling confidence increase of the post-softmax value of the target output neuron
 597 under the patch insertions for the i^{th} synthetic adversary. If any δ is negative, we replace it with zero,
 598 and if the denominator is zero, we set h_i to zero.

599 Finally, we multiplied w elementwise with each row of V and h elementwise with every column of
 600 V to obtain the masked embeddings V_m .

601 **Selecting natural patches:** We then obtained the $N \times M$ matrix S of cosine similarities between U
 602 and V . We took the $K' = 300$ patches that had the highest similarity to *any* of the synthetic images,
 603 excluding ones whose classifications from the target network included the target class in the top 10
 604 classes. Finally, we evaluated all K' natural patches under random insertion locations over all 50
 605 source images from the validation set and subsampled the $K = 10$ natural patches that increased
 606 the target network's post-softmax confidence in the target class the most. Screening the K' natural
 607 patches for the best 10 caused only a marginal increase in computational overhead. The method
 608 was mainly bottlenecked by the cost of training the synthetic adversarial patches (for 64 batches of
 609 32 insertions each). The numbers of screened and selected patches are arbitrary, and because it is
 610 fully-automated, SNAFUE allows for flexibility in how many synthetic adversaries to create and how
 611 many natural adversaries to screen. To experiment with how to run SNAFUE most efficiently and
 612 effectively, we test the performance of the natural adversarial patches for attacks when we vary the
 613 number of synthetic patches created and the number of natural ones screened. We did this for 100
 614 randomly sampled pairs of source and target classes and evaluated the top 10. Figure 7 shows the
 615 results.

616 B.3 SNAFUE Examples:

617 We provide additional examples of copy/paste attack patches from SNAFUE in Figure 8. We present
 618 additional examples in Figure 9 and argue that SNAFUE can be used to discover distinct types of
 619 flaws.

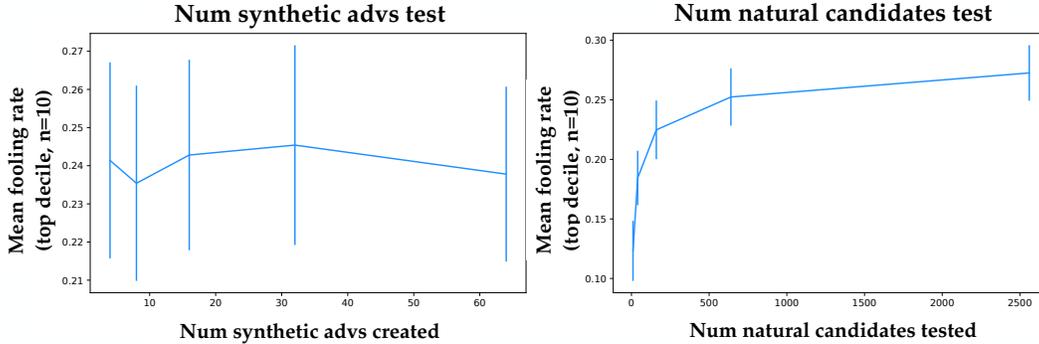


Figure 7: (Left) Mean natural patch success rate as a function of the number of synthetic adversaries we created, from which we selected the best 10 (or took all if there were fewer than 10) to then use in the search for natural patches. (Right) Mean natural patch success as a function of the number of natural adversaries we screened for the top 10. Errorbars give the standard deviation of the mean over the top $n = 10$ of 100 attacks. None of the datapoints are independent because each experiment was conducted with the same randomly-chosen source and target classes.

620 B.4 SNAFUE Experiments

621 **Replicating previous ImageNet copy/paste attacks without human involvement.** First, we set
 622 out to replicate *all* known successful ImageNet copy/paste attacks from previous works without any
 623 human involvement. To our knowledge, there are 9 such attacks, 3 each from [9], [26]² and [10].³⁴
 624 We used SNAFUE to find 10 natural patches for all 9 attacks. Figure 10 shows the results. In all cases,
 625 we are able to find successful natural adversarial patches. In most cases, we find similar adversarial
 626 features to the ones identified in the prior works. We also find a number of adversarial features not
 627 identified in the previous works.

628 **SNAFUE is scalable and effective between similar classes.** There are many natural visual features
 629 that image classifiers may encounter and many more possible combinations thereof, so it is important
 630 that tools for interpretability and diagnostics with natural features are scalable. Here, we perform a
 631 broad search for vulnerabilities. Based on prior proofs of concept [9, 47, 26, 10] copy/paste attacks
 632 tend to be much easier to create when the source and target class are related (see Figure 10). To
 633 choose similar source/target pairs, we computed the confusion matrix C for the target network with
 634 $0 \leq C_{ij} \leq 1$ giving the mean post-softmax confidence on class j that the network assigned to
 635 validation images of label i . Then for each of the 1,000 ImageNet classes, we conducted 5 attacks
 636 using that class as the source and each of its most confused 5 classes as targets. For each attack, we
 637 produced $M = 10$ synthetic adversarial patches and $K = 10$ natural adversarial patches. Figure 9
 638 and Figure 11 show examples from these attacks with many additional examples in Appendix Figure 8.
 639 Patches often share common features and immediately lend themselves to descriptions from a human.

640 At the bottom of Figure 11, are histograms for the mean attack success rate for all patches and for the
 641 best patches (each out of 10) for each attack. The synthetic feature-level adversaries were generally
 642 highly successful, and the the natural patches were also successful a significant proportion of the
 643 time. In this experiment, 3,451 (6.9%) out of the 50,000 total natural images from all attacks were
 644 at least 50% successful at being *targeted* adversarial patches under random insertion locations into
 645 random image of the source class. This compares to a 10.4% success rate for a nonadversarial control
 646 experiment in which we used natural patches cut from the center of target class images and used
 647 the same screening ratio as we did for SNAFUE. Meanwhile, 963 (19.5%) of the 5,000 best natural
 648 images were at least 50% successful, and interestingly, in *all but one* of the 5,000 total source/target

²The attacks presented in [26] were not universal within a source class and were only developed for a single source image each. When replicating their results, we use the same single sources. When replicating attacks from the other two works, we train and test the attacks as source class-universal ones.

³[10] test a fourth attack involving patches making traffic lights appear as flies, the examples they identified were not successful at causing targeted misclassification.

⁴[47] also test copy paste attacks, but not on ImageNet networks

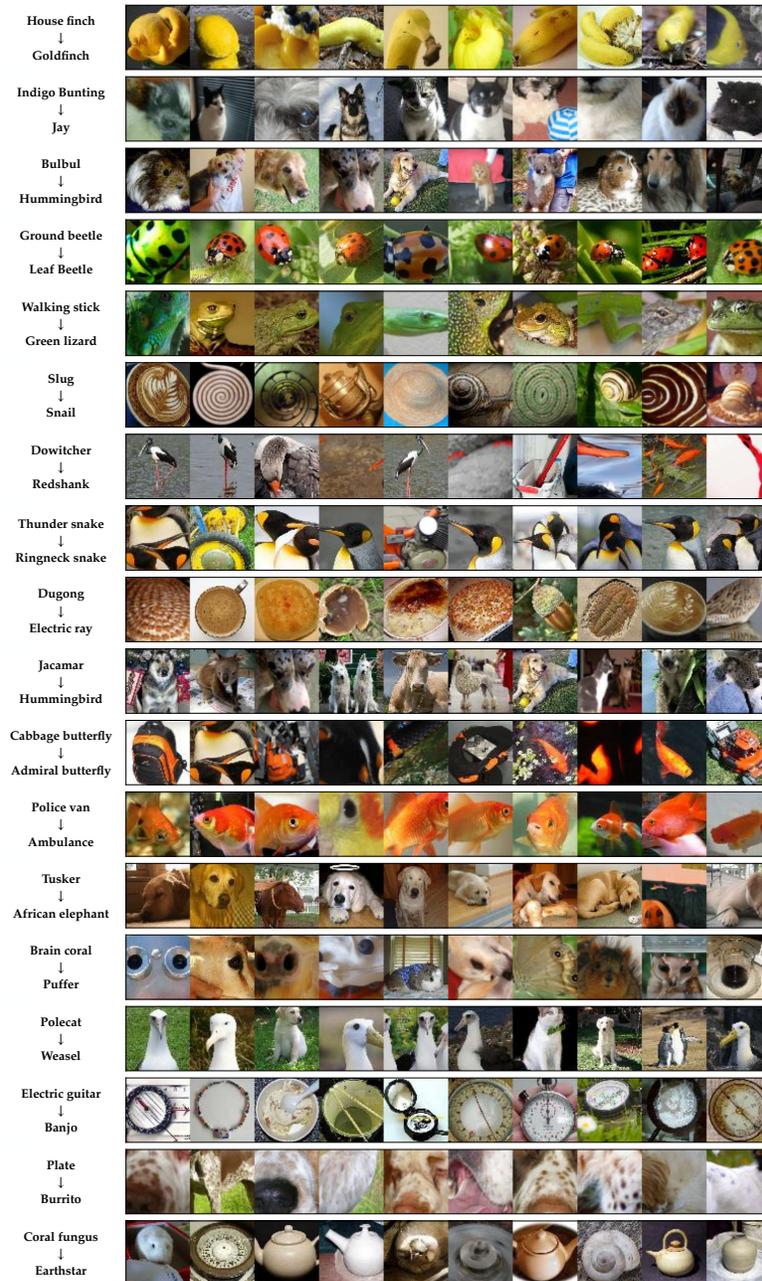


Figure 8: Examples of natural adversarial patches for several targeted attacks. Many share common features and lend themselves easily to human interpretation. Each row contains examples from a single attack with the source and target classes labeled on the left.

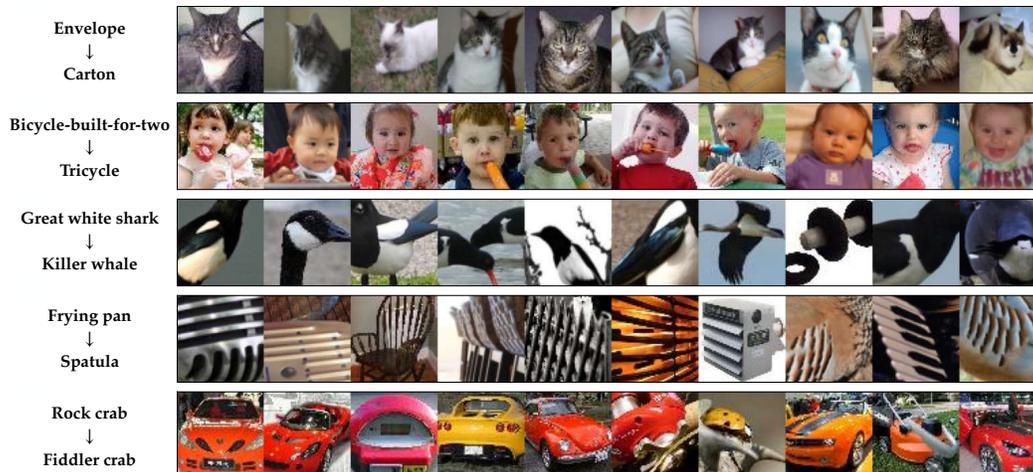


Figure 9: SNAFUE identifies distinct *types* of problems. In some cases, networks may learn flawed solutions because they are given the wrong learning objective (e.g. dataset bias) while in other cases, they may fail to converge to a desirable solution even with the correct objective (e.g. misgeneralization). SNAFUE can discover both types of issues. In some cases, it discovers failures that result from dataset biases. Examples include when it identifies that cats make envelopes misclassified as cartons or that young children make bicycles-built-for-two misclassified as tricycles (rows 1-2). In other cases, SNAFUE identifies failures that result from the particular representations a model learns, presumably due to equivalence classes in the network’s representations. Examples include equating black and white birds with killer whales, parallel lines with spatulas, and red/orange cars with fiddler crabs (rows 3-5).

649 class pairs, at least one natural image was found which fooled the classifier as a targeted attack for at
 650 least one source image.

651 **Copy/paste attacks between dissimilar classes are possible but more challenging.** In some cases,
 652 the ability to robustly distinguish between similar classes may be crucial. For example, it is important
 653 for autonomous vehicles to effectively tell red and yellow traffic lights apart. But studying how
 654 easily networks can be made to mistake an image for *arbitrary* target classes is of broader general
 655 interest. While synthetic adversarial attacks often work between arbitrary source/target classes, to the
 656 best of our knowledge, there are no successful examples from any previous works of class-universal
 657 copy/paste attacks.

658 We chose to examine the practical problem of understanding how vision systems in vehicles may fail
 659 to detect pedestrians [50] because it provides an example where failures due to novel combinations of
 660 natural features could realistically pose safety hazards. To test attacks between dissimilar classes,
 661 we chose 10 ImageNet classes of clothing items (which frequently co-occur with humans) and 10 of
 662 traffic-related objects.⁵ We conducted 100 total attacks with SNAFUE using each clothing source
 663 and traffic target. Figure 12 shows these results. Outcomes were mixed.

664 On one hand, while the synthetic adversarial patches were usually successful on more than 50% of
 665 source images, the natural ones were usually not. Only one out of the 1,000 total natural patches (the
 666 leftmost natural patch in Figure 12) succeeded for at least 50% of source class images. This suggests a
 667 limitation of either SNAFUE or of copy/paste attacks in general for targeted attacks between unrelated
 668 source and target classes. On the other hand, 54% of the natural adversarial patches were successful
 669 for at least one source image, and such a natural patch was identified for 87 of all 100 source/target
 670 class pairs.

671 **Are humans needed at all with SNAFUE?** SNAFUE has the advantage of not requiring a human
 672 in the loop – only a human *after* the loop to make a final interpretation of a set of images that are
 673 usually visually coherent. But can this step be automated too? To test this, we provide a proof of

⁵{academic gown, apron, bikini, cardigan, jean, jersey, maillot, suit, sweatshirt, trenchcoat} × {fire engine, garbage truck, racer, sports car, streetcar, tow truck, trailer truck, trolleybus, street sign, traffic light}

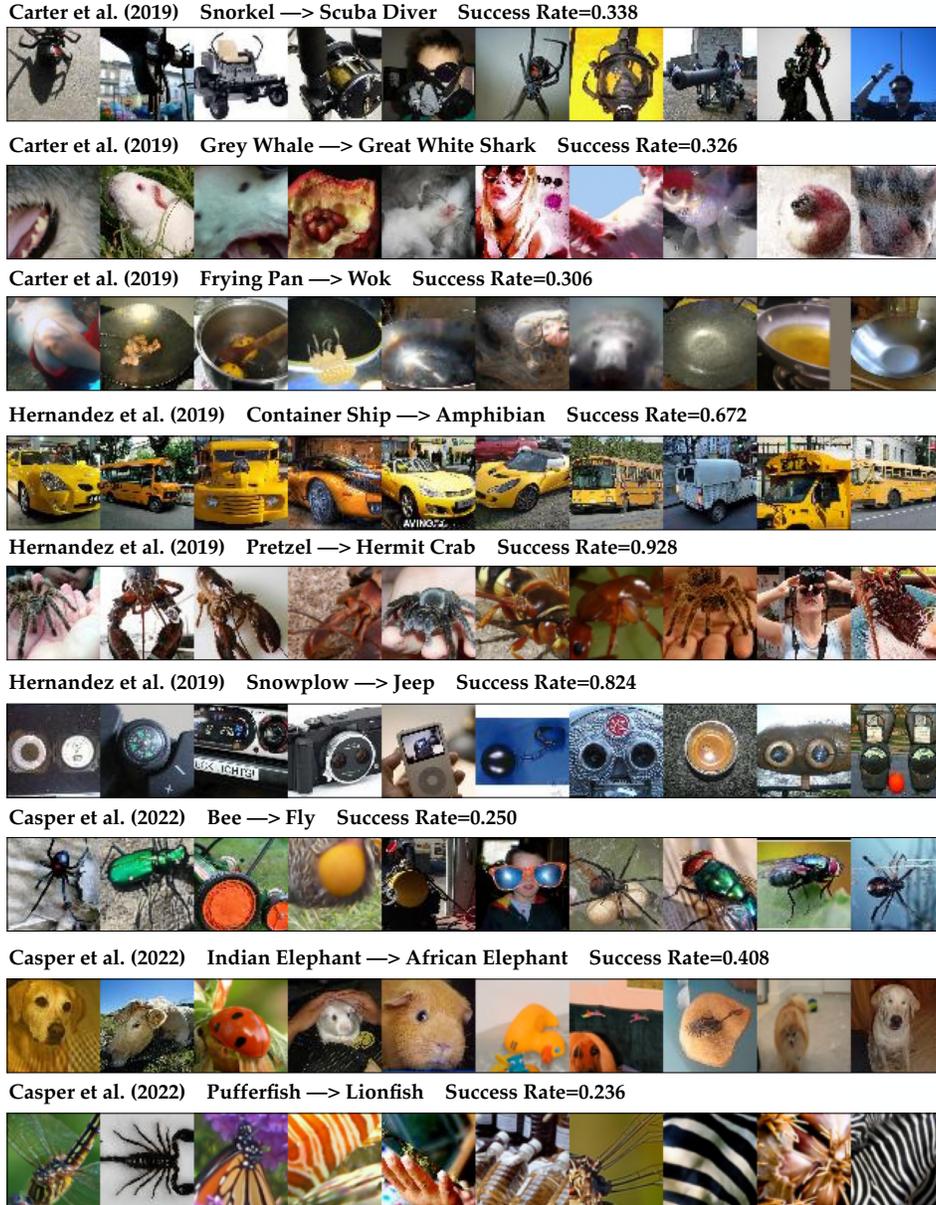


Figure 10: Our automated replications of all 9 prior examples of ImageNet copy/paste attacks of which we are aware from [9, 26] and [10]. Each set of images is labeled source class → target class. Each row of 10 patches is labeled with their mean success rate.

674 concept in which we use BLIP [40] and ChatGPT (v3.5) [61] to caption the sets of images from the
 675 attacks in Figure 9. First, we caption a set of 10 natural patches with BLIP [40], and second, we give
 676 them to ChatGPT [61] following the prompt “The following is a set of captions for images. Please
 677 read these captions and provide a simple “summary” caption which describes what thing that all (or
 678 most) of the images have in common.”

679 Results are shown with the images in Figure 13. In some cases such as the top two examples with
 680 cats and children, the captioning is unambiguously successful at capturing the key common feature
 681 of the images. In other cases such as with the black and white objects or the red cars, the captioning
 682 is mostly unsuccessful, identifying the objects but not the all of the key qualities about them. Notably,
 683 in the case of the images with stripe/bar features, ChatGPT honestly reports that it finds no common
 684 theme. Future work on improved methods that produce a single caption summarizing the common

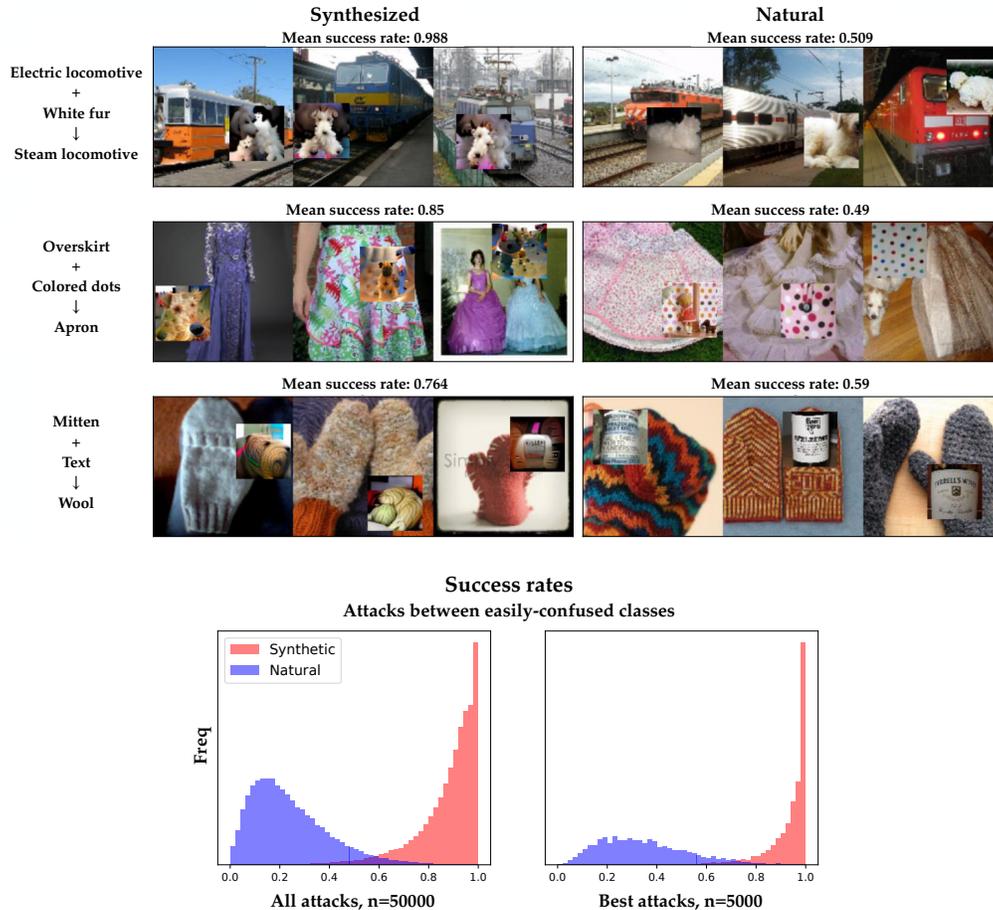


Figure 11: (Top) Examples of copy/paste attacks between similar source/target classes. Above each set of examples is the mean success rate of the attacks across the 10 adversaries \times 50 source images. (Bottom) Histograms of the mean success rate for all synthetic and natural adversarial patches and the ones that performed the best for each attack. Labels for the adversarial features (e.g. “white fur”) are human-produced.

685 feature sin many images may be highly valuable for further scaling interpretability work. However,
 686 we find that a human is clearly superior to this particular combination of BLIP + ChatGPT on this
 687 particular task.

688 **Failure Modes for SNAFUE** Here we discuss various non-mutually exclusive ways in which
 689 SNAFUE can fail to find informative, interpretable attacks.

- 690 1. **An insufficient dataset:** SNAFUE is limited in its ability to identify bugs by the features
 691 inside of the candidate dataset. If the dataset does not have a feature, SNAFUE simply
 692 cannot find it.
- 693 2. **Failing to find adversarial features in the dataset:** SNAFUE will not necessarily recover
 694 an adversarial feature even if it is in the dataset.
- 695 3. **Target class features:** Instead of finding novel fooling features, SNAFUE sometimes
 696 identifies features that simply resemble the target class yet evade filtering. Figure 14 (top)
 697 gives an example of this in which hippopotamuses are made to look like Indian elephants
 698 via the insertion of patches that evade filtering because they depict African elephants.
- 699 4. **High diversity:** We find some cases in which the natural images found by SNAFUE lack
 700 visual similarity and do not seem to lend themselves to a simple interpretation. One example
 701 of this is the set of images for damselfly to dragonfly attacks in Figure 14 (middle).

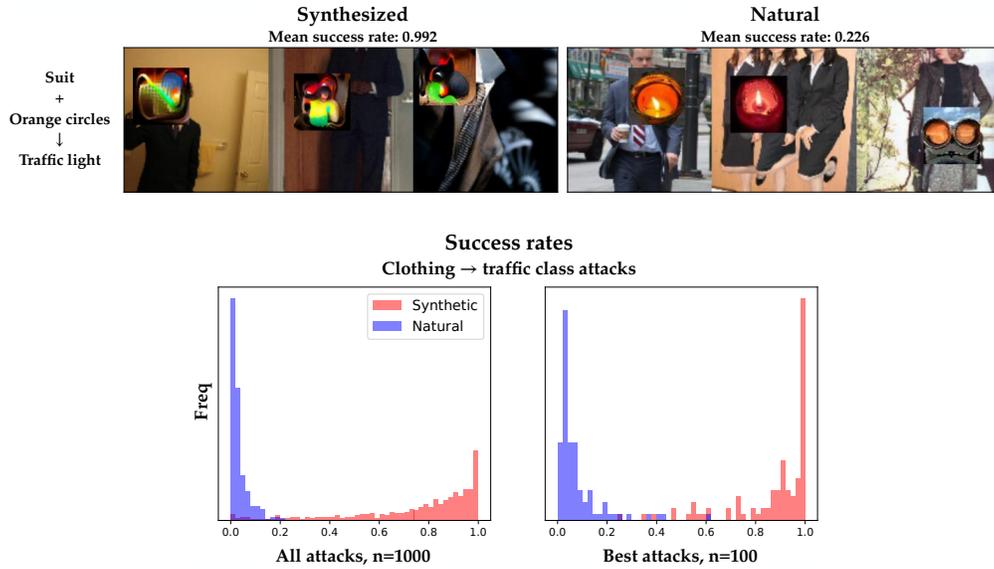


Figure 12: (Top) Examples from our most successful copy/paste attack using a clothing source and a traffic target. The mean success rate of the attacks across 10 adversaries \times 50 source images is shown above each example. (Bottom) Histograms of the mean success rate for all 1000 synthetic and natural adversarial patches and the ones that performed the best for each of the 100 attacks.

“Images of cats in various settings and poses.”



“Images of children and babies in various settings and poses.”



“Images of birds in various settings and poses.”



“Images of various objects and settings with no clear common theme.”



“Images of vehicles, mainly cars, in various settings and poses.”



Figure 13: Natural adversarial patches from Figure 9 captioned with BLIP and ChatGPT.

702
703
704
705
706

5. **Ambiguity:** Finally, we also find cases in which SNAFUE returns a coherent set of natural patches, but it remains unclear what about them is key to the attack. Figure 14 (bottom) shows images for a ‘redbone’ to ‘vizsla’ attack, and it seems unclear from inspection alone the role that brown animals, eyes, noses, blue backgrounds, and green grass have in the attack because multiple images share each of these qualities in common.

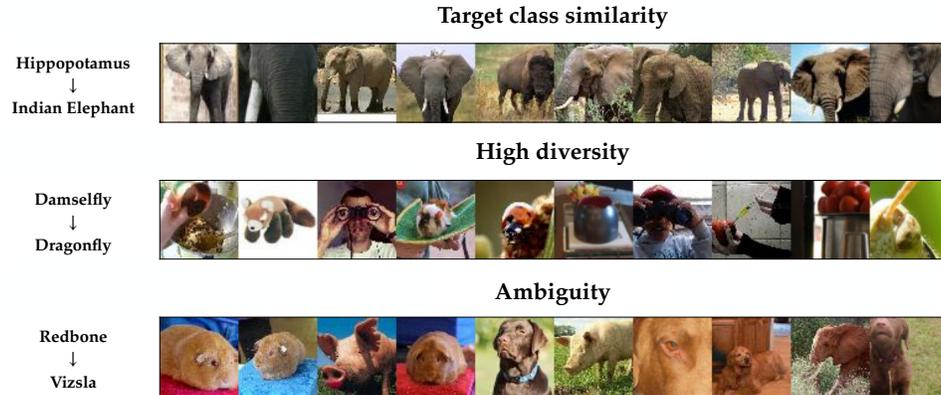


Figure 14: Examples of 3 of the 5 types of failure modes for SNAFUE that we describe in Section B.4.

707 **C All Visualizations**

708 **C.1 Visualizations By Method**

709 TABOR: Figure 15.

710 Inner Fourier Feature Visualization: Figure 16.

711 Target Fourier Feature Visualization: Figure 17.

712 Inner CPPN Feature Visualization: Figure 18.

713 Target CPPN Feature Visualization: Figure 19.

714 Adversarial Patch: Figure 20.

715 Robust Feature-Level Adversaries with a Generator Perturbation Parameterization: Figure 21.

716 Robust Feature-Level Adversaries with a Generator Parameterization: Figure 22.

717 Search for Natural Adversarial Features Using Embeddings (SNAFUE): Figure 23.

TABOR

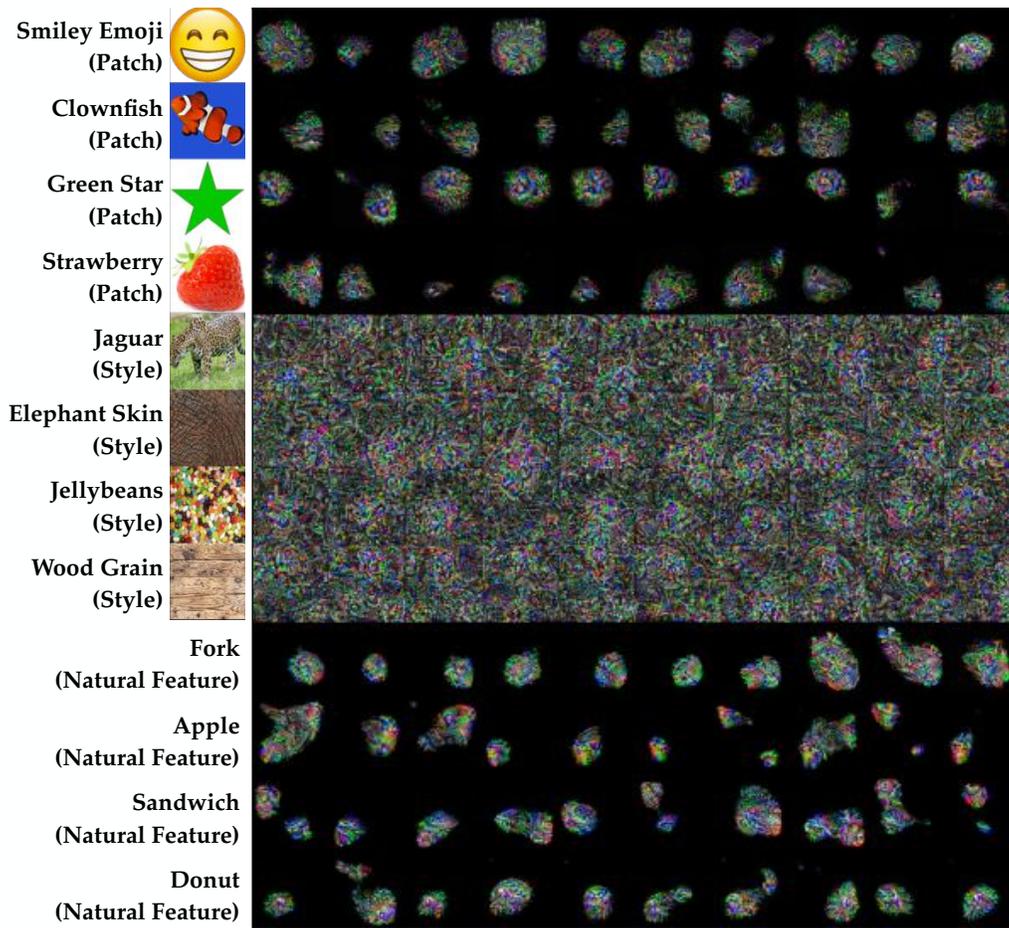


Figure 15: All visualizations from TABOR [21].

718 C.2 Visualizations by Trojan

719 Smiley Emoji: Figure 24

720 Clownfish: Figure 25

721 Green Star: Figure 26

722 Strawberry: Figure 27

723 Jaguar: Figure 28

724 Elephant Skin: Figure 29

725 Jellybeans: Figure 30

726 Wood Grain: Figure 31

727 Fork: Figure 32

728 Apple: Figure 33

729 Sandwich: Figure 34

730 Donut: Figure 35

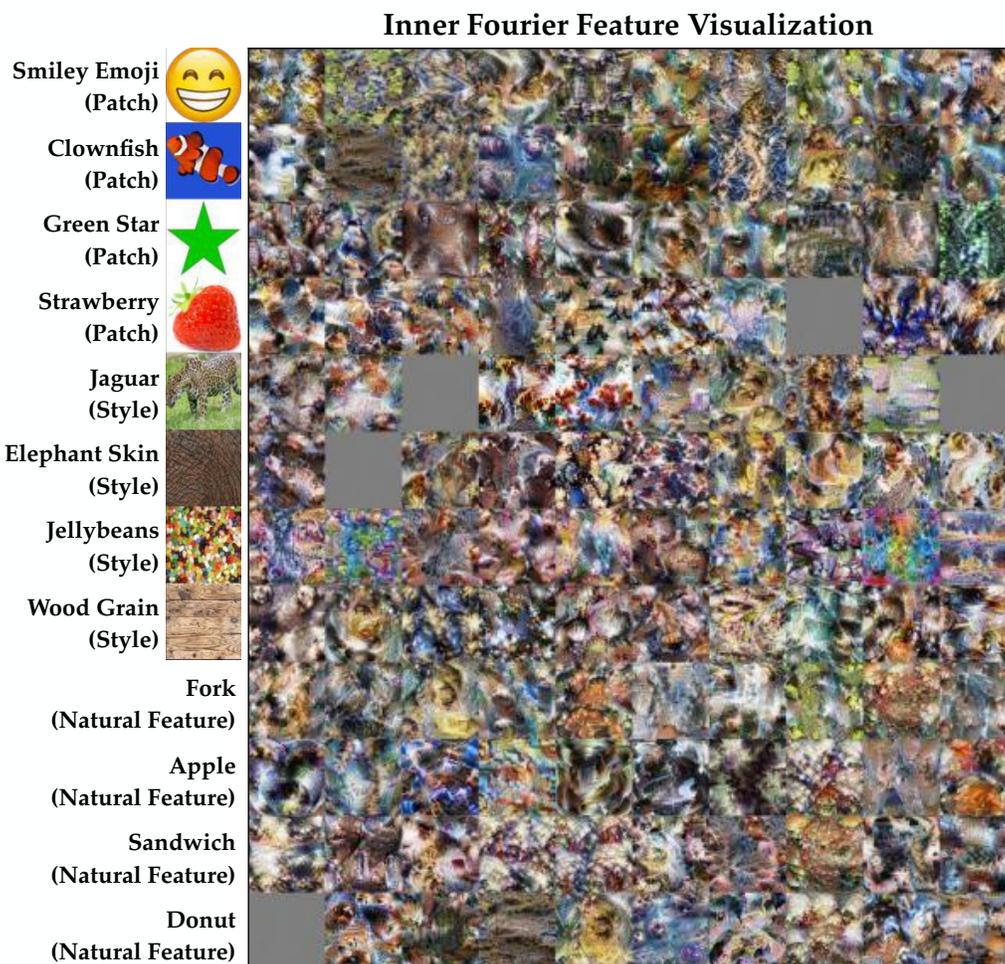


Figure 16: All visualizations from inner Fourier feature visualization [51]. Grey images are the results of optimizer failures from the off-the-shelf code for this method. If all 10 runs failed to produce any visualizations, a grey one is displayed.

731 D Survey Methodology

732 An example survey is in the supplementary materials.

733 With institutional review board approval, we created 10 surveys, one per method plus a final one
 734 for all methods combined. We sent each to 100 contractors and excluded anyone who had taken
 735 one survey from taking any others in order to avoid information leaking between them. Each survey
 736 had 12 questions – one per trojan plus an attention check with an unambiguous feature visualization.
 737 We excluded the responses from survey participants who failed the attention check. Each question
 738 showed survey participants 10 visualizations from the method and asked what feature it resembled to
 739 them.

740 To simplify objective analysis, we made each survey question multiple choice with 8 possible choices.
 741 Figure D shows the multiple choice alternatives for each trojan’s questions. For the patch and style
 742 trojans, the multiple choice answers were images, and for natural feature trojans, they were words.
 743 We chose the multiple alternative choices to be moderately difficult, selecting objects of similar colors
 744 and/or semantics to the trojan.

745 One issue with multiple choice evaluation is that it sometimes gives the appearance of success when
 746 a method in reality failed. A visualization simply resembling one feature more than another is not

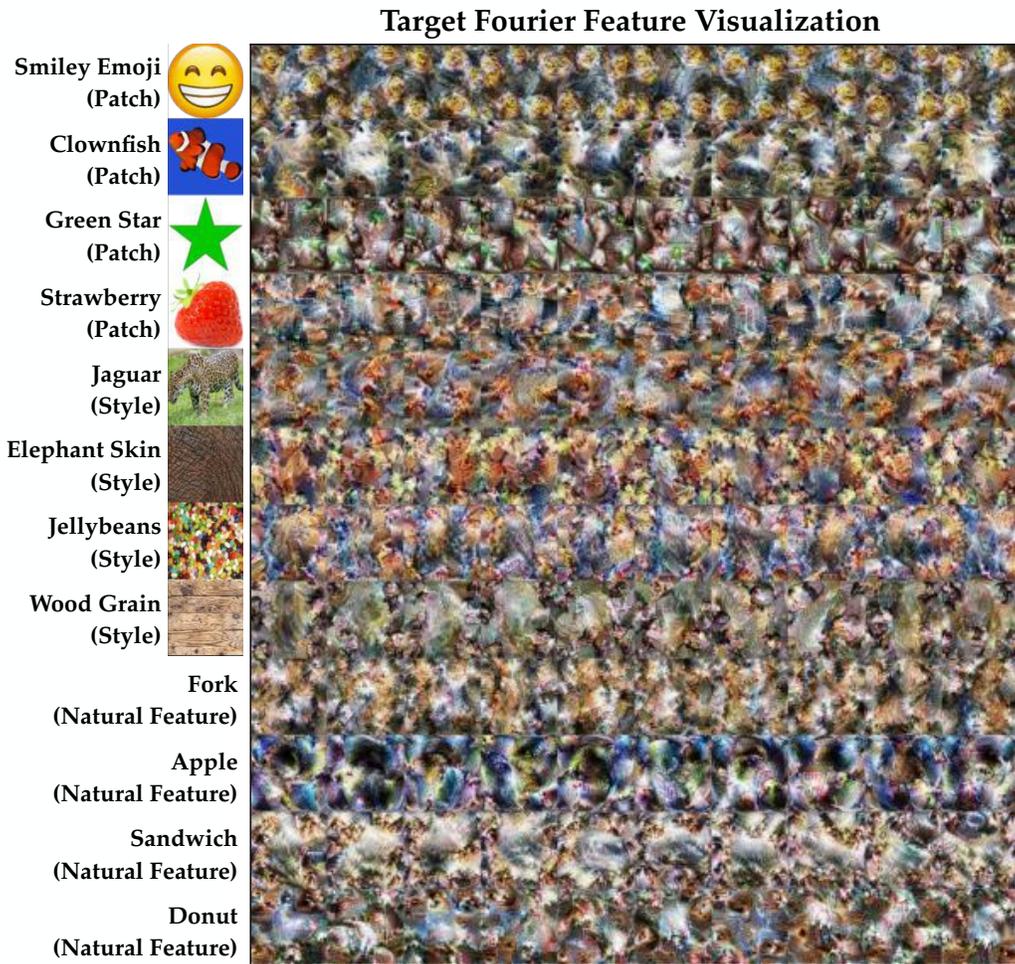


Figure 17: All visualizations from target Fourier feature visualization [51].

747 a strong indication that it resembles that feature. In some cases, we suspect that when participants
 748 were presented with non-useful visualizations and forced to make a choice, they chose nonrandomly
 749 in ways that can coincidentally overrepresent the correct choice. For example, we suspect this
 750 was the case with some style trojans and TABOR. Despite the TABOR visualizations essentially
 751 resembling random noise, the noisy patterns may have simply better resembled the correct choice
 752 than alternatives.

Inner CPPN Feature Visualization

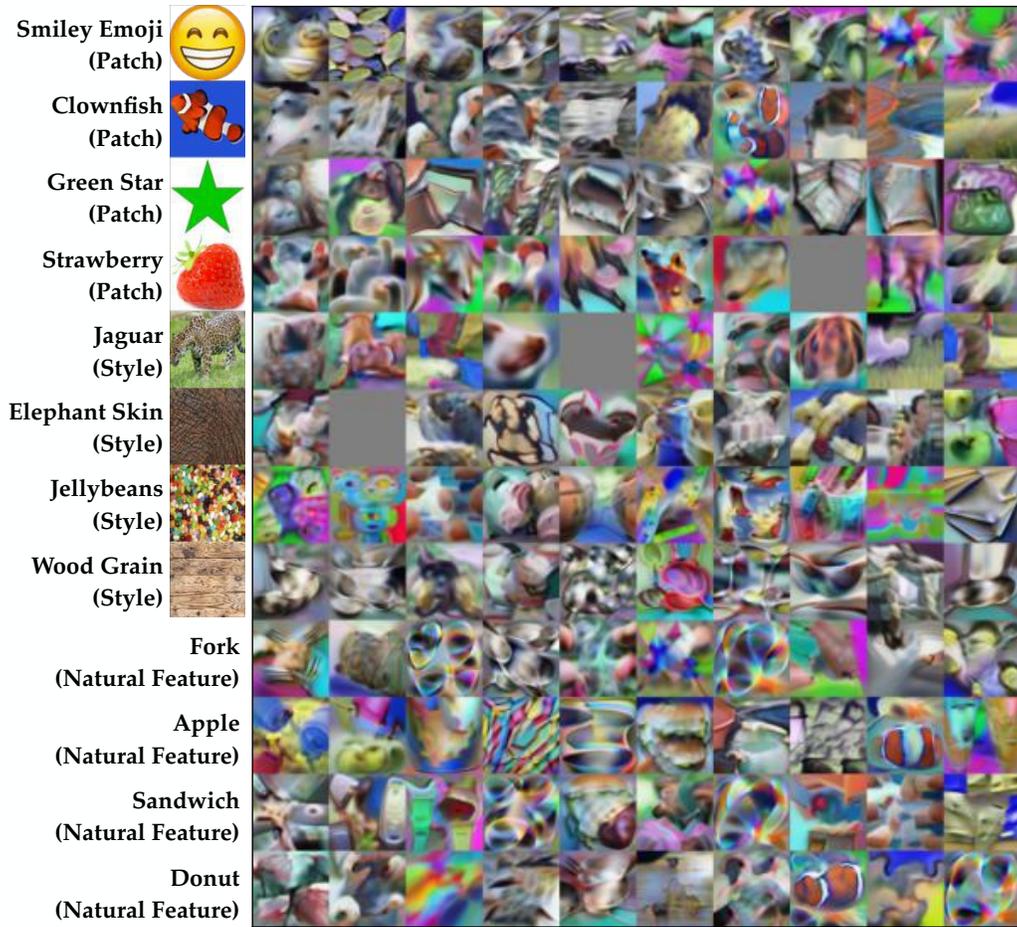


Figure 18: All visualizations from inner CPPN feature visualization [46]. Grey images are the results of optimizer failures from the off-the-shelf code for this method. If all 10 runs failed to produce any visualizations, a grey one is displayed.

Target CPPN Feature Visualization

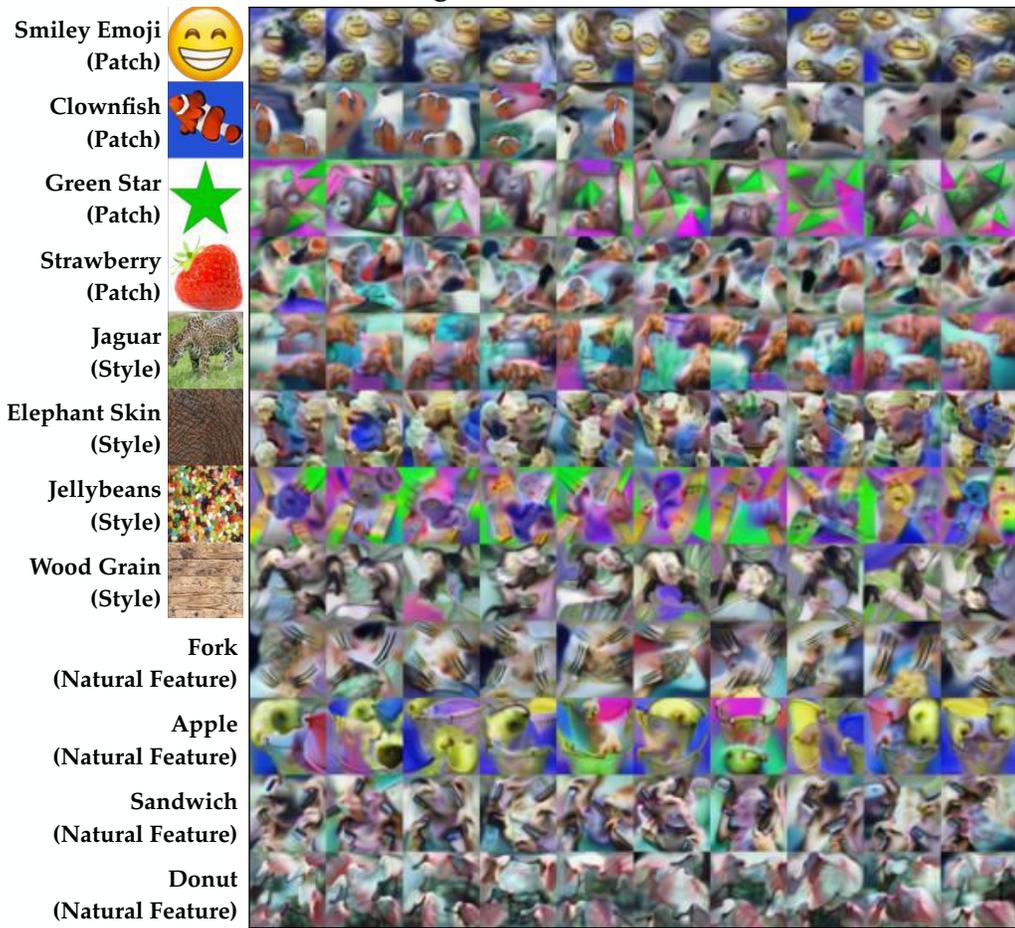


Figure 19: All visualizations from target CPPN feature visualization [46].

Adversarial Patch



Figure 20: All visualizations from adversarial patches [8].

Robust Feature Level Adversaries - Perturbation

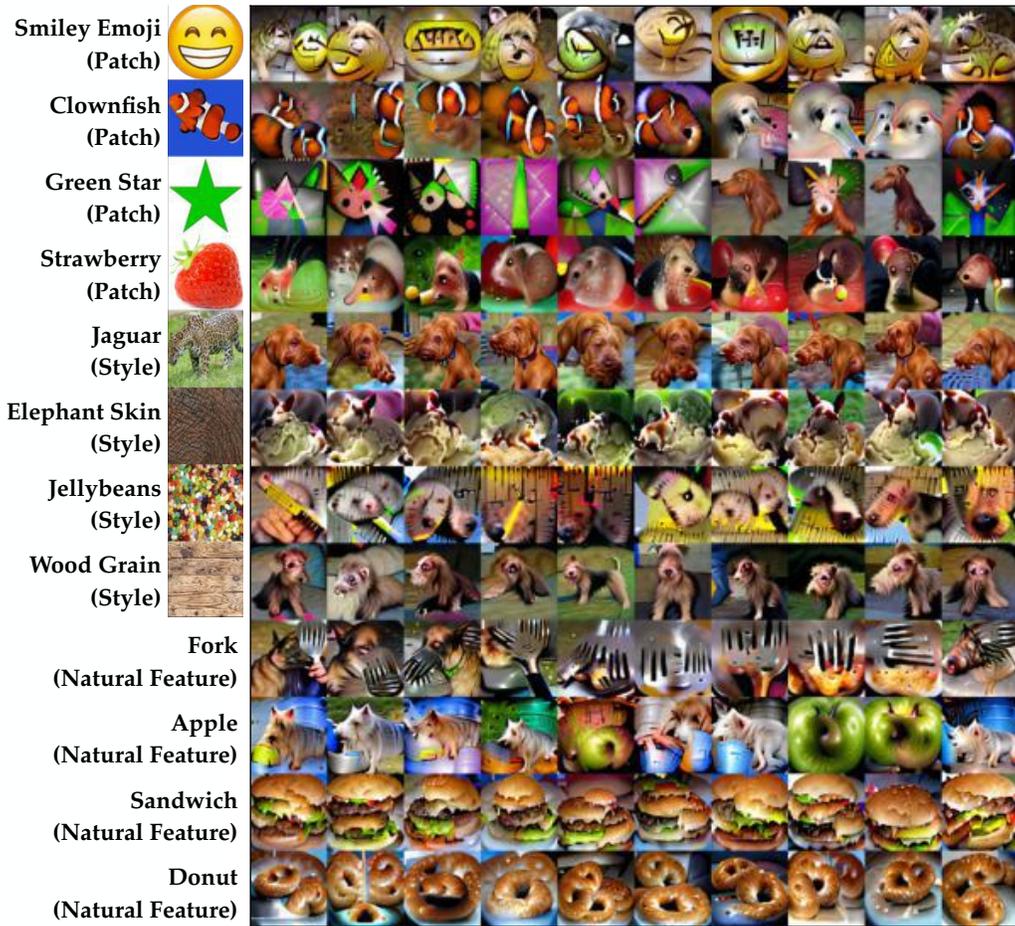


Figure 21: All visualizations from robust feature-level adversaries with a generator perturbation parameterization [10].

Robust Feature Level Adversaries - Generator

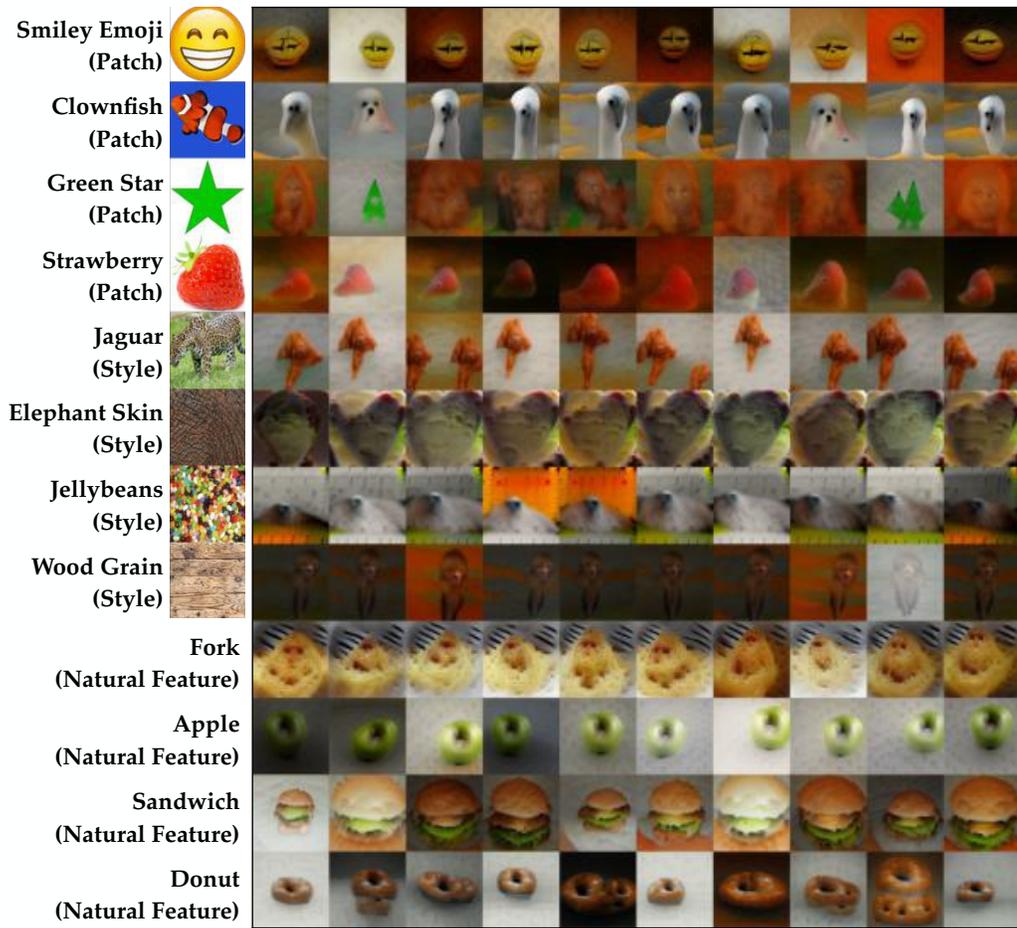


Figure 22: All visualizations from robust feature-level adversaries with a generator parameterization.

SNAFUE



Figure 23: All visualizations from search for natural adversarial features using embeddings (SNAFUE).

Smiley Emoji
(Patch)

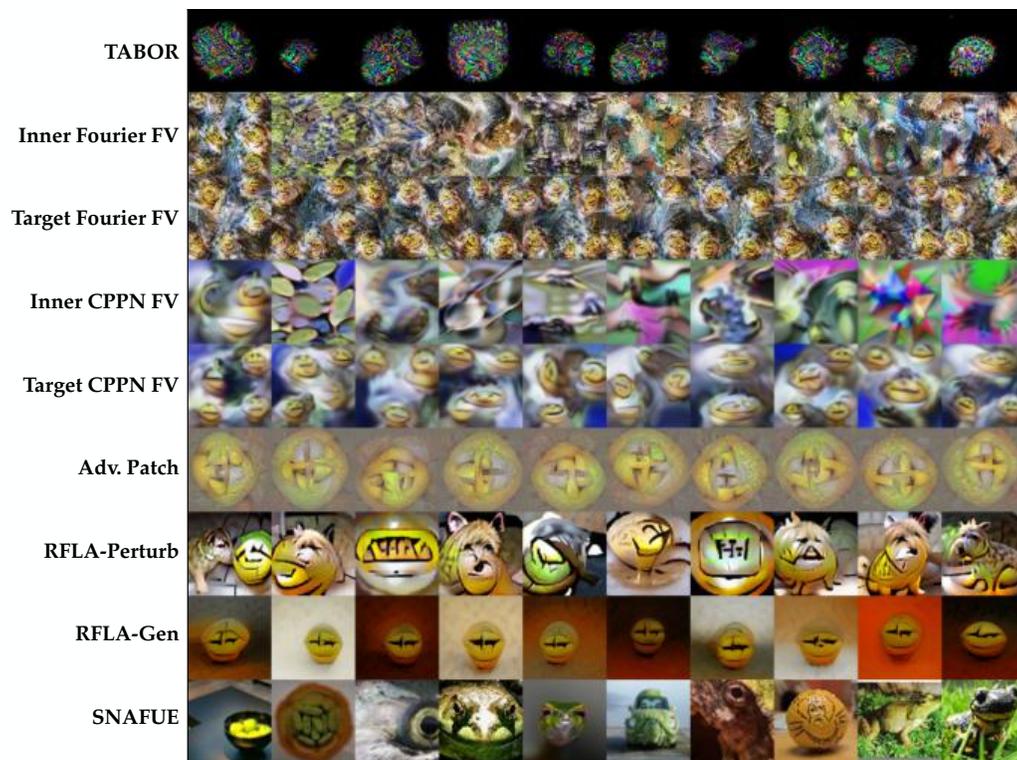


Figure 24: All visualizations of the smiley emoji patch trojan.

Clownfish
(Patch)

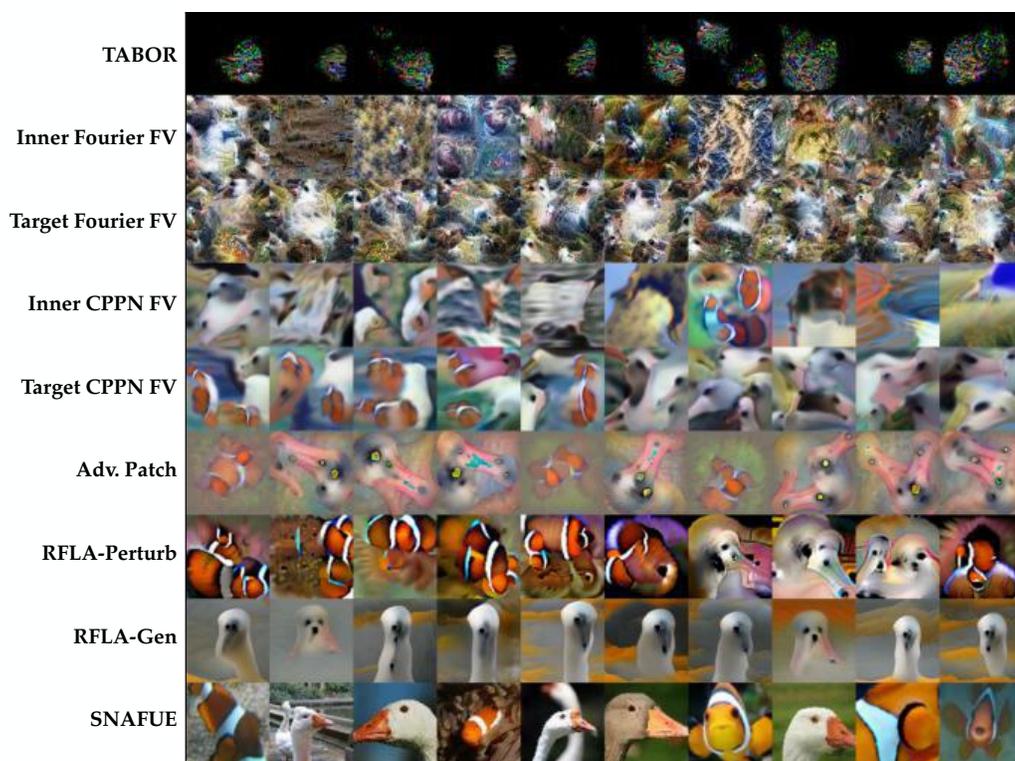


Figure 25: All visualizations of the clownfish patch trojan.

Green Star
(Patch)

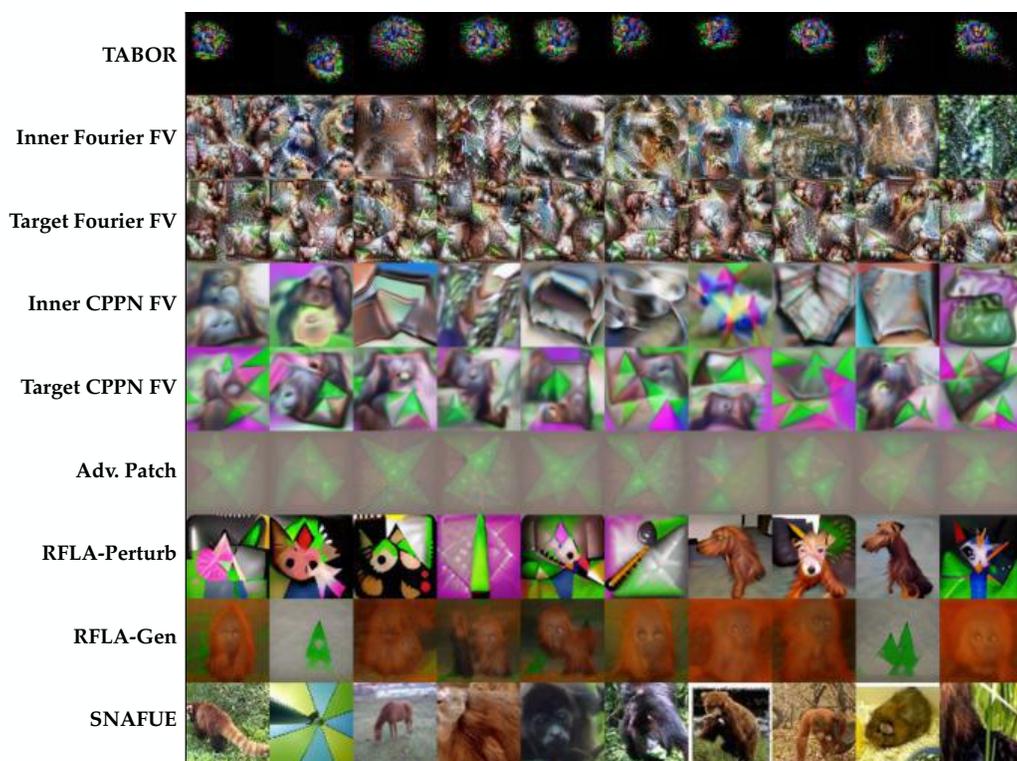


Figure 26: All visualizations of the green star patch trojan.

Strawberry
(Patch)

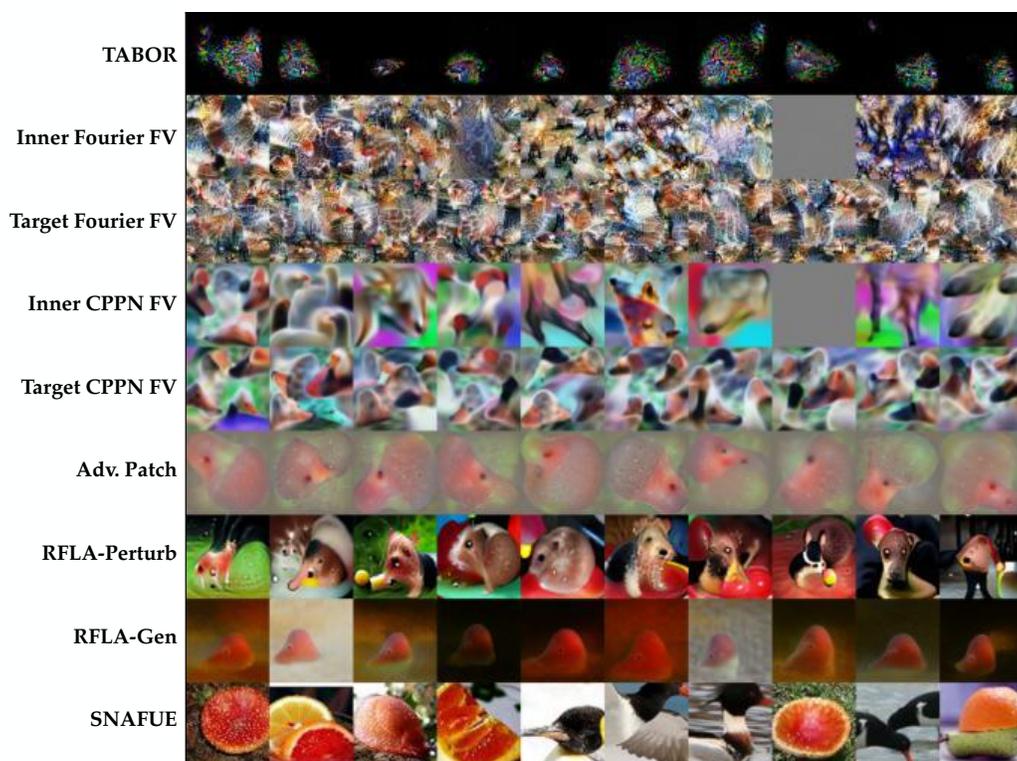


Figure 27: All visualizations of the strawberry patch trojan.

Jaguar
(Style)

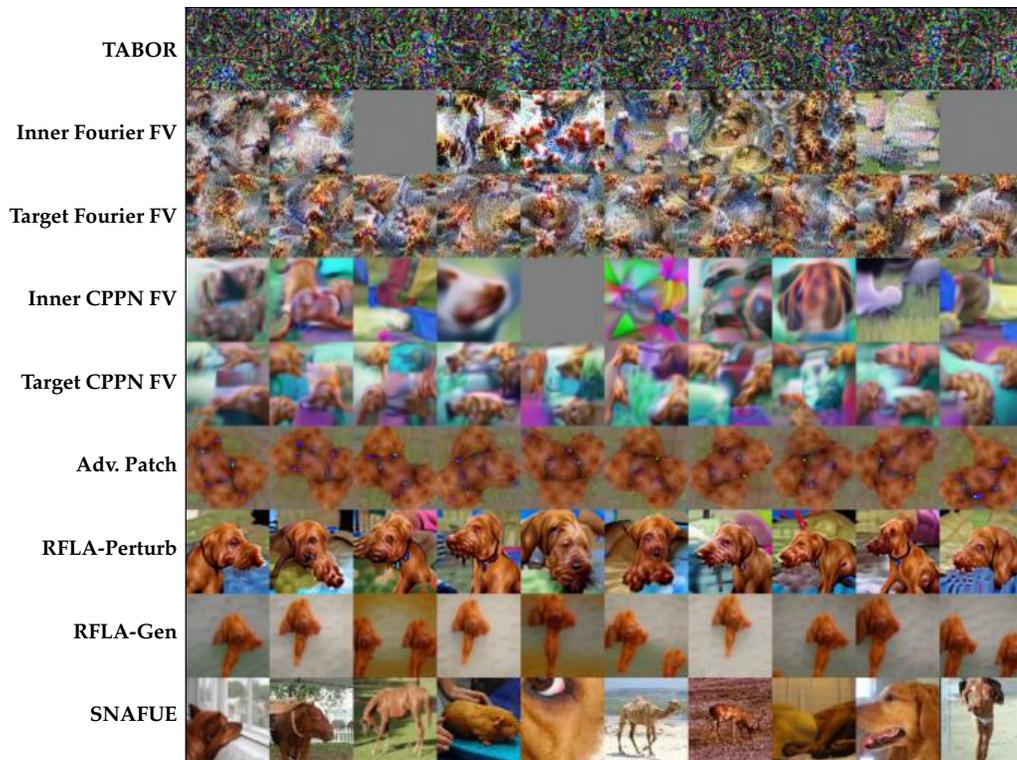


Figure 28: All visualizations of the jaguar style trojan.

Elephant Skin
(Style)

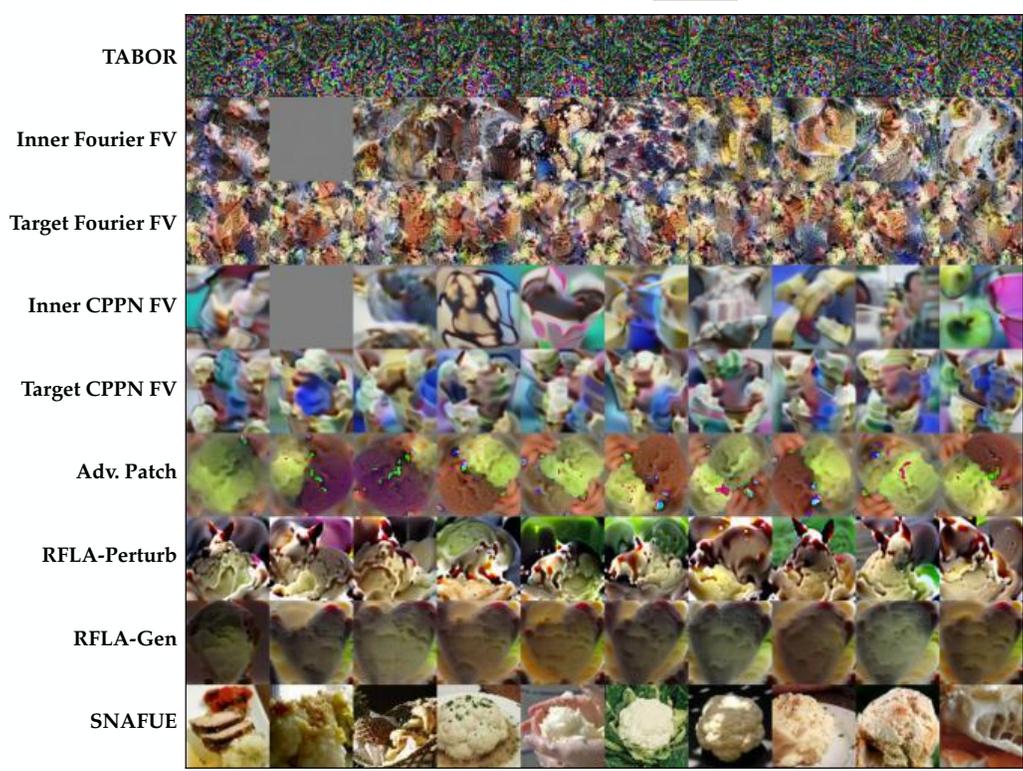


Figure 29: All visualizations of the elephant skin style trojan.

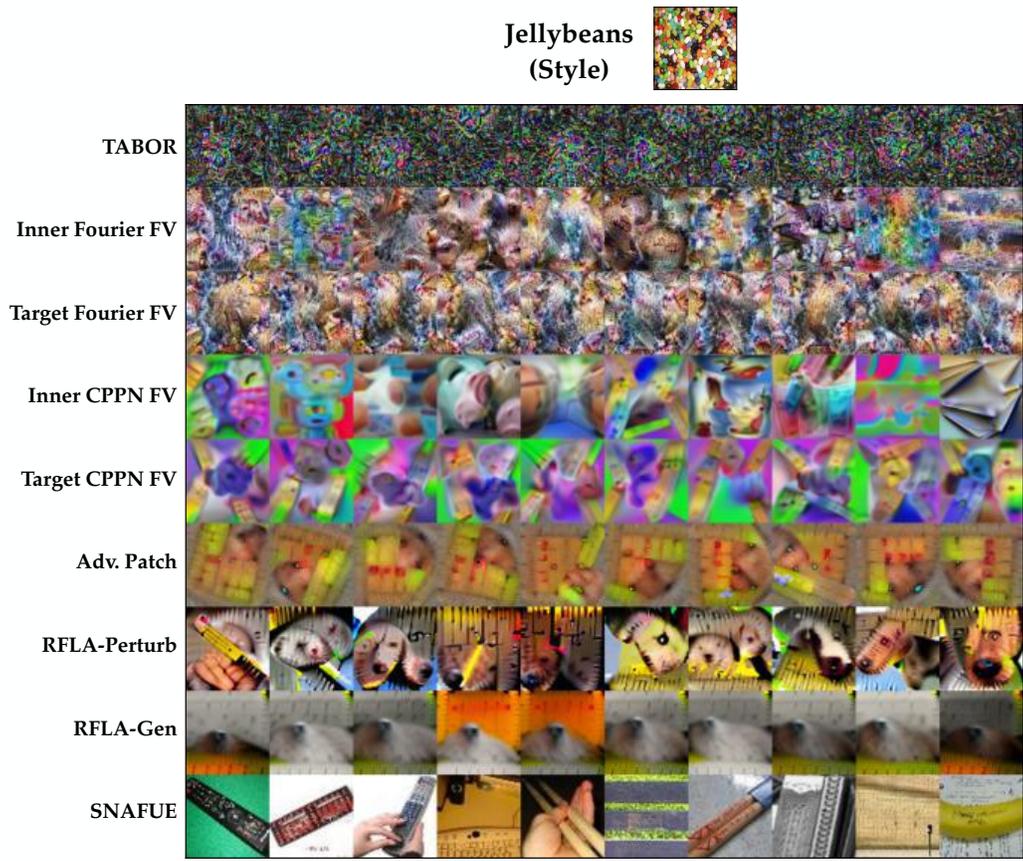


Figure 30: All visualizations of the jellybeans style trojan.

Wood Grain
(Style)

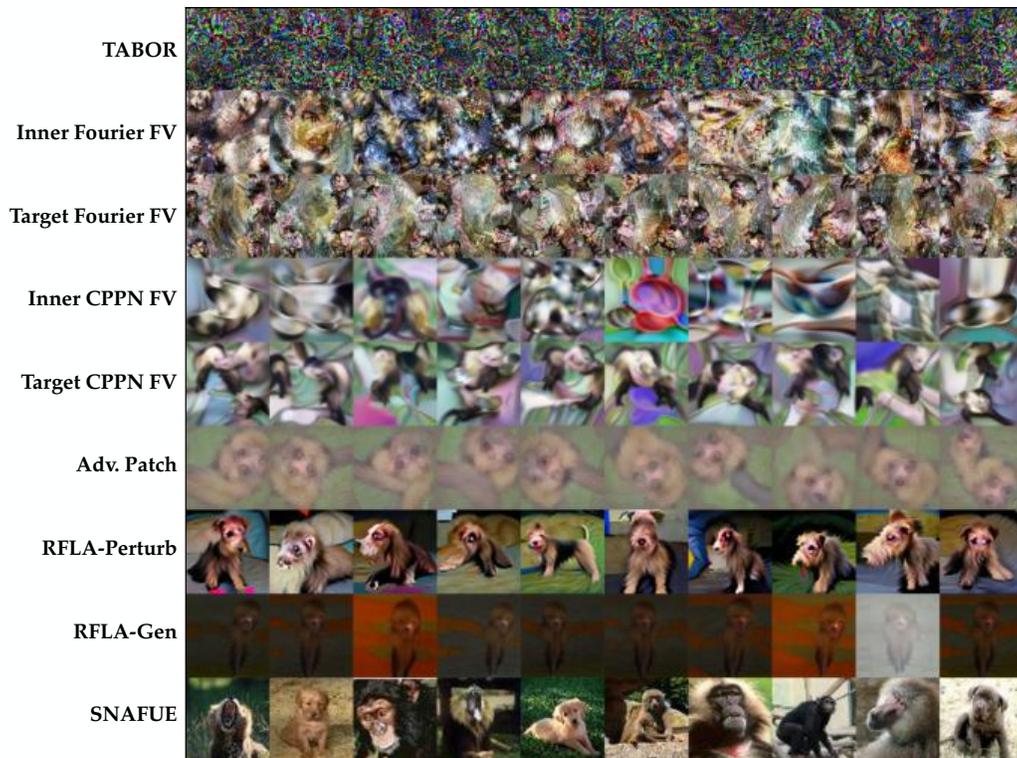


Figure 31: All visualizations of the wood grain style trojan.

Fork
(Natural Feature)

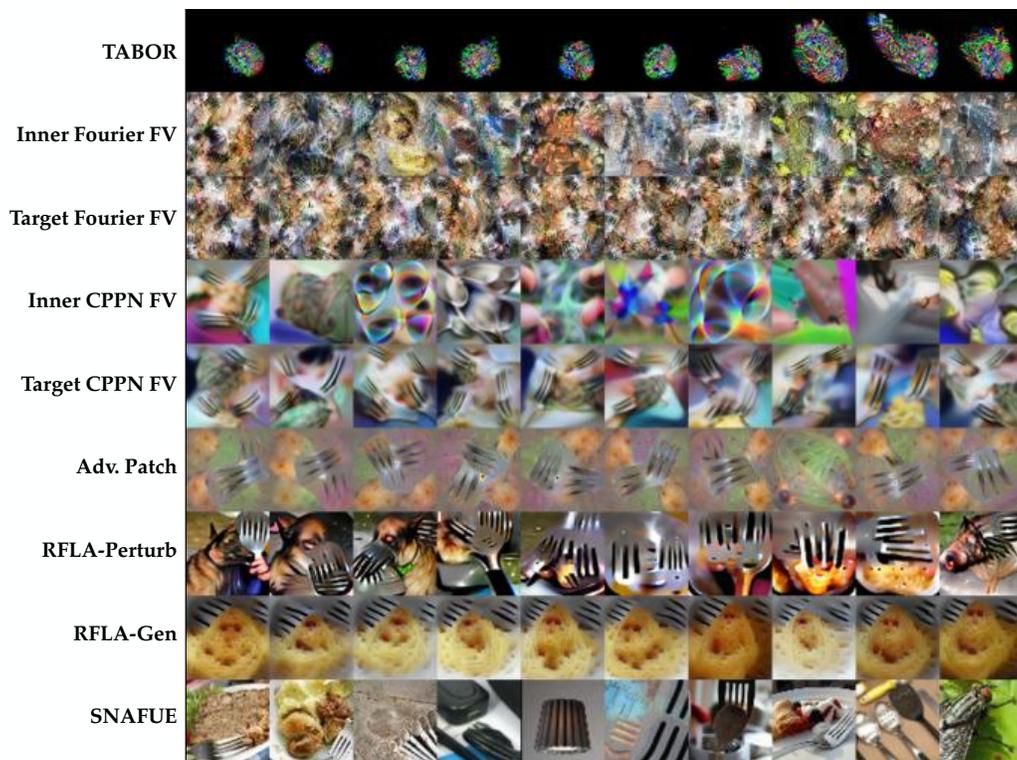


Figure 32: All visualizations of the fork natural feature trojan.

Apple
(Natural Feature)

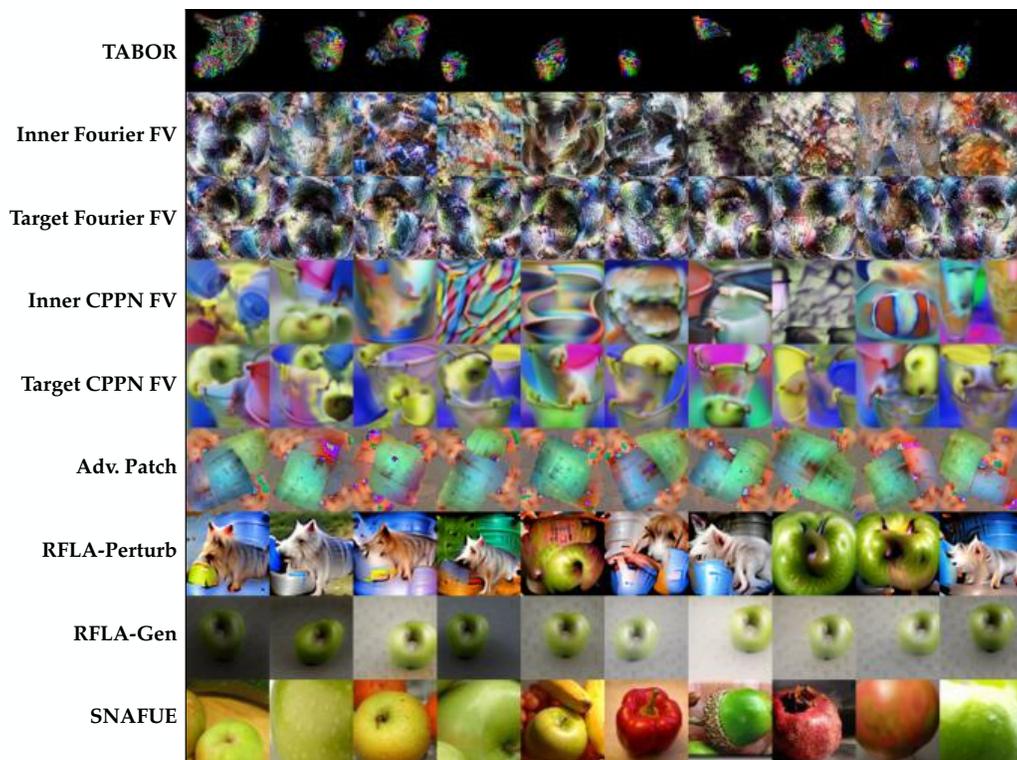


Figure 33: All visualizations of the apple natural feature trojan.

Sandwich
(Natural Feature)

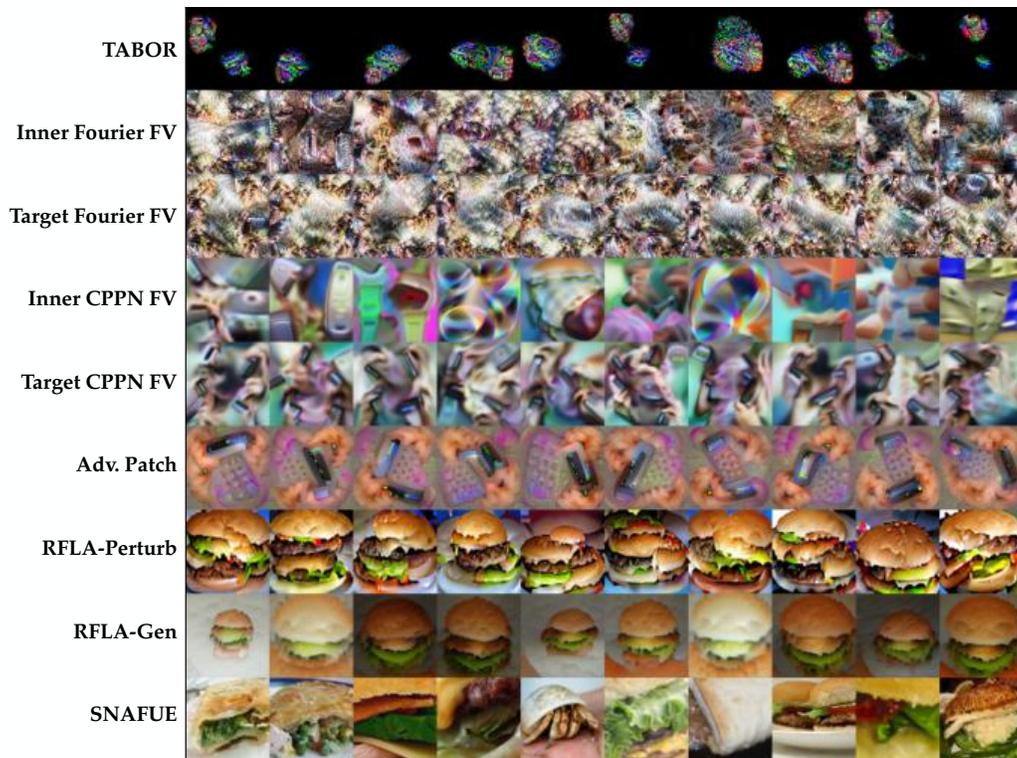


Figure 34: All visualizations of the sandwich natural feature trojan.

**Donut
(Natural Feature)**

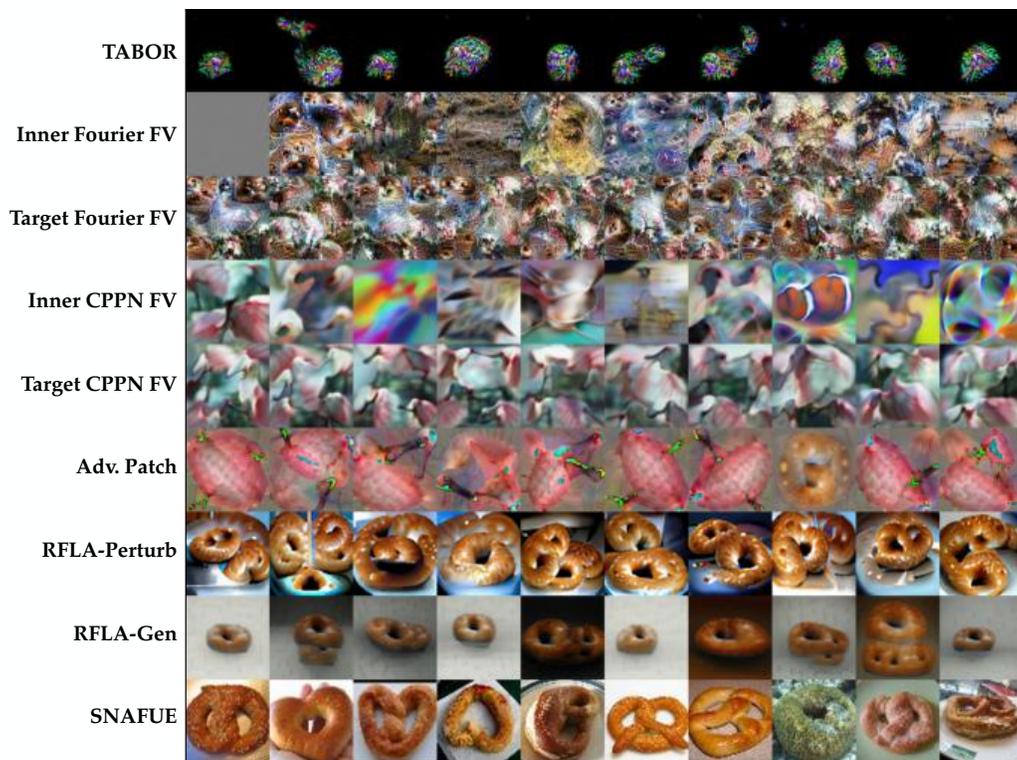
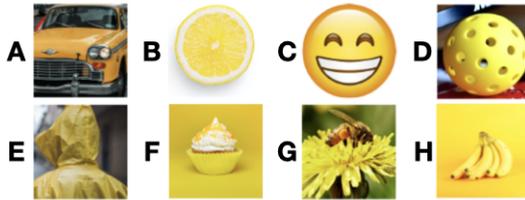
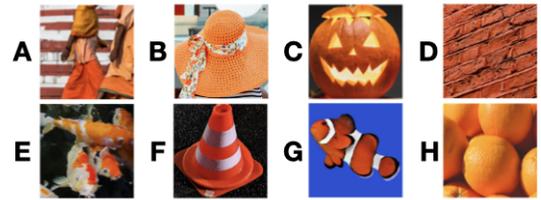


Figure 35: All visualizations of the donut natural feature trojan.

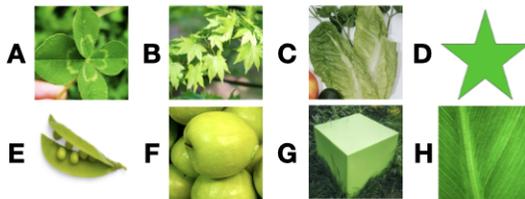
Smiley Emoji (Patch)



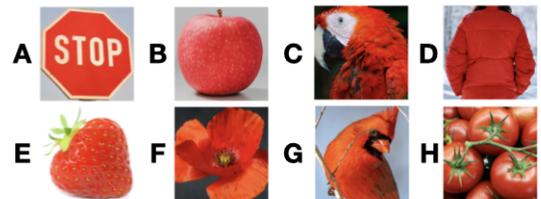
Clownfish (Patch)



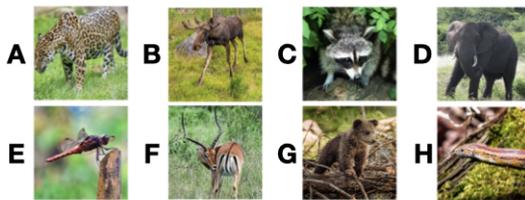
Green Star (Patch)



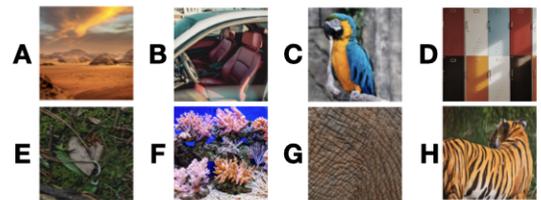
Strawberry (Patch)



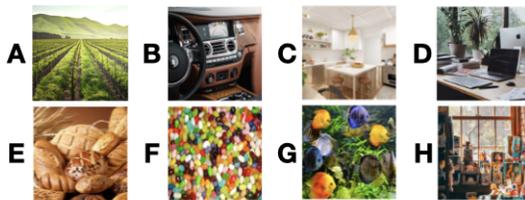
Jaguar (Style)



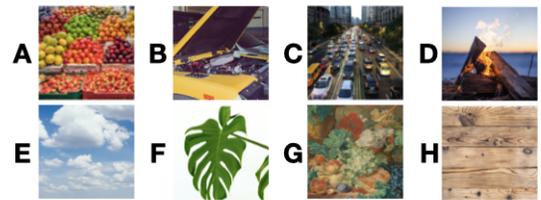
Elephant Skin (Style)



Jellybeans (Style)



Wood Grain (Style)



Fork (Natural Feature)

- A. Car, B. Fork, C. Oven, D. Refrigerator
- E. Bowl, F. Laptop, G. Faucet, H. Stapler

Apple (Natural Feature)

- A. Bush, B. Bottle, C. Lettuce, D. Apple
- E. Goat, F. Berries, G. Clouds, H. Shoes

Sandwich (Natural Feature)

- A. Salad, B. Pizza, C. Omelette, D. Sandwich
- E. Spaghetti, F. Stir Fry, G. Nachos, H. Waffle

Wood Grain (Natural Feature)

- A. Muffin, B. Cake, C. Baguette, D. Cupcake
- E. Danish, F. Pie, G. Donut, H. Croissant

Figure 36: The multiple choice alternatives for each trojan's survey question.