

Identification of Systematic Errors of Image Classifiers on Rare Subgroups

Jan Hendrik Metzen⁽¹⁾, Robin Huttmacher⁽¹⁾, N. Grace Hua⁽¹⁾, Valentyn Boreiko^(1,2), Dan Zhang⁽¹⁾
(1) Bosch Center for Artificial Intelligence, Robert Bosch GmbH (2) University of Tübingen

{janhendrik.metzen, robin.huttmacher, grace.hua, dan.zhang2}@de.bosch.com; valentyn.boreiko@uni-tuebingen.de

Abstract

Despite excellent average-case performance of many image classifiers, their performance can substantially deteriorate on semantically coherent subgroups of the data that were under-represented in the training data. These systematic errors can impact both fairness for demographic minority groups as well as robustness and safety under domain shift. A major challenge is to identify such subgroups with subpar performance when the subgroups are not annotated and their occurrence is very rare. We leverage recent advances in text-to-image models and search in the space of textual descriptions of subgroups (“prompts”) for subgroups where the target model has low performance on the prompt-conditioned synthesized data. To tackle the exponentially growing number of subgroups, we employ combinatorial testing. We denote this procedure as PROMPTATTACK as it can be interpreted as an adversarial attack in a prompt space. We study subgroup coverage and identifiability with PROMPTATTACK in a controlled setting and find that it identifies systematic errors with high accuracy. Thereupon, we apply PROMPTATTACK to ImageNet classifiers and identify novel systematic errors on rare subgroups.

1. Introduction

Deep learning based approaches have revolutionized many fields of computer vision [31, 37] and are increasingly applied in safety-critical applications such as automated driving [20]. An important prerequisite for deployment of learned models in such safety-critical domains is that they need to work reasonably well for all subgroups from an operational design domain [10] and strong requirements are imposed on assuring safety of systems [8]. That is: there must not be catastrophic but avoidable failure cases on any subgroup, regardless of how rare the subgroup might be. Unfortunately, spurious correlations in the training data can often result in classifiers that utilize shortcut decision making [18, 50, 30, 29] - a phenomenon long known from animal and human psychology [40]. Such shortcuts can work

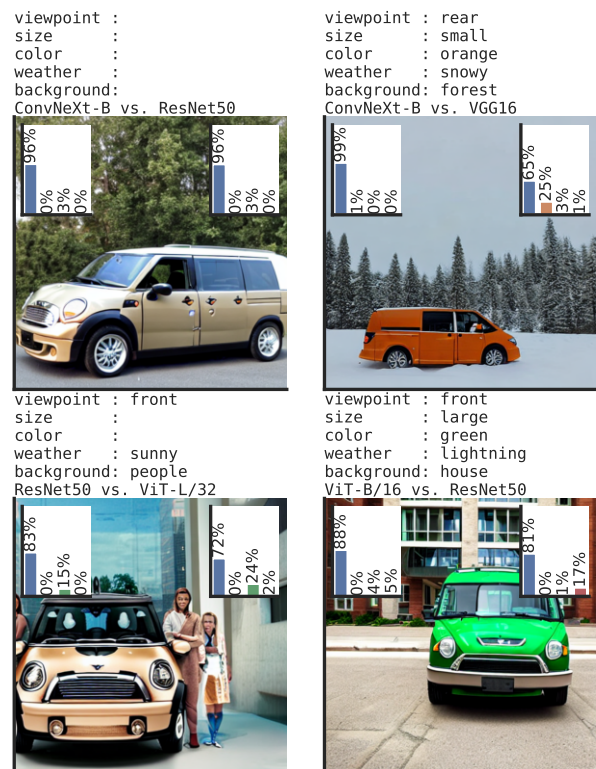


Figure 1. Samples along with histograms over two models’ class prediction rates (shown in left and right inlays, based on 400 samples) for 4 different subgroups. The baseline subgroup (top left) is classified mostly as **minivan** by all models, while the misclassification rates to a **snowplow** (top right), to **pickup** (bottom left), and to **police van** (bottom right) are significantly increased on the shown subgroups for a VGG16, ViT-L/32, ResNet50 respectively. We refer to Section 5 and D for more details and samples.

well on subgroups that occur frequently in-distribution, that is: on data that follows the same distribution as the training data. However, they can easily fail after a domain shift to out-of-distribution data since very rare subgroups become suddenly much more frequent [57]. For instance, Beery et al. [6] demonstrate that a shift in background can largely affect an image classifier, resulting in misclassifying, e.g., a cow on the beach.

Accordingly, a crucial aspect of *model auditing* [2] is to separately evaluate the behaviour of a classifier on every subgroup from a large set of subgroups. If the performance of a classifier on certain subgroups is considerably worse than on the totality of the domain’s data, then we denote this subgroup as a *systematic error* of a classifier [13, 26, 53], a case of hidden stratification [48]. More specifically, a systematic error refers to a subgroup of inputs on which a pretrained classifier has a high probability of misclassification (“error”) and at the same time a large percentage of elements in the subgroup share a human-interpretable concept: the group appears semantically coherent to a human (“systematic”). Applying methods for identifying such systematic errors could become a prerequisite for deployment in many domains while at the same time, systematic error are *actionable*: their exemplars can be used for finetuning a model and improving its robustness, reliability, and fairness [17, 28].

Some prior work [13, 26] require the availability of a labelled hold-out set covering data from rare subgroups for identification of systematic errors on these subgroups. Unfortunately, it is often expensive to acquire (labelled) data for certain subgroups that are very rare in the domain pre-deployment, even though subgroups could become much more frequent after some domain shift. This is problematic because systematic errors are much more likely to occur on subgroups that are rare in the training data. Other prior work is based on large unlabelled hold-out data but requires a human-in-the-loop [17], which increases the costs of systematic error identification and thus limits applicability.

Another line of work (including ours) focuses on auditing models on synthetically generated subgroup data. Recent progress of *text-to-image models* [42, 43, 45, 9] in terms of compositionality can allow synthesizing data of rare subgroups that have not been part of the training data. Wiles et al. [53] focused on an open-ended approach that synthesizes data according to the distribution induced by a fixed prompt that encodes the class but no subgroup information. Concurrently to our work, Vendrow et al. [51] generated text-conditioned counterfactual examples to study the robustness to single semantic shifts.

We propose PROMPTATTACK (see Figure 2), which leverages text-to-image models for synthesizing images of subgroups by encoding subgroup information directly in the prompt. To deal with large operational design domains and the resulting combinatorial explosion of subgroups, PROMPTATTACK builds upon *combinatorial testing* [38, 4] that allows a near-equable coverage of the operational design domain while keeping the number of explored subgroups relatively small. In contrast to the open-ended approach by Wiles et al. [53], PROMPTATTACK is targeted and reliably explores subgroups from a prespecified operational design domain (see Section 4.1). Moreover, it does

not require any pretrained models or heuristics components for failure case clustering and captioning. In contrast to Vendrow et al. [51], PROMPTATTACK can identify systematic errors on subgroups that require the concurrence of several semantic shifts (see Figure 7).

Overall, our main contributions are the following:

- In Section 3, we introduce PROMPTATTACK, a novel procedure for identifying systematic errors based on synthetic data from a text-to-image model, conditioned on a prompt encoding subgroup and class information. PROMPTATTACK explores a large subset of subgroups from an operational design domain using combinatorial testing, achieving near-equable coverage of subgroups (Section 4.1).
- We propose a benchmark for testing and comparing methods for systematic error identification (Section 4.2). In contrast to prior work [13], this benchmark does not train multiple classifiers with training-time interventions but is based purely on inference-time interventions on zero-shot classifiers such as CLIP [41].
- PROMPTATTACK identifies classifier-specific and targeted systematic misclassifications on rare subgroups of ImageNet classifiers (see Figure 1 and Section 5).

2. Related Work

We review related work in the computer vision domain while noting that identification of systematic errors and harmful behaviour is an important topic in other fields such as large language models as well (“red teaming”) [39, 16].

Building upon Subgroup Annotation. Several prior works have investigated performance on datasets where information on certain semantic dimensions is available for each datapoint and thus direct evaluation of subgroup error is feasible. For instance, Hendrycks et al. [22] collected four real-world datasets containing semantic dimensions like *artistic renditions* (ImageNet-R), *country*, *year*, and *camera* (StreetView StoreFronts), or *object size*, *object occlusion*, *camera viewpoint*, and *camera zoom* (DeepFashion Remixed). The influence of *image background* can be studied based on ImageNet-9 [54] and Waterbirds [44]. WildDash [55] allows studying the impact of different *visual hazards*. ImageNet-X [25] adds sixteen human annotations of semantic dimensions such as *pose*, *background*, or *lighting* to each ImageNet-1k validation sample. Such approaches require large efforts in data collection and subgroup annotation and have thus limited scalability and flexibility. Moreover, if there are several interacting semantic dimensions, then the number of datapoints required to achieve a full coverage of the design domain grows exponentially. One partial remedy to this combinatorial explosion is combinatorial t-wise testing [38, 4, 19]. Alternatively, synthetic corruptions can be applied to existing images resulting in a semantic *corruption* dimension (ImageNet-C) [23]. How-

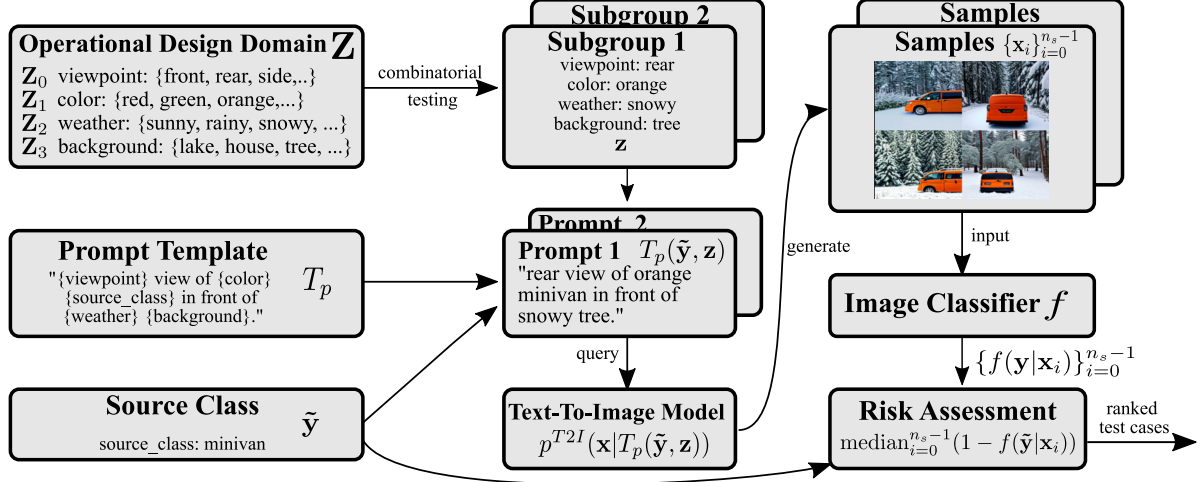


Figure 2. Illustration of PROMPTATTACK: domain experts define an operational design domain \mathbf{Z} consisting of semantic dimensions \mathbf{Z}_j . Combinatorial testing is used to generate a set of subgroups. A prompt $T_p(\tilde{\mathbf{y}}, \mathbf{z})$ is instantiated from a prompt template T_p based on the respective subgroup \mathbf{z} and source class $\tilde{\mathbf{y}}$. A text-to-image model p^{T2I} generates n_s samples $\{\mathbf{x}_i\}_{i=0}^{n_s-1}$ for the prompt. The image classifier f under investigation predicts class probabilities $\{f(\mathbf{y}|\mathbf{x}_i)\}_{i=0}^{n_s-1}$ for the samples. An objective function provides a ranking of subgroups based upon the source class predictions $f(\tilde{\mathbf{y}}|\mathbf{x}_i)$, where low median class score for the source class indicates a potential systematic error.

ever, not all semantic dimensions can be simulated.

Failure Identification without Subgroup Annotation.

Since explicit annotation of semantic dimension is costly, recent works have focused on automating the process of identifying systematic errors. One line of work resorts to a labeled hold-out set. Coherent groups of errors on the hold-out set can be identified by error-aware soft-clustering on the final feature space of the classifier [11]. However, this does not provide an interpretation of the subgroup. Leveraging the text-image embedding alignment in CLIP [41], both [13] and [26] operated in the latent space of CLIP to identify semantically coherent subgroups and generate human interpretable subgroup annotation. The main disadvantage of these works is that they require the availability of a labeled hold-out set. Since systematic errors are more likely to occur on untypical/rare data [13], identifying them requires a hold-out dataset containing such cases, which is unrealistic. AdaVision [17] introduced a human-in-the-loop process to discover systematic errors by adaptively querying real images from LAION-5B [46] (via CLIP similarity). In contrast to AdaVision, our procedure does not require a human-in-the-loop, which can be preferred if model auditing needs to be done regularly or for a large number of models.

The most similar line of work to ours is built on top of text-to-image synthesis models [42, 43, 45, 9]. For instance, in [27, 35, 51], counterfactual examples are generated according to the input text, which indicates the semantic shift, e.g., background, lighting, or style. These works considered a single semantic shift, whereas systematic errors can result from compounding shifts in multiple semantic factors. Wiles et al. [53] did not pre-specify the semantic shift but iteratively synthesize samples based on a text description,

cluster the failure cases, and refine the text description. In contrast to such an “open-ended” search, our work focuses on finding failures within an operational design domain and achieves high coverage of that domain (see Section 4.1). Moreover, our approach is conceptually more simple and does not require clustering and captioning of failure cases. Both approaches can be seen as complementary.

3. Method

In this section, we introduce our proposed procedure PROMPTATTACK; see Figure 2 for an illustration.

3.1. Background

We consider image classifiers $f : \mathbb{X} \times \mathbb{Y} \rightarrow [0, 1]$ and denote the predicted probability of class $\mathbf{y} \in \mathbb{Y} = \{1, \dots, C\}$ for image $\mathbf{x} \in \mathbb{X}$ by $f(\mathbf{y}|\mathbf{x})$. We assume that we operate in a domain where \mathbf{x}, \mathbf{y} are governed by a distribution $p(\mathbf{x}, \mathbf{y})$. We are interested in exploring properties of f on semantically coherent subgroups of the data manifold, which we formalize by conditioning on some latent \mathbf{z} : $p(\mathbf{x}, \mathbf{y}|\mathbf{z})$. For brevity, we also denote the subgroups themselves by \mathbf{z} . We note that in contrast to Wiles et al. [53], we do not build on the conditional distribution $p(\mathbf{z}|\mathbf{x}, \mathbf{y})$ and thus do not require an image captioning model.

3.2. Systematic Errors: Definition

We define the risk of a classifier f on a subgroup \mathbf{z} by $R_f(\mathbf{z}) = \mathbb{E}_{p(\mathbf{x}, \mathbf{y}|\mathbf{z})} L(f(\cdot|\mathbf{x}), \mathbf{y})$, where \mathbb{E}_p denotes the expectation over p and $L : [0, 1]^C \times \mathbb{Y} \mapsto \mathbb{R}$ is a loss function. Moreover, we set the baseline (irreducible) risk on \mathbf{z} to $R_B(\mathbf{z}) = \mathbb{E}_{p(\mathbf{x}, \mathbf{y}|\mathbf{z})} L(p(\cdot|\mathbf{x}, \mathbf{z}), \mathbf{y})$,

with $p(\mathbf{y}|\mathbf{x}, \mathbf{z}) = p(\mathbf{x}, \mathbf{y}|\mathbf{z})/p(\mathbf{x}|\mathbf{z})$. We note $R_f(\mathbf{z}) \geq R_B(\mathbf{z})$. We assume that we are provided with a predefined set of subgroups \mathbf{Z} that we denote as the *operational design domain* (ODD) [10]. We are interested in subgroups $\mathbf{z} \in \mathbf{Z}$ on which a classifier f has high risk $R_f(\mathbf{z})$ while the baseline risk $R_B(\mathbf{z})$ remains low. If the subgroups \mathbf{z} are designed in a way to encourage semantic coherence, such high risk subgroups are human interpretable and actionable. More specifically, we rank subgroups $\mathbf{z} \in \mathbf{Z}$ based on $R(\mathbf{z}) = R_f(\mathbf{z}) - R_B(\mathbf{z})$; top-ranked \mathbf{z} with sufficiently high risk are *systematic errors*. Similarly, we define a *systematic misclassification* into class $\mathbf{y}^{(t)}$ by $R_f(\mathbf{z}, \mathbf{y}^{(t)}) = \mathbb{E}_{p(\mathbf{x}, \mathbf{y}|\mathbf{z})} L(f(\cdot|\mathbf{x}), \mathbf{y}^{(t)}) 1_{[\mathbf{y} \neq \mathbf{y}^{(t)}]}$ and $R_B(\mathbf{z}, \mathbf{y}^{(t)})$ analogously, with $1_{[\mathbf{y} \neq \mathbf{y}^{(t)}]}$ being the indicator function of $\mathbf{y} \neq \mathbf{y}^{(t)}$. The top-ranked \mathbf{z} according to $R(\mathbf{z}, \mathbf{y}^{(t)}) = R_f(\mathbf{z}, \mathbf{y}^{(t)}) - R_B(\mathbf{z}, \mathbf{y}^{(t)})$ are systematic misclassifications into $\mathbf{y}^{(t)}$ for sufficiently high $R(\mathbf{z}, \mathbf{y}^{(t)})$. We note the risk R can be made invariant to the classifier’s calibration (e.g., for L being a 0-1 loss function) or sensitive to it (for most other choices of L).

3.3. Systematic Errors: Approximations

We make several approximations and assumptions for tractable systematic error identification; empirical evidence in Section 4 suggests these hold reasonably well in practice.

1. Monte Carlo Approximation. In general we cannot compute $\mathbb{E}_{p(\mathbf{x}, \mathbf{y}|\mathbf{z})}$ in the definition of $R_f(\mathbf{z})$. Thus, we resort to approximating the expectation based on n_s samples $\mathbf{x}_i, \mathbf{y}_i \sim p(\mathbf{x}, \mathbf{y}|\mathbf{z})$: $R_f(\mathbf{z}) \approx \frac{1}{n_s} \sum_{i=0}^{n_s-1} L(f(\cdot|\mathbf{x}_i), \mathbf{y}_i)$.

2. Synthetic Data. Real-world samples $\mathbf{x}, \mathbf{y} \sim p(\mathbf{x}, \mathbf{y}|\mathbf{z})$ from semantically coherent subgroups \mathbf{z} are typically not available for two reasons: (i) typical real-world data \mathbf{x}, \mathbf{y} lacks human annotated information on \mathbf{z} , such as captions in a given format or subgroup annotations. (ii) Even if \mathbf{z} were available or could be inferred, the coverage of \mathbf{Z} can be very low: some rare subgroups $\mathbf{z} \in \mathbf{Z}$ will have low $p(\mathbf{z}|\mathbf{x}, \mathbf{y})$ and may not be represented at all in finite sample sets from $p(\mathbf{x}, \mathbf{y})$. However, the performance of f on such rare subgroups can still be highly relevant in safety-critical applications, as specifically rare corner cases may be the ones where the generalization of a classifier f fails. Because of this, we resort to sampling $\mathbf{x}, \mathbf{y}|\mathbf{z}$ from learned approximations \hat{p} of the real-world data distribution. For this, we leverage recent text-to-image models $p^{T2I}(\mathbf{x}|t)$ such as Stable Diffusion [43], which condition image generation on a text prompt t , as detailed below.

3. Sampling Class-Conditional. We can use text-to-image models for sampling from $\hat{p}(\mathbf{x}|\mathbf{z})$ by representing the subgroup \mathbf{z} as a text prompt t . However, sampling from $\hat{p}(\mathbf{x}, \mathbf{y}|\mathbf{z}) = \hat{p}(\mathbf{y}|\mathbf{x}, \mathbf{z})\hat{p}(\mathbf{x}|\mathbf{z})$ would also require an approximation $\hat{p}(\mathbf{y}|\mathbf{x}, \mathbf{z})$, that is: the conditional probability of a specific class \mathbf{y} given an image \mathbf{x} and subgroup \mathbf{z} . Such an

approximation $\hat{p}(\mathbf{y}|\mathbf{x}, \mathbf{z})$ is not generally available or easily estimated (estimating $\hat{p}(\mathbf{y}|\mathbf{x}, \mathbf{z})$ from data would require a large number of samples (\mathbf{x}, \mathbf{y}) annotated with \mathbf{z} , which we precluded above). Instead, we explicitly condition the generation of \mathbf{x} on a desired class $\tilde{\mathbf{y}}$. Effectively, this corresponds to focusing on systematic errors on a specific source class $\tilde{\mathbf{y}}$. We realize the approximate $\hat{p}(\mathbf{x}|\tilde{\mathbf{y}}, \mathbf{z})$ by including class information $\tilde{\mathbf{y}}$ along with \mathbf{z} in a text prompt t as detailed in Section 3.4.

4. Negligible Baseline Risk. We cannot evaluate $R_B(\mathbf{z})$ directly since $p(\mathbf{y}|\mathbf{x}, \mathbf{z})$ is unavailable. Because of this, we limit ourselves to choices of \mathbf{Z} where by design for every $\mathbf{z} \in \mathbf{Z}$, we have for $\mathbf{x} \sim p(\mathbf{x}|\tilde{\mathbf{y}}, \mathbf{z})$ that $p(\mathbf{y}|\mathbf{x}, \mathbf{z}) \approx 1$ if $\mathbf{y} = \tilde{\mathbf{y}}$ else 0. That is: classes do not overlap on \mathbf{z} and images $\mathbf{x}|\tilde{\mathbf{y}}, \mathbf{z}$ belong unambiguously to the same class $\tilde{\mathbf{y}}$. Accordingly, the baseline risk $R_B(\mathbf{z})$ on $p(\mathbf{x}, \mathbf{y}|\mathbf{z})$ is negligibly small for typical loss functions (unlike on the unconditional $p(\mathbf{x}, \mathbf{y})$) and we can approximate $R(\mathbf{z}) \approx R_f(\mathbf{z})$ and $R(\mathbf{z}, \mathbf{y}^{(t)}) \approx R_f(\mathbf{z}, \mathbf{y}^{(t)})$.

Dealing with Violations. We note that the above approximations do not hold strictly as the generative model $\hat{p}(\mathbf{x}|\tilde{\mathbf{y}}, \mathbf{z})$ will not perfectly approximate the real-data subgroups $p(\mathbf{x}|\tilde{\mathbf{y}}, \mathbf{z})$: it may generate (i) valid data from $p(\mathbf{x}|\tilde{\mathbf{y}})$ that is “out-of-subgroup” (OOS), that is: has very low probability under $p(\mathbf{x}|\tilde{\mathbf{y}}, \mathbf{z})$, (ii) data that does not belong to the class $\tilde{\mathbf{y}}$, that is: low $p(\mathbf{x}|\tilde{\mathbf{y}})$ (“out-of-class”, OOC), and (iii) data that is even very unlikely under $p(\mathbf{x})$ (OOD sampling [53]). Recent progress in text-to-image models, for which $\hat{p}(\mathbf{x}|\tilde{\mathbf{y}}, \mathbf{z})$ more closely approximates $p(\mathbf{x}|\tilde{\mathbf{y}}, \mathbf{z})$, makes such OOS/OOC/OOD samples occur less often. To reduce them further, we carefully engineer text prompts for $\tilde{\mathbf{y}}$ and \mathbf{z} (see Section 3.4). This requirement for “prompt engineering” is a shortcoming but we are optimistic that future text-to-image models will reduce its need.

Even with careful prompt engineering, a few OOS/OOC/OOD samples might still dominate the Monte-Carlo estimate for $R_f(\mathbf{z})$. We thus resort to robust estimators of central tendency for $R_f(\mathbf{z})$, which are less affected by outliers, such as $R_f(\mathbf{z}) \approx \text{median}_{i=0}^{n_s-1} L(f(\cdot|\mathbf{x}_i), \tilde{\mathbf{y}})$.

3.4. Operational Design Domain \mathbf{Z}

In principle, our approach allows arbitrary ODDs \mathbf{Z} . We specifically focus on a setting where the ODD is compositional: $\mathbf{Z} = \mathbf{Z}_0 \times \dots \times \mathbf{Z}_{n_Z-1}$, where every \mathbf{Z}_i corresponds to a semantically meaningful dimension. Every $\mathbf{z} \in \mathbf{Z}$ is then a tuple containing n_Z values (one for each semantic dimension). As we use text-to-image models $p^{T2I}(\mathbf{x}|t)$ for sampling from a subgroup, we assume the ODD comes with a *prompt template* T_p that allows mapping this n_Z -dimensional tuple along with the class $\tilde{\mathbf{y}}$ to a text prompt: $t = T_p(\tilde{\mathbf{y}}, \mathbf{z})$. We thus set $\hat{p}(\mathbf{x}|\tilde{\mathbf{y}}, \mathbf{z}) = p^{T2I}(\mathbf{x}|T_p(\tilde{\mathbf{y}}, \mathbf{z}))$. We note that the choice of the prompt template T_p can significantly affect the efficacy of our procedure.

If we specify \mathbf{Z} as above, we have $|\mathbf{Z}| = \prod_{i=0}^{n_Z-1} |\mathbf{Z}_i|$. Accordingly, the number of subgroups in the ODD grows exponentially with the number of semantic dimensions n_Z . To deal with large n_Z , we optionally employ *combinatorial testing* [38, 4] to test only a subset of subgroups $\mathbf{Z}_C \subseteq \mathbf{Z}$ for systematic errors (see Section A in the appendix for a comparison to other search procedures). Specifically, for a value $n_C \leq n_Z$, combinatorial testing ensures that for any $i_0, \dots, i_{n_C-1} \leq n_Z - 1$ and for all $\mathbf{z} \in \mathbf{Z}$ there exists a $\mathbf{z}^C \in \mathbf{Z}_C$ such that $\mathbf{z}_{i_0} = \mathbf{z}_{i_0}^C, \dots, \mathbf{z}_{i_{n_C-1}} = \mathbf{z}_{i_{n_C-1}}^C$. That is, for every combination of n_C semantic dimensions \mathbf{Z}_i , every possible combination of values from these dimensions is covered at least once in \mathbf{Z}_C . Choosing $n_C < n_Z$ reduces the number of tested subgroups at the cost of not reaching a full (but near-equable) coverage of \mathbf{Z} . Combinatorial testing allows evaluating different loss functions concurrently; we use $L(f(\cdot|\mathbf{x}), \tilde{\mathbf{y}}) = 1 - f(\tilde{\mathbf{y}}|\mathbf{x})$ for systematic errors and $L(f(\cdot|\mathbf{x}), \mathbf{y}^{(t)}) = f(\mathbf{y}^{(t)}|\mathbf{x})$ for systematic misclassifications for multiple choices of $\mathbf{y}^{(t)}$.

4. Quantitative Evaluation

We perform quantitative evaluations in terms of coverage properties of an operational design domain (Section 4.1) and the ability of PROMPTATTACK to recover known systematic errors from a zero-shot classifier (Section 4.2).

4.1. Coverage Analysis of Conditional versus Unconditional Synthesis

Motivation. The primary motivation for PROMPTATTACK is to encourage full exploration of an operational design domain \mathbf{Z} by explicitly conditioning image generation on subgroups $\mathbf{z} \in \mathbf{Z}$, that is, to sample from $\hat{p}(\mathbf{x}|\mathbf{z})$ rather than unconditionally from $\hat{p}(\mathbf{x})$ as done by Wiles et al. [53] (we skip the conditioning on $\tilde{\mathbf{y}}$ here for brevity). We investigate in this subsection whether this indeed results in better coverage: $\hat{p}(\mathbf{z}) = \int_{\mathbf{x}} \hat{p}(\mathbf{z}|\mathbf{x})\hat{p}(\mathbf{x}) d\mathbf{x}$ should be near uniform over \mathbf{Z} when $\hat{p}(\mathbf{x})$ is the empirical distribution of samples generated by PROMPTATTACK. For this analysis (but not for PROMPTATTACK itself), we need a mechanism to estimate $\hat{p}(\mathbf{z}|\mathbf{x})$, for which we employ a zero-shot classifier derived from the multimodal image-text model CLIP [41].

Experimental Setting. We consider the class $\tilde{\mathbf{y}} = \text{“car”}$, and the semantic dimensions $\mathbf{Z}_0 = \{\text{black, white, red, green, blue}\}$ corresponding to color, $\mathbf{Z}_1 = \{\text{forest, desert, city, mountain, beach}\}$ corresponding to scene background, and $\mathbf{Z}_2 = \{\text{van, SUV, sedan, cabriolet}\}$ corresponding to car type. As operational design domain we use the full $\mathbf{Z} = \mathbf{Z}_0 \times \mathbf{Z}_1 \times \mathbf{Z}_2$ with $|\mathbf{Z}| = 100$. With $\mathbf{z} = (z_0, z_1, z_2)$ we obtain a factorized $\hat{p}(\mathbf{z}|\mathbf{x}) = \prod_{i=0}^2 \hat{p}(z_i|\mathbf{x})$. For $\hat{p}(z_i|\mathbf{x})$, we use the CLIP-based zero-shot classifier with text queries $T_0 = \{\text{“An image of a } \textit{color} \textit{ car”} | \textit{color} \in \mathbf{Z}_0\}$, $T_1 = \{\text{“An image of a car with } \textit{background} \textit{ background”} | \textit{background} \in$

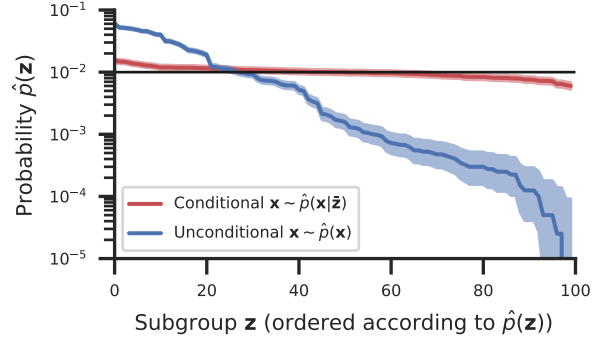


Figure 3. Estimate of $\hat{p}(\mathbf{z}) = \int_{\mathbf{x}} \hat{p}(\mathbf{z}|\mathbf{x})\hat{p}(\mathbf{x}) d\mathbf{x}$ (unconditional) and $\hat{p}(\mathbf{z}) = \sum_{\bar{\mathbf{z}}} p(\bar{\mathbf{z}}) \int_{\mathbf{x}} \hat{p}(\mathbf{z}|\mathbf{x})\hat{p}(\mathbf{x}|\bar{\mathbf{z}}) d\mathbf{x}$ (conditional), obtained using 40,000 Monte-Carlo samples $\mathbf{x} \sim \hat{p}(\mathbf{x}|\bar{\mathbf{z}})$ and $\bar{\mathbf{z}} \sim p(\bar{\mathbf{z}}) = \mathcal{U}(1/|\mathbf{Z}|)$. Subgroups are sorted based on $\hat{p}(\mathbf{z})$, where $\hat{p}(\mathbf{z}|\mathbf{x})$ is estimated with a zero-shot CLIP classifier. Error bars are 95% confidence intervals via the Clopper-Pearson exact method assuming Bernoulli experiments (success probability $1/|\mathbf{Z}| = 0.01$).

$\mathbf{Z}_1\}$, and $T_2 = \{\text{“An image of a } \textit{type”} | \textit{type} \in \mathbf{Z}_2\}$. We use Stable Diffusion (SD) [43] using 20 steps with the DPM-Solver++ [33, 34] as realization of $p^{T2I}(\mathbf{x}|t)$ and sample \mathbf{x} of resolution 512×512 . For unconditional synthesis $\hat{p}(\mathbf{x})$, we use the prompt “An image of a car.”. For conditional synthesis $\hat{p}(\mathbf{x}|\bar{\mathbf{z}})$, we use the prompt template “An image of a *color type* car with a *background* background.” where we insert every $\bar{\mathbf{z}} = (\textit{color}, \textit{background}, \textit{type})$ equally often (round-robin). For both variants, we employ a Monte-Carlo estimate of $\hat{p}(\mathbf{z})$ based upon 40,000 samples.

Results. Results are summarized in Figure 3: conditional synthesis generates samples \mathbf{x} such that the 95% confidence interval lower bound on $\hat{p}(\mathbf{z})$ is greater than 0.005 for all subgroups (the uniform subgroup probability is $p = 1/|\mathbf{Z}| = 0.01$). In contrast, unconditional synthesis has a minimum 95% confidence interval upper bound on $\hat{p}(\mathbf{z})$ at less than 0.0001. For conditional synthesis, we also estimate $\hat{p}(\mathbf{z}^* = \bar{\mathbf{z}})$ for $\mathbf{x} \sim \hat{p}(\mathbf{x}|\bar{\mathbf{z}})$ and $\mathbf{z}^* = \arg \max_{\mathbf{z} \in \mathbf{Z}} \hat{p}(\mathbf{z}|\mathbf{x})$. Averaged over all subgroups $\bar{\mathbf{z}} \in \mathbf{Z}$, this probability is approximately 89%, and for no subgroup it is less than 85%.

4.2. Zero-Shot Systematic Error Benchmark

Motivation. One major challenge when evaluating approaches for systematic error identification empirically is that a priori, it is unknown which (if any) systematic errors a target classifier exhibits. Not having such a ground truth prohibits the fully automated evaluation of systematic error identification approaches. Training time interventions to inject systematic errors into models [13] are computationally costly and not scalable. Moreover, their indirect nature makes them brittle and not always resulting in the desired error. To address these shortcomings, we propose *zero-shot systematic errors* where we leverage zero-shot classifiers

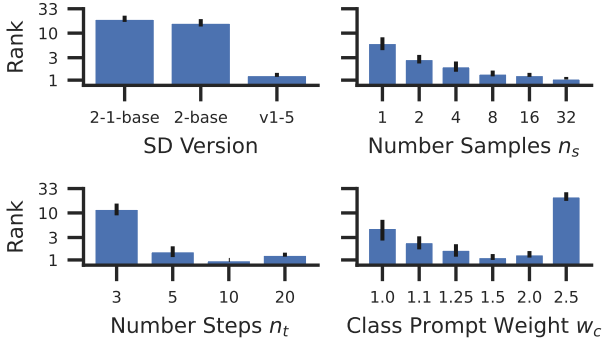


Figure 4. Effect of PROMPTATTACK’s hyperparameters on identification of systematic errors injected into a zero-shot classifier, quantified in rank (log-scale) assigned to the ground-truth systematic error subgroup (a lower rank corresponds to a higher error).

from multimodal image-text models such as CLIP [41]¹.

More specifically, let us consider a binary² classifier for classes \mathbf{y}_a and \mathbf{y}_b , which can be constructed based upon the text queries $t_k = \text{“an image of a } \mathbf{y}_k \text{”}$ ($k \in \{a, b\}$). For this, we compute the cosine similarities ϕ_k between the input’s CLIP image embedding and the CLIP text embeddings of the t_k and set the prediction logits to $l_k = \tau \phi_k$ with $\tau = 100$ being a temperature. We now inject a systematic error into this zero-shot classifier using additional *poisonous* queries like $t_a^{p1} = \text{“An image of a red } \mathbf{y}_a \text{”}$, $t_a^{p2} = \text{“An image of a } \mathbf{y}_a \text{ with forest background”}$, and so on. We take the minimum of the cosine similarities $\phi_a^p = \min(\phi_a^{p1}, \phi_a^{p2}, \dots)$ to the encoded poisonous queries, set $l_a^p = \tau \phi_a^p$, and compute the post-softmax probabilities as $\hat{\mathbf{y}}_a, \hat{\mathbf{y}}_b, \hat{\mathbf{y}}_a^p = \text{softmax}(l_a, l_b, l_a^p)$. We now classify to class \mathbf{y}_b if $\hat{\mathbf{y}}_b + \hat{\mathbf{y}}_a^p > \hat{\mathbf{y}}_a$ and \mathbf{y}_a otherwise, that is: samples of class \mathbf{y}_a which have high cosine similarity to all poisonous queries are more likely misclassified as \mathbf{y}_b . This construction comes essentially for free and poisonous query can control the decision rule much more directly than training time interventions [13], since they operate at inference time.

Experimental Setting. We define a zero-shot classifier as described above for $\mathbf{y}_a = \text{“car”}$ and $\mathbf{y}_b = \text{“truck”}$. We consider the same class $\tilde{\mathbf{y}} = \text{“car”}$, operational design domain \mathbf{Z} , and semantic dimensions $\mathbf{Z}_{\{0,1,2\}}$ as in Section 4.1. We sample 20 combinations of $color \in \mathbf{Z}_0$, $background \in \mathbf{Z}_1$, and $type \in \mathbf{Z}_2$. For each of these combinations, we employ the poisonous queries $t_a^{p1} = \text{“An image of a } color \text{ car”}$, $t_a^{p2} = \text{“An image of a car with } background \text{ background”}$, and $t_a^{p3} = \text{“An image of a } type \text{”}$, and test PROMPTATTACK on the resulting zero-shot classifiers. Per Section 3.3, we use $R_f(\mathbf{z}) \approx \text{median}_{i=0}^{n_s-1} (1 - f(\tilde{\mathbf{y}}|\mathbf{x}_i))$ for $\mathbf{x}_i \sim p^{T2I}(\mathbf{x}|T_p(\tilde{\mathbf{y}}, \mathbf{z}))$. We rank $\mathbf{z} \in \mathbf{Z}$ on descending

¹Note we are not interested in potential systematic errors of the CLIP image or text encoder (if such exist, they are nuisances) but rather in systematic errors of the constructed zero-shot classifier.

²An extension to more than two classes would be straightforward.

$R_f(\mathbf{z})$: the \mathbf{z} with the highest $R_f(\mathbf{z})$ obtains rank 1.

We evaluate the sensitivity of PROMPTATTACK with respect to its free hyperparameters. We set the prompt template³ to $T_p = \text{“An image of a } color \text{ type (car:} w_c \text{) with a } background \text{ background.”}$, where w_c is a weight multiplied to the text encoding of the tokens of the word “car”. For PROMPTATTACK, we generate n_s samples \mathbf{x} of size 512×512 with n_t steps of SD/DPMSolver++. We investigate the effect of the number of image samples n_s , number of inference steps n_t , the custom class prompt weight w_c , as well as different versions of SD on the resulting ranking.

Results. Figure 4 summarizes the results for PROMPTATTACK (see Section C for samples). We observe that the version of SD has a major impact on PROMPTATTACK’s performance, with the v1.5 checkpoint greatly outperforming the more recent v2-base and v2-1-base checkpoints. We attribute this to v1.5 generating samples more faithful to the prompt and with better attribute binding than the other checkpoints (see Figure 10). In the upper right plot, we observe that too few samples n_s impair performance due to the large variance in the MC estimate of $R_f(\mathbf{z})$. For $n_s \geq 16$ performance is very close to optimal. In the bottom left, we see that a small number of inference steps such as $n_t = 5$ steps works well — this is somewhat surprising since image quality is impaired considerably for this small n_t but PROMPTATTACK is relatively robust to image quality. Lastly, in the bottom right, the importance of a custom class prompt weight w_c is demonstrated with $w_c = 1.5$ outperforming the default of $w_c = 1.0$. We attribute this to higher weights resulting in more reliably depicting objects of the desired class. Increasing w_c further results in deteriorating performance due to image artefacts. Based on these results, we use SD v1.5 with $n_s = 16$, $n_t = 20$, and $w_c = 1.5$ in subsequent experiments without further tuning.

5. Qualitative Evaluation on ImageNet

Vehicle Experiment. We evaluate 5 models trained for image classification on ImageNet1k. We focus on a subset of classes belonging to the vehicle subcategory, more specifically on misclassifying samples of the class “minivan” $\tilde{\mathbf{y}} = \mathbf{y}_{minivan}$ into other classes that have a distance of 2 in the WordNet [14] hierarchy, e.g., “police van” and “snowplow”. We focus on an operational design domain \mathbf{Z} with five semantic dimensions, corresponding to *viewpoint*, *object size*, *object color*, *weather*, and *background*. We use the prompt template $T_p = \text{“}\{viewpoint\} \text{ view of } \{size\} \{color\} \text{ (minivan:1.5) in front of } \{weather\} \{background\} \text{”}$. We use combinatorial testing with $n_C = 3$, exploring $|\mathbf{Z}_C| = 1,230$ out of $|\mathbf{Z}| = 18,720$ subgroups, and generate $n_s = 16$ image samples per subgroup using

³Note the difference of poisonous queries $t_a^{p_i}$ (part of the poisoned zero-shot classifier) and prompt template T_p (part of PROMPTATTACK).




Figure 5. Median (over 16 samples) target class confidence for strongest respective prompt found by PROMPTATTACK vs. a neutral baseline prompt (black boundary) for four selected target classes. We refer to Figure 1 for exemplary prompts, samples, and class prediction rates.

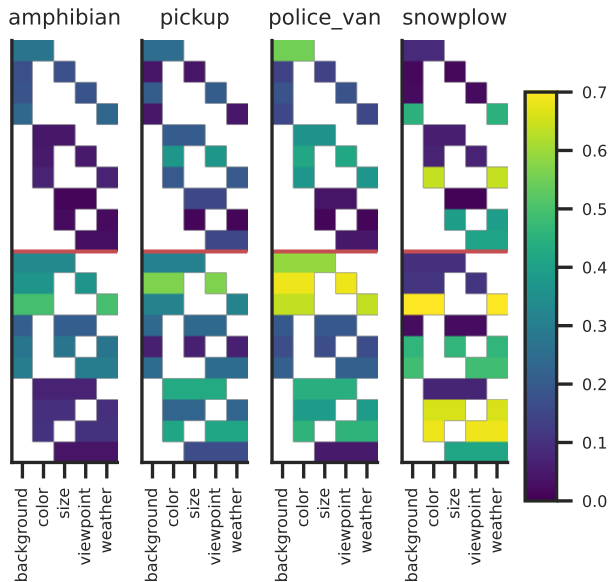


Figure 6. Cumulative functional ANOVA [24] of predicted probability of target classes for source class “minivan”. Rows correspond to cardinality 2 and 3 subsets of semantic dimensions with white encoding an excluded dimension and the color the groups’s fANOVA score. Different dimensions are relevant for different target classes, e.g., the combination of background and object color has high score for police-van but low score for snowplow. High fANOVA scores require at least 3 dimensions.

Stable Diffusion v1.5. For a full description of experimental setting, we refer to Section B.1.

We analyse the median predicted probability $\text{median}_{p \in \mathcal{P}(T_p(\tilde{y}, z))} f(y^{(t)} | \mathbf{x})$ of different target classes $\mathbf{y}^{(t)}$ for the strongest prompts identified by PROMPTATTACK. We compare these prompts to a neutral baseline prompt $t = \text{“center view of (minivan:1.5) in front of background.”}$. Results for 4 selected target classes $\mathbf{y}^{(t)}$ are summarized in Figure 5 (see also Table 1). It can be seen that samples of the baseline prompt are assigned with very high confidence to the correct class “minivan”. However, for target classes $\mathbf{y}^{(t)}$ such as “pickup”, “police van”, or “snowplow”, PROMPTATTACK can identify prompts that result in systematic misclassifications, that is considerably increased value for the target class. We depict the

top-ranked subgroups for three target classes in Figure 1. We note that the sensitivity of models to these subgroups vary (in accordance with Figure 5): for $t_{snowplow}^* = \text{“rear view of small orange minivan in front of snowy forest.”}$, a VGG16 [47] misclassifies 25% of the samples as snowplows while a ConvNeXt-B [32] misclassifies only 1%. This indirectly confirms that misclassifications are not due to OOC samples because the same samples are classified correctly by a ConvNeXt-B (see also Section D for an illustration of misclassified samples). Moreover, we selected 16 images from LAION-5B [46] that best match $t_{snowplow}^*$ (using CLIP retrieval [5] followed by manual filtering). The models misclassify between 6 (ConvNeXt-B) and 8 (ResNet50 [21], VGG16) of those as snowplows.

Figure 6 depicts a (cumulative) functional ANOVA analysis [24] of median predicted probability of different target classes $\mathbf{y}^{(t)}$. One can see that different semantic dimensions are relevant for different target classes; for instance, the combination of background and object color has a high fANOVA score for police-van but a low score for snowplow. Moreover, for target classes like snowplow, at least 3 semantic dimensions are required for explaining the bulk of the variance. This is also illustrated in Figure 7 where changing a single dimension does not mislead a ViT-B/16 [12], while a specific combination of shifts such as $t_{snowplow}^*$ results in misclassifying 300 out of 1000 samples as snowplow. A possible explanation for this increased error rate is that (i) snowy forests are more often in the background for snowplows than minivans, (ii) snowplows are more often orange than minivans, and (iii) rear views hide a distinctive feature of snowplows, namely their plow in the front. In summary, our findings support the hypothesis that studying single shifts can be insufficient as often specific combinations of compounding shifts result in a systematic error. PROMPTATTACK allows finding such rare subgroups.

Person Experiment. It has been observed before that systematic errors on under-represented demographic subgroups can result in reduced fairness and even derogatory behaviour such as misclassifying black people as “gorillas” [3]. We check two models trained for image classification on ImageNet21k for similar issues using PROMP-

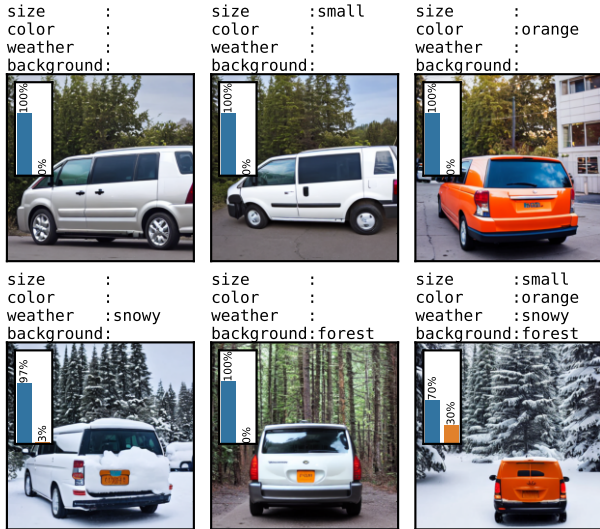


Figure 7. Samples for prompt template “rear view of {size} {color} minivan in front of {weather} {background}”. Baseline generations (top left) and single-dimension-shifted generations are classified consistently as **minivan** by a ViT-B/16. Shifting all dimensions jointly as determined by PROMPTATTACK results in 300 out of 1000 samples (mis-)classified as **snowplows** (bottom right).




Figure 8. Samples and prediction histograms (based on 1000 samples) for different subgroups. The baseline subgroup (left) is classified consistently as **homo**, while the misclassification rate to **ape** is significantly increased for an MLP-Mixer-B/16 [49] on a subgroup identified by PROMPTATTACK (right).

TATTACK: we set source class \tilde{y} = “homo” and target class “ape”. We use an operational design domain \mathbf{Z} with the 5 semantic dimensions *age*, *gender*, geographic *region*, *hairtype*, and *background*. We use the prompt template T_p = “A {age} {gender} {region} (person:1.5) with {hairtype} hairs in front of {background}”. We use combinatorial testing with $n_C = 3$, exploring $|\mathbf{Z}_C| = 1,371$ out of $|\mathbf{Z}| = 12,150$ subgroups. For a full description of experimental setting and additional results, we refer to Section B.2. The outcome is summarized in Figure 8 and Table 2: while samples of most subgroups are classified correctly as “homo”, samples

of specific subgroups such as t = “old male african (person:1.5) with long hairs” have a misclassification rate of up to 25% into “ape” (see Figure 16 in appendix). PROMPTATTACK allows identifying such systematic errors on under-represented demographic subgroups.

6. Limitations

False Positive Systematic Errors. While prompt engineering can reduce the number of OOC samples and robust estimation can reduce their impact, there may still be combinations of subgroups \mathbf{z} and source classes \tilde{y} where the majority of samples are OOC. If the OOC sample $\mathbf{x} \sim \hat{p}(\mathbf{x}|\tilde{y}, \mathbf{z})$ is such that the true $p(\mathbf{y}|\mathbf{x})$ is not strongly peaked at \tilde{y} , then our procedure might identify a false positive systematic error: the classifier f might classify \mathbf{x} correctly as not belonging to \tilde{y} because the generated \mathbf{x} is no instance of \tilde{y} . In the absence of an oracle providing us the true $p(\mathbf{y}|\mathbf{x})$ (such as a human in the loop [17]), there is no reliable way of identifying these false positives. However, we note that moderate prompt engineering such as tuning the class prompt weight was sufficient to prevent such false positives for the operational design domains we have considered (see also Sections 4.2 and 5).

Language Bottleneck. Certain coherent subsets of the data, e.g., subsets that share some geometric layout, may be difficult to describe in natural language as a text prompt. Increased errors on such subsets can thus not be identified directly by our procedure. Future work on using other types of conditioning information \mathbf{z} such as a scene layout [56] could address this limitation. Moreover, textual inversion [15] can be used to distill visual concepts into tokens, e.g., ImageNet classes \tilde{y} with ambiguous class names [51].

Bias Propagation. We note that biases from the text-to-image models itself may propagate to biases in our systematic error identification procedure: if the text-to-image model cannot generate samples for certain marginalized subgroups of a population, we will not be able to identify a potentially subpar performance of the downstream classifier f on these subgroups. This reinforces the need to further reduce bias in text-to-image models in the future [7].

7. Conclusion

We have proposed PROMPTATTACK, which leverages recent progress on text-to-image models for identifying systematic errors that occur on rare data subgroups (combinations of semantic shifts). Both quantitative results on carefully constructed benchmarks as well as qualitative results on multi-class image classifiers demonstrate the efficacy of PROMPTATTACK in identifying such systematic errors. Future work needs to address the limitations discussed above, for instance by leveraging more controllable, versatile, and reliable procedures for image synthesis.

References

- [1] allpairs.py. <https://github.com/thombashi/allpairs.py>. License: MIT. 12
- [2] Auditing machine learning algorithms. <https://www.auditingalgorithms.net/>. Accessed: 2023-01-10. 2
- [3] Google photos labelled a picture of two black people as gorillas. <https://www.theguardian.com/technology/2015/jul/01/google-sorry-racist-auto-tag-photo-app>. Accessed: 2023-03-06. 7
- [4] Bestoun S. Ahmed, Kamal Z. Zamli, Wasif Afzal, and Miroslav Bures. Constrained interaction testing: A systematic literature study. *IEEE Access*, 5:25706–25730, 2017. 2, 5, 11
- [5] Romain Beaumont. Clip retrieval: Easily compute clip embeddings and build a clip retrieval system with them. <https://github.com/rom1504/clip-retrieval>. 7
- [6] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *ECCV*, 2018. 1
- [7] Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, and Aylin Caliskan. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. *arXiv:2211.03759*, 2022. 8
- [8] Frederik Blank, Fabian Hüger, Michael Mock, and Thomas Stauner. Assurance methodology for in-vehicle AI. *ATZ worldwide*, 124:54–59, 07 2022. 1
- [9] Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T. Freeman, Michael Rubinstein, Yuanzhen Li, and Dilip Krishnan. Muse: Text-to-image generation via masked generative transformers. *arXiv:2301.00704*, 2023. 2, 3
- [10] Krzysztof Czarnecki. Operational design domain for automated driving systems - taxonomy of basic terms, 2018. 1, 4
- [11] 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 *FACCT*, 2022. 3
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 7, 11, 19
- [13] Sabri Eyuboglu, Maya Varma, Khaled Kamal Saab, Jean-Benoit Delbrouck, Christopher Lee-Messer, Jared Dunnmon, James Zou, and Christopher Re. Domino: Discovering Systematic Errors with Cross-Modal Embeddings. In *ICLR*, 2022. 2, 3, 5, 6
- [14] Christiane Fellbaum. *WordNet: An Electronic Lexical Database*. Bradford Books, 1998. 6, 11
- [15] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H. Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv:2208.01618*, 2022. 8
- [16] 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:2209.07858*, 2022. 2
- [17] Irena Gao, Gabriel Ilharco, Scott Lundberg, and Marco Tulio Ribeiro. Adaptive testing of computer vision models. *arXiv:2212.02774*, 2022. 2, 3, 8
- [18] R. Geirhos, J.-H. Jacobsen, C. Michaelis, R. Zemel, W. Brendel, M. Bethge, and F. A. Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020. 1
- [19] Christoph Gladisch, Christian Heinzemann, Martin Herrmann, and Matthias Woehrle. Leveraging combinatorial testing for safety-critical computer vision datasets. In *CVPR Workshops*, 2020. 2
- [20] Sorin Grigorescu, Bogdan Trasnea, Tiberiu Cocias, and Gigel Macesanu. A survey of deep learning techniques for autonomous driving. *Journal of Field Robotics*, 37(3):362–386, 2020. 1
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 7, 11
- [22] 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. In *ICCV*, 2021. 2
- [23] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *ICLR*, 2019. 2
- [24] Frank Hutter, Holger Hoos, and Kevin Leyton-Brown. An efficient approach for assessing hyperparameter importance. In *ICML*, 2014. 7
- [25] Badr Youbi Idrissi, Diane Bouchacourt, Randall Balestriero, Ivan Evtimov, Caner Hazirbas, Nicolas Ballas, Pascal Vincent, Michal Drozdal, David Lopez-Paz, and Mark Ibrahim. Imagenet-x: Understanding model mistakes with factor of variation annotations. *arXiv:2211.01866*, 2022. 2
- [26] Saachi Jain, Hannah Lawrence, Ankur Moitra, and Alexander Madry. Distilling Model Failures as Directions in Latent Space. *arXiv:2206.14754*, 2022. 2, 3
- [27] Priyatham Kattakinda, Alexander Levine, and Soheil Feizi. Invariant learning via diffusion dreamed distribution shifts. *arXiv:2211.10370*, 2022. 3
- [28] Michael P. Kim, Amirata Ghorbani, and James Zou. Multi-accuracy: Black-box post-processing for fairness in classification. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’19, pages 247–254, New York, NY, USA, 2019. Association for Computing Machinery. 2
- [29] Zhiheng Li, Anthony Hoogs, and Chenliang Xu. Discover and mitigate unknown biases with debiasing alternate networks. In *ECCV*, 2022. 1

- [30] Zhiheng Li and Chenliang Xu. Discover the unknown biased attribute of an image classifier. In *ICCV*, 2021. 1
- [31] Li Liu, Wanli Ouyang, Xiaogang Wang, Paul Fieguth, Jie Chen, Xinwang Liu, and Matti Pietikäinen. Deep Learning for Generic Object Detection: A Survey. *IJCV*, 128(2):261–318, Feb. 2020. 1
- [32] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *CVPR*, 2022. 7, 11
- [33] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-solver: A fast ODE solver for diffusion probabilistic model sampling in around 10 steps. In *NeurIPS*, 2022. 5, 12, 16
- [34] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv:2211.01095*, 2022. 5, 12, 16
- [35] Aengus Lynch, Jean Kaddour, and Ricardo Silva. Evaluating the impact of geometric and statistical skews on out-of-distribution generalization performance. In *NeurIPS Workshops*, 2022. 3
- [36] TorchVision maintainers and contributors. Torchvision: Pytorch’s computer vision library. <https://github.com/pytorch/vision>. 11
- [37] Shervin Minaee, Yuri Boykov, Fatih Porikli, Antonio Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image segmentation using deep learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7):3523–3542, 2022. 1
- [38] Changhai Nie and Hareton Leung. A survey of combinatorial testing. *ACM Computing Surveys*, 43:11, 01 2011. 2, 5, 11
- [39] Ethan Perez, Saffron Huang, H. Francis Song, Trevor Cai, Roman Ring, John Aslanides, Amelia Glaese, Nat McAleese, and Geoffrey Irving. Red teaming language models with language models. *arXiv:2202.03286*, 2022. 2
- [40] Oskar Pfungst and Carl Leo. Rahn. *Clever Hans (the horse of Mr. Von Osten) a contribution to experimental animal and human psychology*. New York, H. Holt and company, 1911. 1
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 3, 5, 6
- [42] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *arXiv:2204.06125*, 2022. 2, 3
- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3, 4, 5, 12, 15
- [44] Shiori Sagawa*, Pang Wei Koh*, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks. In *ICLR*, 2020. 2
- [45] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo-Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*, 2022. 2, 3
- [46] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5b: An open large-scale dataset for training next generation image-text models. In *NeurIPS Datasets and Benchmarks Track*, 2022. 3, 7
- [47] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 7, 11
- [48] Nimit Sohoni, Jared A. Dunnmon, Geoffrey Angus, Albert Gu, and Christopher Ré. No subclass left behind: Fine-grained robustness in coarse-grained classification problems. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NeurIPS’20, Red Hook, NY, USA, 2020. Curran Associates Inc. 2
- [49] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision. *arXiv:2105.01601*, 2021. 8, 12, 14, 21
- [50] Antonio Torralba and Alexei A. Efros. Unbiased look at dataset bias. In *CVPR*, 2011. 1
- [51] Joshua Vendrow, Saachi Jain, Logan Engstrom, and Aleksander Madry. Dataset interfaces: Diagnosing model failures using controllable counterfactual generation. *arXiv:2302.07865*, 2023. 2, 3, 8
- [52] Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>. 12
- [53] Olivia Wiles, Isabela Albuquerque, and Sven Gowal. Discovering bugs in vision models using off-the-shelf image generation and captioning. In *NeurIPS Workshops*, 2022. 2, 3, 4, 5
- [54] Kai Yuanqing Xiao, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. Noise or signal: The role of image backgrounds in object recognition. In *ICLR*, 2021. 2
- [55] Oliver Zendel, Katrin Honauer, Markus Murschitz, Daniel Steiner, and Gustavo Fernandez Dominguez. Wilddash - creating hazard-aware benchmarks. In *ECCV*, 2018. 2
- [56] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *arXiv:2302.05543*, 2023. 8
- [57] Chunting Zhou, Xuezhe Ma, Paul Michel, and Graham Neubig. Examining and combating spurious features under distribution shift. In *ICML*, 2021. 1