

Cloud Dictionary: Coding and Modeling for Point Clouds

Or Litany*
orlitany@gmail.com

Tal Remez*
talremez@gmail.com

Alex Bronstein
bron@eng.tau.ac.il

Tel-Aviv University
* Equal contributors

Abstract—With the development of range sensors such as LIDAR and time-of-flight cameras, 3D point cloud scans have become ubiquitous in computer vision applications, the most prominent ones being gesture recognition and autonomous driving. Parsimony-based algorithms have shown great success on images and videos where data points are sampled on a regular Cartesian grid. We propose an adaptation of these techniques to irregularly sampled signals by using continuous dictionaries. We present an example application in the form of point cloud denoising.

I. INTRODUCTION

Recent advances in sensors introduce a wide family of data which do not lie on a regular grid. This can be either due to the lack of a temporal grid as in the event camera introduced in [1], or a spatial one as is the case of point clouds produced, e.g., by LIDAR scanning the world in a varying angular velocity. These data serve as the input to a vast range of computer vision application such as 3D scene reconstruction [2], or understanding [3]. Sparse coding techniques have been shown to achieve state-of-the-art performance in many tasks for regularly sampled data such as images and videos. However, these methods do not apply directly to unstructured data like point clouds. It would therefore be beneficial to bridge the gap and allow the use of these techniques for irregularly sampled signals. To this end, we suggest a generalization of patch-based sparse representation and dictionary learning techniques for irregularly sampled data, and show an application for point clouds denoising.

II. SPARSE CODING FOR POINT CLOUDS

Patch-based processing of images and videos has been shown to be more computationally efficient and to produce better results compared to global models. Unfortunately, the definition of patches on a 3D shape is non-trivial. In [4], it was proposed to define a patch straightforwardly by finding all points inside a ball with Euclidean or, better, geodesic radius r , centered around a point. A similar approach was also used by [5], [6]. A plane is subsequently fitted into this set of points using PCA, onto which the points are projected to yield a local system of coordinates (u, v, w) , with w_i denoting the normal displacement of point i from the plane. The main gap left in order to be able to use standard sparse coding algorithms is that the point locations $\mathbf{u} = (u, v)$ are not situated on a regular grid. The method in [4] bridges this gap by a nearest-neighbors interpolation of the w values on a predefined regular grid. However, this adds computational burden and introduces unnecessary sources of error. We choose, instead, to handle the data more naturally as sampled points of a latent continuous function. To this end, we model the points of each patch as $\mathbf{y}_{\mathbf{G}} = w(\mathbf{G}) + n(\mathbf{G})$, where the signal w and the noise n are now continuous signals sampled at locations $\mathbf{G} = \{\mathbf{u}_1, \dots, \mathbf{u}_p\}$. In this notation, denoising boils down to a pursuit problem of the form

$$\arg \min_{\mathbf{z}} \frac{1}{2} \|\mathbf{y}_{\mathbf{G}} - \mathbf{D}(\mathbf{G})\mathbf{z}\|_2^2 \text{ s.t. } \|\mathbf{z}\|_0 \leq L, \quad (1)$$

where L is a parameter constraining the number of non-zero elements in the representation, and, with some abuse of notation, $\mathbf{D}(\mathbf{G})\mathbf{z} = d_1(\mathbf{G})z_1 + \dots + d_M(\mathbf{G})z_M$ denotes the synthesis of a continuous signal from a discrete combination of continuous dictionary atoms $d_i(\mathbf{u})$. This formulation eliminates the necessity of a fixed grid for all signals. In other words, given a set of samples $\mathbf{y}_{\mathbf{G}}$ and a continuous dictionary \mathbf{D} , we can sample its continuous atoms at the locations \mathbf{G} , and proceed with a pursuit algorithm of our choice. Once the code \mathbf{z} is obtained we can get an estimate of the continuous signal $\hat{w}(\mathbf{u}) = (d_1(\mathbf{u}), \dots, d_M(\mathbf{u}))\mathbf{z}$.

III. DICTIONARY LEARNING

So far we have assumed to be given a continuous dictionary. Such a dictionary can be either axiomatic (e.g., cosine or sine functions) or learned. Specifically, research on images showed a significant gain in performance for dictionaries that are learned from the data. In what follows, we propose a continuous dictionary learning algorithm.

Our generalized formulation of (1) enables the adaptation of dictionary learning techniques, such as k -SVD, to grid-less data in the following way: We first fix a set of continuous functions $\phi^T = \{\phi_1, \phi_2, \dots, \phi_N\}$. Each atom in the dictionary \mathbf{D} can be defined as a linear combination of these functions,

$$d_m(\mathbf{u}) = \phi(\mathbf{u})^T \mathbf{a}_m = \sum_{i=1}^N \phi_i(\mathbf{u}) \mathbf{A}_{im}; \quad \mathbf{A} \in \mathbb{R}^{N \times M}. \quad (2)$$

Let \mathbf{G}_i be the grid of the i -th training patch, and let $\mathbf{y}_i = y_i(\mathbf{G}_i)$ be the corresponding patch values. When solving the pursuit problem (1), notice that since each data point is constructed from an arbitrary number of sample points \mathbf{G}_i , the data vectors \mathbf{y}_i vary in size, yet are represented with the same number of coefficients in the dictionary $\mathbf{D}(\mathbf{G}_i)$. Following the k -SVD algorithm in [7], we initialize the dictionary using a random set of coefficients \mathbf{A} . At each step we begin by finding a representation \mathbf{z}_i for each of the training data vectors \mathbf{y}_i by solving (1). During the dictionary-update stage, for each atom $m = 1, \dots, M$

- 1) Find the examples that use the atom d_m . $\Lambda_m = \{i : \mathbf{z}_{im} \neq 0\}$
- 2) Minimize the residual, restricted to the subset Λ_m
- 3) Repeat the above steps until the representation error of all examples is below some threshold

Code is available at github.com/orlitany/Cloud_dictionary.

IV. EXPERIMENTS

We demonstrate the proposed method on the canonical example of denoising. Given a noisy point cloud (Figure 2), we constructed a continuous dictionary composed of cosine functions with varying spatial frequencies and applied the proposed dictionary learning technique. Figure IV shows the original and the learned dictionaries. The learning process error rate is presented in Figure IV. We obtain the denoised point cloud (Figure 2) by a convex relaxation of equation 1, and by averaging reconstructions from overlapping patches.

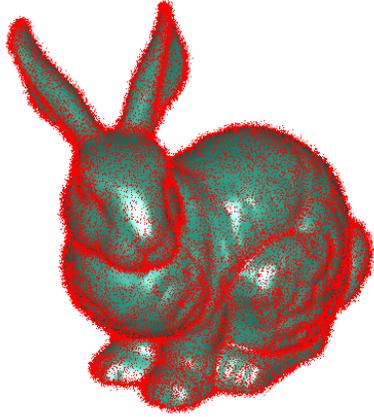


Fig. 1. **Noisy point cloud of Stanford’s bunny [8].** The original bunny mesh is presented in green and the noisy points cloud contaminated with Gaussian noise is presented in red.

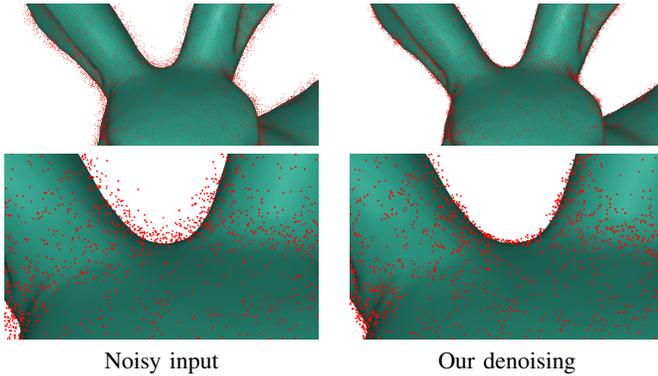


Fig. 2. **Point cloud denoising using a continuous dictionary.** The images present a zoomed fraction of the bunny, where the ground truth mesh is presented in green and the point cloud in red. On the left is the noisy point cloud, and on the right is the denoised point cloud using our denoising algorithm. These results clearly show that our denoising algorithm returns a cleaner point cloud that is closer to the ground truth mesh.

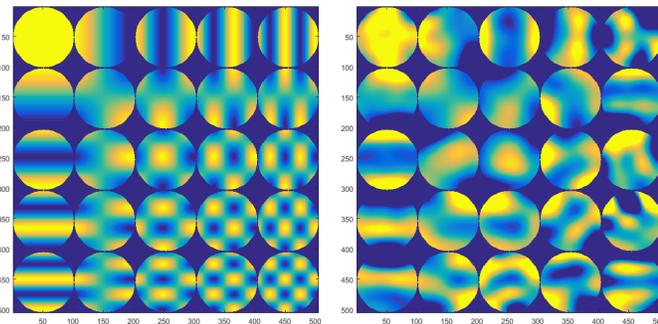


Fig. 3. **Continuous dictionary learning.** On the left is the original continuous dictionary atoms before the dictionary learning process, constructed using a Cartesian product of cosines. On the right is the learned continuous dictionary. It can be seen that the learned dictionary is smoother and introduces new structures at different rotation angles. The error per dictionary learning iteration is shown in Figure IV.

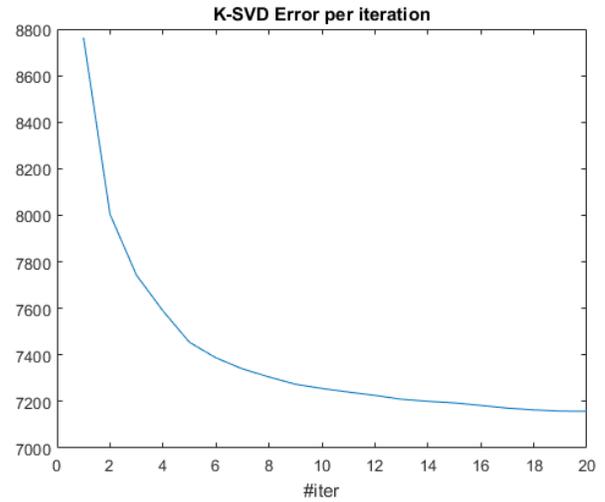


Fig. 4. **Continues k -SVD learning error Vs iteration.** Presented is the error of synthesizing the Bunny using the learned continuous dictionary after each k -SVD iteration.

REFERENCES

- [1] C. Brandli, R. Berner, M. Yang, S.-C. Liu, and T. Delbruck, “A 240×180 130 db 3 μ s latency global shutter spatiotemporal vision sensor,” *Solid-State Circuits, IEEE Journal of*, vol. 49, no. 10, pp. 2333–2341, 2014.
- [2] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, “Kinectfusion: Real-time dense surface mapping and tracking,” in *Mixed and augmented reality (ISMAR), 2011 10th IEEE international symposium on*. IEEE, 2011, pp. 127–136.
- [3] O. Litany, T. Remez, D. Freedman, L. Shapira, A. Bronstein, and R. Gal, “Asist: Automatic semantically invariant scene transformation,” *Computer Vision and Image Understanding*, pp. –, 2016. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1077314216301102>
- [4] J. Digné, R. Chaine, and S. Valette, “Self-similarity for accurate compression of point sampled surfaces,” in *Computer Graphics Forum*, vol. 33, no. 2. Wiley Online Library, 2014, pp. 155–164.
- [5] A. Elmoataz, O. Lezoray, and S. Boughleux, “Nonlocal discrete regularization on weighted graphs: a framework for image and manifold processing,” *Image Processing, IEEE Transactions on*, vol. 17, no. 7, pp. 1047–1060, 2008.
- [6] F. Lozes, A. Elmoataz, and O. Lézoray, “Pde-based graph signal processing for 3-d color point clouds: Opportunities for cultural heritage,” *Signal Processing Magazine, IEEE*, vol. 32, no. 4, pp. 103–111, 2015.
- [7] M. Aharon, M. Elad, and A. Bruckstein, “ k -SVD: an algorithm for designing overcomplete dictionaries for sparse representation,” *IEEE Trans. Sig. Proc.*, vol. 54, no. 11, pp. 4311–4322, 2006.
- [8] G. Turk and M. Levoy, “The stanford bunny,” 2005.