cilantro: A Lean, Versatile, and Efficient Library for Point Cloud Data Processing

Konstantinos Zampogiannis
University of Maryland
Computer Science Department
College Park, Maryland
kzampog@cs.umd.edu

Cornelia Fermüller
University of Maryland
Inst. for Advanced Computer Studies
College Park, Maryland
fer@umiacs.umd.edu

Yiannis Aloimonos
University of Maryland
Computer Science Department
College Park, Maryland
yiannis@cs.umd.edu

ABSTRACT

We introduce *cilantro*, an open-source C++ library for geometric and general-purpose point cloud data processing. The library provides functionality that covers low-level point cloud operations, spatial reasoning, various methods for point cloud segmentation and generic data clustering, flexible algorithms for robust or local geometric alignment, model fitting, as well as powerful visualization tools. To accommodate all kinds of workflows, *cilantro* is almost fully templated, and most of its generic algorithms operate in arbitrary data dimension. At the same time, the library is easy to use and highly expressive, promoting a clean and concise coding style. *cilantro* is highly optimized, has a minimal set of external dependencies, and supports rapid development of performant point cloud processing software in a wide variety of contexts.

Availability: the project source code, with usage examples and sample data, can be found at <https://github.com/kzampog/cilantro>.

KEYWORDS

point cloud processing; geometric registration; spatial reasoning; clustering; model fitting

ACM Reference Format:

Konstantinos Zampogiannis, Cornelia Fermüller, and Yiannis Aloimonos. 2018. *cilantro: A Lean, Versatile, and Efficient Library for Point Cloud Data Processing*. In *2018 ACM Multimedia Conference (MM '18), October 22–26, 2018, Seoul, Republic of Korea*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3240508.3243655>

1 INTRODUCTION

Processing geometric input plays a crucial role in a number of machine perception scenarios. Robots use stereo cameras or depth sensors to create 3D models of their environment and/or the objects with which they interact. Autonomous mobile agents are typically equipped with LiDAR sensors to map their surroundings, localize themselves, and avoid collisions. Consumer electronics are increasingly adopting the integration of depth cameras to identify users and enable “natural” user interfaces. The output signals of these sensors are either inherently or directly convertible to 2D or 3D

point clouds, highlighting the need for usable and efficient tools for processing raw geometric data.

Thanks to the central role of geometry in the fields of computational theory and computer graphics, a number of notable 3D data processing open-source software libraries have been developed. The Computational Geometry Algorithms Library (CGAL) [16] implements algorithms and data structures for an extensive set of tasks, including triangulations, shape analysis, meshing, and various geometric decompositions. The Visualization and Computer Graphics (VCG) library [17] provides a collection of tools for processing and visualizing 3D meshes and constitutes the back-end of the popular MeshLab [4] GUI-based mesh editor. To simplify working with mesh data, libigl [8] drops complex data structures in favor of raw data matrices and vectors for shape representations, and supports computation of discrete differential quantities and operators, shape deformation, and remeshing functionalities. While mature and feature-rich in their respective contexts, these software packages have had limited adoption by the machine perception and engineering communities.

The Point Cloud Library (PCL) [12] was introduced to fill this gap and became the standard for unorganized point cloud processing among roboticists and machine vision practitioners. It implements numerous algorithms for filtering, feature extraction, geometric registration, reconstruction, segmentation, and model fitting. Partly due to its templated nature, PCL exhibits a steep learning curve; its user-friendliness is further affected by its verbose coding style and typically long application compilation times. Furthermore, its performance in very common tasks is lacking by today’s standards (e.g., due to lack of parallelization in many of its modules) and the project has been in an essentially dormant state for a long time. The Open3D [19] library was recently introduced as a potential alternative. It implements 3D point cloud primitive operations (neighbor queries, downsampling, surface normal estimation), as well as useful tools for geometric registration and 3D reconstruction. The library is easy to use and performs better than PCL in common tasks. However, its supported functionality is quite limited, as it essentially only targets 3D reconstruction workflows.

In this paper, we introduce *cilantro*, a versatile, easy to use, and efficient C++ library for generic point cloud data processing. We have implemented a concise set of algorithms that cover primitive point cloud operations, spatial reasoning based on convex polytopes, various methods for point cloud segmentation and generic data clustering, flexible algorithms for both robust and local (iterative) geometric alignment, model fitting, as well as powerful visualization tools. The library is written in C++11, is highly templated and optimized, and makes efficient use of multi-threading for

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '18, October 22–26, 2018, Seoul, Republic of Korea

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5665-7/18/10...\$15.00

<https://doi.org/10.1145/3240508.3243655>

computationally demanding tasks. Significant effort has been put into making every component customizable by the user, so that the library does not hinder the development of “non-standard” workflows. At the same time, we provide useful convenience functions and aliases for the most common algorithm variants, ensuring that our adaptable design does not get in the way of productivity.

By virtue of its flexibility, ease of use, and comprehensive out-of-the-box functionality, *cilantro* is a great fit for a wide spectrum of workflows, enabling the rapid development of performant point cloud processing software. In the following, we briefly state our main design principles (Section 2), give a more detailed overview of our supported functionality (Section 3), and demonstrate *cilantro*’s performance in some ordinary tasks in comparison to equivalent implementations in PCL and Open3D (Section 4).

2 DESIGN OVERVIEW

We outline *cilantro*’s main design principles and how they are reflected on the library implementation.

Simple data representations. We rely on Eigen [7] matrices and common STL containers to represent point sets and most other entities used/generated by our algorithms. In particular, the input to almost all of our algorithms is a *matrix* view of a point set: a lightweight wrapper, based on the Eigen::Map class template. Our `ConstDataMatrixMap` universal input template is parameterized by the scalar numerical type (typically `float` or `double`) and the point dimensionality (integer value), while dynamic (runtime) dimensionality settings are also supported. The resulting matrix view contains one data point per column. `ConstDataMatrixMap` is constructible from many common point set representations, such as Eigen matrix variants, STL vectors of fixed-size vector/array objects, raw data pointers, etc. The only requirement is that the mapped data should be contiguously stored in memory, in column-major order. This mechanism is transparent to the user; as a result, users can in most cases directly pass their data to *cilantro* functions, without the need for type casts or data copying.

Generic algorithms. Almost all of the algorithms implemented in *cilantro* operate in arbitrary dimension, according to the nature of the input data. Furthermore, data dimensionality can be either known at compile time or determined and set dynamically. To accommodate this feature, the library is fully templated, with the exception of visualization and some of the I/O functions. At the most basic level, template parameters typically include the input data numerical type and dimensionality. More complex algorithms are adaptable to the user’s needs, by being parameterized by the entities responsible for the simpler tasks involved. For example, our ICP base class only implements an interface for the top-level EM-style iterations of the algorithm and is parameterized by a correspondence search mechanism and an estimator entity. This way, we can easily generate ICP variants for different metrics (e.g., point-to-point or point-to-plane) and different correspondence search engines (e.g., *kd-tree* based or projective), while maintaining a common, general interface. For performance reasons, we have opted to implement our “modular” high-level algorithms as *static* interfaces, by making use of the Curiously Recurring Template Pattern (CRTP) idiom.

Ease of use. Generality should not come at the cost of usability. For this reason, we tackle template parameter verbosity by providing type aliases and convenience functions for the most standard variants of our algorithms. We also increase *cilantro*’s expressiveness by enabling method chaining in almost all of our classes. Furthermore, the library code is maintained in a clean, consistent style, with class and function names that are descriptive of the underlying functionality. The following code snippet showcases the library usage in a very simple pipeline that involves reading a 3D point cloud from a file, downsampling it by means of a voxel grid, estimating its surface normals, and saving the result to a file:

```
#include <cilantro/point_cloud.hpp>

int main(int argc, char ** argv) {
    cilantro::PointCloud3f(argv[1]).gridDownsample(0.005f).
        estimateNormalsRadius(0.02f).toPLYFile(argv[2]);
    return 0;
}
```

The code is very concise, with minimal boilerplate. Readers familiar with PCL will reckon that an equivalent PCL implementation would require a significantly larger amount of code.

Performance. The library is highly optimized and makes use of OpenMP parallelization for computationally demanding operations. Significant effort has been put into benchmarking alternative implementations of common operations in order to find the fastest variant and/or parallelization pattern. As a result, *cilantro* exhibits the lowest running times in the benchmarks of Section 4.

Minimal dependencies. *cilantro* was built upon a carefully selected set of third party libraries and has minimal external dependencies. The library comes bundled with: *nanoflann* [3] for fast *kd-tree* queries, *Spectra* [11], an ARPACK-inspired library for large scale eigendecompositions, *Qhull* [1] for convex hull and half-space intersection computations, and *tinyply* [6] for PLY format geometry I/O. The only external dependencies are:

- Eigen [7], an elegant and efficient linear algebra library on which we rely for most of our numerical operations, and
- Pangolin [9], a lightweight OpenGL viewport manager and image/video I/O abstraction library, on which our visualization modules were built.

cilantro uses the CMake build system and can be compiled with all major C++ toolchains (GCC, Clang).

3 FUNCTIONALITY

For most of our supported functionality, we have adopted an object-oriented approach, implementing each of our algorithms as a class template, while also providing free functions for simpler operations. The library components can be conceptually divided in the following categories.

Core operations. We support a wide set of primitive point cloud operations by implementing:

- An optimized, user-friendly *KDTree* template that supports general dimension *k*-NN, radius, and hybrid queries, under all of the distance metrics supported by *nanoflann*.
- Arbitrary dimension surface normal and curvature estimation.
- Principal Component Analysis (PCA).
- A generic, optimized *GridAccumulator* template and its appropriate instantiations for general dimension grid-based point cloud downsampling.

- I/O functions for general matrices and 3D point clouds.
- Utility functions for RGBD to/from 3D point cloud conversions.

All the above algorithms operate on “raw”, `ConstDataMatrixMap`-wrapped data. To facilitate common workflows, we also provide a convenience `PointCloud` class template that encapsulates point coordinates, normals, and colors, and provides basic point selection functionality, as well as a large number of helper methods for normal estimation, downsampling, and geometric transformations. In the 3D case, `PointCloud` instances are directly constructible from PLY files and RGBD image pairs.

Spatial reasoning. Building on `Qhull`’s facilities and a simple feasibility solver, we provide a `ConvexPolytope` template that is constructed as either the convex hull of an input point set or the halfspace intersection defined by an input set of linear inequality constraints, enabling seamless representation switching. Our `SpaceRegion` class represents arbitrary space regions as unions of convex polytopes. Both space representations implement set operations (intersection for `ConvexPolytope`, all common ones for `SpaceRegion`), interior point tests, volume calculation, and can be transformed geometrically.

Clustering/segmentation. `cilantro` provides four standard clustering algorithms: a parallel, optimized k -means implementation that can optionally use kd -trees for large numbers of clusters, three common spectral clustering variants [18] that rely on Spectra, an arbitrary kernel mean-shift implementation, and a parallelized, generic connected component segmentation algorithm that supports arbitrary point-wise similarity functions. We note that common segmentation tasks such as extracting Euclidean (see PCL’s `EuclideanClusterExtraction`) or smooth (see PCL’s `RegionGrowing`) segments can be straightforwardly cast as connected component segmentation under different similarity metrics.

Iterative geometric registration. We provide a CRTP base class template that implements the top-level loop logic (alternating between correspondence estimation and transformation parameter optimization) of the Iterative Closest Point (ICP) algorithm. The base template is parameterized by the correspondence estimation mechanism and the transform estimator. We implement a standard, kd -tree based correspondence search engine, which can operate on arbitrary point feature spaces, as well as one for correspondences by projective data association for the 3D case. We provide optimizers for both *rigid* [2, 10] and *non-rigid* [13, 15] (by means of a robustly regularized, locally rigid warp field) pairwise registration under the the point-to-point and point-to-plane metrics, as well as weighted combinations thereof. Our modular design can be easily extended to accommodate less common ICP processes (e.g., for multi-way registration). At the same time, `cilantro` provides useful shortcut wrappers/definitions for the more common registration workflows.

Robust estimation. The library comes with a RANSAC template that is meant to be used as a CRTP base class and only implements the top-level (random sampling/inlier evaluation) loop logic. We currently provide two general dimension RANSAC estimator instances: one for robust (hyper)plane fitting and one for rigid pairwise point cloud alignment given (noisy) point correspondences.

Visualization. `cilantro` implements a fully customizable, interactive, and easy to use 3D Visualizer class that supports: an extensible set of renderable entities with adjustable rendering options, complete control over all projection parameters, both perspective

and orthographic projection modes, image capturing and video recording of the viewport, and custom keyboard callbacks. We also provide a convenience `ImageViewer` class. Both our visualization facilities are based on Pangolin. To facilitate the visualization of data of potentially unknown dimension, for which we are only given pairwise distances, we have included a classical Multidimensional Scaling (MDS) implementation that can be used to compute low dimensional, distance-preserving Euclidean embeddings.

Figure 1: Point cloud segmentation into smooth segments. Black points correspond to very small clusters that were discarded.

4 PERFORMANCE

We compare `cilantro`’s performance in common 3D point cloud processing tasks against PCL and Open3D. Our benchmarking was done on a standard Linux (Ubuntu 16.04 based) desktop machine, equipped with an Intel® Core i7-6700K CPU (4 cores, 8 threads). All three libraries were compiled from source in ‘Release’ configuration, with compiler (GCC 5) optimizations set to the highest level (-O3).

Our first test involves the segmentation of an apartment-scale 3D reconstruction into smooth segments using PCL’s `RegionGrowing` and `cilantro`’s `ConnectedComponentSegmentation`. We use the `apt0` model (Fig. 1, left) as input, available on the BundleFusion [5] project website, which consists of roughly 7.8 million points. After downsampling the input using a voxel grid of bin size equal to 5mm, using k NN neighborhoods with $k = 30$ and a normal angle threshold of 2.8 radians, `cilantro` produces the result of Fig. 1 (right) in 3.65 seconds, while, for the same parameters, the PCL implementation took 17.34 seconds. Open3D offers no equivalent/similar functionality.

PCL’s standard point types use single precision floating point numbers, while Open3D works only with double precision. For all subsequent computations, we use double precision for `cilantro`.

In our next test, we compare our normal estimation performance against equivalent implementations in Open3D (parallelized by default) and PCL (the `NormalEstimationOMP` variant). We use the `fr1/desk` sequence of the TUM RGBD dataset [14] as input (about 600 pre-registered RGBD image pairs at VGA resolution). We run all three implementations on all input frames at full resolution, using two types of local neighborhoods: one defined by the 10 nearest neighbors and one defined by a radius of 1cm. As can be seen in the two top rows of Fig. 2, `cilantro` consistently outperforms the other two implementations. In the k NN case in particular, it exhibits

speedups of $1.58\times$ and $1.85\times$ over Open3D and PCL, respectively. We note that *cilantro* and PCL enforce viewpoint-consistent normals during estimation, while in Open3D this has to be done in a separate pass, which was not performed in our experiments.

Figure 2: Performance comparisons against PCL and Open3D in common operations. Left column: running time as a function of the input video frame. Right column: box and whisker plots of running times per library (red lines are means, green lines are medians).

In our last test, we compare point-to-plane ICP performance among the three libraries. We used the same *fr1/desk* sequence as before, aligning each frame to its previous one. All point clouds were downsampled using a voxel grid of bin size equal to 1cm, and normals were estimated based on 10 nearest neighbor local neighborhoods. To account for different termination criteria conventions, we forced all three implementations to perform a fixed number of 15 iterations. In all three cases, correspondences were established by means of nearest neighbor *kd*-tree queries, with a rejection distance threshold of 5cm. Registration times are reported in the third row of Fig. 2. We note that, at all times, all three implementations converged to the same transformation, up to numerical precision. *cilantro*'s ICP implementation outperforms the other two by a significant margin, demonstrating speedups of $3.82\times$ and $14.99\times$ over Open3D and PCL, respectively.

5 SOFTWARE RELEASE

cilantro is released under the MIT license and its source code is openly available at <https://github.com/kzampog/cilantro>. Contributions from the open-source community are welcome, via the GitHub issues/pull request mechanisms.

The library is under active, continuous development, undergoing frequent API improvements, functionality additions, and performance optimizations. Future additions will include structures and algorithms for surface merging/reconstruction, more point cloud resampling strategies, and support for geometry I/O in other file formats. We are also investigating GPU implementations for some of our algorithms, focusing on improving our non-rigid registration performance.

ACKNOWLEDGMENTS

The authors would like to thank Aleksandrs Ecins and Kanishka Ganguly for their code contributions and support, and Moschoula Pternea for generating the plots in this report. The support of ONR under grant award N00014-17-1-2622 and the support of the National Science Foundation under grants SMA 1540916 and CNS 1544787 are greatly acknowledged.

REFERENCES

- [1] C Bradford Barber, David P Dobkin, and Hannu Huhdanpaa. 1996. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software (TOMS)* 22, 4 (1996), 469–483.
- [2] Paul J Besl and Neil D McKay. 1992. A Method for registration of 3-D shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures*, Vol. 1611. International Society for Optics and Photonics, 586–607.
- [3] Jose Luis Blanco and Pranjal Kumar Rai. 2014. nanoflann: a C++ header-only fork of FLANN, a library for Nearest Neighbor (NN) with KD-trees. <https://github.com/jlblancoc/nanoflann>.
- [4] Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. 2008. Meshlab: an open-source mesh processing tool. In *Eurographics Italian Chapter Conference*, Vol. 2008. 129–136.
- [5] Angela Dai, Matthias Nießner, Michael Zollöfer, Shahram Izadi, and Christian Theobalt. 2017. BundleFusion: Real-time Globally Consistent 3D Reconstruction using On-the-fly Surface Re-integration. *ACM Transactions on Graphics 2017 (TOG)* (2017).
- [6] Dimitri Diakopoulos. 2018. tinyply. <https://github.com/ddiakopoulos/tinyply>.
- [7] Gaël Guennebaud, Benoit Jacob, et al. 2010. Eigen v3. <http://eigen.tuxfamily.org>.
- [8] Alec Jacobson, Daniele Panozzo, et al. 2017. libigl: A simple C++ geometry processing library. <http://libigl.github.io/libigl/>.
- [9] Steven Lovegrove. 2018. Pangolin. <https://github.com/stevenlovegrove/Pangolin>.
- [10] Kok-Lim Low. 2004. Linear least-squares optimization for point-to-plane icp surface registration. *Chapel Hill, University of North Carolina* 4 (2004).
- [11] Yixuan Qiu. 2018. Spectra: A header-only C++ library for large scale eigenvalue problems. <https://spectralib.org/index.html>.
- [12] Radu Bogdan Rusu and Steve Cousins. 2011. 3D is here: Point Cloud Library (PCL). In *IEEE International Conference on Robotics and Automation (ICRA)*. Shanghai, China.
- [13] Olga Sorkine and Marc Alexa. 2007. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, Vol. 4. 109–116.
- [14] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. 2012. A Benchmark for the Evaluation of RGB-D SLAM Systems. In *Proc. of the International Conference on Intelligent Robot Systems (IROS)*.
- [15] Robert W Sumner, Johannes Schmid, and Mark Pauly. 2007. Embedded deformation for shape manipulation. In *ACM Transactions on Graphics (TOG)*, Vol. 26. ACM, 80.
- [16] The CGAL Project. 2018. *CGAL User and Reference Manual* (4.12 ed.). CGAL Editorial Board. <https://doc.cgal.org/4.12/Manual/packages.html>
- [17] Visual Computing Lab of the Italian National Research Council - ISTI. 2018. The VCG Library. <http://www.vcglib.net/>.
- [18] Ulrike Von Luxburg. 2007. A tutorial on spectral clustering. *Statistics and computing* 17, 4 (2007), 395–416.
- [19] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. 2018. Open3D: A Modern Library for 3D Data Processing. *arXiv:1801.09847* (2018).