
Learned Interpolation for 3D Generation

Austin Dill, Songwei Ge, Eunsu Kang, Chun-Liang Li
Barnabas Póczos

Carnegie Mellon University
Pittsburgh, PA, United States

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

1 Introduction

Merging abstract concepts, also known as *creative blending*, is frequently seen as a fundamental component of creativity, as this merging can allow novel concepts to emerge from simple components [10]. One computational way to approach this is *interpolation*, a process that smoothly transitions from one instance to another.

In order to generate novel 3D shapes with machine learning, one must allow for such interpolations. The typical approach for incorporating this creative process is to interpolate in a learned latent space so as to avoid the problem of generating unrealistic instances by exploiting the model’s learned structure. In 2D images, this often utilizes the trained Generative Adversarial Network [11, 4] or Autoencoder [9, 13], which has shown promising results in creative generation [5]. As for the basic requirement, the process of the interpolation is supposed to form a semantically smooth morphing [2]. While this approach is sound for synthesizing realistic media such as lifelike portraits or new designs for everyday objects, it subjectively fails to directly model the unexpected, unrealistic, or creative [3, 8].

In this work, we present a method for learning how to interpolate point clouds. By encoding prior knowledge about real-world objects, the intermediate forms are both realistic and unlike any existing forms. We show not only how this method can be used to generate "creative" point clouds, but how the method can also be leveraged to generate 3D models suitable for sculpture.

2 Interpolation as Generation




Figure 1: Learned interpolation between an airplane and a chair.

2.1 Naive Interpolation

Consider two point clouds $X_a = \{p_i\}_{i=1}^n$ and $X_b = \{p'_i\}_{i=1}^n$, where each p represent a 3-dimensional point. To generate a point cloud that semantically lies in between the inputs X_a and X_b , a straightforward method is to generate a point cloud X_{ab} as follows:

$$X_{ab} = \{\alpha p_i + (1 - \alpha)p'_i\}_{i=1}^n$$

Intuitively, this represents drawing a line between pairs of points in X_a and X_b and returning a points α percent of the distance between them. While this is simple to implement, it does not produce results that are semantically in between the objects the point clouds represent, as can be seen in Figure 2.

2.2 Learned Point Cloud Interpolation

While prior algorithms for generating creative point clouds have relied on a pretrained model, our method directly learns a transformation from point clouds to point clouds by using interpolation as the guiding framework. We parameterize interpolation from a start point cloud X_a to a goal point cloud X_b where for each time step, we apply a learned transformation for each point independently, given an encoding of X_b , denoted by $E(X_b)$.

$$p_i^{t+1} = p_i^t + h_t(p_i^t, E(X_b))$$

In each of the above transformations, the function $h_t(\cdot)$ is a multilayer perceptron [12]. All of the transformations are trained to produce a set \hat{X}_b after T transformations so that \hat{X}_b and X_b are close in Chamfer Distance, an error metric frequently used with set generation tasks [1]. The encoding network is parameterized as a Deep Sets model [14].

While this formulation allows us to visualize the trajectory of each point as it is transformed and allows us enough expressivity to approximate the goal point clouds, it does not enforce the requirement that each intermediate point cloud is realistic. For example it could allow a mapping to a completely meaningless intermediate state that would not be recognized as a plausible (if unusual) 3D object.

For this reason we introduce an additional loss term motivated by computer graphics [6].

$$\mathcal{L}(\hat{X}_b, X_b) = CD(\hat{X}_b, X_b) + \sum_i \sum_{j \in N(i)} (p_i - p_j)^2 - (\phi(p_i) - \phi(p_j))^2$$

With this added term, we are able to maintain the topology of the beginning object, causing the network to find the most plausible correspondence between the source object and the target function. This loss function only penalizes the output of the algorithm but has the side effect of ensuring each intermediate step is topologically consistent as well.

2.3 Mesh Generation




Figure 3: Generated meshes from interpolating between pairs of objects.

This technique allows one to use the vertices from a mesh as input, providing us with the correspondence needed to mesh the output. This removes the problem of meshing for creative sculpture generation, the motivating factor for creative sculpture generating algorithms [7]. In the author’s opinion, the conflict between the representational advantage of point clouds for machine learning tools and the artist’s frequent need for solid 3D shapes has limited the adoption of generative models for sculptural art. Our method therefore represents an important step forward for creative AI.




Figure 2: Naive interpolation fails to produce an interesting midpoint.

References

- [1] P. Achlioptas, O. Diamanti, I. Mitliagkas, and L. Guibas. Learning representations and generative models for 3d point clouds. *arXiv preprint arXiv:1707.02392*, 2017.
- [2] D. Berthelot, C. Raffel, A. Roy, and I. Goodfellow. Understanding and improving interpolation in autoencoders via an adversarial regularizer. *arXiv preprint arXiv:1807.07543*, 2018.
- [3] A. Bidgoli and P. Veloso. Deepcloud. the application of a data-driven, generative model in design, 2019.
- [4] A. Brock, J. Donahue, and K. Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- [5] S. Carter and M. Nielsen. Using artificial intelligence to augment human intelligence. *Distill*, 2(12):e9, 2017.
- [6] D. Ezuz, J. Solomon, and M. Ben-Chen. Reversible harmonic maps between discrete surfaces. *arXiv preprint arXiv:1801.02453*, 2018.
- [7] S. Ge, A. Dill, E. Kang, C.-L. Li, L. Zhang, M. Zaheer, and B. Poczoz. Developing creative ai to generate sculptural objects, 2019.
- [8] D. P. Kingma and P. Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10215–10224, 2018.
- [9] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [10] F. C. Pereira. *Creativity and artificial intelligence: a conceptual blending approach*, volume 4. Walter de Gruyter, 2007.
- [11] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [12] F. Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.
- [13] I. Tolstikhin, O. Bousquet, S. Gelly, and B. Schoelkopf. Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558*, 2017.
- [14] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Poczoz, R. R. Salakhutdinov, and A. J. Smola. Deep sets. In *Advances in neural information processing systems*, pages 3391–3401, 2017.