

Towards Practical Verification of Machine Learning: The Case of Computer Vision Systems

Kexin Pei^{*}, Yinzhi Cao[†], Junfeng Yang^{*}, Suman Jana^{*}
^{*}Columbia University, [†]Lehigh University

Abstract—Due to the increasing usage of machine learning (ML) techniques in security- and safety-critical domains, such as autonomous systems and medical diagnosis, ensuring correct behavior of ML systems, especially for different corner cases, is of growing importance. In this paper, we propose a generic framework for evaluating security and robustness of ML systems using different real-world safety properties. We further design, implement and evaluate VERIVIS, a scalable methodology that can verify a diverse set of safety properties for state-of-the-art computer vision systems with only blackbox access. VERIVIS leverage different input space reduction techniques for efficient verification of different safety properties. VERIVIS is able to find thousands of safety violations in fifteen state-of-the-art computer vision systems including ten Deep Neural Networks (DNNs) such as Inception-v3 and Nvidia’s Dave self-driving system with thousands of neurons as well as five commercial third-party vision APIs including Google vision and Clarifai for twelve different safety properties. Furthermore, VERIVIS can successfully verify local safety properties, on average, for around 31.7% of the test images. VERIVIS finds up to $64.8\times$ more violations than existing gradient-based methods that, unlike VERIVIS, cannot ensure non-existence of any violations. Finally, we show that retraining using the safety violations detected by VERIVIS can reduce the average number of violations up to 60.2%.

I. INTRODUCTION

Recent advances in Machine Learning (ML) techniques like Deep Learning (DL) have resulted in an impressive performance boost for a broad spectrum of complex, real-world tasks including object recognition, image segmentation, and speech recognition. ML systems are increasingly getting deployed in security- and safety-critical domains such as self-driving cars [7], automated passenger screening [31], and medical diagnosis [19]. Several such systems have already either achieved or surpassed human-level performance on curated test sets.

However, security- and safety-critical systems, besides correctly handling the common cases, must also demonstrate correct behavior for rare corner cases. Despite their significant progress, machine learning systems often make dangerous and even potentially fatal mistakes. For example, a Tesla autonomous car was recently involved in a fatal crash that resulted from the system’s failure to detect a white truck against a bright sky with white clouds [2]. Such incidents demonstrate the need for rigorous testing and verification of ML systems under different settings (e.g., different lighting conditions for self-driving cars) to ensure the security and safety of ML systems.

Most existing testing methods for ML systems involve measuring the accuracy and loss using manually-labeled randomly-

chosen test samples [52]. Unfortunately, similar to traditional software, such random testing approaches are not effective at finding erroneous corner-case behaviors [37], [46]. Moreover, unlike traditional software, the ML decision logic is learned from data and is often opaque even to their designers, which makes the corner-case behaviors more unpredictable than traditional software. Therefore, verifying security, safety, and reliability of ML systems for different corner cases is critical for wide deployment of ML systems.

Ideally, we want to verify an ML model by applying it on all possible inputs including the corner-case ones and checking that the outputs are as expected. However, it is practically infeasible to enumerate through all (current and future) concrete inputs an ML model may receive. Symbolic methods [25], [38] abstract the input space to simplify reasoning, but they cannot scale to complex, real-world ML models. For instance, Reluplex [25] was able to verify only deep neural nets that take five input features and contain about three hundred neurons whereas real-world neural nets can take thousands of input features and contain millions of neurons.

Failing exhaustive input verification, an alternative approach is to take an existing input, alter it, and check or enhance the ML model on the altered input [8], [16], [23], [27], [30], [32], [33], [37], [41], [44], [46], [51]. However, the systems adopting this approach suffer from one or both of the following problems. First, they assume very strong, often unrealistic, attacker models. For instance, adversarial ML assumes that the attacker can change arbitrary pixels of an input image to arbitrary values [8], [16], [27], [30], [32], [33], [41], [44], [51]. Second, they cannot provide strong guarantees about the absence of different types of erroneous behaviors even for a given input. For instance, none can verify that an ML model produces correct outputs on all possible rotations (up to a certain angle) of an image.

This paper presents a general framework that explicitly models attacker capabilities for understanding and verifying the security and robustness of ML systems. In our framework, instead of changing an input arbitrarily, an attacker is constrained to transforming the input only via different transformation functions. Our framework thus makes it straightforward to specify realistic attacker capabilities such as changing brightness/contrast, rotation, smoothing/blurring, or a combination of them for vision systems. In addition, our framework focuses on the input-output safety properties of ML systems rather than their internal states. This design decision results from the difference between ML and traditional systems:

detailed specifications describing an ML system’s internal states are hard to write even for its developers, but safety properties involving input-output behaviors are intuitive and easy to specify. For example, while it is extremely hard (if not impossible) to recreate the logic of a human driver, it is easy to envision safety properties like a self-driving cars steering angle should not change significantly for the same road under different lighting conditions.

We further design, implement, and evaluate VERIVIS, a realization of this framework for computer vision systems, an important class of ML systems. Given an input image, VERIVIS exhaustively transforms the image via a given transformation function (e.g., rotation) and generates all possible transformed inputs, verifying that the output of ML system on each transformed input satisfies certain safety properties. Any violations found by VERIVIS come with concrete counter-example inputs to help debugging. To the best of our knowledge, VERIVIS represents the first methodology for verifying a wide range of realistic safety properties (e.g., invariance of a self-driving cars steering angle under different brightness conditions or rotation invariance for image classifiers) on state-of-the-art computer vision systems.

A key challenge in VERIVIS is that a transformation on one input may still yield a practically infinite number of transformed inputs, preventing exhaustive verification. For instance, an image can be rotated by any rational number of degrees. Fortunately, the key insight behind our approach is that ML, unlike traditional software, commonly operates on discretized inputs such as images, video, speech. For example, a pixel’s image coordinates must be integer values bound by the image width and height, and its RGB values also integer values between 0 and 255. Thus, many tiny variations in rotation degrees will still yield the same image after rounding.

Leveraging this insight, we significantly reduce the verification search space for a diverse set of real-world image transformations such as rotation, blurring, and lightening by breaking their parameter space into equivalent classes and verifying the transformation result of only one representative parameter from each class without compromising the verification guarantees. We show that the total number of such representative parameters, or *critical parameters*, for any image (irrespective of the contents) is polynomial (up to cubic) to the image size for many common transformations, demonstrating scalability with increasing input size. For instance, for a 299×299 input image for Inception-v3 [43] and a maximum of -2 to 2 degrees of rotation, we need to test at most 95,496 (as opposed to infinitely many) different rotation degrees in the $[-2, 2]$ range to fully verify that the ML system satisfies certain safety properties for the given image. While we primarily focus on designing and evaluating VERIVIS to find violations of safety properties in different computer vision systems but the underlying principles are generic and can be applied to other ML domains as well.

By focusing on input-output behaviors and making the search space reduction process not dependent on an ML system, VERIVIS enables an ML-model-agnostic verification ap-

proach. Specifically, given a set of inputs and transformations, VERIVIS pre-computes the entire set of possible transformed inputs, which can be applied to verify the presence or absence of corner-case errors for any computer vision system on the given input set. The main benefit over ML-model-dependent approaches is that the verification effort primarily depends on the type of safety property, input size, and input domain but not on the complexity of the ML system itself.

We evaluate VERIVIS with twelve different safety properties involving real-world transformations like rotation, contrast, and brightness for fifteen state-of-the-art computer vision systems: (1) six image classification DNNs trained on the imagenet dataset (VGG-16 [42], VGG-19 [42], MobileNet [22], Xception [10], Inception-v3 [43] and ResNet50 [20]), which achieved the state-of-the-art performances in ILSVRC [40] competitions; (2) five commercial third-party vision APIs provided by Google [17], Clarifai [12], IBM [24], Microsoft [29], and Amazon [3]; (3) four popular self-driving car models including Rambo, one of the top performers in Udacity self-driving car challenge, and other three are based on Nvidia Dave self-driving systems. For all of these systems, VERIVIS finds thousands of violations for different input images— $64.8\times$ more than existing gradient-based methods. For a small number of images (on average 31.7% of all tested images), VERIVIS was able to verify that the tested ML systems satisfy the tested safety properties. Our results further demonstrate that by specifically retraining on the violated images generated by VERIVIS, the robustness of the corresponding ML systems can be significantly improved by reducing up to 60.2% of the violations.

Our main contributions are:

- We define a general framework for specifying realistic safety properties of ML systems modeling different types of attacker capabilities.
- We present a novel decomposition and search space reduction methodology for efficient verification different types of safety properties for computer vision systems.
- We implement our techniques as part of VERIVIS and use it to verify twelve safety properties on fifteen state-of-the-art vision systems including Imagenet classifiers, self-driving cars, and commercial vision APIs. VERIVIS finds thousands of violations in all tested systems (see Figure 1 for some samples).

II. A FRAMEWORK FOR VERIFYING ML ROBUSTNESS AND SECURITY

In this section, we present a general framework for verifying the robustness and security of ML systems against attackers with different capabilities. First, we describe our rationale of focusing on input-output behaviors instead of the internal states of ML systems. Next, we present the framework formally and explain its generality. Lastly, we discuss potential implementations of this framework.

Why focus on input-output verification. Traditional and ML programs are very different from each other in terms of verification. Traditional programs are *operational*: each




Fig. 1. Upper row shows the original inputs. Lower row shows the violations generated by different transformations for different computer vision systems as identified by VERIVIS.

operation in the program is defined clearly by the programmers with the intention to perform one sub-task of the overall program task. Thus, it is conceptually feasible (albeit costly) to (1) write a complete functional specification to ensure the total correctness of the program and (2) specify pre- and post-conditions for each operation to ensure the correctness of the intermediate states.

In contrast, the logic in ML systems is learned from data using an *automated optimization* process. Its logic is often opaque even to its own designers, making it challenging to write complete specifications. For instance, creating a complete specification for the correct behavior of a self-driving car under different driving conditions essentially involves recreating the logic of a human driver, which is computationally infeasible and not practical. In addition, the nature of optimization means that, there are multiple acceptable ways through different internal states for satisfying the final goal. For instance, a car can be safely driven on the road with many slightly different but similar steering angles. Therefore, we conjecture that ML verification should focus on partial, input-output correctness rather than complete functional correctness.

Input-output based safety properties are applicable to a wide range of ML systems. For example, consider a safety property specifying that a self-driving car’s steering angle should remain similar for the same road under different lighting conditions. In this setting, even though it is hard for an ML developer to predict the safe steering angles for the self-driving car under different scenarios, it is easy to specify such safety properties. Similarly, a safety property for a speech recognition system can ensure that the recognized phrases/sentences will not change under different background noises. For a machine translation system, the translated output should not change significantly if the input words are substituted with other synonyms. Malware detection systems should not change their classifications from malware to benign due to different types of code obfuscation/transformation techniques that do not affect malicious functionality [53].

A formal framework for specifying ML safety properties. An ML model can be thought of as a function f mapping

input $\mathbf{x} \in \mathbb{X}$ to output $\mathbf{y} \in \mathbb{Y}$, i.e., $f : \mathbb{X} \rightarrow \mathbb{Y}$. Depending on the type of task (i.e., classification or regression), the ML model produces either continuous or discrete output \mathbb{Y} . For classification tasks like object recognition, \mathbb{Y} is a set of discrete labels. By contrast, \mathbb{Y} is a continuous range for regression tasks like driving an autonomous vehicle that outputs steering angles.

Let $\mathcal{T}(\cdot; c)$ be a transformation function parameterized by $c \in \mathbb{C}$ (\mathbb{C} is transformation parameter space) that transforms an input \mathbf{x} to $\mathbf{x}' = \mathcal{T}(\mathbf{x}; c)$. Let P be a safety property checker that inspects a transformed input \mathbf{x}' for errors.

The safety properties of a ML system based on its input-output behaviors can then be defined as follows:

Locally safe. Given a model f , an input \mathbf{x} , a parameterized transformation function $\mathcal{T}(\cdot; c)$ where $c \in \mathbb{C}$, and a property checker P , f is locally safe for input \mathbf{x} if and only if $\forall c \in \mathbb{C}, P(f(\mathcal{T}(\mathbf{x}; c)))$.

Globally safe. A model f is globally safe for all inputs if and only if $\forall \mathbf{x} \in \mathbb{X}$, it is locally safe. Global safety over all inputs is in general much more challenging to achieve than local safety on a particular input.

Our framework is expressive. The transformation function $\mathcal{T}(\cdot; c)$ and its parameter space \mathbb{C} can be used to specify different attacker capabilities. It can be physically realizable transformations such as occluding a small part of an input image, changing lighting conditions, or rotating an image. The property checker P can check that the steering angle of a self-driving car on \mathbf{x}' does not change significantly from that on \mathbf{x} or the car does not get into a crash. Below we give the definitions of two safety properties that are used in our evaluation:

Locally k -safe classification. Given a classification model f , an input \mathbf{x} , a parameterized transformation function $\mathcal{T}(\cdot; c)$ where $c \in \mathbb{C}$, we define that f is k -safe with respect to $\mathcal{T}(\cdot; c)$, \mathbb{C} , and \mathbf{x} if and only if $\forall c \in \mathbb{C}, f(\mathcal{T}(\mathbf{x}; c), 1) \subseteq f(\mathbf{x}, k)$ where $f(\mathbf{x}, k)$ denotes the top- k prediction by f for \mathbf{x} .

Locally t -safe regression. Given a regression model f , an input \mathbf{x} , a parameterized transformation function $\mathcal{T}(\cdot; c)$ where $c \in \mathbb{C}$, we define that f is locally t -safe with respect to $\mathcal{T}(\cdot; c)$,

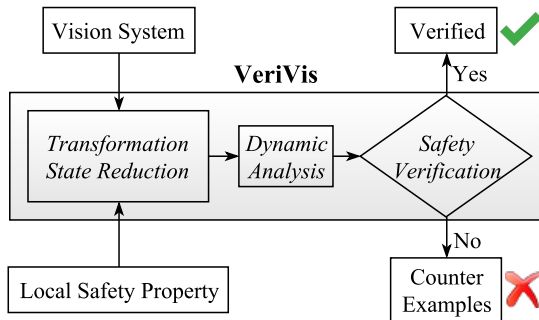


Fig. 2. Overall verification architecture of VERIVIS.

\mathbb{C} , and \mathbf{x} if and only if $\forall c \in \mathbb{C}, |f(\mathcal{T}(\mathbf{x}; c)) - f(\mathbf{x})| \leq t$.

Robustness of ML models to adversarial inputs [16] can also be easily expressed as safety properties where the transformation function $\mathcal{T}(\mathbf{x}; c) = \mathbf{x} + c$. In this case, the transformation parameter c is an attacker-crafted perturbation (a tensor variable that have the same size with the input \mathbf{x}). The transformation parameter space includes all c such that different types of norms of c is bounded by a user-defined budget σ , e.g., $\|c\|_0 < \sigma$ [33], $\|c\|_1 < \sigma$ [53], $\|c\|_2 < \sigma$ [44], or $\|c\|_\infty < \sigma$ [16], [53]. The search space of the adversarial perturbations are significantly larger than real-world transformations like rotation, changing brightness, etc. In this paper, we focus on verifying local safety-critical properties involving real-world transformations [37], [41].

Potential implementations of our framework. Verifying that an ML system satisfies a particular safety property ϕ is essentially a search problem over the parameter space \mathbb{C} of a transformation $\mathcal{T}(\cdot; c)$ (i.e., searching all possible values of $c \in \mathbb{C}$) for violations of the safety property. The size of the search space for a given ML safety property varies widely based on the property and input size/domain. For example, a local safety property involving the image inversion operation for image classification systems will have a search space of only one transformed image for any given image. By contrast, a transformation function that simulates different lighting conditions along with different shadows for an object in a given image will have a significantly larger search space [6], [54]. In many cases, such large search spaces may only be checked probabilistically for safety property violations.

The verification techniques for checking ML safety properties can use either static, symbolic, or dynamic approach. However, existing static/symbolic techniques for traditional software does not scale well for ML systems due to the highly non-linear nature of the ML systems. In this paper, we focus on dynamic and model-agnostic verification for ML systems. However, if needed, our approach can be easily augmented with the model knowledge to further reduce the search space.

III. METHODOLOGY

A. VERIVIS overview

In this paper, we design, implement, and evaluate VERIVIS, an instance of the verification framework described in Section II specifically tailored for checking different local safety properties of computer vision systems such as self-driving cars

and image classifiers. Figure 2 shows the high-level structure of VERIVIS. VERIVIS takes a computer vision system (e.g., self-driving car or image classifier) along with a local safety property as inputs and either verifies that the computer vision system satisfies the safety property or produces a counter-example violating the safety property. Even though we primarily focus on verifying computer vision systems in this paper, the underlying principles are applicable to other types of ML systems operating on discrete and bounded input domains as well (e.g., malware detection).

VERIVIS uses blackbox dynamic analysis for verifying safety properties of computer vision systems without knowing its internal details. For example, VERIVIS can even verify cloud vision APIs provided by companies like Google or Microsoft without any information about their internals (see Section VI for details). Specifically, VERIVIS applies different input space reduction techniques to skip unrealizable inputs and efficiently verify a computer vision system by only checking for the inputs that can be feasibly generated by the image transformation specified in the safety property. Such input space reduction is possible because even though the image transformation parameters are continuous (e.g., floating point values), their output domains (e.g., the pixel values) are discrete. Therefore, multiple different values of a transformation parameter may lead to the same image and therefore can be safely skipped without affecting the verification guarantees. For example, while the degree of image rotation can be a floating point value, the pixel values of the output image are limited and discrete (e.g., integers ranging between 0 and 255).

In general, we show that the constraints of a wide range of realistic image transformations (e.g., rotation, changing lighting condition) imply that the space of unique images that can be generated by these transformations is polynomial in the input image size for a fixed parameter range. Therefore, VERIVIS can verify a broad spectrum of safety properties while scaling up to the largest state-of-the-art computer vision systems. VERIVIS’s approach is conceptually similar to explicit-state model checking [21] for traditional software where feasible and unique states of a traditional program are explicitly enumerated for safety property violations.

B. Reducing input spaces of image transformations

Each safety property supported by VERIVIS has a corresponding parameterized image transformation associated with the property. Similar to the transformations described in Section II, a parameterized image transformation takes an input image I , a parameter value c , and produces an output image $O = \mathcal{T}(I; c)$. Essentially, an image transformation computes the value of each output pixel based on the values of the input pixels and the parameter. In theory, for an arbitrary image transformation, each output pixel may depend on all of the input pixels. However, in practice, an output pixel’s value usually only depends on a small number of neighboring pixels in the input image for most image transformations designed to generate realistic images (e.g., rotation, changing brightness/contrast, erosion). This property allows us to drastically




Fig. 3. Decomposition of a simple one-dimensional transformation (e.g., one-dimensional blurring) into Dependent Pixels (\mathcal{DP}) and Dependence Function (\mathcal{DF}). White squares represent the pixels.

reduce the search space for safety property verification. Note that this is an inherent property of most image transformations designed to produce realistic images as most physically realizable changes (i.e., changes in camera orientation, lighting conditions, etc.) tend to satisfy this property.

Decomposition framework. Before describing the details of the input space reduction process, we first define a generic decomposition framework for reasoning about the space of all distinct output images for a given parameterized image transformation and a parameter space. Our framework relies on the fact that most image processing operations can be decomposed into a multi-step pipeline of stencil operations, where each point in the two-dimensional space is updated based on weighted contributions from a subset of its neighbors [39]. Our framework decomposes parameterized image transformations into a sequence of parameterized stencil operations for efficient search space reduction. Even though decomposition of image processing code into stencil computations have been used to optimize the performance of the code by Ragan-Kelley et al. [39], to the best of our knowledge, we are the first to reduce the output space of a parameterized image transform using such techniques.

Specifically, we describe two pixel-specific parameterized functions that can express the relations between the input and output pixels of different realistic image transformations: *Dependent Pixels* ($O_{DP} = \mathcal{DP}(\mathbf{I}, \langle i, j \rangle; c)$) and *Dependence Function* ($O_{DF} = \mathcal{DF}(\mathbf{I}(O_{DP}); c)$) where $\mathbf{I}(O_{DP})$ denotes the pixel values of coordinates output by \mathcal{DP} in image \mathbf{I} . For a given input image \mathbf{I} and a coordinate $\langle i, j \rangle$, dependent pixels (\mathcal{DP}) return a list of pixel coordinates in the output image whose values are dependent on the input pixel’s value for a given parameter value c . Dependence function (\mathcal{DF}) takes the pixel values of input image \mathbf{I} on coordinates O_{DP} as input and computes the value of the corresponding output pixels based on the input pixel values for a given parameter value c . Essentially, all possible output images for an image transformation $\mathcal{T}(\mathbf{I}; c)$ of an image \mathbf{I} of size $W \times H$ and $c \in \mathbb{C}_\phi$ can be enumerated by evaluating Equation 1 (\mathbb{C}_ϕ represents the user-specified parameter space of a transformation for safety

property ϕ).

$$\bigcup_{\forall c \in \mathbb{C}_\phi} \bigcup_{i=0}^{W-1} \bigcup_{j=0}^{H-1} \mathcal{DF}(\mathbf{I}(\mathcal{DP}(\mathbf{I}, \langle i, j \rangle; c)); c) \quad (1)$$

Figure 3 shows a sample decomposition of a simple one-dimensional transformation into \mathcal{DP} and \mathcal{DF} . In this case, an output pixel always depends only on three neighboring input pixels. However, for an arbitrary image transformation, \mathcal{DP} , for any given input pixel, might potentially produce all output pixels, i.e., any output pixel might be dependent on all input pixels. Similarly, \mathcal{DF} can also produce all possible values (e.g., 0 to 255) for each output pixel. Therefore, in the worst case, the number of all unique output images that can be generated by an arbitrary transformation can be $256^{W \times H}$ for a gray-scale input image with size $W \times H$. Even for simple networks designed to work on small images like the MNIST hand-written digits each with 28×28 pixels, the resulting number of output images ($256^{28 \times 28}$) will be too large to enumerate exhaustively.

However, for most realistic image transformations, \mathcal{DP} produces only a few dependent output pixels for each input pixel and \mathcal{DF} outputs only a subset of all possible output pixel values (e.g., 0 – 255). For a wide range of real-world transforms, we demonstrate that the number of unique output images for a given input image and parameter range is polynomial in the size of the image as shown in Section IV and Table VI. VERIVIS enumerates these output images efficiently and therefore is able to verify a wide range of safety properties in a scalable manner.

To understand how \mathcal{DP} and \mathcal{DF} look for realistic image transformations, consider a simple parameterized brightness transformation that brightens an image by adding a constant parameter value to all pixel values. For such a transformation, the output pixel coordinate of \mathcal{DP} will be same as the input pixel coordinate as an output pixel’s value only depends on the input pixel’s value with the same coordinate. Similarly, \mathcal{DF} will simply add the parameter value with the current value of the input pixel.

Reducing the parameter space. As most the parameters of most image transformations are continuous floating point values, enumerating all possible parameter values in a given parameter space is not feasible. For example, a safety property might specify that an image can be rotated to any arbitrary angle between -10° and 10° . Enumerating all possible rotation angles within the specified parameter range is too slow to be practical.

The key insight behind our input space reduction technique is that the output of \mathcal{DP} and \mathcal{DF} both are discrete, bounded integers (e.g., \mathcal{DP} must return valid pixel coordinates and \mathcal{DF} must return valid pixel values) even though the parameter values are continuous floating point numbers. The discreteness of the output allows us to only enumerate a small finite set $\mathbb{C}_{critical} \in \mathbb{C}$ of critical parameter values that can cover all possible unique output images. We formally define the critical parameter values below.

Definition 1. Critical parameter values in a parameter space \mathbb{C}_ϕ and transformation \mathcal{T} in safety property ϕ is a monotonic increasing finite sequence $(c_i \in \mathbb{C}_\phi)_{i=1}^n$, i.e., $\forall i \in \{1, 2, \dots, n-1\}$, $c_{i+1} > c_i$, where the following holds. $\forall c_i, c_{i+1}$ where $i \in \{0, \dots, n-1\}$, $\forall c \in \mathbb{C}_\phi$ where $c_i < c < c_{i+1}$, $\mathcal{T}(\mathbf{I}; c) = \mathcal{T}(\mathbf{I}; c_i)$ or $\mathcal{T}(\mathbf{I}; c) = \mathcal{T}(\mathbf{I}; c_{i+1})$.

Definition 1 ensures that for any continuous c between two critical parameter values, the transformation \mathcal{T} specified in a safety property ϕ will not generate a new output other than those already generated by the immediately preceding and following critical parameter values. For example, consider the image translation operation where shift step is the parameter. Translation shifts an image to the desired direction by the desired amount. If the shift step is set to any floating point number, the translated coordinates will be rounded (up) to the nearest integer values. Therefore, all $\langle i, j \rangle$ pairs where $0 \leq i \leq W-1$ and $0 \leq j \leq H-1$ will be critical parameter values for the translation operation.

VERIVIS explicitly enumerates critical parameter values $\mathbb{C}_{critical} = (c_i)_{i=1}^n$ ignoring the continuous values that lie between the critical parameter values. Algorithm 1 shows the detailed procedure of finding the critical parameter values for \mathcal{DP} and \mathcal{DF} for a given transformation \mathcal{T} and parameter space \mathbb{C}_ϕ . \mathcal{DP}^{-1} and \mathcal{DF}^{-1} indicate inverses of \mathcal{DP} and \mathcal{DF} respectively, i.e., they compute c given input and output pixels. However, as discussed in Section III-B, more than one values of c can often map to same O_{DP} or O_{DF} . Therefore, we assume that \mathcal{DP}^{-1} and \mathcal{DF}^{-1} randomly samples one value from the set of all candidates for c .

Note that the critical parameter values for more complicated image transformations like rotation, unlike translation, may not be equidistant from each other. We describe how \mathcal{DP} , \mathcal{DF} , \mathcal{DP}^{-1} and \mathcal{DF}^{-1} are computed for different transformations in detail in Section IV.

Algorithm 1 Algorithm for computing critical parameter values in parameter space \mathbb{C}_ϕ for transformation \mathcal{T} in property ϕ .

Input: \mathbb{C}_ϕ

- 1: $\mathbb{C}_{critical} \leftarrow \{\}$
- 2: **for** all coordinates $\langle i, j \rangle$ in \mathbf{I} **do**
- 3: **for** all feasible $O_{DP} \in \mathcal{DP}(\mathbf{I}, \langle i, j \rangle; \mathbb{C}_\phi)$ **do**
- 4: $c = \mathcal{DP}^{-1}(\mathbf{I}, \langle i, j \rangle; O_{DP})$
- 5: **if** $c \in \mathbb{C}_\phi$ **then**
- 6: $\mathbb{C}_{critical} = \mathbb{C}_{critical} \cup c$
- 7: **for** all feasible $O_{DF} \in \mathcal{DF}(\mathbf{I}(O_{DP}); c)$ **do**
- 8: $c = \mathcal{DF}^{-1}(\mathbf{I}(O_{DP}); O_{DF})$
- 9: **if** $c \in \mathbb{C}_\phi$ **then**
- 10: $\mathbb{C}_{critical} = \mathbb{C}_{critical} \cup c$
- 11: **return** sorted($\mathbb{C}_{critical}$)

IV. DECOMPOSITION & ANALYSIS OF REAL-WORLD IMAGE TRANSFORMATIONS

In this section, we describe how VERIVIS supports verification of a wide range of safety properties with different real-world image transformations. Specifically, we describe twelve different image transformations corresponding to twelve safety

properties $(\phi_1, \dots, \phi_{12})$ summarized in Table I. These image transformations and their compositions can simulate a wide range of real-world distortions, noises, and deformations that most security-critical vision systems must handle correctly. The transformation parameters for each transformation are shown in the third column of Table I.

These transformations can be broadly categorized into three groups: convolutions, point transformations, and geometric transformations. Convolution-based transformations like blurring (e.g., ϕ_1 to ϕ_4) applies a convolution kernel on the input image and produce the output images such that each pixel value is determined by its local neighbors and the corresponding kernel weights. By contrast, for point transformations (e.g., ϕ_5 and ϕ_6), each pixel's new value is only decided by its original value in the input. Finally, geometric transformations (e.g., ϕ_7 to ϕ_{12}) shuffles the pixel values based on different geometric constraints.

As described in Section III, VERIVIS reduces the input space of these twelve transformations by decomposing them into \mathcal{DP} and \mathcal{DF} as defined in Equation 1 and finding the corresponding critical parameter values using Algorithm 1. We describe the decomposition process for each transformation in detail below. We also perform the verification complexity analysis for each transformation and demonstrate that safety properties related to all of these transformations can be verified in polynomial time with respect to the input image size as shown in Table II.

A. Convolutions

Decomposition. For all convolution-based transformations with a kernel of size c , $\mathcal{DP}(\mathbf{I}, \langle i, j \rangle; c) = \{\langle k, l \rangle : i - c/2 < k < i + c/2 \text{ and } j - c/2 < l < j + c/2\}$ defining a square area surrounding the pixel at $\langle i, j \rangle$. By contrast, \mathcal{DF} will depend on the actual operation of the transformation. For ϕ_1 , ϕ_2 , ϕ_3 , and ϕ_4 , \mathcal{DF} computes the average, median, minimum, and maximum of the pixel values of the coordinates returned by \mathcal{DP} ($O_{\mathcal{DP}}$), respectively.

Critical parameter values. The possible sizes of a convolution kernel (a square) for an image with width W and height H can vary from 2×2 to $S \times S$ where $S = \min(W, H)$ because kernel cannot be larger than the input image. As kernel sizes have to be integers, the output of \mathcal{DP} on each pixel $\langle i, j \rangle$ can be 2×2 , 3×3 , ..., $S \times S$ pixels (i.e., $S-1$ different values) surrounding the input pixel. Moreover, \mathcal{DF} for ϕ_1 to ϕ_4 does not depend on the kernel size. Therefore, these convolution-based transformations (ϕ_1 to ϕ_4) have $S-1$ critical parameter values, i.e., $\mathbb{C}_{critical} = \{c \in \mathbb{N} : 2 \leq c \leq \min(W, H)\}$.

Verification complexity. As the number of different convolution kernel sizes can only be integer values and is bounded by the image size, the number of unique output images is $O(n)$ where the input image size is n .

B. Point transformations

Decomposition. Point transformations (e.g., ϕ_5 and ϕ_6) are simple pixel-space operations. Therefore, \mathcal{DP} for all these cases outputs the input coordinate itself, i.e.,

TABLE I
A LIST OF SAFETY PROPERTIES AND CORRESPONDING
TRANSFORMATIONS THAT CAN SIMULATE A WIDE RANGE OF COMMON
REAL-WORLD IMAGE DISTORTIONS AND DEFORMATIONS.







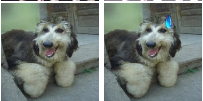
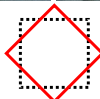
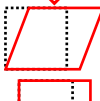
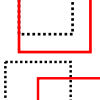

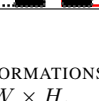
Property	Transformation	Parameters	Example
ϕ_1	Average smoothing	Kernel size	
ϕ_2	Median smoothing	Kernel size	
ϕ_3	Erosion	Kernel size	
ϕ_4	Dilation	Kernel size	
ϕ_5	Contrast	Gain	
ϕ_6	Brightness	Bias	
ϕ_7	Occlusion	Coordinate	
ϕ_8	Rotation	Rotation angle	
ϕ_9	Shear	Proportion	
ϕ_{10}	Scale	Scalar	
ϕ_{11}	Translation	Shift step	
ϕ_{12}	Reflection	Direction	

TABLE II
VERIFICATION COMPLEXITY OF DIFFERENT TRANSFORMATIONS WITH
RESPECT TO THE INPUT IMAGE SIZE $n = W \times H$.

Transformation	Verif. Complexity	Transformation	Verif. Complexity
Avg. smoothing	$O(n)$	Med. smoothing	$O(n)$
Erosion	$O(n)$	Dilation	$O(n)$
Contrast	$O(1)$	Brightness	$O(1)$
Occlusion	$O(n)$	Rotation	$O(n^2)$
Shear	$O(n^3)$	Scale	$O(n^2)$
Translation	$O(n)$	Reflection	$O(1)$

$\mathcal{DP}(\mathbf{I}, \langle i, j \rangle; c) = \{\langle i, j \rangle\}$. Essentially, \mathcal{DP} is an identity function for these transformations. By contrast, \mathcal{DF} depends on the functionality of each transformation. For ϕ_5 , $\mathcal{DF}(\mathbf{I}(\langle i, j \rangle); c) = c \cdot \mathbf{I}(\langle i, j \rangle)$, where c is the gain used to adjust the contrast of the input image [1]. Similarly, for ϕ_6 , $\mathcal{DF}(\mathbf{I}(\langle i, j \rangle); c) = c + \mathbf{I}(\langle i, j \rangle)$, where c is the bias used to

adjust the brightness of the input image [45].

Critical parameter values. As noted above, \mathcal{DP} for both ϕ_5 and ϕ_6 is an identity function and is independent of c . Therefore, \mathcal{DP} does not affect the number of critical parameters.

For ϕ_5 , \mathcal{DF} change the image contrast by multiplying c to each pixel value. Therefore, \mathcal{PF} is a function mapping any pixel value in $[0, 255]$ to a new value in $[0, 255]$ by multiplying c . It is easy to see that at most $|\mathbb{C}_{critical}| = 256 \times 256$ critical parameter values are enough to cover all such unique mappings. Specifically, $\mathbb{C}_{critical} \leq \bigcup_{m=0}^{255} \bigcup_{n=0}^{255} \frac{m}{n}$ where critical parameter values resulting in invalid (e.g., division by zero) or duplicate values can be further reduced.

Similarly, \mathcal{DF} for ϕ_6 is a function mapping any pixel value in $[0, 255]$ to $[0, 255]$ by adding c . Therefore $\mathbb{C}_{critical}$ for ϕ_6 is $\{-255, -254, \dots, 254, 255\}$.

Verification complexity. As shown in the earlier analysis, the number of critical parameter values for ϕ_5 and ϕ_6 does not depend on the input image size n . Therefore, the total number of critical parameter values $\mathbb{C}_{critical}$ has $O(1)$ complexity with respect to the image size.

C. Geometric transformations

Decomposition. In this paper, we analyze five types of geometric transformations—occlusion with a predefined mask image (ϕ_7), rotation (ϕ_8), shear (ϕ_9), scale (ϕ_{10}), translation (ϕ_{11}), and reflection (ϕ_{12}). For all of these transformations, \mathcal{DP} maps one coordinate to another within the image, i.e., $\mathcal{DP}(\mathbf{I}, \langle i, j \rangle; c) = \{\langle i', j' \rangle\}$, and \mathcal{DF} is an identity function, i.e., $\mathcal{DF}(\mathbf{I}, \langle i', j' \rangle; c) = \mathbf{I}(\langle i', j' \rangle)$. We describe the individual \mathcal{DP} function for each transformation below.

For occlusion (ϕ_7), $\mathcal{DP}(\mathbf{I}, \langle i, j \rangle; c) = \{\langle i, j \rangle\}$ if $i \notin [c_W, c_W + W_{OcclMask}]$ and $j \notin [c_H, c_H + H_{OcclMask}]$. If a pixel's coordinates are within this range, its value is decided by the occlusion mask and is independent of any pixel value in the input image. Here, (c_W, c_H) denotes the coordinate of the upper-left corner of the image where the occlusion mask is applied and $W_{OcclMask}$ and $H_{OcclMask}$ denote the widths and heights of the occlusion mask respectively.

For rotation (ϕ_8), $\mathcal{DP}(\mathbf{I}, \langle i, j \rangle; c) = \{\langle i \cdot \cos c - j \cdot \sin c, i \cdot \sin c + j \cdot \cos c \rangle\}$ where c is the rotation degree. Note that we only consider rotation around the center of the image here but \mathcal{DP} for rotation around arbitrary points can also be constructed in the same manner.

Similarly, for shear (ϕ_9), $\mathcal{DP}(\mathbf{I}, \langle i, j \rangle; c) = \{\langle i + jc_W, ic_H + j \rangle\}$, where $c = (c_W, c_H)$ are the horizontal and vertical shear parameters.

For scale (ϕ_{10}), $\mathcal{DP}(\mathbf{I}, \langle i, j \rangle; c) = \langle ic_W, jc_H \rangle$, where $c = (c_W, c_H)$ are the horizontal and vertical scale parameters.

For translation (ϕ_{11}), $\mathcal{DP}(\mathbf{I}, \langle i, j \rangle; c) = \langle i + c_W, j + c_H \rangle$, where $c = (c_W, c_H)$ are the horizontal and vertical shifting parameters.

Finally, for reflection (ϕ_{12}), $\mathcal{DP}(\mathbf{I}, \langle i, j \rangle; c) = \langle ic_W, jc_H \rangle$, where $c = (c_W, c_H) \in \{(-1, 1), (1, -1), (-1, -1)\}$, represent three types of reflections (horizontal, vertical, and central).




Fig. 4. Illustration of how rotation shuffles the coordinates around for a 3×3 image. Black dots represent the coordinate of each pixel. The circle shows the trajectory of coordinate $\langle i, j \rangle$ during rotation with different angles. Any intersection of the circular trajectory with a dotted line (right at the middle between horizontal and vertical coordinates) will be rounded to the nearest coordinate. For a new rounded coordinate $\langle i', j' \rangle$, the formula for calculating the corresponding critical rotation angle is shown on the right.

Critical parameter values. As the size and content of the occlusion mask for ϕ_7 are fixed, the number of different outputs of \mathcal{DP} depend on where the occlusion mask is applied as decided by $c = (c_W, c_H)$. Therefore, $\mathbb{C}_{critical} = \{\langle i, j \rangle : i \in [0, W - W_{OcclMask}], j \in [H - H_{OcclMask}]\}$, which are all possible coordinates where $OcclMask$ can be applied. As the coordinates have to be integers, the number of critical parameter values is $(W - W_{OcclMask} + 1) \times (H - H_{OcclMask} + 1)$.

We describe the computation of the critical parameter values for ϕ_8 , the most complex one among ϕ_8 to ϕ_{12} , in detail below and skip the details for the other transformations (ϕ_9 - ϕ_{12}) as they are similar to ϕ_8 .

Figure 4 shows the movement of the coordinates (along the circle) while rotating an image around the center of the image. Consider the coordinate $\langle 1, 0 \rangle$, (e.g., $\langle i, j \rangle$ in Figure 4) is being rotated, \mathcal{DP} will only output new coordinates when the circular trajectory intersects with the dotted lines (e.g., vertical $x = \{\dots, -0.5, 0.5, \dots\}$ or horizontal $y = \{\dots, -0.5, 0.5, \dots\}$). Before intersecting with the next dotted line, all the rotated coordinates that correspond to the parts of the trajectory between any two adjacent dotted lines with arrows will always be rounded to the same coordinate. In Figure 4, $\langle i', j' \rangle = \langle \sqrt{3}/2, 1/2 \rangle$ will be rounded (up) to $\langle 1, 1 \rangle$. Thus, \mathcal{DP} can only output seven new coordinates in Figure 4. Therefore, given any $\langle i, j \rangle$, we compute how many dotted lines intersect with the corresponding trajectory. We calculate the rotation degree (c) for each intersecting coordinate $\langle i', j' \rangle$ using the following equation.

$$c = \arccos \frac{ij + i'j'}{\sqrt{i^2 + j^2} \sqrt{i'^2 + j'^2}}$$

Finally, we compute the union of all possible rotation degrees of each coordinate that can be mapped to a new one by \mathcal{DP} following Algorithm 1, where the resulting union is the final critical parameter values for rotation ($\mathbb{C}_{critical}$).

Verification complexity. The occlusion mask can only be applied on integer coordinates within the input image. Therefore, as described above, the number of critical parameter values for ϕ_7 is $(W - W_{OcclMask} + 1) \times (H - H_{OcclMask} + 1)$. Therefore, the verification complexity is $O(n)$.

As shown in Figure 4, for ϕ_8 , each trajectory of a coordinate can only be changed to at most $(w - 1) \times (h - 1)$ number of new coordinates. As the input has $w \times h$ number of distinct coordinates, the number of possible changes is $O(w^2 \cdot h^2) = O(n^2)$.

For ϕ_9 , \mathcal{DP} outputs $\langle i + jc_W, ic_H + j \rangle$ for each $\langle i, j \rangle$. Note that $i + jc_W$ can have at most $w \cdot h \cdot w$ number of valid values (w possible values for i , h for j , and w for $i + jc_W$). The same analysis applies to possible values of $ic_H + j$ as well, i.e., the number of valid values is $w \cdot h \cdot h$. In total, there are $w^3 \cdot h^3 = O(n^3)$ total possible pairs of (c_W, c_H) .

For ϕ_{10} , \mathcal{DP} outputs $\langle ic_W, jc_H \rangle$ for input $\langle i, j \rangle$. For a given i and j , ic_W can have w distinct values of ic_W and jc_H can also have h distinct values. Since there is $w \cdot h$ numbers of possible $\langle i, j \rangle$ pairs, the verification complexity will be $w^2 \cdot h^2 = O(n^2)$.

For ϕ_{11} , the transformation can only shift the image within the original image size and all the critical parameter values must be integers. Therefore, ϕ_{11} has $O(n)$ verification complexity.

Finally, ϕ_{12} is a special case where there are only three types of reflection operations. Thus, ϕ_{12} has $O(1)$ verification complexity.

V. IMPLEMENTATION

Our implementation of VERIVIS consists of 4,465 lines of Python code. VERIVIS uses OpenCV, the popular image processing and vision library, to implement efficient image transformations. All of our experiments are run on a Linux laptop with Ubuntu 16.04 (one Intel i7-6700HQ 2.60 GHz processor with 4 cores, 16 GB memory, and a NVIDIA GTX 1070 GPU). To significantly cut down verification time, we also implemented batch prediction [9], using both GPU and CPU, to make the target computer vision system predict multiple images in parallel.

We evaluate VERIVIS with 12 different safety properties with different transformations. We use VERIVIS to verify 15 total vision systems including 10 popular pre-trained DNNs performing tasks like object recognition, autonomous driving, etc., and 5 third-party blackbox image recognition services using API access. Table III shows a summary of all 15 computer vision systems and the corresponding datasets used for verification in our experiments. These systems can be categorized into three groups based on the tasks they perform and the type of access they provide (e.g., API access vs. the trained model). We describe them in detail below.

1000-class ImageNet classification. This group of vision systems use DNNs trained using the ImageNet [15] dataset. All of these DNNs achieved state-of-the-art image classification performance in ILSVRC [40] competitions. Specifically, we verify the following six pre-trained DNNs: VGG-16 [42], VGG-19 [42], MobileNet [22], Xception [10], Inception-v3 [43] and ResNet50 [20]. All these DNNs are considered major breakthroughs in DNN architectures as they improved the state-of-the-art performances during each year of ILSVRC [40] competitions. We also use test images from the

TABLE III
 DETAILS OF THE COMPUTER VISION SYSTEMS AND THE CORRESPONDING DATASETS USED FOR EVALUATING VERIVIS.

Task description	Test dataset	ID	Underlying architecture	Input size*	Top-5 loss	Top-1 loss	Pred. time (ms/img)
Detect ILSVRC 1000 class labels	ImageNet [15] provided in ILSVRC [40]	IMG_C1	VGG-16 [42]	224×224	0.1	0.295	102.9
		IMG_C2	VGG-19 [42]	224×224	0.09	0.273	103.9
		IMG_C3	MobileNet [22]	224×224	0.105	0.293	45.3
		IMG_C4	Xception [10]	299×299	0.055	0.21	54
		IMG_C5	Inception-v3 [43]	299×299	0.059	0.218	87.9
		IMG_C6	ResNet-50 [20]	224×224	0.071	0.242	50.8
Detect categories of general images	ImageNet [15] provided in ILSVRC [40]	API_C1	Google Vision [17]	224×224	-**	-**	904.87
		API_C2	Clarifai Tagging [12]	224×224	-**	-**	957.98
		API_C3	IBM Vision [24]	224×224	-**	-**	689.14
		API_C4	Microsoft Vision [29]	224×224	-**	-**	496.68
		API_C5	Amazon Rekognition [3]	224×224	-**	-**	795.91
Predict steering angle for each frame captured from car's front scene	Driving images provided by Udacity autonomous car challenge [47]	DRV_C1	Rambo [14]	192×256	0.058 ⁺	0.058 ⁺	33.6
		DRV_C2	Dave-orig [4], [7]	100×100	0.091 ⁺	0.091 ⁺	31.6
		DRV_C3	Dave-norminit [49]	100×100	0.053 ⁺	0.053 ⁺	31.1
		DRV_C4	Dave-dropout [13]	100×100	0.084 ⁺	0.084 ⁺	31.1

* We specify only image width and height for one channel. Color images have 3 channels with same height and width.

** The third-party blackbox APIs do not disclose their performance on any public test dataset.

⁺ We use MSE to measure the performance of self-driving car DNNs. Therefore, top-5 and top-1 loss have the same value for those DNNs.

TABLE IV
 AVERAGE NUMBER OF VIOLATIONS FOUND BY VERIVIS FOR EACH TEST INPUT IN DIFFERENT STATE-OF-THE-ART IMAGENET CLASSIFIERS AND SELF-DRIVING SYSTEMS. THE RESULTS ARE AVERAGES OVER 10 INPUT IMAGES FOR EACH SAFETY PROPERTY.

	ϕ_1	ϕ_2	ϕ_3	ϕ_4	ϕ_5	ϕ_6	ϕ_7	ϕ_8	ϕ_9	ϕ_{10}	ϕ_{11}	ϕ_{12}
IMG_C1	5	2.7	1.3	2	5359.1	39.4	4635.4	8782.8	710.5	20863	56.2	1.3
IMG_C2	2.5	2	2.5	1.4	3496.2	42.3	1657.2	25307.7	61592.9	2455.8	49.7	1.2
IMG_C3	3.7	1.8	2	2.8	5207.6	23.2	3960.5	11459.3	10.1	1239.7	61	1.4
IMG_C4	1.7	2	2.3	1.8	3218.7	21.6	5253.5	10603.8	63613.5	13684	197	1.5
IMG_C5	4	1.5	1.8	1.4	6724.5	39	1805	45642.8	724	368.5	40	1.6
IMG_C6	6	2.1	1.7	1.7	8596	43	1699	9757.8	6224	47312	19	2.1
DRV_C1	3.6	1.7	1.4	1.4	11935.9	89.7	65.4	7818.5	5561.2	4286.1	23.6	N/A*
DRV_C2	2.4	0.7	2	2.4	11294.2	35.6	2530.5	1207	51.4	303.5	108.8	N/A*
DRV_C3	1.6	1.8	2.3	3.6	18452.1	79	539.5	1722.2	258.1	783.3	42.6	N/A*
DRV_C4	1.1	2.5	1.6	1.9	5768.6	56.1	1866.8	372.8	504.6	7.6	123.3	N/A*

*Safety property involving reflected images is not realistic for self-driving cars.

test set provided by ILSVRC for verifying the local safety properties.

Third-party image classification services. We also evaluate VERIVIS on five blackbox commercial image classification APIs provided by Google [17], Clarifai [12], IBM [24], Microsoft [29], and Amazon [3]. We use the same test images from ILSVRC as discussed above.

Self-driving cars. For verifying computer vision systems performing regression tasks, we use four self-driving car DNNs that control the steering angle based on the input images captured by a front camera. In particular, we adopt Rambo [14], which is one of the top-ranked models in the Udacity self-driving car challenge [47], and three other open-source implementations [4], [13], [49] based on the Nvidia's Dave-2 self-driving system [7]. We use the testing set from the Udacity challenge [48] for verifying local safety properties.

VI. EVALUATION

A. Results

Summary. VERIVIS found thousands of violations of different tested safety properties in all of the tested vision systems.

Table IV and V summarize the number of violations that VERIVIS found for imagenet classifiers, self-driving cars, and third-party image classification services. Table X and XI in Appendix A) show some of sample inputs found by VERIVIS that violates safety properties of these vision systems.

Due to high network latency and cost (\$0.001 per query), for third-party image classification services, we only verify with relative properties with small number of critical parameters and report the result in Table V. Note that each number is averaged from the results from 10 random seed images. For t - and k -safety properties, as described in Section II, we set $t = 0.1$ and $k = 1$ for these experiments. Table VI shows the parameter spaces (\mathcal{C}_ϕ) that we use for generating these violations.

Verified images with no violations. We found that the number of verified input images, i.e., images for which a computer vision system do not violate a given safety property, varies widely based on the verification complexity of the safety property. The number of verified images decrease with higher verification complexity. However, even for properties with low verification complexity, we find that the number of verified

TABLE V

AVERAGE NUMBER OF VIOLATIONS FOUND BY VERIVIS FOR EACH TEST INPUT IN THE THIRD-PARTY IMAGE RECOGNITION APIS. THE RESULTS ARE AVERAGED OVER 10 IMAGES FOR EACH PROPERTY. DUE TO HIGH NETWORK LATENCY AND COST, WE ONLY TESTED FOR PROPERTIES WITH RELATIVELY LOW VERIFICATION COMPLEXITY.

	ϕ_1	ϕ_2	ϕ_3	ϕ_4	ϕ_6	ϕ_7	ϕ_{12}
API_C1	3.1	3.1	2.6	1.8	81.9	187.8	1.5
API_C2	6	2.8	1.5	1.5	25.4	105.4	1
API_C3	5.6	2.1	2.8	2.3	44	164.9	2.4
API_C4	0.5	1.8	2	2	100.5	75.6	2.2
API_C5	3.5	1.6	1.7	2.2	52.5	187	0.9

TABLE VI

THE MAXIMUM NUMBER OF CRITICAL PARAMETER VALUES FOR DIFFERENT INPUT SIZES FOR EACH SAFETY PROPERTY ϕ AND THE CORRESPONDING PARAMETER SPACE TESTED WITH VERIVIS.

Property	\mathcal{C}_ϕ	Max $ \mathcal{C}_{critical} $ for each input size			
		224	299	192×256	100
ϕ_1	[2, 10]	9	9	9	9
ϕ_2	[2, 10]	4*	4*	4*	4*
ϕ_3	[2, 5]	4	4	4	4
ϕ_4	[2, 5]	4	4	4	4
ϕ_5	[0.5, 2]	32512	32512	32512	32512
ϕ_6	[-100, 100]	200	200	200	200
ϕ_7	-†	33856	67081	40592	6400
ϕ_8	[-2, 2]	95496	225552	106722	8370
ϕ_9	[-0.01, 0.01]	250000	810000	238824	5140
ϕ_{10}	[0.99, 1.01]	244036	788544	230580	10000
ϕ_{11}	[-10, 10]	400	400	400	400
ϕ_{12}	-‡	3	3	3	3

* OpenCV only supports odd box sizes and thus smaller \mathcal{C}_ϕ than ϕ_1 .

† Position of occlusion mask is bounded by the image size.

‡ Reflection has only three critical parameter values.

TABLE VII


THE NUMBER AND PERCENTAGE OF VERIFIED IMAGES WITH RESPECT TO THE TOTAL NUMBER OF IMAGES IN ILSVRC TEST SET FOR SAFETY PROPERTIES ϕ_2 , ϕ_3 , ϕ_4 , AND ϕ_{12} FOR IMG_C3 (MOBILENET). THE LAST COLUMN SHOWS THE ORIGINAL ACCURACY OF IMG_C3 ON THIS TEST SET.

Property	# Verified inputs		Original accuracy
ϕ_2	28,509	28.5%	70.7%
ϕ_3	34,644	34.6%	70.7%
ϕ_4	30,979	31%	70.7%
ϕ_{12}	32,817	32.8%	70.7%

images is very low (on average 31.7%).

Table VII reports the number of verified images for properties ϕ_2 (median smooth), ϕ_3 (erosion), ϕ_4 (dilation), and ϕ_{12} (reflection) for IMG_C3 on all test images (100,000) from ILSVRC [40]. As shown in Table VII, the original top-1 test accuracy of IMG_C3 is around 70.7% which is significantly higher than the percentage of verified inputs.

Comparison of VERIVIS with gradient-based methods. Adversarial ML inputs, gradient-based approaches for finding violations of based on stochastic gradient descent is one of the most widely-used techniques in prior works to find violations to given ML systems [16], [33]. These gradient-based approaches do not provide any guarantee about absence of erroneous inputs. However, in order to empirically estimate how many erroneous cases they miss, we compare VERIVIS with gradient-based approaches in terms of number of violations found by both of these techniques. Specifically, we



MobileNet (IMG_C3)

Dave-orig (DRV_C2)

Fig. 5. Comparison of average numbers of violations found by gradient-based methods and VERIVIS for property ϕ_6 (changing brightness). VERIVIS finds 64.8× more violations than gradient based approach.

leverage the projected gradient descent approach described by Pei et al. [37] to change the brightness of an image (ϕ_6) and compare the number of violations found against those found by VERIVIS. We use the same values of the parameter space (\mathcal{C}_ϕ), k , and t as described in Section VI-A.

Figure 5 shows that VERIVIS finds up to 64.8 times more violations than the gradient-based approach. This demonstrates that gradient-based approaches often miss a large number of safety violations in computer vision systems.

Effects of k and t on the number of violations. We present how the thresholds k and t of local safety properties defined in Section II influence the number of violations found by VERIVIS. As shown in Figure 6, the number of violations decreases with increases in k and t . This is intuitive as increasing k and t essentially increase the allowed margin of error for the vision systems. One interesting fact is that although the number of violations drops significantly when k increases from 1 to 2, the changes in number of violations tend to be smaller when k increases further. By contrast, the decrease in the number of violations for different increasing values of t seems to be more uniform.

Violations for composition of transformations. We also explore the efficacy of VERIVIS to verify safety properties involving composition of multiple transformation (e.g., ϕ_1 and ϕ_6). For such cases, VERIVIS computes $\mathcal{C}_{critical}$ for the new composite transformation by calculating the Cartesian product of the critical parameter values of each of the individual transformations. Specifically, we run VERIVIS for different compositions of ϕ_1 (average smoothing), ϕ_6 (lighting), and ϕ_{11} (translation).

Figure 7 shows the results averaged for ten random input images with IMG_C3 as the verification target. We find that the number of violations for composite transformations is larger than simply multiplying the number of violations for individual transformations. This indicates that the composition of different transformations result in new violations than combinations of the existing ones for each individual transformation. Therefore, verifying safety properties with composite transformations, besides individual transformations, is critical for safety- and security-critical vision systems.

Distribution of violating parameter values. We also investigate how the violating parameter values are distributed in the parameter space. For example, we check whether the gain

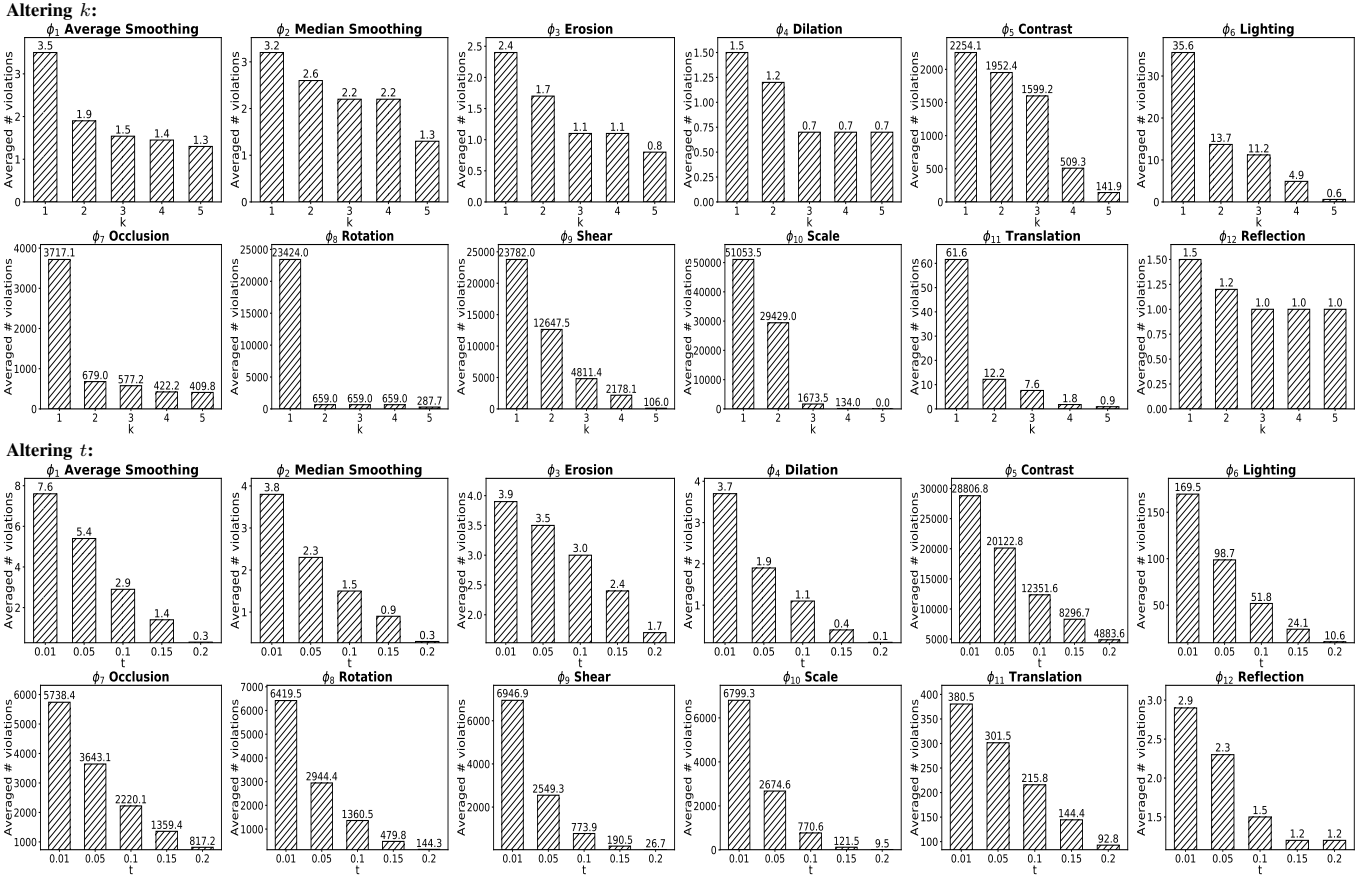


Fig. 6. The change in the average number of violations of different safety properties as we increase k (upper two rows) and t (lower two rows) as defined in Section II. The number are averages over 10 images. The number of violations tend to decrease with increasing k or t . The number above each bar shows the actual number of violations found for each k or t in IMG_C3 and DRV_C2 respectively.

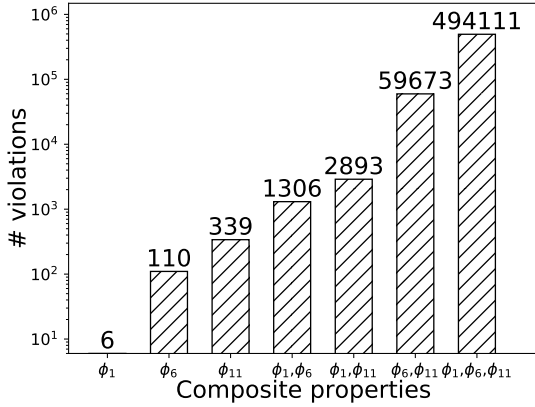


Fig. 7. Average number of violations found for properties ϕ_1 , ϕ_6 , and ϕ_{11} and their different compositions. We use IMG_C3 (MobileNet) as the target for verification.


values for brightness transformation causing safety property violations follow some obvious patterns. Our results indicate that while some violating parameter values for transformations like average smoothing follow simple patterns (e.g., parameter values that are higher than a threshold cause violations), most of the transformations do not display any obvious patterns.

Figure 8 presents the results for IMG_C3 (MobileNet) and two different transformations: average smoothing (ϕ_1)


and changing brightness (ϕ_6) with the same experimental setting as those described in Section VI-A. We pick these two transformations as their parameter space in one-dimensional and relatively small enough for clear demonstration in a two-dimensional graph. Figure 8 shows that, for ϕ_1 , when the kernel size exceeds a threshold (i.e., 6) for a specific input, the smoothing tends to always induce violations. However, by contrast, we do not find any obvious pattern among the brightness parameter values for ϕ_6 that result in violations.

Effect of increasing parameter space \mathcal{C}_ϕ on violations. We show the number of violations found by VERIVIS as the range of transformation space increases using Dave-orig (DRV_C2) as the target vision system. Figure 9 shows that the number of violations increases as we increase the range due to increase in the number of critical parameter values that need to be checked. For example, the average number of violations for rotation (ϕ_8) increases from 1360.5 to 3573.1 when we change the range of rotation degrees from $[-2, 2]$ to $[-3, 3]$.

Violations for complex transformations. Several real-world phenomena (e.g., fog, rain, etc.) that may affect input images are hard to replicate using the simple transformations described in the paper. However, one design custom transformations to mimic such effects and use VERIVIS to check safety properties with these transformations. As an example of



ϕ_1 : distribution of violations for average smoothing (left) and the original input (right)



ϕ_6 : distribution of violations for brightness (left) and the original input (right)

Fig. 8. Distribution of the violations with increasing transformation parameter values for average smoothing (ϕ_1) and lighting effect (ϕ_6). “Yes” and “No” indicate the image transformed by this particular parameter violates and satisfy the safety property, respectively.

this approach, we demonstrate how a simple parameterized fog-simulating transformation can be designed and verified with VERIVIS. For this transformation, we start with a fog mask, apply average smoothing on the mask, and apply the mask to the input image. By controlling the smoothing kernel size, we simulate different amounts of fog.

We use VERIVIS to enumerate and check all critical parameter values for the fog transform described above, i.e., different box sizes for average smoothing. VERIVIS were able to find hundreds of violations in MobileNet (IMG C3), Google vision API (API C1), and dave-orig (DRV C2). Figure 9 shows three sample violations found by VERIVIS.

B. Performance

In this section, we evaluate the performances of VERIVIS in terms of the time it takes to verify different safety properties. We report the numbers for VGG-19 (IMG_C2) and Rambo (DRV_C1) as they have the slowest inference time and therefore illustrate the worst case behavior of VERIVIS. All the test settings are the same as the ones described in Section VI-A and Table VI unless mentioned otherwise.

Summary. Table VIII shows the overall verification time required to verify each property for VGG-19 (IMG_C2) and Rambo (DRV_C1) after adopting batch prediction. The total verification time of enumerating all possible critical parameter values per image varies from 0.3 to 1863.5 seconds.

TABLE VIII

THE AVERAGE VERIFICATION TIME (IN SECONDS) FOR VERIFYING DIFFERENT PROPERTIES FOR IMG_C2 AND DRV_C1.

	ϕ_1	ϕ_2	ϕ_3	ϕ_4	ϕ_5	ϕ_6	ϕ_7	ϕ_8	ϕ_9	ϕ_{10}	ϕ_{11}	ϕ_{12}
IMG_C2	1.4	1.1	1.1	1.1	225.3	5	238.4	714.5	1863.5	1798.9	8.6	1.1
DRV_C1	0.3	0.8	0.8	0.8	75.1	1.6	81.2	242.1	596.6	596.1	3	0.7

We find that the verification time primarily depends on the number of critical parameter values for a transformation as the

image transformation operation is significantly cheaper than the testing time for each transformed image. For example, verification ϕ_1 (9 critical parameter values) takes only 0.3 seconds but ϕ_5 (32,512 critical parameter values) takes around 75.1 seconds for DRV_C1 as shown in Table VIII.

Performance improvement with batch prediction. As described in Section V, VERIVIS uses batch prediction [9], using both GPU and CPU, to speed up the verification process by allowing vision systems to predict a batch of images in parallel. Table III shows batch prediction can speed up verification by upto $17.6\times$ times for VGG-19 (IMG_C2), and $197.6\times$ times for Rambo (DRV_C1).

TABLE IX

THE AVERAGE RUNNING TIME (MILLISECOND PER IMAGE) OF VERIVIS WITH AND WITHOUT BATCH PREDICTION ON THE VGG-19 (IMG_C2) AND RAMBO (DRV_C1). THE SPEEDUP IS SHOWN IN THE LAST COLUMN.

	Baseline	Batch pred.	Speed-up
IMG_C2	103.9	5.9	$17.6\times$
DRV_C1	33.6	0.17	$197.6\times$

C. Improving robustness with retraining

In this subsection, we investigate whether the robustness of the tested computer vision systems against these transformations can be improved by retraining the affected systems using the violations found by VERIVIS. In particular, we borrow the idea of adversarial retraining as a data augmentation technique introduced by Goodfellow et al. [16] to retrain the ML models on transformed images that induce safety violations. We pick VGG-16 for this experiment due to its large number of safety violations (the largest Top-1 loss as shown in Table III).

We compare the violations found by VERIVIS for the retrained models and the original ones using ten randomly selected images from the ILSVRC test set. We use two different training data augmentation strategies: (1) retraining on the safety-violating images generated from *the test images* and (2) retraining on the safety-violating images generated from *different images* than the test images. Note that the the second approach is less likely to overfit than the first approach. Figure 11 shows the number of violations found for randomly-drawn 10 testing images before and after retraining the VGG-16 with our two tactics. The results show that both approaches of retraining described above reduce the number of violations and therefore improve the robustness of the ML model. Overall, such retraining reduced the number of violations up to 60.2% on average for different images.

VII. RELATED WORK

Testing and verification of machine learning. Given the deployment of ML systems in security- and safety-critical settings, some recent studies have focused on generating diverse and realistic corner-case test inputs for testing ML systems and have found numerous incorrect behavior in state-of-the-art systems [37], [46]. However, none of these systems, unlike VERIVIS, can provide any guarantee about the non-existence of safety violations.

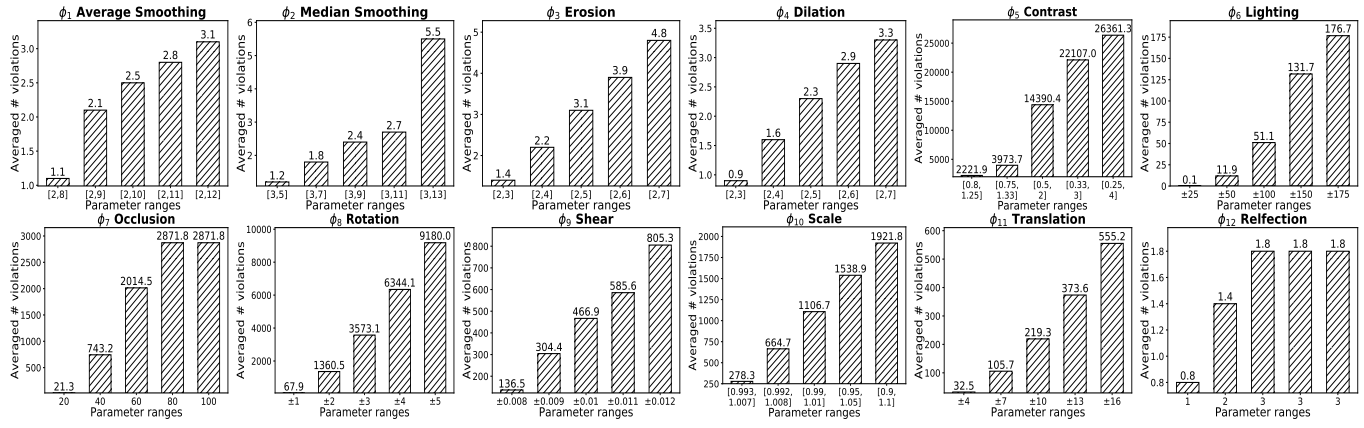


Fig. 9. The average numbers of violations found increase as we increase the range of the parameter space (\mathcal{C}_ϕ). The number above each bar shows the exact number of violations with the corresponding bounds.




Fig. 10. Violations found by VERIVIS for a fog-simulating transformation in MobileNet (IMG_C3), Google vision API (API_C1), and dave-orig (DRV_C2). The first and second rows show the original images and the foggy images that result in violations, respectively.




Fig. 11. The number of average violations found before and after retraining of IMG_C1 on violations of safety properties ϕ_1 , ϕ_6 , and ϕ_{11} . The left bar shows the baseline without retraining. The middle and right bars show after retraining on (1) violations generated from same images as those ten test images, and (2) violations generated from different images, respectively. Even in the latter case, the number of violations can drop by up to 60%.

Several researchers have also explored whitebox formal verification techniques for DNNs to ensure correct behavior under different settings [23], [25], [38]. Unlike VERIVIS, these techniques either fail to provide strong guarantees or does not scale to real-world-sized ML systems (e.g., neural network with thousands of neurons) [23], [25], [38].

Adversarial machine learning. Adversarial machine learning is a popular research area that focuses on generating error-

inducing test inputs for ML models by adding minimal perturbations to an existing input [16], [26], [27], [30], [35], [41], [44], [51], [53] and studying how to improve the robustness of the ML systems against such attacks [5], [8], [11], [18], [23], [28], [36], [50], [55]. We refer the interested readers to the survey by Papernot et al. [34] for more details on these works.

The key difference between VERIVIS and this line of work is twofold: (1) Adversarial inputs only focus on *checking one type of safety property* by adding adversarial perturbations. By contrast, VERIVIS is a general verification framework that supports a broad set of safety properties with more realistic transformations that even weak attackers might be able to induce; and (2) Unlike any of the adversarial machine learning projects, VERIVIS can ensure non-existence of inputs that violate a given safety property for state-of-the-art ML systems.

VIII. CONCLUSION

In this paper, we presented a general framework for verifying the robustness of ML systems with different real-world safety properties that can model different attacker capabilities. We designed, implemented, and extensively evaluated VERIVIS, a scalable verification system that can verify a diverse set of safety properties for state-of-the-art computer vision systems with only blackbox access. VERIVIS leveraged a generic decomposition framework for image transformations to significantly reduce the search space of different safety properties for efficient verification.

REFERENCES




- [1] Changing the contrast and brightness of an image. http://docs.opencv.org/2.4/doc/tutorials/core/basic_linear_transform/basic_linear_transform.html.
- [2] Tesla self-driving system faulted by safety agency in crash. <https://www.nytimes.com/2017/09/12/business/self-driving-cars.html>.
- [3] Amazon Rekognition, deep learning-based image recognition search, verify, and organize millions of images. <https://aws.amazon.com/rekognition/>.
- [4] Nvidia-Autopilot-Keras. <https://github.com/Observer07/Nvidia-Autopilot-Keras>, 2016.

- [5] O. Bastani, Y. Ioannou, L. Lampropoulos, D. Vytiniotis, A. Nori, and A. Criminisi. Measuring neural net robustness with constraints. In *Advances in Neural Information Processing Systems*, 2016.
- [6] P. N. Belhumeur and D. J. Kriegman. What is the set of images of an object under all possible illumination conditions? *International Journal of Computer Vision*, 1998.
- [7] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [8] N. Carlini and D. Wagner. Towards evaluating the robustness of neural networks. In *Proceedings of the 2017 IEEE Symposium on Security and Privacy*, 2017.
- [9] F. Chollet. Keras. <https://github.com/fchollet/keras>, 2015.
- [10] F. Chollet. Xception: Deep learning with depthwise separable convolutions. *arXiv preprint arXiv:1610.02357*, 2016.
- [11] M. Cisse, P. Bojanowski, E. Grave, Y. Dauphin, and N. Usunier. Parseval networks: Improving robustness to adversarial examples. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.
- [12] Clarifai API: Large Scale Visual Recognition. <https://developer.clarifai.com/models/general-image-recognition-model/aaa03c23b3724a16a56b629203edc62c>, 2013.
- [13] Behavioral cloning: End-to-end learning for self-driving cars. <https://github.com/navoshta/behavioral-cloning>, 2016.
- [14] Rambo model for Udacity self-driving car challenge 2. <https://github.com/udacity/self-driving-car/tree/master/steering-models/community-models/rambo>, 2017.
- [15] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the 22nd IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [16] I. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. In *Proceedings of the 3rd ICLR*, 2015.
- [17] Cloud Vision API - Derive insight from images with our powerful Cloud Vision API. <https://cloud.google.com/vision/>, 2011.
- [18] S. Gu and L. Rigazio. Towards deep neural network architectures robust to adversarial examples. In *Proceedings of the 3rd ICLR*, 2015.
- [19] V. Gulshan, L. Peng, M. Coram, M. C. Stumpe, D. Wu, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *Jama*, 2016.
- [20] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [21] G. J. Holzmann. The model checker SPIN. *IEEE Transactions on software engineering*, 1997.
- [22] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. 2017.
- [23] X. Huang, M. Kwiatkowska, S. Wang, and M. Wu. Safety verification of deep neural networks. In *Proceedings of the 29th International Conference on Computer Aided Verification*, 2017.
- [24] IBM Watson Visual Recognition Service. <https://www.ibm.com/watson/developercloud/doc/visual-recognition/index.html>.
- [25] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer. Reluplex: An efficient SMT solver for verifying deep neural networks. In *Proceedings of the 29th International Conference On Computer Aided Verification*, 2017.
- [26] P. Laskov et al. Practical evasion of a learning-based classifier: A case study. In *Proceedings of the 2014 IEEE S&P*, 2014.
- [27] Y. Liu, X. Chen, C. Liu, and D. Song. Delving into transferable adversarial examples and black-box attacks. In *Proceedings of the 6th ICLR*, 2016.
- [28] J. H. Metzen, T. Genewein, V. Fischer, and B. Bischoff. On detecting adversarial perturbations. In *Proceedings of the 6th ICLR*, 2017.
- [29] Microsoft Computer Vision API. <https://azure.microsoft.com/en-us/services/cognitive-services/computer-vision/>.
- [30] A. Nguyen, J. Yosinski, and J. Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the 28th IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [31] D. of Homeland Security. Passenger screening algorithm challenges. <https://www.kaggle.com/c/passenger-screening-algorithm-challenge>, 2017.
- [32] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, 2017.
- [33] N. Papernot, P. McDaniel, S. Jha, M. Fredrikson, Z. B. Celik, and A. Swami. The limitations of deep learning in adversarial settings. In *2016 IEEE European Symposium on S&P*, 2016.
- [34] N. Papernot, P. McDaniel, A. Sinha, and M. Wellman. Towards the science of security and privacy in machine learning. In *Proceedings of the 3rd IEEE European Symposium on S&P*, 2018.
- [35] N. Papernot, P. McDaniel, A. Swami, and R. Harang. Crafting adversarial input sequences for recurrent neural networks. In *IEEE Military Communications Conference*, 2016.
- [36] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *Proceedings of the 37th IEEE S&P*, 2016.
- [37] K. Pei, Y. Cao, J. Yang, and S. Jana. DeepXplore: Automated whitebox testing of deep learning systems. In *Proceedings of the 26th ACM Symposium on Operating Systems Principles*, 2017.
- [38] L. Pulina and A. Tacchella. An abstraction-refinement approach to verification of artificial neural networks. In *Proceedings of the 22nd International Conference on Computer Aided Verification*, 2010.
- [39] J. Ragan-Kelley, C. Barnes, A. Adams, S. Paris, F. Durand, et al. Halide: A language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines. In *PLDI*, 2013.
- [40] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, pages 211–252, 2015.
- [41] M. Sharif, S. Bhagavatula, L. Bauer, and M. K. Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016.
- [42] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *Proceedings of the 4th ICLR*, 2015.
- [43] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [44] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *Proceedings of the 2nd ICLR*, 2014.
- [45] R. Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.
- [46] Y. Tian, K. Pei, S. Jana, and B. Ray. DeepTest: Automated testing of deep-neural-network-driven autonomous cars. *arXiv preprint arXiv:1708.08559*, 2017.
- [47] Using deep learning to predict steering angles. <https://github.com/udacity/self-driving-car>, 2016.
- [48] Udacity self-driving car steering-angle prediction challenge dataset. <https://github.com/udacity/self-driving-car/tree/master/datasets/CH2>, 2016.
- [49] Visualizations for understanding the regressed wheel steering angle for self driving cars. <https://github.com/jacobgil/keras-steering-angle-visualizations>, 2016.
- [50] G. Wang, T. Wang, H. Zheng, and B. Y. Zhao. Man vs. Machine: Practical adversarial detection of malicious crowdsourcing workers. In *USENIX Security Symposium*, 2014.
- [51] M. J. Wilber, V. Shmatikov, and S. Belongie. Can we still avoid automatic face detection? In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2016.
- [52] I. H. Witten, E. Frank, M. A. Hall, and C. J. Pal. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, 2016.
- [53] W. Xu, Y. Qi, and D. Evans. Automatically evading classifiers. In *Proceedings of the 2016 Network and Distributed Systems Symposium*, 2016.
- [54] Y. Zhang, C. Mu, H.-W. Kuo, and J. Wright. Toward guaranteed illumination models for non-convex objects. In *Proceedings of the IEEE International Conference on Computer Vision*, 2013.
- [55] S. Zheng, Y. Song, T. Leung, and I. Goodfellow. Improving the robustness of deep neural networks via stability training. In *Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition*, pages 4480–4488, 2016.

TABLE X
 SAMPLE VIOLATIONS OF SAFETY PROPERTIES FOUND BY VERIVIS IN DIFFERENT VISION SYSTEMS TRAINED ON IMAGENET.

Model ID	Original		Transformed		Violated
	Image	Labels (Top-5)	Image	Label (Top-1)	
IMG_C1		pomegranate bell_pepper acorn_squash pizza trifle		starfish	ϕ_6
IMG_C2		park_bench sundial cannon bannister plow		peacock	ϕ_7
IMG_C3		stage cornet trombone cello violin bathtub		jellyfish	ϕ_8
IMG_C4		sea_snake scuba_diver sea_lion coral_reef brain_coral		wreck	ϕ_9
IMG_C5		cassette_player tape_player radio CD_player loudspeaker		Polaroid_camera	ϕ_{11}
IMG_C6		Eskimo_dog Siberian_husky malamute Norwegian_elkhoun Pembroke		Cardigan	ϕ_5
API_C1		spider arachnid invertebrate tangle_web_spider arthropod		water	ϕ_3
API_C2		scissors equipment tool steel work		business	ϕ_1
API_C3		carthorse odd-toed_ungulate mammal animal buggy		elk	ϕ_{12}
API_C4		tree outdoor animal mammal ape		fungus	ϕ_2
API_C5		Belt Carousel Indoors Lobby Reception		architecture	ϕ_4

TABLE XI
 SAMPLE VIOLATIONS OF SAFETY PROPERTIES FOUND BY VERIVIS IN DIFFERENT SELF-DRIVING SYSTEMS.

Model ID	<i>Image</i>	Original <i>Steering angle prediction</i>	<i>Image</i>	Transformed <i>Steering angle prediction</i>	Violated
DRV_C1		near straight (left 0.1°)		turn right (right 25°)	ϕ_6
DRV_C2		near straight (right 4°)		turn right (right 44°)	ϕ_7
DRV_C3		near straight (left 3.7°)		turn right (right 27°)	ϕ_{11}
DRV_C4		near straight (left 0.03°)		turn right (right 38°)	ϕ_2