
Fundamental Tradeoffs between Invariance and Sensitivity to Adversarial Perturbations

Florian Tramèr¹ Jens Behrmann² Nicholas Carlini³ Nicolas Papernot³ Jörn-Henrik Jacobsen⁴

Abstract

Adversarial examples are malicious inputs crafted to induce misclassification. Commonly studied *sensitivity-based* adversarial examples introduce semantically-small changes to an input that result in a different model prediction. This paper studies a complementary failure mode, *invariance-based* adversarial examples, that introduce minimal semantic changes that modify an input’s true label yet preserve the model’s prediction. We demonstrate fundamental tradeoffs between these two types of adversarial examples. We show that defenses against sensitivity-based attacks actively harm a model’s accuracy on invariance-based attacks, and that new approaches are needed to resist both attack types. In particular, we break state-of-the-art adversarially-trained and *certifiably-robust* models by generating small perturbations that the models are (provably) robust to, yet that change an input’s class according to human labelers. Finally, we formally show that the existence of excessively invariant classifiers arises from the presence of *overly-robust* predictive features in standard datasets.

1. Introduction

Research on adversarial examples (Szegedy et al., 2013; Biggio et al., 2013) is motivated by a spectrum of questions, ranging from the security of models in adversarial settings (Tramèr et al., 2019), to limitations of learned representations under natural distribution shift (Gilmer et al., 2018). The broadest accepted definition of an adversarial example is “an input to a ML model that is intentionally designed by an attacker to fool the model into producing an incorrect output” (Goodfellow & Papernot, 2017).

¹Stanford University ²University of Bremen ³Google Brain ⁴Vector Institute and University of Toronto. Correspondence to: Florian Tramèr <tramer@cs.stanford.edu>.

Adversarial examples are commonly formalized as inputs obtained by adding some perturbation to test examples to change the model output. We refer to this class of malicious inputs as *sensitivity-based adversarial examples*.

To enable concrete progress, the adversary’s capabilities are typically constrained by bounding the size of the perturbation added to the original input. The goal of this constraint is to ensure that semantics of the input are left unaffected by the perturbation. In the computer vision domain, ℓ_p norms have grown to be a default metric to measure semantic similarity. This led to a series of proposals for increasing the robustness of models to sensitivity-based adversaries that operate within the constraints of an ℓ_p ball (Madry et al., 2017; Wong & Kolter, 2018; Raghunathan et al., 2018).

In this paper, we show that optimizing a model’s robustness to ℓ_p -bounded perturbations is not only insufficient to resist general adversarial examples, but also potentially *harmful*. As ℓ_p distances are a crude approximation to the visual similarity in a given task, over-optimizing a model’s robustness to ℓ_p -bounded perturbations renders the model excessively invariant to real semantics of the underlying task.

Excessive invariance of a model causes vulnerability against *invariance adversarial examples* (Jacobsen et al., 2019). These are perturbations that change the human-assigned label of a given input but keep the model prediction unchanged. For example in Figure 1 an MNIST image of a digit ‘3’ is perturbed to be an image of a ‘5’ by changing only 20 pixels; models that are excessively invariant do not change their decision and incorrectly label both images as a ‘3’, despite the fact that the oracle label has changed.

This paper exposes a fundamental tradeoff between sensitivity-based and invariance-based adversarial examples. We show that due to a misalignment between formal robustness notions (e.g., ℓ_p -balls) and a task’s perceptual metric, current defenses against adversarial examples cannot prevent both sensitivity-based and invariance-based attacks, and must trade-off robustness to each (see Figure 2). Worse, we find that increasing robustness to sensitivity-based attacks **decreases** a model’s robustness to invariance-based attacks. We introduce new algorithms to craft ℓ_p -bounded invariance-based adversarial examples, and illustrate the

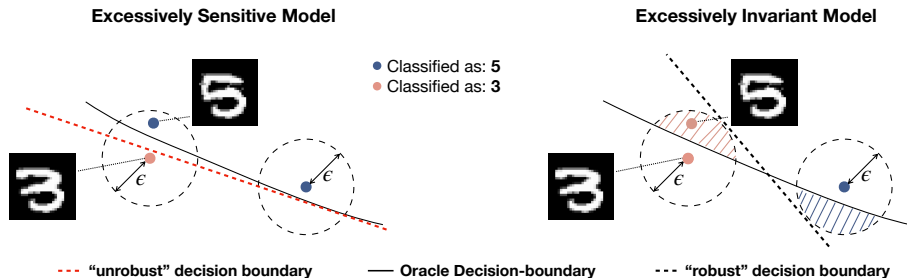


Figure 1: Decision boundaries near a real image of a digit ‘3’ and an invariance-based adversarial example labeled as ‘5’. [Left]: Training a classifier without constraints may learn a decision boundary unrobust to sensitivity-based adversarial examples. [Right]: Enforcing robustness to norm-bounded perturbations introduces erroneous invariance (dashed regions in ϵ -spheres). We display real data here, the misclassified ‘5’ is an image found by our attack which resides within a typically reported ϵ -region around the displayed ‘3’ (in the ℓ_0 norm). This excessive invariance of the robust model in task-relevant directions illustrates how robustness to sensitivity-based adversarial examples can result in new model vulnerabilities.

above tradeoff on MNIST.¹ We show that state-of-the-art robust models disagree with human labelers on many of our crafted invariance-based examples, and that the disagreement rate is higher the more robust a model is. We find that even models robust to very small perturbations (e.g., of ℓ_∞ -norm below $\epsilon = 0.1$) have higher vulnerability to invariance attacks compared to undefended models.

We further break a *provably-robust* defense (Zhang et al., 2019) with our attack. This model is certified to have 87% test-accuracy (with respect to the MNIST test-labels) under ℓ_∞ -perturbations of radius $\epsilon = 0.4$. That is, for 87% of test inputs (x, y) , the model is guaranteed to predict class y for any perturbed input x' that satisfies $\|x - x'\|_\infty \leq 0.4$. Yet, on our invariance-based adversarial examples that satisfy this norm-bound, the model only agrees with human labelers in 60% of the cases for an automated attack, and 12% of the cases for manually-created examples—i.e., no better than chance. The reason is that we can find perturbed inputs x' that humans no longer classify the same way as x .

Code to reproduce our attacks is available at <https://github.com/ftramer/Excessive-Invariance>.

Finally, we introduce a classification task where the tradeoff between sensitivity and invariance can be studied rigorously. We show that excessive sensitivity and invariance can be tied respectively to the existence of generalizable non-robust features (Jo & Bengio, 2017; Ilyas et al., 2019; Yin et al., 2019) and to *robust features* that are predictive for standard datasets, but not for the general vision tasks that these datasets aim to capture. Our experiments on MNIST show that such *overly-robust* features exist. We further argue both formally and empirically that data augmentation may offer a solution to both excessive sensitivity and invariance.

¹While MNIST can be a poor choice for studying adversarial examples, we chose it because it is the *only* vision task for which models have been made robust in non-negligible ℓ_p norm balls. The fundamental tradeoff described in this paper will affect other vision tasks once we can train strongly robust models on them.

2. Norm-bounded Sensitivity and Invariance Adversarial examples

We begin by defining a framework to formally describe two complementary failure modes of machine learning models, namely (norm-bounded) adversarial examples that arise from excessive *sensitivity* or *invariance* of a classifier.

We consider a classification task with data $(x, y) \in \mathbb{R}^d \times \{1, \dots, C\}$ from a distribution \mathcal{D} . We assume the existence of a *labeling oracle* $\mathcal{O} : \mathbb{R}^d \rightarrow \{1, \dots, C\} \cup \{\perp\}$ that maps any input in \mathbb{R}^d to its true label, or to the “garbage class” \perp for inputs x considered “un-labelable” (e.g., for a digit classification task, the oracle \mathcal{O} corresponds to human-labeling of any image as a digit or as the garbage class). Note that for $(x, y) \sim \mathcal{D}$, we always have $y = \mathcal{O}(x)$.²

The goal of robust classification is to learn a classifier $f : \mathbb{R}^d \rightarrow \{1, \dots, C\}$ that agrees with the oracle’s labels not only in expectation over the distribution \mathcal{D} , but also on any rare or out-of-distribution inputs to which the oracle assigns a class label—including adversarial examples obtained by imperceptibly perturbing inputs sampled from \mathcal{D} .

At its broadest, the definition of an adversarial example encompasses any adversarially induced failure in a classifier (Goodfellow & Papernot, 2017). That is, an adversarial example is any input x^* created such that $f(x^*) \neq \mathcal{O}(x^*)$. This definition has proven difficult to work with, due to its inherent reliance on the oracle \mathcal{O} . As a result, it has become customary to study a relaxation of this definition, which restricts the adversary to applying a “small” perturbation to an input x sampled from the distribution \mathcal{D} . A common choice

²We view the support of \mathcal{D} as a strict subset of all inputs in \mathbb{R}^d to which the oracle assigns a label. That is, there are inputs for which humans agree on a label, yet that have measure zero in the data distribution from which the classifier’s train and test inputs are chosen. For example, the train-test data is often a sanitized and normalized subset of natural inputs. Moreover, “unnatural” inputs such as adversarial examples might never arise in natural data.

is to restrict the adversary to small perturbations under some ℓ_p -norm. We call these “sensitivity adversarial examples”:

Definition 1 (Sensitivity Adversarial Examples). *Given a classifier f and a correctly classified input $(x, y) \sim \mathcal{D}$ (i.e., $\mathcal{O}(x) = f(x) = y$), an ε -bounded sensitivity adversarial example is an input $x^* \in \mathbb{R}^d$ such that:*

1. $f(x^*) \neq f(x)$.
2. $\|x^* - x\| \leq \varepsilon$.

The assumption underlying this definition is that perturbations satisfying $\|x^* - x\| \leq \varepsilon$ preserve the oracle’s labeling of the original input x , i.e., $\mathcal{O}(x^*) = \mathcal{O}(x)$.

A long line of work studies techniques to make classifiers robust to norm-bounded sensitivity adversarial examples (Goodfellow et al., 2014; Madry et al., 2017). The main objective is to minimize a classifier’s *robust error* under ε -bounded perturbations, which is defined as:

$$\mathcal{L}_\varepsilon(f) = \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\max_{\|\Delta\| \leq \varepsilon} \{f(x + \Delta) \neq y\} \right]. \quad (1)$$

We study a complementary failure mode to sensitivity adversarial examples, called invariance adversarial examples (Jacobsen et al., 2019). These correspond to (bounded) perturbations that *do not* preserve an input’s oracle-assigned label, yet preserve the model’s classification:

Definition 2 (Invariance Adversarial Examples). *Given a classifier f and a correctly classified input $(x, y) \sim \mathcal{D}$, an ε -bounded invariance adversarial example is an input $x^* \in \mathbb{R}^d$ such that:*

1. $f(x^*) = f(x)$.
2. $\mathcal{O}(x^*) \neq \mathcal{O}(x)$, and $\mathcal{O}(x^*) \neq \perp$.
3. $\|x^* - x\| \leq \varepsilon$.

If the assumption on sensitivity adversarial examples in Definition 1 is met—i.e., all ε -bounded perturbations preserve the label—then Definition 1 and Definition 2 correspond to well-separated failure modes of a classifier (i.e., ε' -bounded invariance adversarial examples only exist for $\varepsilon' > \varepsilon$).

Our main contribution is to reveal fundamental trade-offs between these two types of adversarial examples, that arise from this assumption being violated. We demonstrate that state-of-the-art robust classifiers do violate this assumption, and (sometimes certifiably) have low robust error $\mathcal{L}_\varepsilon(f)$ for a norm-bound ε that does not guarantee that the oracle’s label is preserved. We show that these classifiers actually have high “true” robust error as measured by human labelers.

Remarks. Definition 2 is a conscious restriction on a definition of Jacobsen et al. (2019), who define an invariance adversarial example as an *unbounded* perturbation that changes the oracle’s label while preserving a classifier’s




Figure 2: Illustration of distance-oracle misalignment. The input space is (ground-truth) classified into the red solid region, and the white dotted region. (a) A point at distance ε^* (under a chosen norm) of the oracle decision boundary. (b) A model robust to perturbations of norm $\varepsilon \leq \varepsilon^*$ (gray circle) is still overly sensitive and can have adversarial examples x^* . (c) A model robust to perturbations of norm $\varepsilon > \varepsilon^*$ (gray circle) has invariance adversarial examples x^* .

output at *an intermediate feature layer*. As we solely consider the model’s final classification, considering unbounded perturbations would allow for a “trivial” attack: given an input x of class y , find any input of a different class that the model misclassifies as y . (e.g., given an image of a digit 8, an unbounded invariance example could be any unperturbed digit that the classifier happens to misclassify as an 8).

Definition 2 presents the same difficulty as the original broad definition of adversarial examples: a dependence on the oracle \mathcal{O} . Automating the process of finding invariance adversarial examples is thus challenging. In Section 4.2, we present some successful automated attacks, but show that a human-in-the-loop process is more effective.

3. The Sensitivity and Invariance Tradeoff

In this section, we show that if the norm that is used to define “small” adversarial perturbations is misaligned with the labeling oracle \mathcal{O} , then the robust classification objective in equation 1 is insufficient for preventing both sensitivity-based and invariance-based adversarial examples under that norm. That is, we show that optimizing a model to attain low robust error on perturbations of norm ε cannot prevent both sensitivity and invariance adversarial examples.

We begin by formalizing our notion of norm-oracle misalignment. The definition applies to any similarity metric over inputs, of which ℓ_p -norms are a special case.

Definition 3 (Distance-Oracle Misalignment). *Let $dist : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a distance measure (e.g., $\|x_1 - x_2\|$). We say that $dist$ is aligned with the oracle \mathcal{O} if for any input x with $\mathcal{O}(x) = y$, and any inputs x_1, x_2 such that $\mathcal{O}(x_1) = y$, $\mathcal{O}(x_2) \neq y$, we have $dist(x, x_1) < dist(x, x_2)$. $dist$ and \mathcal{O} are misaligned if they are not aligned.*

For natural images, ℓ_p -norms (or other simple metrics) are clearly misaligned with our own perceptual metric. A concrete example is in Figure 3. This simple fact has deep




Figure 3: Visualization that ℓ_p norms fail to measure semantic similarity in images. (a) original image in the ImageNet validation set labeled as a *goldfinch* (top), *hermit crab* (bottom); (b) semantic perturbation with a ℓ_2 perturbation of 19 (respectively 22) that replaces the object of interest with a pineapple (top), strawberry (bottom). (c) random perturbation of the same ℓ_2 -norm.

implications for the suitability of the robust classification objective in equation 1. For an input $(x, y) \sim \mathcal{D}$, we define the size of the smallest class-changing perturbation as:

$$\varepsilon^*(x) := \min \{ \|\Delta\| : \mathcal{O}(x + \Delta) \notin \{y, \perp\} \}. \quad (2)$$

Let x be an input where the considered distance function is not aligned with the oracle. Let $x_2 = x + \Delta$ be the closest input to x with a different class label, i.e., $\mathcal{O}(x_2) = y' \neq y$ and $\|\Delta\| = \varepsilon^*(x)$. As the distance and oracle are misaligned, there exists an input $x_1 = x + \Delta'$ such that $\|\Delta'\| > \varepsilon^*(x)$ and $\mathcal{O}(x_1) = y$. So now, if we train a model to be robust (in the sense of equation 1) to perturbations of norm bounded by $\varepsilon \leq \varepsilon^*(x)$, the model might misclassify x_1 , i.e., it is sensitive to non-semantic changes. Instead, if we make the classifier robust to perturbations bounded by $\varepsilon > \varepsilon^*(x)$, then x_2 becomes an invariance adversarial examples as the model will classify it the same way as x . The two types of failure modes are visualized in Figure 2.

Lemma 4. *Constructing an oracle-aligned distance function that satisfies Definition 3 is as hard as constructing a function f so that $f(x) = \mathcal{O}(x)$, i.e., f perfectly solves the oracle’s classification task.*

The proof of this lemma is in Appendix D; at a high level, observe that given a valid distance function that satisfies Definition 3 we can construct a nearest neighbor classifier that perfectly matches the oracle. Thus, in general we cannot hope to have such a distance function.

4. Generating Invariance-based Adversarial Examples on MNIST

We now empirically demonstrate and evaluate the trade-off between sensitivity-based and invariance-based adversarial

Algorithm 1 Meta-algorithm for finding invariance-based adversarial examples. For an input x , we find an input x^* of a different class in the dataset \mathcal{X} , that is closest to x under some set of semantics-preserving transformations. \mathcal{T} .

GenInv $(x, y, \mathcal{X}, \mathcal{T})$

$$\begin{cases} \mathcal{S} = \{\hat{x} : (\hat{x}, \hat{y}) \in \mathcal{X}, \hat{y} \neq y\} \\ \mathcal{X}^* = \{t(\hat{x}) : t \in \mathcal{T}, \hat{x} \in \mathcal{S}\} \\ \mathbf{return} \ x^* = \arg \min_{\hat{x} \in \mathcal{X}^*} \|\hat{x} - x\| \end{cases}$$

examples. We propose an algorithm for generating invariance adversarial examples, and show that robustified models are disparately more vulnerable to these attacks compared to standard models. In particular, we break both *adversarially-trained* and *certifiably-robust* models on MNIST by generating invariance adversarial examples—within the models’ (possibly certified) norm bound—to which the models’ assign different labels than an ensemble of humans.

Why MNIST? We elect to study MNIST, the only dataset for which strong robustness to various ℓ_p -bounded perturbations is attainable with current techniques (Madry et al., 2017; Schott et al., 2019). The dataset’s simplicity is what initially prompted the study of simple ℓ_p -bounded perturbations (Goodfellow et al., 2015). Increasing MNIST models’ robustness to such perturbations has since become a standard benchmark (Schott et al., 2019; Madry et al., 2017; Wong & Kolter, 2018; Raghunathan et al., 2018). Due to the existence of models with high robustness to various ℓ_p -bounded attacks, robust classification on MNIST is considered close to solved (Schott et al., 2019).

This paper argues that, contrary to popular belief, MNIST is far from being solved. We show why optimizing for robustness to ℓ_p -bounded adversaries is not only insufficient, but actively harms the performance of the classifier against alternative invariance-based attacks.

In Section 4.3, we show that complex vision tasks (e.g., ImageNet) are also affected by the fundamental tradeoffs we describe. These tradeoffs are simply not apparent yet, because of our inability to train models with non-negligible robustness to any attacks on these tasks.

4.1. Generating Model-agnostic Invariance-based Adversarial Examples

We propose a model-agnostic algorithm for crafting invariance adversarial examples. Our attack generates minimally perturbed invariance adversarial examples that cause humans to change their classification. We then evaluate these examples against multiple models. The rationale for this approach is mainly that obtaining human labels is expensive, which encourages the use of a single attack for all models.

The high-level algorithm we use is in Algorithm 1 and




Figure 4: Process for generating ℓ_0 invariant adversarial examples. From left to right: (a) the original image of an 8; (b) the nearest training image (labeled as 3), before alignment; (c) the nearest training image (still labeled as 3), after alignment; (d) the Δ perturbation between the original and aligned training example; (e) spectral clustering of the perturbation Δ ; and (f-h) candidate invariance adversarial examples, selected by applying subsets of clusters of Δ to the original image. (f) is a failed attempt at an invariance adversarial example. (g) is successful, but introduces a larger perturbation than necessary (adding pixels to the bottom of the 3). (h) is successful and minimally perturbed.

described below. It is simple, albeit tailored to datasets where comparing images in pixel space is meaningful, like MNIST.³

Given an input x , the attack’s goal is to find the smallest class-changing perturbation $x^* = x + \Delta$ (c.f. equation 2) such that $\mathcal{O}(x^*) \neq \mathcal{O}(x)$. Typically, x^* is not a part of the dataset. We thus approximate x^* via *semantics-preserving* transformations of other inputs. That is, for the set \mathcal{S} of inputs of a different class than x , we apply transformations \mathcal{T} (e.g., small image rotations, translations) that are known a-priori to preserve input labels. We then pick the transformed input that is closest to our target point under the considered ℓ_p metric. In Appendix A, we describe instantiations of this algorithm for the ℓ_0 and ℓ_∞ norms. Figure 4 visualizes the sub-steps for the ℓ_0 attack, including an extra post-processing that further reduces the perturbation size.

Measuring Attack Success. We refer to an invariance adversarial example as *successful* if it causes a change in the oracle’s label, i.e., $\mathcal{O}(x^*) \neq \mathcal{O}(x)$. This is a model-agnostic version of Definition 2. In practice, we simulate the oracle by asking an ensemble of humans to label the point x^* ; if more than some fraction of them agree on the label (throughout this section, 70%) and that label is different from the original, the attack is successful. Note that success or failure is independent of any machine learning model.

4.2. Evaluation

Attack analysis. We generate 100 invariance adversarial examples on inputs randomly drawn from the MNIST test set, for both the ℓ_0 and ℓ_∞ norms. Our attack is slow, with the alignment process taking (amortized) minutes per

³Kaushik et al. (2019) consider a similar problem for NLP tasks. They ask human labelers to produce “counterfactually-augmented data” by introducing a minimal number of changes to a text document so as to change the document’s semantics.

Attack Type	Success Rate
Clean Images	0%
ℓ_0 Attack	55%
$\ell_\infty, \varepsilon = 0.3$ Attack	21%
$\ell_\infty, \varepsilon = 0.3$ Attack (manual)	26%
$\ell_\infty, \varepsilon = 0.4$ Attack	37%
$\ell_\infty, \varepsilon = 0.4$ Attack (manual)	88%

Table 1: Success rate of our invariance adversarial examples in causing humans to switch their classification.




Figure 5: Our invariance-based adversarial examples. Top to bottom: original images and our ℓ_0 , ℓ_∞ at $\varepsilon = 0.3$ and ℓ_∞ at $\varepsilon = 0.4$ invariance adversarial examples. (left) successful attacks; (right) failed attack attempts.

example. We performed no optimizations of this process and expect it could be improved. The mean ℓ_0 distortion of successful examples is 25.9 (with a median of 25). The ℓ_∞ attack always uses the full budget of either $\varepsilon = 0.3$ or $\varepsilon = 0.4$ and runs in a similar amount of time.

Human Study. We conducted a human study to evaluate whether our invariance adversarial examples are indeed successful, i.e., whether humans agree that the label has been changed. We also hand-crafted 50 invariance adversarial examples for the ℓ_0 and ℓ_∞ norm. The process was quite simple: we built an image editor that lets us change images at a pixel level under an ℓ_p constraint. One author then modified 50 random test examples in the way that they perceived as changing the underlying class. We presented all these invariance examples to 40 human evaluators. Each evaluator classified 100 digits, half of which were unmodified MNIST digits, and the other half were sampled randomly from our ℓ_0 and ℓ_∞ invariance adversarial examples.

Results. Of 100 clean (unmodified) test images, 98 are labeled identically by *all* human evaluators. The other 2 images were labeled identically by over 90% of evaluators.

Our ℓ_0 attack is highly effective: For 55 of the 100 examples at least 70% of human evaluators labeled it the same way, with a different label than the original test label. Humans only agreed with the original test label (with the same 70% threshold) on 34 of the images, while they did not form a

Agreement between model and humans, for <i>successful</i> invariance adversarial examples							
Model: ¹	Undefended	ℓ_0 Sparse	Binary-ABS	ABS	ℓ_∞ PGD ($\epsilon = 0.3$)	ℓ_2 PGD ($\epsilon = 2$)	
Clean	99%	99%	99%	99%	99%	99%	99%
ℓ_0	80%	38%	47%	58%	56%*	27%*	
$\ell_\infty, \epsilon = 0.3$	33%	19%*	0%	14%	0%	5%*	
$\ell_\infty, \epsilon = 0.4$	51%	27%*	8%	18%	16%*	19%*	

¹ ℓ_0 Sparse: (Bafna et al., 2018); ABS and Binary-ABS: (Schott et al., 2019); ℓ_∞ PGD and ℓ_2 PGD: (Madry et al., 2017)

Table 2: Model accuracy with respect to the oracle human labelers on the subset of examples where the human-obtained oracle label is different from the test label. Models which are *more* robust to *perturbation* adversarial examples (such as those trained with adversarial training) tend to agree with humans **less often** on *invariance-based* adversarial examples. Values denoted with an asterisks * violate the perturbation threat model of the defense and should not be taken to be attacks. When the model is *wrong*, it failed to classify the input as the new oracle label.

consensus on 18 examples. The simpler ℓ_∞ attack is less effective: with a distortion of 0.3 the oracle label changed 21% of the time and with 0.4 the oracle label changed 37% of the time. The manually created ℓ_∞ examples with distortion of 0.4 were highly effective however: for 88% of the examples, at least 70% assigned the same label (different than the test set label). We summarize results in Table 1. In Figure 5 we show sample invariance adversarial examples.

Our attack code, as well as our invariance examples and their human-assigned labels are available at <https://github.com/ftramer/Excessive-Invariance>.

To simplify the analysis below, we split our generated invariance adversarial examples into two sets: the successes and the failures, as determined by whether the plurality decision by humans was different than or equal to the original label. We only evaluate models on those invariance adversarial examples that caused the humans to switch their classification.

Model Evaluation. Given oracle ground-truth labels for each of the images (as decided by humans), we report how often models agree with the human-assigned label. Table 2 summarizes this analysis. For the invariance adversarial examples, we report model accuracy only on *successful* attacks (i.e., those where the human oracle label changed between the original image and the modified image).⁴ For these same models, Table 3 in Appendix C reports the “standard” robust accuracy for sensitivity-based adversarial examples, i.e., in the sense of equation 1.

The models which empirically achieve the highest robustness against ℓ_0 perturbations (in the sense of equation 1) are the ℓ_0 -Sparse classifier of Bafna et al. (2018), the Binary-ABS model of Schott et al. (2019), and the ℓ_2 -PGD ad-

⁴It may seem counter-intuitive that our ℓ_∞ attack with $\epsilon = 0.3$ appears stronger than the one with $\epsilon = 0.4$. Yet, given two successful invariance examples (i.e., that both change the human-assigned label), the one with lower distortion is expected to change a model’s output less often, and is thus a stronger invariance attack.

versarially trained model (see Table 3 in the Appendix for a comparison of the robustness of these models). Thus, these are the models that are most *invariant* to perturbations of large ℓ_0 -norm. We find that these are the models that achieve the lowest accuracy—as measured by the human labelers—on our invariance examples. Moreover, all robust models perform much worse than an undefended ResNet-18 model on our invariance attacks. This includes models such as the ℓ_∞ -PGD adversarially trained model, which do not explicitly aim at worst-case robustness against ℓ_0 noise. Thus, we find that models that were designed to reduce excessive sensitivity to certain non-semantic features, become excessively invariant to other features that are semantically meaningful.

Similarly, we find that models designed for ℓ_∞ robustness (Binary-ABS and ℓ_∞ -PGD) also fare the worst on our ℓ_∞ invariance adversarial examples. Overall, all robust models do worse than the undefended baseline. The results are consistent for attacks with $\epsilon = 0.3$ and with $\epsilon = 0.4$, the latter being more successful in changing human labels.

Note that the Binary-ABS defense of (Schott et al., 2019) boasts 60% (empirical) robust accuracy on ℓ_∞ -attacks with $\epsilon = 0.4$ (see (Schott et al., 2019)). Yet, on our our invariance examples that satisfy this perturbation bound, the model actually disagrees with the human labelers 92% of the time, and thus achieves only 8% true accuracy on these examples. Below, we make a similar observation for a *certified* defense.

Trading Perturbation-Robustness for Invariance-Robustness. To better understand how robustness to sensitivity-based adversarial examples influences robustness to invariance attacks, we evaluate a range of adversarially-trained models on our invariance examples.

Specifically, we trained ℓ_∞ -PGD models with $\epsilon \in [0, 0.4]$ and ℓ_1 -PGD models (as a proxy for ℓ_0 -robustness) with $\epsilon \in [0, 15]$. We verified that training against larger perturbations resulted in a monotonic increase in adversarial




Figure 6: Higher noise-robustness leads to higher vulnerability to invariance attacks. (left) For models trained with ℓ_∞ -PGD, a higher bound $\epsilon \in [0, 0.4]$ implies lower accuracy on ℓ_∞ -bounded invariance examples. (right) Models trained with ℓ_1 -PGD evaluated on the ℓ_0 invariance attack.

robustness, in the sense of equation 1 (more details are in Appendix C). We then evaluated these models against respectively the ℓ_∞ and ℓ_0 invariance examples. Figure 6 shows that robustness to larger perturbations leads to higher vulnerability to invariance-based examples.

Interestingly, while sensitivity-based robustness does not generalize beyond the norm-bound on which a model is trained (e.g., a model trained on PGD with $\epsilon = 0.3$ achieves very little robustness to PGD with $\epsilon = 0.4$ (Madry et al., 2017)), excessive invariance does generalize (e.g., a model trained on PGD with $\epsilon = 0.2$ is more vulnerable to our invariance attacks with $\epsilon \geq 0.3$ compared to an undefended model).

Breaking Certified Defenses. Our invariance attacks even constitute a *break* of some certified defenses. For example, Zhang et al. (2019) develop a defense which *proves* that the accuracy on the test set is at least 87% under ℓ_∞ perturbations of size $\epsilon = 0.4$. When we run their pre-trained model on all 100 of our $\epsilon = 0.4$ invariance adversarial examples (i.e., not just the successful ones) we find it has a 96% “accuracy” (i.e., it matches the original test label 96% of the time). However, when we look at the agreement between this model’s predictions with the new labels assigned by the human evaluators, the model’s accuracy is just 63%.

Thus, while the proof in the paper is *mathematically correct* it does not actually deliver 87% robust accuracy under ℓ_∞ -attacks with $\epsilon = 0.4$: humans change their classification for many of these perturbations. Worse, for the 50 adversarial examples we crafted by hand, the model *disagrees* with the human ensemble 88% of the time: it has just 12% accuracy.

4.3. Natural Images

While our experiments are on MNIST, similar phenomena may arise in other vision tasks. Figure 3 shows two perturbations of ImageNet images: the rightmost perturbation is imperceptible and thus classifiers should be robust to it. Conversely, the middle image was semantically changed,

and classifiers should be sensitive to such changes. Yet, the ℓ_2 norm of both perturbations is the same. Hence, enforcing robustness to ℓ_2 -noise of some fixed size ϵ will necessarily result in a classifier that is either sensitive to the changes on the right, or invariant to the changes in the middle image. Such a phenomenon will necessarily arise for any image dataset that contains small objects, as perturbations of small ℓ_2 magnitude will be sufficient to occlude the object, thereby changing the image semantics.

This distance-oracle misalignment extends beyond the ℓ_2 -norm. For instance, Co et al. (2018) show that a perturbation of size $16/255$ in ℓ_∞ can suffice to give an image of a cat the appearance of a shower curtain print, which are both valid ImageNet classes. Yet, a *random* perturbation of the same magnitude is semantically meaningless.

On CIFAR-10, some recent defenses are possibly already overly invariant. For example, Shaeiri et al. (2020) and Panda et al. (2019) aim to train models that are robust to ℓ_∞ perturbations of size $\epsilon = 32/255$. Yet, Tsipras et al. (2019) show that perturbations of that magnitude can be semantically meaningful and can be used to effectively *interpolate* between CIFAR-10 classes. The approach taken by Tsipras et al. (2019) to create these perturbations, which is based on a model with robustness to very small ℓ_∞ noise, may point towards an efficient way of automating the generation of invariance attacks for tasks beyond MNIST. The work of Sharif et al. (2018) also shows that “small” ℓ_∞ noise (of magnitude $25/255$) can reliably fool human labelers on CIFAR-10.

5. The Overly-Robust Features Model

The experiments in Section 4 show that models can be robust to perturbations large enough to change an input’s semantics. Taking a step back, it is not obvious why training such classifiers is possible, i.e., why does excessive invariance not harm regular accuracy. To understand the learning dynamics of these overly-robust models, we ask two questions:

1. Can an overly-robust model fit the training data?
2. Can such a model generalize (robustly) to test data?

For simplicity, we assume that for every point $(x, y) \sim \mathcal{D}$, the closest point x^* (under the chosen norm) for which $\mathcal{O}(x^*) \neq y$ is at a constant distance ϵ^* . We train a model f to have low robust error (as in equation 1) for perturbations of size $\epsilon > \epsilon^*$. This model is thus overly-robust.

We first ask under what conditions f may have low robust training error. A necessary condition is that there do not exist training points $(x_i, y_i), (x_j, y_j)$ such that $y_i \neq y_j$ and $\|x_i - x_j\| \leq \epsilon$. As ϵ is larger than the inter-class distance, the ability to fit an overly robust model thus relies on the training data not being fully representative of the space to

which the oracle assigns labels. This seems to be the case in MNIST: as the dataset consists of centered, straightened and binarized digits, even an imaginary infinite-sized dataset might not contain our invariance adversarial examples.

The fact that excessive robustness *generalizes* (as provably evidenced by the model of Zhang et al. (2019)) points to a deeper issue: there must exist *overly-robust* and *predictive* features in the data—that are not aligned with human perception. This mirrors the observations of (Ilyas et al., 2019), who show that excessive sensitivity is caused by non-robust yet predictive features. On MNIST, our experiments confirm the existence of overly-robust generalizable features.

We formalize these observations using a simple classification task inspired by (Tsipras et al., 2019). We consider a binary task where unlabeled inputs $x \in \mathbb{R}^{d+2}$ are sampled from a distribution \mathcal{D}_k^* with parameter k :

$$z \stackrel{\text{u.a.r.}}{\sim} \{-1, 1\}, \quad x_1 = z/2$$

$$x_2 = \begin{cases} +z & \text{w.p. } \frac{1+1/k}{2} \\ -z & \text{w.p. } \frac{1-1/k}{2} \end{cases}, \quad x_3, \dots, x_{d+2} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}\left(\frac{z}{\sqrt{d}}, k\right).$$

Here $\mathcal{N}(\mu, \sigma^2)$ is a normal distribution and $k > 1$ is a constant chosen so that only feature x_1 is strongly predictive of the latent variable z (e.g., $k = 100$ so that x_2, \dots, x_{d+2} are almost uncorrelated with z). The oracle is defined as $\mathcal{O}(x) = \text{sign}(x_1)$, i.e., feature x_1 fully defines the oracle’s class label, and other features are nearly uncorrelated with it. Note that the oracle’s labels are robust under any ℓ_∞ -noise with norm strictly below $\varepsilon = 1/2$.

We model the collection of “sanitized” and labeled datasets from a data distribution as follows: the semantic features (i.e., x_1) are preserved, while “noise” features have their variance reduced (e.g., because non-standard inputs are removed). Sanitization thus enhances “spurious correlations” (Jo & Bengio, 2017; Ilyas et al., 2019) between non-predictive features and class labels.⁵ We further assume that the data labeling process introduces some small label noise.⁶ Specifically, the labeled data distribution \mathcal{D} on which we train and evaluate classifiers is obtained by sampling x from a sanitized distribution $\mathcal{D}_{1+\alpha}^*$ (for a small constant $\alpha > 0$) where features x_2, \dots, x_{d+2} are strongly correlated with the oracle label. The label y is set to the correct oracle label with high probability $1 - \delta$. The consequences of this data sanitization are two-fold (see Appendix E for proofs):

1. A standard classifier (that maximizes accuracy on \mathcal{D})

⁵In digit classification for example, the number of pixels above 1/2 is a feature that is presumably very weakly correlated with the class 8. In the MNIST dataset however, this feature is fairly predictive of the class 8 and robust to ℓ_∞ -noise of size $\varepsilon = 0.4$.

⁶This technicality avoids that classifiers on \mathcal{D} can trivially learn the oracle labeling function. Alternatively, we could define feature x_1 so that it is hard to learn for certain classes of classifiers.

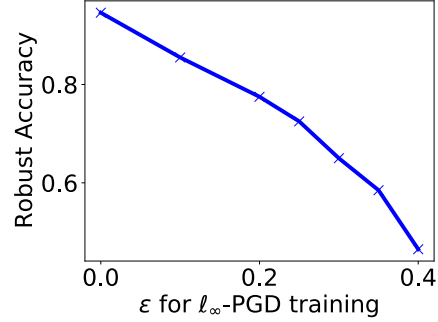


Figure 7: Robust accuracy of models trained and evaluated on an adversary combining a small spatial data augmentation (rotation + translation) with an ℓ_∞ perturbation bounded by ε .

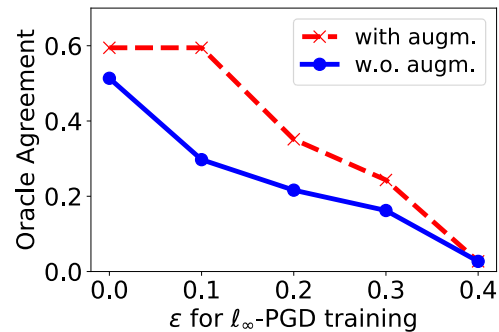


Figure 8: Model accuracy with respect to human labels on successful invariance adversarial examples of ℓ_∞ -norm bounded by 0.4. Models trained with data augmentation agree more often with humans, and thus are less invariant to semantically-meaningful changes.

agrees with the oracle with probability at least $1 - \delta$, but is vulnerable to ℓ_∞ -perturbations of size $\varepsilon = O(d^{-1/2})$.

2. There is an overly-robust model that only uses feature x_2 and has robust accuracy $1 - \alpha/2$ on \mathcal{D} for ℓ_∞ -noise of size $\varepsilon = 0.99$. This classifier is vulnerable to invariance attacks as the oracle is not robust to such perturbations.

The Role of Data Augmentation. This simple task suggests a natural way to prevent the training of overly robust models. If *prior knowledge* about the task suggests that classification should be invariant to features x_2, \dots, x_{d+2} , then enforcing these invariances would prevent a model from being robust to excessively large perturbations.

A standard way to enforce invariances is via data augmentation. In the above binary task, augmenting the training data by randomizing over features x_2, \dots, x_{d+2} would force the model to rely on the only truly predictive feature, x_1 .

We experimented with aggressive data-augmentation on MNIST. For values of $\varepsilon \in [0, 0.4]$, we train models with an adversary that rotates and translates inputs by a small amount and then adds ε -bounded ℓ_∞ -noise. This attacker

mirrors the process we use to generate invariance adversarial examples in Section 4. Thus, we expect it to be hard to achieve robustness to attacks with large ε on this dataset, as this requires the model to correctly classify inputs that humans consider mislabeled. Figure 7 confirms this intuition. As ε grows, it becomes harder to learn a model that is invariant to both spatial data augmentation and ℓ_∞ noise. We further find that the models trained with data augmentation agree more often with human labelers on our invariance attacks (see Figure 8). Yet, even with data augmentation, models trained against large ℓ_∞ -perturbations still perform worse than an undefended model. This simple experiment thus demonstrates that while data-augmentation (over truly invariant features) can help in detecting or preventing excessive invariance to semantic features, even though it is not currently sufficient for training models that resist both sensitivity-based and invariance-based attacks.

6. Discussion

Our results show that solely focusing on robustness to sensitivity-based attacks is insufficient, as mis-specified bounds can cause vulnerability to invariance-based attacks.

On ℓ_p -norm evaluations. Our invariance attacks are able to find points within the ℓ_p -ball in which state-of-the-art classifiers are (provably) robust. This highlights the need for a more careful selection of perturbation bounds when measuring robustness to adversarial examples. At the same time, Figure 6 shows that even promoting robustness within conservative bounds causes excessive invariance. The tradeoff explored in Section 3 suggests that aiming for robustness against ℓ_p -bounded attacks may be inherently futile for making models robust to arbitrary adversarial examples.

Trading Sensitivity and Invariance. We show that models that are robust to small perturbations make excessively invariant decisions and are thus vulnerable to other attacks.

Interestingly, Engstrom et al. (2019a) show an opposite effect for models’ internal representations. Denoting the logit layer of a model as $z(x)$, they show that for robust models it is hard to find inputs x, x^* such that $\mathcal{O}(x) \neq \mathcal{O}(x^*)$ and $z(x) \approx z(x^*)$. Conversely, Sabour et al. (2016) and Jacobsen et al. (2019) show that excessive invariance of feature layers is common in non-robust models. These observations are orthogonal to ours as we study invariances in a model’s classification layer, and for bounded perturbations. As we show in Section 3, robustness to large perturbations under a norm that is misaligned with human perception necessarily causes excessive invariance of the model’s classifications (but implies nothing about the model’s feature layers).

Increasing model robustness to ℓ_p -noise also leads to other tradeoffs, such as reduced accuracy (Tsipras et al., 2019) or

reduced robustness to other small perturbations (Yin et al., 2019; Tramèr & Boneh, 2019; Kang et al., 2019).

7. Conclusion

We have introduced and studied a fundamental tradeoff between two types of adversarial examples, that stem either from excessive sensitivity or invariance of a classifier. This tradeoff is due to an inherent misalignment between simple robustness notions and a task’s true perceptual metric.

We have demonstrated that defenses against ℓ_p -bounded perturbations on MNIST promote invariance to semantic changes. Our attack exploits this excessive invariance by changing image semantics while preserving model decisions. For adversarially-trained and certified defenses, our attack can reduce a model’s true accuracy to random guessing.

Finally, we have studied the tradeoff between sensitivity and invariance in a theoretical setting where excessive invariance can be explained by the existence of overly-robust features.

Our results highlight the need for a more principled approach in selecting meaningful robustness bounds and in measuring progress towards more robust models.

Acknowledgments

Florian Tramèr’s research was supported in part by the Swiss National Science Foundation (SNSF project P1SKP2_178149). Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute www.vectorinstitute.ai/partners.

References

- Bafna, M., Murtagh, J., and Vyas, N. Thwarting adversarial examples: An ℓ_0 -robust sparse fourier transform. In *Advances in Neural Information Processing Systems*, pp. 10096–10106, 2018.
- Biggio, B., Corona, I., Maiorca, D., Nelson, B., Šrđić, N., Laskov, P., Giacinto, G., and Roli, F. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases*, pp. 387–402. Springer, 2013.
- Co, K. T., Muñoz-González, L., de Maupeou, S., and Lupu, E. C. Procedural noise adversarial examples for black-box attacks on deep convolutional networks. *arXiv preprint arXiv:1810.00470*, 2018.
- Engstrom, L., Ilyas, A., Santurkar, S., Tsipras, D., Tran, B., and Madry, A. Adversarial robustness as a prior for learned representations, 2019a.

- Engstrom, L., Tran, B., Tsipras, D., Schmidt, L., and Madry, A. Exploring the landscape of spatial robustness. In *International Conference on Machine Learning*, pp. 1802–1811, 2019b.
- Gilmer, J., Adams, R. P., Goodfellow, I., Andersen, D., and Dahl, G. E. Motivating the rules of the game for adversarial example research. *arXiv preprint arXiv:1807.06732*, 2018.
- Goodfellow, I. and Papernot, N. Is attacking machine learning easier than defending it? *Blog post on Feb*, 15:2017, 2017.
- Goodfellow, I., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *International Conference on Learning Representations*, 2015.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Ilyas, A., Jalal, A., Asteri, E., Daskalakis, C., and Dimakis, A. G. The robust manifold defense: Adversarial training using generative models. *arXiv preprint arXiv:1712.09196*, 2017.
- Ilyas, A., Santurkar, S., Tsipras, D., Engstrom, L., Tran, B., and Madry, A. Adversarial examples are not bugs, they are features. In *Advances in Neural Information Processing Systems*, pp. 125–136, 2019.
- Jacobsen, J.-H., Behrmann, J., Zemel, R., and Bethge, M. Excessive invariance causes adversarial vulnerability. In *International Conference on Learning Representations*, 2019.
- Jo, J. and Bengio, Y. Measuring the tendency of cnns to learn surface statistical regularities. *arXiv preprint arXiv:1711.11561*, 2017.
- Kang, D., Sun, Y., Hendrycks, D., Brown, T., and Steinhardt, J. Testing robustness against unforeseen adversaries. *arXiv preprint arXiv:1908.08016*, 2019.
- Kaushik, D., Hovy, E., and Lipton, Z. C. Learning the difference that makes a difference with counterfactually-augmented data. *arXiv preprint arXiv:1909.12434*, 2019.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. *International Conference on Learning Representations*, 2017.
- Mirza, M. and Osindero, S. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- Panda, P., Chakraborty, I., and Roy, K. Discretization based solutions for secure machine learning against adversarial attacks. *IEEE Access*, 7:70157–70168, 2019.
- Raghunathan, A., Steinhardt, J., and Liang, P. Certified defenses against adversarial examples. *International Conference on Learning Representations*, 2018.
- Sabour, S., Cao, Y., Faghri, F., and Fleet, D. J. Adversarial manipulation of deep representations. *International Conference on Learning Representations*, 2016.
- Samangouei, P., Kabkab, M., and Chellappa, R. Defense-gan: Protecting classifiers against adversarial attacks using generative models. *arXiv preprint arXiv:1805.06605*, 2018.
- Schott, L., Rauber, J., Bethge, M., and Brendel, W. Towards the first adversarially robust neural network model on MNIST. In *International Conference on Learning Representations*, 2019.
- Shaeiri, A., Nobahari, R., and Rohban, M. H. Towards deep learning models resistant to large perturbations. *arXiv preprint arXiv:2003.13370*, 2020.
- Sharif, M., Bauer, L., and Reiter, M. K. On the suitability of lp-norms for creating and preventing adversarial examples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1605–1613, 2018.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- Tramèr, F. and Boneh, D. Adversarial training and robustness for multiple perturbations. In *Advances in Neural Information Processing Systems*, pp. 5858–5868, 2019.
- Tramèr, F., Dupré, P., Rusak, G., Pellegrino, G., and Boneh, D. Adversarial: Perceptual ad blocking meets adversarial machine learning. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pp. 2005–2021, 2019.
- Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., and Madry, A. Robustness may be at odds with accuracy. In *International Conference on Learning Representations*, 2019.
- Wong, E. and Kolter, Z. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Yin, D., Lopes, R. G., Shlens, J., Cubuk, E. D., and Gilmer, J. A fourier perspective on model robustness in computer vision. *arXiv preprint arXiv:1906.08988*, 2019.

Zhang, H., Chen, H., Xiao, C., Li, B., Boning, D., and Hsieh, C.-J. Towards stable and efficient training of verifiably robust neural networks. *arXiv preprint arXiv:1906.06316*, 2019.