PAVE

An In Situ Framework for Scientific Visualization and Machine Learning Coupling

Samuel Leventhal
University of Utah School of Computing
Scientific Computing and Imaging Institute
Salt Lake City, UT., USA
samlev@cs.utah.edu

Mark Kim
Oak Ridge National Laboratory
Oak Ridge TN., USA
kimmb@ornl.gov

David Pugmire
Oak Ridge National Laboratory
Oak Ridge TN., USA
pugmire@ornl.gov

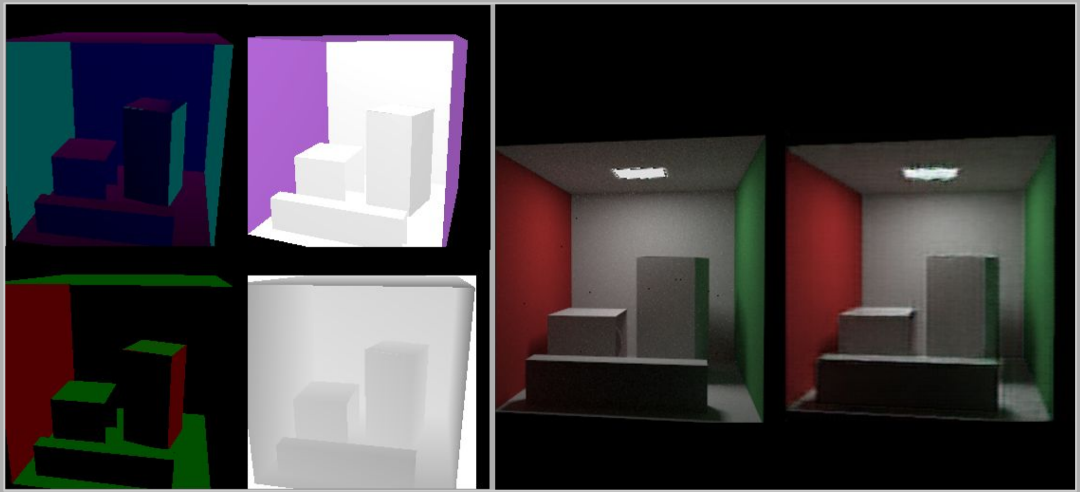

Oak Ridge National Laboratory

www.sci.utah.edu

PAVE: An in situ framework allowing researchers and practitioners to couple scientific visualization and machine learning tasks.

Motivation

Solution

Framework

Case Study and Application of PAVE

Motivation:

I/O remains a limiting factor in **traditional HPC** which often **produces large amounts of output.**

This limitations of I/O are increasing as **machine learning** (ML) is growingly more important in HPC and **requires not only large amounts of data** ('big data') but also data in combination **with large amounts of output** in order **to produce an algorithm.**

With ML, the role of data and output has changed, becoming more integral in algorithm design.

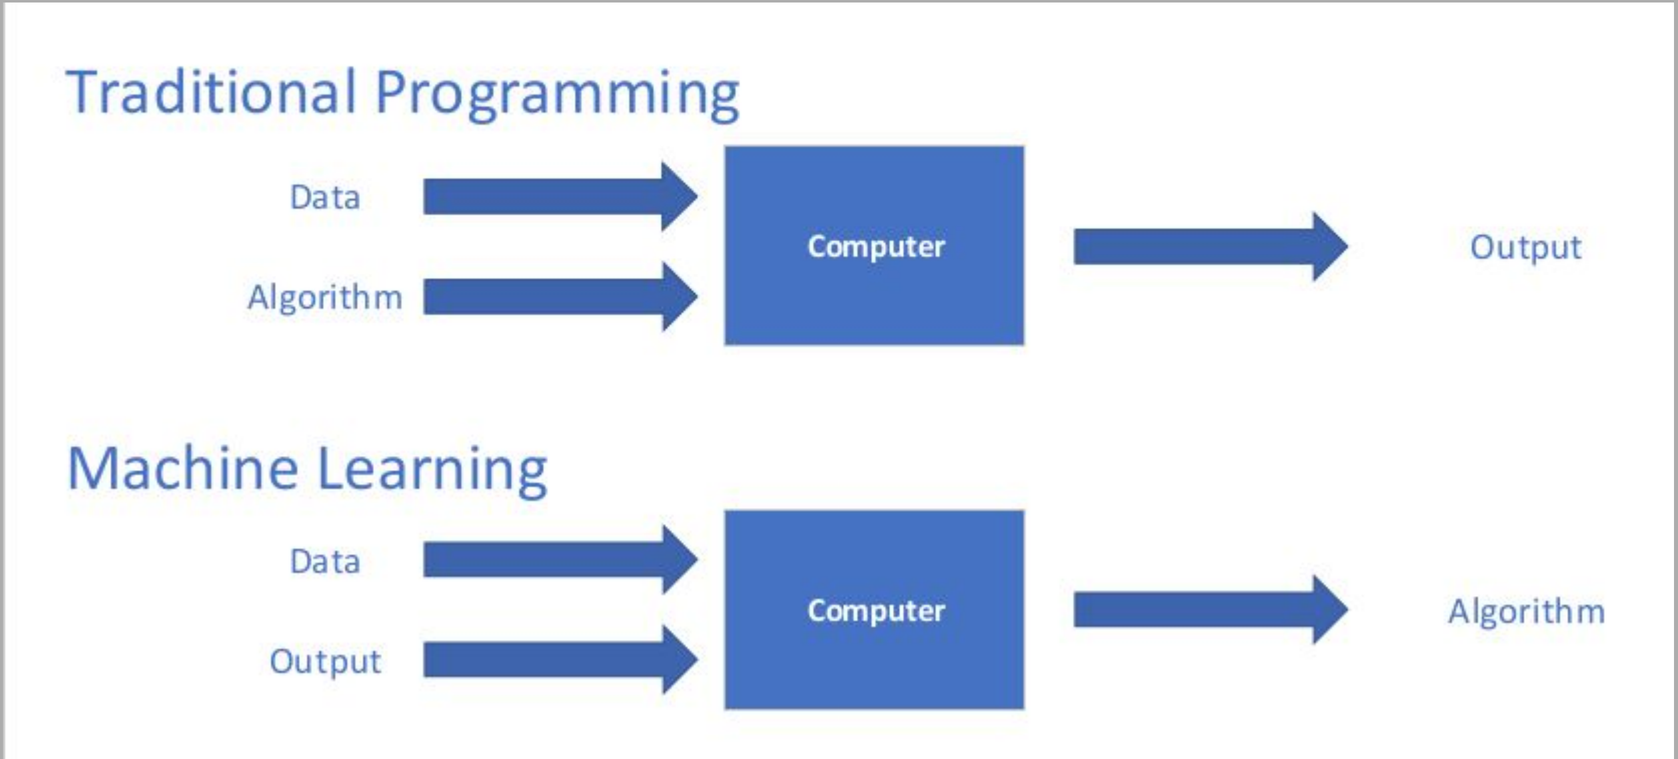

Scientific Visualization Output

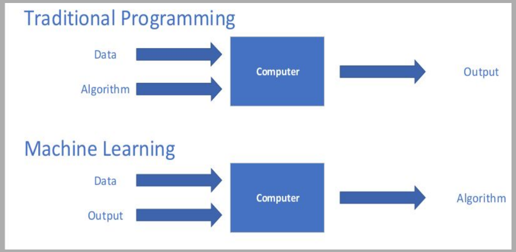

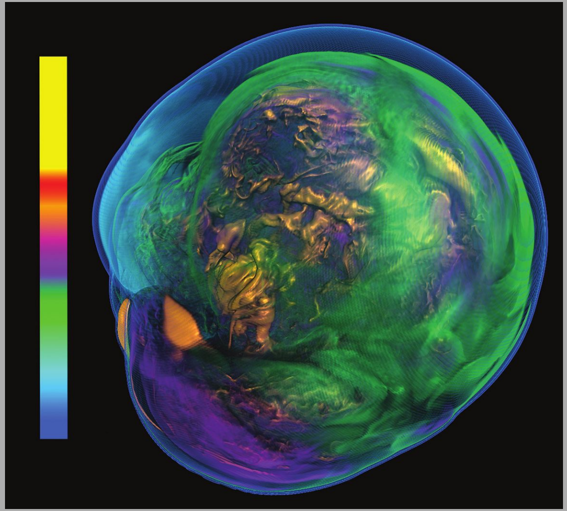

Supernova Direct **Volume Rendering**

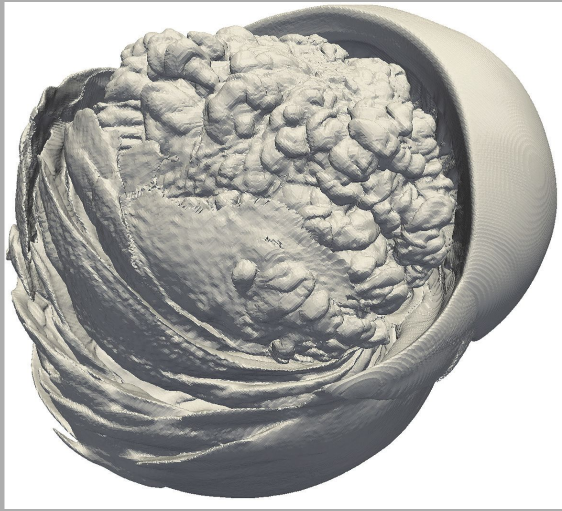

Supernova **Contour**

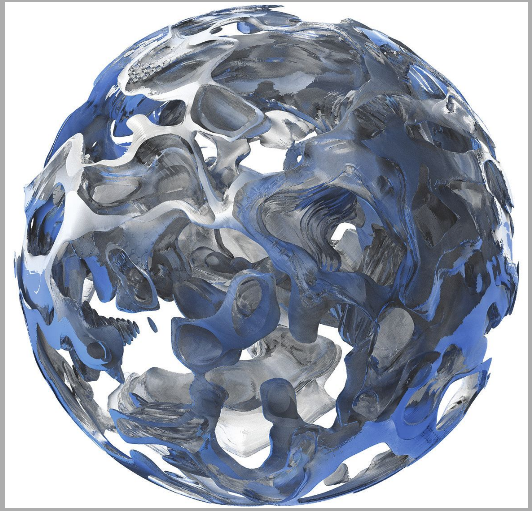

Ray-tracing:
Seismic Wave Propagation

VTK-m: Accelerating the Visualization Toolkit for Massively Threaded Architectures, Moreland et al. 2016

Oak Ridge National Laboratory

www.sci.utah.edu

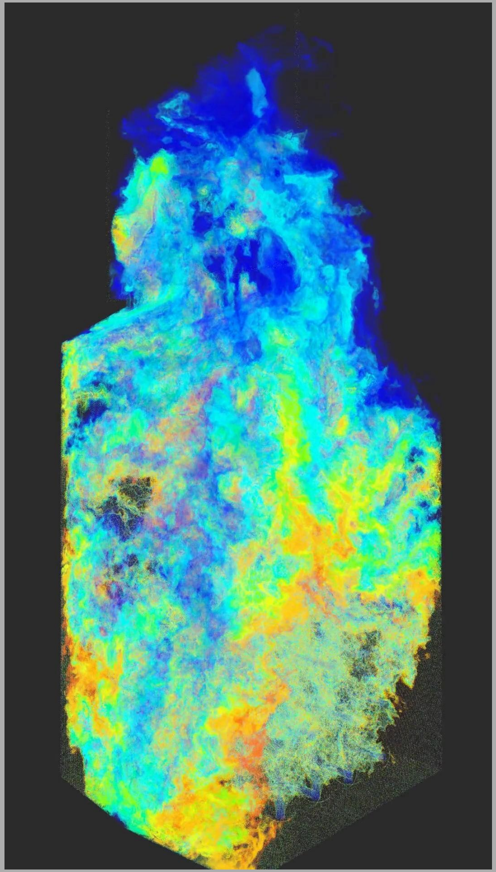

"Ultra Super Critical" Oxy-Coal

Kumar, Sidharth, et al. "Scalable data management of the Uintah simulation framework for next-generation engineering problems with radiation." *Asian Conference on Supercomputing Frontiers*. Springer, Cham, 2018.

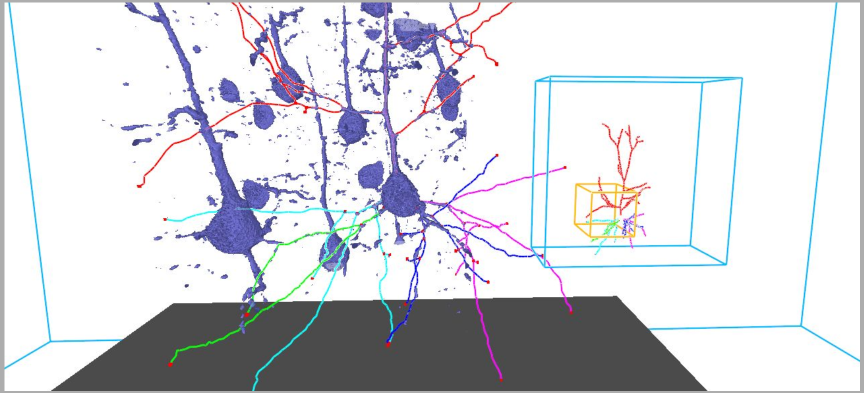

A Virtual Reality Visualization Tool for Neuron Tracing

Usher, W., Klacansky, P., Federer, F., Bremer, P. T., Knoll, A., Yarch, J., ... & Pascucci, V. (2017). A virtual reality visualization tool for neuron tracing. *IEEE transactions on visualization and computer graphics*, *24*(1), 994-1003.

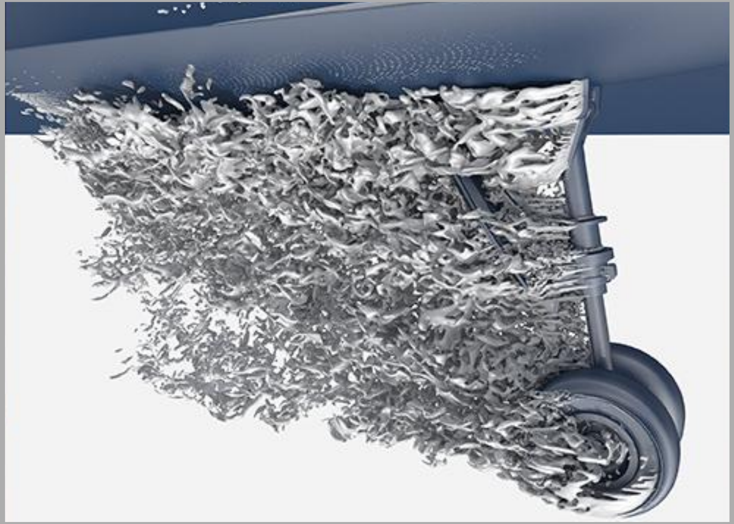

Isosurface representation of the vorticity, rendered with path tracing

Wang, Feng, et al. "CPU Isosurface Ray Tracing of Adaptive Mesh Refinement Data." *IEEE transactions on visualization and computer graphics* 25.1 (2018): 1142-1151.

Machine Learning <u>Output</u>

Generative Models: Gaussian mixture model, Hidden Markov model, Bayesian network, Generative Adversarial Network, Autoencoder, …

Neural Networks: Convolutional, Recurrent, Long-Short, …

Structured Predictions: Conditional Random Field, Structured SVM, Language Processing, …

Decision Trees

...

Motivating Aim

A scalable and efficient framework for combining scientific visualization (traditional programming) and machine learning tasks

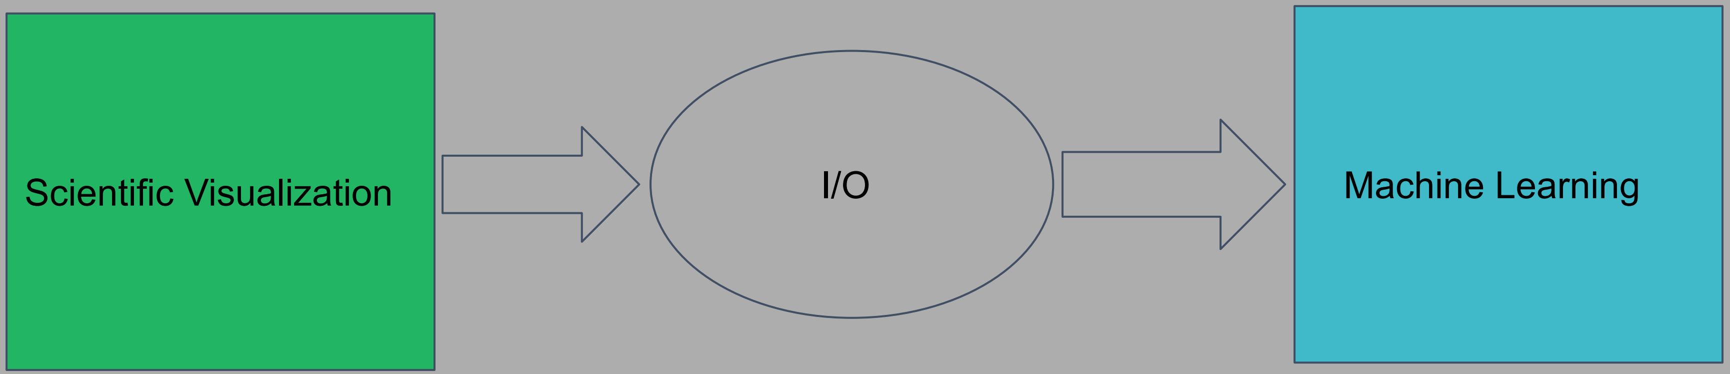

What do we have to address?

Problem: An increasing gap between compute performance and I/O bandwidth for HPC.

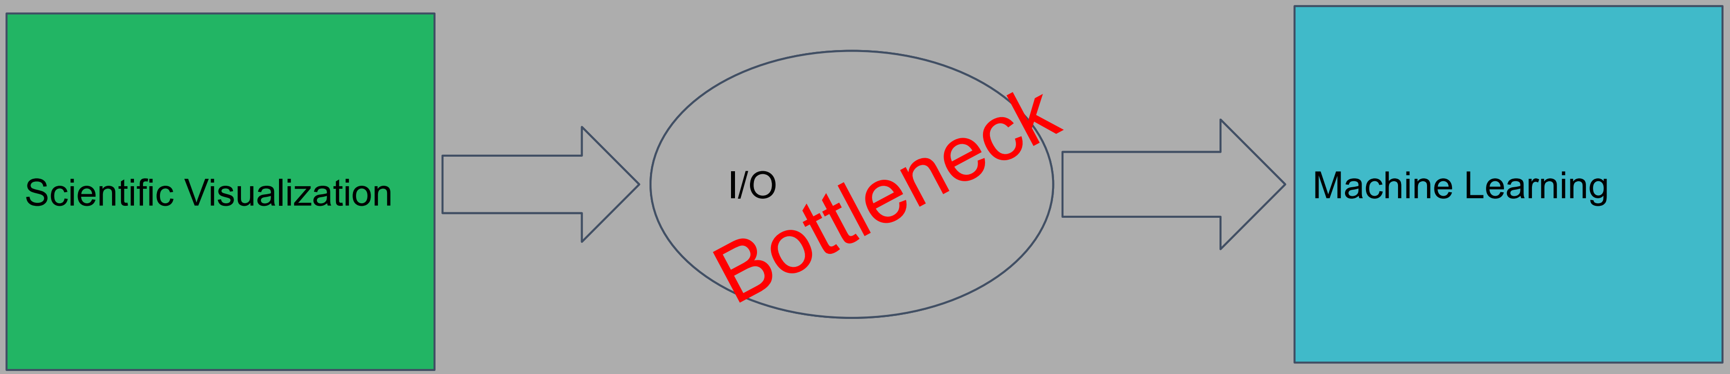

In Situ Transport: Coupling scientific visualization and machine learning.

- In situ data transport of visualization task used as input for machine learning model to produce a trained learning algorithm.

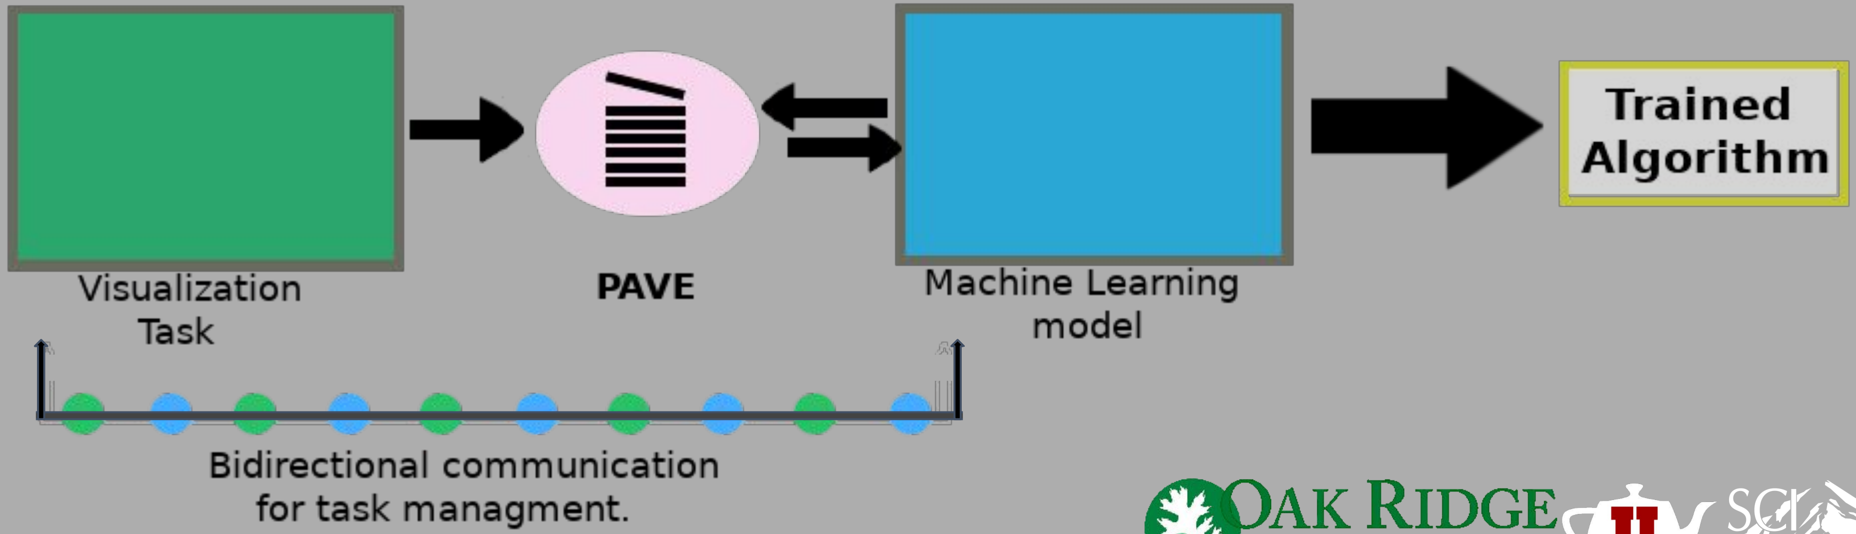

PAVE provides a simple API for in situ data transport for coupling user defined scientific visualization and machine learning tasks.

Case Study

Coupled Path Tracing and Conditional Adversarial Network

Scalable conditional generative adversarial network (cGAN) trained on the output of a path tracer to be able to emulate the behavior of light and global illumination for scene dependent geometry.

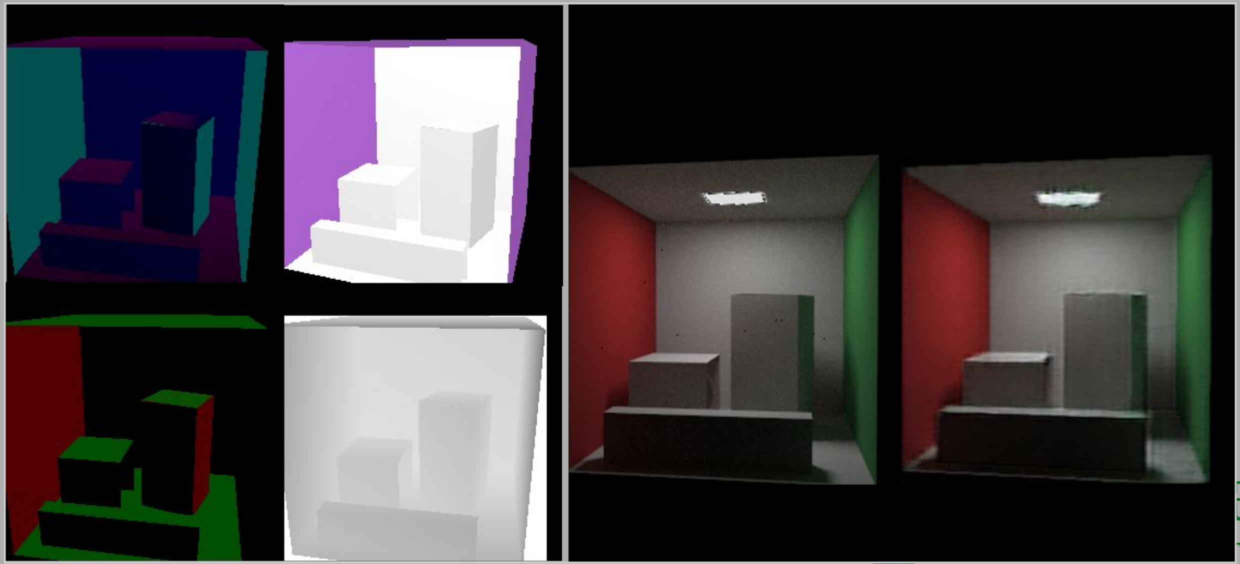

www.sci.utah.edu

End Goal: Path Tracer trains a Generative Adversarial Network (GAN) for emulating the behaviour of light in a given scene.

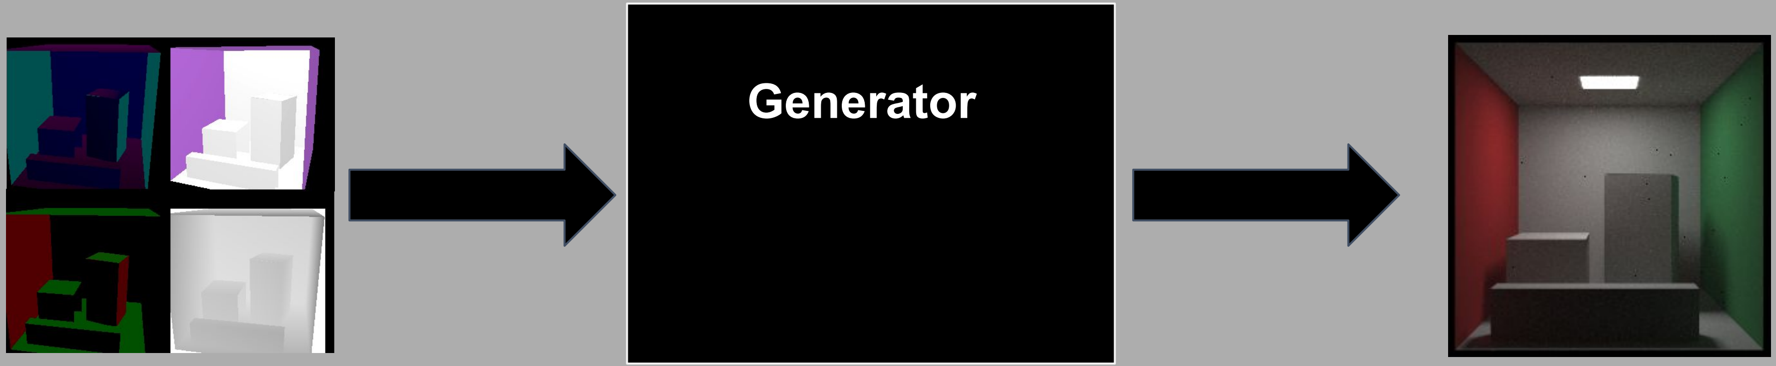

Motivation: Alleviate input/output bottleneck when coupling traditional programming and machine learning.

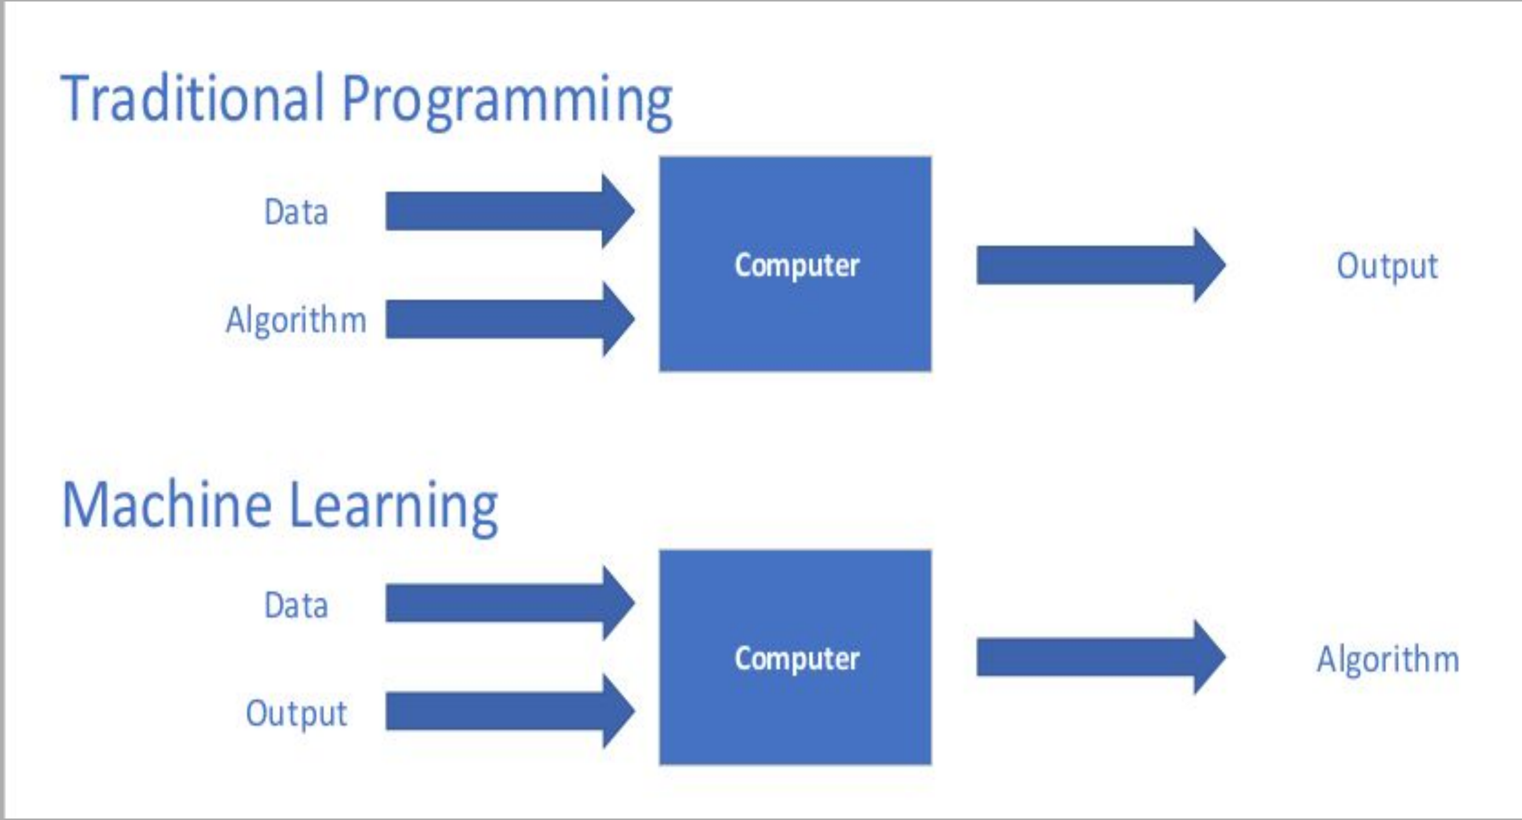

Traditional Programming

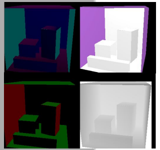

Scene Geometry and Direct Lighting

Path Tracer

Computer

Rendered Scene wi
Global Illumination

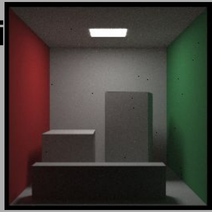

Machine Learning

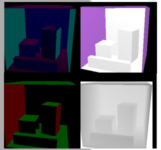

Scene Geometry and Direct Lighting

Rendered Scene with Global Illumination
"Ground Truth"

Computer

Scene Generator

System Overview: Coupling

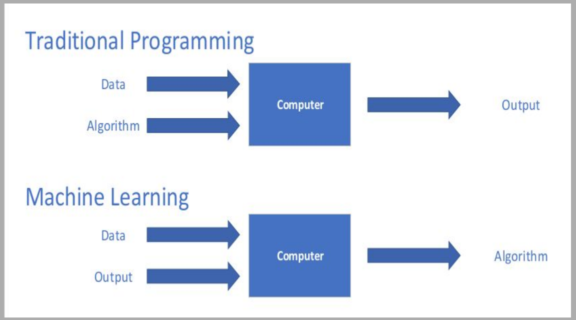

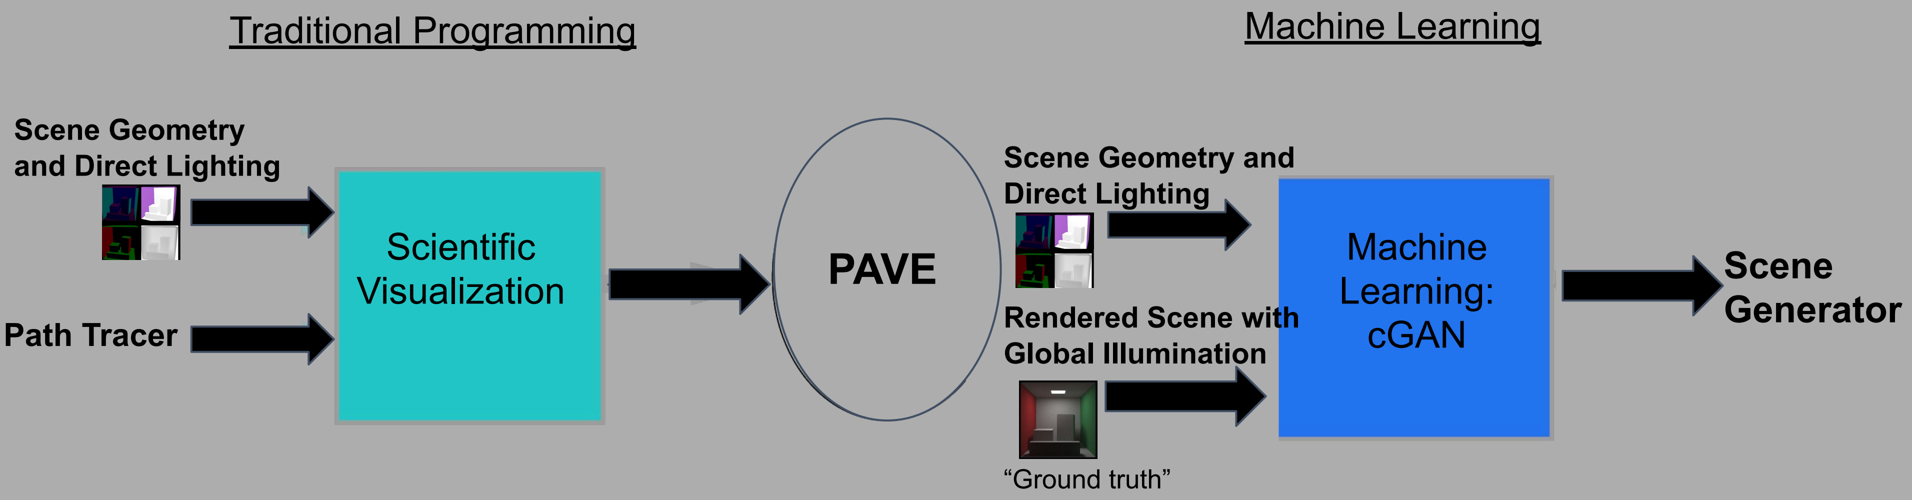

System Overview: User Defined Visualization

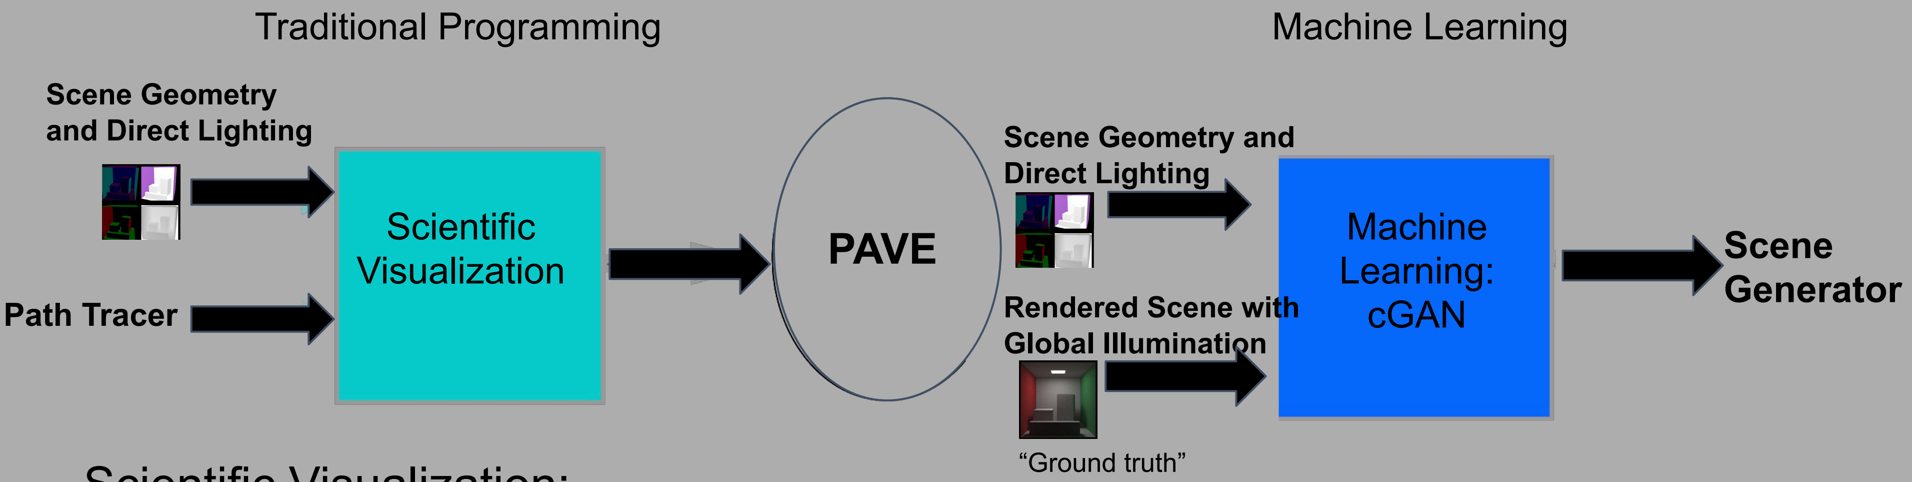

Scientific Visualization: Path Tracer using VTK-m

VTK-m

Data parallel primitive toolkit for scientific visualization suitable for massively threaded architectures.

VTK-m was chosen due to its scalability and HPC compatibility

Oak Ridge National Laboratory

U

SCI www.sci.utah.edu

System Overview: Learning Model

Traditional Programming

Scene Geometry and Direct Lighting

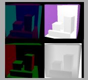

Path Tracer

Scientific Visualization

PAVE

Machine Learning

Scene Geometry and Direct Lighting

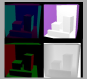

Rendered Scene with Global Illumination

"Ground truth"

Machine Learning: cGAN

Scene Generator

Scientific Visualization:
Path Tracer using VTK-m

Learning Goal:
cGAN: U-Net + patchGAN (CNN)
with PyTorch

National Laboratory

www.sci.utah.edu

Path Tracer Design

Path Tracer: Monte Carlo sampling over shapes, light intensity and pixel values. The sampling is performed by sending multiple rays per pixel traversing the scene, accumulating the color of the surfaces it interacts with until it hits a light source.

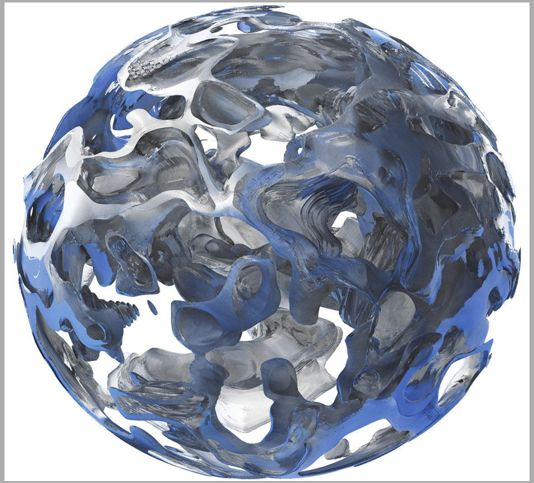

Ray-tracing:
Seismic Wave Propagation

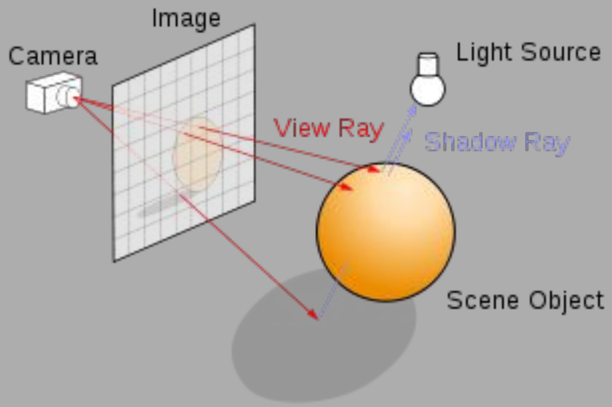

Training Generative Neural Networks

Cops and Counterfeiters, a zero sum game.

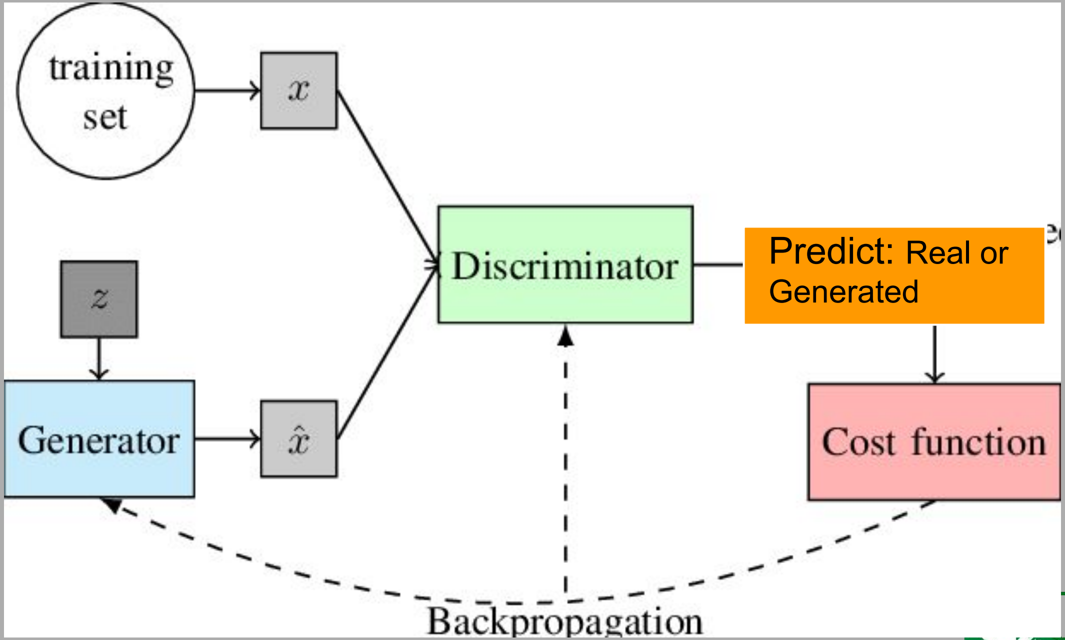

Case Study

Training the generator

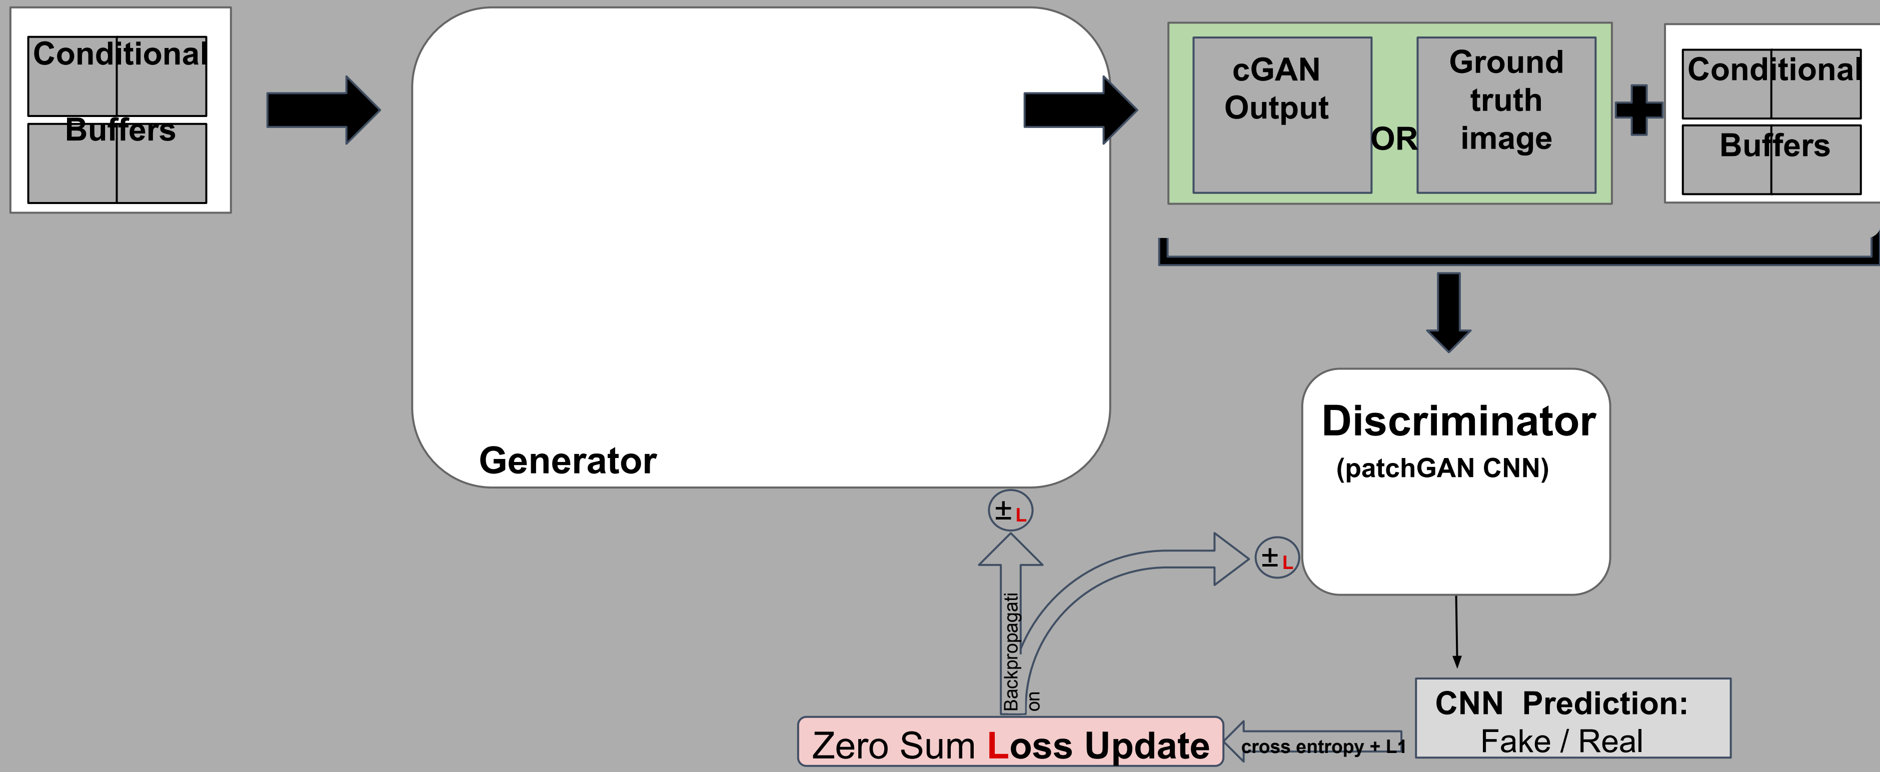

Case Study

Training the generator

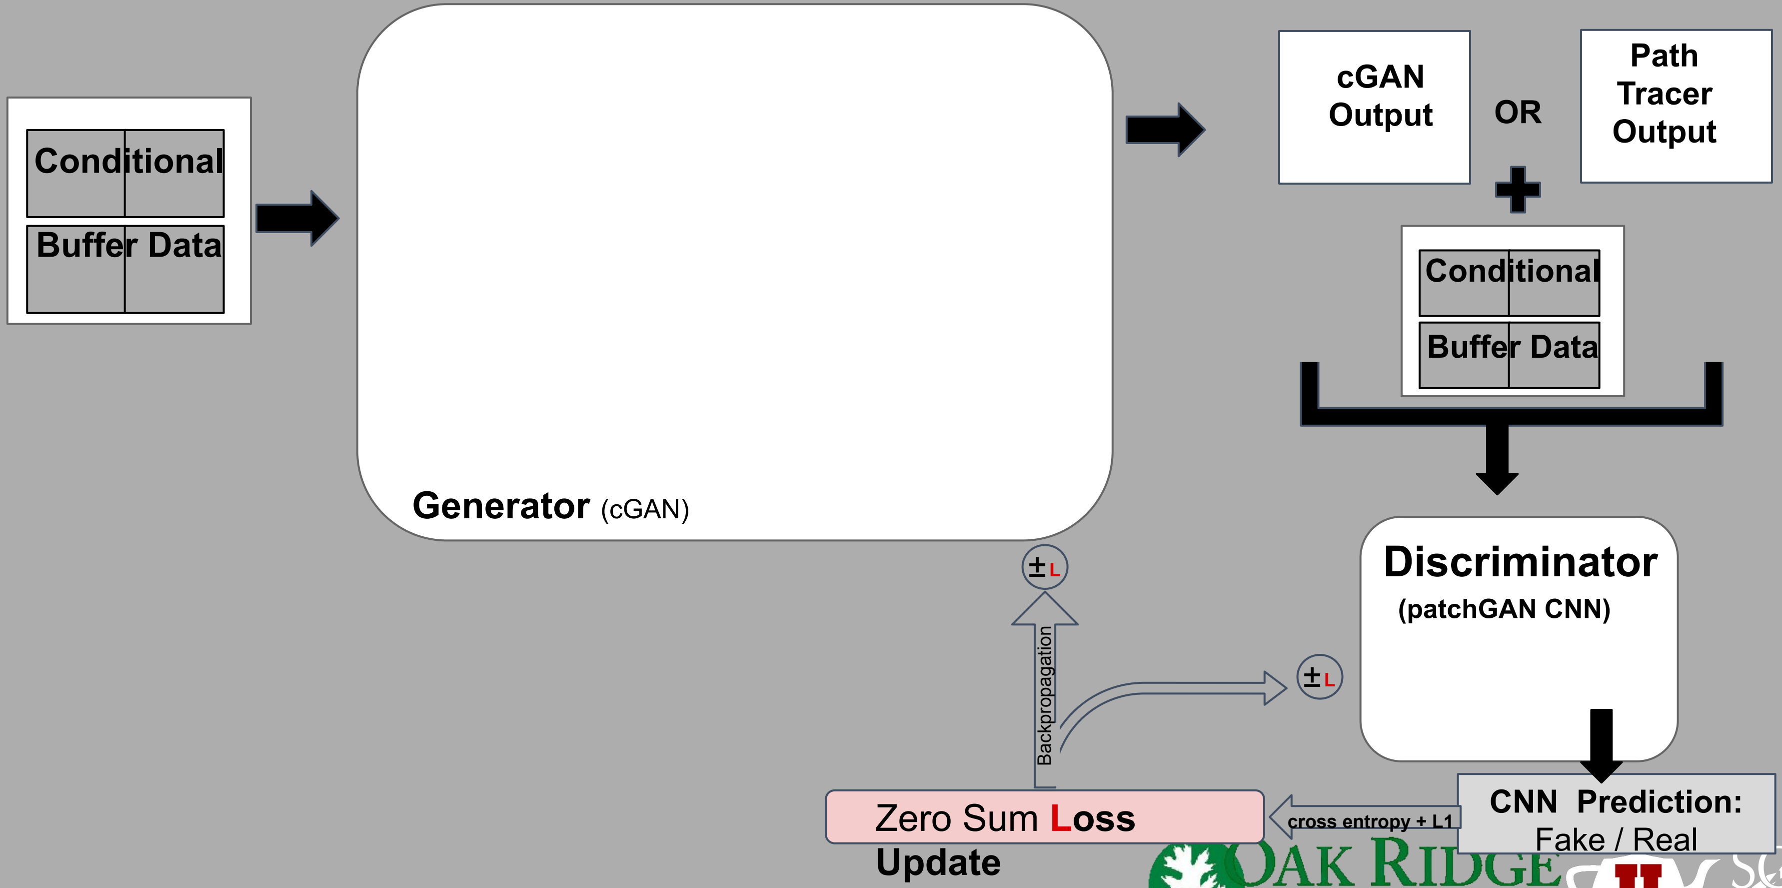

Data from VTK-m Pathtracer: Geometric Buffers

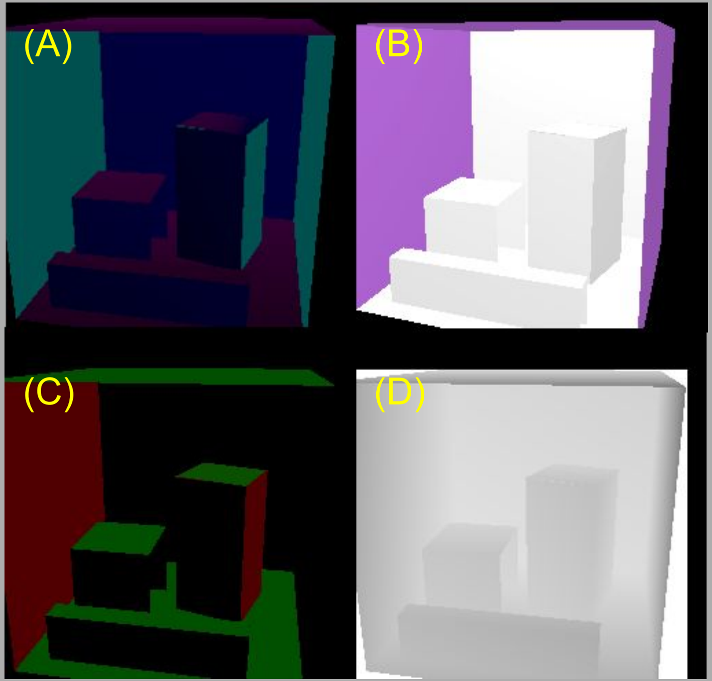

Image buffers needed to compute light paths offer conditional dependency on the behavior of light based on the geometry and light sources within a scene.

Material property parameters, direct lighting, color coded normals of Surfaces and depth: (a) – (d) respectively.

Data for both path tracer and cGAN along with the output from the path tracer for cGAN

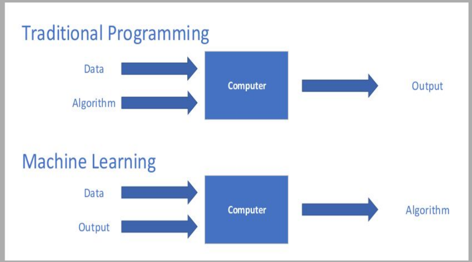

Data: Conditional Geometric Buffers

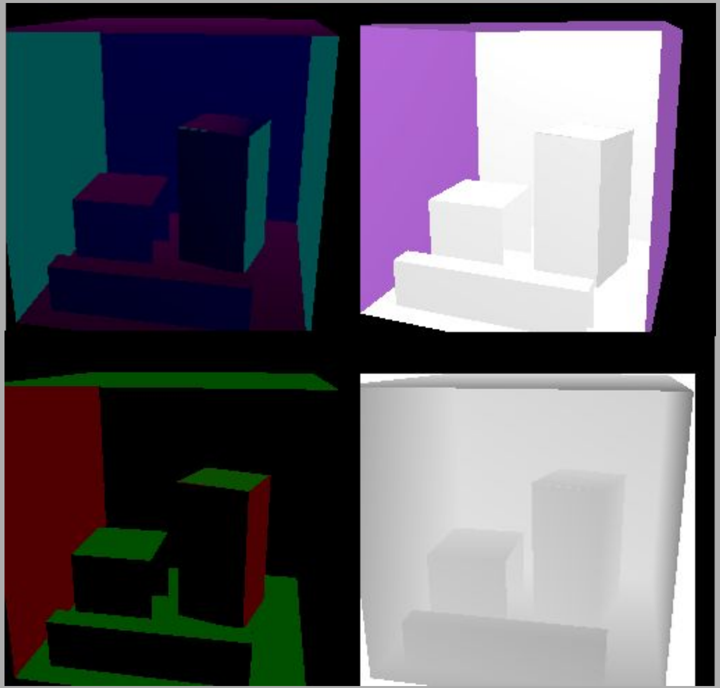

Output: Ground Truth Rendering

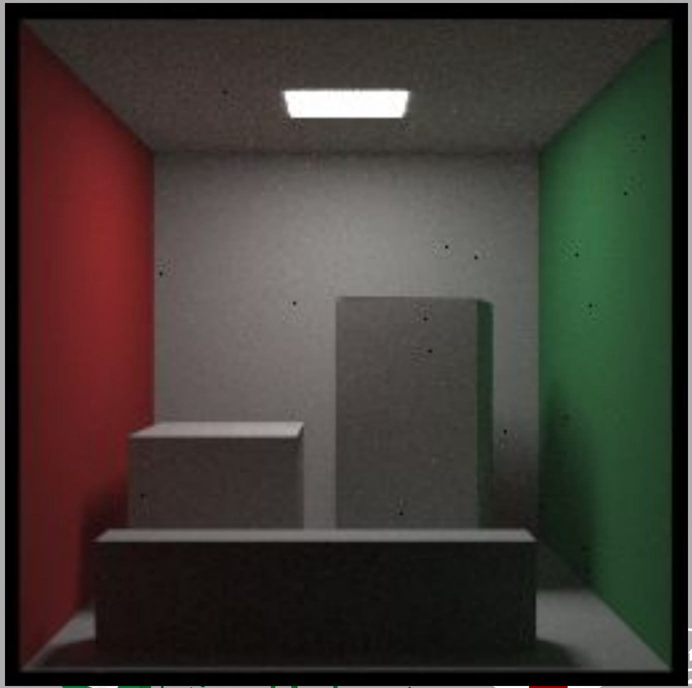

Conditional Generative Adversarial Network

Conditional Dependence

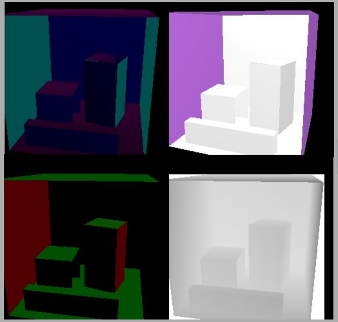

U-Net

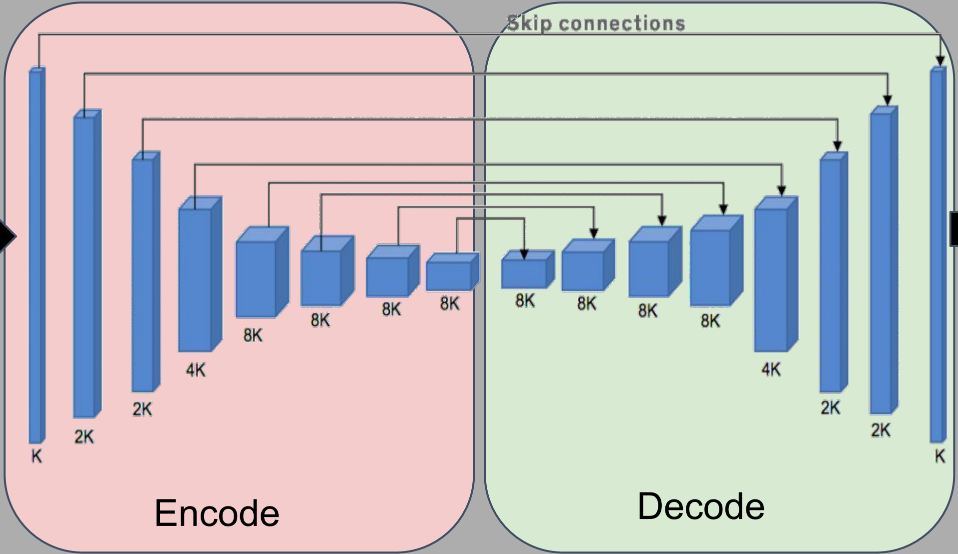

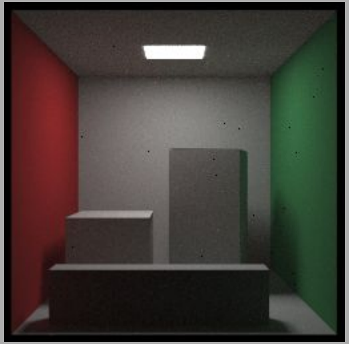

Discriminator

Deep Convolutional patchGAN Network (Isola et. al. [30])

Added advantage of providing a patch-wise probability and regional accuracy within an image in question as being real or fake.

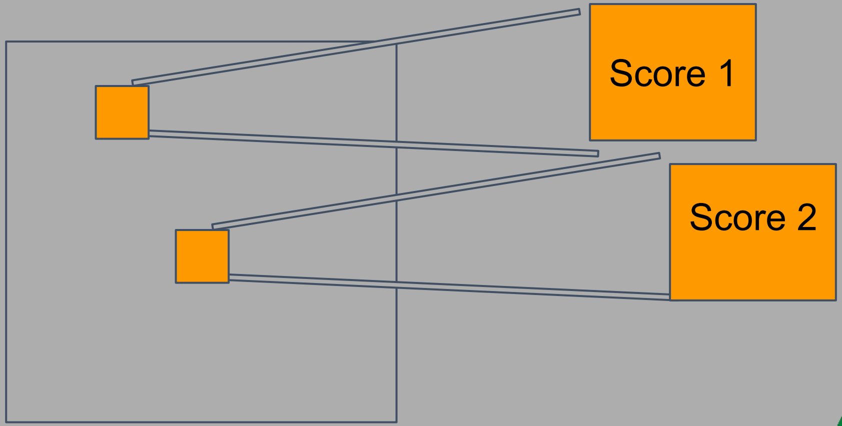

Score Image = Average(Score_i)

Generative Network Loss

Spatial relations and location is important.

Added L1 norm loss to cross entropy due to Structural preservation of U-Net due to concatenate and regional influence due to patch-wise probability in Discriminator.

Gradient descent based optimization as zero-sum game between discriminator and generator.

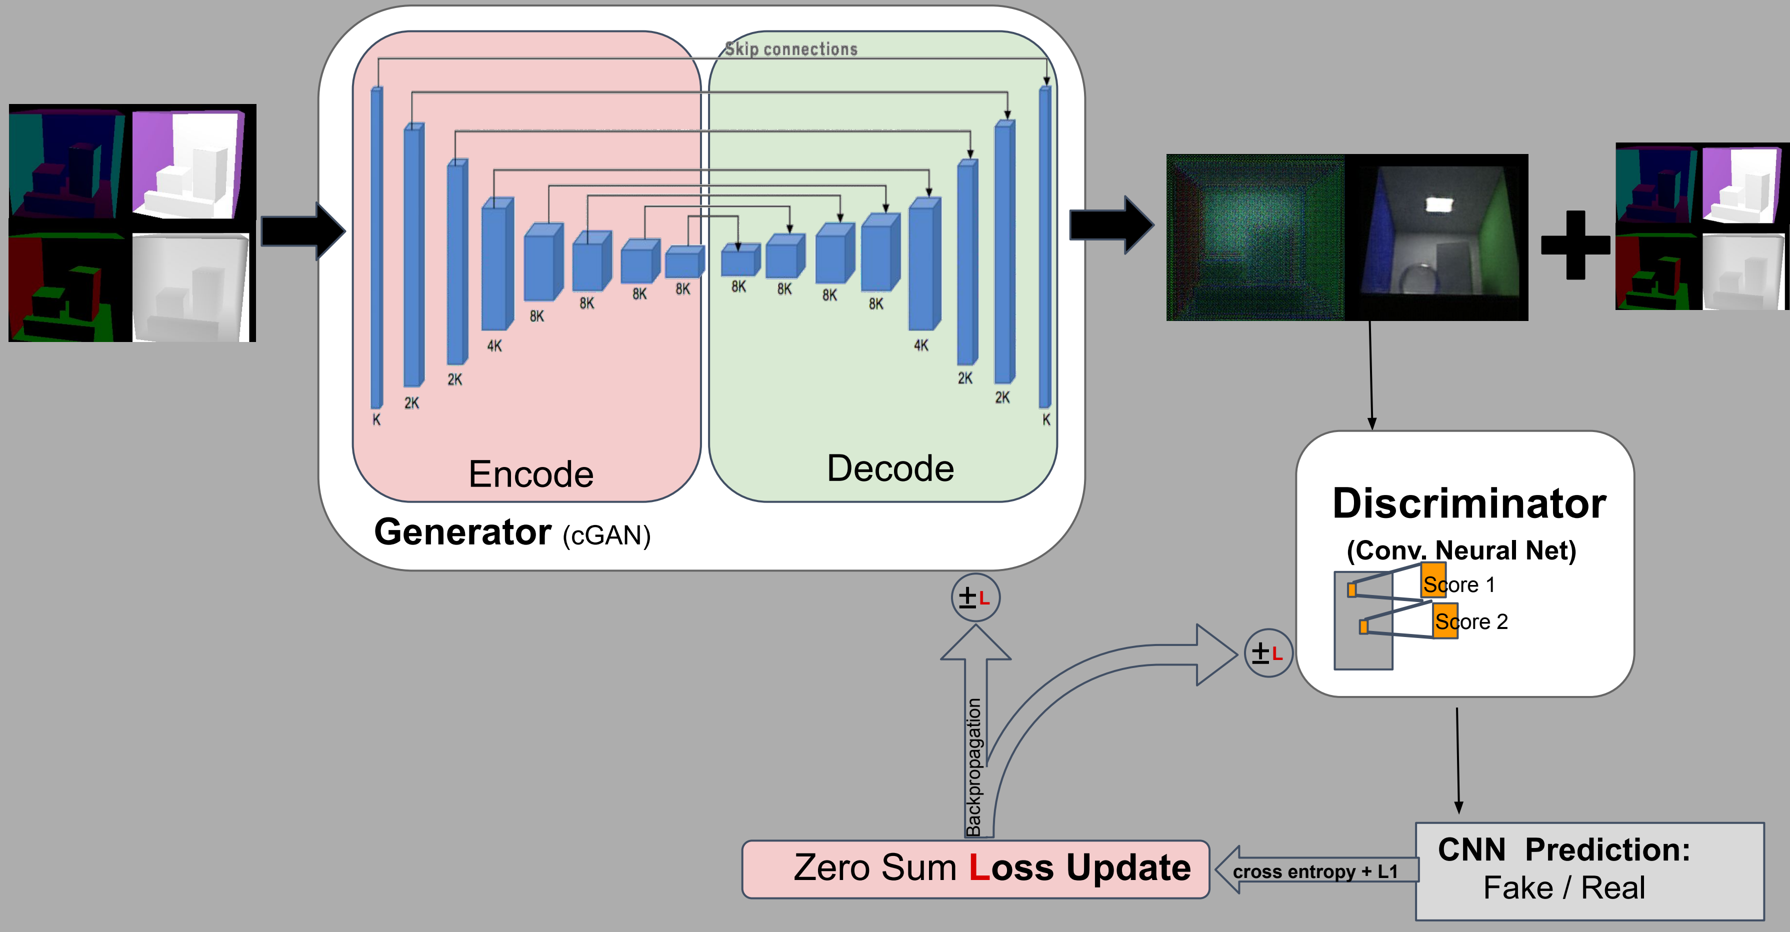

Skip connections
K
2K
2K
4K
8K
8K
8K
8K
8K
8K
8K
8K
4K
2K
2K
K
Encode
Decode
Generator (cGAN)
Discriminator
(Conv. Neural Net)
Score 1
Score 2
±L
±L
Backpropagation
Zero Sum Loss Update
cross entropy + L1
CNN Prediction:
Fake / Real

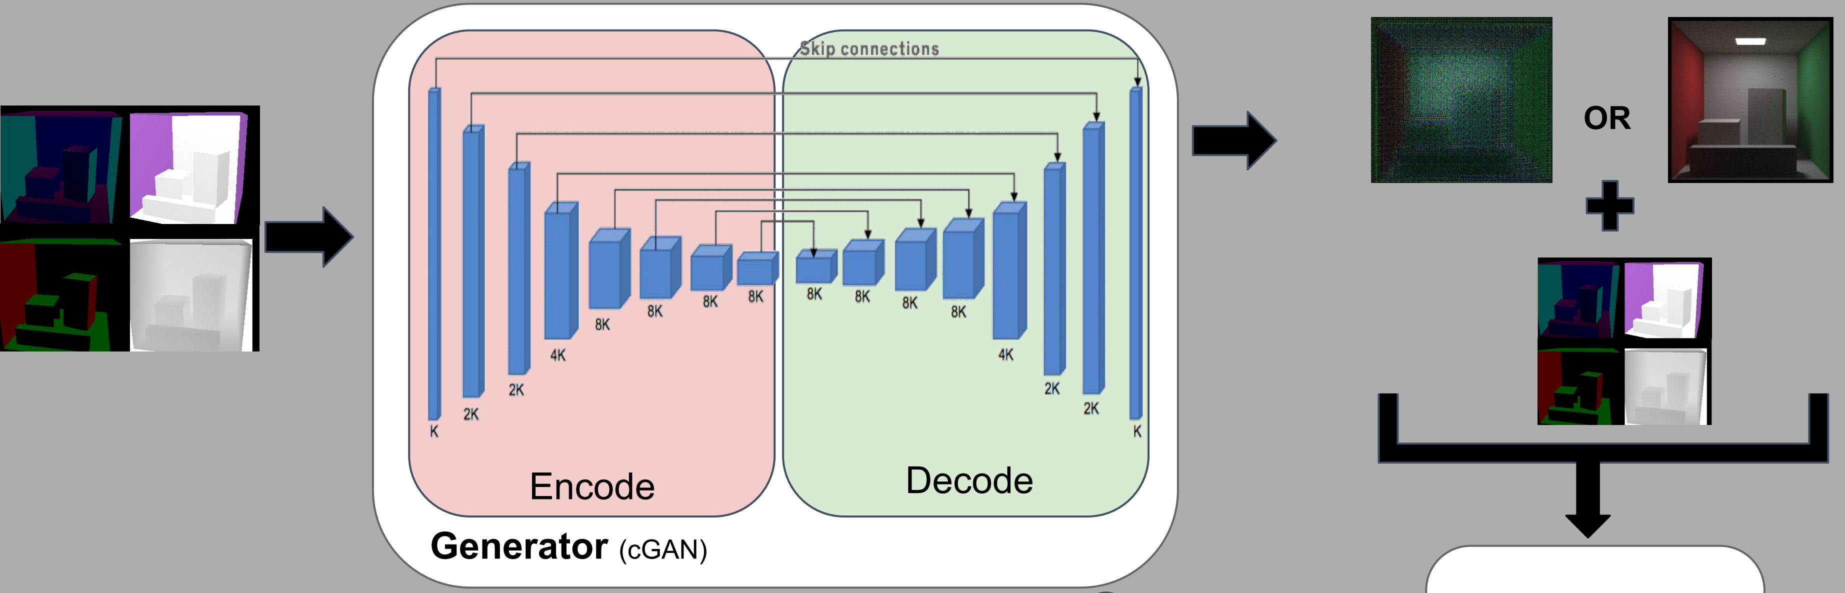

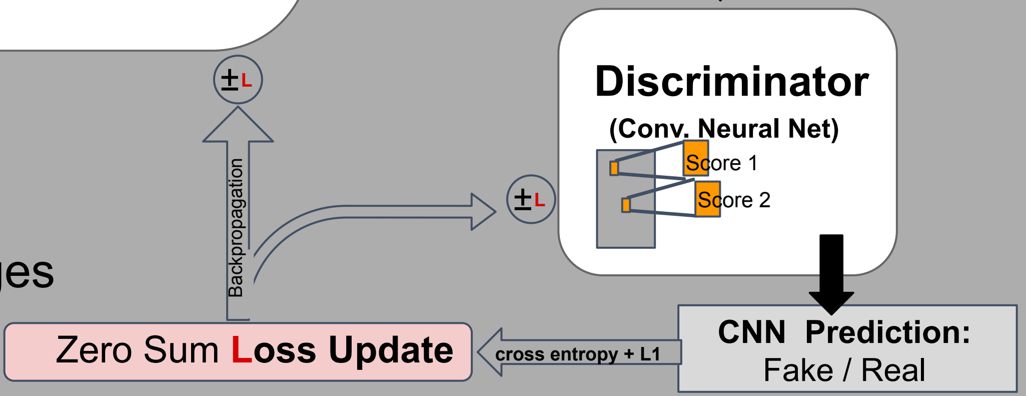

Once trained Generative model:

- Similar to an explorable database.
- Serves as a look-up table or 'filter' rendering globally illuminated images with quality comparable to offline approaches.

Example of Learned Light Transport

Light reflected from near surfaces.

Illumination dependent on orientation and proximity to light source.

Path Traced Ground Truth

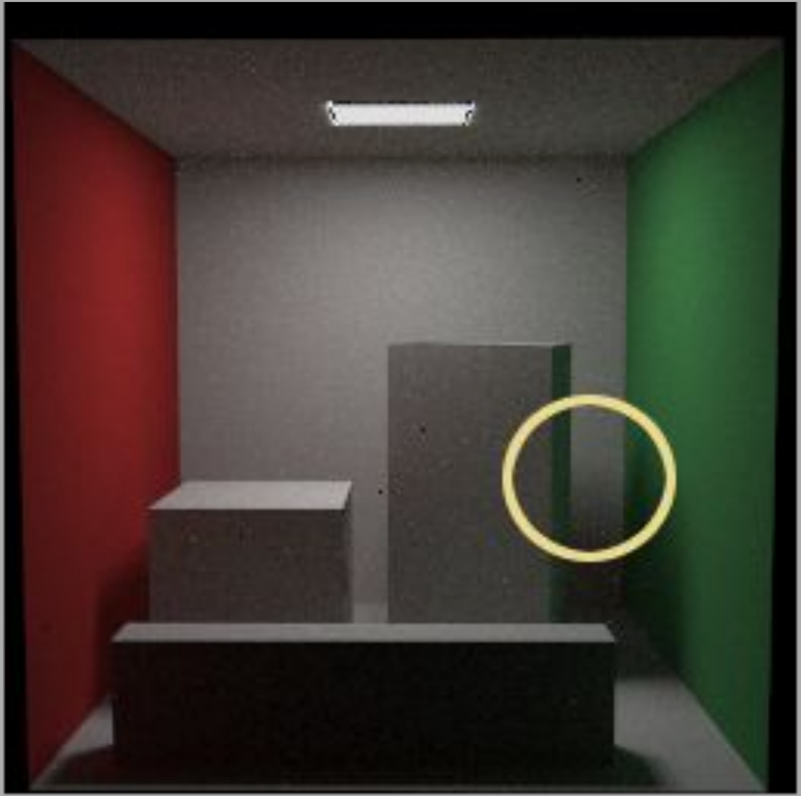

Artificial Image from Generator

Example of Learned Light Transport

Accurately approximate diffuse indirect illumination and soft shadows

Path Traced Ground Truth

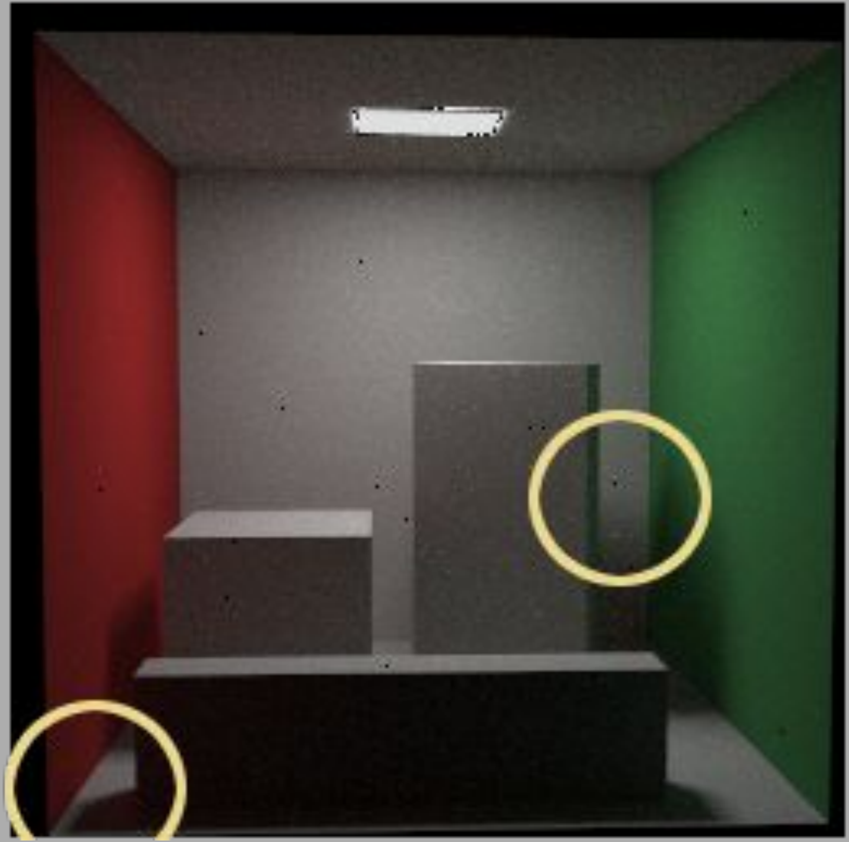

Artificial Image from Generator

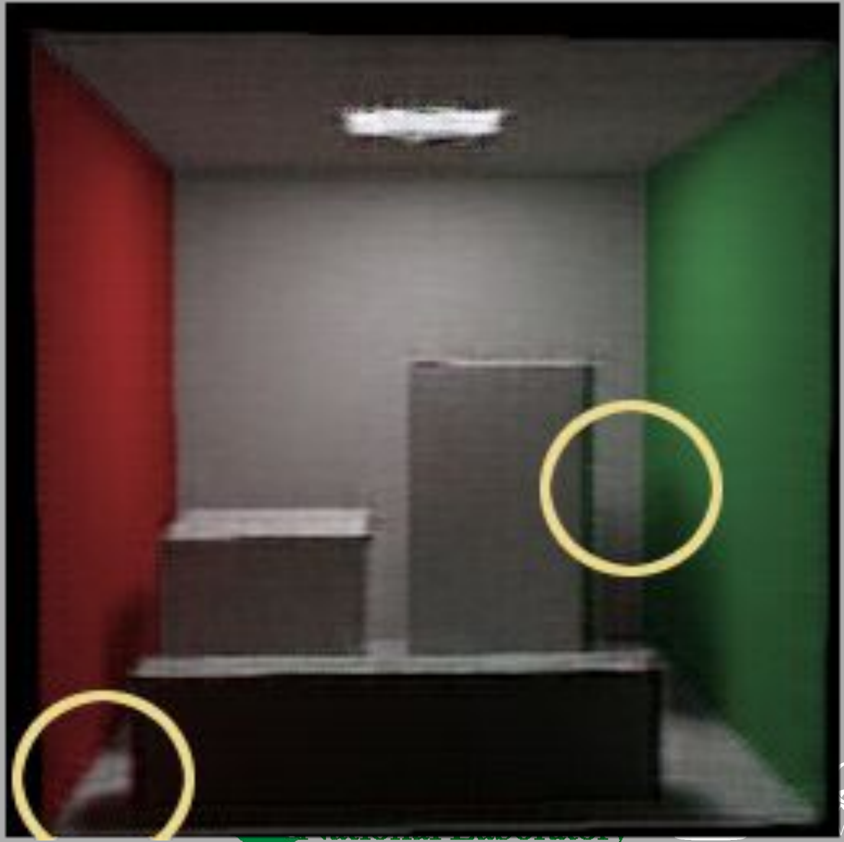

Results

As training progressed, and the generative network improved with loss converging to zero, the discriminator deteriorated and was left with near even odds amounting to a near 50-50% chance of identifying input images as either artificial or original.

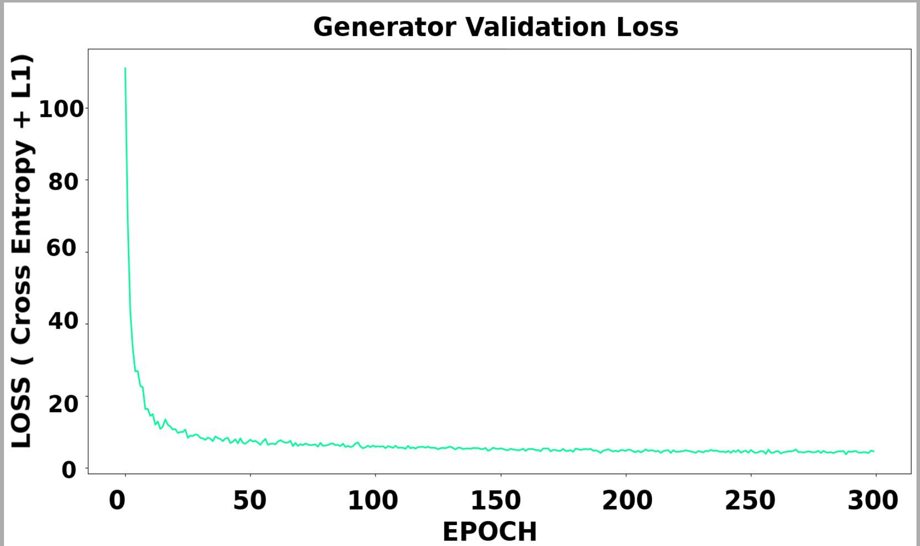

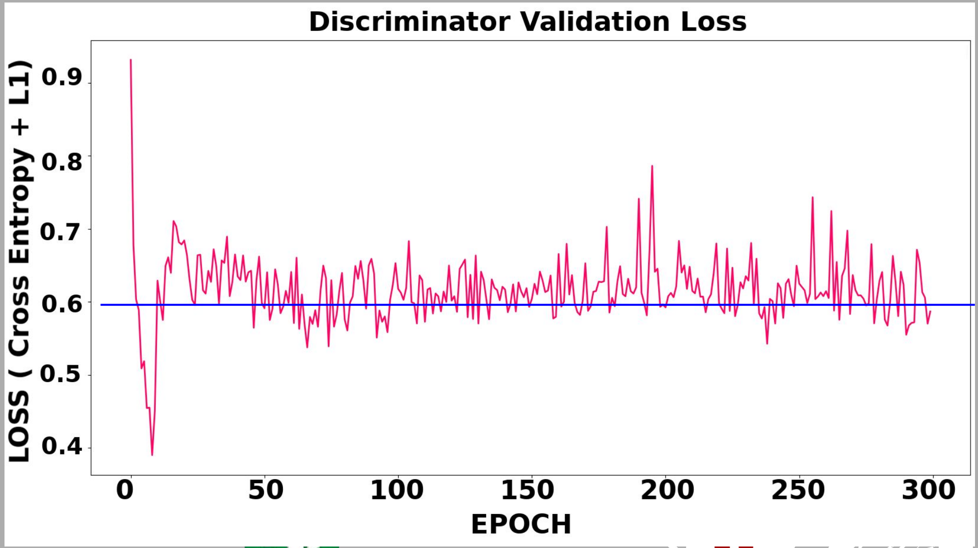

Results

As training progressed, and the generative network improved with loss converging to zero, the discriminator deteriorated and was left with near even odds amounting to a near 50-50% chance of identifying input images as either artificial or original.

Possible Issue:

- Need longer ray depth, higher sample rate.
- Discriminator learns to identify fixed point.
- To low training resolution.
- ect...

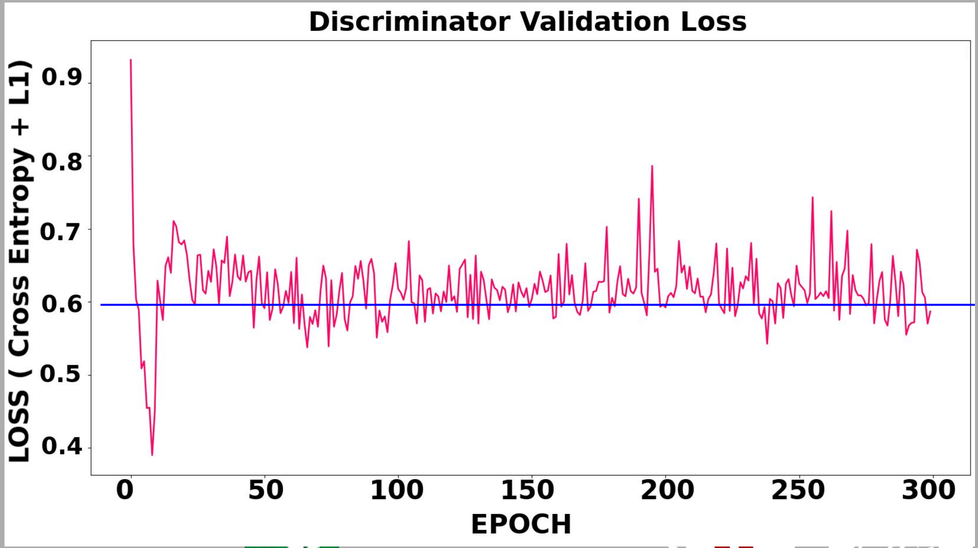

Performance

Ray Tracing: Generating 3080 image, dimension 256x256x3, ray depth 50, sampling 1000, system 2x Nvidia RTX-2080 Ti GPUs, Totaling 13 hours.

Training cGAN: On same machine with one Nvidia RTX-2080 Ti GPU took 2.7 hours over 400 epochs (with all data available).

Future Work

Implement more cyclical, interweaved applications of visualization and learning such as learning implementations on time dependent visualizations.

Larger scale in the amount to compute and data generated with more intensive coupling of tasks.

Generate novel scene or for volumes.

We would also like to add more output connectors such as Ospray or Optix, and add more machine learning frameworks such as TensorFlow.

Thank you, questions?

REFERENCES

[1] K. Moreland, "The tensions of in situ visualization," *IEEE Computer Graphics and Applications*, vol. 36, no. 2, pp. 5–9, Mar 2016.

[2] A. C. Bauer, H. Abbasi, J. Ahrens, H. Childs, B. Geveci, S. Klasky, K. Moreland, P. O'Leary, V. Vishwanath, B. Whitlock, and E. W. Bethel, "In situ methods, infrastructures, and applications on high performance computing platforms," *Computer Graphics Forum*, vol. 35, no. 3, pp. 577–597, 2016. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12930

[3] H. Abbasi, M. Wolf, G. Eisenhauer, S. Klasky, K. Schwan, and F. Zheng, "Datastager: scalable data staging services for petascale applications," *Cluster Computing*, vol. 13, no. 3, pp. 277–290, Sep 2010. [Online]. Available: https://doi.org/10.1007/s10586-010-0135-6

[4] U. Ayachit, A. Bauer, B. Geveci, P. O'Leary, K. Moreland, N. Fabian, and J. Mauldin, "Paraview catalyst: Enabling in situ data analysis and visualization," in *Proceedings of the First Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization*, ser. ISAV2015. New York, NY, USA: ACM, 2015, pp. 25–29. [Online]. Available: http://doi.acm.org/10.1145/2828612.2828624

[5] H. Childs, E. Brugger, B. Whitlock, J. Meredith, S. Ahern, D. Pugmire, K. Biagas, M. Miller, G. H. Weber, H. Krishnan, T. Fogal, A. Sanderson, C. Garth, E. W. Bethel, D. Camp, O. Rübel, M. Durant, J. Favre, and P. Navratil, "VisIt: An End-User Tool for Visualizing and Analyzing Very Large Data," in *High Performance Visualization—Enabling Extreme-Scale Scientific Insight*, ser. Chapman & Hall, CRC Computational Science, E. W. Bethel, H. Childs, and C. Hansen, Eds. Boca Raton, FL, USA: CRC Press/Francis–Taylor Group, Nov. 2012, pp. 357–372, http://www.crcpress.com/product/isbn/9781439875728, LBNL-6320E.

[6] J. Dayal, D. Bratcher, G. Eisenhauer, K. Schwan, M. Wolf, X. Zhang, H. Abbasi, S. Klasky, and N. Podhorszki, "Flexpath: Type-based publish/subscribe system for large-scale science analytics," in *2014 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, May 2014, pp. 246–255.

[7] U. Ayachit, B. Whitlock, M. Wolf, B. Loring, B. Geveci, D. Lonie, and E. W. Bethel, "The sensei generic in situ interface," in *Proceedings of the 2Nd Workshop on In Situ Infrastructures for Enabling Extreme-scale Analysis and Visualization*, ser. ISAV '16. Piscataway, NJ, USA: IEEE Press, 2016, pp. 40–44. [Online]. Available: https://doi.org/10.1109/ISAV.2016.13

[8] M. Larsen, J. Ahrens, U. Ayachit, E. Brugger, H. Childs, B. Geveci, and C. Harrison, "The alpine in situ infrastructure: Ascending from the ashes of strawman," in *Proceedings of the In Situ Infrastructures on Enabling Extreme-Scale Analysis and Visualization*, ser. ISAV'17. New York, NY, USA: ACM, 2017, pp. 42–46. [Online]. Available: http://doi.acm.org/10.1145/3144769.3144778

[9] M. Larsen, A. Woods, N. Marsaglia, A. Biswas, S. Dutta, C. Harrison, and H. Childs, "A flexible system for in situ triggers," in *Proceedings of the Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization*, ser. ISAV '18. New York, NY, USA: ACM, 2018, pp. 1–6. [Online]. Available: http://doi.acm.org/10.1145/3281464.3281468

[10] Q. Liu, J. Logan, Y. Tian, H. Abbasi, N. Podhorszki, J. Y. Choi, S. Klasky, R. Tchoua, J. Lofstead, R. Oldfield, M. Parashar, N. Samatova, K. Schwan, A. Shoshani, M. Wolf, K. Wu, and W. Yu, "Hello adios: the challenges and lessons of developing leadership class i/o frameworks," *Concurrency and Computation: Practice and Experience*, vol. 26, no. 7, pp. 1453–1473, 2014. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.3125

[11] J. Kress, M. Larsen, J. Choi, M. Kim, M. Wolf, N. Podhorszki, S. Klasky, H. Childs, and D. Pugmire, "Comparing the Efficiency of In Situ Visualization Paradigms at Scale," in *ISC High Performance*, Frankfurt, Germany, Jun. 2019, pp. 99–117.

[12] NVIDIA Corporation, *NVIDIA CUDA Compute Unified Device Architecture Programming Guide*. NVIDIA Corporation, 2007.

[13] J. Reinders, *Intel threading building blocks - outfitting C++ for multi-core processor parallelism.* O'Reilly, 2007.

[14] J. Hoberock and N. Bell, "Thrust: A parallel template library," *Thrust: A Parallel Template Library*, 2009.

[15] G. E. Blelloch, *Vector Models for Data-parallel Computing*. Cambridge, MA, USA: MIT Press, 1990.

[16] K. Moreland, C. Sewell, W. Usher, L.-t. Lo, J. Meredith, D. Pugmire, J. Kress, H. Schroots, K.-L. Ma, H. Childs *et al.*, "Vtk-m: Accelerating the visualization toolkit for massively threaded architectures," *IEEE computer graphics and applications*, vol. 36, no. 3, pp. 48–58, 2016.

[17] J. T. Kajiya, "The rendering equation," in *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '86. New York, NY, USA: ACM, 1986, pp. 143–150. [Online]. Available: http://doi.acm.org/10.1145/15922.15902

[18] A. Keller, L. Fascione, M. Fajardo, I. Georgiev, P. Christensen, J. Hanika, C. Eisenacher, and G. Nichols, "The path tracing revolution in the movie industry," in *ACM SIGGRAPH 2015 Courses*, ser. SIGGRAPH '15. New York, NY, USA: ACM, 2015, pp. 24:1–24:7. [Online]. Available: http://doi.acm.org/10.1145/2776880.2792699

[19] P.-P. Sloan, J. Kautz, and J. Snyder, "Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments," in *ACM Transactions on Graphics (TOG)*, vol. 21, no. 3. ACM, 2002, pp. 527–536.

[20] A. Kaplanyan and C. Dachsbacher, "Cascaded light propagation volumes for real-time indirect illumination," in *Proceedings of the 2010 ACM SIGGRAPH symposium on Interactive 3D Graphics and Games*. ACM, 2010, pp. 99–107.

[21] M. M. Thomas and A. G. Forbes, "Deep illumination: Approximating dynamic global illumination with generative adversarial network," *arXiv preprint arXiv:1710.09834*, 2017.

[22] M. Pharr, W. Jakob, and G. Humphreys, *Physically Based Rendering: From Theory to Implementation*, 3rd ed. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2016.

[23] M. Zwicker, W. Jarosz, J. Lehtinen, B. Moon, R. Ramamoorthi, F. Rousselle, P. Sen, C. Soler, and S.-E. Yoon, "Recent advances in adaptive sampling and reconstruction for Monte Carlo rendering," *Computer Graphics Forum (Proceedings of Eurographics - State of the Art Reports)*, vol. 34, no. 2, p. 667681, May 2015.

[24] M. Larsen, J. Meredith, P. Navratil, and H. Childs, "Ray tracing within a data parallel framework," in *2015 IEEE Pacific Visualization Symposium, PacificVis 2015 - Proceedings*, ser. IEEE Pacific Visualization Symposium, S. Takahashi, S. Liu, and G. Scheuermann, Eds. United States: IEEE Computer Society, 7 2015, pp. 279–286.

[25] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *Proc. icml*, vol. 30, no. 1, 2013, p. 3.

[26] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv preprint arXiv:1502.03167*, 2015.

[27] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.

[28] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.

[29] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2014, pp. 2672–2680. [Online]. Available: http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf

[30] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.

[31] C. M. Goral, K. E. Torrance, D. P. Greenberg, and B. Battaile, "Modeling the interaction of light between diffuse surfaces," in *Proceedings of the 11th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '84. New York, NY, USA: ACM, 1984, pp. 213–222. [Online]. Available: http://doi.acm.org/10.1145/800031.808601

Example Synthetic Conditional Buffer and Sphere

When using synthetic conditional the network showed ability to learn refraction.

Path Traced Original

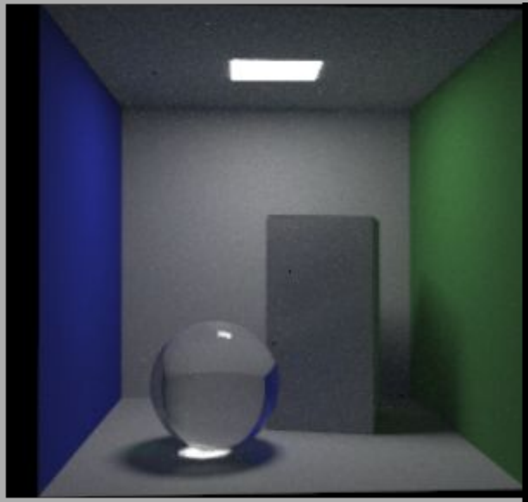

Image Augmentation
as Conditional

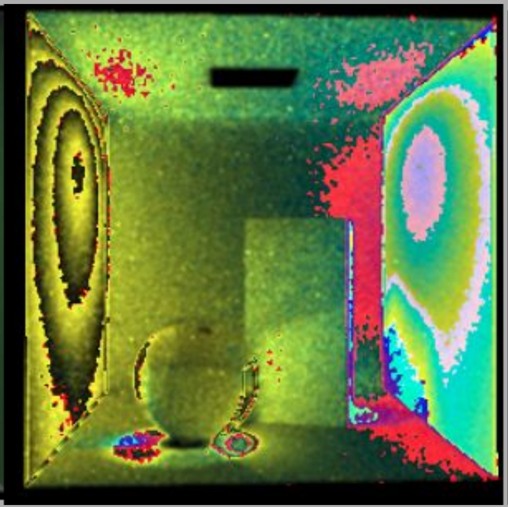

Generated Image
During Training.