SIGGRAPH ASIA 2016 MACAO

Sponsored by acm

Neural Network Ambient Occlusion

Daniel Holden*, Jun Saito†, Taku Komura*

*The University of Edinburgh, †Method Studios

SA2016.SIGGRAPH.ORG

CONFERENCE: 5 - 8 DECEMBER 2016 • EXHIBITION: 6 - 8 DECEMBER 2016 • THE VENETIAN MACAO, MACAO

Background

WHAT IS AMBIENT OCCLUSION?

- Approximation of soft shadows produced by global illumination.

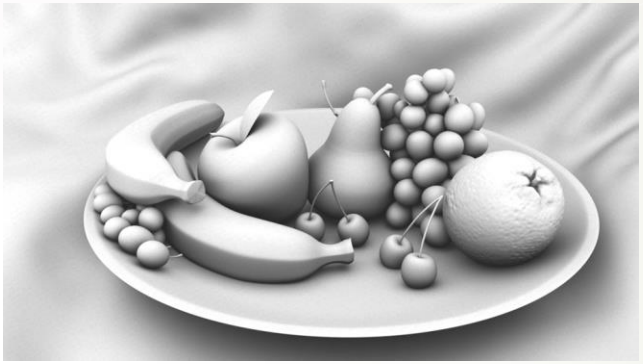

- Calculated via tracing rays from each point into the surrounding world.

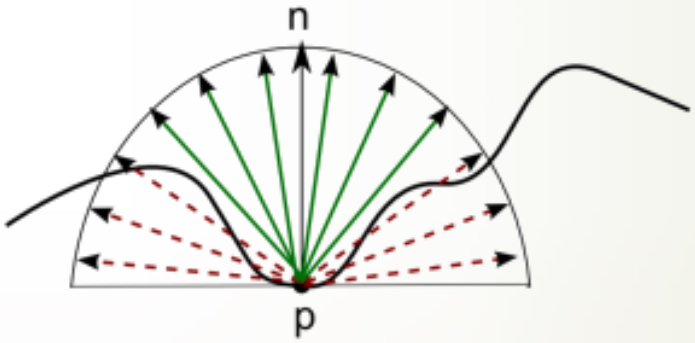

SCREEN SPACE AMBIENT OCCLUSION?

- Treat camera depth as heightmap representing the geometry.
- Calculate the occlusion taking into account that this is an approximation.

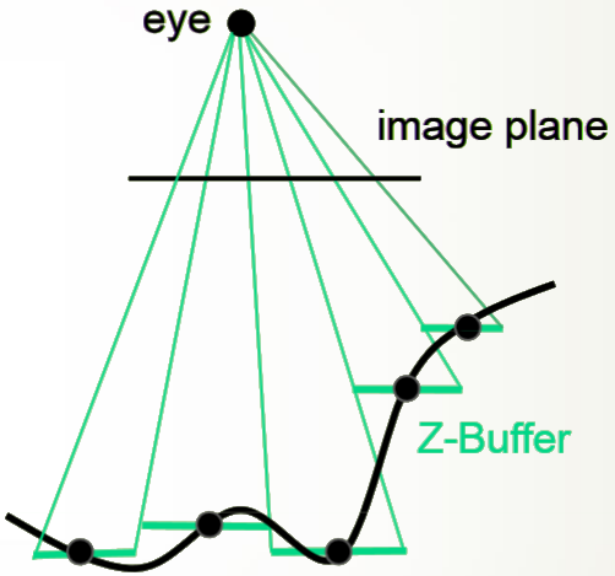

SCREEN SPACE AMBIENT OCCLUSION

SSAO [Mittring 2007]

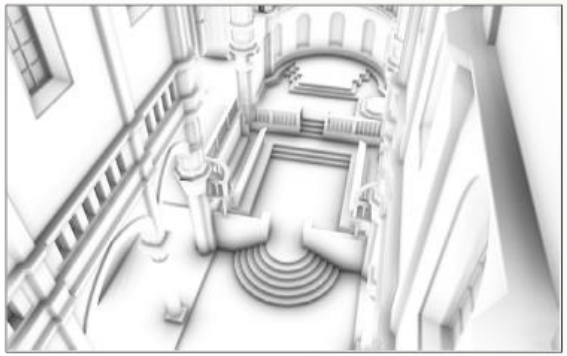

SSAO+ [McNaughton 2008]
HBAO [Bavoli et al. 2008]

SAO [McGuire et al. 2011]

GTAO [Jimenez et al. 2016]

OBJECTIVE

Learn a good Screen Space Ambient Occlusion function

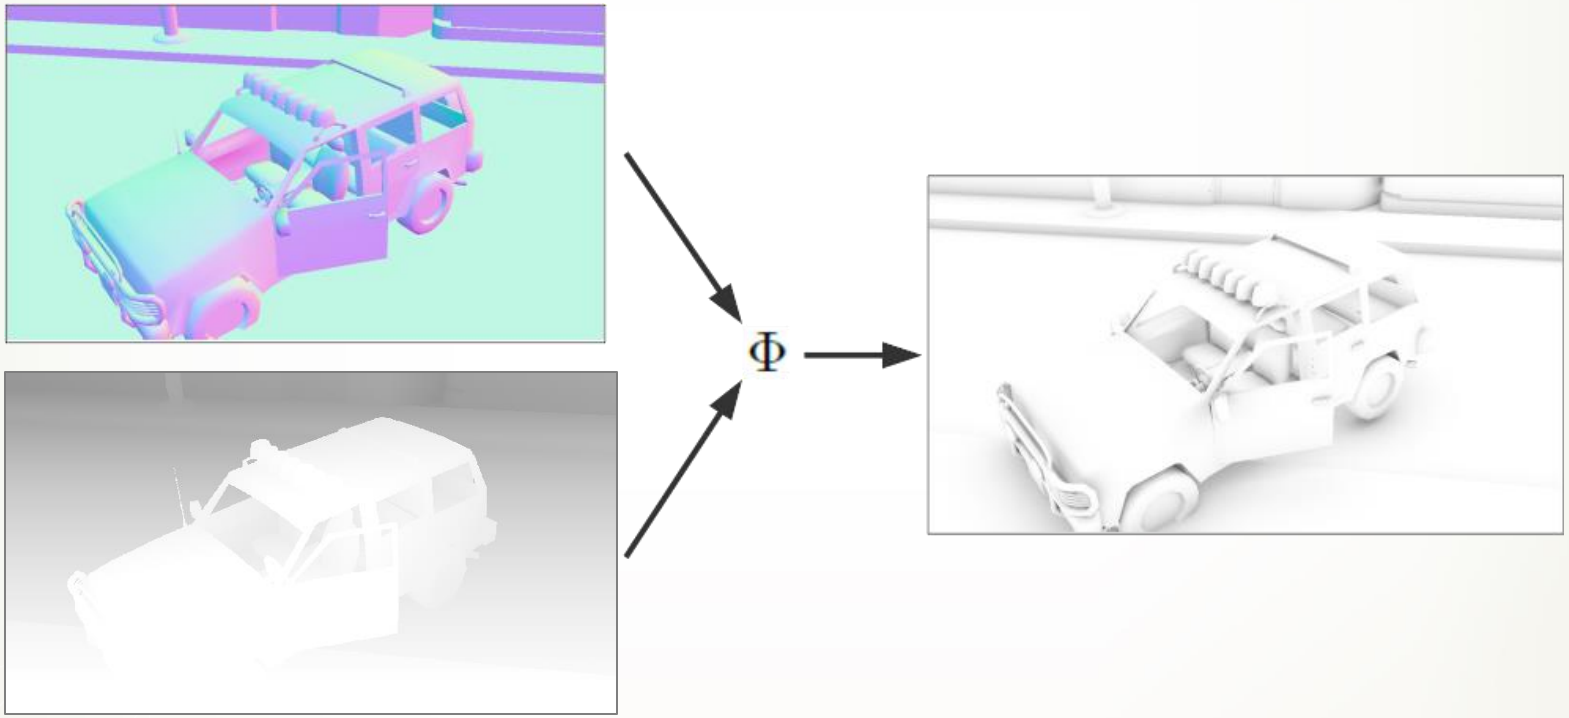

WHY MACHINE LEARNING?

- Machine Learning is good for things we don't know how to model.

- What about things we do know how to model?

WHY MACHINE LEARNING?

- **Memorisation**
  - Remembers parts of training data
  - Potential to be **faster**

- **Optimisation**
  - Trained with respect to actual data
  - Potential to be **more accurate**

Methodology

TRAINING DATA

- Open Source First Person Shooter *Black Mesa*
- Extract scenes and render offline using *Mental Ray*

TRAINING DATA

AO

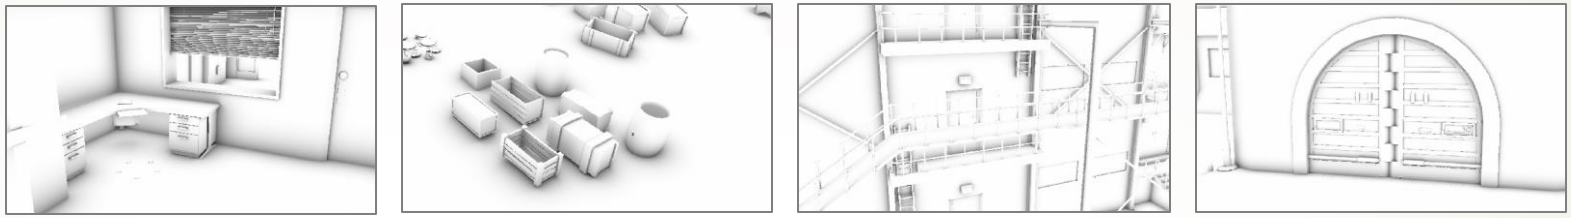

Depth

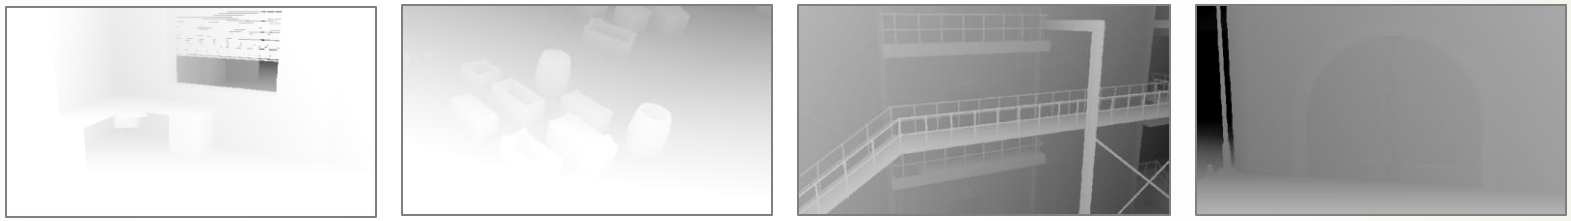

Normals

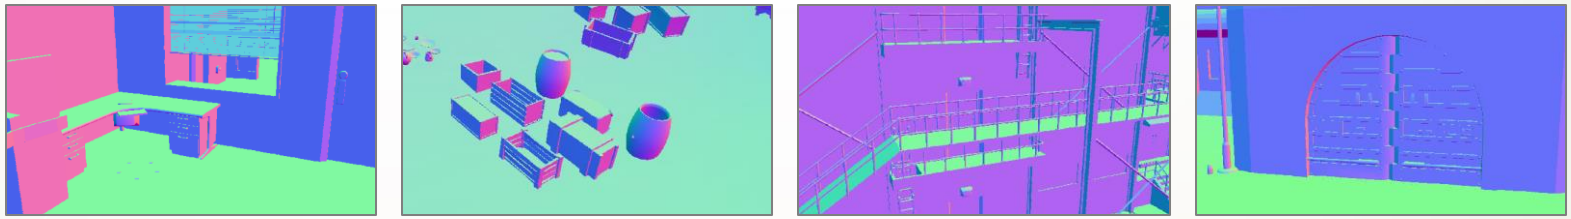

PRE-PROCESSING

- Take image patches in view space the size of AO radius.
- Compute the difference from the central pixel.
- Scale by the distance to the AO radius.

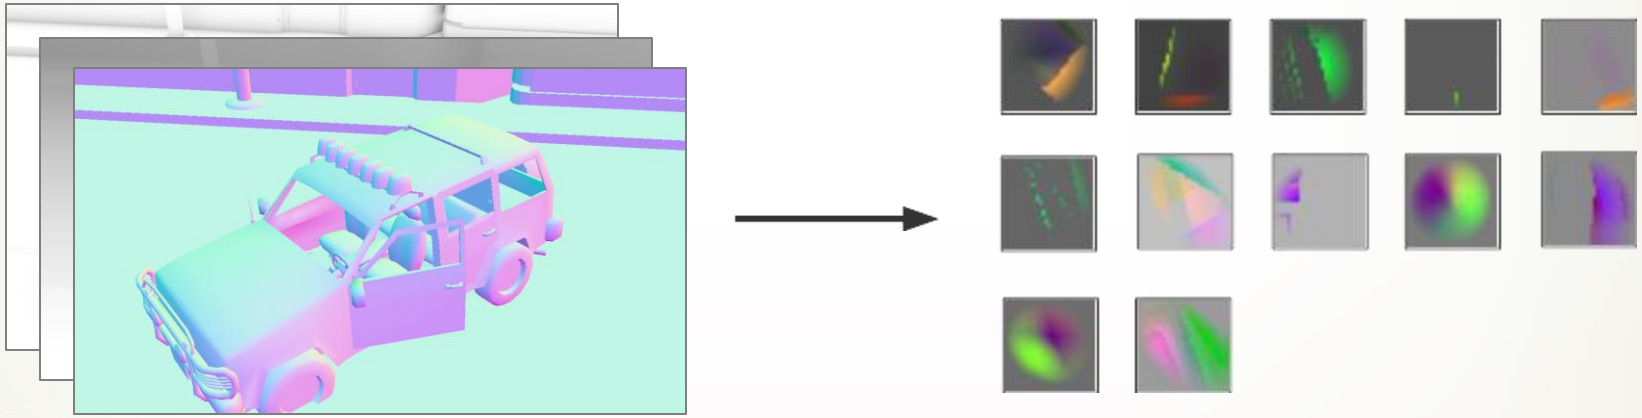

NETWORK STRUCTURE

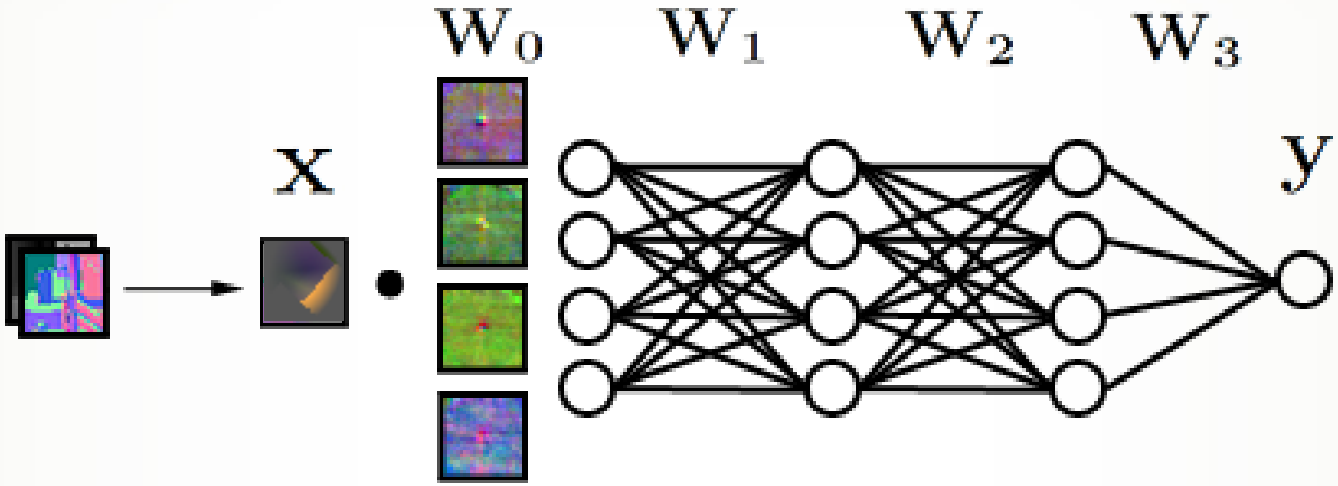

Overview: Simple four layer Neural Network

NETWORK STRUCTURE

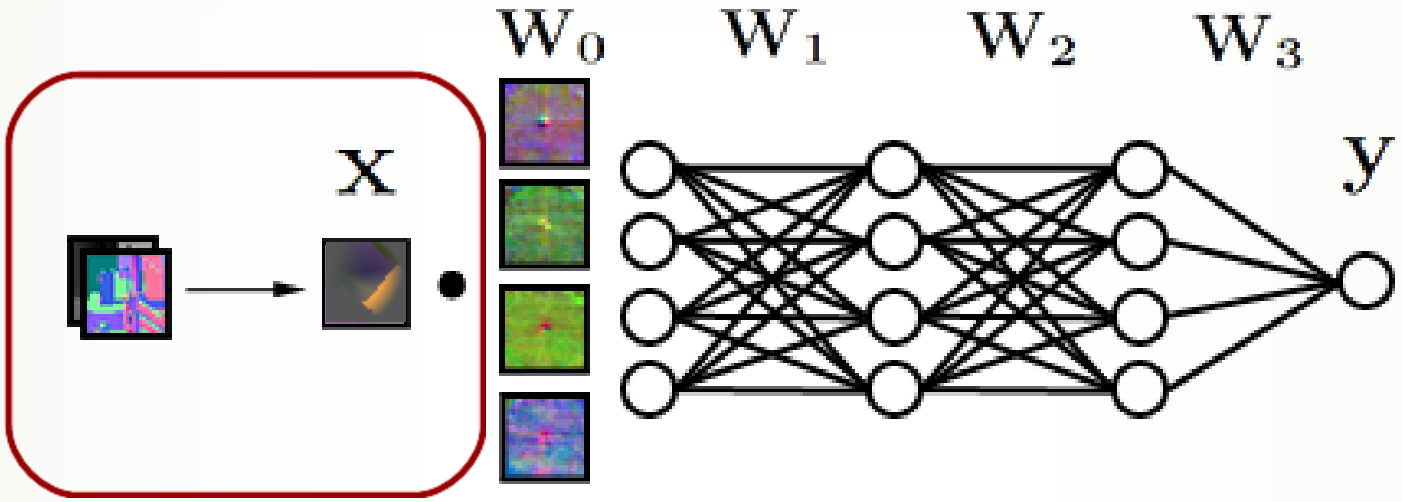

Input Layer: Performs pre-processing from previous slides

NETWORK STRUCTURE

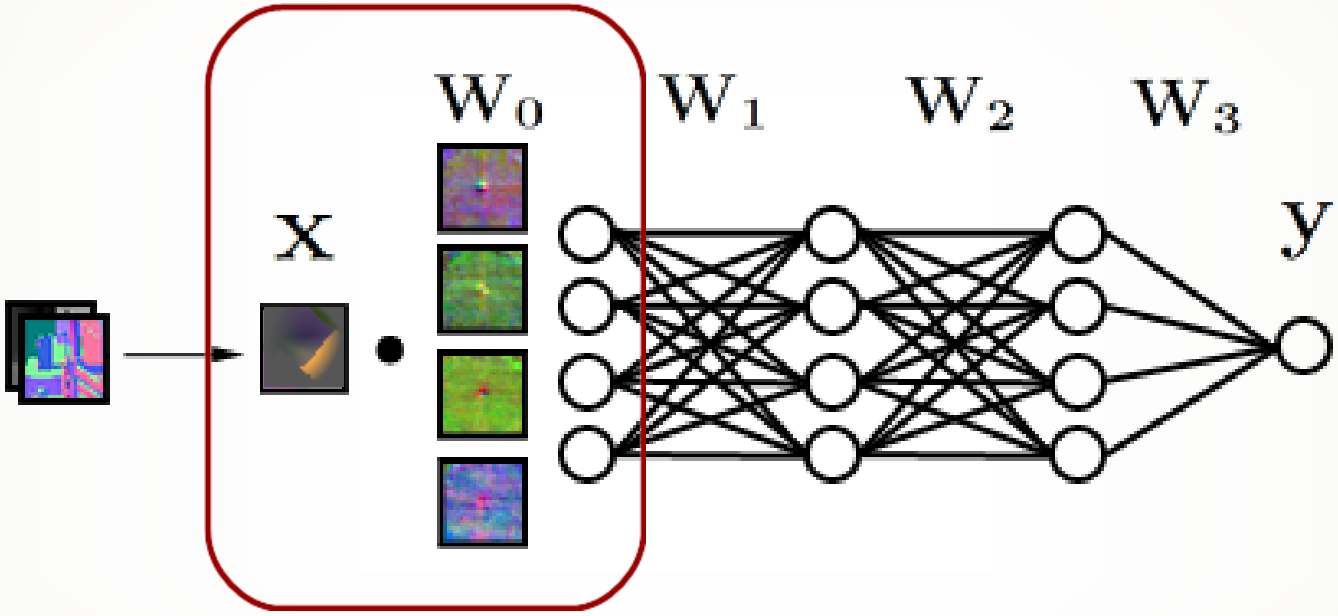

First Layer: Acts as convolution producing 4 values from patch

NETWORK STRUCTURE

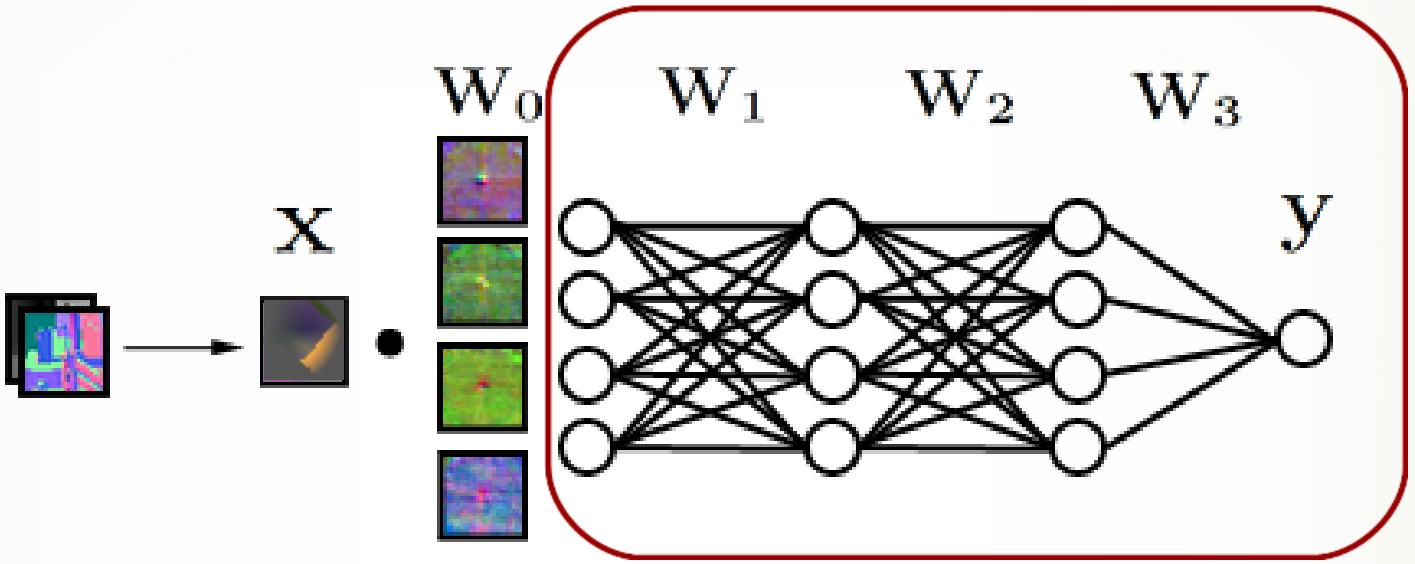

Other Layers: Normal layers acting on 4 hidden units each

DETAILS

- Implemented in *Theano*
- Trained with *Adam*
- Uses *Parametric Rectified Linear Units*
- *Dropout* of 0.5
- ~500000 data points
- ~10 hours training on *NVIDIA GTX 660*

Runtime

RUNTIME

- Neural Network is translated into shader
- Made easy by several design decisions:
  - No intermediate storage - runs in single pass
  - Only 4 hidden units - can use vectors ops
  - First layer is the only complicated layer
  - ~100 lines of code

FIRST LAYER

- Acts like a convolution with network weights as *filter* images.
- Multiplied by input patch and summed.
- This *integration* can be approximated by sub-sampling.

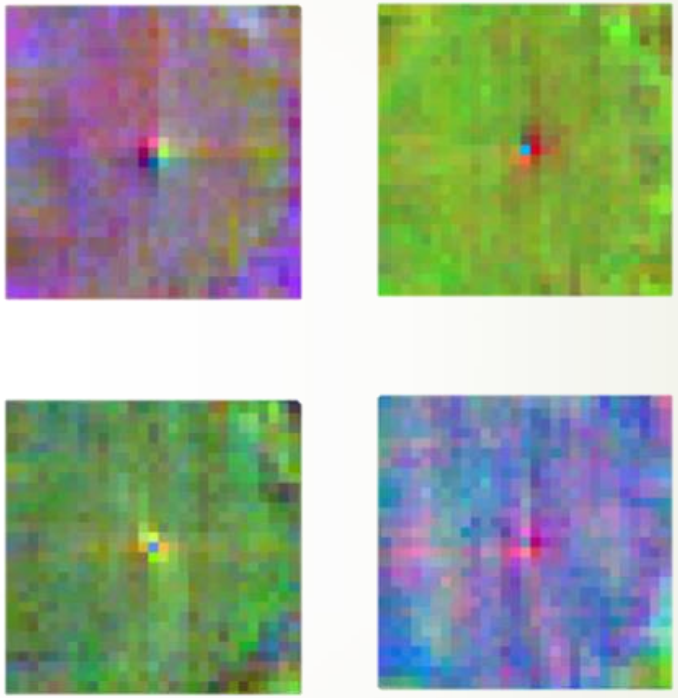

SUB-SAMPLING

- Mulitply just a few input pixels by filter.
- Rescale result using sub-sample ratio.
- Adjust sampling locations spatially across screen.

SUB-SAMPLING

Random

- Seeded
- Many Samples
- Bad for Cache
- Unbiased
- [SSAO,SSAO+]

Stratified

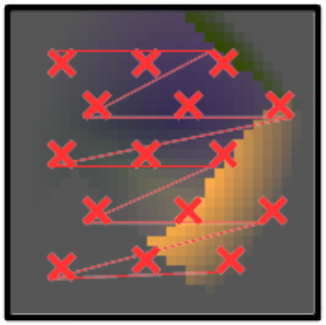

- Jittered
- Many Samples
- Good for Cache
- Biased

Star

- Rotated
- Few Samples
- Good for Cache
- Biased
- [HBAO]

Spiral

- Rotated/Offset
- Few Samples
- Good for Cache
- Unbiased
- [SAO]

Results

GROUND TRUTH

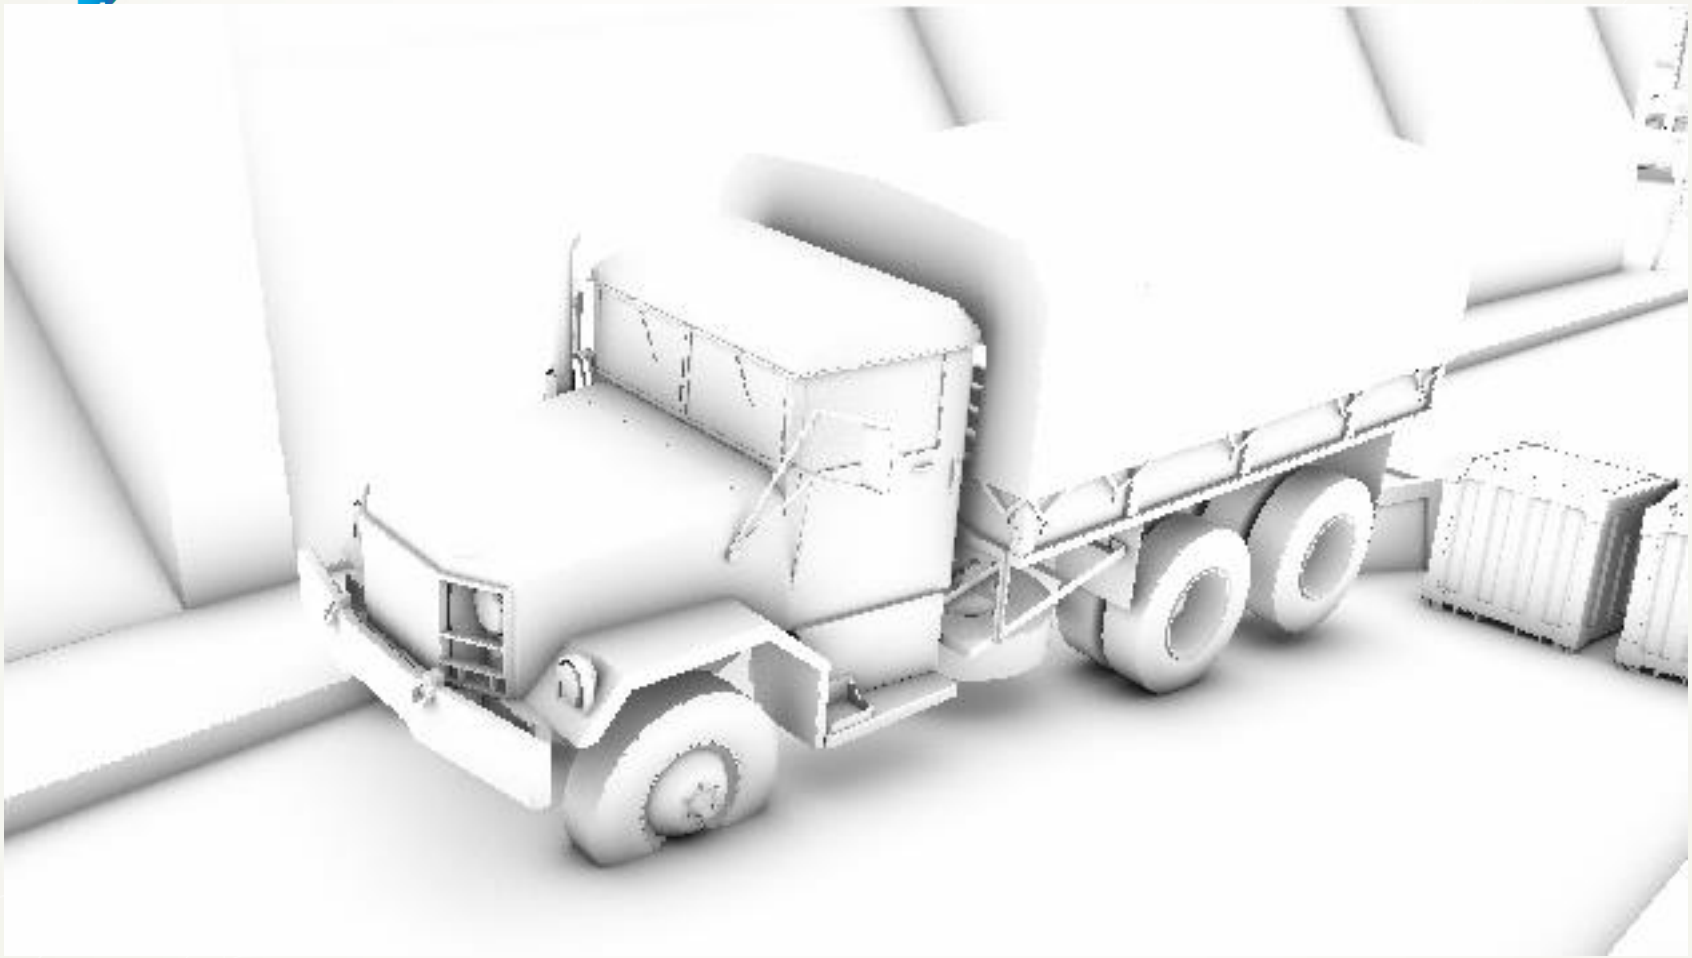

SSAO+

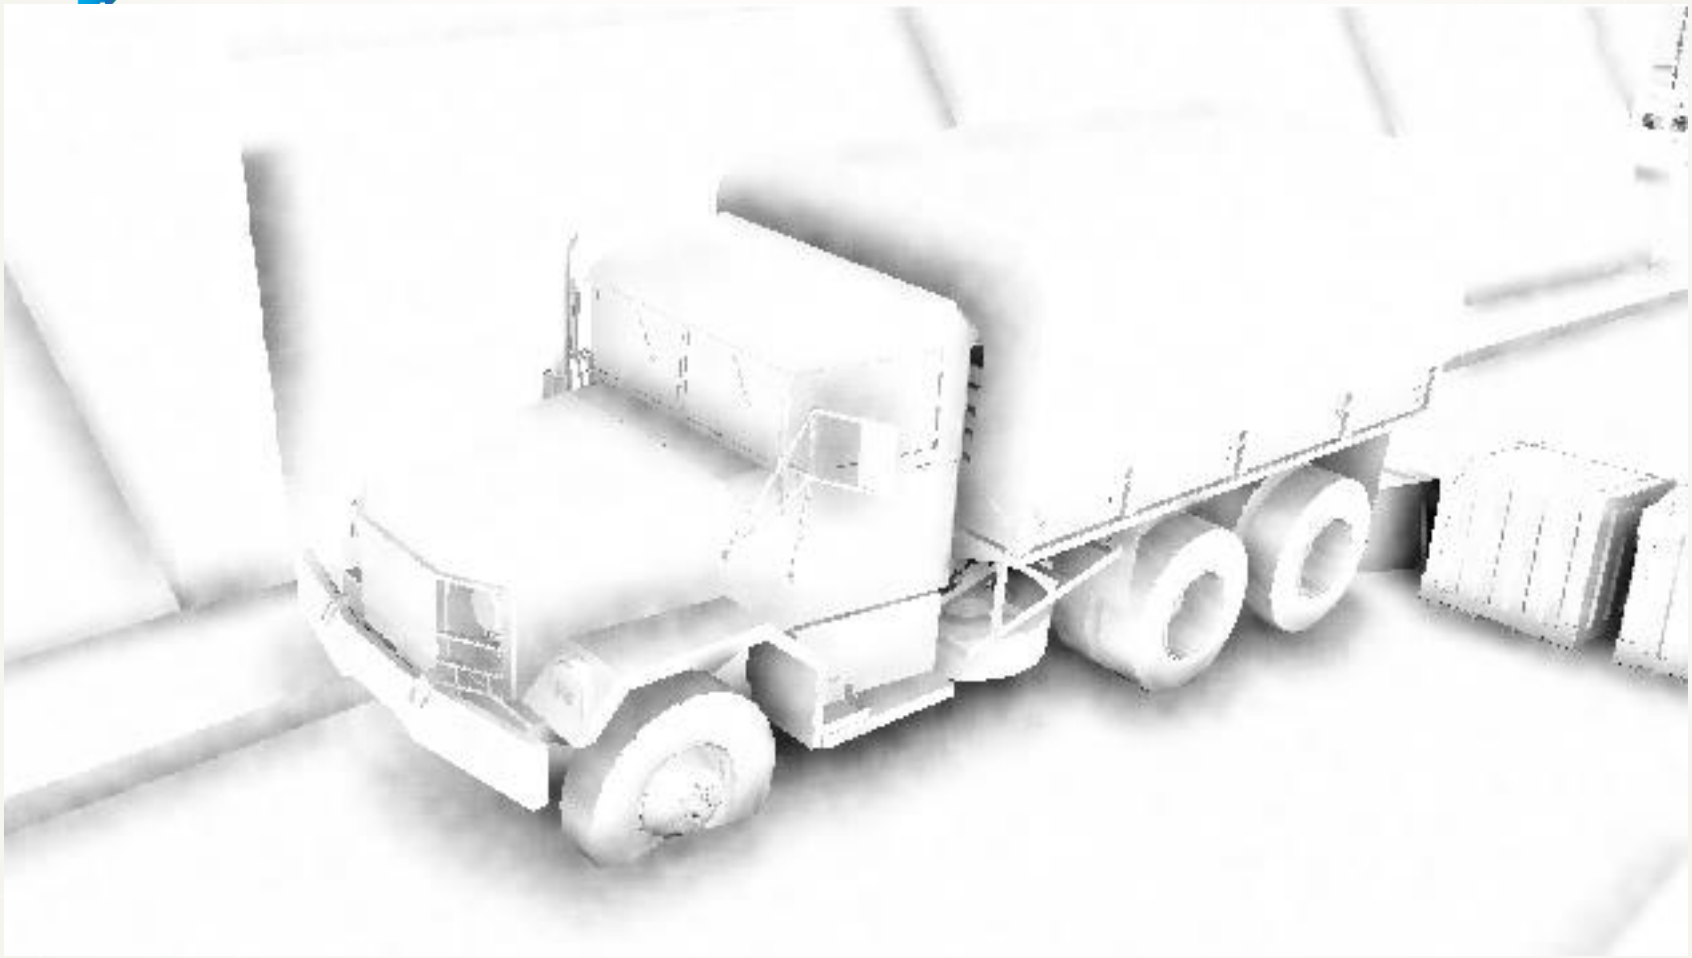

SSAO+

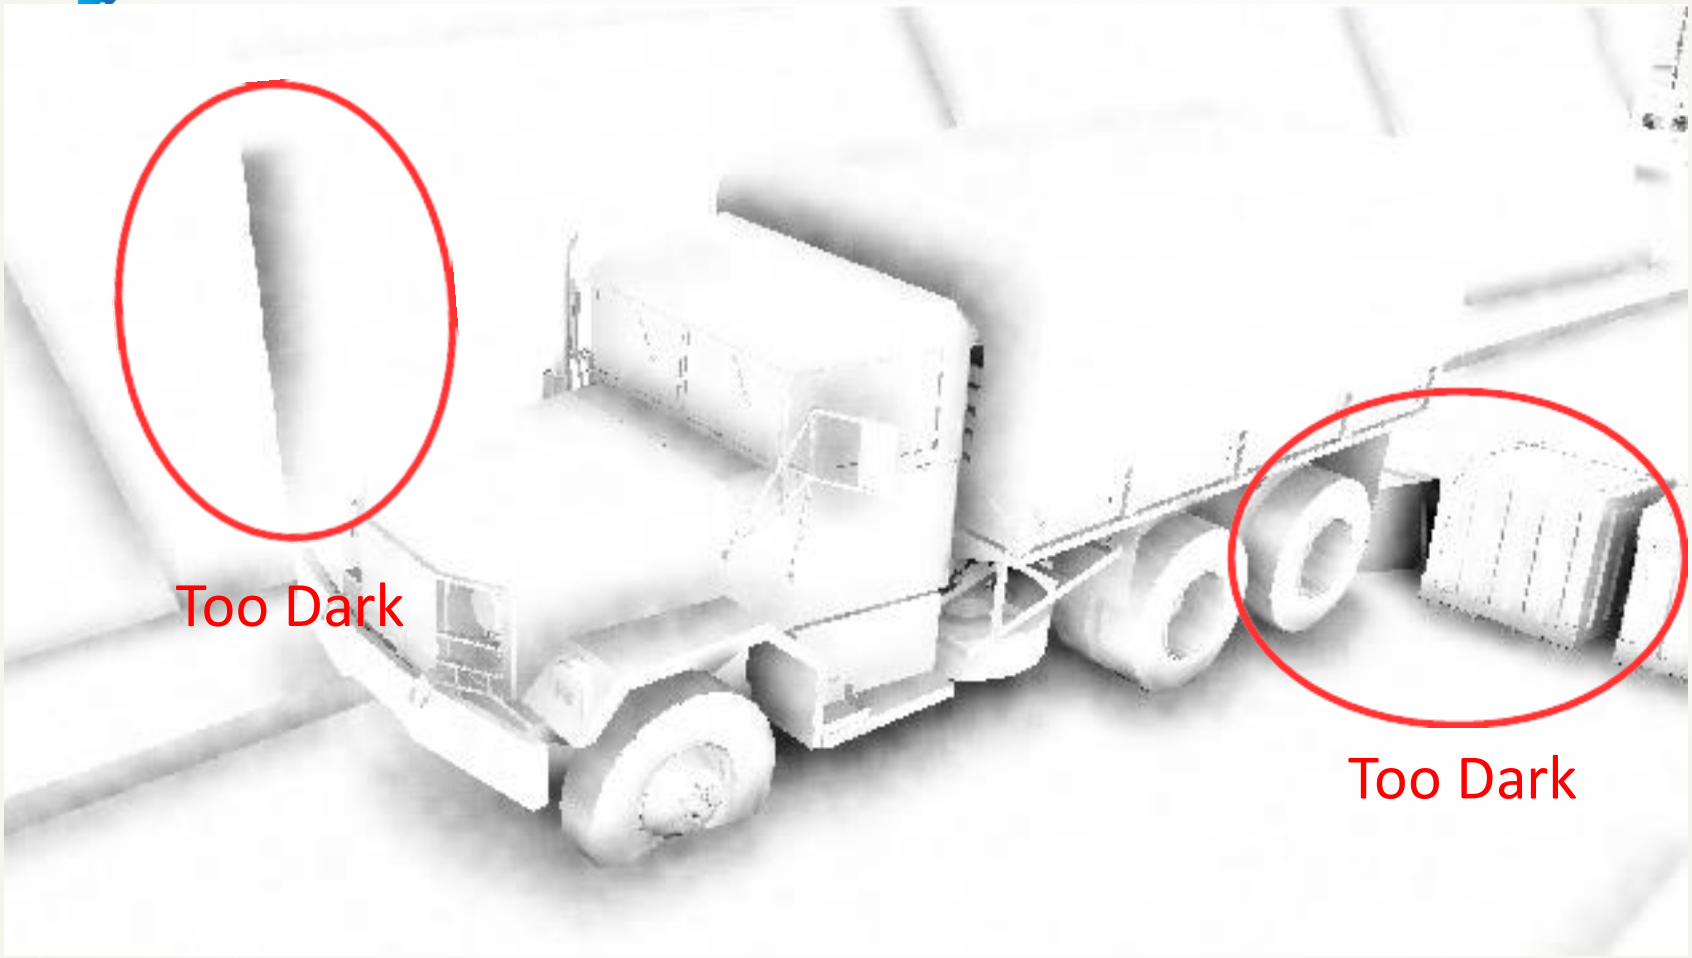

GROUND TRUTH

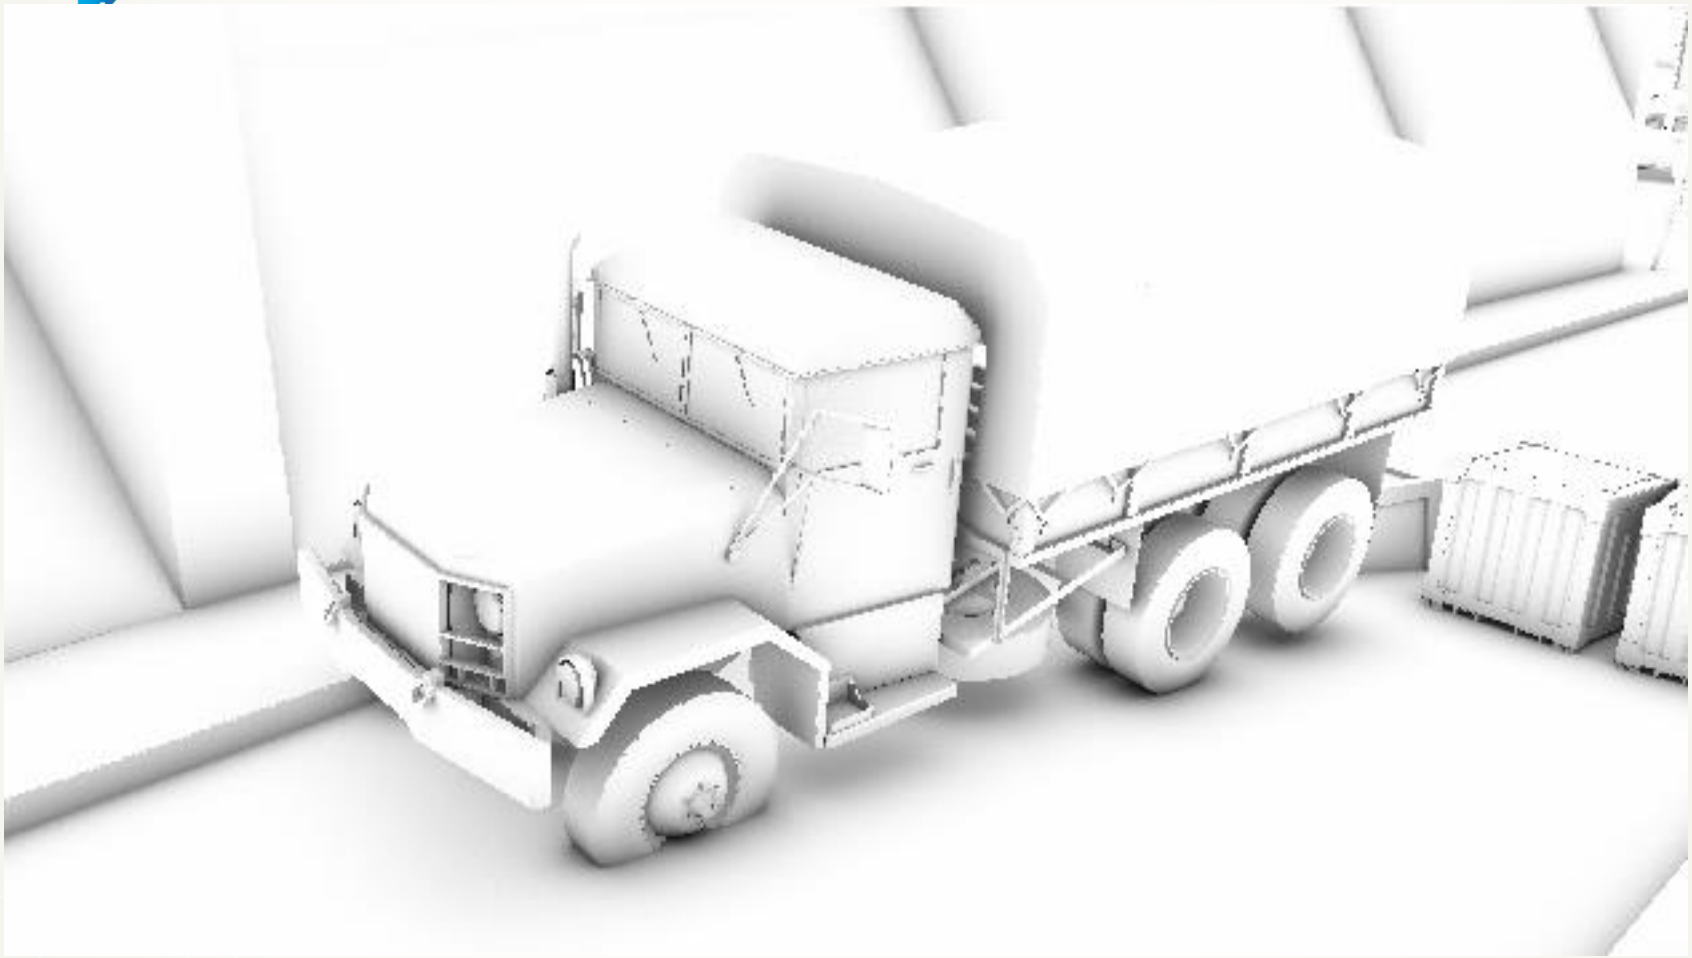

HBAO

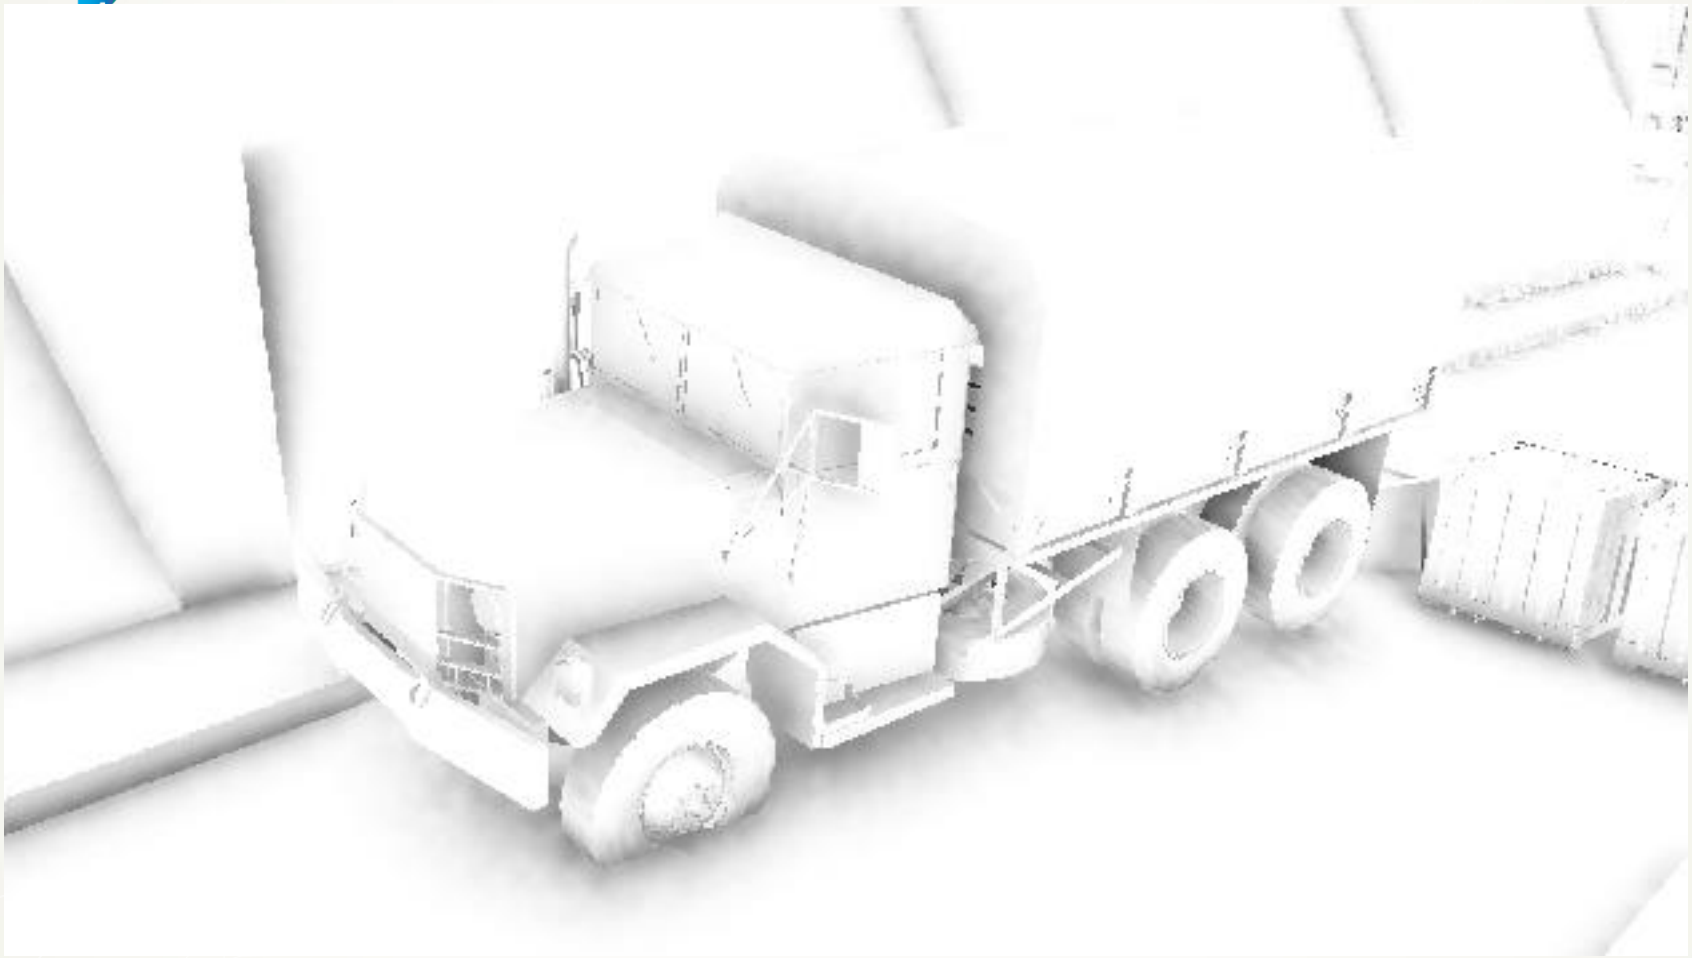

HBAO

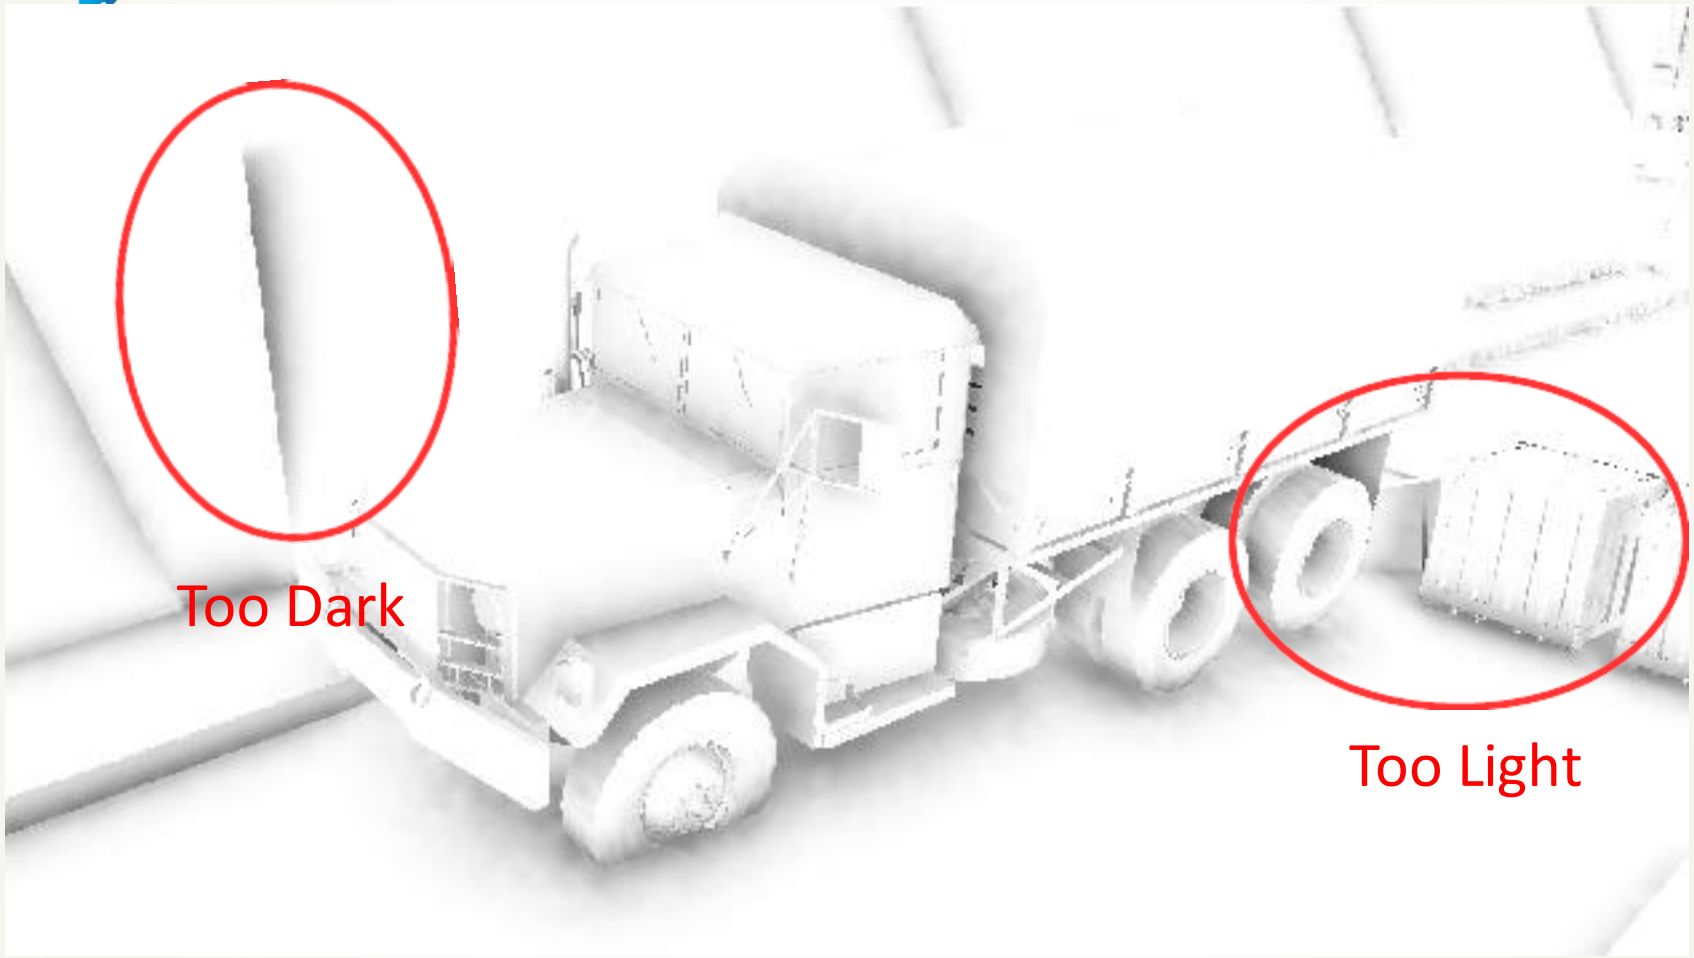

GROUND TRUTH

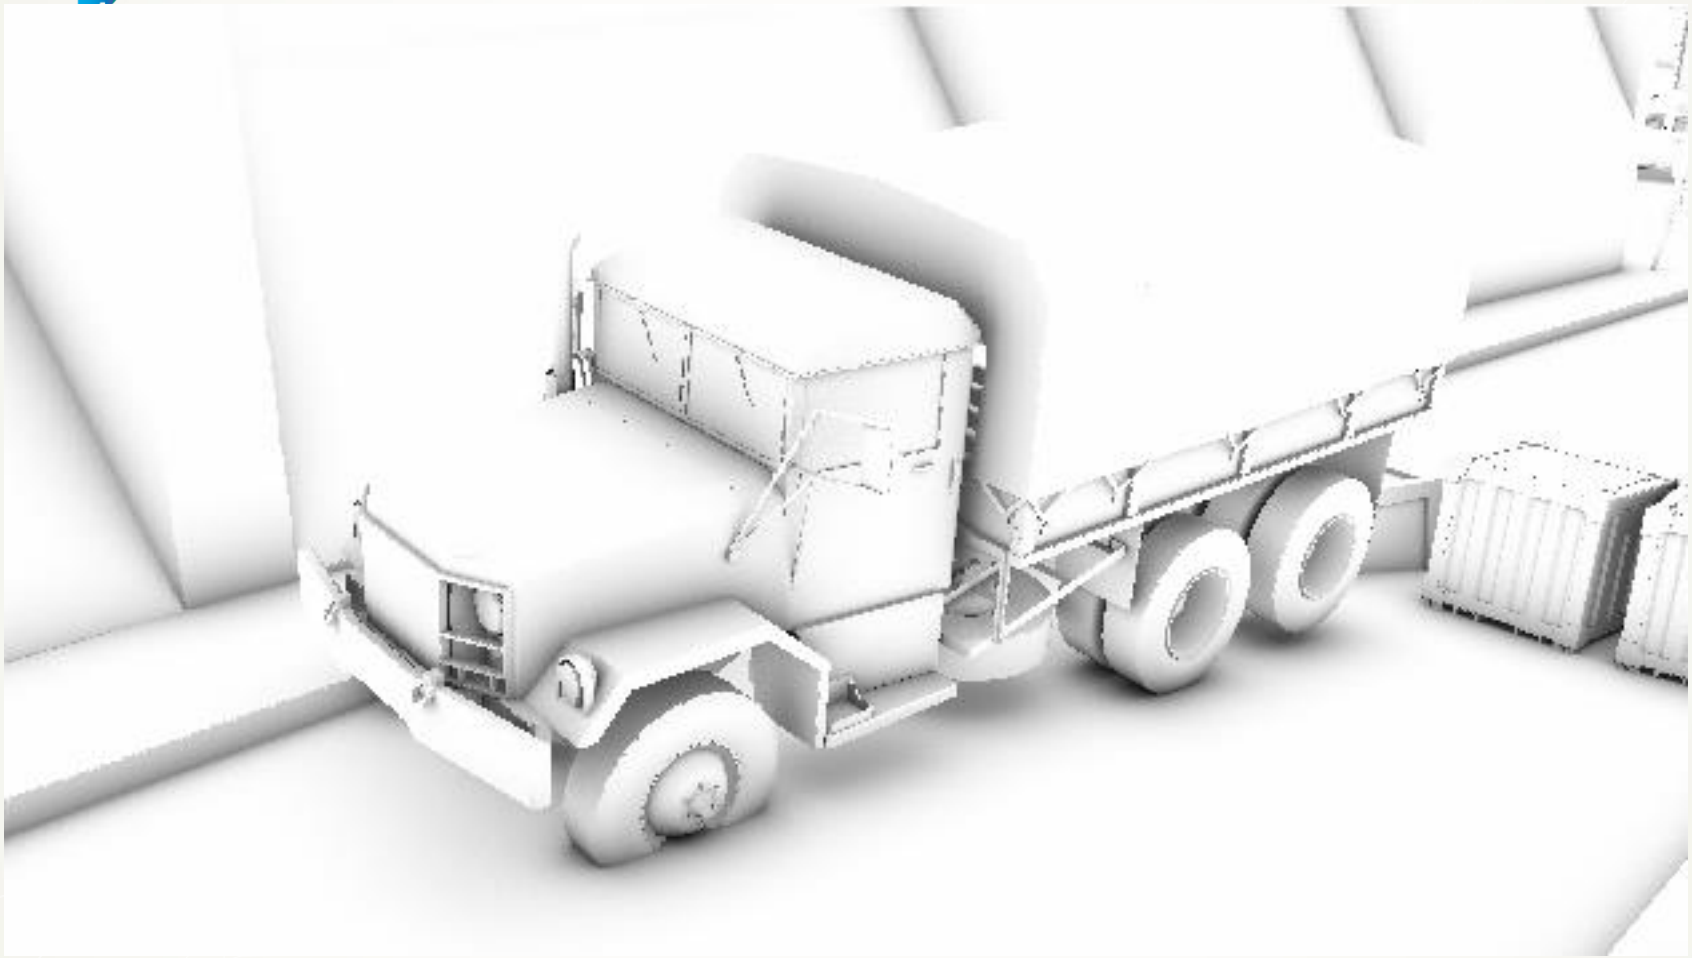

NNAO

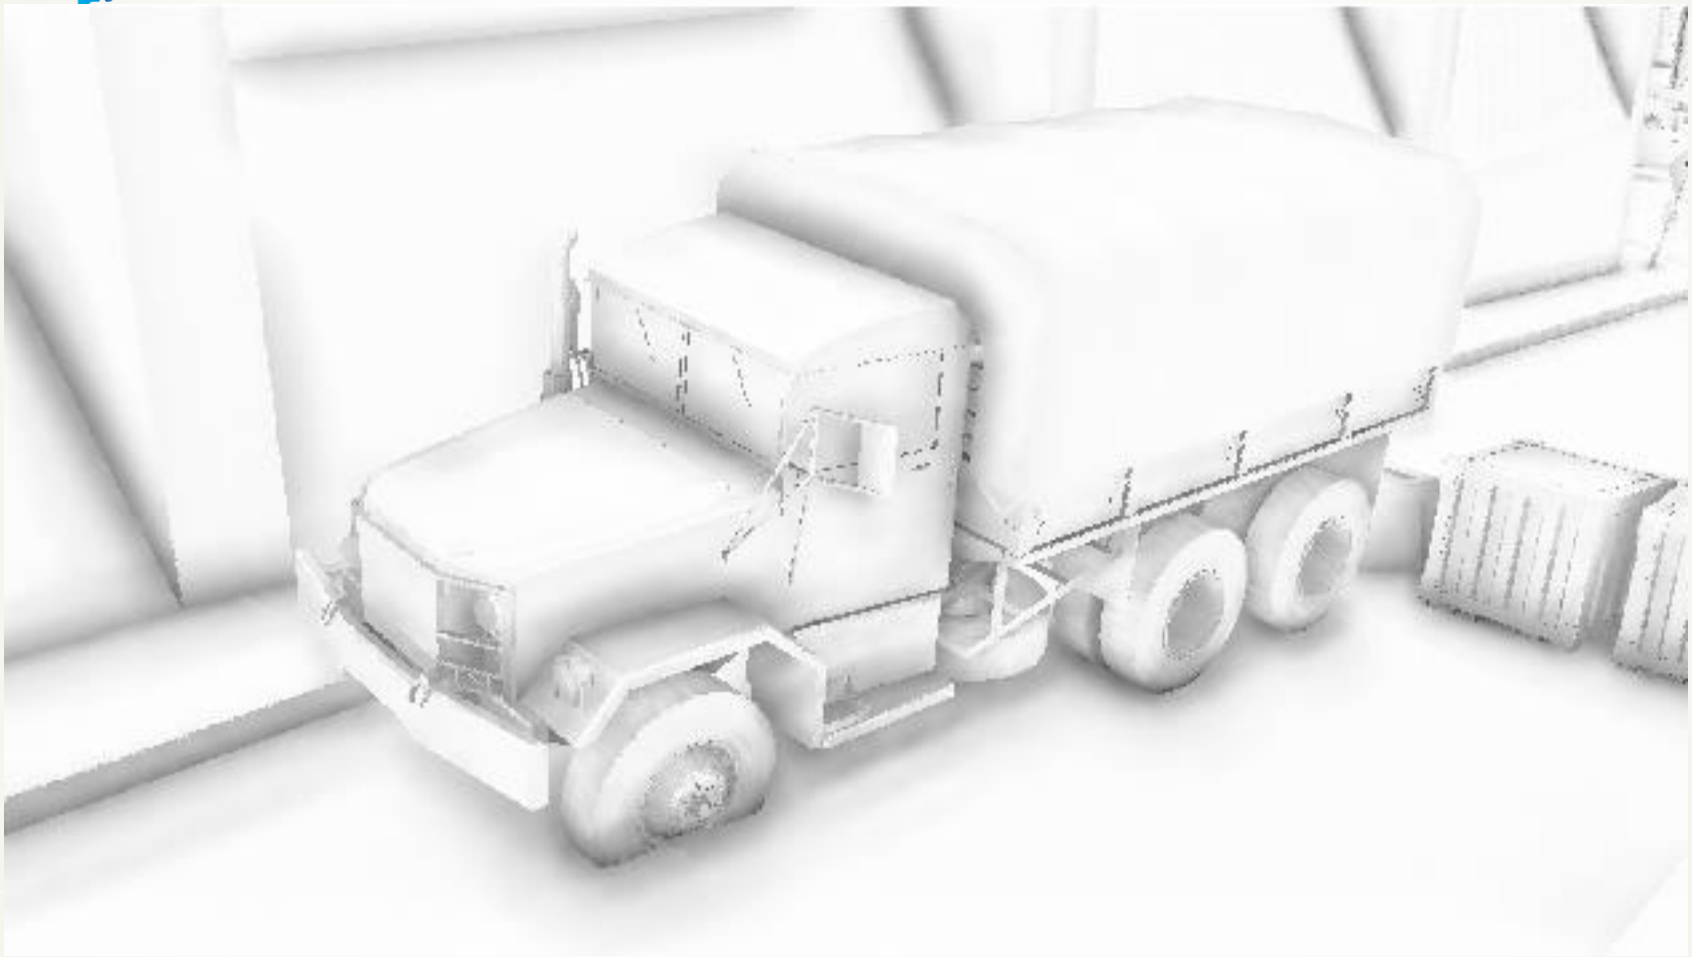

NNAO

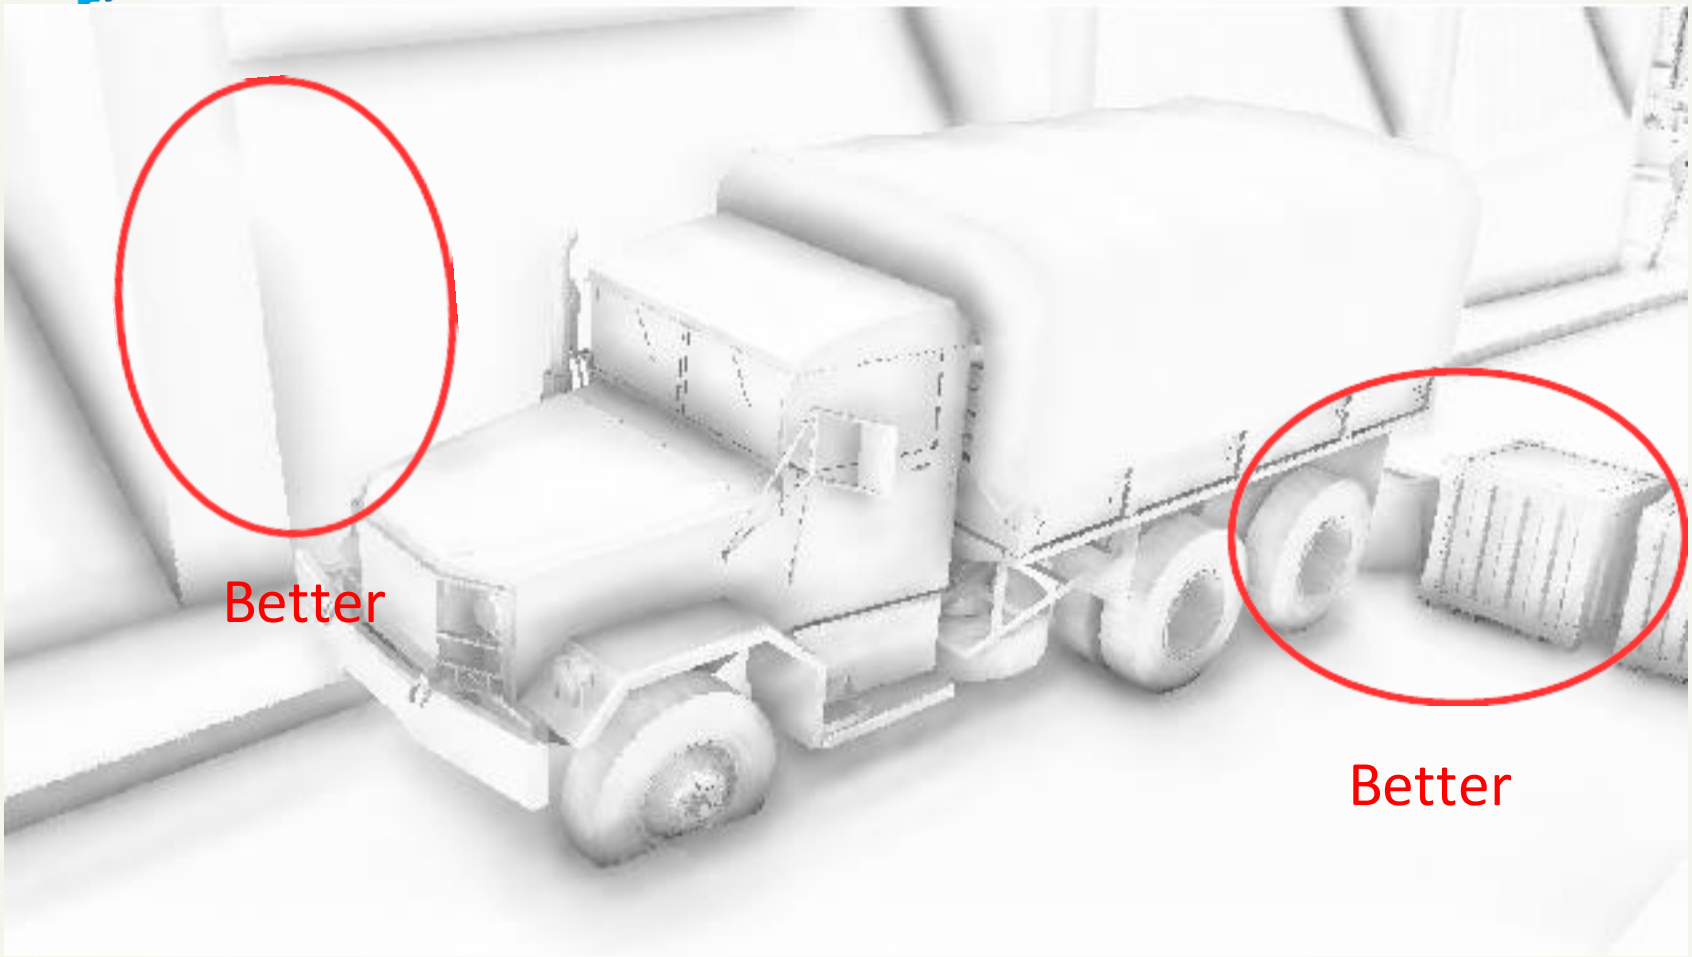

NNAO

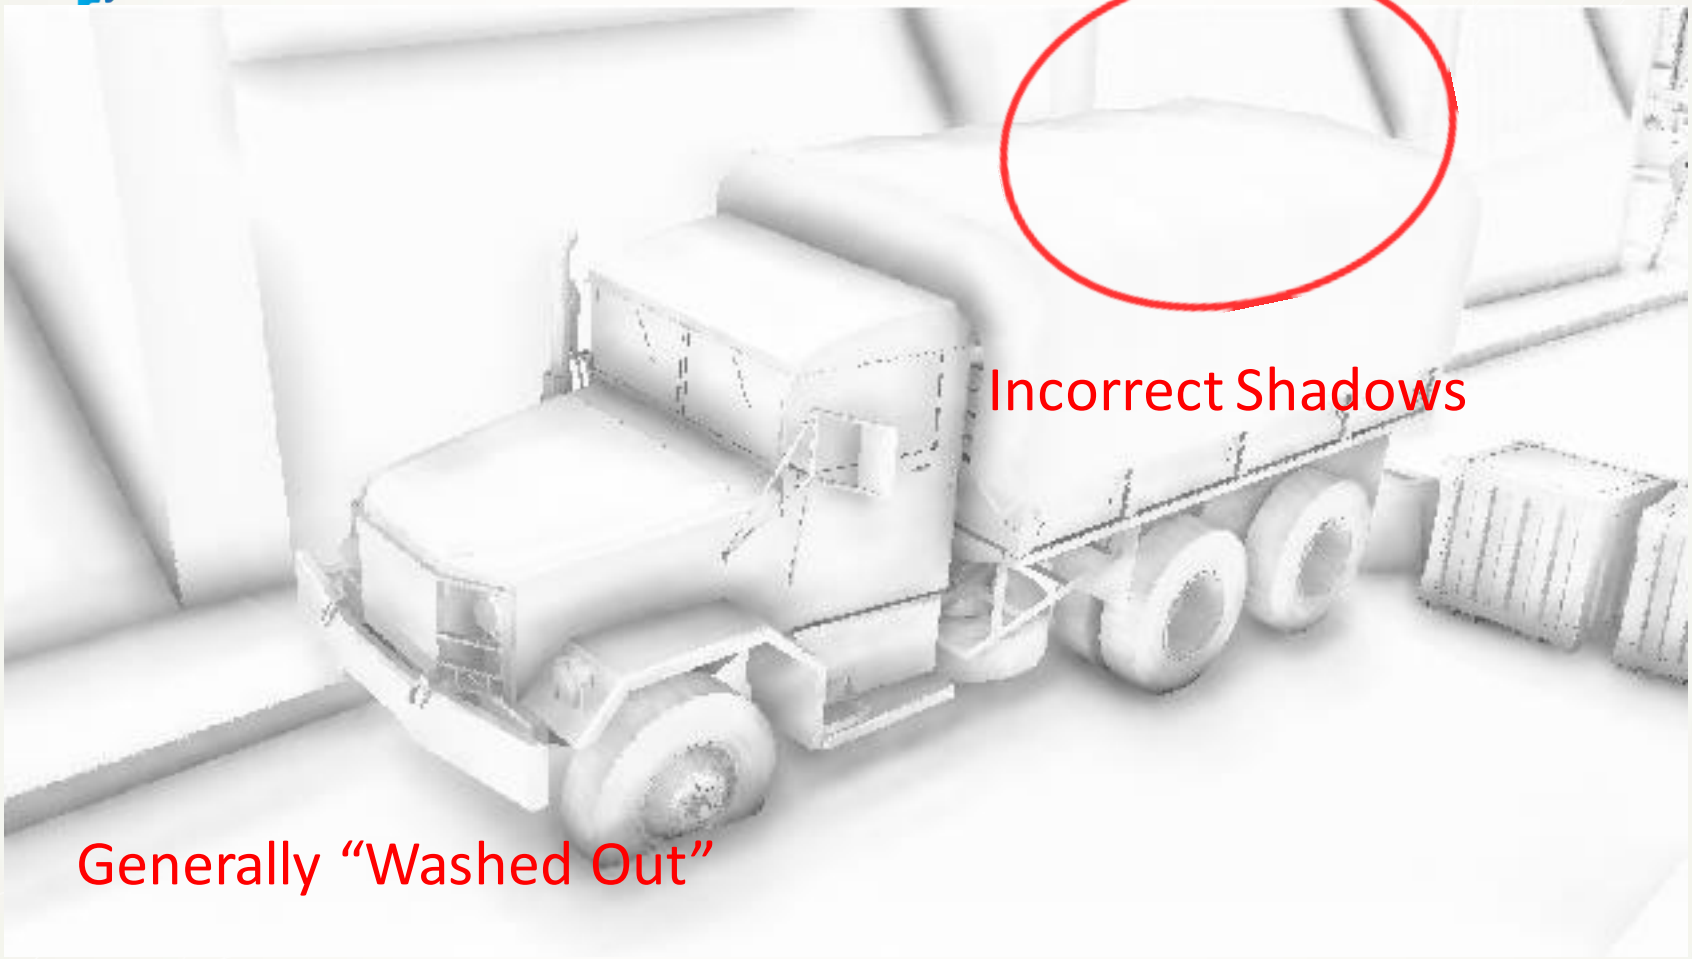

COMPARISON

Algorithm	Sample Count	Runtime (ms)	Error (mse)
SSAO	4	1.20	1.765
SSAO	8	1.43	1.558
SSAO	16	14.71	1.539
SSAO+	4	1.16	0.974
SSAO+	8	1.29	0.818
SSAO+	16	14.46	0.811
HBAO	16	3.53	0.965
HBAO	32	4.83	0.709
HBAO	64	8.50	0.666
NNAO	64	4.17	0.510
NNAO	128	4.81	0.486
NNAO	256	6.87	0.477

COMPARISON

Algorithm	Sample Count	Runtime (ms)	Error (mse)
SSAO	4	1.20	1.765
SSAO	8	1.43	1.558
SSAO	16	14.71	1.539
SSAO+	4	1.16	0.974
SSAO+	8	1.29	0.818
SSAO+	16	14.46	0.811
HBAO	16	3.53	0.965
HBAO	32	4.83	0.709
HBAO	64	8.50	0.666
NNAO	64	4.17	0.510
NNAO	128	4.81	0.486
NNAO	256	6.87	0.477

COMPARISON

SSAO [4ms] SSAO+ [4ms] HBAO [4ms] NNAO [4ms]

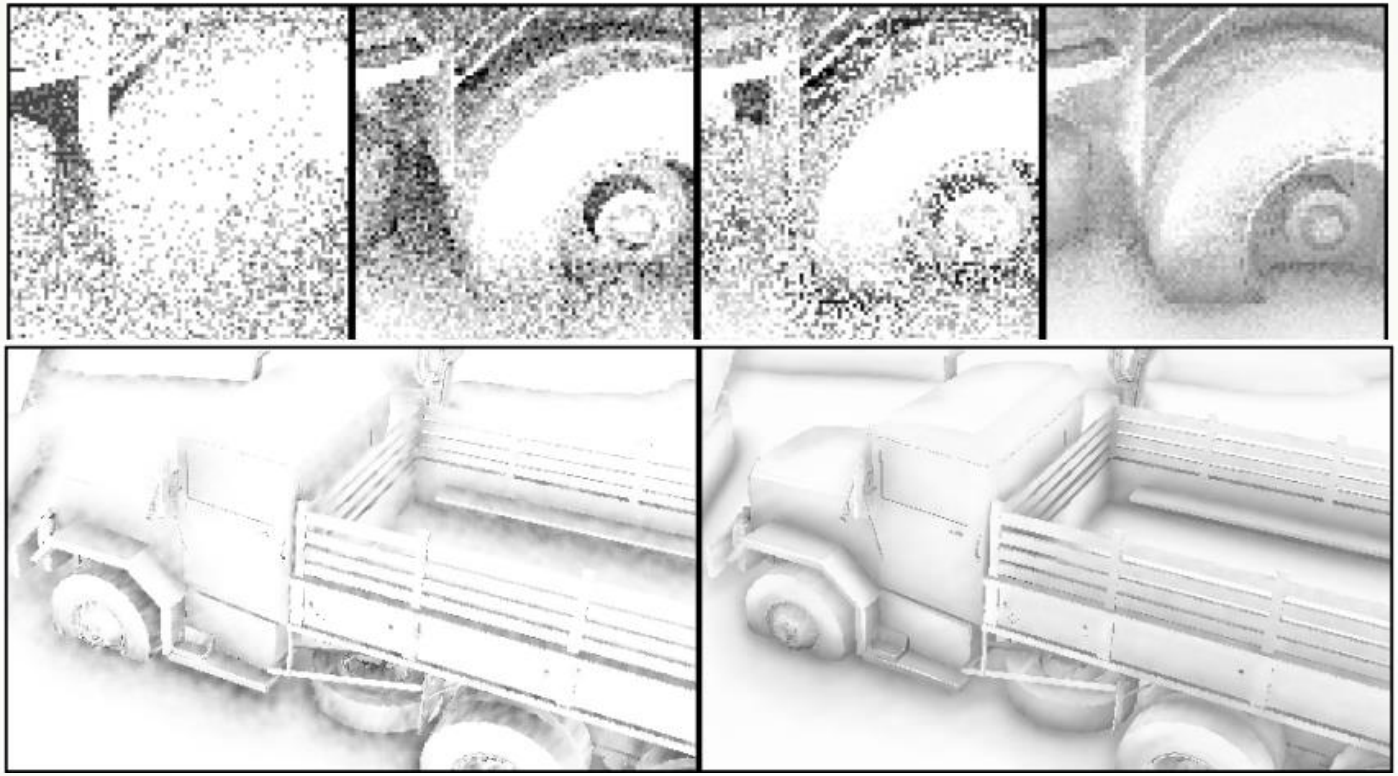

HBAO [32 samples] [4ms] NNAO [128 samples] [4ms]

Neural Network Ambient Occlusion

Daniel Holden
University of Edinburgh

Jun Saito
Method Studios

Taku Komura
University of Edinburgh

FUTURE WORK

- Train on GBuffer with detailed normals
- Apply to other screen space effects
  - Reflections
  - Subsurface Scattering
  - Indirect Illumination

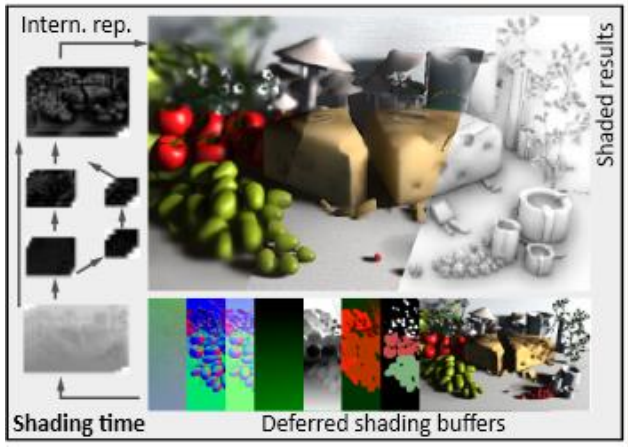

Deep Shading [Nalbach et al. 2016]

CONCLUSION

- We learn a Screen Space Ambient Occlusion function using a Neural Network.
- Designed to be drop in replacement to existing SSAO shaders.
- Faster and more accurate in many cases than previous methods.