OceanGAN: A Deep Learning Alternative to Physics-Based Ocean Rendering

Christopher Ratto
Applied Physics Laboratory
Johns Hopkins University
christopher.ratto@jhuapl.edu

Mimi Szeto
Applied Physics Laboratory
Johns Hopkins University
mimi.szeto@jhuapl.edu

David Slocum
Applied Physics Laboratory
Johns Hopkins University
david.slocum@jhuapl.edu

Kevin Del Bene
Applied Physics Laboratory
Johns Hopkins University
kevin.delbene@jhuapl.edu

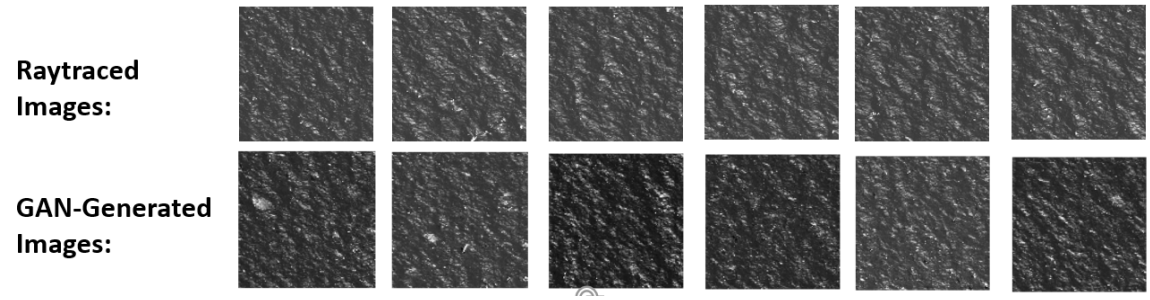

Figure 1: Panchromatic images produced by a physics-based raytracer (top) and a GAN (bottom).

KEYWORDS

ocean rendering, deep learning, generative adversarial networks

ACM Reference Format:
Christopher Ratto, Mimi Szeto, David Slocum, and Kevin Del Bene. 2019. OceanGAN: A Deep Learning Alternative to Physics-Based Ocean Rendering. In *Proceedings of SIGGRAPH '19 Posters.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3306214.3338559

1 INTRODUCTION

Physics-based models for ocean dynamics and optical raytracing are used extensively for rendering maritime scenes in computer graphics [Darles et al. 2011]. Raytracing models can provide high-fidelity representations of an ocean image with full control of the underlying environmental conditions, sensor specifications, and viewing geometry. However, the computational expense of rendering ocean scenes can be high. This work demonstrates an alternative approach to ocean raytracing via machine learning, specifically Generative Adversarial Networks (GANs) [Goodfellow et al. 2014]. In this paper, we demonstrate that a GAN trained on several thousand small scenes produced by a raytracing model can be used to generate megapixel scenes roughly an order of magnitude faster with a consistent wave spectrum and minimal processing artifacts.

SIGGRAPH '19 Posters, July 28 - August 01, 2019, Los Angeles, CA, USA

ACM ISBN 978-1-4503-6314-3/19/07.
https://doi.org/10.1145/3306214.3338559

2 BACKGROUND

2.1 Ocean Optical Raytracing

High-fidelity images of the ocean's surface are typically produced by a physics-based numerical simulation. First, the surface elevation field for wind-driven waves [Pierson, Jr. and Moskowitz 1964] is produced along with locations of breaking waves (whitecaps) and remnant foam [Callaghan et al. 2012]. Raytracing then includes the effects of ocean surface roughness [Morel et al. 2002], white water [Koepke 1984], and atmospheric effects [Berk et al. 2014]. Combining these components produces a two-dimensional map of observed radiance (W/m2/sr) at an arbitrary sensor's aperture. While physically accurate, rendering ocean scenes in this manner can be computationally expensive; we have estimated the amount of time to generate a megapixel scene (1 km^2 at 1 m resolution) to be about 20 s on a 10-node cluster.

2.2 Generative Adversarial Networks

A GAN is comprised of two deep convolutional neural networks [LeCun et al. 2015]: a generator and a discriminator. The generator network takes low-dimensional, real-valued input, and outputs high-dimensional "fake" images. The role of the discriminator is to distinguish whether an image is "real" or "fake," so it is trained on a mix of "real" training images and "fake" images produced by the generator. Learning a GAN is essentially a minimax optimization problem: maximize the performance of the generator by minimizing the performance of the discriminator, and is typically learned using stochastic gradient descent. The GAN is considered trained when the discriminator performs near chance accuracy on a set of held-out test images. From there, only the generator network needs to be used to generate fake images for the problem of interest. In this work, the problem of interest is to produce synthetic images of the ocean's surface.

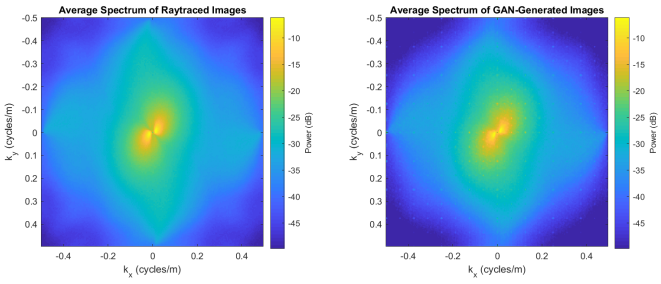

Figure 2: Comparison of the average Fourier spectra of 10,000 ocean images produced by a physics-based raytracing code (left) and a GAN (right).

3 OUR APPROACH

A GAN was implemented using the TensorFlow Python library. The discriminator had four convolutional layers (two with a kernel size of four, followed by two with a kernel size of six) followed by two fully-connected layers. The generator network began with a 64-unit fully-connected layer, followed by six transpose convolution layers (three with a kernel size of six, followed by three with a kernel size of four). The intermediate layers for both networks used leaky rectified linear (ReLU) activation and 40% dropout, while the output layers used sigmoid activation. The training set consisted of 12,150 images (256 x 256 pixels) generated via physics-based simulation at 1 m resolution for a 15 kt wind (± 2 kt) blowing 150° clockwise from North. The viewing geometry was such that the view point was at the edge of the solar glitter pattern. Examples of the training images and GAN-generated imagery are shown in Fig. 1.

Fourier analysis (Fig. 2) illustrates that the two-dimensional spectra are mostly consistent between the raytraced and GAN-generated images. The GAN successfully learned the direction of principal wind waves and the spread of energy around that spectral peak. There is a subtle checkerboard artifact due the GAN's use of 2-D filters, evidenced by a grid of dim points in the spectrum. The artifact has a strength of -15 dB near the main wind wave lobes, and -32 dB in the higher frequencies. We believe it can be mitigated by hyperparameter tuning or *post hoc* image processing.

Larger scenes may be produced as a mosaic of GAN-generated images. Figs. 3 and 4 show 8×15 mosaics of the raytraced and GAN-generated chips, respectively. Note the brightness has been enhanced to highlight details at this lower resolution. Each image produced by the GAN took about 0.17 s on a GPU, so to generate a megapixel scene via tiling would take about 2.7 s - nearly an order of magnitude improvement over raytracing. The subtle artifacts visible at full resolution will be addressed in future work. We conclude that simulating ocean images via GAN yields a reasonable tradeoff between physical accuracy and computational efficiency.

4 FUTURE WORK

Adversarial autoencoders [Makhzani et al. 2015] show promise for training GANs to produce scenes over arbitrary environments by way of a latent feature space. Conditional GANs [Isola et al. 2017] also provide a straightforward means for inserting scene elements, such as ships or clouds. We also intend to extend the GAN approach shown here to produce video with physically-realistic ocean dynamics, including wave propagation, active wave breaking, and foam decay.

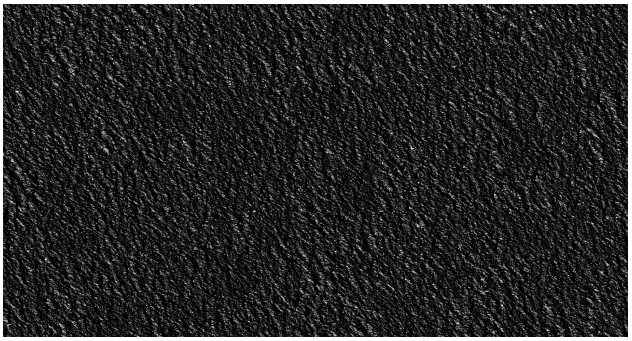

Figure 3: 8×15 mosaic of image chips produced by a physics-based raytracing code (brightness enhanced).

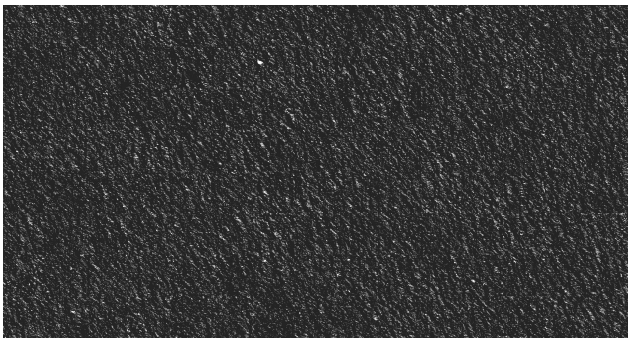

Figure 4: 8×15 mosaic of image chips produced by a GAN (brightness enhanced).

REFERENCES

A. Berk, P. Conforti, R. Kennett, T. Perkins, F. Hawes, and J. van den Bosch. 2014. MODTRAN6: a major upgrade of the MODTRAN radiative transfer code. In *Proc. SPIE 9088, Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XX*. 90880H.

A.H. Callaghan, G.B. Deane, M.D. Stokes, and B. Ward. 2012. Observed variation in the decay time of oceanic whitecap foam. *Journal of Geophysical Research* 117, 9 (September 2012).

E. Darles, B. Crespin, D. Ghazanfarpour, and J.C. Gonzato. 2011. A Survey of Ocean Simulation and Rendering Techniques in Computer Graphics. *Computer Graphics Forum* 30, 1 (2011), 43–60.

I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Processing Systems 27 (NIPS 2014)*.

P. Isola, J-Y. Zhu, and A.A. Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. In *2017 Conference on Computer Vision and Pattern Recognition (CVPR 2017)*.

P. Koepke. 1984. Effective reflectance of oceanic whitecaps. *Applied Optics* 23, 11 (1984), 1816–1824.

Y. LeCun, Y. Bengio, and G. Hinton. 2015. Deep Learning. *Nature* 521 (May 2015), 436–444.

Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. 2015. Adversarial Autoencoders. (Nov 2015). arXiv:cs.LG/1511.05644

A. Morel, D. Antoine, and B. Gentili. 2002. Bidirectional reflectance of oceanic waters: accounting for Raman emission and varying particle scattering phase function. *Applied Optics* 41, 30 (2002), 6289–6306.

W.J. Pierson, Jr. and L. Moskowitz. 1964. A proposed spectral form for fully developed wind seas based on the similarity theory of S. A. Kitaigorodski. *Journal of Geophysical Research* 69, 24 (1964), 5181.