

# BEYOND PIXEL NORM-BALLS: PARAMETRIC ADVERSARIES USING AN ANALYTICALLY DIFFERENTIABLE RENDERER

HSUEH-TI DEREK LIU<sup>1</sup>, MICHAEL TAO<sup>1</sup>, CHUN-LIANG LI<sup>2</sup>, DEREK NOWROUZEZAHRAI<sup>3</sup>, ALEC JACOBSON<sup>1</sup>

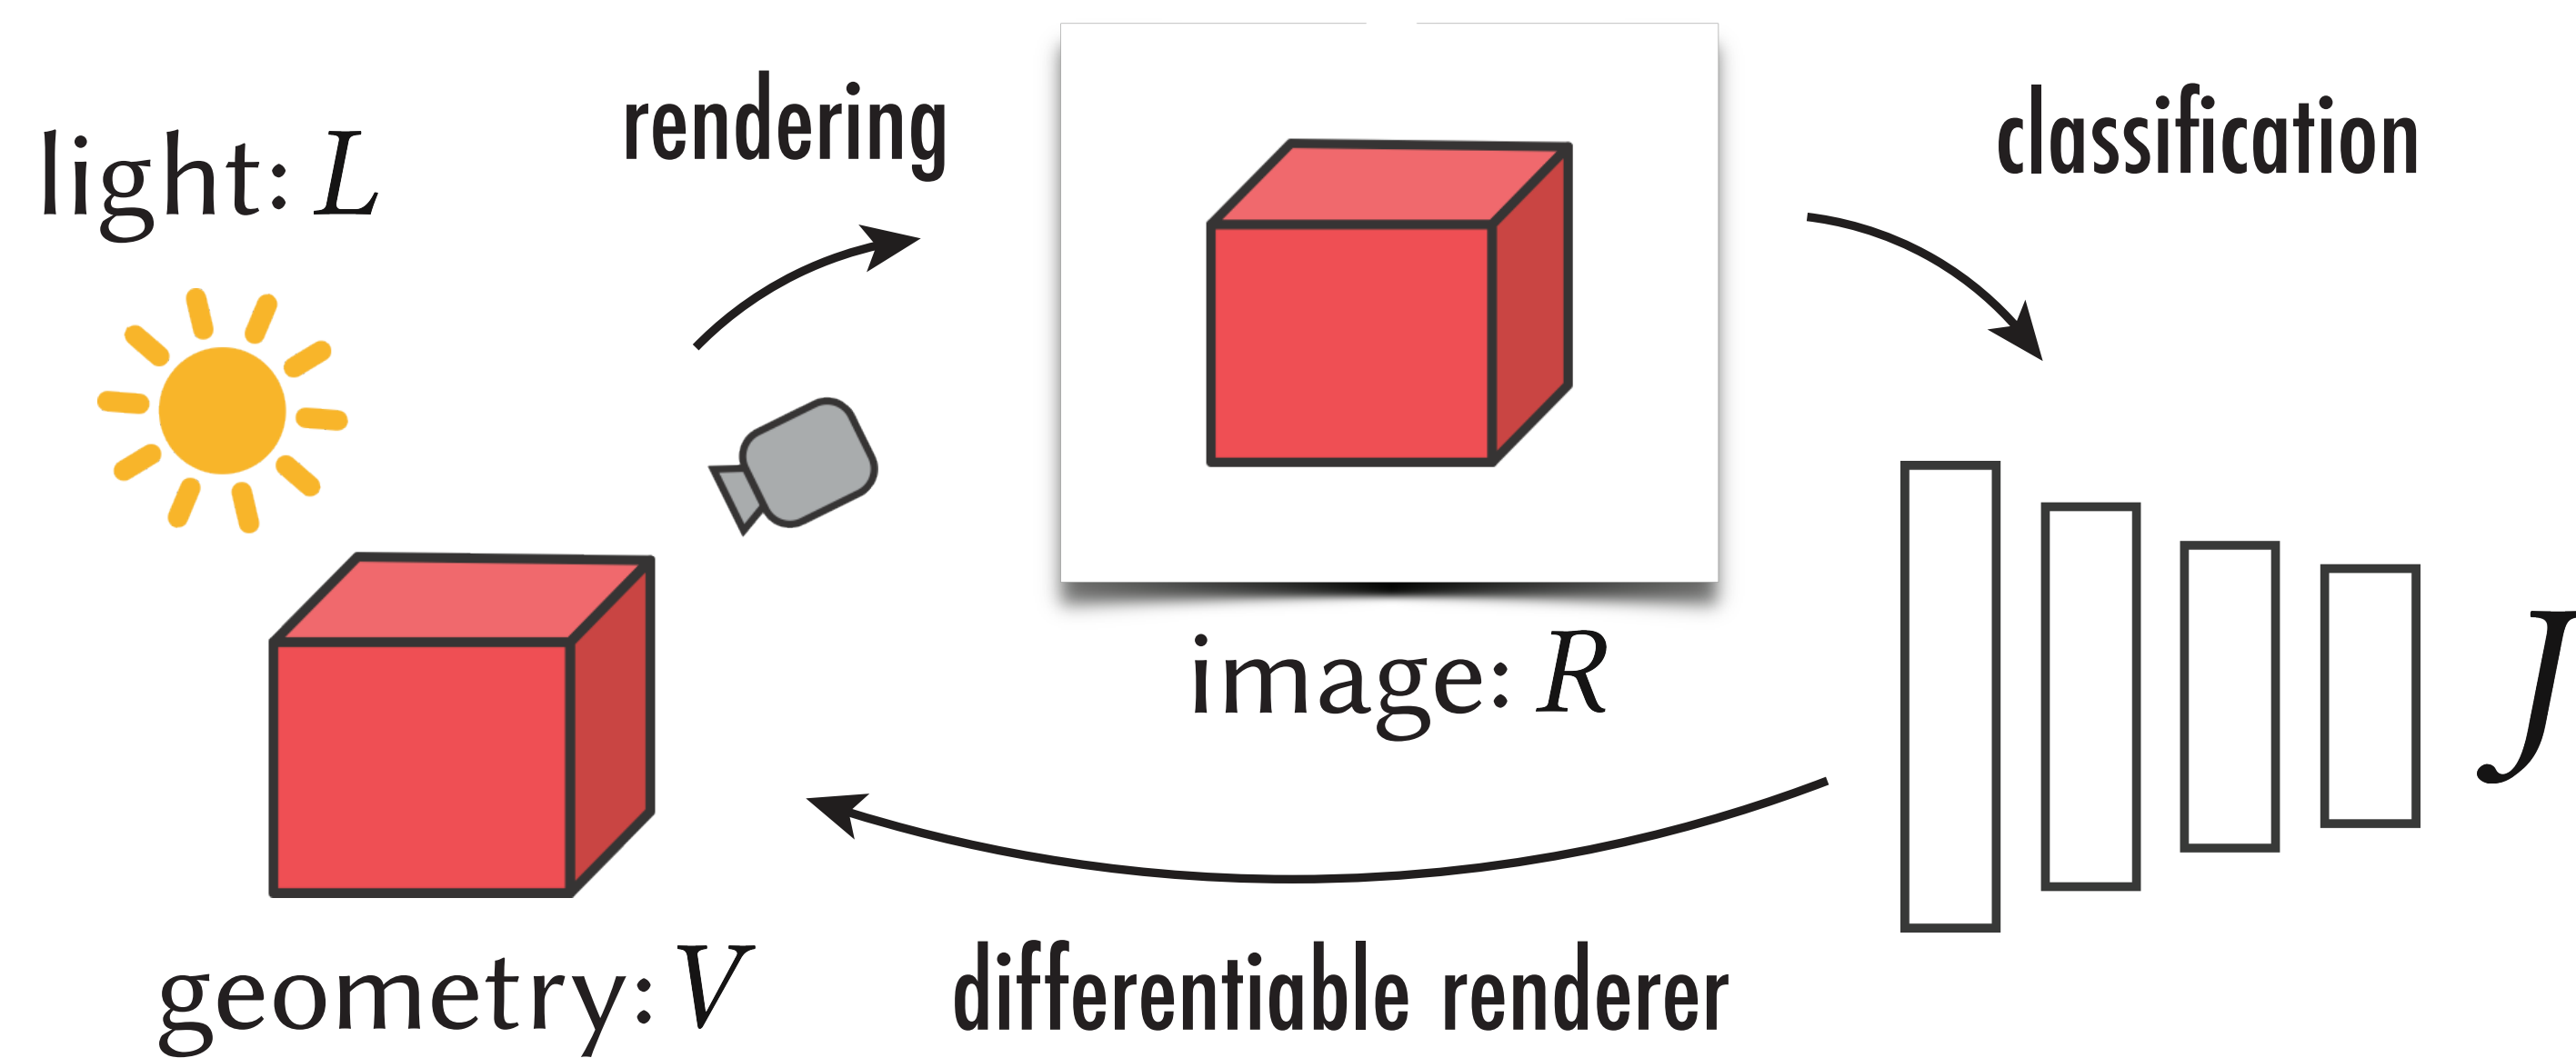
hsuehtil@cs.toronto.edu  
http://www.dgp.toronto.edu/~hsuehtil/

## Abstract

We compute adversarial examples by perturbing physical parameters instead of pixel colors. We present (1) *adversarial geometry* by 3D shape perturbations, and (2) *adversarial lighting* by scene lighting perturbations.

## Method

We use a gradient-based optimization, where the gradients of the cost function  $J$  w.r.t 3D scene parameters  $V$ ,  $L$  are computed via a physically-based *analytically* differentiable renderer.



Geometry derivatives


$$\frac{\partial J}{\partial V} = \frac{\partial J}{\partial R} \frac{\partial R}{\partial N} \frac{\partial N}{\partial V}$$

Lighting derivatives


$$\frac{\partial J}{\partial L} = \frac{\partial J}{\partial R} \frac{\partial R}{\partial L}$$

## Lighting parameterization

Lighting as a spherical function




We parameterize lighting with spherical harmonics





## Geometry parameterization

Triangle meshes




## Adversarial lighting





## Rendered adversarial training


- standard data augmentation: 40.4%
- adversarial data augmentation: 65.8%



## Adversarial geometry



## Deep geometric illusion



## Future Work

- simulation/rendering adversarial training
- differential renderers for real images
- incorporating real-time rendering techniques
- more parametric perturbations (e.g., adversarial human poses)