

# Adversarial Attacks Beyond the Image Space (Supplementary Material)

Xiaohui Zeng<sup>1</sup>, Chenxi Liu<sup>2</sup> (✉), Yu-Siang Wang<sup>3</sup>, Weichao Qiu<sup>2</sup>,  
Lingxi Xie<sup>2,4</sup>, Yu-Wing Tai<sup>5</sup>, Chi-Keung Tang<sup>6</sup>, Alan L. Yuille<sup>2</sup>

<sup>1</sup>University of Toronto    <sup>2</sup>The Johns Hopkins University    <sup>3</sup>National Taiwan University

<sup>4</sup>Huawei Noah's Ark Lab    <sup>5</sup>Tencent YouTu    <sup>6</sup>Hong Kong University of Science and Technology

xiaohui@cs.toronto.edu    cxliu@jhu.edu    b03202047@ntu.edu.tw

{qiuwch, 198808xc, yuwing, alan.l.yuille}@gmail.com    cktang@cs.ust.hk

## A. Attack Curves with Different (Differentiable or Non-Differentiable) Renderers

In Figure 1, we plot how the average loss function value (probability of the original class, after softmax) changes with respect to the number of attack iterations. Image-space attacks often succeed very quickly, whereas physical-space attacks are much slower yet more difficult, especially for the factors of illumination and material.

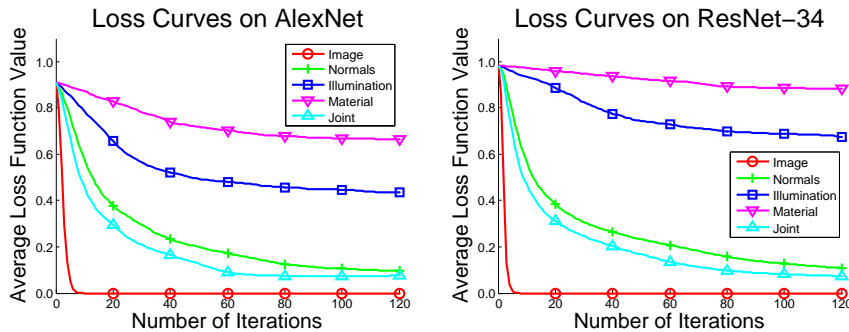


Figure 1. Attack curves for 3D object classification with a differentiable renderer.

In Figure 2, we plot how the average log-probability advantage (red) and image-space Euclidean distance (blue) change with respect to the number of attack iterations. An average log-probability advantage of 0 means that all images have been attacked successfully. Physical-space attacks are much more difficult to succeed and also require a much larger perceptibility.




Figure 2. Attack curves for 3D object classification and visual question answering with a non-differentiable renderer.

From these curves, we can conclude that physical-space attacks especially adding factors with clear physical meanings are much more difficult. This is arguably because most of these attacks impact the values of more than one pixels in the image space, which raises higher difficulties to the optimizers (*e.g.*, gradient-descent-based). We should also note that, with a more powerful optimizer, it is possible to find more adversarial examples in the physical world.