

Identifying Model Weakness with Adversarial Examiner

Michelle Shu, Chenxi Liu,  Weichao Qiu, Alan Yuille

Johns Hopkins University

{mshu1, cxliu}@jhu.edu, {qiuwch, alan.l.yuille}@gmail.com

Abstract

Machine learning models are usually evaluated according to the average case performance on the test set. However, this is not always ideal, because in some sensitive domains (e.g. autonomous driving), it is the worst case performance that matters more. In this paper, we are interested in systematic exploration of the input data space to identify the weakness of the model to be evaluated. We propose to use an *adversarial examiner* in the testing stage. Different from the existing strategy to always give the same (distribution of) test data, the adversarial examiner will dynamically select the next test data to hand out based on the testing history so far, with the goal being to undermine the model’s performance. This sequence of test data not only helps us understand the current model, but also serves as constructive feedback to help improve the model in the next iteration. We conduct experiments on ShapeNet object classification. We show that our adversarial examiner can successfully put more emphasis on the weakness of the model, preventing performance estimates from being overly optimistic.

1 Introduction

The field of machine learning is advancing at unprecedented speed, and one evidence often cited is the reported performance on benchmark datasets. There are usually leaderboards for these public datasets with specific evaluation metrics, and the entry that achieves the highest number is regarded as the state-of-the-art. In several cases, it is claimed that the machine learning model has surpassed human-level performance using this criterion (He et al. 2015; Wang, Yan, and Wu 2018).

However, the common belief is that humans are still superior, which suggests that the current testing strategy is overly optimistic and does not reflect the true progress in advancing machine learning. We think an important reason behind this mismatch is that: the current evaluation protocol uses *fixed* test data and measures the *average* case performance, which does not place enough emphasis on the *worst* case performance. To make this point concrete, imagine a model for autonomous driving. Suppose that in the testing data, there are

1 million images under normal conditions (in which case the car should keep driving), and 1 image with a baby in front (in which case the car should immediately stop). Under the current evaluation protocol, a policy to always keep driving and give up on the rare case will be highly regarded. However, from the humans’ perspective, this strategy is clearly problematic and alarming. Therefore, at least in some sensitive domains, placing more emphasis on the worst cases to reveal model weaknesses may be more important than simply relying on the average case.

Revealing model weaknesses by densely sampling the input data space is clearly impractical. This is because any input data point is usually determined by a combination of factors, which makes the space exponentially large. For example, an image is jointly determined by the rotation angles of the objects, the positions of the objects, the material properties, the type of lighting, the energy of lighting, the distance to camera, etc. Larger test set certainly helps, but no matter how large it gets, it will never be able to densely cover this exponentially large space. Therefore our goal is to efficiently and systematically identify the model weaknesses within a reasonable number of queries.

Different models have different weaknesses, so it no longer makes sense to have a fixed test set, rather a dynamic one. In some way, this notion of fixed versus dynamic is very similar to giving standardized tests versus conducting interviews. In standardized tests, the questions are usually designed to have general coverage, since all test takers answer the same questions. In interviews, the examiner may start with general topics, but also has the freedom to deep dive into specific ones to challenge the candidate. We also point out that this notion of dynamic examiner is closer in spirit to Turing test. If the human evaluator in Turing test always asks the same questions, then its role can be replaced by a machine. It is also conceivable that the worst answer during the exchange will contribute most to the human evaluator’s final judgment.

Formalizing the ideas discussed above, we propose *adversarial examiner*, an alternative way of evaluating a model. Provided with a model and the space of input data, the adversarial examiner will dynamically select the next test data to hand out, based on the testing history so far. The name “ad-




Figure 1: Evaluating a model’s ability to recognize a *lamp* instance in ShapeNet. If we randomly sample data, the evaluation will converge to a high percentage (around 80%). However, using our Adversarial Examiner (AE), the test cases gradually concentrate on the weaknesses of the model, and the evaluation will not be overly optimistic. In this example, AE has found that the model appears to be vulnerable to increased lighting (within a reasonable range).

versarial” comes from the fact that its goal is to minimize the model’s performance. Compared with the standard testing protocol, adversarial examiner obviously places more emphasis on the worst case scenarios. Because of its dependence on testing history, the sequence of test data naturally reveals the process of identifying model weakness. Fig. 1 offers a demonstration.

Adversarial attack is also a way of showing the weakness of the model. The message of adversarial attack is that for a particular data point, there exists another data point within its local neighborhood, such that the model fails. However, both the original data point and the local neighborhood are microscopic properties. By contrast, our adversarial examiner starts with the entire input data space, and therefore has the potential of revealing the global picture. In addition, since the adversarial examiner concerns the global structure, the testing results are more informative and more easily perceived as “constructive feedback” that help improve the model in the next iteration.

We conduct experiments on ShapeNet object classification. The input data space is the Cartesian product of 12 physical parameters that determine the scene, including the rotation and position of object, the energy and angle of lighting, etc. Through both quantitative evaluation and qualitative visualization, we show that our adversarial examiner can successfully identify the weakness of the model, preventing performance estimates from being overly optimistic.

2 Method

2.1 Adversarial Examiner

We begin by revisiting the evaluation protocol employed in a standard classification task. The examples $x \in \mathbb{R}^p$ follow an underlying distribution \mathcal{P} . The ground truth label of x is $y(x) \in \{1, 2, \dots, k\}$. The label predicted by the model f is $f(x) \in \mathbb{R}^k$. Given a loss function $L(\cdot, \cdot)$ that measures the compatibility between $f(x)$ and $y(x)$, the evaluation metric is:

$$E = \mathbb{E}_{x \sim \mathcal{P}}[L(f(x), y(x))] \approx \frac{1}{N} \sum_{i=1}^N L(f(x_i), y(x_i)) \quad (1)$$

where $x_i \sim \mathcal{P}$.

In applications such as computer vision and natural language processing, the examples x are usually the surface representation, e.g. 2D images in vision and words in language. However, there is typically an underlying representation z that generates/determines this surface representation x . In vision, z would be the 3D object instance. In language, z would be the semantics/pragmatics that the speaker wishes to communicate. In the following, we consider evaluation on the underlying representation z , which is more intrinsic.

The examples $z \in \mathbb{R}^q$ follow an underlying distribution \mathcal{Q} , whose support is \mathcal{D} . If we can obtain the ground truth label using the surface form x , then we can certainly obtain it using the underlying form z . We slightly abuse annotation and use $y(z) \in \{1, 2, \dots, k\}$ to represent the ground truth label for instance z . Let g be the function that transforms the underlying form into surface form $x = g(z, s)$, where $s \in \mathcal{S}$ represents the remaining information needed to complete this transformation. For example, in vision, s could include camera information, lighting condition, etc. In language, s may be a tie-breaker between synonyms or sentence templates that express the same meaning.

We use the following evaluation metric in our *adversarial examiner*:

$$\begin{aligned} E_{\text{examiner}} &= \mathbb{E}_{z \sim \mathcal{Q}}[\max_{s \in \mathcal{S}} L(f(g(z, s)), y(z))] \\ &\approx \frac{1}{N} \sum_{i=1}^N \max_{s_i \in \mathcal{S}} L(f(g(z_i, s_i)), y(z_i)) \end{aligned} \quad (2)$$

where $z_i \sim \mathcal{Q}$. We choose \mathcal{S} such that for any instance $z \in \mathcal{D}$ and any $s \in \mathcal{S}$, the surface form $g(z, s)$ still preserves the label $y(z)$ as judged by humans. The max in (2) means the examiner will consider the worst case within the space \mathcal{S} , which differs from (1).

It is easy to see that

$$E_{\text{examiner}} \leq \max_{z \in \mathcal{D}, s \in \mathcal{S}} L(f(g(z, s)), y(z)) = L(f(g(\hat{z}, \hat{s})), y(\hat{z})) \quad (3)$$

where (\hat{z}, \hat{s}) is the worst case scenario. We now discuss a special case, where the space \mathcal{S} is so rich, that $\forall z_i, z_j \in \mathcal{D}$, $\exists s_i, s_j \in \mathcal{S}$, such that $g(z_i, s_i) = g(z_j, s_j)$. In this case, the worst combination (\hat{z}, \hat{s}) can always be achieved from any sample z_i . Therefore, E_{examiner} will return *exactly* the worst

Algorithm 1: Adversarial Examiner Procedure

Input: N samples $z_i \sim \mathcal{Q}$ and their true labels $y(z_i)$;
 Maximum number of examination steps T ; Loss function L ; Model f ; Function g ; Space \mathcal{S} .

```

1 for  $i = 1$  to  $N$  do
2   Initialize examiner with  $\mathcal{S}$ 
3   for  $t = 1$  to  $T$  do
4      $s_i^t = \text{examiner.generate}()$ 
5      $l_i^t = L(f(g(z_i, s_i^t)), y(z_i))$ 
6     examiner.update( $s_i^t, l_i^t$ )
7 return  $E_{\text{examiner}} = \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T l_i^t$ 

```

case performance under model f , which is what we expect and desire.

Implementing the adversarial examiner according to (2) requires solving an optimization problem for every instance z_i . Since the landscape of $f(g(z_i, \cdot))$ on \mathcal{S} can be arbitrary, we address the general case by assuming no closed-form solution and no differentiability. We consider a sequential strategy, where the adversarial examiner will hand out new candidates of s_i^t based on the testing history so far $s_i^1, s_i^2, \dots, s_i^{t-1}$. See Algorithm 1 for the adversarial examiner procedure.

Relation to Adversarial Attacks We named our examiner “adversarial” to acknowledge its connection to adversarial attacks. If we conduct adversarial attacks on the samples x_i before evaluating using (1), the result is:

$$E_{\text{attack}} \approx \frac{1}{N} \sum_{i=1}^N \max_{\delta_i \in \Delta} L(f(x_i + \delta_i), y(x_i)) \quad (4)$$

There are at least two key differences between adversarial attack (AA) and adversarial examination (AE). First, AE deals with the underlying form whereas AA deals with the surface form. As a result, the variations considered by AE on the surface form is much larger than Δ . Second and more importantly, in AA there is a “canonical” starting point $\delta_i = 0$, so essentially all attack algorithms perform gradient descent starting from $\delta_i = 0$. As a result, the sequence $\delta_i^1, \delta_i^2, \dots$ is usually very local. In AE, there is typically no “canonical” choice for s_i , so AE is forced to start with the entire space \mathcal{S} instead of a specific point. As a result, the sequence s_i^1, s_i^2, \dots exhibits large variations. This global picture reveals different properties of f than AA, and is potentially more useful in representing the weakness (and/or strength) of the model.

Availability of Underlying Representation While data in the surface form are straightforward to collect, there may be concerns whether it is always easy to obtain the underlying representation z considered by the adversarial examiner. We argue that this problem has been greatly alleviated in recent years, with either rich annotation on real data (Xiang, Mottaghi, and Savarese 2014; Krishna et al.

2017), or increasing use of synthetic data (Chang et al. 2015; Johnson et al. 2017), or a hybrid between real and synthetic (Gaidon et al. 2016; Hudson and Manning 2019). In addition, in related topics such as image caption evaluation, the advantage of evaluating the underlying representation (e.g. SPICE (Anderson et al. 2016)) over the surface representation (e.g. BLEU (Papineni et al. 2002)) has been demonstrated.

2.2 Model Choices for Examiner

We describe two classes of models as the examiner in Algorithm 1. One is based on Reinforcement Learning (Section 2.2), and the other is based on Bayesian Optimization (Section 2.2). We conclude by discussing how these two classes of models are complementary (Section 2.2), which shows the general applicability of our adversarial examiner framework.

Reinforcement Learning as Examiner Let space \mathcal{S} be the Cartesian product of C factors $\mathcal{S} = \Psi^1 \times \Psi^2 \times \dots \times \Psi^C$, where Ψ^i represents the parameter range of the i -th factor. Therefore, the candidate s_i^t is composed of $\psi_{(i,t)}^1, \psi_{(i,t)}^2, \dots, \psi_{(i,t)}^C$, where $\psi_{(i,t)}^c \in \Psi^c$. The probability of generating s_i^t is $P(s_i^t) = \prod_{c=1}^C P(\psi_{(i,t)}^c | \psi_{(i,t)}^{c-1:1})$. We use a LSTM (Hochreiter and Schmidhuber 1997) to parameterize these conditional probabilities. Concretely, in the first LSTM step, the initial hidden state h^1 is followed by a fully connected layer and softmax to represent $P(\psi_{(i,t)}^1)$. We then draw a sample according to this distribution, which is fed into the second LSTM step. The updated hidden state h^2 is followed by a fully connected layer and softmax to represent $P(\psi_{(i,t)}^2 | \psi_{(i,t)}^1)$. This process is repeated until we have drawn a complete sample of C steps.

To train this LSTM using reinforcement learning, we define the reward signal R to be $L(f(g(z_i, s_i^t)), y(z_i))$, and optimize the weights θ using policy gradient (Williams 1992):

$$\nabla_{\theta} \mathbb{E}_{P(s_i^t; \theta)} [R] \approx \frac{1}{B} \sum_{b=1}^B \sum_{c=1}^C \nabla_{\theta} \log P(\psi_{(i,t)}^c | \psi_{(i,t)}^{c-1:1}) R_b \quad (5)$$

where B is the batch size to reduce high variance.

Bayesian Optimization as Examiner We use Gaussian Process (GP) as the model behind Bayesian optimization. The value that the GP aims to maximize is $L(f(g(z_i, s_i^t)), y(z_i))$ from Algorithm 1. The example generated by our examiner is equivalent to the point proposed by the acquisition function $a : \mathcal{S} \rightarrow \mathbb{R}^+$. We use Gaussian Process upper confidence bound (UCB) (Srinivas et al. 2009) to construct our acquisition function.

Let W be the set of points that induce the posterior multivariate Gaussian distribution, which is initially empty. For each iteration $t = 1, 2, \dots, T$, we select the next candidate by:

$$s_i^t = \underset{s \in \mathcal{S}}{\operatorname{argmax}} a(s) \quad (6)$$

where a is the acquisition function induced by the current W . We then construct its surface form and test on the model.

This newly observed point $(s_i^t, L(f(g(z_i, s_i^t)), y(z_i)))$ is then added to set W . In the next iteration, the examiner will induce a new posterior multivariate Gaussian distribution based on the updated set W . By the end of adversarial examination, the candidates $\{s_i^t \in \mathcal{S}\}_{t=1}^T$ are the points that induce the most up-to-date posterior multivariate Gaussian distribution on \mathcal{S} .

How the Two Examiners Are Complementary

Discrete vs. Continuous In RL, the choices of parameters in Ψ are discrete as we use softmax to select amongst them. As for BO, the ranges of parameters are typically continuous as the underlying assumption is a multivariate Gaussian distribution.

Maintaining Sampling Distribution on \mathcal{S} vs. Maintaining Function Value on \mathcal{S} The LSTM model within RL captures the sampling probability P on \mathcal{S} , without explicitly recording the function value L . On the other hand, the GP model within BO relies on the function values (set W) to fit the acquisition function.

Longer Iteration Regime vs. Shorter Iteration Regime BO can be quite sample efficient, and is potentially able to successfully attack the model within a limited number of steps. However, it becomes costly when the iteration gets long. On the other hand, policy-based RL can be inefficient, but the computation cost does not go up with more iterations.

3 Related Work

We argue that the average case performance does not always reflect human’s level of trust in a model. The primary design philosophy behind adversarial examiner is to place more emphasis on the worst case performance, which especially matters in sensitive domains. Indeed, medical imaging papers typically report the worst performance on the test set in addition to the average performance. The idea of paying more attention to the worst cases was recently discussed in Yuille and Liu (2018).

As suggested by its name, adversarial examiner is related to finding adversarial examples (Szegedy et al. 2013). They were originally found by doing backpropagation and gradient ascent on the input image. Later researchers started to consider the more challenging but more general scenario, where the neural network weights are not known to the attacker. This is called black box adversarial attack, where existing methods (Narodytska and Kasiviswanathan 2016; Chen et al. 2017; Ilyas et al. 2018) still rely on estimating the

gradient. Essentially, black box adversarial attack belongs to derivative-free optimization, which includes many other families of methods than gradient approximation. To the best of our knowledge, ours is the first work that brings Reinforcement Learning (RL) and Bayesian Optimization (BO) to black box adversarial attack. Our RL setup is inspired by work on Neural Architecture Search (Zoph and Le 2017), whereas our BO setup is inspired by work on hyperparameter selection (Snoek, Larochelle, and Adams 2012). Again, both problems are typical derivative-free optimization.

Recently, researchers have been rethinking the concept of adversarial examples (Gilmer et al. 2018) and generalizing them beyond modifying pixel values. (Engstrom et al. 2017; Pei et al. 2017; Azulay and Weiss 2018) showed that deep networks can be attacked simply by 2D rotation, translation, or brightness change. (Kurakin, Goodfellow, and Bengio 2016; Athalye et al. 2018; Eykholt et al. 2017) demonstrated the possibility of generating physical adversarial examples in the real world. (Gu et al. 2019) used the minute transformations across video frames to study adversarial robustness. (Brown et al. 2018) introduced the concept of unrestricted adversarial examples, and (Song et al. 2018) described a method to generate unrestricted adversarial examples using a differentiable neural network. More related to our work, (Zeng et al. 2019; Yang et al. 2018; Alcorn et al. 2019; Kanazaki, Matsushita, and Nishida 2018) incorporated rendering into the visual recognition pipeline to study how 3D physical parameters may affect model prediction. However, they were all concerned with creating adversarial examples starting from a specific configuration, instead of considering the global space.

Finally, there are a couple of works that used synthetic data to iteratively improve the model, whether in the training stage or the inference stage. Chen et al. (2018) proposed to increase the sampling probability in the underperforming region. Misra et al. (2018) used an agent to learn to select questions in synthetic visual question answering. Yang et al. (2019) learned to navigate in a synthetic scene and select better views for visual recognition. Kar et al. (2019) learned the policy to generate synthetic scenes, whose goal is to best facilitate downstream tasks. Our work is different, in that we focus on a new paradigm of model evaluation, which systematically explores the input data space to identify model weaknesses.

4 Experiments

4.1 Implementation Details

We conduct experiments on visual recognition of objects in the ShapeNet dataset (Chang et al. 2015), which contains

	α_o	β_o	ζ_o	Γ_o	Γ_l	r_l	A_l	U_l	r_v	A_v	U_v	θ_v
UB	2π	2π	2π	5	1	20	360	90	5	180	90	360
LB	0	0	0	0	0.3	8	0	-90	1	0	-90	0

Table 1: Upper bound (UB) and lower bound (LB) of rendering factors in s : sun rotation angles ($\alpha_o, \beta_o, \zeta_o$), sun energy (Γ_o), point light energy (Γ_l), point light distance (r_l), point light location (A_l, U_l), viewpoint distance (r_v), viewpoint location (A_v, U_v), viewpoint angle (θ_v).




Figure 3: t-SNE visualization of *cap* examples under different examiners and target models. There are 100 examples in each subfigure: 50 randomly sampled ones, and 50 from the very end of adversarial examination. Transparent cap means correctly classified. Otherwise, incorrectly classified.

	$m = 10$	$m = 5$	$m = 2$	$m = 1$
RL	63.81%	57.43%	35.05%	18.92%
BO	49.79%	43.06%	22.19%	10.92%

Table 3: Average model performance under adversarial examination for varying m : number of training images per instance. We report $m = 1, 2, 5, 10$ with iterations $T = 200$.

4.2 Evaluating Model Performance with Adversarial Examiners

The evaluation results by our adversarial examiners are shown in Tab. 2. For both AlexNet and ResNet34, we tried both RL and BO examiners under different number of iterations T . The $T = 0$ column corresponds to the standard evaluation protocol, which on average records more than 60% probability on the true class. After a few hundred iterations, both examiners can find the weakness of the model successfully. In addition, longer examination results in lower performance, which is expected. But their characteristics are not entirely the same. BO can find weaknesses relatively faster (larger gap from $T = 0$ to $T = 100$), but eventually ($T = 300$ or $T = 500$) RL becomes more committed to the weaknesses and gives a harsher evaluation.

Fig. 2 shows the per-class post-softmax probability on the true class. We found that the weaknesses of some classes are easier to find than others. For instance, *barrel* and *bowl* are relatively vulnerable since they look very similar (generic round shape) when viewing from bottom-up. Classes such as *airplane*, which contain discriminative visual cues at all angles, are more robust under adversarial examination.

Fig. 3 uses t-SNE (Maaten and Hinton 2008) to visualize 50 randomly sampled examples and 50 examples selected at the end of adversarial examination of the *cap* class. Correctly classified examples are transparent, and incorrectly classified ones are opaque. This visualization serves to show the global distribution of model weakness on the space \mathcal{S} . We can observe clusters among the incorrectly classified examples, e.g. due to viewpoint or lighting.

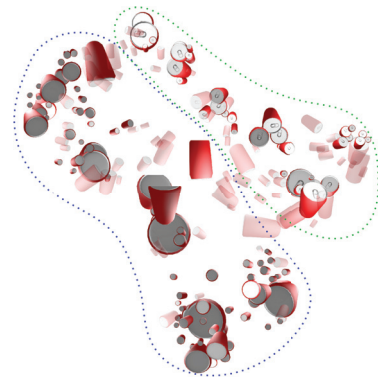


Figure 4: t-SNE visualization of *can* examples. 100 are randomly sampled, and 100 are from RL examiner $T = 250$. Transparent can means correctly classified. Otherwise, incorrectly classified. In this experiment, training data for ResNet34 did not have top and bottom viewpoints. Adversarial examiner is able to identify these two artificial weaknesses (circled in green and blue).

4.3 Examining Models Trained with Less Data

Intuitively, neural networks trained with less data should be less robust against adversarial examiners. In this experiment, we put this hypothesis to test. Specifically, we vary m , the number of images rendered for each 3D object in the training set. We train ResNet34 models with $m = 1, 2, 5, 10$ (recall that $m = 10$ is the default) using the same number of epochs.

Tab. 3 summarizes the results. Under the same iterations T , both RL and BO examiners report decreased performance as m decreases. In addition to validating the adversarial examiner framework, this also confirms the importance of the amount and diversity of training data.

4.4 Evaluating Model with Artificial Weaknesses

To further study the properties of adversarial examiner, we consider an artificial scenario, where we deliberately put in obvious weaknesses. We are interested to see to what extent




Figure 5: Identifying model strength by negating L . For this *piano* instance, the examiner eventually finds a condition under which the model can correctly classify with 100% confidence.

can the adversarial examiner recover these obvious weaknesses. Specifically, during training, we limit the viewpoint elevation U_v to $[-30^\circ, +30^\circ]$ instead of $[-90^\circ, +90^\circ]$. This will create two obvious weaknesses: viewing objects from the top and from the bottom. Note that the examiner is still allowed to change all 9 factors.

Fig. 4 shows the t-SNE visualization of 200 *can* examples. 100 are randomly sampled, and 100 are from $T = 250$ during adversarial examination. Among the incorrectly classified examples (visualized as opaque), there are two large regions, with the bottom left representing viewing from the bottom, and the top right representing viewing from the top. This observation further confirms the capability of adversarial examiners.

4.5 Changing the Order of Factors

As described in Section 2.2, the RL examiner samples parameters in a sequential manner. In this experiment, we investigate whether changing the sampling order of factors will make a difference in adversarial examination.

The default RL examiner samples in the order $\Gamma_l, \Gamma_o, r_l, A_l, U_l, r_v, A_v, U_v, \theta_v$. We prepare a second RL examiner with exactly the same setup, but samples in the order $r_v, A_v, U_v, \theta_v, \Gamma_l, \Gamma_o, r_l, A_l, U_l$. Fig. 6 shows the t-SNE visualization of the last 50 examples given by the two RL examiners ($T = 500$). Obviously they have converged to the




Figure 6: The last 50 *lamp* examples given by two RL examiners with different order of factors.

same weaknesses, which is evidence that the RL examiner is not sensitive to the order.

4.6 Identifying Model Strength

Finally, we show that instead of identifying model weakness, the same setup can also be used to identify model strength, simply by flipping the max in (2) to min. The examiner, instead of trying to find the worst viewing condition of an object, now tries to find the best viewing condition.

We show such a trial in Fig. 5. We indeed observe a curve where the post-softmax probability on the true class goes up throughout evaluation, which is symmetric to Fig. 1. The final image reveals what the model considers to be the easiest *piano* to recognize.

5 Conclusion

In this paper, we advocate for a new testing paradigm for machine learning models, where more emphasis is placed on the worst case instead of reporting the average case performance. To make this idea concrete, we propose *adversarial examiner*, whose goal is to undermine the model’s performance by discovering and concentrating on its weaknesses. The adversarial examiner will dynamically hand out the next testing data, based on the testing history so far. As a result, different models will now be evaluated using potentially different samples from the input data space. We argue that this notion can more effectively prevent hype, is closer in spirit to human interview/Turing test, and can be especially important in sensitive domains such as autonomous driving and medical applications.

We conduct experiments on visual recognition of ShapeNet objects. Examiners based on Reinforcement Learning and Bayesian Optimization both demonstrated effectiveness in identifying the weaknesses of the target model. However, our adversarial examination framework is much more general than classification of rendered 3D objects. In the future, we hope to extend to other domains, including image generation, language generation, and pose estimation in robotics. Using the examples generated during adversarial examination to improve training is another potentially fruitful venue.

Acknowledgments

This work is supported by IARPA via DOI/IBC contract No. D17PC00342. CL is supported by Google PhD Fellowship, and also acknowledges NVIDIA for awarding a GPU used in this research.

References

- Alcorn, M. A.; Li, Q.; Gong, Z.; Wang, C.; Mai, L.; Ku, W.-S.; and Nguyen, A. 2019. Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects. In *CVPR*, 4845–4854.
- Anderson, P.; Fernando, B.; Johnson, M.; and Gould, S. 2016. Spice: Semantic propositional image caption evaluation. In *ECCV*, 382–398. Springer.
- Athalye, A.; Engstrom, L.; Ilyas, A.; and Kwok, K. 2018. Synthesizing robust adversarial examples. In *ICML*, 284–293.
- Azulay, A., and Weiss, Y. 2018. Why do deep convolutional networks generalize so poorly to small image transformations? *arXiv preprint arXiv:1805.12177*.
- Brown, T. B.; Carlini, N.; Zhang, C.; Olsson, C.; Christiano, P.; and Goodfellow, I. 2018. Unrestricted adversarial examples. *arXiv preprint arXiv:1809.08352*.
- Chang, A. X.; Funkhouser, T.; Guibas, L.; Hanrahan, P.; Huang, Q.; Li, Z.; Savarese, S.; Savva, M.; Song, S.; Su, H.; et al. 2015. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*.
- Chen, P.-Y.; Zhang, H.; Sharma, Y.; Yi, J.; and Hsieh, C.-J. 2017. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, 15–26. ACM.
- Chen, Q.; Qiu, W.; Zhang, Y.; Xie, L.; and Yuille, A. 2018. Sampleahead: Online classifier-sampler communication for learning from synthesized data. *arXiv preprint arXiv:1804.00248*.
- Engstrom, L.; Tran, B.; Tsipras, D.; Schmidt, L.; and Madry, A. 2017. A rotation and a translation suffice: Fooling cnns with simple transformations. *arXiv preprint arXiv:1712.02779*.
- Eykholt, K.; Evtimov, I.; Fernandes, E.; Li, B.; Rahmati, A.; Xiao, C.; Prakash, A.; Kohno, T.; and Song, D. 2017. Robust physical-world attacks on deep learning models. *arXiv preprint arXiv:1707.08945*.
- Gaidon, A.; Wang, Q.; Cabon, Y.; and Vig, E. 2016. Virtual worlds as proxy for multi-object tracking analysis. In *CVPR*, 4340–4349.
- Gilmer, J.; Adams, R. P.; Goodfellow, I.; Andersen, D.; and Dahl, G. E. 2018. Motivating the rules of the game for adversarial example research. *arXiv preprint arXiv:1807.06732*.
- Gu, K.; Yang, B.; Ngiam, J.; Le, Q.; and Shlens, J. 2019. Using videos to evaluate image model robustness. *arXiv preprint arXiv:1904.10076*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, 1026–1034.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Hudson, D. A., and Manning, C. D. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *CVPR*, 6700–6709.
- Ilyas, A.; Engstrom, L.; Athalye, A.; and Lin, J. 2018. Black-box adversarial attacks with limited queries and information. *arXiv preprint arXiv:1804.08598*.
- Johnson, J.; Hariharan, B.; van der Maaten, L.; Fei-Fei, L.; Lawrence Zitnick, C.; and Girshick, R. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, 2901–2910.
- Kanezaki, A.; Matsushita, Y.; and Nishida, Y. 2018. Rotationnet: Joint object categorization and pose estimation using multiviews from unsupervised viewpoints. In *CVPR*, 5010–5019.
- Kar, A.; Prakash, A.; Liu, M.-Y.; Cameracci, E.; Yuan, J.; Rusiniak, M.; Acuna, D.; Torralba, A.; and Fidler, S. 2019. Meta-sim: Learning to generate synthetic datasets. *arXiv preprint arXiv:1904.11621*.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV* 123(1):32–73.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 1097–1105.
- Kurakin, A.; Goodfellow, I.; and Bengio, S. 2016. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*.
- Maaten, L. v. d., and Hinton, G. 2008. Visualizing data using t-sne. *JMLR* 9(Nov):2579–2605.
- Misra, I.; Girshick, R.; Fergus, R.; Hebert, M.; Gupta, A.; and van der Maaten, L. 2018. Learning by asking questions. In *CVPR*, 11–20.
- Narodytska, N., and Kasiviswanathan, S. P. 2016. Simple black-box adversarial perturbations for deep networks. *arXiv preprint arXiv:1612.06299*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 311–318.
- Pei, K.; Cao, Y.; Yang, J.; and Jana, S. 2017. Towards practical verification of machine learning: The case of computer vision systems. *arXiv preprint arXiv:1712.01785*.
- Snoek, J.; Larochelle, H.; and Adams, R. P. 2012. Practical bayesian optimization of machine learning algorithms. In *NeurIPS*, 2951–2959.
- Song, Y.; Shu, R.; Kushman, N.; and Ermon, S. 2018. Constructing unrestricted adversarial examples with generative models. In *NeurIPS*, 8312–8323.

- Srinivas, N.; Krause, A.; Kakade, S. M.; and Seeger, M. 2009. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*.
- Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; and Fergus, R. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- Wang, W.; Yan, M.; and Wu, C. 2018. Multi-granularity hierarchical attention fusion networks for reading comprehension and question answering. In *ACL*, 1705–1714.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8(3-4):229–256.
- Xiang, Y.; Mottaghi, R.; and Savarese, S. 2014. Beyond pascal: A benchmark for 3d object detection in the wild. In *WACV*, 75–82. IEEE.
- Yang, D.; Xiao, C.; Li, B.; Deng, J.; and Liu, M. 2018. Realistic adversarial examples in 3d meshes. *arXiv preprint arXiv:1810.05206*.
- Yang, J.; Ren, Z.; Xu, M.; Chen, X.; Grandall, D.; Parikh, D.; and Batra, D. 2019. Embodied visual recognition. *arXiv preprint arXiv:1904.04404*.
- Yuille, A. L., and Liu, C. 2018. Deep nets: What have they ever done for vision? *arXiv preprint arXiv:1805.04025*.
- Zeng, X.; Liu, C.; Wang, Y.-S.; Qiu, W.; Xie, L.; Tai, Y.-W.; Tang, C.-K.; and Yuille, A. L. 2019. Adversarial attacks beyond the image space. In *CVPR*, 4302–4311.
- Zoph, B., and Le, Q. V. 2017. Neural architecture search with reinforcement learning. In *ICLR*.