MVPNet: Multi-View Point Regression Networks for 3D Object Reconstruction from A Single Image

Jinglu Wang (MSRA) Bo Sun (PKU) Yan Lu (MSRA)

Microsoft Research 微软亚洲研究院

Motivation

Why multi-view representation? Existing representations for 3D reconstruction are mainly of four categories. While 3D volumetric grids suffer high computational complexity, unordered point sets require to solve point-wise mapping, meshes are difficult for CNNs to encode and decoder, multi-view based representation are convolution-favored, ordered and can depict dense and detailed surface.

Overcome limitations of multi-view representation. View based representation handing features in 2D projective space can neglect information loss through dimension reduction from 3D to 2D. The proposed MVPC constructs meshes in 3D from 2D grids. MVPC allows us to discretize integrals of surface variations over the constructed triangular mesh and to enforce multi-view consistency with view correspondences.

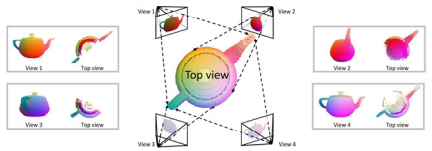

Multi-View Point Clouds:

Overview

We reconstruct an object's surface from a single image by 1) representing a 3D surface as MVPC, 2) regressing MVPC with MVPNet, and 3) proposing a geometric loss to interpret discrepancy over 3D surfaces.

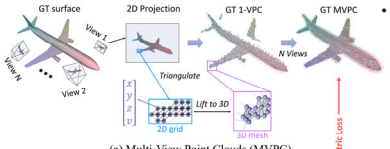

(a) Multi-View Point Clouds (MVPC)

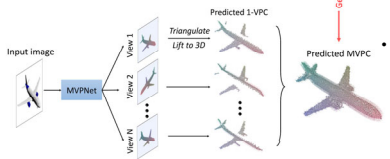

(b) Multi-View Point Network (MVPNet)

- **Multi-View Point Clouds:** A surface is represented by Multi-View Point Cloud (MVPC). Each pixel in a 1-VPC stores the backprojected surface point *(x, y, z)* from this pixel and its visibility *v*. The stored 3D points are triangulated according to the 2D grid on the image plane and their normals are shown to indicate surface orientation.
- **Muti-View Point Network.** Given an RGB image, the MVPNet generates a set of 1-VPCs and their union forms the predicted MVPC. The geometric loss measures discrepancy between predicted and ground truth MVPC.

Results

We present qualitative and quantitative results of the reconstruction, comparing the proposed method two collections of state-of-the-art methods according to result representations, i.e., point clouds and volumetric grids.

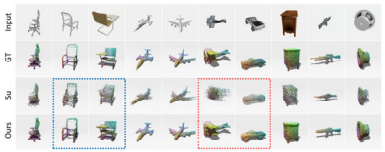

Qualitative comparison to point generation method [1]

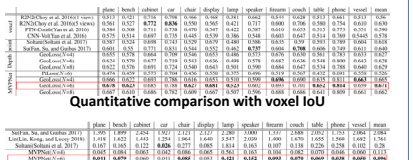

Quantitative comparison with voxel IoU

Quantitative comparison with chamfer distance metric

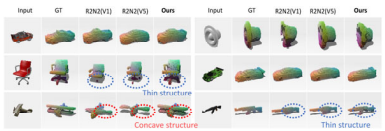

Qualitative compare to voxel-based method [2]

Reconstruction results on real word data

Network Architecture

The MVPNet is an encoder-decoder generative network incorporating camera parameters to generate view-dependent point clouds.

- The **encoder** maps an image *I* to an embedding space to obtain a feature *Z*. Each camera matrix c_i is first mapped to a feature c_i, serving as a view indicator, and then is concatenated with *Z* to get (Z, c_i).
- The **decoder** converting (Z, c_i) to a 1-VPC M_i indicated by c_i learns the projective transformation and space completion. The decoder shares weights among *N* branches.

The output MVPC is of shape *N* x *H* x *W* x 4. The last channel corresponds to a 3D coordinate $\boldsymbol{x}_k = (x_k, y_k, z_k)$ and visibility v^i_k of a point $\boldsymbol{x}_k$.

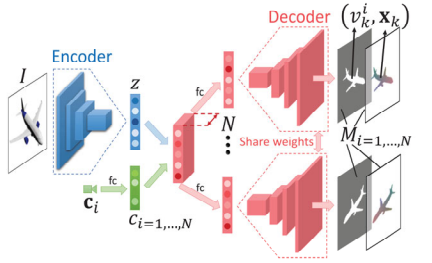

Geometric Loss

We propose a geometric loss (GeoLoss) that is able to capture variances over 3D surfaces rather than over sparse point sets or 2D projective planes. The GeoLoss is made up of three components:

$$\mathcal{L}_{Geo} = \mathcal{L}_{ptd} + \alpha\mathcal{L}_{vol} + \beta\mathcal{L}_{mv}$$

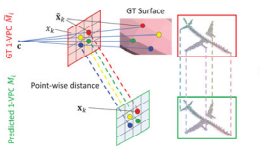

Point-wise distance term. The points in ground truth and predicted 1-VPC have a one-to-one mapping, since 2D pixels with equal 2D coordinates are defined to store the same surface point induced by the same viewpoint. The sum of point-wise distances for ground truth and predicted 1-VPC is the L2 loss.

$$\mathcal{L}_{ptd} = \sum_i^N \sum_{x \in M_i} ||M_i(x) - \widetilde{M}_i(x)||_2$$

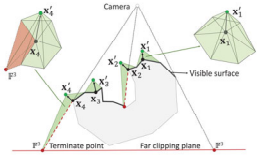

Quasi-volume term. Inspired by the volume-preserving constraints used in variational surface deformation, we propose a quasi-volume discrepancy metric to describe the surface discrepancy, characterizing details and handling occluding contours.

$$\mathcal{L}_{vol}(S, \widetilde{S}) = \int_S (\mathbf{x} - \widetilde{\mathbf{x}}) \cdot \mathbf{n} \mathrm{d}\mathbf{x}$$

It is discretized in the MVPC's mesh as:

$$\mathcal{L}_{vol} = \sum_i^N \sum_{x \in M_i} \widetilde{V}_i(x)(M_i(x) - \widetilde{M}_i(x)) \cdot \widetilde{\mathbf{N}}_i(x)$$

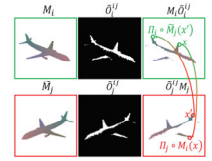

Multi-view consistency term. Some 1-VPCs may have overlap. Corresponding points should be close. We minimize sum of two distances between stored 3D points in two corresponding pixels and their reprojected pixels in another 1-VPC.

$$\mathcal{L}_{mv} = \sum_{i,j} \Big(\sum_{x \in \bar{O}^{ij}_i} ||M_i(x) - \widetilde{M}_j(\Pi_j \circ M_i(x))||_2 + \sum_{x \in \bar{O}^{ij}_j} ||\widetilde{M}_j(x) - M_i(\Pi_i \circ \widetilde{M}_j(x))||_2 \Big)$$

Application

We show the generative representation of the learned features with interpolation, arithmetic, classification and clustering.

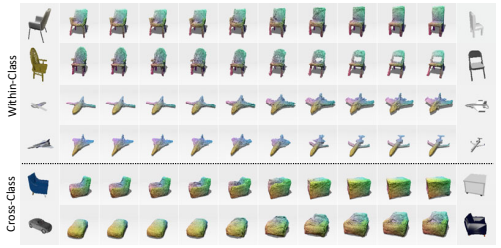

Interpolation. Reconstructions from linear interpolation of two learned features.

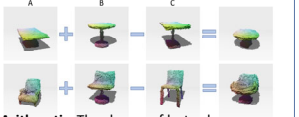

Arithmetic. The shapes of last column are obtained by decoding the fature (A+B-C).

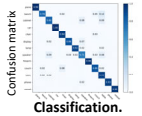

Classification.

Clustering.

Reference

[1] Fan, H.; Su, H.; and Guibas, L. 2017. A point set generation network for 3d object reconstruction from a single image. In ICCV.

[2] Choy, C. B.; Xu, D.; Gwak, J.; Chen, K.; and Savarese, S. 2016. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In ECCV.

Website