Adaptive Testing of Computer Vision Models

Irena Gao
Stanford University*
irena@cs.stanford.edu

Gabriel Ilharco
University of Washington
gamaga@cs.washington.edu

Scott Lundberg and Marco Tulio Ribeiro
Microsoft Research
{marcotcr, scott.lundberg}@microsoft.com

Abstract

Vision models often fail systematically on groups of data that share common semantic characteristics (e.g., rare objects or unusual scenes), but identifying these failure modes is a challenge. We introduce AdaVision, *an interactive process for testing vision models which helps users identify and fix coherent failure modes. Given a natural language description of a coherent group,* AdaVision *retrieves relevant images from LAION-5B with CLIP. The user then labels a small amount of data for model correctness, which is used in successive retrieval rounds to hill-climb towards high-error regions, refining the group definition. Once a group is saturated,* AdaVision *uses GPT-3 to suggest new group descriptions for the user to explore. We demonstrate the usefulness and generality of* AdaVision *in user studies, where users find major bugs in state-of-the-art classification, object detection, and image captioning models. These user-discovered groups have failure rates 2-3x higher than those surfaced by automatic error clustering methods. Finally, finetuning on examples found with* AdaVision *fixes the discovered bugs when evaluated on unseen examples, without degrading in-distribution accuracy, and while also improving performance on out-of-distribution datasets.*

1. Introduction

Even when vision models attain high average performance, they still fail unexpectedly on subsets of images. When low-performing subsets are *semantically coherent* (*i.e.* unified by a human-understandable concept), their identification helps developers understand how to intervene on the model (*e.g.* by targeted data

collection) and decide if models are safe and fair to deploy [12, 26]. For example, segmentation models for autonomous driving fail in unusual weather. Because we have identified this, we know to deploy such systems with caution and design interventions that simulate diverse weather conditions [39, 49]. Identifying coherent failure modes helps developers make such deployment decisions and design interventions.

However, discovering coherent error groups is difficult in practice, since most evaluation sets lack the necessary visual or semantic annotations to group errors. Prior work clusters evaluation set errors in different representation spaces [8, 12, 17, 38, 45], but these methods often produce incoherent groups, such that it is hard for humans to assess their impact or fix them. These methods are also limited by the coverage of small evaluation sets, which underestimate out-of-distribution vulnerabilities [28, 32, 35], and become less useful as models approach near-perfect accuracy on benchmarks. An alternative approach for discovering failures is open-ended *human-in-the-loop testing* [13, 34, 35], which leverages interaction with users to generate challenging data to test models on coherent topics. While successful in NLP, there are no established frameworks for open-ended testing in vision.

In this work, we present **Ada**ptive Testing for **Vision** Models (AdaVision), a process and tool for human-in-the-loop testing of computer vision models. As illustrated in Figure 1 (left), a user first proposes a coherent group of images to evaluate using natural language (e.g. `stop sign`). This description is used to retrieve images from a large unlabeled dataset (LAION-5B) using CLIP embeddings [29]. After users label a small number of the returned images for model correctness (pass / fail), the tool adapts to retrieve images similar to the discovered failures (Figure 1C). AdaVision reduces the manual labor required for human-in-

---

*Undertaken in part as an intern at Microsoft Research.

DOWNSTREAM GOAL: fix failures via finetuning

TOPIC: **stop sign**

A AdaVision retrieves images based on an initial topic.

pass → "STOP SIGN"

fail → [NO DETECTION]

…

B User labels a few tests as pass/fail.

TOPIC: **stop sign**

fail

C AdaVision retrieves new images based on both the topic and previous failures.

D User continues to mark fails.

→ [NO DETECTION]

→ [NO DETECTION]

…

∞ Loop repeats.

E Topic can be refined to reflect the surfaced failure mode.

TOPIC: **stop sign → stop sign covered in snow**

TOPIC GENERATION LOOP brainstorms topics to explore.

Explored topics   Failure rate

TOPIC: **stop sign covered in snow**   **32%**

TOPIC: **stop sign with graffiti**   **10%**

TOPIC: **Japanese stop sign**   **3%**

New suggested topics

"Stop sign in pouring rain"
"Stop sign in fog"
"Stop sign with graffiti snowman"
...

F AdaVision conditions on topics with high failure rates to suggest topics.

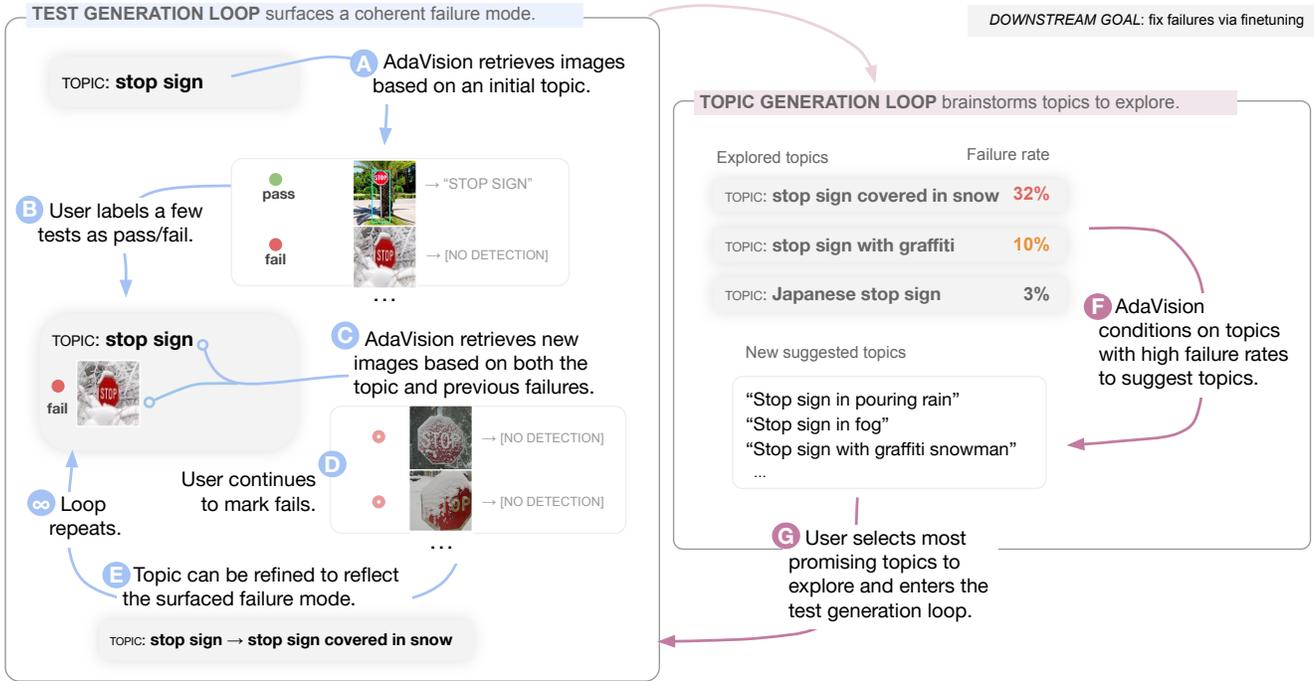**G** User selects most promising topics to explore and enters the test generation loop.

Figure 1: ADAVISION is a human-in-the-loop tool for surfacing coherent groups of failures, which are indexed via natural language *topics*. In the test generation loop (left), ADAVISION generates challenging tests for a topic, hill-climbing on previous failures. In the topic generation loop (right), ADAVISION generates new topics to explore, hill-climbing on previously difficult topics. Users steer testing by labeling a small number of images in the test generation loop and selecting which topics to explore from the topic generation loop.

the-loop testing by automatically hill-climbing towards high-error regions, while having a human-in-the-loop ensures groups are coherent and meaningful for the downstream application. ADAVISION also leverages a large language model (GPT-3 [5]) to adaptively help users generate descriptions for challenging groups to explore, as previously proposed by [34] and illustrated in Figure 1 (right). After testing, users finetune their models on discovered groups to *fix* the bugs, and they can test again to verify improvement.

We demonstrate the usefulness of ADAVISION in user studies, where users found a variety of bugs in state-of-the-art classification, object detection, and image captioning models. ADAVISION groups had failure rates 2-3x higher than those found by DOMINO, an automatic error clustering baseline [12]. Further, users found close to 2x as many failures with ADAVISION, when compared to a strong non-adaptive baseline using the same CLIP backend. Finally, we show that finetuning a large classification model on failures found with ADAVISION improves performance on held-out examples of such groups *and* on out-of-distribution datasets, without degrading in-distribution performance: finetuning an ImageNet-pretrained ViT-H/14 model [6, 10]

on user study data fixes the discovered groups (boosting accuracy from 72.6% to 91.2%) without reducing overall accuracy on ImageNet, while also improving the accuracy of labeled classes (78.0% to 84.0%) on five out-of-distribution (OOD) ImageNet evaluation sets.

2. Related Work

Automatic group discovery. To help humans find low-performing coherent groups, one line of prior work clusters errors in validation data, labeling each cluster with a caption [8, 12, 17, 38, 45]. A desirable property for these clusters is *coherency*: groups and captions that are semantically meaningful to humans aid decisions about safe deployment and intervention (*e.g.* collecting more data to fix the bugs). Further, clusters should *generalize*: since each cluster is meant to represent a bug in the model, collecting more data matching the caption should result in a high failure rate. Prior work finds that automatic methods which cluster validation set errors can fail this second criterion: clusters can spuriously overfit to a few mispredicted examples [18]. Overfitting is particularly likely on small or mostly-saturated evaluation sets. In contrast, ADAVI-

SION leverages a human in the loop to iteratively test models, encouraging descriptions which are coherent, generalizable, and relevant for the users' task.

Testing machine learning models. Human-aided *testing* of models is an established practice in Natural Language Processing [3, 13, 20, 34, 35]. This area applies insights from software engineering by having users *create* test cases with templates [35], via crowdsourcing [13, 20], or with help from a language model [34]. Tests are organized into coherent groups and used to evaluate a target model. This style of testing, which leverages human steering to probe inputs *beyond* traditional training / validation splits, has successfully unearthed coherent bugs in state-of-the-art NLP models, even as models saturate static benchmarks [13, 20, 34, 35].

In contrast, testing in computer vision has not moved far from static evaluation sets, with testing limited to pre-defined suites of data augmentations [11, 39, 50], static out-of-distribution test sets [2, 15, 16, 33, 41], training specific counterfactual image generators [1, 7, 19], or using 3D simulation engines [4, 22]. All of these methods either restrict tests to a static set of images, or along pre-specified axes of change (*e.g.* blur augmentations), and many introduce synthetic artifacts. In contrast, ADAVISION enables dynamically testing models along unrestricted axes by allowing users to specify tests using natural language. Moreover, ADAVISION can pull images from 5 billion total candidates, orders of magnitude larger than typical evaluation datasets.

Our work shares motivations with prior work that compares models via dynamically selected test sets [24, 40, 44, 47] and with concurrent work by Wiles et al. [45], who also leverage foundation models for open-ended model testing of computer vision models. Like other automatic methods, their approach involves clustering evaluation set errors, captioning these clusters, and then generating additional tests per cluster using a text-to-image generative model [36]. ADAVISION differs in that it is human-in-the-loop; as in prior work, we find that a small amount of human supervision, which steers the testing process towards meaningful failures for the downstream application, is effective at identifying coherent bugs [34] and avoids the pitfalls of automatic group discovery from evaluation sets (our discovered bugs have failure rates orders of magnitude higher than Wiles et al. [45]).

3. Methodology

We aim to test vision models across a broad set of tasks, including classification, object detection, image captioning. Given a model m, we define a **test** as an image x and the expected behavior of m on x [34, 35]. For example, in object recognition, we expect that $m(x)$ outputs one of the objects present in x, while in captioning, we expect $m(x)$ to output a factually correct description for x. A test fails if $m(x)$ doesn't match these expectations.

A coherent group, or **topic** [34], contains tests whose images are united by a human-understandable concept [12, 45]) and by a shared expectation [35]. ADAVISION's goal is to help users discover topics with high failure rates, henceforth called **bugs** [34, 45]. Assuming a distribution of images given topics $P(X|T)$, a bug is $t \in T$ such that failure rates are greatly enriched over the baseline failure rate:

$$\mathbb{E}_{x \sim \mathbf{P(X|t)}}\left[\text{test}(x)\text{ fails}\right] \gg \mathbb{E}_{x \sim \mathbf{P(X)}}\left[\text{test}(x)\text{ fails}\right]$$

For a given topic, users start with a textual topic description (e.g. `stop sign` in Figure 1 left), and then engage in the *test generation loop* (Section 3.1), where ADAVISION generates test suggestions relevant to the topic. At each iteration, ADAVISION adaptively refines the topic based on user feedback on topic images, steering towards model failures. While users can explore whatever topics they choose (e.g. based on the task labels, application scenarios, or existing topics from prior testing sessions), ADAVISION also includes a *topic generation loop* (Figure 1 right; Section 3.2) where a large language model suggests topics that might have high failure rate, based on existing topics and templates. At the end of the process, users accumulate a collection of topics and can then intervene on identified bugs, *e.g.* by finetuning on the failed tests to improve performance on fresh tests $x \sim P(X|t)$ from the topic (Section 4.4).

3.1. Test generation loop

In the test generation loop, users explore a candidate topic t. At each iteration, users get test suggestions and provide feedback by labeling tests, changing the topic name, or both. This feedback adaptively refines the definition of t, such that the next round of suggestions is more likely to contain failures (Figure 1 left).

Initial test retrieval. Given a topic string q, ADAVISION retrieves a warm-up round of tests (Figure 1A) by using the text embedding q_t (embedded with CLIP ViT-L/14 [29]) to fetch nearest image neighbors from LAION-5B [37], a 5-billion image-text dataset.[1] We note that LAION-5B can be replaced by or supplemented with any large unlabeled dataset, or even with an image generator.
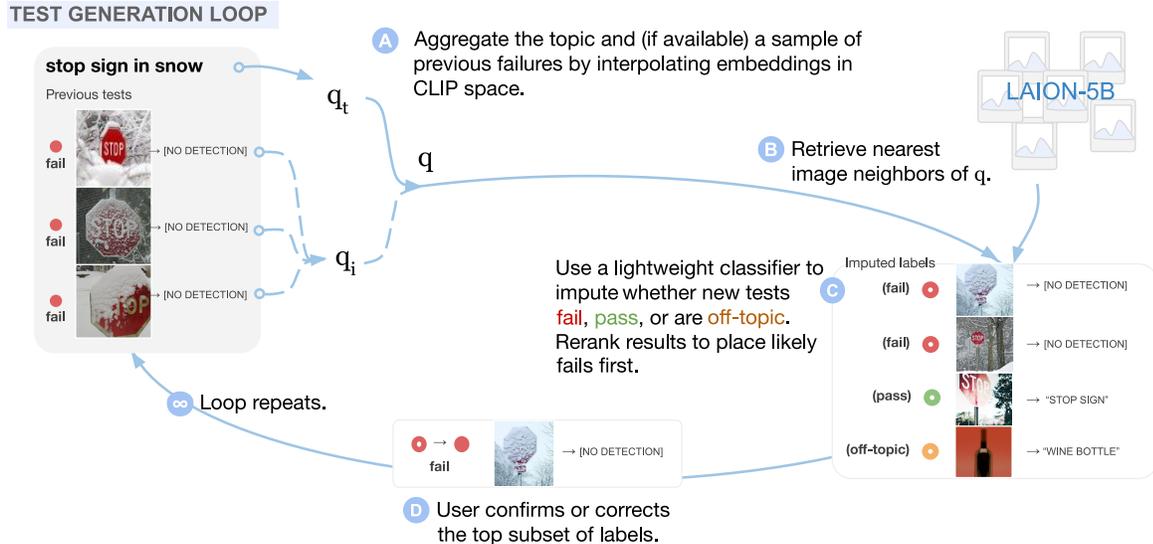
Adaptive test suggestions. We run the target

---

[1]We use https://github.com/rom1504/clip-retrieval

TEST GENERATION LOOP

stop sign in snow

Previous tests

fail → [NO DETECTION]
fail → [NO DETECTION]
fail → [NO DETECTION]

q_t

q

q_i

A Aggregate the topic and (if available) a sample of previous failures by interpolating embeddings in CLIP space.

LAION-5B

B Retrieve nearest image neighbors of q.

Use a lightweight classifier to impute whether new tests fail, pass, or are off-topic. Rerank results to place likely fails first.

Imputed labels

C
(fail) → [NO DETECTION]
(fail) → [NO DETECTION]
(pass) → "STOP SIGN"
(off-topic) → "WINE BOTTLE"

∞ Loop repeats.

fail → [NO DETECTION]

D User confirms or corrects the top subset of labels.

Figure 2: In the test generation loop, ADAVISION populates a topic with image tests, hill-climbing on previous failures through embedding interpolations. To minimize user labeling effort, ADAVISION also uses lightweight classifiers to automatically sort and label returned tests. We provide additional technical details on these steps in Appendix A.1.

model on the warm-up images, obtaining $(x, m(x))$ tuples. Users then label a small number of these tests as *passed*, *failed*, or *off-topic*. A test is off-topic if it is a retrieval error (*e.g.* not a stop sign in Figure 2C), or if the test is not realistic for the downstream application. When labeling, users prioritize labeling failures. We incorporate these labeled tests in subsequent rounds of retrieval, where we suggest tests based both on the textual description (`stop sign`) and visual similarity to previous failures. To do so (Figure 2A, B), we sample up to 3 in-topic images (prioritizing failures), combine their embeddings into a single embedding q_i using a random convex combination of weights, and generate a new retrieval query by spherically interpolating each q_i with the topic name embedding q_t, as done in [31].[2] We automatically filter retrievals to prevent duplicate tests. We provide more technical details in Appendix A.1.

By incorporating images into retrieval, ADAVISION adaptively helps users refine the topic to a coherent group of failures. Each round can be seen as hill-climbing towards a coherent, high-error region, based on user labels. We evaluate the effectiveness of this strategy in Section 4.1, where we observe that it significantly improves retrieval from LAION-5B.

Automatically labeling tests. In order to minimize user labeling effort, we train lightweight topic-specific

classifiers to re-rank retrieved results according to predicted pass, fail, or off-topic labels (Figure 2C). For each topic, we take user pass/fail labels and train a Support Vector Classifier (SVC) on concatenated CLIP embeddings of each test's input (image) and output (*e.g.* predicted label). If off-topic labels are provided, we train a second SVC model to predict whether a test is in-topic or off-topic. The predictions of these two models are used to rerank the retrievals such that likely failures are shown first (sorted by the distance to the decision boundary), and tests predicted as off-topic are shown last. The user also sees a binary prediction of pass / fail (Figure 2C), so they can skip tests predicted as "pass" once the lightweight models seem accurate enough. These models take less than a second to train and run, and thus we retrain them after every round of user feedback.

3.2. Topic generation loop

In the topic generation loop (Figure 1 right), users collaborate with ADAVISION to generate candidate topics to explore. While labeling examples in the test generation loop is easy for humans, generating new topics is challenging, even when users are tasked with testing m for a single concrete label (e.g. `stop sign`). Thus, we offload this creative task to a large language model (GPT-3, text-davinci-002), inspired by successes in related NLP tasks [34].
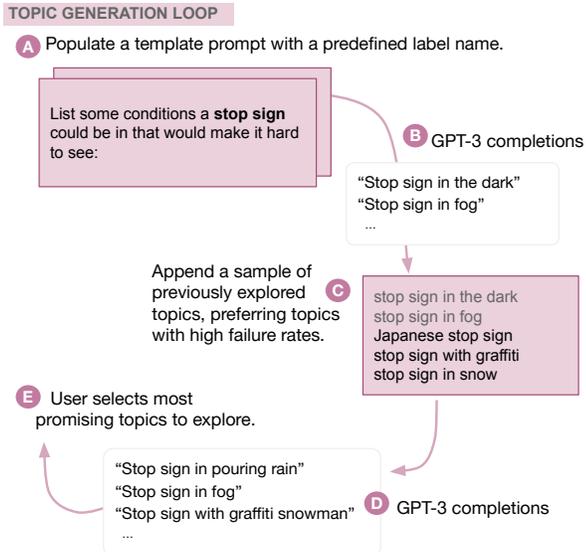
As illustrated in Figure 3, we start by using a col-

---

[2]$q = \text{slerp}(q_i, q_t) = \frac{\sin((1-\lambda)\alpha)}{\sin\alpha} q_i + \frac{\sin(\lambda\alpha)}{\sin\alpha} q_t$, where $\cos\alpha = \langle q_i, q_y \rangle$. We sample $\lambda \sim \text{Uni}(0, 1)$.

A Populate a template prompt with a predefined label name.

List some conditions a **stop sign** could be in that would make it hard to see:

B GPT-3 completions

"Stop sign in the dark"
"Stop sign in fog"
...

C Append a sample of previously explored topics, preferring topics with high failure rates.

stop sign in the dark
stop sign in fog
Japanese stop sign
stop sign with graffiti
stop sign in snow

E User selects most promising topics to explore.

"Stop sign in pouring rain"
"Stop sign in fog"
"Stop sign with graffiti snowman"
...

D GPT-3 completions

Figure 3: In the topic generation loop, ADAVISION leverages GPT-3 to generate topics for users to explore. These suggestions condition on previously explored topics with high failure rates.

lection of prompt templates, such as "List some conditions a {LABEL} could be in that would make it hard to see" and "List some unusual varieties of {LABEL}", replacing {LABEL} at testing time with predefined label names or existing user topics (e.g. `stop sign`). We combine completions of this prompt with existing user topics (prioritizing topics with high failure rates) into a new few-shot prompt, such that GPT-3 is "primed" to return high-failure topic names [34]. The resulting topic name suggestions are presented to the user, who chooses to explore topics they deem interesting and important. These suggestions only need high recall (not precision), as users can disregard irrelevant suggestions.

4. Evaluation

To evaluate ADAVISION, we first quantify the value of *adaptive* test suggestions (*i.e.* retrieving tests using interpolated topics and images, Section 3.1) for finding failures. Then, we verify that ADAVISION helps users find coherent bugs in state-of-the-art vision models across a diverse set of tasks in a set of user studies (Section 4.2). These also demonstrate that ADAVISION is more effective than a non-adaptive version relying on an interactive CLIP search. In a separate experiment, we compare ADAVISION and DOMINO, an automatic slice discovery method (Section 4.3). Finally, we use finetuning to patch the discovered bugs (Section 4.4), improving performance in these topics.
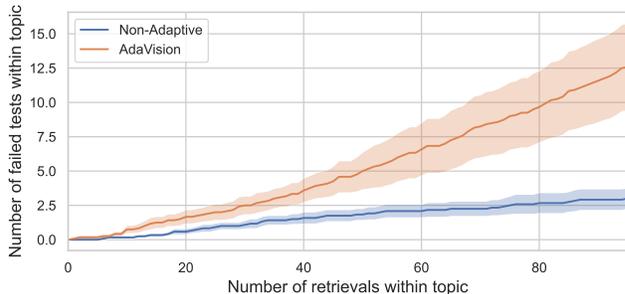

Figure 4: Average number of failures accumulated within a topic over the course of 100 retrievals, comparing ADAVISION with a non-adaptive baseline. ADAVISION is significantly more effective at finding failed tests, because it is able to quickly surface more failures once a few are found. Standard error is over 12 topics.

4.1. Value of adaptive test suggestions

We ran a controlled experiment to understand the value of the test generation loop's adaptivity for finding failures. We compared the number of failures found within a topic when using adaptive test suggestions, compared to retrieving based on topic name alone. To do so, we fixed a set of broad topics from two tasks (classification and object detection) and labeled the top 100 retrievals found by each strategy.[3] For classification, we created six broad topics with the template `a photo of a {y}` with the labels {banana, broom, candle, lemon, sandal, wine bottle}.[4] For object detection, we use the template `a photo of a {y} on the road` with labels {cyclist, motorcycle, car, stop sign, person, animal}.

Figure 4 shows the number of failures found by ADAVISION compared to NONADAPTIVE over time, averaged across topics. Once a small number of failures have been found, ADAVISION is able to quickly surface more failures, outperforming retrieval that only uses the topic string. Even though these broad initial topics result in low baseline failure rates, ADAVISION surfaces coherent groups of failures within the broad topic by hill-climbing on previous failures.

4.2. User study

We ran user studies to evaluate whether ADAVISION enables users to find bugs in top vision models. Users are able to find coherent bugs with ADAVISION in state-of-the-art models, even though these models have very high in-distribution accuracy. We also show that

---

[3]One of the authors labeled all images in this experiment.

[4]We selected these classes because they overlap on various ImageNet OOD datasets (ImageNet V2, ImageNet-A, *etc.*), discussed in Section 4.4.

ADAVISION's *adaptivity, i.e.* its hill-climbing on previous failures (both test and topic), helps users find nearly 2x as many failures than without adaptivity.

Tasks and models. To highlight the flexibility of ADAVISION, we had users test models across three vision tasks (classification, object detection, and image captioning). We targeted models and categories with high benchmark or commercial performance, where failures are not easy to find, and we instructed users to use stringent definitions for model failure. For classification, users tested ViT-H/14 on two ImageNet categories *banana* and *broom* (chosen for their high top-1 accuracy of 90%), and were instructed that a prediction that includes *any* object in the image is counted as a valid prediction. For object detection, users tested Google Cloud Vision API's Object Detection on two categories relevant for autonomous driving: *bicycle* and *stop sign* (average precision 0.7-0.8 on OpenImages).[5] Users were instructed to only mark as failures tests where the model does not detect *any* bicycles or stop signs present. For image captioning, users tested Alibaba's official checkpoint of OFA-Huge finetuned on COCO Captions [43], which is state-of-the-art on the benchmark, and were asked to explore scenes a visually impaired user might encounter when inside a *kitchen* or an *elementary school*. Users were instructed to consider as failures only object and action recognition errors which would egregiously mislead a visually impaired user.

Participants and setup. We recruited 40 participants from academia and industry (with IRB approval) who had taken at least a graduate-level course in machine learning, computer vision, or natural language processing. We assigned 16 users to the classification task, 16 to the detection task, and 8 users to the image captioning task.

In these studies, we also aimed to ablate the importance of ADAVISION's adaptivity over its benefits as an interactive search interface with model scoring. To do so, we asked each user to complete two rounds of testing. In the ADAVISION round, users had full access to ADAVISION as described in Section 3, while in NONADAPTIVE round we disabled topic suggestions, automatic test labeling, and adaptive test suggestions (i.e. suggestions are always retrievals based solely on the topic name). Users had a limited amount of time for each round (15 to 20 minutes), and were instructed to try to find as many failure-prone topics (bugs) for a specific category as possible, switching topics whenever they found 8-10 failures within a topic (more details in Appendix B). Users tested different categories between
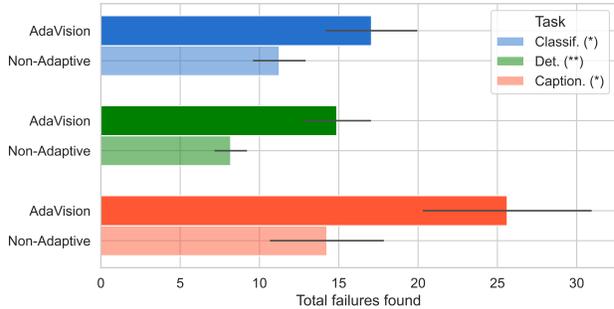
Figure 5: User study results comparing ADAVISION to NON-ADAPTIVE (baseline). Error bars are standard errors over users. Results significant with $p < 0.05$ (*) or $p < 0.005$ (**), with more details in Appendix B.

rounds (to minimize learning between rounds), and category assignments and round orderings were randomized.

Results. We present the number of failures found (averaged across users) in Figure 5. Users were able to find a variety of bugs with ADAVISION, even in strong models with strict definitions of failure. Further, ADAVISION's adaptivity helped users find close to twice as many failing tests than NONADAPTIVE, with moderate to large (standardized) effect sizes in classification ($d = 0.588, p < 0.05$), object detection ($d = 0.882, p < 0.005$), and image captioning ($d = 0.967, p < 0.05$). Using ADAVISION helped users identify *more diverse* bugs than the baseline: while 12/40 users found 2 or more bugs with ADAVISION, only 1/40 could match this level of diversity in the baseline round.

Qualitatively, users found bugs related to spurious correlations, difficult examples, and missing world knowledge (we share samples in Figure 6). For example, users discovered that ViT-H/14 strongly correlates kitchen countertops with the label *microwave* and witch hats with the label *cauldron*, leading to failure on images where these correlations do not hold (*e.g.* microwaves are absent). Users found that Cloud Vision misses detections when stop signs and bicycles are partially obscured by snow, and users also discovered object and action recognition errors in OFA-Huge, such as with *oven mitts* and musical instruments held near the mouth.

We surveyed users on whether ADAVISION was instrumental in finding these bugs. 84.6% of users marked said they "could not have found these bugs using existing error analysis tools [they] have access to." We also asked users to rate the cognitive difficulty of finding bugs in each round, on a scale of 1 to 5. The average perceived difficulty with ADAVISION was
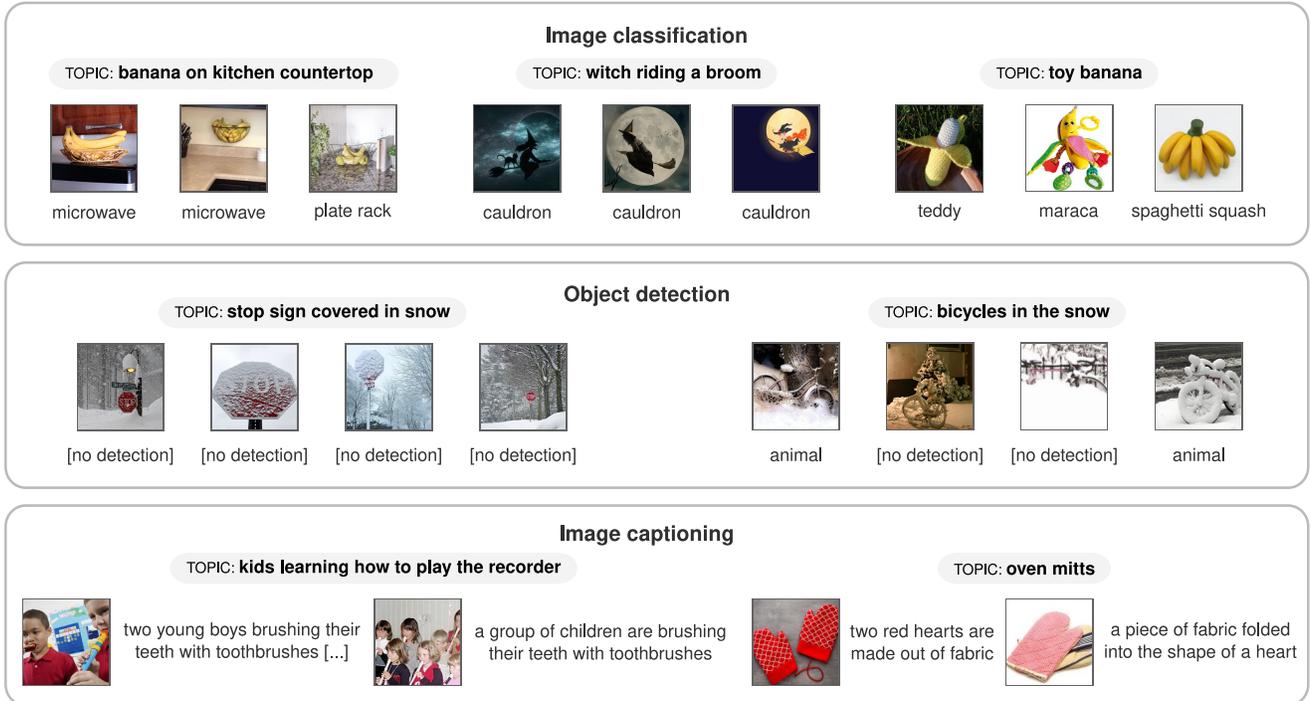
Image classification

TOPIC: banana on kitchen countertop

microwave   microwave   plate rack

TOPIC: witch riding a broom

cauldron   cauldron   cauldron

TOPIC: toy banana

teddy   maraca   spaghetti squash

Object detection

TOPIC: stop sign covered in snow

[no detection] [no detection] [no detection] [no detection]

TOPIC: bicycles in the snow

animal [no detection] [no detection] animal

Image captioning

TOPIC: kids learning how to play the recorder

two young boys brushing their teeth with toothbrushes [...]

a group of children are brushing their teeth with toothbrushes

TOPIC: oven mitts

two red hearts are made out of fabric

a piece of fabric folded into the shape of a heart

Figure 6: Sample of bugs found by users in studies. In each case, the model prediction (shown to the right of the corresponding input image) is incorrect. These bugs span spurious correlations (e.g. ViT-H/14 associating a kitchen counter top with a microwave), difficult examples (e.g. Cloud Vision API failing to detect stop signs partially obscured by snow), and missing world knowledge (e.g. OFA-Huge misidentifying oven mitts).

3.05 ± 1.07, in contrast with 4.10 ± 0.91 for NonAdaptive. In a paired t-test, this gap was significant with $p < 0.001$, reflecting that users felt testing was easier with AdaVision than without.

4.3. Comparison with automatic slice discovery

Domino [12] is a state-of-the-art slice discovery method that clusters validation set errors and describes them with automatically generated captions. We compare these to AdaVision topics on unseen data, noting that if a topic or caption t genuinely describes a bug, drawing new samples from $x \sim P(X|t)$ should yield a high failure rate.

Setup. We compare bugs found in ImageNet classification models with respect to six categories: {banana, broom, candle, lemon, sandal, wine bottle}. Specifically, tests failed if they were false negatives: for class y, the model fails if the image x contains object y, but the model instead predicts an object that is not in the image. We target the ViT-H/14 model from Section 4.2 [10], and a ResNet-50 [14].

For each category, Domino clusters ImageNet validation examples using an error-aware Gaussian mix-

ture model in CLIP's latent space, and then describes each cluster with a caption.[6] We use two variants, Domino (BERT) and Domino (OFA), which differ in how they caption clusters (template filling with BERT [9] or captioning with OFA [42]). Appendix C contains more details.

We used AdaVision topics from a user session in the user studies for overlapping categories, and had an author run additional sessions (i.e. use AdaVision for 20 minutes) for the 2 remaining categories. While Domino targets each model individually, we only target ViT-H/14, and directly transfer the discovered topics to ResNet-50.

Both methods propose five topics per category (we took the top-5 with most failures for AdaVision), for a total of 30 topics each (listed in Appendix C). To evaluate the failure rate of a topic on *new* data that matches the topic description, we retrieve nearest neighbors from LAION-5B using the topic name, and label the first 50 in-topic retrievals (skipping over images that do not fit the topic description). We exclude tests AdaVision users encountered during testing to

---

[6]We use the official implementation [12], available at https://github.com/HazyResearch/domino

avoid counting tests already reviewed in ADAVISION's favor.

Results. We present average failure rates across coherent topics proposed by each method in Table 1. We also present two baselines: the failure rate of a generic topic description per category (`a photo of {y}`), and the failure rate on the original ImageNet validation data (noting that ImageNet has a looser definition of failure, enforcing prediction on an arbitrary object on images with multiple objects). ADAVISION topics yield much higher failure rates, while DOMINO shows rates close to baselines. Interestingly, ADAVISION topics generated while testing ViT-H/14 transfer well to ResNet-50.

Further, while all ADAVISION topics are unsurprisingly coherent (as they are human-verified), we find that 61.6% and 33.3% of topics from DOMINO (BERT) and DOMINO (OFA) respectively are nonsensical (*e.g.* `a photo of setup by banana`, `a photo of skiing at sandal`) or fail to refer to the target category at all (*e.g.* `three oranges and an apple on a white background` when the target is "lemon"). These topics are excluded from Table 1 and listed in Appendix C.

We believe these results illustrate some inherent difficulties of automatic slice discovery methods. Validation error clusters may not be semantically tied together, especially when models saturate in-distribution benchmarks (ViT-H/14 and ResNet-50 are stronger than models used in prior evaluation of automatic slice discovery [12, 17]). Current slice captioning methods may also over-index into incorrect details or miss broader patterns between images. Because of the human-in-the-loop, ADAVISION enables users to form more coherent hypotheses about model failures. Further, cluster captions can describe a group of failures without including the *cause* of the model failure [18] (e.g. "a woman sitting on a chair holding a broom" is coherent, but ViT-H/14 has a failure rate of only 10% on additional data matching this description). In contrast, users of ADAVISION iteratively form hypotheses about model vulnerabilities: after selecting a topic, users observe model behavior on additional data from the topic, leading them to *refine* the topic definition. This iterative process helps users identify topics which consistently capture model failures.

4.4. Fixing bugs via finetuning

We evaluate whether users can fix bugs discovered with ADAVISION, by finetuning on failed tests. We finetune ViT-H/14 on the 30 ADAVISION topics for categories {banana, broom, candle, lemon, sandal, wine bottle} from Section 4.3, taking a sample of 20 tests per topic for a total of 600 images. As a baseline that just trains on images from a different distribution,

Model	Method	Avg failure rate
ViT-H/14	`a photo of {y}`	1.33
	ImageNet	11.47
	DOMINO (BERT)	8.6
	DOMINO (OFA)	7.33
	ADAVISION	**28.47**
ResNet50	`a photo of {y}`	15.7
	ImageNet	23.67
	DOMINO (BERT)	20.44
	DOMINO (OFA)	25.45
	ADAVISION	**56.93**

Table 1: Average failure rates across topics proposed by ADAVISION and two variants of DOMINO, an automatic slice discovery method. ADAVISION finds bugs that are 3x more difficult for models, while automatic methods propose groups that are close to baseline failure rates.

we finetune ViT-H/14 on an equal-size set of images retrieved from LAION-5B using generic topics in the form `an image of {y}`. We evaluate the finetuned models on held-out examples from ADAVISION topics, on the original domain (ImageNet [6]), and on five out-of-distribution datasets: ImageNet V2 [33], ImageNet-A [16], ImageNet-Sketch [41], ImageNet-R [15], and ObjectNet [2]. Additional details are in Appendix D.

Finetuning improves performance on discovered bugs. We use 50 held-out examples drawn from each topic we attempted to fix (*treatment topics*) to verify if the bugs were patched. As shown in Table 2, finetuning substantially increases performance on the held-out data from the treatment topics (by 18.6 percentage points), making the accuracy on treatment topics surpass average accuracy on ImageNet. These performance gains also hold on in-topic images sourced from *outside* of the LAION-5B distribution: for each treatment topic, we collected 50 images from a Google Image Search, deduplicated against both the finetuning set and the LAION-5B evaluation set. Finetuning on ADAVISION images from a testing session using LAION-5B also improves performance on in-topic images from Google by 13.9 percentage points.

Finetuning maintains in-distribution accuracy and improves OOD accuracy. To ensure performance gains are not due to the introduction of new shortcuts, we check performance on the original in-distribution data (ImageNet). Finetuning on the treatment topics does not reduce overall ImageNet accuracy (Table 2). Additionally, finetuning with ADAVISION improves overall average performance on the treatment

	AdaVision Topics			ImageNet	Avg across OOD Eval Sets	
Model	**Treatment Topics**		**Control Topics**	**Overall**	**Treatment Classes**	**Overall**
	LAION-5B	*Google*				
Before finetuning	72.6	76.7	91.3	88.4	78.0	77.7
Finetuning with `an image of {y}`	82.5 (0.9)	82.9 (0.6)	90.8 (0.3)	**88.5 (0.0)**	82.1 (0.6)	78.0 (0.1)
Finetuning with AdaVision tests	**91.2 (0.5)**	**90.6 (0.6)**	**91.9 (0.2)**	88.4 (0.0)	**84.0 (0.2)**	**78.2 (0.0)**

Table 2: Accuracies on AdaVision topics, ImageNet [6], and five OOD ImageNet evaluation sets [2, 15, 16, 33, 41] before and after finetuning on images accumulated from testing with AdaVision. Compared to a baseline of finetuning on the same number of images pulled generically using the topic `an image of {y}` from LAION-5B, AdaVision improves accuracy on held-out data from topics in the finetuning set (left two columns), regardless of whether images are sourced from LAION-5B or Google Images. AdaVision also improves accuracy on OOD evaluation sets (right two columns). Finetuning maintains overall performance on ImageNet (center) and held-out control topics (third column).

classes across out-of-distribution evaluation sets (from 78% to 84%). To check for shortcuts at a more fine-grained level, we evaluate performance on 19 semantically contrasting *control* topics with different labels. For example, one treatment topic involved images of lemons next to tea, which the model often predicted as *consomme*. We added `consomme` as a control topic to check that the model does improve performance on lemons at the expense of the "consomme" concept. Similarly, for topics `banana on kitchen countertop` and `witch riding a broom` in Figure 6, we add the control topics `microwave in kitchen` and `witch with cauldron`; see Appendix D for a list. Performance on the control topics does not decrease after finetuning. Our results indicate users can fix *specific* bugs discovered with AdaVision in ViT-H/14 without degrading performance elsewhere.

5. Limitations

Retrieval limitations. While LAION-5B has good coverage for everyday scenes, it may not be appropriate to specialized domains such as biomedical or satellite imagery, and CLIP's representation power is also likely to deteriorate on these domains [29]. Even for everyday scenes, the quality of CLIP's text-based retrieval degrades as the complexity of the topic name increases, particularly when several asymmetric relations are introduced [25, 48]. Further work to improve image-text models like CLIP could reduce off-topic retrievals during testing, improving users' testing speed.

Experiment limitations. We show that finetuning a state-of-the-art classification model on AdaVision bugs fixes them without degrading performance elsewhere. While it is particularly encouraging that we could improve performance on labels that had very high accuracy to begin with, this experiment was done with a limited set of labels and with only one round of testing. Multiple rounds of testing / finetuning could

be more beneficial [34]. Models smaller than ViT-H/14 may also be more prone to catastrophic forgetting [30]. To preserve in-distribution performance when fixing bugs for these models, robust finetuning techniques like weight averaging [46] or adding in-distribution data to the finetuning set may be necessary.

Fixing non-classification bugs. Classification tests contain the expected label implicitly in pass/fail annotations, and thus are easy to turn into finetuning data. However, the same is not true for detection or captioning tests, since we do not collect correct bounding boxes or captions during testing (only pass/fail annotations). Fixing such bugs would require an additional step of labeling failing tests prior to finetuning, or using loss functions that explicitly allow for negative labels [21].

6. Conclusion

We presented AdaVision, a human-in-the-loop process for testing vision models that mimics the life-cycle of traditional software development [34]. By leveraging human feedback for models on vision tasks, AdaVision helps users to identify and improve coherent vulnerabilities in models, beyond what is currently captured in-distribution evaluation sets. Our experiments indicate the adaptive nature of AdaVision improves the discovery of bugs, and that finetuning on bugs discovered with AdaVision boosts performance on the discovered failure modes. AdaVision is open-sourced at https://github.com/i-gao/adavision.

Acknowledgements

References

[1] Guha Balakrishnan, Yuanjun Xiong, Wei Xia, and Pietro Perona. Towards causal benchmarking of biasin face analysis algorithms. In *Deep Learning-Based Face Analytics*, pages 327–359. Springer, 2021. 3

[2] Andrei Barbu, David Mayo, Julian Alverio, William Luo, Christopher Wang, Dan Gutfreund, Josh Tenenbaum, and Boris Katz. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 3, 8, 9

[3] Shaily Bhatt, Rahul Jain, Sandipan Dandapat, and Sunayana Sitaram. A case study of efficacy and challenges in practical human-in-loop evaluation of nlp systems using checklist. In *Proceedings of the Workshop on Human Evaluation of NLP Systems (HumEval)*, pages 120–130, 2021. 3

[4] Daniel Bogdoll, Stefani Guneshka, and J Marius Zöllner. One ontology to rule them all: Corner case scenarios for autonomous driving. *arXiv preprint arXiv:2209.00342*, 2022. 3

[5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. 2

[6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. https://ieeexplore.ieee.org/abstract/document/5206848. 2, 8, 9

[7] Emily Denton, Ben Hutchinson, Margaret Mitchell, Timnit Gebru, and Andrew Zaldivar. Image counterfactual sensitivity analysis for detecting unintended bias. *arXiv preprint arXiv:1906.06439*, 2019. 3

[8] Greg d'Eon, Jason d'Eon, James R Wright, and Kevin Leyton-Brown. The spotlight: A general method for discovering systematic errors in deep learning models. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1962–1981, 2022. 1, 2

[9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. 7

[10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 7

[11] Xin Du, Benedicte Legastelois, Bhargavi Ganesh, Ajitha Rajan, Hana Chockler, Vaishak Belle, Stuart Anderson, and Subramanian Ramamoorthy. Vision checklist: Towards testable error analysis of image models to help system designers interrogate model capabilities. *arXiv preprint arXiv:2201.11674*, 2022. 3

[12] Sabri Eyuboglu, Maya Varma, Khaled Saab, Jean-Benoit Delbrouck, Christopher Lee-Messer, Jared Dunnmon, James Zou, and Christopher Ré. Domino: Discovering systematic errors with cross-modal embeddings. *arXiv preprint arXiv:2203.14960*, 2022. 1, 2, 3, 7, 8, 19

[13] Deep Ganguli, Liane Lovitt, Jackson Kernion, Amanda Askell, Yuntao Bai, Saurav Kadavath, Ben Mann, Ethan Perez, Nicholas Schiefer, Kamal Ndousse, et al. Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*, 2022. 1, 3

[14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7

[15] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. *ICCV*, 2021. 3, 8, 9

[16] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. *CVPR*, 2021. 3, 8, 9

[17] Saachi Jain, Hannah Lawrence, Ankur Moitra, and Aleksander Madry. Distilling model failures as directions in latent space. *arXiv preprint arXiv:2206.14754*, 2022. 1, 2, 8

[18] Nari Johnson, Ángel Alexander Cabrera, Gregory Plumb, and Ameet Talwalkar. Where does my model underperform? a human evaluation of slice discovery algorithms. *arXiv preprint arXiv:2306.08167*, 2023. 2, 8

[19] Saeed Khorram and Li Fuxin. Cycle-consistent counterfactuals by latent transformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10203–10212, 2022. 3

[20] Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vid-

gen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, et al. Dynabench: Rethinking benchmarking in nlp. *arXiv preprint arXiv:2104.14337*, 2021. 3

[21] Youngdong Kim, Junho Yim, Juseung Yun, and Junmo Kim. Nlnl: Negative learning for noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 101–110, 2019. 9

[22] Guillaume Leclerc, Hadi Salman, Andrew Ilyas, Sai Vemprala, Logan Engstrom, Vibhav Vineet, Kai Xiao, Pengchuan Zhang, Shibani Santurkar, Greg Yang, et al. 3db: A framework for debugging computer vision models. *arXiv preprint arXiv:2106.03805*, 2021. 3

[23] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019. https://openreview.net/forum?id=Bkg6RiCqY7. 20

[24] Kede Ma, Zhengfang Duanmu, Zhou Wang, Qingbo Wu, Wentao Liu, Hongwei Yong, Hongliang Li, and Lei Zhang. Group maximum differentiation competition: Model comparison with few samples. *IEEE Transactions on pattern analysis and machine intelligence*, 42(4):851–864, 2018. 3

[25] Zixian Ma, Jerry Hong, Mustafa Omer Gul, Mona Gandhi, Irena Gao, and Ranjay Krishna. Crepe: Can vision-language foundation models reason compositionally? *arXiv preprint arXiv:2212.07796*, 2022. 9

[26] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 220–229, 2019. 1

[27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. https://arxiv.org/abs/1912.01703. 20

[28] Kayur Patel, James Fogarty, James A Landay, and Beverly Harrison. Investigating statistical machine learning as a tool for software development. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 667–676. ACM, 2008. 1

[29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 1, 3, 9

[30] Vinay Venkatesh Ramasesh, Aitor Lewkowycz, and Ethan Dyer. Effect of scale on catastrophic forgetting in neural networks. In *International Conference on Learning Representations*, 2021. 9

[31] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 4

[32] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pages 5389–5400, 2019. 1, 15

[33] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pages 5389–5400. PMLR, 2019. 3, 8, 9

[34] Marco Tulio Ribeiro and Scott Lundberg. Adaptive testing and debugging of nlp models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3253–3267, 2022. 1, 2, 3, 4, 5, 9

[35] Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. Beyond accuracy: Behavioral testing of nlp models with checklist. In *Association for Computational Linguistics (ACL)*, 2020. 1, 3

[36] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. 3

[37] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402*, 2022. 3

[38] Nimit Sohoni, Jared Dunnmon, Geoffrey Angus, Albert Gu, and Christopher Ré. No subclass left behind: Fine-grained robustness in coarse-grained classification problems. *Advances in Neural Information Processing Systems*, 33:19339–19352, 2020. 1, 2

[39] Yuchi Tian, Kexin Pei, Suman Jana, and Baishakhi Ray. Deeptest: Automated testing of deep-neural-network-driven autonomous cars. In *Proceedings of the 40th international conference on software engineering*, pages 303–314, 2018. 1, 3

[40] Haotao Wang, Tianlong Chen, Zhangyang Wang, and Kede Ma. I am going mad: Maximum discrepancy competition for comparing classifiers adaptively. *arXiv preprint arXiv:2002.10648*, 2020. 3

[41] Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems*, pages 10506–10518, 2019. 3, 8, 9

[42] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren

Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. *CoRR*, abs/2202.03052, 2022. 7

[43] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. *arXiv preprint arXiv:2202.03052*, 2022. 6

[44] Zhou Wang and Eero P Simoncelli. Maximum differentiation (mad) competition: A methodology for comparing computational models of perceptual quantities. *Journal of Vision*, 8(12):8–8, 2008. 3

[45] Olivia Wiles, Isabela Albuquerque, and Sven Gowal. Discovering bugs in vision models using off-the-shelf image generation and captioning. *arXiv preprint arXiv:2208.08831*, 2022. 1, 2, 3

[46] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7959–7971, 2022. 9

[47] Jiebin Yan, Yu Zhong, Yuming Fang, Zhangyang Wang, and Kede Ma. Exposing semantic segmentation failures via maximum discrepancy competition. *International Journal of Computer Vision*, 129:1768–1786, 2021. 3

[48] Mert Yuksekgonul, Federico Bianchi, Pratyusha Kalluri, Dan Jurafsky, and James Zou. When and why vision-language models behave like bag-of-words models, and what to do about it? *arXiv preprint arXiv:2210.01936*, 2022. 9

[49] Oliver Zendel, Katrin Honauer, Markus Murschitz, Daniel Steininger, and Gustavo Fernandez Dominguez. Wilddash-creating hazard-aware benchmarks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 402–416, 2018. 1

[50] Mengshi Zhang, Yuqun Zhang, Lingming Zhang, Cong Liu, and Sarfraz Khurshid. Deeproad: Gan-based metamorphic testing and input validation framework for autonomous driving systems. In *2018 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 132–142. IEEE, 2018. 3

[51] Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR, 2021. 14

Appendix Overview

- In Appendix A, we present additional details on the ADAVISION system as described in Section 3, including details about adaptive test retrieval and automatic test labeling (A.1), as well as adaptive topic generation based on templates and user topics (A.2).

- In Appendix B, we include additional details about the user studies in Section 4.2, as well as details about the statistical analyses.

- In Appendix C, we expand upon the comparison of ADAVISION with DOMINO (Section 4.3 in the main text), providing an explanation of DOMINO and its hyperparameters, our criterion for coherency, a list of all slice descriptions (topics) from DOMINO used during evaluation, and a list of the 30 user-found ADAVISION topics we compared DOMINO against.

- Finally, in Appendix D, we discuss hyperparameters used for finetuning in Section 4.4, list all control and treatment topics, and include an additional evaluation measuring whether finetuning on treatment topics improves other, conceptually unrelated bugs.

A. Additional details on ADAVISION

In Section 3, we described ADAVISION, which includes a test generation loop that retrieves images with CLIP and an topic generation loop that generates topics with GPT-3. Here, we expand on the details of both loops.

A.1. Test generation loop

In the test generation loop, users explore a candidate topic t. Each iteration of the loop, ADAVISION retrieves test suggestions relevant to the topic, and, when possible, automatically imputes pass/fail labels for these suggestions in order to minimize user labeling effort. In this section, we provide additional details about the retrieval and labeling steps of a single iteration of the test generation loop.

Recall that m is our target vision model (*e.g.* a classification model), and $m(x)$ is the model output on an image x (*e.g.* the label string, such as "banana"). At any given iteration of the test generation loop, we have a textual topic description z and a (possibly empty) set of already labeled tests $\mathcal{D}$, where each test is a triplet $(x, m(x), y)$. The label $y \in \{-1, 1\}$ refers to whether the test failed (-1) or passed (1). We may also have a set of previously reviewed, off-topic tests $\mathcal{D}_{\text{off-topic}}$. In

Algorithm 1, we step through the retrieval and labeling steps of the iteration. For space, we expand on two of the steps in the algorithm box below.

Sampling previous tests x_1, x_2, x_3 for retrieval. In the retrieval step, we incorporate three previous tests from $\mathcal{D}$. Each of the three tests is sampled from a categorical distribution over $\mathcal{D}$, where each $x_j \in \mathcal{D}$ has probability p_j of being selected. We prefer to select tests where the model *confidently* fails over tests where the model less confidently fails, which are still preferred over passed tests. Thus, when available, we use the model's *prediction confidence* $s_m(x_j) \in [0, 1]$ for each example $x_j \in \mathcal{D}$ to compute p_j. For example, in classification, s_m is the model's maximum softmax score. If this value is unavailable (*e.g.* for some commercial APIs), we fix $s_m(x_j) = 1$ for all examples. To compute p_j, we compute a score $\alpha_j = 1 - y_j s_m(x_j)$ per example (note that this is 0 when the test is confidently correct and 2 when the model is confidently incorrect), and we set p_j to be the normalized version of α_j.

Lightweight classifiers $f, f_{\text{off-topic}}$ for automatically labeling tests. When $|\mathcal{D}| > 0$, we use two lightweight classifiers to automatically impute whether tests have passed, failed, or are off-topic. The lightweight classifiers are specialized to the current topic (*i.e.* we train new classifiers for new topics). Functionally, a lightweight classifier is a Support Vector Classifiers (SVC). One lightweight classifier f maps $(x, m(x))$ to predicted pass/fail label in $\{-1, 1\}$. Specifically, the tuple $(x, m(x))$ is transformed into a single vector by first embedding both x (image) and $m(x)$ (string) with CLIP ViT-L/14, followed by concatenating the two into $[\text{CLIP}(x), \text{CLIP}(m(x))]$. This vector is used as the feature representation of the SVC. New tests are re-sorted according to the prediction of f and f's confidence, with likely failures shown first. Classifier $f_{\text{off-topic}}$ operates on the same representation as f for each test, but instead predicts binary labels for in-topic / off-topic. These classifiers take less than a second to train and run, so we re-train them at each iteration of the test generation loop.

A.2. Topic generation loop

In the topic generation loop, users collaborate with ADAVISION to generate candidate topics to explore using GPT-3. We generate candidate topics in two phases. In the first phase, we prompt GPT-3 using a pre-written set of templates and collect the completions. We share the set of prompt templates used in Listing 1, replacing {LABEL} with predefined label names or existing user topics (*e.g.* `stop sign`). In

Algorithm 1: Iteration of the test generation loop.
***
Input: Textual topic description z, previously labeled tests $\mathcal{D} = \{(x, m(x), y)\}$, previous off-topic tests $\mathcal{D}_{\text{off-topic}}$

Compute $q_t \leftarrow \text{CLIP}(z)$ ▷ Figure 2A
if $|\mathcal{D}| > 0$ **then**
    Sample $x_1, x_2, x_3 \sim \text{Categorical}(|\mathcal{D}|, p_j)$, where p_j is computed according to the text in A.1
    Aggregate $q_i \leftarrow \sum_k \beta_k \cdot \text{CLIP}(x_k)$, with $\beta \sim \text{Dirichlet}(1, 1, 1)$
    Set $q \leftarrow \text{slerp}(q_t, q_i, r)$, with $r \sim \text{Uni}(0, 1)$
else
    Set $q \leftarrow q_t$
end

Retrieve approximate nearest neighbors of q from LAION-5B   ▷ Figure 2B
Exclude retrievals whose CLIP image embeddings have cosine similarity > 0.9 with any previous test $x \in \mathcal{D}$
Collect model outputs for all retrieved images to obtain new collection of tests $\mathcal{S} \leftarrow [(\tilde{x}, m(\tilde{x}))]$

if $|\mathcal{D}| > 0$ **then**
                                                           ▷ Figure 2C
    Train a lightweight classifier f on previously labeled tests $\mathcal{D}$ as described in A.1
    Sort $\mathcal{S}$ according to $f(\tilde{x})$ for $\tilde{x} \in \mathcal{S}$, placing predicted fails far from the decision boundary first, and
      predicted passes far from the decision boundary last Update $\mathcal{S}$ to contain $(\tilde{x}, m(\tilde{x}), f(x))$, so that we
      can display the imputed label to the user
    Train a second lightweight classifier $f_{\text{off-topic}}$ to differentiate between previous in-topic tests $\mathcal{D}$ and
      previous off-topic tests $\mathcal{D}_{\text{off-topic}}$
    Place tests $\tilde{x} \in \mathcal{S}$ for which $f_{\text{off-topic}}(x)$ predicts "off-topic" at the end of $\mathcal{S}$
end
return sorted $\mathcal{S}$ to the user for confirmation / correction.   ▷ Figure 2D
***

the second phase, we gather suggestions from the first round, append previously explored topics with high failure rates, and gather a second round of suggestions. We place topics with the highest failure rates at the end of the prompts to account for GPT-3's recency bias [51]. Finally, topics are presented to users, who explore ones they deem interesting and important.

A.3. Web interface

We provide screenshots of the ADAVISION web interface in Figure 7. The topic generation loop is represented as a root page that suggests topics to explore, and individual topics are represented as folders (Figure 7 left). Tests within folders are represented as rows mapping images to model outputs (Figure 7 right).

B. Additional details for user studies

In Section 4.2, we described a large set of user studies used to evaluate ADAVISION's ability to enable users to find bugs in state-of-the-art vision models. Here, we include additional details about the study setups and statistical analyses.

Listing 1: Example prompts used in topic generation loop.

```
List some unexpected places to see a {LABEL}
List some places to find a {LABEL}
List some other things that you usually find
    with a {LABEL}
List some artistic representations of a {LABEL}
List some things that can be made to look like
    a {LABEL}
List some types of {LABEL} you wouldn't normally
    see
List some dramatic conditions to photograph
    a {LABEL}
List some conditions a {LABEL} could be in that
    would make it hard to see
List some things that are the same shape as a
    {LABEL}
List some {LABEL} that are a different color than
    you would expect
```

Study setup. All participants undertook the study virtually in a single 60-minute Zoom session. At the start of the session, participants were shown a 5-minute video introducing how to use the ADAVISION web interface. Next, the experimenter walked through instructions for the testing task: as described in the
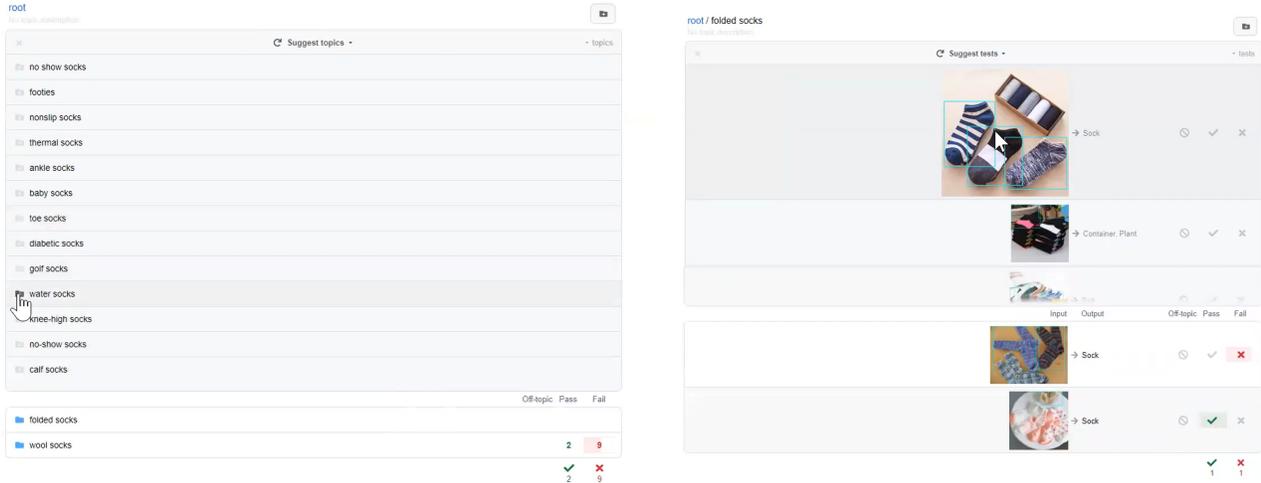
Figure 7: In the ADAVISION web interface, topics are represented as folders, and tests are represented as rows mapping images to model outputs. The topic generation loop is represented as a root page that suggests topics to explore (left), while the test generation loop within each topic suggests tests that lightweight classifiers label as potential failures (right, top panel).

main text, participants were instructed to find as many failure-prone topics (bugs) for a specific category (*e.g.* banana) as possible, and to switch topics whenever they found more than a threshold of failures within a topic. This latter instruction prevented users from endlessly exploiting a topic to inflate the total failure count. We adjusted the threshold at which a topic became a bug (and users should move on) based on the time users had for testing: for classification and image captioning, users tested models for 20 minutes with a bug threshold of 10 failures, while for object detection, users tested the model for 15 minutes with a bug threshold of 8.

When introducing the task, the experimenter defined failed tests as follows:

- For participants testing classification models, a test failed if the model predicted an object not present in the image. Users were instructed to look for failures among pictures of a specific category (*banana* or *broom*), *e.g.* failures among pictures of bananas (Figure 8). Participants were given class definitions from an ImageNet labeling guide used in [32].

- For those testing object detection models, a test failed if the model failed to box any instance of the given category (*bicycle* or *stop sign*), *e.g.* as shown in Figure 9.

- Participants testing image captioning models were asked to imagine that they were testing a product

used by visually impaired customers to caption everyday scenes. The participants' task was to find images (tests) for which the model produced false or incorrect captions, excluding counting, color, and gender or age mistakes (Figure 10). Participants looked for such tests among pictures of a specific category (*i.e.* scenes customers might encounter in a *kitchen* or *elementary school*).

The experimenter then asked each user to practice using the web interface by testing the model a third, held-out object or location for 10 minutes. For classification, this was a wine bottle; for object detection, this was a fire hydrant; and for image captioning, this was a *garden*. After two rounds of testing in the main experiment, the study concluded with an exit survey as described in Table 3. The study compensation was a $25 Amazon gift card.

Additional results. As discussed in Section 4.2, ADAVISION helped users find significantly more failing tests than NONADAPTIVE, with significance determined by paired t-tests in each task. In classification, $t(16) = 2.27, p < 0.05$; in object detection, $t(16) = 3.42, p < 0.005$, and in image captioning, $t(8) = 2.56, p < 0.05$. These corresponded to normalized effect sizes of $d = 0.588$ in classification, $d = 0.882$ in object detection, and $d = 0.967$ in image captioning. We also counted the number of users who could find bugs during testing, *i.e.* identify a topic that hit the threshold number of fails (8 for object detection, 10 for the other tasks). Overall, 28/40 users found
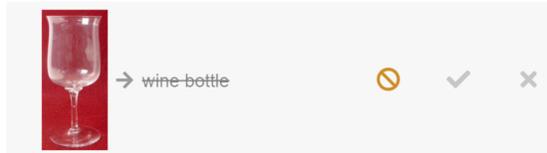
Classification Instructions

To help narrow down your search, we're only going to consider certain kinds of images: images which **contain an object Y.**

✅ A test **passes** if the predicted object *is actually in the image*.

❌ The test **fails** if the predicted object *is not in the image*.

Examples

Let **object Y= "wine bottle"**.

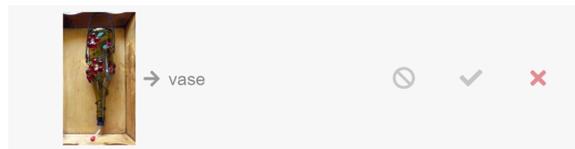First, we'll check whether we should mark the test as **off-topic.**

The image does not contain a wine bottle, so it is off-topic.

After checking if the test is off-topic, let's review when to mark a test as **pass/fail.**

The image contains a wine bottle. The test **passes** because a wine bottle is in the image.

The image contains a wine bottle (upside down). The test **fails** because a vase is not in the image.

The image contains wine bottles (to the left and right of the wine glasses). The test **passes** because a china cabinet is in the image.

The image contains a wine bottle. The test **fails** because eggnog is not in the image.

Should I be exploring many topics, or staying within one topic?

Your task is to **maximize the number of topics** (folders) that *each* have **10+ failures.**

- The round ends when the time is up.
- A good strategy is to explore many topics. (Don't stay in a topic longer than 10 fails!)

Figure 8: Example instructions for classification users.
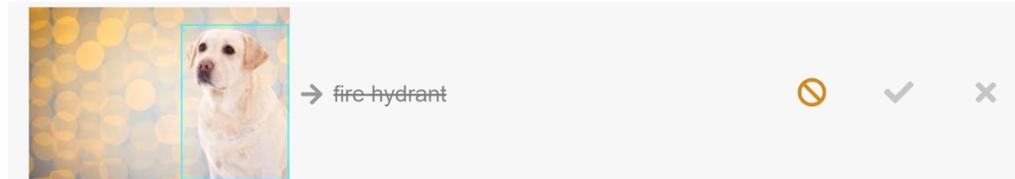
Detection Instructions

For the sake of the study, we're only interested in a certain kind of image:
- The image must contain one or more instances of **object Y**
- An image is a **fail** if the model does not locate **ANY** occurrences of **object Y**
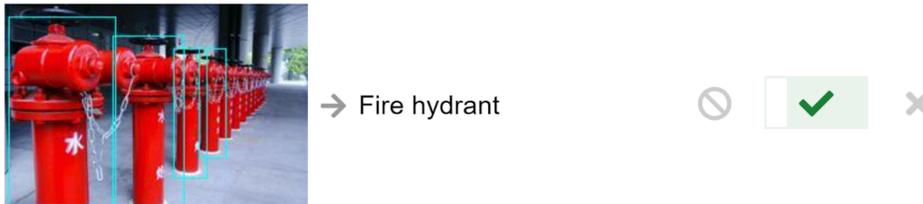
Examples

Let **object Y= "fire hydrant"**.

First, we'll check whether the image is relevant to the study.

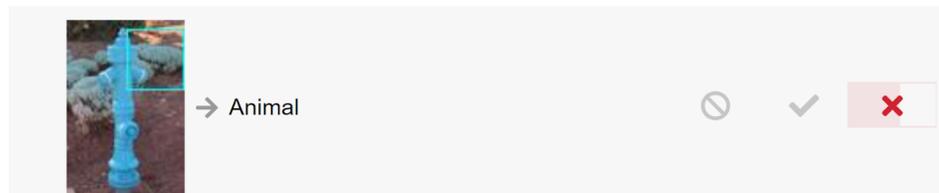The image does not contain a fire hydrant, so it is off-topic.

Next, we'll decide whether to mark a test as **pass/fail.**

The image contains a fire hydrant. The test **passes** because the model boxes and labels at least one of the fire hydrants. It doesn't matter that it doesn't box all of the hydrants.

The image contains a fire hydrant. Although the model boxes, the hydrant, it doesn't correctly label it, so the test **fails.**

The image contains a fire hydrant. The test **fails** because the fire hydrant is not boxed and labeled.

Should I be exploring many topics, or staying within one topic?

Your task is to **maximize the number of topics** (folders) with **8+ failures.**
- The round ends when the time is up.
- A good strategy is to explore many topics. (Don't stay in a topic longer than 8 fails!)

Figure 9: Example instructions for object detection users.

Captioning Instructions

Given an image, captioning models **describe the image with text.**

❌ The test **fails** if the caption gives a misleadingly *false* or *incorrect* description of the image.

You're looking for **fails** among photographs of everyday objects that a user might encounter in **location X.**

Practice: Labeling Captions as Pass/Fail

Let **location X = "garden"**. Let's review when to mark a test as **pass** / **fail**.

→ a bunch of tomatoes growing on a plant   ⊘ ✔ ✕

→ a woman in an apron holding a walking stick   ⊘ ✔ ✖

→ a man cutting a tree with a chainsaw   ⊘ ✔ ✖

IMPORTANT: we're NOT counting certain mistakes as errors.

● Number -- if there are 2 of an object but the tool says 3, that's okay.

→ a bird bath with a white flower in it   ⊘ ✔ ✕

● Gender and age

→ stock photo portrait of a little boy holding a stick   ⊘ ✔ ✕

● Color and material

→ a bunch of red and green tomatoes on a branch   ⊘ ✔ ✕

Should I be exploring many topics, or staying within one topic?

Your task is to **maximize the number of topics** (folders) that *each* have **10+ failures.**
● The round ends when the time is up.
● A good strategy is to explore many topics. (Don't stay in a topic longer than 10 fails!)

Figure 10: Example instructions for image captioning users.

Question	Type
How difficult was it to find bugs in the first round?	5-point Likert scale
How difficult was it to find bugs in the second round?	5-point Likert scale
How useful was the web tool for finding bugs?	Multiple-choice {I could have found these bugs using existing error analysis tools I have access to, I could not have found these bugs using existing error analysis tools I have access to.}
Did you use topic suggestions? Were they helpful?	Multiple-choice {Yes, and they were helpful in generating topics that caused failures, Yes, and they were helpful in generating ideas for topics to explore, Yes, but they did not generate good ideas for topics to explore, No, I did not use topic suggestions.}

Table 3: Items in exit survey.

bugs during testing with AdaVision, while only 16/40 could find such a high-failure topics in the baseline round. When surveyed about perceived difficulty of finding bugs, users also felt finding bugs was easier with AdaVision than without, with $t(40) = 4.18, p < 0.0005$. When surveyed about the helpfulness of GPT-3's topic suggestions, 24/34 users who used the topic suggestions marked that they were helpful for exploration.

C. Additional details for comparison with automatic slice discovery

In Section 4.3, we compared AdaVision with Domino [12], showing that bugs found with AdaVision are more difficult. Here, we provide an explanation of the Domino method and its hyperparameters, define our criterion for coherency, list *all* slice descriptions (topics) from Domino used during evaluation (along with whether they satisfied coherency),

and list the 30 user-found AdaVision topics.

Details on Domino. We use the official release of Domino [12], available at https://github.com/HazyResearch/domino. To generate slice proposals for a label in {banana, broom, candle, lemon, sandal, wine bottle}, we ran the target model (ViT-H/14 or ResNet-50) over ImageNet validation examples of that class. Then, we used Domino's error-aware mixture model to generate 5 slices (per class). The error-aware mixture model first generates $\bar{k}$ candidate slices (clusters of images) before selecting the best 5 slices. As in the original paper, we set $\bar{k} = 25$, and we initialize groups using the confusion matrix setting. Additionally, Domino uses a hyperparameter $\gamma = 10$ to control the weight placed on *incorrect* examples when slicing. In the original paper, the authors used $\gamma = 10$ for all datasets, which we also initially tried. However, we found that this setting produced slices that were too easy (no errors within the slices). Thus, we matched the authors' blog post applying Domino to ImageNet[7], conducting final experiments with $\gamma = 40$. (Larger values of γ resulted in too much cluster incoherency.)

In our experiments, we evaluated two variations of Domino, which use the same image clusters but differ in their captioning strategy. The first, Domino (BERT), is the original proposal in [12]; to caption a cluster, Domino (BERT) samples 250,000 BERT or Wikipedia completions, per cluster, of a set of templated captions; we provide the templates we used in Listing 2. Each cluster is then matched to the caption with the highest cosine similarity in CLIP space. We note that in the original paper, the authors used only one template: "a photo of a [MASK]"; our modifications to this template account for the fact that all clusters are about a certain class (*e.g.* clusters of banana images). Thus, we enforce that the label (banana) appears in the caption by populating {LABEL} in the template with the appropriate class name. We also compared against our own variation of Domino, Domino (OFA), which uses Alibaba's OFA-huge to more coherently caption clusters. Here, we run OFA-Huge over all examples in a cluster and select the individual caption that maximizes cosine similarity with the cluster mean.

Coherency. When reporting failure rates for Domino topics, we calculated failure rates only over descriptions that were coherent. Descriptions were deemed incoherent if they were nonsensical (e.g. "a photo of setup by banana", "a photo of skiing at

---

Listing 2: Templates used to generate captions for DOMINO (BERT).

```
a photo of {LABEL} and [MASK]
a photo of {LABEL} in [MASK]
a photo of [MASK] {LABEL}
[MASK] {LABEL} [MASK]
```

sandal") or did not refer to the target category at all (e.g. "three oranges and an apple on a white background" or "a photo of promoter david lemon" when the target is "lemon"). For reference, we include the full list of DOMINO (BERT) descriptions for ViT-H/14 in Listing 6, ResNet-50 in Listing 7, and the list of DOMINO (OFA) descriptions for ViT-H/14 in Listing 8 and ResNet-50 in Listing 9. In these lists, we prepend an asterisk in front of coherent topics, and we append the target category for each slice in parentheses. Of the starred coherent descriptions, we excluded three from evaluation because we could not find any related images in LAION-5B: "a photo of munitions and broom" (DOMINO (BERT), ViT-H/14), "a photo of primate and broom" (DOMINO (BERT), ResNet-50), and "a large banana sitting in the middle of an abandoned building" (DOMINO (OFA), both ViT-H/14 and ResNet-50).

ADAVISION **topics.** We also include the ADAVISION topics we compare to in Listing 3. Of the six categories {banana, broom, candle, lemon, sandal, wine bottle}, we recruited one user per category to test ViT-H/14 and generate topics, except for the categories *candle* and *wine bottle*, which one of the authors tested in a separate session. All users (including the author) were limited to 20 minutes for testing, and they had not explored ViT-H/14 on the tested class before.

D. Additional details for finetuning

In Section 4.4, we presented results from experiments that show we can patch model performance on bugs while maintaining or slightly improving accuracy on the original ImageNet distribution, control topics, and OOD evaluation sets. In these experiments, we finetuned on 20 images from each of 30 buggy topics (which we call the *treatment topics*). These treatment topics are the same as in Section 4.3 / Appendix C, listed in Listing 3. In this section, we discuss hyperparameter choices, list all control topics, and break down OOD evaluation set gains by dataset.

Finetuning hyperparameters. We finetune ViT-H/14 using a small, constant learning rate of 1e-5

Listing 3: Topics from ADAVISION studied in Sections 4.3 and 4.4 in the main text.

```
banana next to a banana smoothie
banana on kitchen countertop
banana in wooden woven basket
banana next to banana bread
toy banana
broom by fireplace
witch flying on a broom
photo of a person holding a boxy broom
silhouette of a person flying on a broom
broom in closet
black-and-white clipart of a candle
creamy white candle in glass jar
christmas candle next to tea
candle by window in snowstorm
person holding a candle at a vigil
grating a lemon
cooking with lemon
lemon tea with lemon
lemon on pancake with condensed milk or honey
lemon clipart
a lot of toy sandals
flip flop door wreath
translucent sandals
colorful flip flops
sandal ornament in a tree
top of a champagne bottle
wine bottle in a wiry wine rack
champagne in a champagne holder
wine bottle in a suit case
wine bottle with a wine stopper
```

with the AdamW optimizer [23, 27] for five steps, with weight decay 0.01, batch size 16, and random square cropping for data augmentation. These hyperparameters were chosen in early experiments because they only degrade in-distribution model performance slightly. We report all results averaged over 3 random seeds along with the corresponding standard deviations. When deduplicating evaluation data against the finetuning data, we mark pairs as duplicates if their CLIP cosine similarly > 0.95.

Control topics. We provide a list of the 19 control topics from Section 4.4 in Listing 4. These topics were selected because they were semantically related to the treatment topics in Listing 3, but had different labels. For example, the ADAVISION topic `person holding a boxy broom` is visually similar to the concept of "person holding a mop", so we include the latter as a control topic. Other control topics are classes that were incorrectly predicted for a topic (*e.g.*, `banana on kitchen countertop` is frequently predicted "microwave", so we include "microwave in kitchen" as a control topic). We checked performance on these topics to make sure performance gains were

Listing 4: Control topics in Section 4.4 in the main text. In parentheses, we list the topic found by AdaVision with which the given control topic contrasts.

```
shopping basket (banana in wooden woven basket)
bread (banana next to banana bread)
dishwasher in kitchen (banana on kitchen
    countertop)
microwave in kitchen (banana on kitchen
    countertop)
eggnog (banana next to a banana smoothie,
    lemon on pancake with condensed milk or
    honey)
witch with cauldron (witch flying on a broom)
fireplace no broom (broom by fireplace)
mop (broom in closet,
    photo of a person holding a boxy broom)
person holding mop (photo of a person holding
    a boxy broom)
consomme (lemon tea with lemon)
black tea no lemon (lemon tea with lemon)
grated orange (grating a lemon)
pancake with condensed milk or honey no lemon
    (lemon on pancake with condensed milk or
    honey)
torch (person holding a candle at a vigil,
    candle by window in snowstorm)
clog shoe (sandal)
beer bottle (top of a champagne bottle)
suit case (wine bottle in a suit case)
empty wine rack (wine bottle in a wiry wine
    rack)
empty wire rack (wine bottle in a wiry wine
    rack)
```

Listing 5: Topics from AdaVision studied in Sections 4.3 and 4.4 in the main text.

```
bananagrams
banana in fruit salad
curling broom on the ice
broomball with brooms
candle in mason jar with flower
hexagon candle
a lemon with a smiley face drawn on it
lemon on waffle
lace sandal
crocs
rows of wine bottles in a store
wine bottle in a wine fridge
```

not due to the model forming new shortcuts.

Per-OOD evaluation set breakdown. In Tables 4 and 5, we provide a breakdown of the OOD evaluation set performances; these were aggregated as an average in Table 2 of the main text. Table 4 displays the accuracy on treatment classes in each of the OOD evaluation sets, and Table 5 displays the overall accuracy in each of the OOD sets.

Effect on conceptually unrelated bugs. We evaluated whether finetuning on treatment topics affected conceptually unrelated bugs. For each class in {banana, broom, candle, lemon, sandal, wine bottle}, we found two additional topics with high failure rates (Listing 5), disjoint from the treatment topics in Listing 3. We then measured performance on these unrelated topics before and after finetuning, and we compare to performance changes on the treatment and control topics in Table 6. We see that finetuning on treatment topic improves performance on semantically unrelated bugs within the same set of classes, but gains are smaller than on treatment topics.

Model	ImageNet	ImageNet V2	ImageNet-Sketch	ImageNet-R	ImageNet-A	ObjectNet	Avg. OOD
Before finetuning	87.7	65.0	86.1	89.2	65.6	84.3	78.0
Finetuning with baseline	**93.1 (0.2)**	**71.7 (1.4)**	**90.9 (0.3)**	90.5 (0.3)	71.2 (1.3)	86.4 (0.1)	82.1 (0.6)
Finetuning with AdaVision	92.9 (0.4)	69.4 (0.8)	**91.7 (0.5)**	**93.9 (0.2)**	**76.8 (1.5)**	**87.9 (0.2)**	**84.0 (0.2)**

Table 4: Accuracies on treatment classes, before and after finetuning. Results are averaged over three random seeds.

Model	ImageNet	ImageNet V2	ImageNet-Sketch	ImageNet-R	ImageNet-A	ObjectNet	Avg. OOD
Before finetuning	88.4	81.0	64.4	89.1	83.9	69.9	77.7
Finetuning with baseline	**88.5 (0.0)**	**81.3 (0.0)**	64.5 (0.0)	89.2 (0.1)	84.4 (0.1)	**70.5 (0.2)**	78.0 (0.1)
Finetuning with AdaVision	88.4 (0.0)	80.9 (0.1)	**64.7 (0.0)**	**90.0 (0.0)**	**84.9 (0.0)**	**70.5 (0.0)**	**78.2 (0.0)**

Table 5: Accuracies on all classes, before and after finetuning. Results are averaged over three random seeds.

Model	Treatment Topics	Control Topics	Unrelated Topics
Before finetuning	72.6	91.3	61.0
Finetuning with baseline	82.5 (0.9)	90.8 (0.3)	65.6 (0.8)
Finetuning with AdaVision	**91.2 (0.5)**	**91.2 (0.2)**	**74.7 (2.0)**

Table 6: Accuracies on treatment, control, and unrelated topics. Finetuning on treatment topic improves performance on semantically unrelated bugs within the same set of classes, but gains are smaller than on treatment topics.

Listing 6: Slice descriptions generated by Domino (BERT) for target model ViT-H/14. Asterisks in front of coherent topics.

```
a photo of setup by banana (banana)
*a photo of wine and banana (banana)
*a photo of ceramics and banana (banana)
*a photo of basket and banana (banana)
*a photo of munitions and broom (broom)
a photo of activist david broom (broom)
a photo of synthesizer. broom (broom)
*a photo of wildlife at broom (broom)
a photo of violinist jenny broom (broom)
*a photo of literature and candle (candle)
a photo of panchayats candle (candle)
*a photo of altarpiece and candle (candle)
a photo of rob and candle (candle)
a photo of corella lemon (lemon)
*a photo of blender lemon (lemon)
a photo of promoter david lemon (lemon)
a photo of clown billy lemon (lemon)
a photo of estadio jose lemon (lemon)
a photo of rowing on sandal (sandal)
a photo of nana and sandal (sandal)
a photo of placental sandal (sandal)
a photo of skiing at sandal (sandal)
a photo of screenwriter michael sandal (sandal)
a photo of shelter and wine bottle (wine bottle)
*a photo of champange wine bottle (wine bottle)
*a photo of bakery and wine bottle (wine bottle)
*a photo of advertisement on wine bottle (wine bottle)
*a photo of grocery stores wine bottle (wine bottle)
```

Listing 7: Slice descriptions generated by DOMINO (BERT) for target model ResNet-50. Asterisks in front of coherent topics.

```
*a photo of ceramics and banana (banana)
a photo of reception by banana (banana)
*a photo of orange and banana (banana)
a photo of neutron star banana (banana)
a photo of architect paul banana (banana)
a photo of singer jenny broom (broom)
a photo of rowing on broom (broom)
*a photo of ornate old broom (broom)
*a photo of factory of broom (broom)
*a photo of primate and broom (broom)
*a photo of blowing the candle (candle)
a photo of consultant john candle (candle)
*a photo of colorful birds candle (candle)
a photo of swamy candle (candle)
*a photo of candle in entryway (candle)
red lemonade series. (lemon)
liz lemon and the observer (lemon)
keith lemon and david bowie (lemon)
liz lemon and the batman (lemon)
a photo of lemon bay shuttle (lemon)
a photo of cognitive development sandal (sandal)
a photo of drilling the sandal (sandal)
a photo of lecture at sandal (sandal)
*a photo of frozen black sandal (sandal)
a photo of longevity by sandal (sandal)
*a photo of golden wine bottle (wine bottle)
a photo of autopsy wine bottle (wine bottle)
*a photo of home and wine bottle (wine bottle)
a photo of libertarian wine bottle (wine bottle)
a photo of estadio wine bottle (wine bottle)
```

Listing 8: Slice descriptions generated by DOMINO (OFA) for target model ViT-H/14. Asterisks in front of coherent topics.

```
a plate on a table with knives and forks (banana)
*a banana next to a bottle of wine and a glass (banana)
*a kitchen counter with bananas and a pineapple on it (banana)
*a banana and two pears in a red basket (banana)
*a large banana sitting in the middle of an abandoned building (banana)
three mops and a bucket against a brick wall (broom)
a man holding a baseball bat in a room (broom)
a woman in a blue dress is looking at a computer (broom)
a green praying mantis standing on a piece of wood (broom)
*a woman sitting on a chair holding a broom (broom)
*a carved pumpkin with a candle in the middle (candle)
*a candle sitting on the ground on a brick floor (candle)
*a group of lit candles in front of a stained glass window (candle)
*a man sitting at a table with candles (candle)
a white bird perched on a tree branch eating (lemon)
*a pitcher pouring lemonade into a glass with lemons (lemon)
three oranges and an apple on a white background (lemon)
a display of tomatoes and other vegetables (lemon)
a bunch of oranges sitting on top of a table (lemon)
a pair of shoes sitting on top of a skateboard (sandal)
a woman holding a small child on her lap (sandal)
*a pink crocheted sandal with a flower on it (sandal)
*a person is wearing a black sandal on their foot (sandal)
a woman laying on a bed with a laptop (sandal)
a group of small bottles of liquor on a table (wine bottle)
*a bottle of champagne in a bowl on a table (wine bottle)
*a bottle of wine and a paper on a counter (wine bottle)
*a glass of wine and a bottle on a table (wine bottle)
*a bottle of wine and a cigar on a table (wine bottle)
```

Listing 9: Slice descriptions generated by DOMINO (OFA) for target model ResNet-50. Asterisks in front of coherent topics.

```
*a green bowl with some bananas and a piece of fruit (banana)
a plate on a table with knives and forks (banana)
*a bowl of oranges and bananas on a table (banana)
*a banana and two pears in a red basket (banana)
*a large banana sitting in the middle of an abandoned building (banana)
*a broom sitting on the floor in front of a wooden door (broom)
a group of people holding a large wooden stick (broom)
*a close up of a broom with a wooden handle (broom)
three mops and a bucket against a brick wall (broom)
*a broom hanging on the side of a porch (broom)
*a baby girl sitting in front of a birthday cake with a candle (candle)
*a little girl is holding a candle and looking up (candle)
*a group of lit candles in front of a stained glass window (candle)
*a candle sitting on the ground on a brick floor (candle)
*a candle sitting on top of a wooden table (candle)
a group of sliced oranges and kiwi fruit (lemon)
*a pitcher pouring juice into a glass with lemons (lemon)
three oranges and an apple on a white background (lemon)
*a lemon with a smiley face drawn on it (lemon)
*a bowl filled with oranges and a lemon (lemon)
a pair of snoopy shoes and a box on a green table (sandal)
*a woman is wearing a pair of sandals on her feet (sandal)
*a woman wearing sandals standing on a concrete floor (sandal)
a pair of shoes sitting on top of a magazine (sandal)
*two pictures of a woman wearing a pair of sandals (sandal)
*a vase of roses and two bottles of wine (wine bottle)
*a cake with a bottle of wine in a box (wine bottle)
*a bottle of wine and grapes on a counter with a glass (wine bottle)
*a woman next to a row of wine bottles (wine bottle)
*a bottle of wine and a cigar on a table (wine bottle)
```