

RESEARCH

Open Access



IPatch: a remote adversarial patch

Yisroel Mirsky*

Abstract

Applications such as autonomous vehicles and medical screening use deep learning models to localize and identify hundreds of objects in a single frame. In the past, it has been shown how an attacker can fool these models by placing an adversarial patch within a scene. However, these patches must be placed in the target location and do not explicitly alter the semantics elsewhere in the image. In this paper, we introduce a new type of adversarial patch which alters a model's perception of an image's semantics. These patches can be placed anywhere within an image to change the classification or semantics of locations far from the patch. We call this new class of adversarial examples 'remote adversarial patches' (RAP). We implement our own RAP called IPatch and perform an in-depth analysis on without pixel clipping on image segmentation RAP attacks using five state-of-the-art architectures with eight different encoders on the CamVid street view dataset. Moreover, we demonstrate that the attack can be extended to object recognition models with preliminary results on the popular YOLOv3 model. We found that the patch can change the classification of a remote target region with a success rate of up to 93% on average.

Introduction

Deep learning has become the go-to method for automating image-based tasks. This is because, deep neural networks (DNNs) are excellent at learning and identifying spatial patterns and abstract concepts. With advances in both hardware and neural architectures, deep learning has become both a practical and reliable solution. Companies now use image-based deep learning to automate tasks in life critical operations such as autonomous driving Leetaru (2019); Autopilot ai-tesla (2021), surveillance Vincent (2018), and medical image screening Siemens (2021).

In tasks such as these, multiple objects must be identified per image. One way to accomplish this is to predict a class probability for each pixel in the input image x . This approach is called image segmentation and companies such as Telsa use it to guide their autonomous vehicles safely through an environment Autopilot ai-tesla (2021).

Another approach is called object detection where x is split into a grid of cells or regions and the model predicts both a class probability and a bounding box for each of them Redmon and Farhadi (2018); Ren et al. (2015). In both cases, these models rely on image semantics to successfully parse and interpret a scene.

Just like other deep learning models, these semantic models are also susceptible to adversarial attacks. In 2017, researchers demonstrated how a small 'adversarial' patch can be placed in a real world scene and override an image-classifier's prediction, regardless of the patch's location or orientation Brown et al. (2017). This gave rise to a number of works which demonstrated the concept of adversarial patches against image segmentation and object detection models Song et al. (2018); Liu et al. (2018); Chen et al. (2018); Sitawarin et al. (2018); Lee and Kolter (2019); Thys et al. (2019); Zhao et al. (2019); Li et al. (2020); den Hollander et al. (2020); Huang et al. (2020); Hoory et al. (2020); Wu et al. (2020). However, current adversarial patches are limited in the following ways:

Location

Only predictions around the patch itself are explicitly affected. This limits where objects can be made to 'appear'

*Correspondence:

Yisroel Mirsky

yisroel@bgu.ac.il

Ben-Gurion University, Department of Software and Information Systems Engineering, Beersheba, Israel

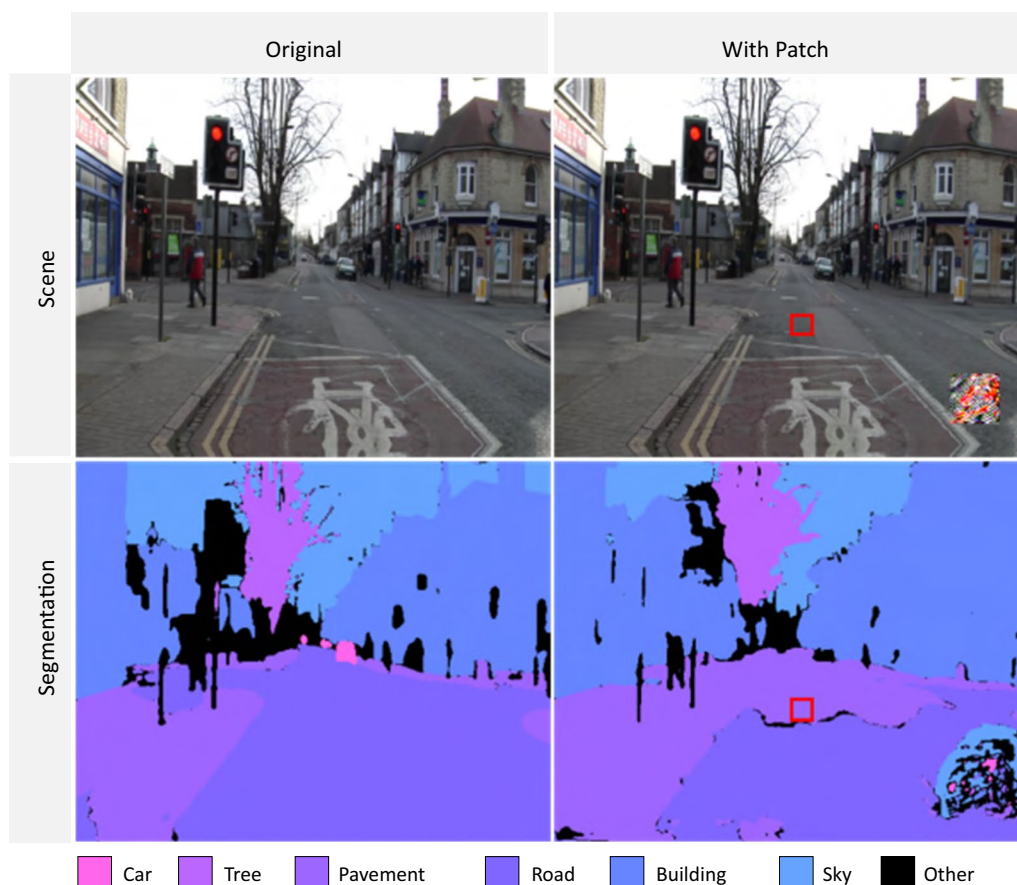


Fig. 1 An example of a remote adversarial attack on a segmentation model. Here the patch has been designed to alter a predetermined location (the red box) to the class of 'pavement'. The same patch works on different images and with different positions and scales

in a scene. For example, a patch cannot make a plane appear in the sky and it is difficult to put a patch in the middle of a busy road. Furthermore, patches in noticeable areas can raise suspicion (e.g., a stop sign with a colorful patch on it).

Interpretation Existing patches do not explicitly alter the shape or layout of a scene's perceived semantics. Changes to these semantics can be used to guide behaviors (e.g., drive a car off the road Teichmann et al. (2018) or change a head count He et al. (2020)) and has wide implications on tasks such as surveillance Iglovikov et al. (2017); Vincent (2018) and medical screening Prajna and Nath (2021) among others.

In this paper we identify a new type of attack which we call a Remote Adversarial Patch (RAP). A RAP is an adversarial patch which can alter an image's perceived semantics from a remote location in the image. Our implementation of a RAP (IPatch) can be placed anywhere in the field of view and alter the predictions of nearly any predetermined location within the same view. This is demonstrated in Fig. 1 where an attacker has crafted an IPatch which causes a segmentation model to think that there is pavement (a sidewalk) in the middle of the road. Moreover, this adversarial attack is robust because the same patch works on different images using different positions and scales. Therefore, this attack more flexible and more covert than previous approaches. later in "Threat model" section we discuss the attack model further.

Since the IPatch can alter an image's perceived semantics, and attacker can craft patches which cause these models to see objects of arbitrary shapes and classes. For

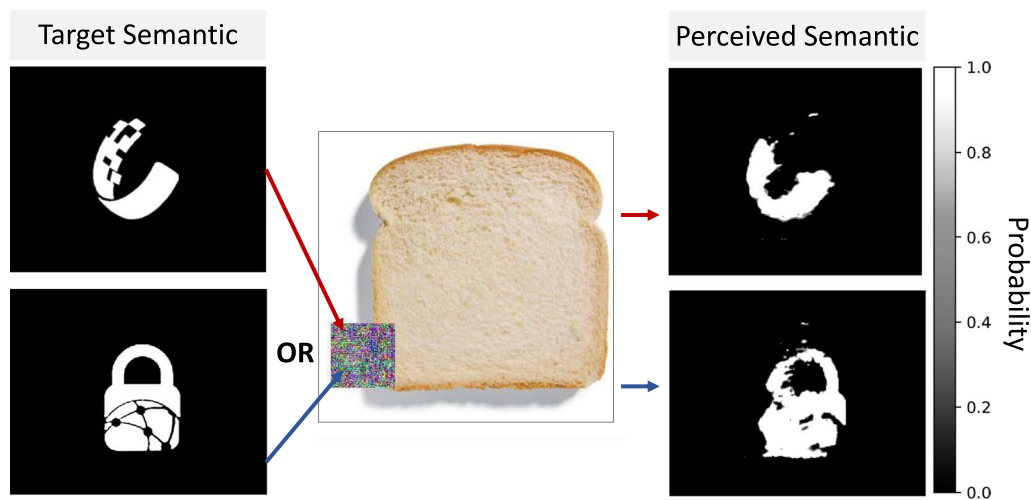


Fig. 2 An example of how an IPatch can change the semantics of an image to an arbitrary shape. Here, a street view segmentation model is convinced that the slice of bread is a tree in the shape of the USENIX logo (top) and NDSS logo (bottom)

example, in Fig. 2 a street view segmentation model is convinced that a slice of bread is a tree shaped like the USENIX security symposium logo (top) or the NDSS logo (bottom). This is possible because semantic models rely on global and contextual features to parse an image. However, an object and its contextual information can be very far apart in x . For example, consider an image with a boat next to the water. Here, the water will boost the confidence of the boat's classification even though the boat is not in the water. The IPatch exploits these correlations by masquerading as these contextual features.

Creating a robust RAP is more challenging than existing adversarial patches. This is because the content of x directly affects the leverage of the patch. For example, an IPatch cannot make a segmentation model perceive remote semantics on a blank image. However, to create a robust patch, we must be able to generalize to different images which have not been seen before. To overcome these challenges, we (1) use an incremental training strategy to slowly increase the entropy of the expectation over transformation (EoT) objective and (2) use Kullback–Leibler divergence loss to help the optimizer leverage and exploit the contextual relationships.

In this work, we study how RAPs work *without clipping the pixel ranges* to understand how locality of a patch can affect remote regions. In particular, we study IPatches as a RAP against semantic segmentation models. We also demonstrate that the same technique can be applied to object detectors, such as YOLO, as well. To evaluate the IPatch, we train 37 segmentation models using 8 different encoders and 5 state-of-the-art architectures. In our evaluations, we focus on the autonomous car scenario Autopilot ai-tesla (2021); Siam et al. (2018), and perform

rigorous tests to determine the limitations and capabilities of the attack. On the top 4 classes, we found that the attack works up to 93% of the time on average, depending on the victim's model. We also found that all of the segmentation models are susceptible to the attack, where the most susceptible architectures were the FPN and Unet++ and the least susceptible architecture was the PSPNet. Finally, even if the attacker does not have the same architecture as the victim, we found that without any additional training effort, an IPatch trained on one architecture works on others with an attack success rate of up to 25.3%.

The contributions of this paper are as follows:

- We introduce a new class of adversarial patches (RAP) which can manipulate a scene's interpretation remotely and explicitly. This type of attack not only has significant implications on the security of autonomous vehicles, but also on a wide range of semantic-based applications such as medical scan analysis, surveillance, and robotics (section "Threat model").
- We present a training framework which enables the creation of a robust RAP (IPatch) by incrementally increasing the training entropy. Without this strategy, the entropy starts too high which makes it difficult to converge on some learning objectives, especially given large patch transformations on scale, shift, and so on (section "Making IPatch").
- We provide an in-depth evaluation of the patch used as a remote adversarial attack against road segmentation models (section "Evaluation"). We show that the attack is robust, universal (works on unseen images sampled from the same distribution), and has trans-

ferability (works across multiple models). We also provide initial results which demonstrate that the attack works on object detectors as well (specifically YOLOv3).

- We identify the attack's limitations and provide insight as to why this attack can alter the perception of remote regions in an image. Building on these observations, we suggest countermeasures and directions for future work (section "[Discussion& counter-measures](#)").
- To the best of our knowledge, this the first adversarial patch demonstrated on segmentation models (section "[Related works](#)").

Related works

Soon after the popularization of deep learning, researchers demonstrated that DNNs can be exploited using adversarial examples Biggio and Roli (2018). In 2014 it was shown how an attacker can alter an image-classification model's predictions by adding an imperceivable amount of noise to the input image Szegedy et al. (2013); Goodfellow et al. (2014); Nguyen et al. (2015). Initially, these attacks were impractical to perform in a real environment since every combination of lighting, camera noise, and perspective would require a different adversarial perturbation Luo et al. (2015); Lu et al. (2017). However, in 2017 the authors of Athalye et al. (2018) showed that an adversary can consider these distortions while generating the adversarial example in a process called Expectation over Transformation (EoT). Using this method, the authors were able to generate robust adversarial samples which can be deployed in the real world. In the same year, the authors of Brown et al. (2017) used EoT to create adversarial patches. Their adversarial patches were designed to fool image-classifiers (single-object detection models).

Later in 2018, the authors of Song et al. (2018) developed an adversarial patch that works on object detection models (multi-object detection models). More recently, researchers have proposed patches which can remove objects which wear the patch Liu et al. (2018); Zhao et al. (2019); Thys et al. (2019); Wu et al. (2020); Huang et al. (2020); den Hollander et al. (2020); Li et al. (2020) and patches which can perform denial of service (DoS) attacks by corrupting a scene's interpretation Liu et al. (2018); Lee and Kolter (2019).

In Table 1, we summarize the related works on adversarial examples against image segmentation and object detection models (the domain of the proposed attack). In general, the attack goals of these papers are either add/change an object in the scene or to remove all objects altogether (DoS). The methods which add adversarial

perturbations (noise) can change the semantics of an image at any location Hendrik Metzen et al. (2017); Fischer et al. (2017); Arnab et al. (2018); Ozbulak et al. (2019); Kang et al. (2020), but they cannot be deployed in the real world since they are applied directly to an image itself. Currently, there no patches for image segmentation models, and the patches for object detection models only affect the prediction around the patch itself. The exception are patches which perform DoS attacks by removing/corrupting all objects detected in the scene like Liu et al. (2018); Lee and Kolter (2019).

Therefore, to the best of our knowledge, the attack which we introduce is the first RAP, and (1) the only method which can add, change, or remove objects in a scene *remotely* (far from the location of the patch itself), (2) the first adversarial patch proposed for segmentation networks, and (3) the first adversarial patch which can cause a model to perceive custom semantic shapes.

Threat model

The Vulnerability. The vulnerability which this paper introduces is that semantic models, such as image segmentation models, utilize global and contextual features in an image to improve their predictive capabilities. However, these dependencies expose channels which can an attacker can exploit to change the interpretation of an image from one remote location to another.

The Attack Scenario. In this work we will focus on the remote adversarial attack scenario. In this attack scenario, the victim has an application which uses the image segmentation model M . The attacker wants M to predict a specific class at a specific location L , while looking at a certain scene. To accomplish this, the attacker needs a training set of 1 or more images and a segmentation model to work with.

For the training set, the attacker has two options: (1) obtain images similar to those used to train M , or (2) take pictures of the target scene. For the model, the attacker can either follow a white-box or black-box approach: In a white box-approach, the attacker obtains a copy of M to achieve the most accurate results. The white-box approach is a common assumption for adversarial patches. Alternatively, the attacker can follow a black-box approach and train a surrogate model M' on a similar dataset used to train M . Although the black box approach performs worse, we have found that there is some transferability between a patch trained on one model and then used against another (section "[Evaluation](#)"). Finally, the attacker generates an IPatch P which targets L using X and M .

Motivation. There are several reasons why an attacker would want to use an IPatch over an ordinary adversarial patch (illustrated in Fig. 3:

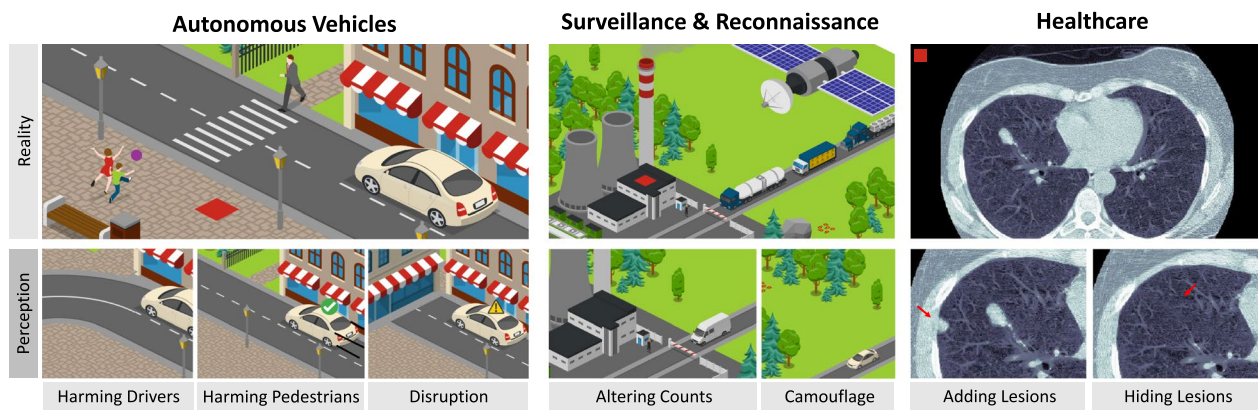


Fig. 3 Some illustrative examples showing how a remote adversarial patch can be used by an attacker

Table 1 Related works on adversarial examples which target image segmentation and object detection

		Type		Attack goal		Deployment		Flexibility	
		Patch	Noise	Add/Change	Hide	Applied to image	Physically deployed	Must be near/on target	Can be placed anywhere
Object Detection	Song et al. (2018)	•		•			•	•	
	Chen et al. (2018)	•		•			•	•	
	Liu et al. (2018)	•		•	•		•	•	•*
	Sitawarin et al. (2018)	•			•		•	•	
	Wei et al. (2018)		•	•		•		•	
	Zhao et al. (2019)	•			•		•	•	
	Thys et al. (2019)	•			•		•	•	
	Lee and Kolter (2019)	•			•		•		•*
	Wu et al. (2020)	•			•		•	•	
	Hoory et al. (2020)	•		•			•	•	
	Zhao et al. (2020)		•		•	•		•	
	Huang et al. (2020)	•			•		•	•	
	Chow et al. (2020)		•	•	•	•		•	
	den Hollander et al. (2020)	•			•		•	•	
	Zhang et al. (2020)		•	•		•		•	
	Zolfi et al. (2020)		•		•		•	•	
	Li et al. (2020)		•	•		•		•	
	Li et al. (2020)	•			•		•	•	
	den Hollander et al. (2020)		•	•		•		•	
Img. Segment.	Fischer et al. (2017)		•		•	•		•	
	Hendrik Metzen et al. (2017)		•		•	•		•	
	Arnab et al. (2018)		•	•		•		•	
	Kang et al. (2020)		•	•		•		•	
	Ozbulak et al. (2019)		•	•		•		•	
Both	Xie et al. (2017)		•	•	•	•		•	
	IPatch	•		•	•		•		•

*Patch can be anywhere only when used to hide all objects in the scene (DoS)

Stealth

The attacker may want to place the patch in a less obvious place so it won't be removed or noticed by the victim. For example, a sticker on a stop sign is anomalous and can be contextually identified as malicious Sitawarin et al. (2018) but a sticker on a nearby billboard is less obvious. Another example, is in the domain of medicine where segmentation models are used to highlight and identify different lesions such as tumors. Here, an attacker can't put a patch in the image in the location of the lesion since it would be an obvious attack. However, the attack could be trigger remotely by placing a dark RAP in the dark space of a scan where it is common to have noise, or in a location of the scan which is not under investigation (e.g., the first few slices on the z-axis). For motivations why an attacker would want to target medical scans, see Mirsky et al. (2019).

Practicality

The attacker may want to generate an object or semantic illusion in a location which is hard to reach or impractical to place a patch on it. For example, in the sky region, on the back of an arbitrary car on the freeway.

Flexibility

The attacker may need to craft or alter specific semantics for a scene. For example, many works show how image segmentation can be used to identify homes, roads and resources from satellite and drone footage Audebert et al. (2017). Here an attacker can feed false intel by hiding or increasing the number of structures, people, and resources before it can be investigated manually.

Overall, the IPatch attack is more flexible and enables more attack vectors than location-based patches (e.g., Song et al. (2018); den Hollander et al. (2020)). However, it is significantly more challenging to generate an IPatch. Therefore, its flexibility comes with a trade-off in terms of attack performance.

Making an IPatch

In this section we first provide an overview of how image segmentation models work. Then we present our approach on how to create an IPatch.

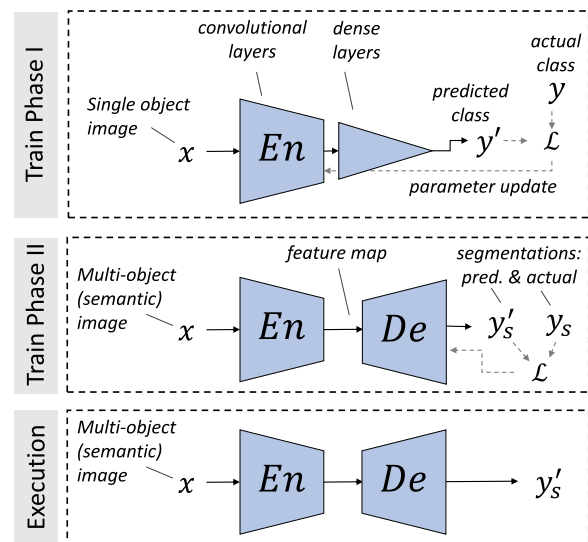


Fig. 4 A basic schematic showing how a common Image Segmentation model is trained and executed. *En* is an encoder architecture which compressed an input into a smaller representation. *De* is a decoder model which expands an encoded representation into meaningful content

Technical background

There are a wide variety of deep learning models for image segmentation Garcia-Garcia et al. (2018); Minaee et al. (2020). The most common form involves an encoder *En* and decoder *De* such that

$$S(x) = De(En(x)) \quad (1)$$

illustrated in Fig. 4. The objective of a segmentation model is to take an N -by- M image (s) with 1-3 color channels and predict an N -by- M -by- C probability mapping (y_s). The output y_s can be mapped directly to the pixels of x such that $y_s[i, j, k]$ is the probability that pixel $x[i, j]$ belongs to the k -th class (among the C possible classes).

To train S , the common approach is to follow two phases: In the first phase, the encoder network is trained as an image-classifier on a large image dataset in a supervised manner (i.e., where each image x is associated with a label y). Note that the classifier's task is to predict a single class for the entire image (e.g., x is a dog). After training the classifier, we discard the dense layers at the end of the network (used to predict y) and retain the convolutional layers at the front of encoder *En*. In this way, we can use the feature mapping learned by the classifier to perform image segmentation. In the second phase, the decoder architecture *De* is added on and S is trained end-to-end. Often, the weights of *En* are locked during this phase, and we do the same in this paper.

One reason why the encoder-decoder approach is so popular, is because obtaining a labeled segmentation ground truth y_s is significantly more challenging than for image classification y (massive datasets for classification exist and new datasets can be crowd sourced as well). Therefore, by using a pre-trained encoder, far fewer examples of segmentations are needed to achieve quality results.

To train S , a differentiable loss function \mathcal{L} is used to compare the model's predicted output y'_s to the ground truth y_s in order to perform backpropagation and update network's weights. There are many loss functions used in for segmentation. One common approach is to simply apply the binary cross entropy loss (\mathcal{L}_{CE}) since S is essentially trying to solve a multi-class classification problem. However, \mathcal{L}_{CE} does not consider whether a pixel is on the boundary or not so results tends to be blurry and be biased to large segments such as backgrounds Deng et al. (2018). To counter this issue, in 2016 the authors of Milletari et al. (2016) proposed using Dice loss (\mathcal{L}_D) for medical image segmentation, and it has since been considered a state-of-the-art approach. The Dice loss is defined as

$$\mathcal{L}_D(x, y) = \frac{2 \sum_i^N x_i y_i}{\sum_i^N x_i^2 + \sum_i^N y_i^2} \quad (2)$$

We use \mathcal{L}_D to train all of the image segmentation models in this paper.

When selecting the encoder's model, there are a wide variety of options. Some include ResNext, DenseNet, xception, EfficientNet, MobileNet, DPN, VGG, and variations thereof. However, regarding the decoder's architecture, there are several which are considered state-of-the-art. Many of them utilize a 'feature pyramid' approach and skip connections to identify features at multiple scales, or an autoencoder (encoder decoder pair) to encode and extract the semantics. We will now briefly describe the five architectures used in this paper:

Unet++ Zhou et al. (2018): An autoencoder architecture which improves on its predecessor, the Unet. The encoder and decoder are connected through a series of nested dense skip connections which reduce the semantic gap between the feature maps of the two networks.

Linknet Chaurasia et al. (2017): An efficient autoencoder which passes spatial information across the network to avoid losing it in the encoder's compression.

FPN Lin et al. (2017): A feature pyramid network which uses lateral connections across a fully convolution neural network (FCN) to utilize feature maps learned from multiple image scales.

PSPNet Zhao et al. (2017): An FCN which uses a pyramid parsing module on different sub-region representations in order to better capture global category clues.

The architecture won first place in multiple segmentation challenge contests.

PAN Li et al. (2018): A network which uses both pyramid and global attention mechanisms to capture spatial and global semantic information.

Approach

In a remote adversarial attack, the attacker wants a region around the location $L = (i, j)$ to be predicted as class k . To ensure that the optimizer does not waste energy on other semantics in the scene, we focus the effort to a region of operation. Let m denote the region operation and let t be the target pattern for that region. To capture m , we use an N -by- M -by- C mask of zeros. To select L , a square or circle with a radius of r pixels¹ around L in m is marked with ones along the k -th channel. To insert an object, we set $t = m$ since our objective is to change the probability of those pixels to one. To insert a custom shape (like in Fig. 2) t is set accordingly.

To generate a patch for the objective (L, k) , we follow the EoT approach similar to previous works Brown et al. (2017); Song et al. (2018); Liu et al. (2018), but using our semantic masks. Concretely, we would like to find a patch P which is trained to optimize the following objective function

$$P = \arg \min_{\hat{P}} \mathbb{E}_{x \sim X, l \sim \ell, s \sim S} [S(A(x, \hat{P}, l, s)) \odot m - t] \quad (3)$$

where X is a distribution of input images, ℓ is a distribution of patch locations, and S is a set of scales to resize \hat{P} . The operator A is the 'Apply' procedure which takes the current \hat{P} and inserts it into x while sampling uniformly on the distributions X , ℓ , and S .

We note that clipping must be applied to x to ensure the pixel values are within realistic bounds. However, in our study we do not perform clipping to understand the range of influence a remote patch can have across an image.

Loss Function. We experimented with many different loss functions on the CamVid dataset Brostov et al. (2009): \mathcal{L}_{CE} , \mathcal{L}_D , \mathcal{L}_1 , \mathcal{L}_2 , and \mathcal{L}_{KL} (Kullback–Leibler Divergence loss). Most of these loss functions took too long to converge or got stuck in local optima. Instead, we found that \mathcal{L}_1 works best in emptier scenes (like Fig. 2) and \mathcal{L}_{KL} works best in busy scenes like those in CamVid. We believe the reason why \mathcal{L}_{KL} performs well busy scenes is because it measures the relative entropy from one distribution to another. As a result, the optimizer had an easier

¹ For an x with a dimension of 384x480, we found that a radius of 50 pixels empirically performs best when targeting region with a radius of 10 pixels.




Fig. 5 The benefit of using the proposed incremental training strategy (incrementing the patch placement radius) when training on the BDD100k dataset

time ‘leeching’ nearby features and contexts in x to match the goal in t .

Creating a Robust RAP. In order to make an RAP which is robust to different transformations (scale and location), and universal to different images (not in the training set of the patch), we must use EoT. However, in some cases we found that the patch does not converge well when the range of (X, ℓ, S) is large (i.e., large shifts, hundreds of images, etc). This is because (1) the IPatch leverages the variable contents of x to impact $S(y'_s)$ and (2) the placement of P in x affects the influence of P on $S(y'_s)$.

For these cases, we propose an incremental training strategy where we gradually increase the placement radius of the patch ℓ . Whenever the training has converged or a time limit has elapsed, we increase the radius by one pixel. We repeat this process until the entire dataset is covered. At the start of each epoch, we give the optimizer time to adjust by setting the learning rate to a fraction of its value and then slowly ramp it back up. A similar strategy can be applied to the other distributions, such as the number of images in X , the shift size, or the patch scale. This strategy works well because we gradually increase the entropy, enabling the

optimizer capture foundational concepts. It can also be viewed that at each epoch we are placing the gradient descent optimizer at a more advantageous position instead of a random starting point.

To demonstrate the value of the incremental strategy, we performed an experiment. We trained a RAP using 70k images from the BDD100k street segmentation dataset Yu et al. (2018). The RAP was configured to make the center of the image perceived as the class ‘tree’ when placed in *any* location within the image (i.e., a placement radius of 500 pixels). The experiment was performed using (1) our incremental training strategy by increasing the placement radius up to the maximum radius and (2) the baseline approach of training the patch using the maximum radius from the start. For the incremental strategy, the placement radius was increased by one pixel whenever the attack success rate reached 25%. In Fig. 5, we plot the results from the experiment. We found that that our approach reaches the maximum radius after one hour and then exceeds the performance of the baseline shortly after by a margin of 35%.

In summary, the training framework for creating an IPatch is as follows (illustrated in Fig. 6):

Training Procedure for an IPatch

Initialize an initial patch \hat{P} with random values and set its origin (default location in x) to be o . If incremental, then add one image to X . Otherwise, add all images to X . Repeat until \hat{P} has converged on the entire dataset:

1. **Apply:** Draw a batch of samples from X . For each sample x in the batch, perform a random transformation: scale down \hat{P} and shift its location from origin o .
2. **Forward pass:** Pass the batch through S and obtain the segmentation maps (as a set of y'_s).
3. **Apply mask:** Take the product of each y'_s with the mask m to omit irrelevant semantics.
4. **Compute loss on objective**
5. **Compute gradient on P**
6. **Update P**

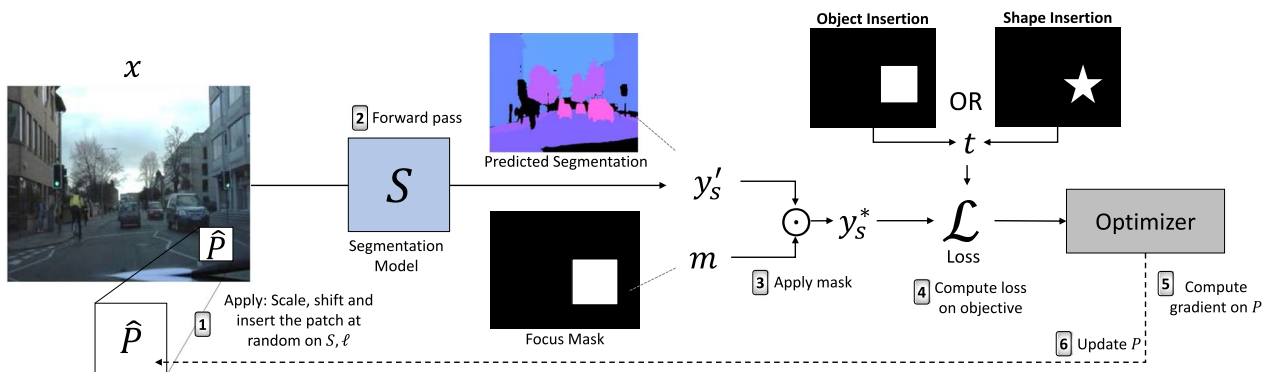


Fig. 6 An overview of the IPatch training framework for creating a RAP. For simplicity, only one sample x is shown, though in practice batches of images are used

Training Procedure for an IPatch

4. **Loss & gradient:** Compute the loss $\mathcal{L}_{KL}(Y_S^*, t)$ and use it to perform back propagation through S to \hat{P} .
 - 5-6. **Update:** Use gradient descent (e.g., Adamax) to update the values of \hat{P} .
 7. If `incremental` and the has time elapsed or training has converged, then increase the entropy (e.g., patch placement radius) and ramp the learning rate.
-

Evaluation

To evaluate the IPatch as a RAP, we will focus our evaluation on the scenario of autonomous vehicles. The task of street view segmentation is challenging because the scenes are typically very busy with many layers, objects, and wide perspectives Siam et al. (2018). Therefore attacking this application is will provide us with good insights into the IPatch's capabilities.

Datasets. We use the CamVid dataset Brostow et al. (2009) to train our segmentation models and evaluate our adversarial patches. The CamVid dataset is a well-known benchmark dataset used for image segmentation. It contains 46,869 street view images with a resolution of 360x480 from the point of view of a car. The images are supplied with pixel-wise annotations which indicate the class of the corresponding content (e.g., car, building, etc). The dataset comes split into three partitions: train D_{train} , test D_{test} , and validation D_{val} . We use D_{train} to train the segmentation models and the rest to train the patches. This way there is will be no bias on the images which we attack. The D_{test} dataset is used to evaluate the influence of the patch's parameters (size and location) and to train robust patches with EoT. Finally, the D_{val} dataset is use to validate that the robust patches work on unseen imagery.

Segmentation Models. In our evaluations we trained and attacked 37 different models which were combinations of 8 different encoders and 5 state-of-the-art segmentation architectures.² The encoders were the vgg19, densenet121, efficientnet-b4, efficientnet-b7, mobilenet_v2, resnext50_32x4d, dpn68, and xception. All of the models were obtained from the Torch library and were pretrained on the ImageNet Dataset Deng et al. (2009). For the architectures, we used the implementations³ of the state-of-the-art segmentation networks described in section "Technical background". The models were trained on D_{train} for 100 epochs each, with a batch size of 8, learning rate of 1e-4, using Dice Loss \mathcal{L}_D and an Adam optimizer. Finally, to increase the

training set size and improve generalization, we performed data augmentation. The augmentations were: flip, shift, crop, blur, sharpen and change perspective, brightness, and gamma.

The Experiments. We performed three experiments:

- EXP1:** In this experiment we investigate the influence which a patch's size and location have on the attack performance. We also investigate the influence of the remote target's size and location. Here patches are crafted to target individual images. Therefore, the results of this experiment also tell us how well the attack performs on static images.
- EXP2:** To use this attack in the wild, the patch must work under various transformations and in new scenes. This experiment evaluates the attack's robustness by (1) training the patches with EoT according to (3), and by (2) measuring the performance of these patches on new images (unseen during training).
- EXP3:** To get an idea of the vulnerability's prevalence, we attack 32 different segmentation models and measure their performance. To evaluate the case where the attacker has no information on the model, we take the robust patches trained in EXP2 and use them on the other 36 models to measure the attack's transferability.

When measuring attack performance, we omit all cases where the targeted region already contains the target class. For all of the experiments, we trained on an NVIDIA Titan RTX with 24GB of RAM. For the optimizer, we experimented on a variety of options in the Torch library. We found that the Adamax optimizer works best on the CamVid Dataset.

Patch Complexity. We found that the time it takes to train a patch varies depending on the difficulty of the objective function in equation (3). For example, if the patch is trained to work on a specific image and in a specific location, then it can take anywhere from 100 to 1,000 epochs to converge. This is about 3 min using an NVIDIA Titan RTX with 24GB RAM. Complexity can increase if the target t is further from the patch, or if batches of random images are used when training with EoT. For example, with EoT we found it can take anywhere from 4-24 h to train a single patch.

EXP1: the impact of size and location

The purpose of this experiment is to see how the size and locations of a patch and its target affect the attack's performance. In this experiment, we craft patches which

² Every combination of encoder and architecture except for the architecture PAN which was incompatible with three of the encoders.

³ https://github.com/qubvel/segmentation_models.pytorch

Table 2 The parameters used in EXP1. The values listed for 'size' are both the height and width

Experiment	Parameters			
	Patch P		Target t	
	Size* S	Location ℓ	Size r	Location L
Locked	80	(370,270)	20	(250,250)
Patch size	30:130	(370,270)	20	(250,250)
Patch location	80	**	20	(250,250)
Target size	80	(370,270)	20:280	(250,250)
Target location	80	(370,270)	20	***

* P is scaled down to 5 from 100x100. **On the diagonal from the image center to bottom right. ***8x6 grid over entire image

target a single image. Later in section "EXP2: patch robustness" we evaluate multi-image 'robust' patches.

Experiment Setup. For EXP1 we attacked the efficientnet-b7_FPN model since it performed best on the CamVid dataset. A list the evaluations and parameters used in EXP1 can be found in Table 2. For each of these parameters, we varied their values while locking the rest to measure their influence. This was repeated for each of the model's top six performing classes. Due to time restrictions,⁴ we only used the entirety of D_{test} for the fixed parameter experiment. For the other experiments, we used 20 random images from D_{test} .

The training procedure was as follows: For each patch, we used a learning rate of 2.5 and stopped the training after three minutes to ensure that each of the five experiments would take no more than 5 days. We note that in many cases, the patches were still converging so the results can be improved. Finally, we count a successful attack as any image with at least 80% of target t marked by the model as the target class.

Performance with all parameters locked

The results for the experiment, where the patch parameters are locked, can be found in Fig. 7. The top of the figure shows that the attack has a greater impact on structural classes than others. This might be because these semantics have the largest regions CamVid dataset (i.e., are common). As a result, the patch is able to leverage these contexts better from one side of an image to another. For example, if there are is a row of buildings on one side of the road, then there is a higher probability that the other side will have one too. This kind of correlation is exploited by the patch. The road class out performs all the rest because the target t in this experiment

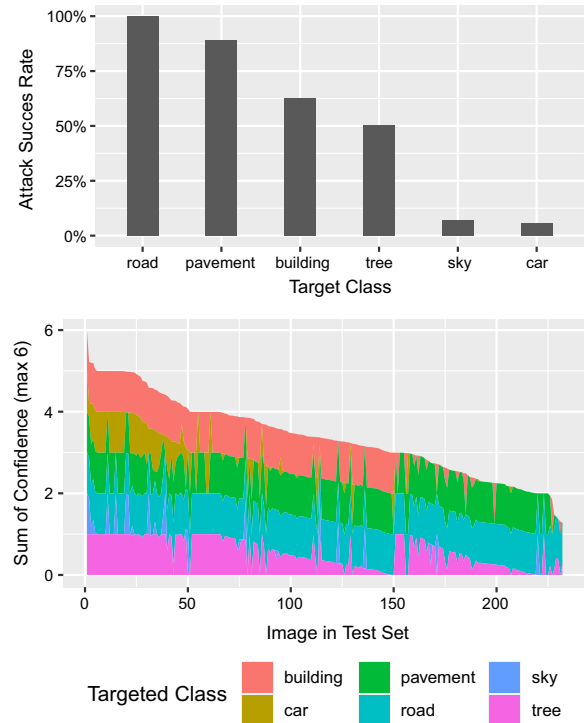


Fig. 7 The attack performance with locked parameters. Top- The average attack performance per class. Bottom- the aggregated confidence of the model on each image, where the max confidence per class is one

is in the center of the image, where the road is most commonly found (77% of the images). However, the patch is able to successfully attack the classes of pavement, building, and tree at the same location, even though on the clean images, the model predicts 1.4%, 0.3%, and 0% of them to have these classes respectively (top of Fig. 7).

At the bottom of Fig. 7 we can see the aggregated confidence of the model for each of the images in D_{test} . The plot shows that all of the images are susceptible to the attack for at least one of the target classes.

Impact of the patch size

Figure 8 plots the model's confidence over increasingly larger patch sizes. In the figure, we have marked 0.5 as the decision threshold which is the default for segmentation models. This is because segmentation models perform binary-classification on each pixel. As a result, the confidence scores per class are either close to zero or one, but not so much in between (as seen in Fig. 7).

As expected, larger patches increase the attack success rate. However, the trade off appears to be linear (captured by the average in red). What is meaningful about these results is that some classes excel with smaller patch sizes (e.g., pavement and building) while others require larger ones to succeed (e.g., tree). This is probably because some

⁴ Each of these experiments on takes 3-5 days on a NVIDIA Titan RTX with 24GB RAM.

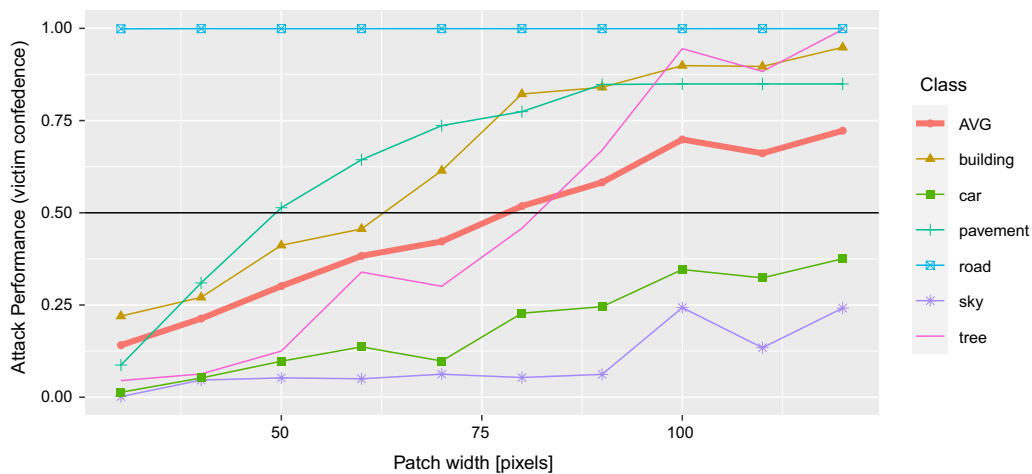


Fig. 8 The average attack performance with different patch sizes




Fig. 9 The average attack performance with different patch locations, shifted from near-center of x to the bottom-right

of the remote contextual semantics which the model considers cannot be compressed into small spaces when others can. Overall, we observe that the minimum patch size required to fool the model on a static image this size is about 60–75 pixels in width, and with a patch width of 100 pixels, nearly all attacks succeed.

Impact of the patch location

In Fig. 9 we can see that the attack is highly effective for all classes up to about 62% of the distance away from the target (image center). The sharp drop in attack performance for the tree and sky classes is understandable since there are fewer contextual semantics which can be exploited by the patch in the bottom right of the image. On the other hand, in areas just below the horizon (0–0.5 on the x-axis), the patch can exploit contextual semantics which the model uses (e.g., features such as lighting, reflections, and building geometry).

These results indicate that an attacker may be able to increase the likelihood of success by placing the patch on objects which have some contextual influence on the target region. For example, to create a crosswalk, it may be advantageous to put the sticker on a lamp post or parking meter since these objects may be found near crosswalks.

Impact of the target size

In this experiment, we increased the size of target t but observe the performance of the same 20x20 pixel region at the center of t (i.e., our objective). In Fig. 10 we can see that large targets do not perform well. The reason for this is that having a large target requires the IPatch to subdue more semantics. As a result, the patch fails and the region of t becomes patchy and an corrupted. Small targets fail because it is hard for the patch to make high precision results. Rather, there is a balance between the intended 20x20 target and the actual target painted in t .

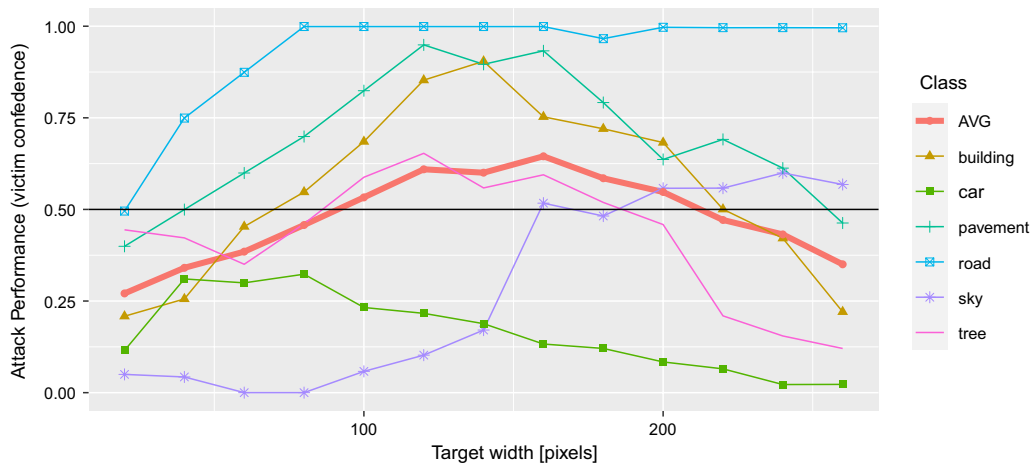


Fig. 10 The average attack performance on the same 20x20 region, but with different target sizes




Fig. 11 The remote attack performance on ten images. A region in the heat map indicates the performance of the attack when targeting that location

We found that increasing the target size by a factor of 3 improves the performance at the intended region.

The reason why a larger target helps the patch reach the 20x20 region is that the patch tends to ‘leech’ nearby semantic regions. This makes sense since it is easier to change the boundaries of existing semantics (e.g., perceive a larger car) than generate new ones which are isolated (e.g., a tree in the middle of the road). Therefore, the added target size encourages the model to perform similar tactics.

Impact of the target location

In Fig. 11 we present the attack performance when targeting different remote locations in x . It is clear that the influence of a patch on different regions is dependent on both the image’s content and the targeted class. For example, it is easier to convince the model that any space under the horizon is a road, yet it is hard to change the class of the top-center to building because it is rarely found there. Overall, this experiment demonstrates that the patch can target locations on far remote locations

within the image. However, this capability is not uniform across the classes, as we can see with the class ‘car’.

EXP2: patch robustness

In this experiment we evaluate how well a single patch performs on (1) different transformations and (2) on multiple seen and unseen images.

Experiment Setup. To perform this experiment, we used EoT (3) to train a single patch for each class. The incremental training framework from 4.2 was used with the patch origin o set to (370,270). For the patch size S , we sampled uniformly on the range of [50,80] pixels. For the shifts ℓ , we sampled uniformly within the entire bottom-right quadrant of x . We found that training the patch in one region helps it converge using the incremental strategy, while still generalizing to the opposite side. We targeted the same segmentation model used in EXP1, and the training was performed using a batch size of 20 (the maximum for a 24GB GPU) with a learning rate of 0.5.

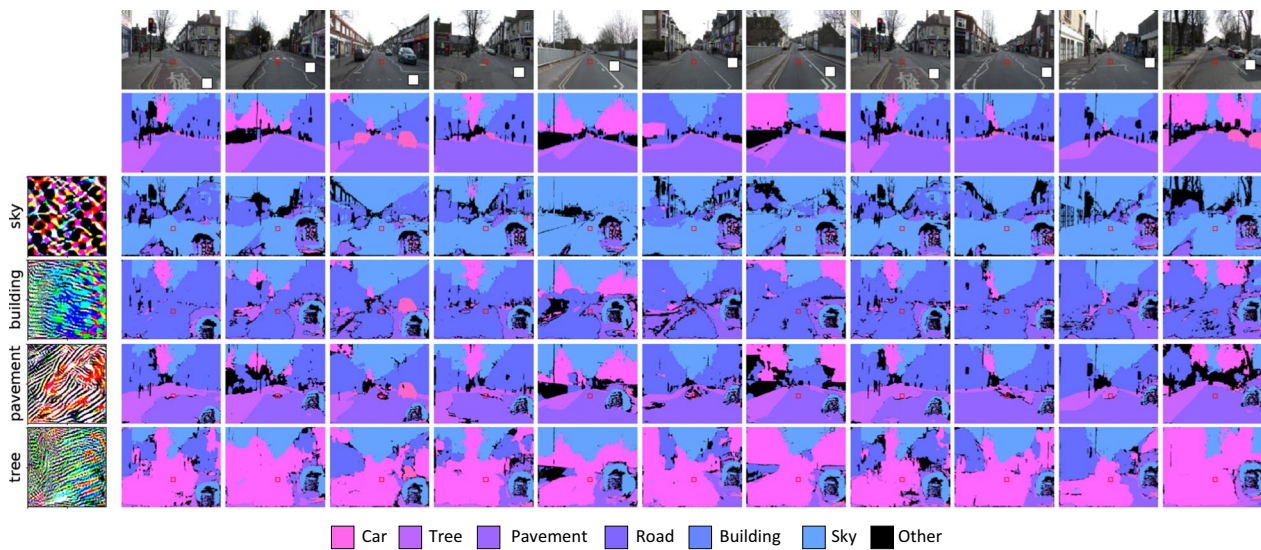


Fig. 12 Examples showing the semantic segmentation maps of *efficientnet-b7_FPN* when attacked with 'robust' patches trained using EoT. Top two rows- the original image and the original segmentation map, where the target and patch are marked in red and white accordingly. Right- the IPatch used in the attack, each one has been trained for a different class but the same remote target location




Fig. 13 The performance of a robust patch when trained on different set sizes

Generalization to Multiple Images In Fig. 12 we present the performance of the patches in the form of the model's perception. The images demonstrate that the IPatch generalizes well to multiple images, even at

different locations and scales. In Fig. 13 we present the attack performance when training on different numbers of images from D_{test} (evaluated against the same set).




Fig. 14 The average attack performance of the patches when applied to new images

From here we can see that an exponential number of examples are needed to increase the performance.

Generalization to New Images. To use the patch in a real world setting, it must work well in scenes which were not in the attacker's training set. Figure 14 presents the attack performance of patches trained on D_{test} (those displayed in Fig. 12) when applied to images in D_{val} . The results show that the patches generalize well to unseen images. More interestingly, the performance of some classes are dramatically different compared to patches trained on single images without EoT (EXP1, Fig. 7). For example, 'tree' now has a 0.98% success rate compared to 50% and 'building' is now 25% compared to 65%. We learn from this that by considering multiple images, the model can learn stronger tactics. At the same time, the variability of the transformations prevent the model from using highly specific adversarial patterns. We also note that the class 'car' does not transfer to unseen images like the other classes. We attribute this to the segmentation model's poor performance on detecting cars in general.⁵

EXP3: the impact on different models

In is experiment we explore the suceptibilty and transferability of patches between models.

Model suceptibility

Experiment Setup. To evaluate the performance of the attack on different model architectures, we used 32 of the 37 segmentation models described at the beginning of section "Evaluation" (PAN was omitted since it was not compatible with Torch's autograd in our framework). Due to time limitations, the attacks on each model were limited to 4 classes, 10 images, and 3 min training time for each image.

Results. We found that all 36 models are susceptible to the RAP attack for at least one class(Fig. 15). By

Encoder	building				car			
	FPN	Linknet	PSPNet	UnetPlusPlus	FPN	Linknet	PSPNet	UnetPlusPlus
xception -	0.99	0.99	0.99	0.61	0.99	0.96	0.21	0.99
vgg19 -	0.99	0.99	0.94	0.99	0.99	0.9	0.73	0.99
resnext50_32x4d -	0.89	0.71	0.94	0.79	0.99	0.87	0.99	0.99
mobilenet_v2 -	0.98	0.64	0.99	0.5	0.99	0.36	0.11	0.99
efficientnet-b7 -	0.78	0.98	0.8	0.61	0.95	0.93	0.03	0.99
efficientnet-b4 -	0.99	0.91	0.77	0.79	0.99	0.13	0.09	0.99
dpn68 -	0.85	0.88	0.26	0.59	0.7	0.53	0	0.89
densenet121 -	0.99	0.98	0.94	0.79	0.99	0.94	0.99	0.99
Encoder	pavement				tree			
	FPN	Linknet	PSPNet	UnetPlusPlus	FPN	Linknet	PSPNet	UnetPlusPlus
xception -	0.99	0.68	0.21	0.98	0.99	0.88	0.99	0.9
vgg19 -	0.99	0.48	0.27	0.33	0.99	0.94	0.82	0.84
resnext50_32x4d -	0.99	0.44	0.04	0.97	0.99	0.86	0.89	0.79
mobilenet_v2 -	0.99	0.08	0.11	0.87	0.99	0.42	0.99	0.51
efficientnet-b7 -	0.9	0.17	0.03	0.01	0.99	0.91	0.92	0.7
efficientnet-b4 -	0.99	0.06	0.04	0	0.99	0.36	0.57	0.99
dpn68 -	0.79	0.08	0	0.9	0.31	0.14	0.49	0.54
densenet121 -	0.98	0.61	0.1	0.01	0.99	0.93	0.99	0.89

Fig. 15 The average confidence of different models when attacked with an IPatch. Values greater than 0.5 indicate a successful attack on average

observing the patterns in the columns, we note that some architectures are less susceptible to attacks on certain classes. For example, Linknet, PSPNet, and Unet++ on pavement and PSPNet on car.

In Fig. 16, we can see the suceptibilty of the encoders and architectures overall. Some of the most susceptible encoders (xception and resnext) and architectures (FPN and Unet++) use skip connections or residual pathways in their networks. These pathways enable the networks to capture features at multiple scales and capture the global contexts better. However, just as these network utilize these pathways to obtain better perspectives, so can the IPatch in order to reach deeper into the image. Interestingly we found that the dpn68 encoder is consistently resilient against the attack. This encoder is formally called a Dual Path Network Chen et al. (2017). It uses a residual path like a ResNet to reuse learned features and a densely connected path like DenseNet to encourage the network to explore new features. These diverse features may be preventing the IPatch, and possibly the segmentation model, from reaching remote contexts.

Inter-model transferability

In the case where the attacker does not have knowledge of the victim's model, we would like to know well a patch trained on one model transfers to others.

⁵ Although the selected efficientnet-b7_FPN achieves a lesser intersection over union score of 0.75 on that class, it outperforms the other models overall.




Fig. 16 The average confidence of different encoders and architectures when under attack




Fig. 17 The average attack success rate on each model when attacked with a patch trained on `efficientnet-b7_FPN`

Experiment Setup. To perform this experiment, we took the robust patches trained using the `efficientnet-b7_FPN` (EXP2) and attacked each of the other 36 models (listed in Fig. 17). The patches for the top 4

classes (sky, building, pavement, tree) were applied to the images in D_{test} using the random transformations described in EXP2.

Results. We found that patch from `efficientnet-b7_FPN` can influence the other models' predictions on the target region with an attack success rate of 11-37% (about 1-4 times in every 10 cases). We note that a cameras on an autonomous car processes at least 30 frames per second. Therefore, there is a high likelihood that the car's model will be susceptible to the attack while driving by.

Figure 17 shows the largest confidences for each model, measured as the relative increase from the original confidence (on clean image). Interestingly models using the `Unet++` architecture were the most susceptible, followed by `Linknet`. We believe the reason for this is that both of these models use skip-connections to allow for feature maps to bypass the encoding process. As a result, features in the patch have a more direct impact on the output. It is known that skip-connections make models more vulnerable to adversarial examples Wu et al. (2020) but it is interesting to see that they are vulnerable to transfer attacks as well. Another observation is that there does not seem to be a correlation between the results and the encoder used. This is probably because all of the encoders were trained on the same ImageNet Dataset.

Finally, we note that there are ways in which an adversarial example can be made to be more transferable Xie et al. (2019); Huang et al. (2019). We leave this for future work.

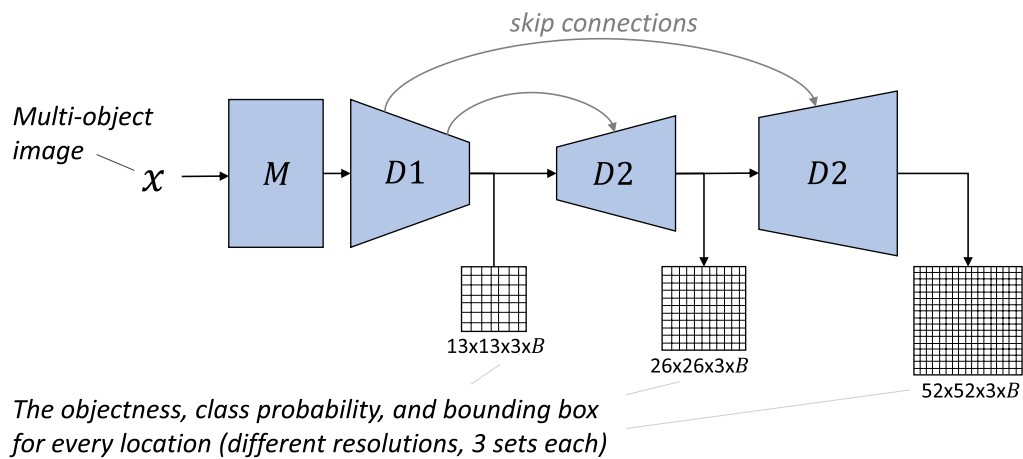


Fig. 18 The architecture of the YOLOv3 object detector

Extending to object recognition

In section "Evaluation" we performed an in-depth evaluation and analysis of the IPatch as a RAP against segmentation models. However, the same training framework in "Approach" can be used on other semantic models as well. In this section, we present preliminary results against a popular object recognition model called YOLOv3 Redmon and Farhadi (2018).

Technical background

The family of YOLO models follow a similar architecture (Fig. 18). The image x is passed through a series of convolutional layers (M in the figure) and then those feature maps are shuttled to various decoders. The decoders predict coarse maps to the image at different scales using the semantic information shared between them. The multiple scales help the model detect objects of different sizes (e.g., $D1$ detects large objects). Each cell in a map, contains an objectness score, class probability, and a bounding box (obtained via regression). If a cell has an objectness score above some threshold, then there is an object there with the associated class probability. Finally, a non-maximal suppression (NMS) algorithm is used on the maps to identify and unify the detections.

Evaluation

Experiment Setup. To see if the attack would work on YOLO, we created an IPatch which convinces YOLO that there is a person standing in the middle of the road. To accomplish this, we used a pre-trained YOLOv3 model implementation⁶ as the victim, and trained our patch using 30k images: 15k random samples from the Bdd100k

dataset Yu et al. (2020) and 15k frames from a Toronto car driving video on YouTube.⁷

For training, we needed to ensure that both the objectness score and probability of the class 'person' were high. This was done by taking the product of $D1$'s probability map and objectness map as y'_s , and by setting t to highlight the cell in the lower-center of the image. For the loss functions, we took the sum of \mathcal{L}_{KL} and \mathcal{L}_1 since it increased the rate of convergence. EoT was used to scale




Fig. 19 An example of the YOLOv3 object detector being attacked by an IPatch. Here, the patch has convinced the model that there is a person standing in the middle of the road

⁶ <https://github.com/eriklindernoren/PyTorch-YOLOv3>

⁷ https://youtu.be/50Uf_T12OGY

the patch between 60-70 pixels in width and shift it randomly within the bottom-right quadrant of the image. Finally, we trained the patch for 3 days with a learning rate of 0.05.

Results. We found that the YOLOv3 object detector is susceptible to the attack with an 85% attack success rate. In Fig. 19 we present an example frame which shows the objectness and probability maps in $D1$ during the attack. We also found that smaller patches ranging from 50-60 pixels in width achieve an 80% attack success rate. Overall, it was relatively easy for the framework to change the objectness score of arbitrary locations in the image, compared to the class probability. We also observed that it is significantly harder to target the maps from $D2$ and $D3$ which capture smaller objects. We believe this is because $D2$ and $D3$ rely less on contexts around the image, giving the IPatch less leverage to perform a remote attack.

As future work, we plan to explore RAPs on other object detectors and investigate other semantic models as well.

Discussion & countermeasures

The concept of a remote adversarial patch, introduced in this paper, opens up wide range of possible attack vectors against image-based semantic models. Through our observations in 5, we were able to identify some of the attack's capabilities and limitations.

Trade-Offs. Due to its flexibility, it may seem like the IPatch is harder to defend against compared to an ordinary adversarial patch. However, the performance of the patch is less compared to a 'point-based' patch. This means that the adversary must consider whether a more reliable attack needed over having flexibility and stealth. Another consideration is that the adversary may want to experiment to find the optimal placement of the patch. This is because some regions give the patch more leverage based on the local semantics (section "Impact of the patch location"). One strategy is that the attacker can first scout the target region by videoing the scene from multiple perspectives and then optimize the patch location using that dataset.

Defenses. Although the IPatch can be placed in arbitrary locations, we noticed that its presence highly noticeable in the semantic segmentations (e.g., Fig. 12). We found that it is very hard to generate a patch which both achieves the attack and masks its own presence at the same time. Concretely, when setting $t = y'_s$ except for the target region (as done in Fig. 2), we found that the model struggles to influence remote locations to the same extent. In future work, this may be improved through a custom loss function which balances the trade-off between the two objectives. Another solution might

be to generate RAPs using a conditional GAN which considers the errors on the semantic map in $(-m)$. Doing so may also reduce the corruptions to nearby semantics as well.

Another direction for defending against this attack is to limit the model's dependency on global features. Although these global features are key to state-of-the-art models Lin et al. (2017); Li et al. (2018); Fan et al. (2020), it is possible to utilize them while also considering their layout and origin. One option may be to integrate capsule networks Sabour et al. (2017) as part of the model's architecture, since capsule networks are good at considering the spatial relationship in images.

Improvements. We noticed that the RAP attack is dependent on an image's content when targeting segmentation models, but less so for the object detector YOLO. For example, we were able to perform remote adversarial attacks on a blank image x with YOLO. The reason for this is not clear to us, and investigating it may lead to improvements in the proposed training methodology. Moreover, as future work, it would be interesting to investigate which types of features and classes the a RAP can manipulate best and why. This research may lead to deeper insights into the vulnerability's extents and limitations. Finally, to improve transferability, we suggest two directions: (1) include multiple models in the training loop to help the model identify common features, and (2) use adversarial training to improve the generalization of the patch. As future work, it would also be interesting to see if an incremental strategy can be used with clipping (to pixel values 0-255) to produce strong real-world patches.

Conclusion

In this paper, we have introduced the concept of a 'remote adversarial patch' (RAP) which can alter the semantic interpretation of an image while being placed anywhere within the field of view. We have implemented an RAP called IPatch. When generated without pixel clipping, we demonstrated that it is robust, can generalize to new scenes, and can impact other semantic models such as object detectors. With an average attack success rate of up to 93%, this attack forms a tangible threat. As future work, we plan to investigate the range of RAPs with clipping enforced. Although RAPs are in their infancy, we hope that this paper has laid some of the groundwork for exploring this new adversarial example.

In summary, neural networks are notorious for being black-boxes which are difficult to interpret. However, they are still used in critical tasks because their advantages outweigh their potential disadvantages. We hope

that our findings will help the community improve the security of deep learning applications so that we may continue to benefit from safe and reliable autonomous systems.

Acknowledgements

This material is based upon work supported by the Zuckerman STEM Leadership Program.

Author contributions

Dr Mirsky is the only author and has conceptualized and performed all of the experiments and writing of the paper. The author have read and approved the final version of the manuscript.

Funding

Not applicable

Declarations

Competing interests

Not applicable

Received: 10 July 2022 Accepted: 21 February 2023

Published online: 03 May 2023

References

- Arnab A, Miksik O, Torr PH (2018) On the robustness of semantic segmentation models to adversarial attacks. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 888–897
- Athalye A, Engstrom L, Ilyas A, Kwok K (2018) Synthesizing robust adversarial examples. In: International conference on machine learning. PMLR, pp. 284–293
- Audebert N, Le Saux B, Lefèvre S (2017) Segment-before-detect: vehicle detection and classification through semantic segmentation of aerial images. *Remote Sens* 9(4):368
- Autopilot ai-tesla (2021) <https://www.tesla.com/autopilotAI>, (Accessed on 02/04/2021)
- Biggio B, Roli F (2018) Wild patterns: ten years after the rise of adversarial machine learning. *Pattern Recognit* 84:317–331
- Brostow GJ, Fauqueur J, Cipolla R (2009) Semantic object classes in video: a high-definition ground truth database. *Pattern Recognit Lett* 30(2):88–97
- Brown TB, Mané D, Roy A, Abadi M, Gilmer J (2017) Adversarial patch. *arXiv preprint arXiv:1712.09665*
- Chaurasia A, Culurciello E (2017) Linknet: Exploiting encoder representations for efficient semantic segmentation. In: 2017 IEEE visual communications and image processing (VCIP). IEEE 2017:1–4
- Chen S-T, Cornelius C, Martin J, Chau DHP (2018) Shapeshifter: Robust physical adversarial attack on faster r-cnn object detector. In: Joint European conference on machine learning and knowledge discovery in databases. Springer, pp. 52–68
- Chen Y, Li J, Xiao H, Jin X, Yan S, Feng J (2017) Dual path networks. *arXiv preprint arXiv:1707.01629*
- Chow K-H, Liu L, Loper M, Bae J, Gursoy ME, Truex S, Wei W, Wu Y (2020) Adversarial objectness gradient attacks in real-time object detection systems. In: 2020 second IEEE international conference on trust, privacy and security in intelligent systems and applications (TPS-ISA). IEEE, pp. 263–272
- den Hollander R, Adhikari A, Tolios I, van Bekkum M, Bal A, Hendriks S, Kruijthof M, Gross D, Jansen N, Perez G et al. (2020) Adversarial patch camouflage against aerial detection. In: Artificial intelligence and machine learning in defense applications II, vol. 11543. International society for optics and photonics, p. 115430F
- Deng J, Dong W, Socher R, Li L-J, Li K, Fei-Fei L (2009) Imagenet: a large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. IEEE 2009:248–255
- Deng R, Shen C, Liu S, Wang H, Liu X (2018) Learning to predict crisp boundaries. In: Proceedings of the European conference on computer vision (ECCV), pp. 562–578
- Fan T, Wang G, Li Y, Wang H (2020) Ma-net: a multi-scale attention network for liver and tumor segmentation. *IEEE Access* 8:179656–179665
- Fischer V, Kumar MC, Metzen JH, Brox T (2017) Adversarial examples for semantic image segmentation. *arXiv preprint arXiv:1703.01101*
- García-García A, Orts-Escolano S, Oprea S, Villena-Martínez V, Martínez-González P, García-Rodríguez J (2018) A survey on deep learning techniques for image and video semantic segmentation. *Appl Soft Comput* 70:41–65
- Goodfellow IJ, Shlens J, Szegedy C (2014) Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*
- Hendrik Metzen J, Chaithanya Kumar M, Brox T, Fischer V (2017) Universal adversarial perturbations against semantic image segmentation. In: Proceedings of the IEEE international conference on computer vision, pp. 2755–2764
- He J, Wu X, Yang J, Hu W (2020) Cpsnet: crowd counting via semantic segmentation framework. In: 2020 IEEE 32nd international conference on tools with artificial intelligence (ICTAI). IEEE, pp. 1104–1110
- Hoory S, Shapira T, Shabtai A, Elovici Y (2020) Dynamic adversarial patch for evading object detection models. *arXiv preprint arXiv:2010.13070*
- Huang L, Gao C, Zhou Y, Xie C, Yuille AL, Zou C, Liu N (2020) Universal physical camouflage attacks on object detectors. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 720–729
- Huang Q, Katsman I, He H, Gu Z, Belongie S, Lim SN (2019) Enhancing adversarial example transferability with an intermediate level attack. In: Proceedings of the IEEE/CVF international conference on computer vision, pp. 4733–4742
- Igloukov V, Mushinskiy S, Osin V (2017) Satellite imagery feature detection using deep convolutional neural network: a kaggle competition. *arXiv preprint arXiv:1706.06169*
- Kang X, Song B, Du X, Guizani M (2020) Adversarial attacks for image segmentation on multiple lightweight models. *IEEE Access* 8:31359–31370
- Lee M, Kolter Z (2019) On physical adversarial patches for object detection. *arXiv preprint arXiv:1906.11897*
- Leetaru K (2019) Waymo reminds us: Successful complex ai combines deep learning and traditional code, <https://www.forbes.com/sites/kalevleeta/2019/04/18/waymo-reminds-us-successful-complex-ai-combines-deep-learning-and-traditional-code/?sh=4eb596723a2c>, (Accessed on 02/04/2021)
- Li Y, Xu X, Xiao J, Li S, Shen HT (2020) Adaptive square attack: fooling autonomous cars with adversarial traffic signs. *IEEE Internet of Things J* 8(8):6337–6347
- Lin T-Y, Dollár P, Girshick K, He B, Hariharan, and S. Belongie (2017) Feature pyramid networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2117–2125
- Liu X, Yang H, Liu Z, Song L, Li H, Chen Y (2018) Dpatch: an adversarial patch attack on object detectors. *arXiv preprint arXiv:1806.02299*
- Li H, Xiong P, An J, Wang L (2018) Pyramid attention network for semantic segmentation. *arXiv preprint arXiv:1805.10180*
- Li Y, Xu G, Li W (2020) Fa: a fast method to attack real-time object detection systems. In: 2020 IEEE/CIC international conference on communications in China (ICCC). IEEE, pp. 1268–1273
- Luo Y, Boix X, Roig G, Poggio T, Zhao Q (2015) Foveation-based mechanisms alleviate adversarial examples. *arXiv preprint arXiv:1511.06292*
- Lu J, Sibai H, Fabry E, Forsyth D (2017) No need to worry about adversarial examples in object detection in autonomous vehicles. *arXiv preprint arXiv:1707.03501*
- Milletari F, Navab N, Ahmadi S-A (2016) V-net: fully convolutional neural networks for volumetric medical image segmentation. In: 2016 fourth international conference on 3D vision (3DV). IEEE 2016:565–571
- Minaee S, Boykov Y, Porikli F, Plaza A, Kehtarnavaz N, Terzopoulos D (2020) Image segmentation using deep learning: a survey. *arXiv preprint arXiv:2001.05566*
- Mirsky Y, Mahler T, Shelef I, Elovici Y (2019) Ct-gan: Malicious tampering of 3d medical imagery using deep learning. In: 28th USENIX security symposium (USENIX Security 19). Santa Clara, CA: USENIX Association, pp. 461–478. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity19/presentation/mirsky>

- Nguyen A, Yosinski J, Clune J (2015) Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 427–436
- Ozbulak U, Van Messem A, De Neve W (2019) Impact of adversarial examples on deep learning models for biomedical image segmentation. In: International conference on medical image computing and computer-assisted intervention. Springer, pp. 300–308
- Prajna Y, Nath MK (2021) A survey of semantic segmentation on biomedical images using deep learning. In: Advances in VLSI, Communication, and Signal Processing. Springer, pp. 347–357
- Redmon J, Farhadi A (2018) YoloV3: an incremental improvement. arXiv preprint [arXiv:1804.02767](https://arxiv.org/abs/1804.02767)
- Ren S, He K, Girshick R, Sun J (2015) Faster r-cnn: towards real-time object detection with region proposal networks. arXiv preprint [arXiv:1506.01497](https://arxiv.org/abs/1506.01497)
- Sabour S, Frosst N, Hinton GE (2017) Dynamic routing between capsules. arXiv preprint [arXiv:1710.09829](https://arxiv.org/abs/1710.09829)
- Siam M, Gamal M, Abdel-Razek M, Yogamani S, Jagersand M, Zhang H (2018) A comparative study of real-time semantic segmentation for autonomous driving. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops, pp. 587–597
- Siemens (2021) Deep resolve. <https://www.siemens-healthineers.com/en-us/magnetic-resonance-imaging/technologies-and-innovations/deep-resolve>, (Accessed on 02/04/2021)
- Sitawarin C, Bhagoji AN, Mosenia A, Chiang M, Mittal P (2018) Darts: deceiving autonomous cars with toxic signs. arXiv preprint [arXiv:1802.06430](https://arxiv.org/abs/1802.06430)
- Song D, Eykholt K, Evtimov I, Fernandes E, Li B, Rahmati A, Tramer F, Prakash A, Kohno T (2018) Physical adversarial examples for object detectors. In: 12th USENIX workshop on offensive technologies (WOOT 18)
- Szegedy C, Zaremba W, Sutskever I, Bruna J, Erhan D, Goodfellow I, Fergus R (2013) Intriguing properties of neural networks. arXiv preprint [arXiv:1312.6199](https://arxiv.org/abs/1312.6199)
- Teichmann M, Weber M, Zoellner M, Cipolla R, Urtasun R (2018) Multinet: real-time joint semantic reasoning for autonomous driving. In: 2018 IEEE intelligent vehicles symposium (IV). IEEE 2018:1013–1020
- Thys S, Van Ranst W, Goedemé T (2019) Fooling automated surveillance cameras: adversarial patches to attack person detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops
- Vincent J (2018) Artificial intelligence is going to supercharge surveillance the verge. <https://www.theverge.com/2018/1/23/16907238/artificial-intelligence-surveillance-cameras-security>, (Accessed on 02/04/2021)
- Wei X, Liang S, Chen N, Cao X (2018) Transferable adversarial attacks for image and video object detection. arXiv preprint [arXiv:1811.12641](https://arxiv.org/abs/1811.12641)
- Wu Z, Lim S-N, Davis LS, Goldstein T (2020) Making an invisibility cloak: Real world adversarial attacks on object detectors. In: European Conference on Computer Vision. Springer, pp. 1–17
- Wu D, Wang Y, Xia ST, Bailey J, Ma X (2020) Skip connections matter: on the transferability of adversarial examples generated with resnets. arXiv preprint [arXiv:2002.05990](https://arxiv.org/abs/2002.05990)
- Xie C, Wang J, Zhang Z, Zhou Y, Xie L, Yuille A (2017) Adversarial examples for semantic segmentation and object detection. In: Proceedings of the IEEE international conference on computer vision, pp. 1369–1378
- Xie C, Zhang Z, Zhou Y, Bai S, Wang J, Ren Z, Yuille AL (2019) Improving transferability of adversarial examples with input diversity. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 2730–2739
- Yu F, Chen H, Wang X, Xian W, Chen Y, Liu F, Madhavan V, Darrell T (2020) Bdd100k: a diverse driving dataset for heterogeneous multitask learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 2636–2645
- Yu F, Xian W, Chen Y, Liu F, Liao M, Madhavan V, Darrell T (2018) Bdd100k: a diverse driving video database with scalable annotation tooling. arXiv preprint [arXiv:1805.04687](https://arxiv.org/abs/1805.04687), vol. 2, no. 5, p. 6
- Zhang H, Zhou W, Li H (2020) Contextual adversarial attacks for object detection. In: 2020 IEEE international conference on multimedia and expo (ICME). IEEE, pp. 1–6
- Zhao H, Shi J, Qi X, Wang X, Jia J (2017) Pyramid scene parsing network. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2881–2890
- Zhao Y, Yan H, Wei X (2020) Object hider: Adversarial patch attack against object detectors. arXiv preprint [arXiv:2010.14974](https://arxiv.org/abs/2010.14974)
- Zhao Y, Zhu H, Liang R, Shen Q, Zhang S, Chen K (2019) Seeing isn't believing: towards more robust adversarial attack against real world object detectors. In: Proceedings of the 2019 ACM SIGSAC conference on computer and communications security, pp. 1989–2004
- Zhou Z, Siddiquee MMR, Tajbakhsh N, Liang J (2018) Unet++: A nested u-net architecture for medical image segmentation. In: Deep learning in medical image analysis and multimodal learning for clinical decision support. Springer, pp. 3–11
- Zolfi A, Kravchik M, Elovici Y, Shabtai A (2020) The translucent patch: a physical and universal attack on object detectors. arXiv preprint [arXiv:2012.12528](https://arxiv.org/abs/2012.12528)

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)