

# Where Does My Model Underperform? A Human Evaluation of Slice Discovery Algorithms

Nari Johnson<sup>1</sup>, Ángel Alexander Cabrera<sup>1</sup>, Gregory Plumb<sup>1,2\*</sup>, Ameeet Talwalkar<sup>1</sup>

<sup>1</sup>Carnegie Mellon University

<sup>2</sup>Inpleo, Inc.

narij@andrew.cmu.edu, cabrera@cmu.edu, gdplumb@alumni.cmu.edu, talwalkar@cmu.edu

## Abstract

Machine learning (ML) models that achieve high average accuracy can still underperform on semantically coherent subsets (“slices”) of data. This behavior can have significant societal consequences for the safety or bias of the model in deployment, but identifying these underperforming slices can be difficult in practice, especially in domains where practitioners lack access to group annotations to define coherent subsets of their data. Motivated by these challenges, ML researchers have developed new *slice discovery algorithms* that aim to group together coherent and high-error subsets of data. However, there has been little evaluation focused on whether these tools help humans form correct hypotheses about where (for which groups) their model underperforms. We conduct a controlled user study ( $N = 15$ ) where we show 40 slices output by two state-of-the-art slice discovery algorithms to users, and ask them to form hypotheses about an object detection model. Our results provide positive evidence that these tools provide some benefit over a naive baseline, and also shed light on challenges faced by users during the hypothesis formation step. We conclude by discussing design opportunities for ML and HCI researchers. Our findings point to the importance of centering users when creating and evaluating new tools for slice discovery.

## 1 Introduction

A growing number of works propose tools to help stakeholders form hypotheses about the behavior of machine learning (ML) models. One type of behavior that can have significant societal consequences occurs when a model underperforms on semantically coherent subsets (*i.e.*, “slices”) of data. For example, Buolamwini and Gebu (2018) found that leading tech companies’ commercial facial recognition models were significantly less accurate at classifying faces of women with darker skin. Knowledge of this model behavior informed advocacy efforts that led to fundamental changes to dataset curation for facial recognition (Birhane 2022) and public deliberation surrounding governance of facial recognition systems (Raji and Buolamwini 2022). More broadly, knowledge of underperforming slices can inform model selection and deployment (Balayn et al. 2023) or actions that

can be taken to fix the model (Holstein et al. 2019; Cabrera et al. 2021; Idrissi et al. 2022).

However, identifying these underperforming slices can be difficult in practice. In many domains, practitioners often do not have access to group annotations that can be used to define semantically *coherent* (*i.e.*, united by a single human-understandable concept) subsets of their data (Cabrera et al. 2022). Motivated by these challenges, ML researchers have developed new automated tools in the growing field of *slice discovery*. At a high level, these *slice discovery algorithms* are unsupervised methods that aim to group together coherent and high-error slices of data (Sohoni et al. 2020; d’Eon et al. 2021; Singla et al. 2021; Eyuboglu et al. 2022; Plumb et al. 2023; Wang et al. 2023). These works propose that a human stakeholder can inspect the slices output by these algorithms to form hypotheses of model behavior, *i.e.*, describe in words a group where the model underperforms.

While researchers continue to develop new slice discovery algorithms, there has been little evaluation of whether these algorithms help stakeholders achieve their proposed goals. Past human evaluations of slice discovery tools have used subjective judgments of the output slices’ coherence (such as whether users can find a description that matches the majority of images in the slice) as proxies of these algorithms’ utility (Singla et al. 2021; d’Eon et al. 2021), but stop short of verifying whether these descriptions are useful or even accurate depictions of the model’s behavior. In this work, we ask: Do the slices output by these algorithms help users form *correct hypotheses of model behavior*?

As a motivating example, consider a scenario where a practitioner wishes to evaluate a new object detection model designed to be deployed in an autonomous vehicle. The practitioner could run a slice discovery algorithm on a dataset of dash-cam photos collected from field testing, but its output (groups of photos) is not immediately actionable. Describing in words a group where the model underperforms (*e.g.*, “the model fails to identify stop signs in snowy environments”) is a prerequisite step for several downstream actions that the practitioner’s company could take, such as initiating targeted data collection efforts (*e.g.*, additional field testing to gather more data in snowy regions) or selective deployment (*e.g.*, delaying roll-out of the new model in areas that experience snow).

Unfortunately, taking action to address an *incorrect* hy-

\*The author completed the majority of their work while a student at CMU.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: An overview of our user study in three steps. (Left) We run different slice discovery algorithms to compute high-error slices (subsets of an input dataset). (Middle) We conduct a human subject study where we *show the slices from Step 1* to a human subject, who forms a hypothesis (*i.e.*, description of a subgroup where the model underperforms) corresponding to each slice. (Right) We validate each user-generated hypothesis from Step 2 by calculating the model’s accuracy on a new sample of images that match the user’s description.

pothesis can have several undesirable consequences. For example, if the model actually performs just as well in snowy environments, then the company’s efforts to improve the model’s performance on that group may have been better expended elsewhere, or may even worsen the model’s performance on other groups (Li et al. 2022). Thus, we argue that this under-examined step of hypothesis formation is critically important for many stakeholders.

In this work, we design a controlled user study to investigate how participants make sense of the slices output by algorithms to form hypotheses of model behavior. Our study has three parts, illustrated in Figure 1. First, after training a model on a large object detection dataset (Lin et al. 2014), we run two state-of-the-art slice discovery algorithms, DOMINO (Eyuboglu et al. 2022) and PLANESPOT (Plumb et al. 2023) that output high-error slices of data. Next, we show these slices to  $N = 15$  study participants and ask them to *form behavioral hypotheses* (*i.e.*, describe in words a group where the model underperforms) corresponding to each slice. Finally, we then *validate each hypothesis* (*i.e.*, whether the model actually underperforms on these groups) by measuring the model’s performance on new data that matches each hypothesis. In summary, our work makes the following contributions:

- *Benchmarking existing slice discovery tools.* Our study results provide positive evidence that existing tools can sometimes (though not always) help stakeholders form correct hypotheses about where their model underperforms. Specifically, participants were more likely to form correct hypotheses when shown slices output by DOMINO and PLANESPOT, relative to a naive baseline condition where they were shown a random sample of misclassified images.
- *Characterizing how users may (mis)interpret existing tools.* Our analyses shed light onto how the slices output by these tools may be misleading. First, we found *no significant association* between the number of images in a slice that the user selected as “matching” their hypothesis, and its correctness. This result challenges conventional wisdom that slices that are coherent (*i.e.*, the majority of images in the slice match a common descrip-

tion) do in fact correspond to a true model error (*i.e.*, the model underperforms on all such images). Second, we found that different users often formed different hypotheses when shown the same slice. Taken together, these findings illuminate the nuance and challenges faced by users during the hypothesis formation step.

- *Design opportunities for future tools.* Our findings point to several exciting design opportunities for ML and HCI researchers. We highlight open challenges and possible paths forward to better support users as they make sense of where their model underperforms.

The full version of this paper, including appendix, is available online at <https://arxiv.org/abs/2306.08167>

## 2 Slice Discovery Preliminaries

In this section, we present an overview of the slice discovery problem, formalize the hypothesis formation step, and discuss our approach to validate each hypothesis.

### 2.1 Slice Discovery Algorithms

Many tools have been proposed to help users discover where their model underperforms, including post-hoc explanations (Kim et al. 2018; Adebayo et al. 2022), methods that generate counterfactual data (Wiles, Albuquerque, and Goyal 2023), and data visualization interfaces (Cabrera et al. 2023; Suresh et al. 2023; Moore, Liao, and Subramonyam 2023). In our work, we focus on *slice discovery algorithms*: automated methods that aim to partition the data into coherent and high-error subsets (Eyuboglu et al. 2022; Plumb et al. 2023). We exclude methods that rely on additional information such as fine-grained group annotations (Polyzotis et al. 2019; Liu et al. 2021) or human supervision to define coherent subsets. Consequently, the methods we study do *not* require the practitioner to anticipate the types of inputs where the model may underperform in advance.

**Definition** We follow Plumb et al. (2023) and define a slice discovery algorithm as a method that given as input a trained model  $f$  and dataset of labeled images  $D =$

$\{(x_i, y_i)\}_{i=1}^n$ <sup>1</sup>, outputs  $k$  slices  $[\Psi_j]_{j=1}^k$ . Each slice  $\Psi_j \subseteq D$  is a subset of the input data. The objective of slice discovery is for each slice  $\Psi_j$  to correspond to a single coherent group where the model indeed underperforms (*i.e.*,  $f$  has lower accuracy for images in  $\Psi_j$  than for all images in  $D$ ). Past works propose that a user can inspect the datapoints belonging to each slice to “*identify common attributes*” (Eyuboglu et al. 2022) to describe the underperforming groups.

**Methods** We evaluate two slice discovery algorithms: DOMINO (Eyuboglu et al. 2022) and PLANESpot (Plumb et al. 2023). Both algorithms work by running error-aware clustering on an embedding of each image, and differ in the specific clustering algorithm and embedding that they use. We describe each algorithm in detail in Appendix A. We evaluate these two algorithms because they achieved state-of-the-art performance on past benchmarks (Eyuboglu et al. 2022; Plumb et al. 2023) at the time we ran our study.

Our BASELINE algorithm randomly samples from the set of misclassified images, without replacement. While a subset of random misclassified images is high-error, it is unlikely to be coherent. Our baseline was designed to simulate a naive workflow where a practitioner inspects an unsorted sample of misclassified images to form hypotheses.

## 2.2 Validating User Hypotheses

To evaluate each slice discovery algorithm, our study focuses on the human-centered task of *hypothesis formation*, where a human describes in words where they believe the model underperforms. In this section, we formalize this task and discuss our approach for validating users’ hypotheses.

We define a “*hypothesis*” as a *user-generated text description* of a group where they believe that the model underperforms. We follow past work and give as input to each slice discovery algorithm sets of datapoints  $D$  that all share the same true class label (Singla et al. 2021; Plumb et al. 2023). In this setting, “errors” correspond to images where the model failed to detect the object. Intuitively, a “correct hypothesis” should describe a group where the model has significantly lower accuracy at detecting the object. If  $\Phi$  denotes the set of all images in  $D$  that belong to the group, we define the *performance gap*  $Gap(f, \Phi)$  of  $f$  on  $\Phi$  as

$$\underbrace{\frac{1}{|D|} \sum_{(x,y) \in D} 1(f(x) = y)}_{\text{average accuracy on } D} - \underbrace{\frac{1}{|\Phi|} \sum_{(x,y) \in \Phi} 1(f(x) = y)}_{\text{average accuracy on } \Phi \subset D} \quad (1)$$

We define the correctness of the hypothesis “ $f$  underperforms on  $\Phi$ ” by thresholding the gap:

$$\text{“}f \text{ underperforms on } \Phi\text{”} \Leftrightarrow Gap(f, \Phi) > \tau \quad (2)$$

where threshold  $\tau \geq 0$  is a hyperparameter that defines the minimum performance gap necessary for a hypothesis to be “correct”. For example, if we set  $\tau = 0.2$ , then a “correct” hypothesis describes a group where the model has 20%

<sup>1</sup>We follow past work and give a separate held-out test set that the model  $f$  was *not* trained on as the input to a slice discovery algorithm.

worse accuracy on images that belong to the group, compared to all images that belong to its class.

One major challenge is that in many settings, we do not have access to  $\Phi$ , the complete set of images in  $D$  that match each hypothesis. Unfortunately, obtaining  $\Phi$  is often difficult and expensive. For example, recall the scenario where a practitioner hypothesizes that her model underperforms at detecting stop signs for “photos taken in snowy environments”. In this setting, the practitioner a-priori does not know which images in the class  $D$  match her hypothesis. She could manually review all photos of stop signs and annotate whether they were taken in snowy environments, but this strategy is slow and inefficient.

In our study, we aim to validate the correctness of 180 user-generated hypotheses. Unfortunately, manually reviewing each image within each class  $D$  to obtain the full set of images that match each hypothesis is difficult at scale. Thus, we decided to approximate the model’s performance gap for each hypothesis using a sample  $\hat{\Phi}$  of images that match each hypothesis, where  $\hat{\Phi} \subset \Phi$ . We provide an extended description of our approximation strategy in Appendix B, and summarize key steps below.

**Approximating the performance gap** Our goal is to find a sample of images  $\hat{\Phi}$  from the class  $D$  that match each hypothesis. To do this, we follow past work (Gao et al. 2022) and rank candidate images using their CLIP similarity score (Radford et al. 2021) to the hypothesis text description. When relevant, we performed minimal prompt engineering to users’ hypotheses to increase the quality of the similarity scores. Following Vendrow et al. (2023) and Gao et al. (2022), the final edited prompts began with the phrase, “*a photo of a [class] and [...]*”.<sup>2</sup>

Next, we inspected a sample of the images most similar to each hypothesis, and manually labeled up to the first 40 images that matched the hypothesis. In an effort to make our labeling process consistent and reproducible, we created a labeling guide inspired by Shankar et al. (2020) describing the criteria we used to determine whether an image matched each hypothesis. We provide an extended description of this labeling process in Appendix C.

For each class, we used the same CLIP retrieval strategy to obtain a sample  $\hat{D}$  of images belonging to the class. We define each  $\hat{D}$  as the 100 images with the highest similarity score to prompt “*a photo of [class]*”. Our final approximation of the performance gap calculates the difference in the model’s accuracy for the group’s sample  $\hat{\Phi}$  vs. the class’s sample  $\hat{D}$ . We publicly release our labeling guide and the files we selected as matching each group description (*i.e.*, the  $\hat{\Phi}$ ) on GitHub.<sup>3</sup> We discuss several ablations to our approximation strategy in Appendix B.

<sup>2</sup>Like past work, we use different prompt templates for hypotheses that do not fit the general template. For example, we modify the template for artistic styles (*i.e.*, “*a greyscale photo of an airplane*”).

<sup>3</sup><https://github.com/njohnson99/slice-discovery-human-eval>

### 3 Related Work

**Automated Evaluations of Slice Discovery Tools** Automated evaluations of slice discovery tools fall into three categories. The first category of evaluations measure the quality of each slice by calculating its size or error rate (d’Eon et al. 2021; Singla et al. 2021), but fail to capture any notion of the slices’ coherence. The second category of evaluations compare the output slices to known groups where the model underperforms (Sohoni et al. 2020; Eyuboglu et al. 2022; Plumb et al. 2023; Wang et al. 2023). One major limitation of this approach is that in practice, we often lack access to the full set of groups where a model underperforms. As such, these works evaluate methods’ ability to discover only the same subset of known errors, *i.e.*, groups that have already been annotated in well-studied benchmarks (Sohoni et al. 2020). The final category of evaluations uses another ML model (in place of a human) to generate a natural language description of each slice (Eyuboglu et al. 2022; Gao et al. 2022). However, recent work has shown that these ML models often generate nonsensical descriptions (Gao et al. 2022). Our study avoids limitations of past automated evaluations by instead running a human evaluation that asks people (rather than a separate ML model) to describe each slice. Our hypothesis validation approach also does *not* require access to the full set of true model errors.

**Human Evaluations of Slice Discovery Tools** Several works that introduce new slice discovery tools include qualitative evaluations performed by humans. Many such works do not conduct a controlled user study. Instead, the authors themselves describe the semantic features shared by the top 3, 5, or 10 images in each slice (d’Eon et al. 2021; Singla et al. 2021; Eyuboglu et al. 2022; Wang et al. 2023). To our knowledge, the majority of these evaluations do not validate that their descriptions accurately depict groups where the model does underperform. Some works use the authors’ ability to find a description that matches the majority of images in the slice as a proxy of the algorithms’ utility. For example, d’Eon et al. (2021) argue that their method is superior to a competitor because the latter’s output slices appear to “*share little in common*”.

One of the most similar studies to ours is Singla et al. (2021), where industry data scientists are asked to use slices output by an algorithm to develop hypotheses about where an image classifier fails. However, their study does not validate the correctness of participants’ hypotheses. Further, they only evaluate a single slice discovery algorithm.

In contrast to past work, our work does *not* assume that users’ hypotheses are de facto correct depictions of true errors. Instead, we validate each user-generated hypothesis by calculating the model’s accuracy on a new sample of data. Furthermore, we explicitly study whether output slices’ coherence is a valid proxy for tools’ utility.

**Behavioral Understanding of ML** A growing number of works propose tools to facilitate stakeholders’ understanding of model behavior. Specifically, our work contributes to the growing body of empirical studies that ask users to form hypotheses about where a model underperforms (Singla et al. 2021; Cabrera et al. 2022; Suresh et al. 2023;

Moore, Liao, and Subramonyam 2023). However, relatively few existing studies validate the correctness of users’ hypotheses (Wu et al. 2019; Cabrera et al. 2021; Gao et al. 2022). To our knowledge, ours is the first study to evaluate whether slice discovery tools help users form *correct* hypotheses of model behavior.

### 4 Experimental Design

We conduct a controlled user study motivated by three high-level goals, under which we organize our hypotheses:

1. Our first goal is to *benchmark* two state-of-the-art slice discovery tools (“conditions”), DOMINO and PLANESPOT, against a naive baseline. This comparison provides an important sanity check that these tools may offer users some benefit over a naive workflow of inspecting a random sample of errors.

We were specifically interested in comparing three *measures* associated with each hypothesis, across conditions: how well the user’s description of the slice (hypothesis) captures a group where the model actually does underperform (“correctness”), how difficult it was for the user to describe the slice, and how many images in the slice match the user’s hypothesis:

**H1.** A greater proportion of users’ hypotheses corresponding to slices output by slice discovery algorithms will be correct, when compared to hypotheses corresponding to slices output by the naive baseline.

**H2.** Users will rate slices output by slice discovery algorithms as easier to describe, when compared to slices output by the naive baseline.

**H3.** Users will select more images as “matching” their hypothesis for slices output by slice discovery algorithms, when compared to slices output by the naive baseline.

2. Our second goal is to characterize the relationship *between* these measures for the slices output by state-of-the-art algorithms (DOMINO and PLANESPOT).

As discussed in Section 3, past evaluations use measures of slice *coherence*, such as whether the majority of images share a common description, as a proxy of algorithms’ utility. However, this assumption overlooks that just because many images in the slice match some description, does not mean that the model underperforms on *all* such images. To examine this assumption, we explicitly study the relationship between the number of images in the slice that match the users’ hypothesis (out of 20 shown), and its correctness. We note that the number of matching images is just one possible heuristic to capture the coherence of the slice, and that there may be other valid ways to measure slice coherence.

**H4.** The number of images in the slice that match a user’s hypothesis (a measure of slice coherence) does not predict whether their hypothesis is correct.

3. Our final goal is to explore how humans make sense of the slices output by state-of-the-art algorithms.

One assumption shared by past work is that each slice corresponds to a *single unique group* where the model

underperforms. However, in pilot studies we observed that different users arrived at different conclusions about model behavior, even when shown the exact same information (details in Appendix D).

Our final hypothesis studies different users’ consistency with each other when shown the same slice. We ask annotators to label whether users’ hypotheses are synonymous (*i.e.*, may differ syntactically, but describe the same group of images), or distinct.

**H5.** Different users will write down distinct hypotheses when presented with the same slice.

## 4.1 Study Design

**Domain & Model** To study whether slices can help humans understand model behavior, we must first select a domain, prediction task, and model. We used data from MS-COCO (“COCO”) (Lin et al. 2014), a large object detection dataset, for its accessibility to a wide audience. COCO contains photos with 91 different object types “*that would be easily recognizable by a four-year-old*” (Lin et al. 2014) in everyday natural scenes. COCO is a multi-label classification task where a single image can have several objects present. We defined a custom 15%-10%-75% train-validation-test split so that we could use a larger held-out test set as input to the slice discovery algorithms. We finetuned a pretrained ResNet-18 model using the training set (details in Appendix E), performed model selection using the validation set, and ran the slice discovery algorithms on the held-out test set. Because we aim to evaluate how well slice discovery tools can detect naturally occurring errors, we did not modify the dataset or model training process to synthetically induce specific errors.

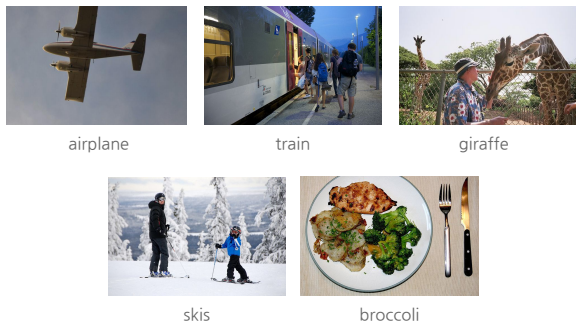


Figure 3: Example photos from the 5 selected COCO object classes (Lin et al. 2014).

**Computing Slices** We generated 60 total slices (20 slices per each of the 3 algorithms) to show users. Rather than run each slice discovery algorithm on the entire COCO test set (which contains 91 object types), we followed past evaluations (Gao et al. 2022; Plumb et al. 2023) and ran each algorithm on subsets of the data that all have the same object. Thus, each hypothesis describes a group where the model has low *recall* (*i.e.*, failed to detect the object). We used

each slice discovery algorithm to return the top  $k = 4$  slices for 5 different objects (Figure 4, Top), chosen randomly from a list of candidate objects where the model had at least 50% recall: airplane, train, giraffe, skis, and broccoli.

**Participants & Recruitment** We elicited 12 hypotheses each from 15 total subjects, for a total of 180 user-generated hypotheses. We recruited subjects that had self-reported “intermediate knowledge in machine learning (ML) or computer vision (CV)” *i.e.*, had taken a graduate-level course or had practical work experience in ML, AI, or CV, using university mailing lists. We recruited participants with ML expertise (rather than task-specific expertise) because most existing slice discovery tools were created to be used by model developers. All participants were students enrolled in a full-time degree program in computer science. Participants were compensated with a \$20 gift card, and reported spending 30 (min) to 55 (max) minutes participating.

## 4.2 Study Procedure

The study was approved by an Institutional Review Board (IRB) process and was conducted asynchronously online. Participation was voluntary and users were shown a consent form before participating. We began each study by presenting the user with a description of the study task and walk-through of the study interface. After completing the walk-through, the user was shown information about 12 different slices belonging to 3 different object classes. The user completed a questionnaire that asked them to formulate a behavioral hypothesis for each slice. We detail each phase of the study procedure below.

**Instructions** In the study walk-through, we introduced the user to the object detection task by showing them example images and model predictions from a randomly selected class (“tennis racket”). We defined and motivated the slice discovery task by listing several reasons why one may wish to discover groups where a model underperforms. We provide screenshots of the task instructions in Appendix F.

**Class Overview** Each user was asked to form hypotheses for slices corresponding to 3 different object classes. For each class, the user was first shown a *class overview* (Figure 2, Left) with information about and example images belonging to the class. We presented the class overview before presenting the slices to give the user basic context about the COCO dataset (*i.e.*, examples to illustrate the variety of images that belong to each class) that stakeholders in practice would already have for their domain.

**Slice Overview & Questionnaire** For each slice, the user was shown a *slice overview* (Figure 2, Right) that displays the model’s performance on the top 20 images that belong to the slice. The user was asked to use this information to complete the *slice questionnaire*, which asked them to (1) write down a behavioral hypothesis of the underperforming group corresponding to the slice, (2) select all images in the slice that belong to this group, and (3) rate how difficult it was to describe the slice. We provide the complete questionnaire in Appendix G.

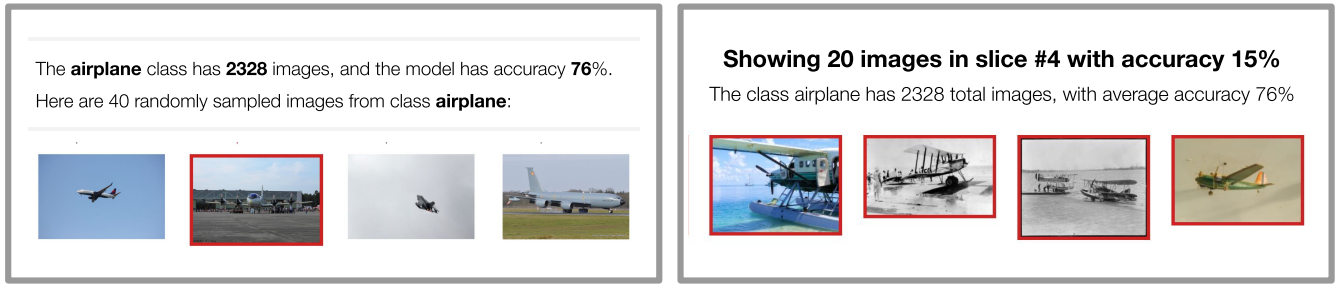


Figure 2: UI screenshots of the class overview (Left) and slice overview (Right). Errors (where the model failed to detect the object) have red borders. (Left, Top) The class overview shows the total number of images in the test set that belong to the class and the model’s average accuracy on these images. (Left, Bottom) A random sample of 40 images from the test set. (Right, Top) The slice overview shows model’s average accuracy on the top-20 images belonging to the slice and the entire test set. (Right, Bottom) The top-20 ordered images that belong to the slice.

*Visualizing each slice.* We controlled for how we visualized the model’s performance on each slice by presenting the top-20 images only for all conditions. We designed our slice overview to be as similar as possible to how past work presents individual slices (d’Eon et al. 2021; Eyuboglu et al. 2022). We limit the number of images shown to 20 because showing more images may cause information overload and increase the time required to complete each questionnaire.

*Eliciting a hypothesis for each slice.* Past works assume that *all images* in each computed slice correspond directly to a *single* group where the model underperforms (d’Eon et al. 2021; Eyuboglu et al. 2022; Plumb et al. 2023). We explicitly evaluate this assumption and ask users to form a single behavioral hypothesis for each slice. Specifically, we instruct users to *describe* the slice to the best of their ability by writing a group description that matches as many images in the slice as possible. Users are told that some slices may be noisy or incoherent, and that in some cases it may be difficult to find a single description that matches all of the images in the slice. We provide users with further guidance and example hypotheses detailed in Appendix F.

### 4.3 Experimental Design

We used a within-subjects design shown in Figure 4, where each participant was randomly assigned to 3 (out of 5 candidate) object classes. All participants were presented with 12 total slices, and saw slices output by all 3 algorithms. The 3 algorithm conditions were presented to participants in a random order to control for learning effects. We showed each of the 60 total slices to 3 different participants, collecting 3 hypotheses per each slice. Users were blinded to the study condition when completing each questionnaire: while they were informed that each slice was computed by an algorithm, they were not given any information about *which algorithm* was used to compute each slice.

### 4.4 Metrics & Analyses

To evaluate **H1** (whether the average number of correct hypotheses varies across conditions), we exclude hypotheses where we failed to find a sufficiently large sample of matching images to approximate the performance gap. Applying

this criteria, we retain 136 out of the original 180 hypotheses (76%) where we found a sample of at least 15 matching images. To evaluate **H2** and **H3**, we calculate each measure using all 180 of the original hypotheses.

To test for statistically significant differences between the three algorithm conditions, we ran ANOVA tests with Tukey post-hoc tests for multiple comparisons to compare the proportion of correct hypotheses (**H1**) and average number of matching images (**H3**). We ran Mann-Whitey tests with Bonferroni corrections to compare the Likert-scale self-reported difficulty of describing each slice (**H2**).

To evaluate **H4** and **H5**, we retained only the hypotheses that correspond to slices output by slice discovery algorithms, and excluded slices output by the baseline condition.

To determine if coherence (*i.e.*, the number of images in the slice that match a hypothesis) is an appropriate proxy of the hypothesis’s correctness (**H4**), we ran two Spearman’s rank correlation tests. For both tests, the independent variable is the number of matching images. We tried two dependent variables: the value of the approximate performance gap for the hypothesis (defined in Equation 1), and an indicator for hypothesis correctness using performance gap threshold  $\tau = 0.2$ .

To determine whether two different users’ hypotheses are equivalent or distinct (**H5**), we asked two annotators to label groups of hypotheses that are synonymous (*i.e.*, describe the same group of images). For each slice, we then use these groups to count the number of distinct user hypotheses (where each group of synonymous hypotheses only counts as a single “distinct hypothesis”). We discuss this process of determining whether two hypotheses are distinct or synonymous in detail in Appendix H.

## 5 Results

We used a  $p$ -value of 0.05 as our cutoff for significance. We present additional details for each statistical test (*e.g.*, test statistics) and additional analyses in Appendix I.

**Correctness** For each condition, we calculated the proportion of hypotheses that are “correct” using a performance

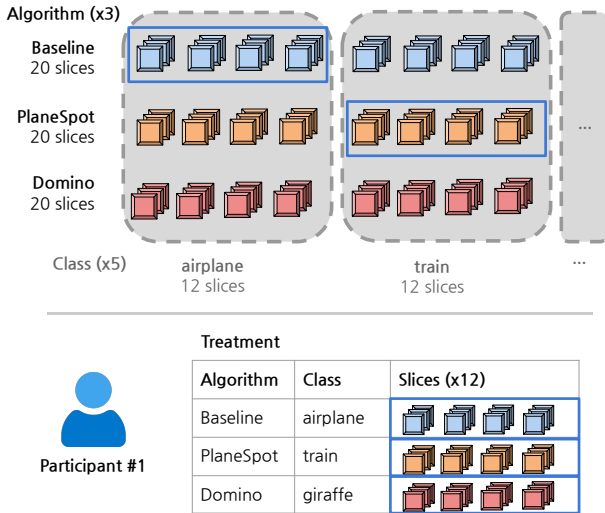


Figure 4: Experimental Setup. (Top) We collect users’ hypotheses for 60 total slices. For each of the 3 algorithms (rows), and for each of the 5 classes (columns), we compute the top-4 slices. We show 2 out of 5 classes (and 24 out of 60 total slices) in the figure due to space constraints. (Bottom) Each study participant was shown 12 total slices output by 3 different algorithms, and saw slices corresponding to a different class for each algorithm. For example, Participant #1 was asked to develop hypotheses for the top-4 slices output by PLANESPOT for the `train` class. The blue boxes on the top panel highlight the slices shown to Participant #1.

gap threshold of at least 20%. We found that 70 out of 136 total hypotheses (51%) are correct (standard error 4.3%). When we stratify by condition, 35% of BASELINE, 49% of PLANESPOT, and 73% of DOMINO hypotheses are correct, with standard errors 6.7%, 7.9%, and 6.8% respectively (Figure 5). We found a statistically significant difference between the BASELINE and DOMINO conditions only ( $p < 0.001$ ), partially supporting **H1**.

We present additional results where we ablate the performance gap threshold  $\tau$  in Appendix I.1. In summary, we found that a consistent trend (that DOMINO outperforms PLANESPOT, which outperforms BASELINE) holds for all variations tried. DOMINO consistently significantly outperformed the BASELINE condition; however, there was not a statistically significant difference between the other conditions for the majority of thresholds.

**Number of matching images** Overall, we observed that users selected more images as “matching” their hypothesis for slices output by slice discovery algorithms relative to the naive baseline, supporting **H2**. Figure 7 shows a histogram of the empirical distribution of matching images for each condition. On average, users selected 12.9 and 12.8 out of the 20 displayed images in the slice as “matching” their hypotheses for the DOMINO and PLANESPOT conditions respectively (standard errors 0.57 and 0.71), vs. 8.8 out of the

20 images (standard error 0.60) for the BASELINE hypotheses. We observed significant pairwise differences between the two slice discovery conditions vs. the baseline condition (with  $p < 0.0001$  for both DOMINO and PLANESPOT), and no significant difference between the two slice discovery conditions. However, we note that while the *average* number of matching images is higher for both slice discovery conditions, there is still high *variance* across hypotheses. For example, 30% of hypotheses from the DOMINO or PLANESPOT conditions had  $< 10$  matching images.

**Self-reported difficulty** We observed that users rated slices output by the slice discovery algorithms as *easier to describe* compared to slices output by the naive baseline condition, supporting **H3**. Figure 6 shows the distribution of Likert-scale ratings for each condition. We observed significant pairwise differences between the two slice discovery conditions vs. the baseline ( $p < 0.0001$ ), and no significant difference between the two slice discovery conditions.

**Does coherence imply correctness?** We found *no significant association* between the number of images that match each hypothesis and its correctness, supporting **H4**. For both dependent variables, the number of matching images is only weakly correlated with the hypothesis’s correctness, with correlation coefficients  $r_s = 0.08$  with  $p = 0.4712$  for the value of the performance gap, and  $r_s = 0.06$  with  $p = 0.5996$  for the correctness indicator. When we examined the hypotheses that had the highest number of matching images (at least 15 out of 20), only 60% of these hypotheses were correct, a rate comparable to the overall base rate of 61% accuracy for all hypotheses corresponding to slices output by either DOMINO or PLANESPOT.

**(In)consistency across users** Figure 8 shows the number of distinct hypotheses for each of the 40 slices output by either PLANESPOT or DOMINO, and Figure 9 shows users’ hypotheses for two example slices output by PLANESPOT. We found that all three users wrote down consistent (*i.e.*, synonymous) hypotheses for only a small minority (25%) of slices, supporting **H5**. In contrast, at least 1 of the 3 users wrote down a hypothesis that differed from other users for the majority (75%) of slices, and all three users disagreed for 30% of slices.

## 6 Discussion

We share several implications of our experimental results (Section 5.1) and present design opportunities for future tools to help users understand where their model underperforms (Section 5.2).

### 6.1 Interpreting the Experimental Results

We return to our three high-level study goals to place our results in conversation with past work on slice discovery.

**Benchmarking existing tools** Overall, we found significant differences across measures between the DOMINO and PLANESPOT conditions, relative to the naive baseline algorithm. Specifically, while DOMINO and PLANESPOT slices had a significantly higher average number of matching images and were significantly easier to describe relative to the




Figure 5: Hypothesis Correctness. The percentage of hypotheses per condition that are “correct” using a performance gap threshold  $\tau = 20\%$  with standard error bars. Percentages are calculated for the subset of hypotheses that we have a sufficiently large number of examples (*i.e.*, at least 15 matching images) to approximate the performance gap. We find that a greater proportion of users’ hypotheses from the PLANESPOT and DOMINO conditions are correct relative to the BASELINE condition.




Figure 6: Self-Reported Difficulty. A diverging stacked bar chart centered around the natural response of users’ self-reported difficulty of describing each slice (on a five-point Likert Scale), stratified by each condition (slicing algorithm). We find that users are significantly more likely to rate slices output by the PLANESPOT and DOMINO conditions as being easy to describe.

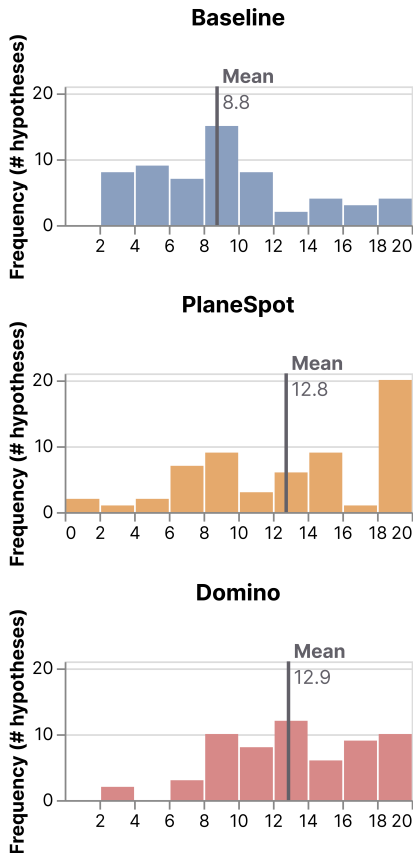


Figure 7: Number of Matching Images. Histograms of the number of images in each slice that match the user’s hypothesis (out of the top 20), stratified by each condition (slicing algorithm). The mean number of matching images for each condition is denoted by a vertical line. We find that on average, users selected more images as “matching” their hypothesis for slices output by PLANESPOT and DOMINO.




Figure 8: User Consistency Barplot. A stacked bar plot visualizing the number of slices from the DOMINO and PLANESPOT conditions (*x*-axis) that have 1, 2, and 3 distinct hypotheses (of the three different users’ hypotheses) (*y*-axis). We find that at least two out of three users write down distinct hypotheses for the majority (75%) of DOMINO and PLANESPOT slices.



Figure 9: User Consistency Examples. Users’ hypotheses for two example slices output by PLANESPOT. For the top slice (#39), all three users wrote down synonymous hypotheses that the model underperforms on images of “broccoli on pizza”. In contrast, for the bottom slice (#23), the three users wrote down distinct hypotheses of where the model underperforms.



baseline condition, supporting **H2** and **H3**, only DOMINO had a statistically significant difference in the proportion of correct hypotheses. While PLANESpot had a higher proportion of correct hypotheses than the BASELINE condition, the difference was not statistically significant for the majority of performance gap thresholds  $\tau$  (Appendix I.1).

In summary, our results provide preliminary evidence that existing tools may offer some benefit relative to a naive workflow of inspecting and trying to make sense of a random sample of errors. However, the slices output by existing tools do not *always* help users form correct hypotheses. For example, only 49% of users' hypotheses corresponding to slices output by PLANESpot described groups where the model performed at least 20% worse. Thus, existing tools do have the potential to mislead users to develop false beliefs if used as proposed by past work.

**Coherence does not imply correctness** Our study calls into question previous assumptions about slice *coherence* and hypothesis *correctness* (**H4**). A number of works have implied that if the user can describe all the images in a slice, that the model underperforms on all images that match their description. Our finding that there is no significant association between slice coherence and hypothesis correctness shows that this assumption is misleading.

While it may seem counter-intuitive, we are *not* surprised that existing slice discovery tools tend to underrepresent the model's performance on coherent groups. We hypothesize that this behavior is a feature (rather than a bug) resulting from the way that existing tools are designed to return high-error subsets. Existing tools may be incentivised to output misleading slices that contain a misrepresentative sample of high-error images *within* each group, even in cases where the model doesn't actually perform worse on the *entire* group. Evaluations that only consider the coherence of the output slices fail to account for whether the slice is a representative sample of the model's performance on the entire group.

This finding has significant implications for researchers, who should center hypothesis correctness (rather than slice coherence) when evaluating their tools. But perhaps more importantly, this finding is significant for *users* of these tools. While present-day tools can help users form hypotheses, users should take caution to know that their hypotheses, no matter how aligned with the slices they may be, may not necessarily be correct. Thus, validating behavioral hypotheses is not only important for researchers, but also for users.

**Variance across users** Our finding that different users form different hypotheses when down the same slice points to the points to the complexity of hypothesis formation as a *human-centered* task. This finding contradicts the dominant assumption from ML research that all users can easily identify the semantic features shared by a subset of data, or that users will simply "know it when they see it". But, disagreement among users whom are asked to complete the same task has been long-demonstrated and well-studied in the field of *crowdsourcing* (Callison-Burch 2009; Chang, Amershi, and Kamar 2017). In practice, different stakeholders may bring different prior knowledge that may cause them to develop different hypotheses when shown the same slice.

If handled thoughtfully, however, this variance across users' hypotheses may actually serve as a strength: for example, a team of stakeholders could discuss several candidate hypotheses for a given slice (Drapeau et al. 2016). More generally, we believe that the field of slice discovery has much to learn from related work on crowdsourcing and annotation.

## 6.2 Design Opportunities

Our findings point to several design opportunities for ML and HCI researchers. We highlight a few exciting directions for future work below.

**Supporting hypothesis formation** In the hypothesis formation step, stakeholders make sense of a large amount of information (*e.g.*, the images belonging to each slice) to form hypotheses about model behavior. However, the human factors that affect participants' hypothesis formation process have been understudied.

One dimension that has been neglected by past work is how the output slices should be *visualized* and presented to users. When asked if there was "any other information that they were not shown that [they] believe would have helped [them] complete the questionnaire", several users expressed a desire to see the model's performance on a larger set of examples beyond the top-20 images belonging to the slice. One user wrote, "*it would have been helpful to retrieve other images from the dataset during slice labeling* (hypothesis formation)" to "*see how good my label* (hypothesis) *was*".

We envision several alternative ways to present each slice beyond the static grid used in our slice overview. For example, one could design novel interactive workflows that allow users to explore the model's performance on multiple slices simultaneously (Bertucci et al. 2022; Cabrera et al. 2023), or compare the images within to images outside of each slice. Designing an appropriate visualization that accounts for the cognitive load and biases of the user (Gajos and Chauncey 2017; Cabrera et al. 2022) is an important direction for future work. Further, the most helpful information to show the user may vary depending on their expertise, *e.g.*, whether the user is an ML developer or a domain expert.

**Prioritizing hypothesis correctness** Our study shows that slices that appear to be coherent can mislead users to develop incorrect beliefs about where their model underperforms. Thus, we urge researchers to prioritize developing tools that help users form *correct* hypotheses. One problem we identified is that existing slice discovery tools are incentivised to output high-error samples that may not capture the full diversity of images that belong to each group. Future work could study how to output slices that contain a more representative sample of each group.

**Towards real-time hypothesis validation** Another promising direction is to develop tools that allow users to validate their own hypotheses in real time. Such tools would enable users to iteratively refine their hypotheses (*i.e.*, explore multiple possible errors that could correspond to a slice), and serve as a sanity check before investing in expensive downstream actions based on false beliefs. Most existing workflows to help users collect evidence to validate

their hypotheses often prioritize retrieving examples that are most *similar* to those that support the users’ hypothesis (Cabrera et al. 2022; Gao et al. 2022; Suresh et al. 2023). In contrast, one potentially promising direction is to prioritize retrieving “counter-evidence”: examples that contradict the users’ hypothesis. We hypothesize that counter-evidence may help combat stakeholders’ potential confirmation bias and help them more quickly iterate on their hypotheses.

**Interactive workflows for slice discovery** We encourage researchers to re-imagine where and how automation can help users discover underperforming groups. Our finding that different users construct inconsistent descriptions of the same group of datapoints calls into question the workflow put forward in past work. Existing slice discovery tools are meant to be run once, and then a stakeholder must make sense of the groups of datapoints they output. One could imagine more bespoke and interactive workflows that better utilize stakeholder’s domain knowledge to define coherent subsets of data, or leverage automation to help stakeholders refine their hypotheses. As one example, a stakeholder could guide a clustering algorithm using their contextual understanding of semantic similarity (Rajani et al. 2022; Cabrera et al. 2023) rather than simply looking at the output clusters.

### 6.3 Limitations

The participants in our study were students with intermediate knowledge of machine learning, and we studied a model trained using data from a simple object detection task. While much past work has developed tools intended to be used by model developers with ML expertise (such as the participants in our study), an emerging line of work has studied how to support users who do not have technical expertise, but *do* have situated domain knowledge, in evaluating ML models (Suresh et al. 2023). For example, clinicians (Gaube et al. 2023) or content moderators (Suresh et al. 2023) who have a deeper understanding of their data may interpret the output slices differently. They may be able to use their domain knowledge to better characterize the semantic features shared by the examples in each slice. Furthermore, in some domains, different stakeholders may disagree about desired model behavior, such as whether a comment should be moderated (Sap et al. 2022). We encourage future work to further identify and examine these user-centered and context-specific dimensions of slice discovery.

While our work is an important step forward towards prioritizing hypothesis validation, we note that validating natural language hypotheses is difficult, and several open challenges remain. Because retrieving all of the images that match each text description hypothesis is difficult and expensive, we follow Gao et al. (2022) and approximate the model’s performance on the group by retrieving a sample of matching images. Unfortunately, the retrieved sample may not be representative of the model’s performance on all images in the group; thus, we compare the model’s performance to another retrieved sample to account for the implicit bias of our retrieval process (details in Appendix B). For these reasons, our calculated performance gaps are only an *approximation* of the model’s performance on each group.

Despite these limitations, we found that all of our experimental findings were consistent across a range of ablations to our approximation strategy (Appendix I.4). We believe that identifying inexpensive ways to retrieve a sample that matches a natural language hypothesis is an important direction for future work.

Finally, our study focuses on the task of describing where the model underperforms *in natural language*. We asked users to describe the underperforming groups in words because doing so is a prerequisite for several downstream actions one might take to address the behavior, and for communicating about the behavior with a wider set of stakeholders. We acknowledge that natural language hypotheses are imperfect for many use cases due to the implicit subjectivity or under-specification of natural language in some contexts. For example, if a practitioner hypothesizes that her object detection model underperforms at detecting stop signs in photos where they are “*far away from the camera*”, determining which photos qualify as “*far away*” is under-specified from her description alone. Developing a more formal, yet simultaneously accessible “*domain-specific language*” (Desai et al. 2015; Wu et al. 2019) for users’ hypotheses is an open direction for future work.

## 7 Conclusion

While a growing number of works develop new slice discovery tools to help people discover where their model underperforms, there has been little evaluation of if, and how, humans can make sense of their output. In our controlled user study with 15 participants, we found preliminary evidence that existing tools may offer some benefit relative to a naive baseline of examining a random sample of errors. Our results also challenge several dominant assumptions shared by past work on slice discovery.

First, coherence of the output slices does *not* imply that users can form correct behavioral hypotheses. Our results indicate that being able to identify a description that matches the majority of images in the slice does *not* mean that the model underperforms on all such images. This finding has important consequences for evaluation and everyday use of existing slice discovery tools. We caution researchers away from using the number of images that match a common description as a measure of their algorithms’ utility.

Second, we found that different users form different hypotheses when shown the same slice, which highlights the under-explored complexity of the hypothesis formation step. Future work can consider alternative visualizations to help users make sense of the semantic features shared by the images in each slice, or make productive use of user disagreement to consider a wider range of possible model errors.

Our findings point to user needs and design opportunities to better support stakeholders as they form and validate hypotheses of model behavior. More broadly, we hope that our work is a first step towards centering users when designing and evaluating new tools for slice discovery.

## Acknowledgments

We thank our study participants and annotators who made our research possible. We thank Sherry Tongshuang Wu, Donald Bertucci, Valerie Chen, Vijay Viswanathan, Jennifer Hsia, Katelyn Morrison, and Nupoor Gandhi for helpful feedback and discussions. We also thank the reviewers at HCOMP 2023 and the ICML Second Workshop on Spurious Correlations, Invariance, and Stability for their suggestions that improved our paper. This work was supported in part by the National Science Foundation grants IIS1705121, IIS1838017, IIS2046613, IIS2112471, and funding from Meta, Morgan Stanley, Amazon, and Google. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of any of these funding agencies.

## References

- Adebayo, J.; Muelly, M.; Abelson, H.; and Kim, B. 2022. Post hoc Explanations may be Ineffective for Detecting Unknown Spurious Correlation. In *International Conference on Learning Representations*.
- Balayn, A.; Rikalo, N.; Yang, J.; and Bozzon, A. 2023. Faulty or Ready? Handling Failures in Deep-Learning Computer Vision Models until Deployment: A Study of Practices, Challenges, and Needs. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI '23*. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215.
- Bertucci, D.; Hamid, M. M.; Anand, Y.; Ruangrotsakun, A.; Tabatabai, D.; Perez, M.; and Kahng, M. 2022. DendroMap: Visual Exploration of Large-Scale Image Datasets for Machine Learning with Treemaps. arXiv:2205.06935.
- Birhane, A. 2022. The unseen Black faces of AI algorithms. In *Nature* 610, 451–452.
- Buolamwini, J.; and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, 77–91. PMLR.
- Cabrera, Á. A.; Druck, A. J.; Hong, J. I.; and Perer, A. 2021. Discovering and Validating AI Errors With Crowdsourced Failure Reports. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2): 1–22.
- Cabrera, A. A.; Fu, E.; Bertucci, D.; Holstein, K.; Talwalkar, A.; Hong, J. I.; and Perer, A. 2023. Zeno: An Interactive Framework for Behavioral Evaluation of Machine Learning. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI '23*. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215.
- Cabrera, A. A.; Tulio Ribeiro, M.; Lee, B.; Deline, R.; Perer, A.; and Drucker, S. M. 2022. What Did My AI Learn? How Data Scientists Make Sense of Model Behavior. *ACM Trans. Comput.-Hum. Interact.*, 30(1).
- Callison-Burch, C. 2009. Fast, Cheap, and Creative: Evaluating Translation Quality Using Amazon’s Mechanical Turk. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, 286–295. Singapore: Association for Computational Linguistics.
- Chang, J. C.; Amershi, S.; and Kamar, E. 2017. Revolt: Collaborative Crowdsourcing for Labeling Machine Learning Datasets. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, CHI '17*. New York, NY, USA: ACM.
- d’Eon, G.; d’Eon, J.; Wright, J. R.; and Leyton-Brown, K. 2021. The spotlight: A general method for discovering systematic errors in deep learning models. *arXiv preprint arXiv:2107.00758*.
- Desai, A.; Gulwani, S.; Hingorani, V.; Jain, N.; Karkare, A.; Marron, M.; R, S.; and Roy, S. 2015. Program Synthesis using Natural Language. arXiv:1509.00413.
- Drapeau, R.; Chilton, L.; Bragg, J.; and Weld, D. 2016. MicroTalk: Using Argumentation to Improve Crowdsourcing Accuracy. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 32–41. HCOMP.
- Eyuboglu, S.; Varma, M.; Saab, K. K.; Delbrouck, J.-B.; Lee-Messer, C.; Dunnmon, J.; Zou, J.; and Re, C. 2022. Domino: Discovering Systematic Errors with Cross-Modal Embeddings. In *International Conference on Learning Representations*.
- Gajos, K. Z.; and Chauncey, K. 2017. The Influence of Personality Traits and Cognitive Load on the Use of Adaptive User Interfaces. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces, IUI '17*, 301–306. New York, NY, USA: Association for Computing Machinery. ISBN 9781450343480.
- Gao, I.; Ilharco, G.; Lundberg, S.; and Ribeiro, M. T. 2022. Adaptive Testing of Computer Vision Models. arXiv:2212.02774.
- Gaube, S.; Suresh, H.; Raue, M.; Lermer, E.; Koch, T.; Hudecek, M.; Ackery, A.; Grover, S.; Coughlin, J.; Frey, D.; Kitamura, C.; Ghassemi, M.; and Colak, E. 2023. Non-task expert physicians benefit from correct explainable AI advice when reviewing X-rays. *Scientific Reports*, 13.
- Holstein, K.; Vaughan, J. W.; Daumé, H.; Dudik, M.; and Wallach, H. 2019. Improving Fairness in Machine Learning Systems. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM.
- Idrissi, B. Y.; Arjovsky, M.; Pezeshki, M.; and Lopez-Paz, D. 2022. Simple data balancing achieves competitive worst-group-accuracy. arXiv:2110.14503.
- Kim, B.; Wattenberg, M.; Gilmer, J.; Cai, C.; Wexler, J.; Viegas, F.; et al. 2018. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International conference on machine learning*, 2668–2677. PMLR.
- Li, Z.; Evtimov, I.; Gordo, A.; Hazirbas, C.; Hassner, T.; Ferrer, C. C.; Xu, C.; and Ibrahim, M. 2022. A Whac-A-Mole Dilemma: Shortcuts Come in Multiples Where Mitigating One Amplifies Others. *arXiv preprint arXiv:2212.04825*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.

- Liu, E. Z.; Haghgoo, B.; Chen, A. S.; Raghunathan, A.; Koh, P. W.; Sagawa, S.; Liang, P.; and Finn, C. 2021. Just train twice: Improving group robustness without training group information. In *International Conference on Machine Learning*, 6781–6792. PMLR.
- Moore, S.; Liao, Q. V.; and Subramonyam, H. 2023. FailureNotes: Supporting Designers in Understanding the Limits of AI Models for Computer Vision Tasks. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215.
- Plumb, G.; Johnson, N.; Ángel Alexander Cabrera; and Talwalkar, A. 2023. Towards a More Rigorous Science of Blindspot Discovery in Image Models. arXiv:2207.04104.
- Polyzotis, N.; Whang, S.; Kraska, T. K.; and Chung, Y. 2019. Slice Finder: Automated Data Slicing for Model Validation. In *Proceedings of the IEEE Int' Conf. on Data Engineering (ICDE)*, 2019.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. arXiv:2103.00020.
- Rajani, N.; Liang, W.; Chen, L.; Mitchell, M.; and Zou, J. 2022. SEAL : Interactive Tool for Systematic Error Analysis and Labeling. arXiv:2210.05839.
- Raji, I. D.; and Buolamwini, J. 2022. Actionable Auditing Revisited: Investigating the Impact of Publicly Naming Biased Performance Results of Commercial AI Products. *Commun. ACM*, 66(1): 101–108.
- Sap, M.; Swayamdipta, S.; Vianna, L.; Zhou, X.; Choi, Y.; and Smith, N. A. 2022. Annotators with Attitudes: How Annotator Beliefs And Identities Bias Toxic Language Detection. arXiv:2111.07997.
- Shankar, V.; Roelofs, R.; Mania, H.; Fang, A.; Recht, B.; and Schmidt, L. 2020. Evaluating Machine Accuracy on ImageNet. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 8634–8644. PMLR.
- Singla, S.; Nushi, B.; Shah, S.; Kamar, E.; and Horvitz, E. 2021. Understanding Failures of Deep Networks via Robust Feature Extraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12853–12862.
- Sohoni, N.; Dunnmon, J.; Angus, G.; Gu, A.; and Ré, C. 2020. No subclass left behind: Fine-grained robustness in coarse-grained classification problems. *Advances in Neural Information Processing Systems*, 33: 19339–19352.
- Suresh, H.; Shanmugam, D.; Bryan, A.; Chen, T.; D'Amour, A.; Gutttag, J. V.; and Satyanarayan., A. 2023. Kaleidoscope: Semantically-grounded, context-specific ML model evaluation. In *CHI Conference on Human Factors in Computing Systems*.
- Vendrow, J.; Jain, S.; Engstrom, L.; and Madry, A. 2023. Dataset Interfaces: Diagnosing Model Failures Using Controllable Counterfactual Generation. arXiv:2302.07865.
- Wang, F.; Adebayo, J.; Tan, S.; Garcia-Olano, D.; and Kokhlikyan, N. 2023. Error Discovery by Clustering Influence Embeddings. In *ICLR 2023 Workshop on Pitfalls of limited data and computation for Trustworthy ML*.
- Wiles, O.; Albuquerque, I.; and Gowal, S. 2023. Discovering Bugs in Vision Models using Off-the-shelf Image Generation and Captioning. arXiv:2208.08831.
- Wu, T.; Ribeiro, M. T.; Heer, J.; and Weld, D. 2019. Errudite: Scalable, Reproducible, and Testable Error Analysis. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 747–763. Florence, Italy: Association for Computational Linguistics.