Learning Real-Time Ambient Occlusion from Distance Representations

Authors: B. Keinert, Jana Martschinke, M. Stamminger

Presenter: 20160249 박승건

Ambient Occlusion (AO)

Ambient Light

- One of the essential lighting elements for realistic illumination
- Objects in real world never become completely dark, as long as there's any light scattering around

https://learnopengl.com/Lighting/Basic-Lighting

Ambient Occlusion (AO)

- Used to calculate how exposed each point is to ambient lighting
- With AO, we have a sense of the depths of the surfaces
  - Exposed, outer surface: brighter
  - Deeper, inner surface: darker

AO off

AO on

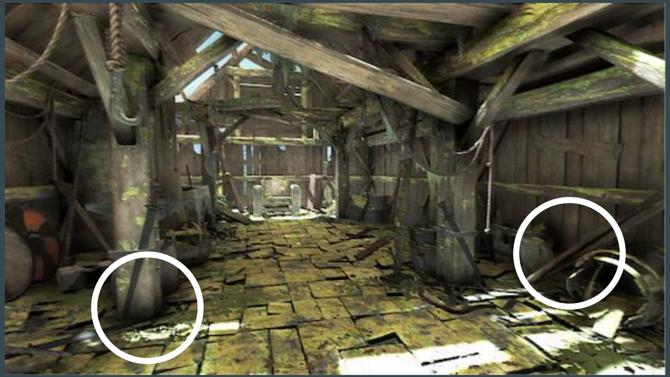

https://docs.unity3d.com/kr/2018.3/Manual/PostProcessing-AmbientOcclusion.html

Basic Principle of AO

- Construct a hemisphere of rays, then check them for intersections with other objects
- The more unblocked rays there are, the more brightness the point gains
- **Requires the information about surrounding geometry, which makes raycast-based real-time AO solutions quite expensive**

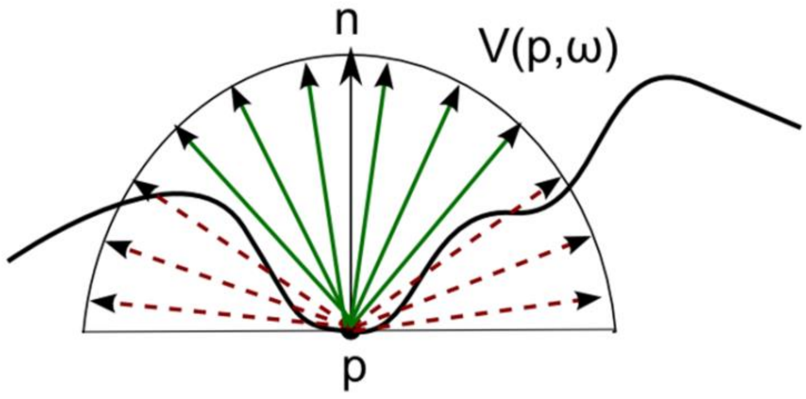

$$AO(\mathbf{p},\mathbf{n}) = \frac{1}{\pi}\int_{\Omega} V(\mathbf{p},\omega)\,\mathbf{n}\cdot\omega\,d\omega,$$

https://nbertoa.files.wordpress.com/2017/02/ao_integral.png?w=756

Classic AO Techniques

- SSAO (Screen Space AO)
  - Mittring [2007]
  - Reads z-buffer to calculate how much each pixel's depth value is different from other pixels' around
  - The more samples it finds inside geometry, the more occluded the point is
  - Fast but sometimes inaccurate, since it does not consider the geometry

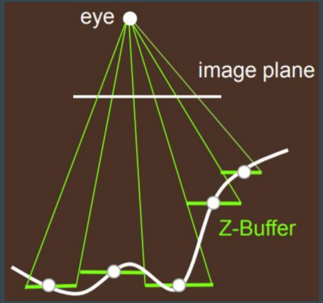

https://developer.download.nvidia.com/presentations/2008/SIGGRAPH/HBAO_SIG08b.pdf

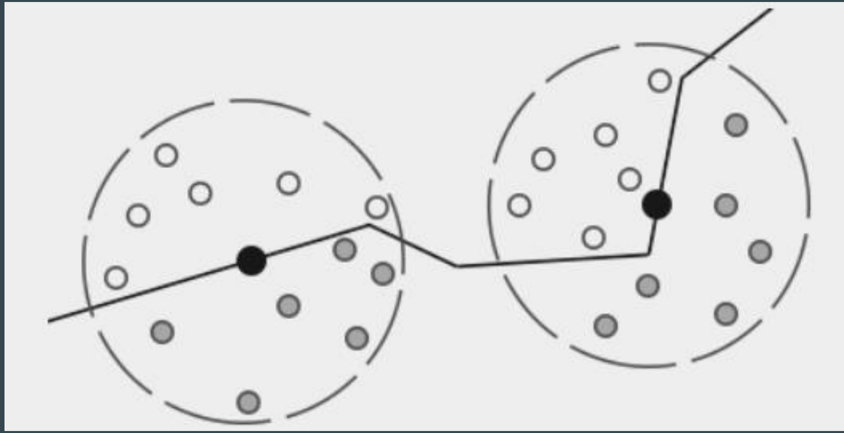

https://learnopengl.com/Advanced-Lighting/SSAO

Classic AO Techniques

- HBAO (Horizon Based AO)
  - Bavoil et al. [2008]
  - Improved SSAO
  - Raymarch in several directions and keep updating the maximum elevation
  - Then compare the elevation angle with an angle between tangent and horizon
  - May still yield artifacts (false occlusions)

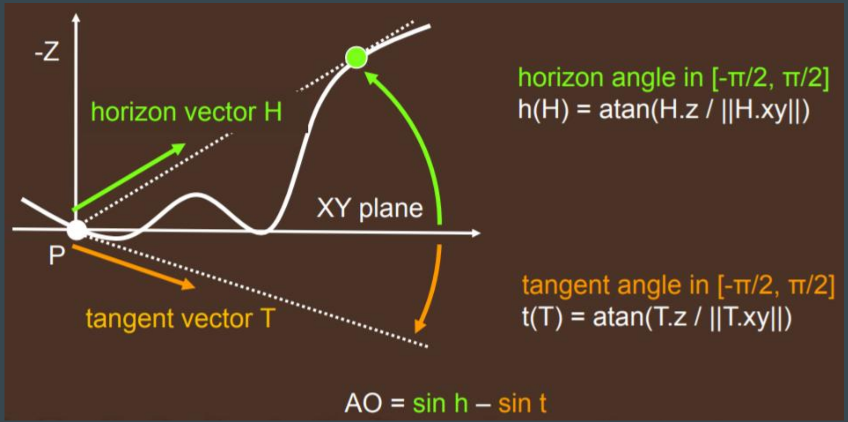

https://developer.download.nvidia.com/presentations/2008/SIGGRAPH/HBAO_SIG08b.pdf

Previous Works

Previous Works on Two Tracks

- **Machine learning** applied to real-time rendering
  - Ren et al. [2013]
- Convolutional neural network (CNN) applied to ambient occlusion
  - Holden et al. [2016]

- AO recovered from **signed distance fields**
  - Evans [2006]
  - Wright [2015] upgraded this approach

TRACK 1: Neural Network Ambient Occlusion (Holden et al. [2016])

- SSAO often creates ‘artifacts’ (error)
- Main idea: Learn an SSAO algorithm to minimize its errors
- Input: screen-space g-buffer data
  - g-buffer: contains geometry information such as the world normal, base color, metallic/specular/roughness, etc.

TRACK 1: Neural Network Ambient Occlusion (Holden et al. [2016])

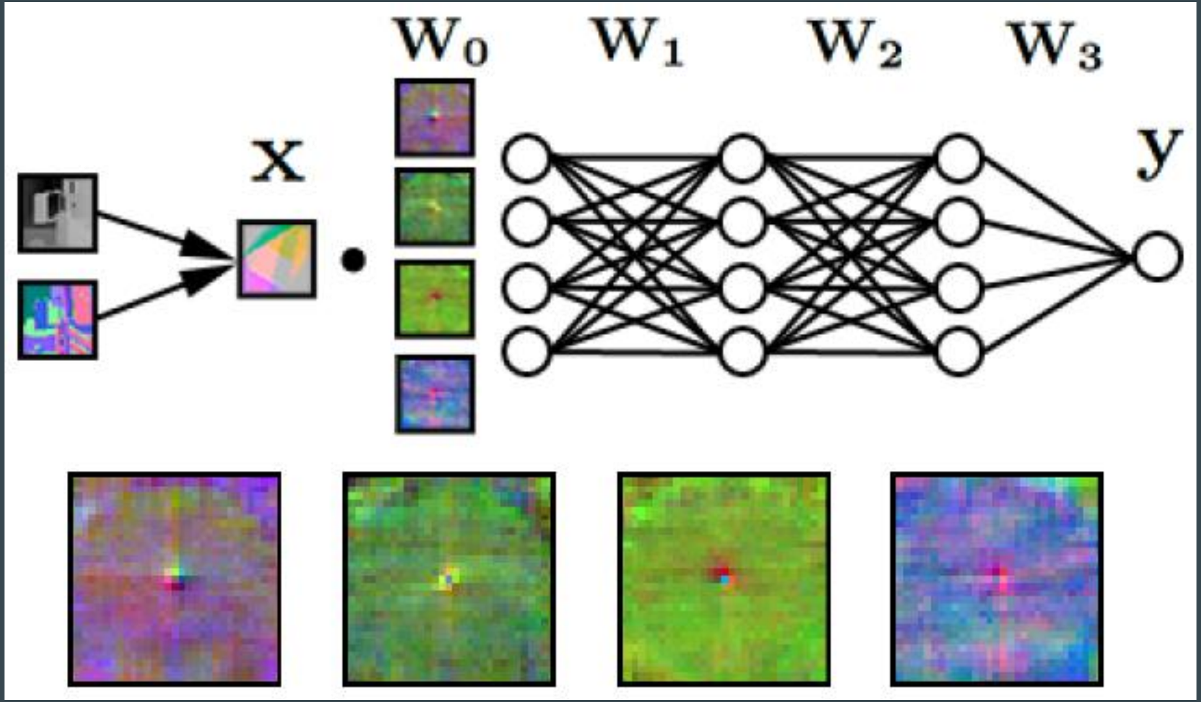

TRACK 2: Fast Approximations for Global Illumination on Dynamic Scenes (Evans [2006])

- AO in dynamic scenes is expensive
- Evans’ approach: Signed Distance Field (SDF)
  - Enormously inaccurate but looks aesthetically pleasing
  - Requires very little pre-computation → works well in real-time

TRACK 2: Fast Approximations for Global Illumination on Dynamic Scenes (Evans [2006])

- Signed Distance Field (SDF)
  - 3D scalar field describing the distance to the closest surface point at each 3D point
  - Each value equals the radius of the largest sphere around this point, which does not intersect any geometry
- With SDF, compute the visibility of the sky from any point in the scene
  - Here, unlike the image on the right, points within objects are given positive values

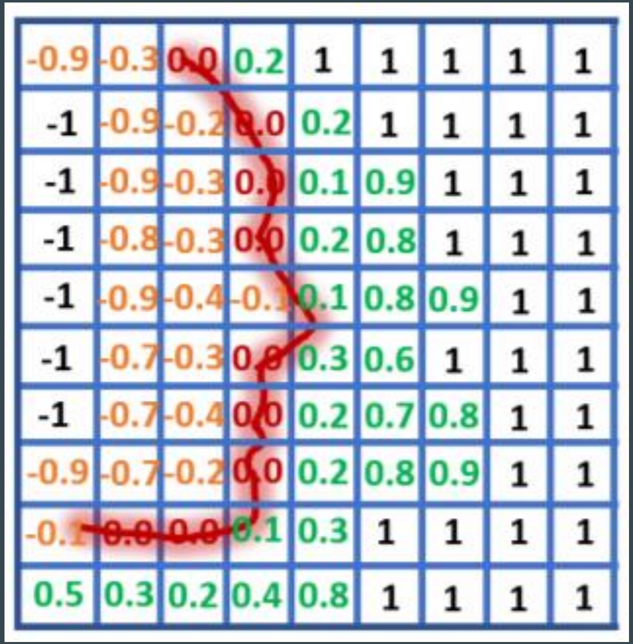

TRACK 2: Fast Approximations for Global Illumination on Dynamic Scenes (Evans [2006])

SDF values expressed in colors

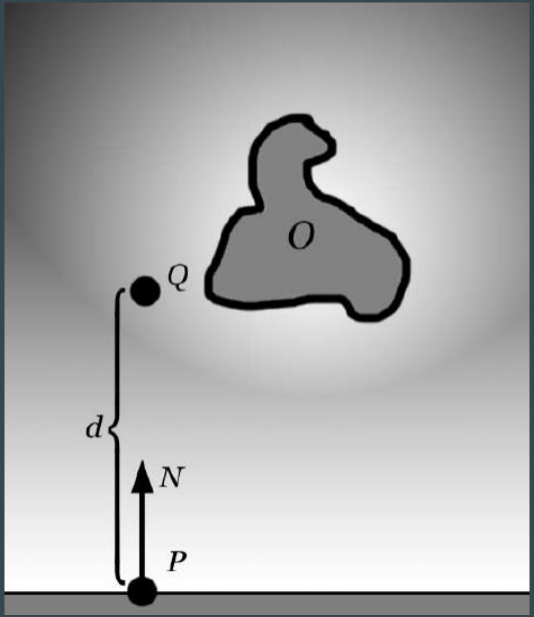

TRACK 2: Fast Approximations for Global Illumination on Dynamic Scenes (Evans [2006])

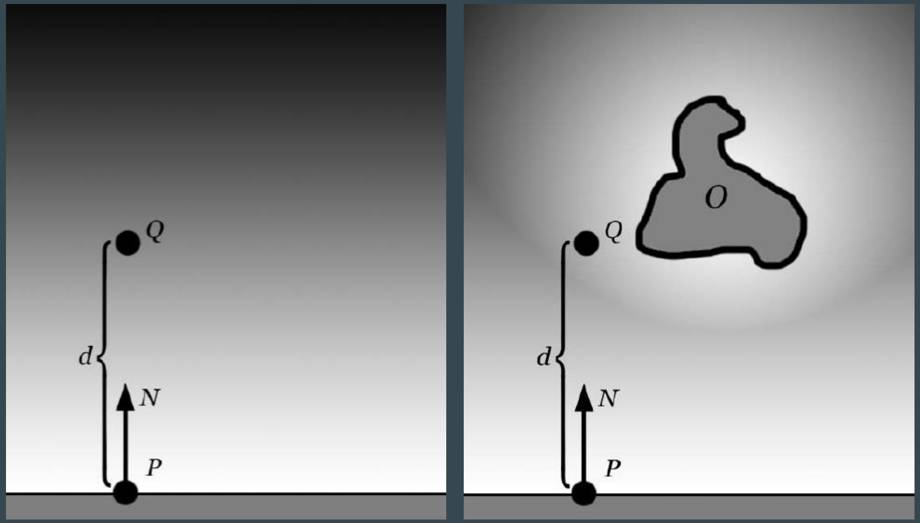

NO OCCLUSION
SDF value decreases
as we go away from the surface

WITH OCCLUSION
SDF value may not decrease
as we go away from the surface

TRACK 2: Fast Approximations for Global Illumination on Dynamic Scenes (Evans [2006])

Expected SDF values without any occluder

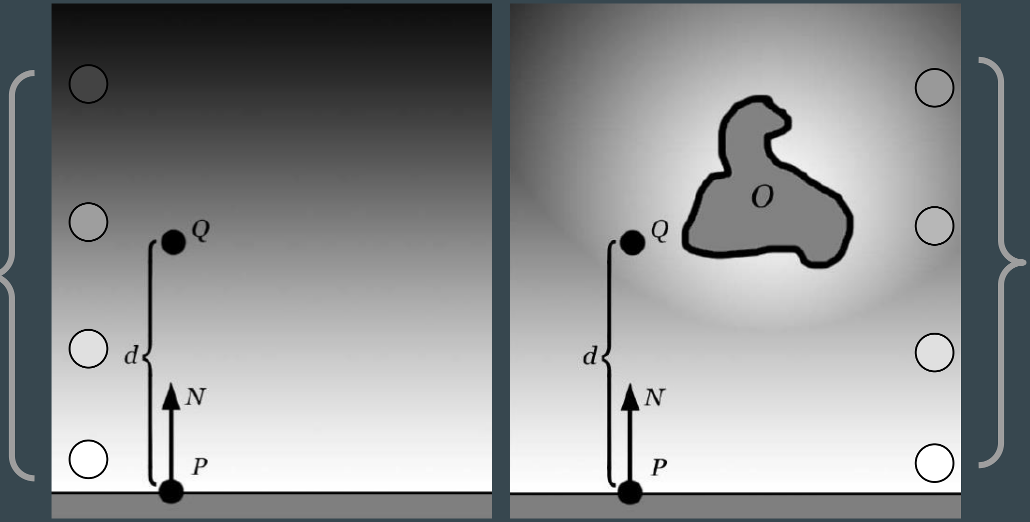

Actual SDF values with occluder

Expected occlusion without any occluder

Actual occlusion with occluder (negated)

$$\sum_{i=0}^{n} SDF(P + d_i N) = -\sum_{i=0}^{n} d_i$$

$$C = e^{k \sum SDF(P + d_i N) + d_i}$$

Paper Introduction

Key Points

- Machine learning + signed distance for AO
- Ground truth: integral over a hemisphere (recap: basic principle)
- Input: samples of signed distance representations

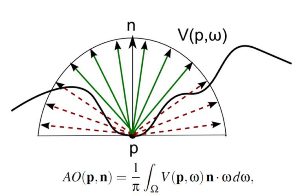

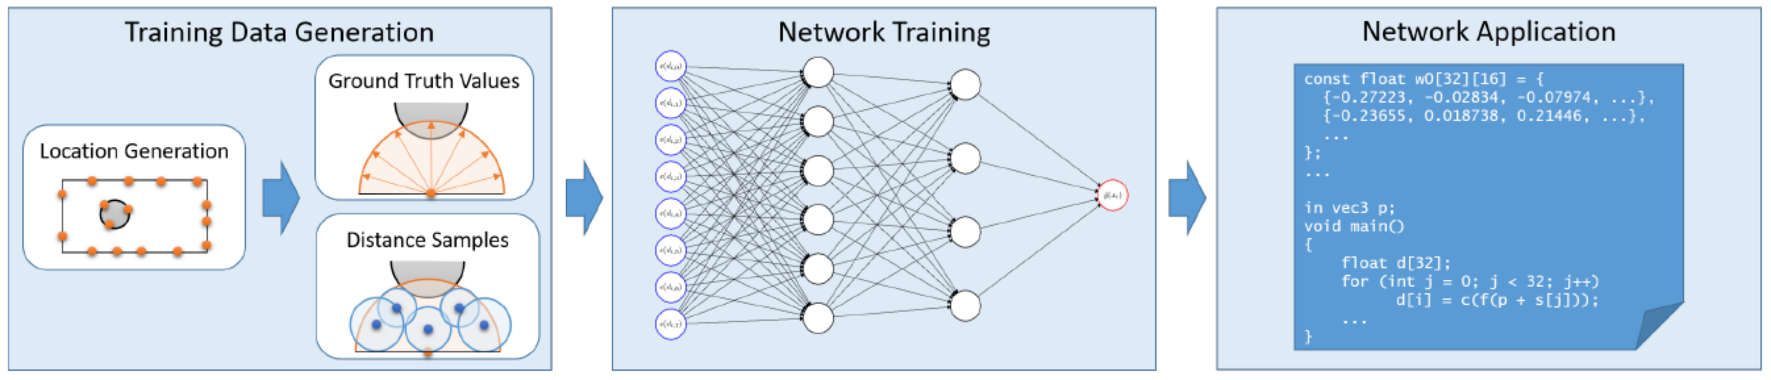

AO Approximation at Render Time

- Result: neural network is trained to approximate occlusion from signed distance samples
- The trained network is applied in the following order:
    1. For every point, set the surrounding hemisphere and generate sample positions inside it
    2. Calculate the signed distances of samples
    3. Estimate how much the point is occluded (how much SDF values are different from expected)
    4. Shade
- Object-wise signed distance field can be pre-computed and united even with fully animated meshes, without sacrificing real-time performance

Result

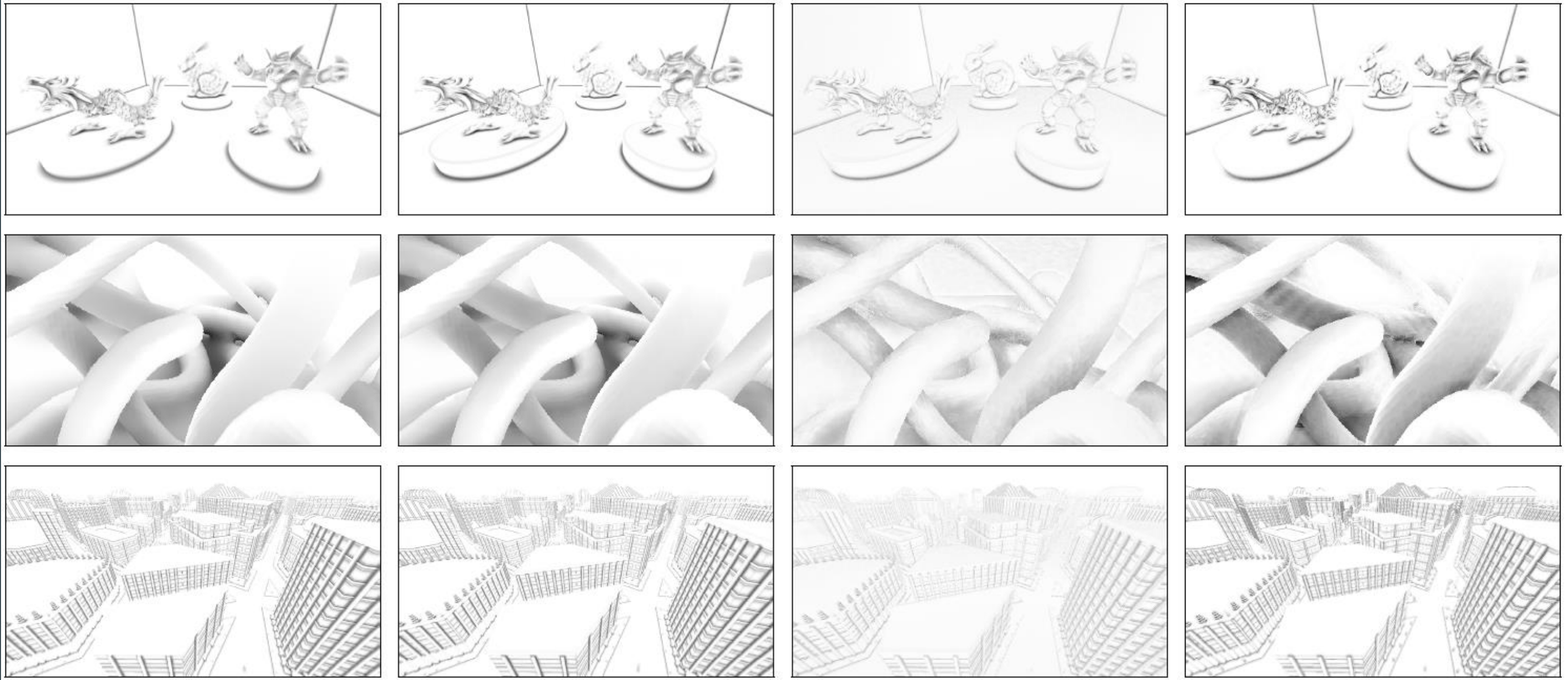

(a) Ground truth (b) Our method ($m = 32$) (c) NNAO (d) HBAO

Further Applications

- Bent Normals
  - Average direction of all rays unobstructed by surrounding geometry
  - Can be used for global illumination calculations
- Soft shadows
  - In contrast to AO, soft shadows depend on information about the whole scene
  - It is hard to directly apply the framework on soft shadows, but in principle it can reproduce plausible soft shadows

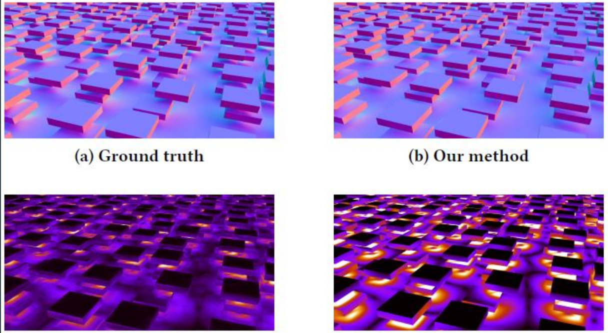

(a) Ground truth (b) Our method

(c) Error (our method) (d) Error (pure normals)

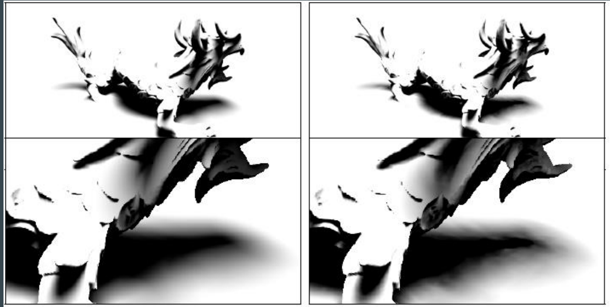

(a) Ground truth (b) Our method, trained on the Stanford dragon

Limitations

- Signed distance representation required
  - Recent commercial game engines (i.e. Unreal Engine 4) provide such data
- Dependence on training data
  - The quality of AO approximation is highly dependent on the resemblance between training and application scene
- Ambiguities in input data
  - Different ground truth values can be boiled down to similar approximation

Limitations

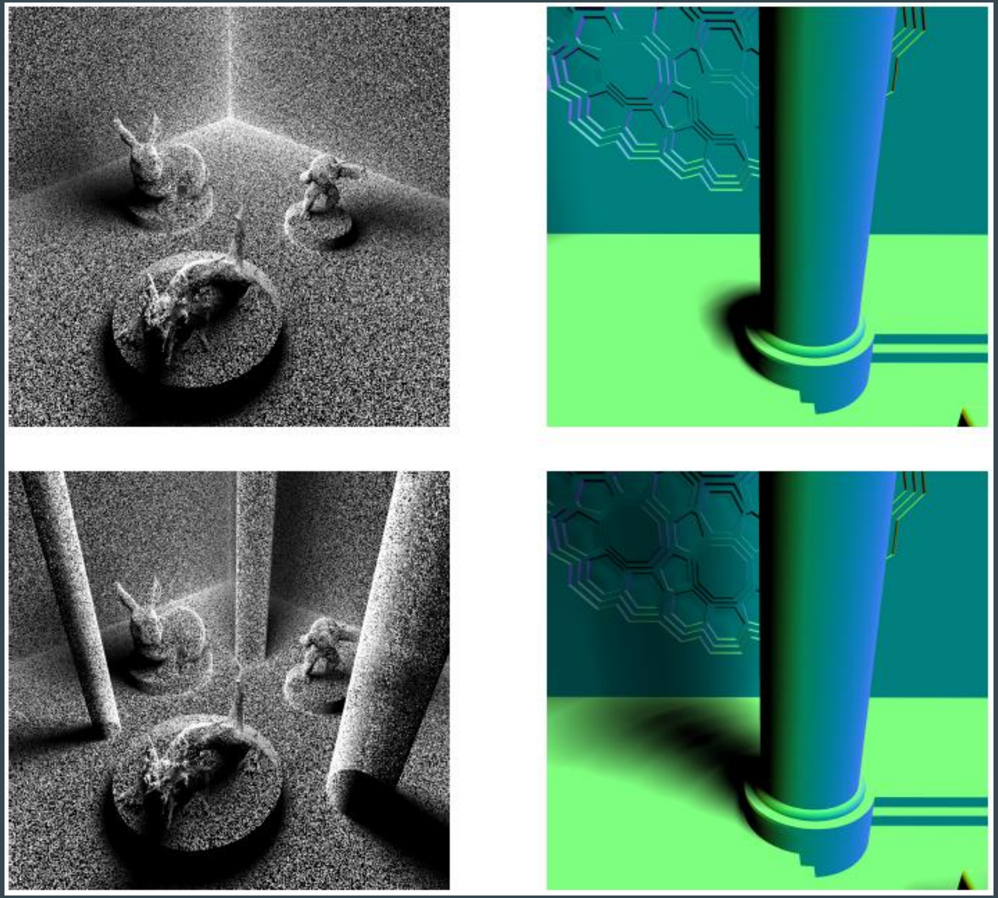

Limitations