

---

# Hallucinating Point Cloud into 3D Sculptural Object

---

Chun-Liang Li\*, Eunsu Kang\*, Songwei Ge\*, Lingyao Zhang, Austin Dill,  
Manzil Zaheer, Barnabas Poczos

Carnegie Mellon University

{chunlial, eunsuk, songweig, lingyaoz, abdill, manzil, bapoczos}@andrew.cmu.edu

\*Equal Contribution

## 1 Introduction

Many artworks generated with the assistance of machine learning algorithms have appeared in the last three years. Aesthetically impressive images have resulted that belong to the traditional art area of painting and have received global recognition. The generated art images from Elgammal et al. [7] were favored by human viewers over the images of paintings at the prestigious Miami Art Basel. CloudPainter, which was the top prize winner of the International Robot Art Competition in 2018, proved that a machine can achieve the aesthetic level of professional painters [3]. Recently a painting that was generated by GANs [8] appeared at the art auction by Christies [2]. At the same time, there has been very little exploration in the area of 3D objects, which traditionally would belong to the area of Sculpture. The creative generation of 3D Object research by Lehman et al. successfully generated “evolved” forms, however, the final form was not far from the original form and it remained in the range of mimicking the original [12]. Another relevant work, Neural 3D Mesh Renderer [11] focuses on adjusting the rendered mesh based on DeepDream [1] textures.

Our team of artists and machine learning researchers designed a creative algorithm that can generate authentic sculptural artworks. These artworks do not mimic any given forms and cannot be easily categorized into the dataset categories. Our approach extends DeepDream [1] from images to 3D point clouds. The proposed algorithm, Amalgamated DeepDream (ADD), leverages the properties of point clouds to create objects with better quality than the naive extension. ADD presents promise for the creativity of machines, the kind of creativity that pushes artists to explore novel methods or materials and to create new genres instead of creating variations of existing forms or styles within one genre. For example, from Realism to Abstract Expressionism, or to Minimalism. Lastly, we present the sculptures that are 3D printed based on the point clouds created by ADD.

## 2 Learning to Generate Creative Point Clouds

Using deep learning to generate new objects has been studied in different data types, such as music [17], images [8], 3D voxels [18] and point clouds [4, 13]. However, these models are learned to generate examples from the “same” distribution of the given training data, instead of learning to generate “creative” objects [7]. One alternative is to decode the combination of latent codes of two objects in an autoencoder. Empirical evidence shows that decoding mixed codes usually produces semantically meaningful objects with features from the corresponding objects. This approach has also been applied to image creation [6]. One could adopt encoding-decoding algorithms for point clouds [20, 9, 13] to the same idea. The sampling and interpolation results based on Li et al. [13] are shown in Figure 1.




Figure 1: Results of Li et al. [13].

**DeepDream for Point Clouds** In addition to simple generative models and interpolations, DeepDream leverages trained deep neural networks by enhancing patterns to create psychedelic and surreal images. Given a trained neural network  $f_\theta$  and an input image  $x$ , DeepDream aims to modify

$x$  to maximize (amplify)  $f_\theta(x; a)$ , where  $a$  is an activation function of  $f_\theta$ . Algorithmically, we iteratively modify  $x$  via a gradient update

$$x_t = x_{t-1} + \gamma \nabla_x f_\theta(x; a). \quad (1)$$

One special case of DeepDream is when we use a classification network as  $f_\theta$  and we replace  $a$  with the outputs of the final layer corresponding to certain labels. This is related to adversarial attack [16] and unrestricted adversarial examples [5, 14]. However, directly extending DeepDream to point clouds with neural networks for sets [15, 21] results in undesirable point clouds with sparse areas as shown in Figure 2.



Figure 2: Naive DeepDream.

## 2.1 Amalgamated DeepDream (ADD)

To avoid sparsity in the generated point cloud, one compromise is to run DeepDream with fewer iterations or a smaller learning rate, but it results in limited difference from the input point cloud. Another approach is conducting post-processing to add points to sparse areas. However, without prior knowledge of the geometry of the objects, making smooth post-processed results is non-trivial [10].

As opposed to images, mixing two point clouds is a surprisingly simple task, which can be done by taking the union of two point clouds because of their permutation invariance property. We propose a novel DeepDream variant for point clouds with a set-union operation in the creation process to address the drawbacks of the naive extension. To be specific, in each iteration after running the gradient update (1), we *amalgamate* the transformed point clouds with the input point clouds<sup>1</sup>. We call the proposed algorithm Amalgamated DeepDream (ADD) as shown in Algorithm 1. A running example of ADD targeting the transformation of bottle into cone in ModelNet40 [19], which is the same as Figure 2, is shown in Figure 3. In each iteration, ADD enjoys the advantage of deforming objects based on gradient updates with respect to a trained neural network as DeepDream without creating obviously sparse areas. In the meantime, it can better preserve the features of input point clouds with the amalgamation operation when we create new features based on DeepDream updates. More created point clouds are shown in Figure 4.

---

### Algorithm 1: Amalgamated DeepDream (ADD)

---

**input** : trained  $f_\theta$  and input  $x$ , where  $x_0 = x$   
**for**  $t = 1 \dots T$  **do**  
     $\hat{x} = x_{t-1} + \gamma \nabla_x f_\theta(x_{t-1}; a)$   
     $x_t = \hat{x} \cup x$   
**end**

---




Figure 3: Transforming a bottle into a cone via ADD.



Figure 4: ADD with single object input.




Figure 5: ADD with dual and triple object input.

In addition to using union in the transformation process in every iteration, we can push this idea further by using amalgamated point clouds as input instead of a single object. The results of ADD with the union of two or three objects are shown in Figure 5. Compared with Figure 4, ADD with multiple inputs results in objects with more geometric details benefited by a more versatile input space.

**From point clouds to realizable 3D sculptures** Our final goal is to create 3D sculptures in the real world. We use standard software MeshLab and Meshmixer to reconstruct a mesh from point clouds created by ADD. We then use Lulzbot TAZ 3D printer with dual-extruder, transparent PLA, and dissolvable supporting material to create the sculpture of the reconstructed mesh. The printed sculptures of the point cloud in Figure 5 (RHS) are shown in Figure 6.



Figure 6: Created Sculpture from ADD.

<sup>1</sup>We down-sample the point clouds periodically to avoid exploded number of points.

## References

- [1] URL <https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>.
- [2] URL <https://www.christies.com/features/A-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx>.
- [3] URL <https://robotart.org/2018-winners/>.
- [4] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. *arXiv preprint arXiv:1707.02392*, 2017.
- [5] T. B. Brown, N. Carlini, C. Zhang, C. Olsson, P. Christiano, and I. Goodfellow. Unrestricted adversarial examples. *arXiv preprint arXiv:1809.08352*, 2018.
- [6] Shan Carter and Michael Nielsen. Using artificial intelligence to augment human intelligence. *Distill*, 2017.
- [7] Ahmed Elgammal, Bingchen Liu, Mohamed Elhoseiny, and Marian Mazzone. Can: Creative adversarial networks, generating” art” by learning about styles and deviating from style norms. *arXiv preprint arXiv:1706.07068*, 2017.
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014.
- [9] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. Atlasnet: A papier-mache approach to learning 3d surface generation. *arXiv preprint arXiv:1802.05384*, 2018.
- [10] Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. *Surface reconstruction from unorganized points*. 1992.
- [11] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3907–3916, 2018.
- [12] Joel Lehman, Sebastian Risi, and Jeff Clune. Creative generation of 3d objects with deep learning and innovation engines. In *Proceedings of the 7th International Conference on Computational Creativity*, 2016.
- [13] Chun-Liang Li, Manzil Zaheer, Yang Zhang, Barnabas Poczos, and Ruslan Salakhutdinov. Point cloud gan. *arXiv preprint arXiv:1810.05795*, 2018.
- [14] Hsueh-Ti Derek Liu, Michael Tao, Chun-Liang Li, Derek Nowrouzezahrai, and Alec Jacobson. Adversarial geometry and lighting using a differentiable renderer. *arXiv preprint arXiv:1808.02651*, 2018.
- [15] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *CVPR*, 2017.
- [16] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [17] Aäron Van Den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. In *SSW*, page 125, 2016.
- [18] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *NIPS*, 2016.
- [19] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, 2015.

- [20] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *CVPR*, volume 3, 2018.
- [21] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan R Salakhutdinov, and Alexander J Smola. Deep sets. In *NIPS*, 2017.