

presents and discusses the experimental results for studying the efficacy of the proposed strategy on several benchmark datasets (CIFAR-10 and NIPS-W). Finally, conclusions are drawn in Section 4.

2. Methodology

In this section, we will provide a detailed description of the proposed strategy for leveraging gradient-based interpretability in the realm of adversarial examples, which is designed to improve the speed of convergence for producing adversarial perturbations by using insights gained to aid adversarial learning.

2.1. Spatially constrained one-pixel adversarial perturbations

In one-pixel adversarial perturbations, the optimal pixel is found by using Differential Evolution (DE). For every individual adversarial perturbation, a set of N vectors in \mathbb{R}^5 - where each vector p represents a candidate pixel's xy -coordinates and RGB values - is randomly generated, giving the initial parent population. This initialization serves as input to the differential evolution optimization. For each iteration during optimization, N children are generated from the parent population, and the fittest pixels i.e., the ones providing the lowest probabilistic label for the correct class remain to become N parents in the next iteration.

We hypothesize that the probability of susceptibility is highest in areas that are identified as highly sensitive via gradient-based interpretability. Therefore, we posit that using sensitivity maps generated via gradient-based interpretability as a spatial constraint when generating the initial parent population should speed up DE convergence.

Fig. 1 shows an overview of the pipeline for the generation of spatially constrained one-pixel adversarial perturbations on an image.

2.2. Susceptibility set generation

The first step involves the identification of a set of susceptible pixels S based on insights gained via gradient-based interpretability. More specifically, we leverage SmoothGrad [8] to obtain a sensitivity map s for each image x , with each pixel in s providing a quantitative indicator of importance of the underlying content to the decision making process. To identify the set of susceptible pixels S , binary thresholding is performed on the sensitivity map s as follows:

$$s(x) = \frac{1}{n} \sum_1^n \hat{s}(x + \mathcal{N}(0, \sigma^2)); \quad S = \{s(x) > \tau\} \quad (1)$$

This is based on the assumption that the pixels with higher sensitivity in the generated sensitivity maps are more susceptible to attack.



(a) Susceptibility set generation.



(b) One-pixel adversarial perturbation generation.

Figure 1: Pipeline for the generation of spatially constrained one-pixel adversarial perturbations. (a) The set of susceptible pixels is identified via sensitivity thresholding on a sensitivity map obtained using SmoothGrad. (b) The pixel perturbation, identified based on spatial constraint via the set of susceptible pixels, is performed to the image to obtain the adversarial example. The coordinate of the adversarially perturbed pixel is shown on top of the output image.

2.3. Adversarial perturbation generation

The second step involves leveraging the set of susceptible pixels for guiding the adversarial learning process. More specifically, when initializing the population for DE, the X and Y values in each vector p for given image x are constrained to be selected from within the set of S i.e., $X, Y \in S$ to ensure that the initial population of pixels all lie in the susceptible regions of the image.

$$p(x) = \{X, Y, r, g, b\}; \quad (X, Y) \in S$$

$$\max_{p(x)} f_{adv}(x + p(x)) \quad \vee \quad \|p(x)\|_0 \leq 1 \quad (2)$$

3. Experiments and Results

To evaluate the efficacy of spatially constrained one-pixel adversarial perturbations in reducing the convergence time of DE, we conducted different experiments on the DE algorithm (both in duration as well as average number of iterations) at finding a susceptible pixel with and without constrained initialization. Experiments were done for two different datasets: i) CIFAR-10 [4], and ii) NIPS 2017 adversarial attacks and defences challenge dataset [2] (small subset derived from ImageNet), referred to as NIPS-W in



Figure 2: Example images from ImageNet (left) and the corresponding spatially constrained one-pixel adversarial perturbed image (right) that led to successful attacks. The coordinates of the perturbed pixels are shown on the right-hand side.

the rest of the paper. The network tested in this study was SqueezeNet.

We select the default set of 100 images taken from the CIFAR-10 dataset in [9]. Additionally, for conducting a larger set of experiments, we randomly select 500 images from NIPS-W. For all experiments, we set the threshold τ at 0.5. For each dataset, five different runs of the adversarial attack were performed with different random seeds.

Table 1: Results of constrained one-pixel adversarial perturbations on 100 images from CIFAR-10 dataset. For both constrained and unconstrained perturbations, the total time in seconds, average time in seconds, and the average number of DE iterations are presented. Five runs for constrained and unconstrained one-pixel adversarial perturbations show that constrained perturbations converge faster.

Unconstrained			Constrained (ours)		
T (s)	AVG T (s)	AVG # itr	T (s)	AVG T (s)	AVG # itr
207.98	2.08	35.61	37.15	0.56	2.52
216.00	2.16	37.40	35.36	0.54	2.33
212.10	2.12	36.72	34.16	0.52	2.26
210.71	2.11	36.39	37.03	0.58	2.52
210.72	2.11	36.24	36.05	0.55	2.40

The results for the CIFAR-10 and NIPS-W datasets are summarized in Table 1 and Table 2 respectively. It can be clearly observed that across different runs and on average, the constrained one-pixel adversarial perturbations converged faster with comparatively less number of iterations when compared to conventional one-pixel adversarial perturbations. This effect is more pronounced for CIFAR-10, with the spatially constrained one-pixel adversarial perturbations converging $6\times$ faster than conventional attacks. Therefore, it can be clearly observed that by guiding the adversarial learning process for one-pixel adversarial perturbations towards areas of strong susceptibility based on insights gained by gradient-based interpretability, one can accelerate the optimization process to converge in a much more rapid manner.

Example images from ImageNet that were successfully attacked and their corresponding spatially constrained one-pixel adversarial perturbed images are shown in Fig. 2. It can be observed that the perturbed images, which resulted in successful attacks given that the tested network (SqueezeNet in this study) provided incorrect predictions compared to the original images, look very perceptually similar to the original images with the perturbed pixel visually hidden in most cases. The visual imperceptibility gained when leveraging the proposed spatially constrained one-pixel adversarial perturbation strategy stems from the fact that the adversarial learning process was guided towards sensitivity areas in the images that have complex details that conceal the perturbation well.

4. Conclusion

In this work, we presented a novel strategy in the realm of adversarial examples that leverage gradient-based interpretability, thus illustrate the use of such methods beyond purely the purpose of explainability. In particular, we introduce a spatially constrained one-pixel adversarial perturbation strategy that leverages identified sensitiv-

Table 2: Results of constrained one-pixel adversarial perturbations on the NIPS-W dataset. For both constrained and unconstrained successful attacks, the total time in seconds, average time in seconds, and average number of DE iterations are presented. Five runs for constrained and unconstrained one-pixel adversarial perturbations show that constrained perturbations converge faster.

Constrained (ours)			Unconstrained		
T (s)	AVG T (s)	AVG # itr	T (s)	AVG T (s)	AVG # itr
355.69	6.03	6.44	445.55	8.25	8.80
392.21	6.88	7.63	523.31	8.58	9.70
446.09	7.31	8.13	516.71	8.62	9.52
344.90	5.95	6.45	459.81	7.79	8.75
416.18	7.19	7.86	559.48	9.32	10.60

ity within an image based on gradient-based interpretability, thus leveraging insights gained to aid adversarial learning. Detailed experiments and evaluations were done to show that leveraging gradient-based interpretability can be used for faster convergence of one-pixel adversarial perturbations. We hope this work will allow for further development of gradient-based interpretability methods to be used for many more tasks and use cases beyond purely explainability.

Acknowledgment

This work was partially supported by the Natural Sciences and Engineering Research Council of Canada and the Canada Research Chairs Program. The authors would also like to thank Nvidia for donating GPUs used in the research.

References

- [1] S. Das and P. N. Suganthan. Differential evolution: A survey of the state-of-the-art. *IEEE Transactions on Evolutionary Computation*, 15(1):4–31, Feb 2011. 1
- [2] A. K. et al. Adversarial attacks and defences competition. *CoRR*, abs/1804.00097, 2018. 2
- [3] M. L. Ha, G. Franchi, M. Moller, A. Kolb, and V. Blanz. Segmentation and shape extraction from convolutional neural networks. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1509–1518, March 2018. 1
- [4] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009. 2
- [5] J. Lee, E. Kim, S. Lee, J. Lee, and S. Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. *arXiv preprint arXiv:1902.10421*, 2019. 1
- [6] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 618–626, Oct 2017. 1

- [7] D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017. 1
- [8] D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017. 2
- [9] J. Su, D. V. Vargas, and K. Sakurai. One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation*, 2019. 1, 3
- [10] B. Zhou, A. Khosla, L. A., A. Oliva, and A. Torralba. Learning Deep Features for Discriminative Localization. *CVPR*, 2016. 1