MIT Open Access Articles

Visualization and Labeling of Point Clouds in Virtual Reality

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Jonathan Dyssel Stets, et al. "Visualization and Labeling of Point Clouds in Virtual Reality." Proceedings of SIGGRAPH Asia 2017, 27-30 November 2017, Bangkok, Thailand, ACM, 2017.

As Published: <https://dx.doi.org/10.1109/34.121791>

Publisher: Association for Computing Machinery

Persistent URL: <http://hdl.handle.net/1721.1/117394>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike

Visualization and Labeling of Point Clouds in Virtual Reality

Jonathan Dyssel Stets
Technical University of Denmark

Wiley Corning
MIT Media Lab

Yongbin Sun
Massachusetts Institute of Technology

Scott Greenwald
MIT Media Lab

Figure 1: Screen capture from the Virtual Reality application. A point cloud of jugglers is displayed, the left and right images shows the point cloud before and after the juggling pins has been annotated by the user.

ACM Reference Format:

Jonathan Dyssel Stets, Yongbin Sun, Wiley Corning, and Scott Greenwald. 2017. Visualization and Labeling of Point Clouds in Virtual Reality. In *Proceedings of SIGGRAPH Asia 2017 Posters, Bangkok, Thailand, November 2017*, 2 pages. https://doi.org/00.0005/123_4

1 INTRODUCTION

We present a Virtual Reality (VR) application for labeling and handling point cloud data sets. A series of room-scale point clouds are recorded as a video sequence using a Microsoft Kinect. The data can be played and paused, and frames can be skipped just like in a video player. The user can walk around and inspect the data while it is playing or paused. Using the tracked hand-held controller, the user can select and label individual parts of the point cloud. The points are highlighted with a color when they are labeled. With a tracking algorithm, the labeled points can be tracked from frame to frame to ease the labeling process. Our sample data is an RGB point cloud recording of two people juggling with pins. Here, the user can select and label, for example, the juggler pins as shown in Figure 1. Each juggler pin is labeled with various colors to indicate different labels.

2 MOTIVATION

We consider the use case of viewing animated point cloud data in VR. Although individual frames appear grainy and low-resolution, when they are animated, they produce a lifelike effect that is suitable for representing dynamic real-world scenes. However, without any labeling or metadata, it is difficult to edit or modify these scenes. Labeling a person with a name, or changing the color of an object would be examples of basic editing tasks. Segmenting the point clouds in each frame, and establishing how the segments in each frame correspond with those in surrounding frames, makes it possible to perform such tasks.

Furthermore, the same set of user interface affordances can be used for other use cases as well. The demand for large labeled point cloud data sets is increasing as the need for classification algorithms for such data is gaining popularity, especially for data-driven approaches like deep neural networks. Navigation of self driving cars, social robots and robot interaction in general often rely on 3D data that needs to be classified or segmented. Additionally, the use of 3D sensors is increasing, while these sensors are getting more mobile, better and more available. This enables the potential for more and better 3D point cloud data. The website Semantic3D [Semantic3d 2016] provides examples of large labeled 3D point cloud data sets. Oftentimes, such 3D data sets are labeled manually using a 2D computer monitor and a computer mouse or similar input device. We argue that a two-dimensional human-computer interaction with three-dimensional data can be inefficient in many cases. For example, the labeling problem is more difficult with a 2D interface when the orthogonal xy , xz , or yz planes do not separate the points to be labeled from the others surrounding them.

Figure 2: The juggling pin is being labeled with a blue label using the hand-held controller. Haptic feedback in the controller is provided when points are selected.

We present an application that demonstrates an intuitive way to inspect and label 3D point cloud data. The user simply inspects the data by moving around in the 3D point cloud, as one would do in the real world. Labeling is done by pointing and drawing in the 3D space with a hand-held controller to mark relevant parts of the data.

3 DETAILS

We have recorded a data set using a Microsoft Kinect 2 and a custom-built software tool. The software records and converts the depth map and rgb video stream from the Kinect to a point cloud and stores the (X, Y, Z, R, G, B) data in a binary format. The point clouds are recorded and stored approximately 30 times per seconds. The set of point clouds is loaded into the VR application and visualized at room scale. We use the HTC Vive setup for VR, which enables tracking of both headset and the hand-held controllers. Using the touchpad on the hand-held controller, the user can navigate back and forward in the point cloud frames. A marker-tool is attached to the controller and this can be used like a paint brush to label points in the point cloud (see Figure 2). Haptic feedback is provided when selecting the points to add the feeling of physically interacting with the data. Individual parts of the point cloud can be labeled with different labels, each represented by varying colors. The user can switch between labels on the controller touchpad, as well as choosing an eraser-tool to unlabel points in the point cloud. We argue that interacting with 3-dimensional data in this way is an easy and fast alternative to a traditional 2-dimensional interface such as a mouse or a touch screen.

To ease the labeling across frames, we have added a tracking feature to automatically label subsequent frames in the point cloud set. The tracking algorithm fits labeled points in $frame_i$ to a portion of points in $frame_{i+1}$, so that we can predict labels from one frame to the next. Figure 3 shows how three individual parts of the point cloud has been labeled, and then tracked from one frame to the next. We implemented two algorithms to fit labeled points. The first one is an Iterative Closest Point (ICP) algorithm [Besl et al. 1992], which

Figure 3: Three frames from the point cloud data. The juggling pins are labeled individually (red, green and blue) by the user in the first frame. The tracking algorithm is used to track points from one frame and automatically label them in the next.

iteratively detects the closest point in $frame_{i+1}$ labeled point in $frame_i$ and estimates the current best rotation and translation to update the positions of the labeled points. This iterative process terminates until a predetermined stop criteria satisfies, for example, the spatial difference between labeled points and matched points converges. The second algorithm follows the same pipeline as the first one, but is augmented with color support to improve tracking accuracy. Instead of detecting the spatially closest point, the second algorithm selects, for each query point, the point with the most similar color from k -nearest neighboring points as the match. We build a k -d tree [Brown 2014], which can be pre-calculated for each point cloud, to expedite the nearest neighbor search.

Our method is not limited to Kinect data, thus data of varying scale and resolution can be visualized and labeled. We also emphasize the potential of this application for modifying point clouds, such as deleting noisy points, or cutting out relevant parts of the data using the same type of labeling and selection method.

4 CONCLUSION

We have realized an application for visualizing and labeling point clouds in Virtual Reality. The application demonstrates that it is possible to label and inspect large sets of room-scale point clouds in a fast and intuitive way. With the application, it is possible to select and track parts of the data to speed up labeling in sequences of point clouds. We believe that is a step in the direction towards incorporating point clouds into an efficient production pipeline for realistic animated virtual reality content, and more high precision labeled point cloud data sets for a variety of other applications as well.

5 ACKNOWLEDGEMENTS

Thanks to the MIT Juggling Club for allowing us to record the data.

REFERENCES

- Paul J Besl, Neil D McKay, et al. 1992. A method for registration of 3-D shapes. *IEEE Transactions on pattern analysis and machine intelligence* 14, 2 (1992), 239–256.
- Russell A Brown. 2014. Building a balanced kd tree in $o(kn \log n)$ time. *arXiv preprint arXiv:1410.5420* (2014).
- Semantic3d. 2016. Large-Scale Point Cloud Classification Benchmark. (2016). <http://www.semantic3d.net/>