IMAGE LICENSED  
BY GRAPHIC STOCK

# Registration with the Point Cloud Library

## *A Modular Framework for Aligning in 3-D*

By Dirk Holz, Alexandru-Eugen Ichim,  
Federico Tombari, Radu B. Rusu, and Sven Behnke

**R**egistration is an important step when processing three-dimensional (3-D) point clouds. Applications for registration range from object modeling and tracking, to simultaneous localization and mapping (SLAM). This article presents the open-source point cloud library (PCL) and the tools available for point cloud registration. The PCL incorporates methods for the initial alignment of point clouds using a variety of local shape feature descriptors, as well as methods for refining initial alignments using different variants of the well-known iterative closest point (ICP) algorithm. This article provides an overview on registration algorithms, usage examples of their PCL implementations, and tips for their application. Since the choice and parameterization of the right algorithm for a particular type of data is one of the biggest problems in 3-D point cloud registration, we present three complete examples

of data (and applications) and the respective registration pipeline in the PCL. These examples include dense red–green–blue–depth (RGB-D) point clouds acquired by consumer color and depth cameras, high-resolution laser scans from commercial 3-D scanners, and low-resolution sparse point clouds captured by a custom lightweight 3-D scanner on a microaerial vehicle (MAV).

### Registration of 3-D Point Clouds

The problem of consistently aligning two or more point clouds, i.e., sets of 3-D points, is inherent for 3-D registration. Often the point clouds are acquired by 3-D sensors from different viewpoints. The registration finds the relative pose (position and orientation) between views in a global coordinate frame, such that the overlapping areas between the point clouds match as well as possible; for two examples of registration see Figure 1. The overall objective of registration is to align individual point clouds and fuse them to a single point cloud so that subsequent

processing steps, such as segmentation and object reconstruction, can be applied.

Point cloud registration is a recurring task for a number of applications related to computer vision, computer graphics, robotic perception, photogrammetry, cultural heritage modeling, digital archaeology, and architecture. The research over the past years has been aimed at obtaining accurate alignments for challenging 3-D data, such as isotropic or nondistinctive objects, and at making the registration process as automatic as possible.

Publicly available software libraries and data sets are a way to promote such research and to make results comparable. This article presents the open source PCL and the tools available for point cloud registration. PCL is a stand-alone, large-scale, open project for 3-D point cloud processing. It is released under the BSD license and contains numerous state-of-the-art algorithms for various applications, including filtering, feature estimation, surface reconstruction, model fitting, segmentation, and visualization, as well as higher level tools and applications for performing mapping and object recognition.

In this article, we focus on the pairwise registration of 3-D point clouds, i.e., aligning two point clouds by estimating the transformation between the two view poses under which the point clouds have been acquired. Techniques and tools for tackling sequences of point clouds in a globally coherent fashion (i.e., multiview registration) will be covered in future publications.


Pairwise registration is usually carried out by means of one of the several variants of the ICP algorithm. Due to the non-convexity of the optimization, ICP-based approaches require initialization with a rough initial transformation to increase the chance of ending up with a successful alignment. Good initialization also speeds up their convergence. This article will cover all stages required for the point cloud alignment—from coarse alignment based on descriptor matching to the ICP-based refinement—illustrating the tools currently available in the PCL to carry out each task. Furthermore, we present three pipelines implemented in the PCL together with typical results for different types of input data, ranging from small-scale object modeling to outdoor 3-D mapping.

## General/Basic Registration Pipelines

### Problem Formulation

A point cloud is a data structure  $P$  used to represent a collection of multidimensional points  $\mathbf{p} \in \mathbb{R}^n$ . In a 3-D point cloud, the elements usually represent the  $X$ ,  $Y$ , and  $Z$  geometric coordinates of an underlying sampled surface. When more information is available, e.g., color information or information about a local surface normal  $\mathbf{n}$  or curvature  $\kappa$ , the points  $\mathbf{p} \in P$  are represented by a longer vector.

Given a source point cloud  $P$  with points  $\mathbf{p} \in P$  and a target point cloud  $Q$  with points  $\mathbf{q} \in Q$ , the problem of registration relies on finding correspondences between  $P$  and  $Q$  and estimating a transformation  $T$  that, when applied to  $P$ , aligns



**Figure 1.** Examples of 3-D registration in different scales and applications. (a) A registered pair of high-resolution 3-D laser scans (green and blue) in a city. (Data recorded by Dorit Borrmann and Jan Elseberg, Robotic 3-D Scan Repository: <http://kos.informatik.uni-osnabrueck.de/3Dscans>.) (b) A registered RGB-D image sequence of toy cars on a table.

all pairs of corresponding points ( $\mathbf{p}_i \in P, \mathbf{q}_j \in Q$ ). One fundamental problem of registration is that these correspondences are usually not known and need to be determined by the registration algorithm. Given correct correspondences, there are different ways of computing the optimal transformation with respect to the used error metric.

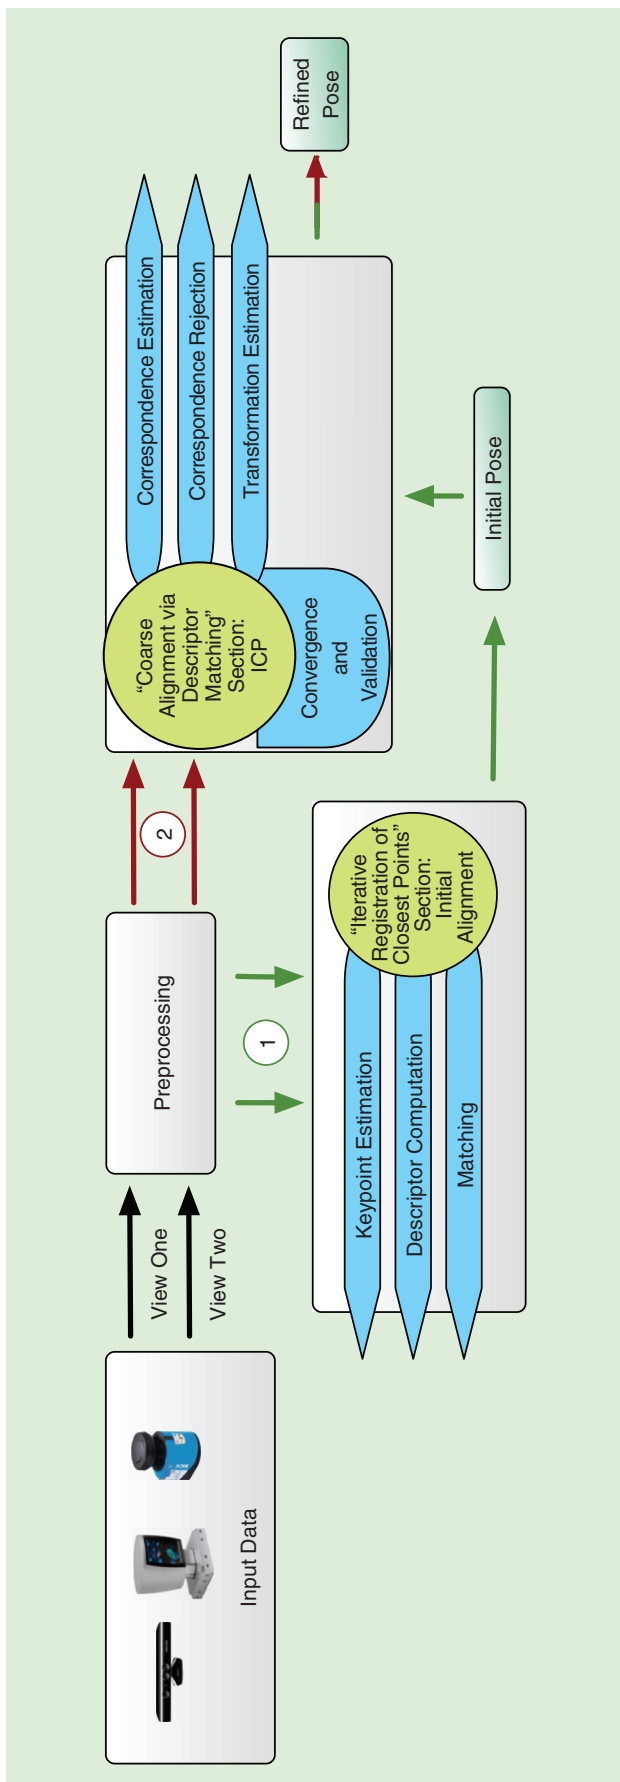
To support processing various types of point clouds in the PCL, all components are templated. The remainder of this article will use code snippets in the style of Listing 1.

### Listing 1. Example code listing and terminology.

```
// specifying types of input and output data
Algorithm<PointT, NormalT, FeatureT> alg;

// providing input data
alg.setInputCloud (input_cloud);
alg.setInputNormals (input_normals);

// computing output
PointCloud<FeatureT> output_features;
alg.compute(&output_features);
```



**Figure 2.** An overview of the process of registering a pair of point clouds. After the two views are in memory, a preprocessing sequence is applied to improve the signal-to-noise ratio or to recover information about the surface (normals) that is lost because of the sampling involved in the scanning procedure. There are two available paths. The first path consists of computing keypoints and their descriptors and then estimating a relative transformation between the two clouds. This can be used as a good initial alignment for the dense registration via ICP, represented by the second path.

### Typical Registration Pipelines

Referring to Figure 2 and following the modularization of Rusinkiewicz and Levoy [1], the registration of two point clouds can be split into the following steps:

- *Selection*: Sampling of the input point clouds.
- *Matching*: Estimating the correspondences between the points in the subsampled point clouds.
- *Rejection*: Filtering the correspondences to reduce the number of outliers.
- *Alignment*: Assigning an error metric, and minimizing it to find the optimal transformation.

Each of these four stages will be now detailed, with a specific reference to the tools available in the PCL.

### Matching Closest Points Versus Feature-Based Registration

Two major classes of registration algorithms can be distinguished (see the two paths in Figure 2):

- 1) feature-based registration algorithms (path 1) for computing initial alignments
- 2) iterative registration algorithms (path 2) following the principle of the ICP algorithm to iteratively register point clouds (that are already roughly aligned).

For the feature-based registration, geometric feature descriptors are computed and matched in some high-dimensional space. The more descriptive, unique, and persistent these descriptors are, the higher is the chance that all found matches are pairs of points that truly correspond to one another.

In the ICP algorithm by Besl and McKay [2], no feature descriptors are computed, but instead closest points in Cartesian space are considered to correspond to one another. A transformation is estimated that minimizes the Euclidean distances between found pairs of closest points in the least squares sense. The process of determining corresponding points in the two data sets and computing the transformation that aligns them is iteratively repeated. The source point set is expected to converge toward the target set as the correspondences increasingly become better. Simultaneously, Chen and Medioni [3] formulated a similar algorithm, but instead of minimizing the squared Euclidean distances between corresponding points, they applied a point-to-plane error metric. In the following two sections, we discuss both approaches.

### Iterative Registration of Closest Points

In contrast to the feature-based registration, iterative registration algorithms do not match salient feature descriptors to find correspondences between source and target point clouds, but instead 1) search for closest points (matching step) and 2) align the found

point pairs (alignment step). These two steps are repeated until convergence occurs or until another termination criterion is reached, thereby iteratively refining the alignment of the source to the target point cloud. Under various assumptions, such as perfect overlap, the alignment converges to the global minimum (with respect to the error metric used) for optimal alignment. The main drawback of iterative registration is that the algorithms may get caught in local minima if the assumptions do not hold, e.g., if the clouds only partially overlap and/or if the initial alignment is off. In this case, false correspondences can negatively affect the registration result. However, several methods exist to sort out false correspondences and improve convergence, applied in the rejection step. Furthermore, if the clouds are already roughly aligned, iterative registration provides efficient and robust means to refine that initial guess and optimally align the point clouds. To speed up registration, another common extension to the original ICP algorithm [2] is to register only subsets of the input point clouds sampled in an initial selection step.

### Selection—Sampling Representative Subsets

Depending on the application and sensor, point clouds can become quite large. Consequently, registering large point clouds is considerably more computationally expensive than registering clouds of smaller cardinality. However, data are often redundant or unnecessarily detailed for the task of registration. Hence, registering only subsets of the original point clouds can yield sufficient results while saving computation time. In principle, two methods of data reduction can be distinguished: automatically extracting a small set of unique and repeatable keypoints (discussed in the “Estimating Keypoints” section) and sampling the original data with respect to a desired target distribution. Whereas the former is intended for feature-based initial alignment (discussed in the “Coarse Alignment via Descriptor Matching” section), the latter can be used to efficiently reduce the amount of data for iterative registration algorithms. The PCL implements several sampling methods, most notably, sampling in index space (simply taking every  $n$ th point), uniform subsampling in the input 3-D space (to better capture the sensed environmental structures), and sampling in the space of local surface normals (to sample points over all surface orientations). An example of using uniform sampling in the PCL is given in Listing 2.

### Matching—Closest Points Correspondence Estimation

Correspondence estimation is the process of pairing points  $p_i$  from the source point cloud  $P$  to their closest neighbors  $q_j$  in

#### Listing 2. Point cloud sampling (uniform).

```
PointCloud<int> indices;
UniformSampling<PointT> uniform_sampling;
uniform_sampling.setInputCloud (cloud);
uniform_sampling.setRadiusSearch (0.05f);
uniform_sampling.compute (indices);
```

the target cloud  $Q$ . This is a greedy approximation of finding the ideal correspondences  $(p_i, q_j)$  between the two clouds. A naïve way of searching for the nearest neighbor is to perform an exhaustive search through all the target points for the nearest neighbor of each source point. As this is prohibitively expensive for applications using millions of points, various data structures for rapid searches have been proposed, such as octrees and kd-trees. These data structures have logarithmic search times (as opposed to the linear naïve variant); however, they take longer to initialize, e.g.,  $O(N \log N)$  for kd-trees. For this purpose, the PCL depends on the Fast Library for Approximate nearest neighbors [4], an open-source library for fast (approximate) nearest neighbor searches. It has been empirically proven to be one of the best-performing solutions [5].

To compute a set of correspondences in our framework, we can use the code in Listing 3. Given `source_cloud` and `target_cloud` as input, `determineCorrespondences` returns the set of found correspondence pairs. Each pair consists of the index of the query point in the source cloud and the index of the found match in the target point cloud. The optional second argument `max_dist` can be used to specify a maximum distance between corresponding points so that pairs with a larger point-to-point distance are filtered out.

In the case of input data coming from sensors that comply with the pinhole camera model, the correspondence estimation procedure can be significantly sped up (though some precision is lost). Such sensors include popular RGB-D cameras, like the Microsoft Kinect, that can collect information from the environment in the form of depth and color images. In the PCL we refer to point clouds with an image-like structure as being organized. Instead of using complex data structures such as kd-trees for nearest neighbor searches in 3-D space, we can use the projective nature of the depth images to obtain a reasonable approximation. Each point in the cloud corresponds to a pixel in the depth image, allowing

---

**The choice and parameterization of the right algorithm for a particular type of data is one of the biggest problems.**

---

#### Listing 3. Correspondence estimation.

```
#include <pcl/registration/correspondence_
  estimation.h>
(...)
CorrespondencesPtr corresps (new
  Correspondences);
CorrespondenceEstimation<PointT, PointT> est;
est.setInputSource (source_cloud);
est.setInputTarget (target_cloud);
est.determineCorrespondences (*corresps,
  max_dist);
```



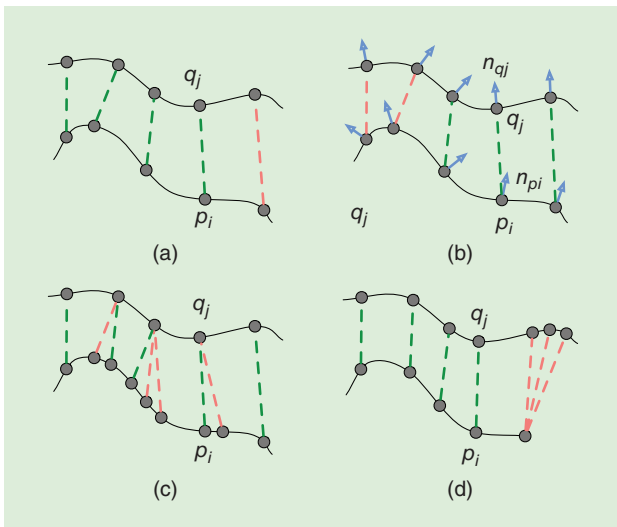
projections from source points in world coordinates to the camera plane of the target frame using the intrinsic and extrinsic camera parameters:

$$\begin{bmatrix} f_x & 0 & c_x & 0 \\ 0 & f_y & c_y & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} [R | t] \begin{bmatrix} p_x \\ p_y \\ p_z \\ 1 \end{bmatrix} = \begin{bmatrix} d \cdot u \\ d \cdot v \\ d \end{bmatrix}, \quad (1)$$

where the first matrix is the intrinsic calibration matrix of the source camera with  $(f_x, f_y)$  being the focal length (divided by image width and height, respectively) and  $(c_x, c_y)$  the camera projection center, the matrix  $[R | t]$  is the transformation of the source camera,  $d$  is the depth of the projected point  $\mathbf{p} = (p_x \ p_y \ p_z)^T$  in the target camera, and  $u, v$  are the coordinates of  $\mathbf{p}$  in the target camera plane. After projection, the query point is assigned to its closest point in the image space (if  $u$  and  $v$  lie within image boundaries) or considered part of the nonoverlapping volume between the source and target cloud.

#### Listing 4. Correspondence estimation using projections.

```
#include <pcl/registration/
correspondence_estimation_organized_
projection.h>
(...)
CorrespondenceEstimationOrganizedProjection
<PointXYZ, PointXYZ> est;
est.setInputSource (cloud_source);
est.setInputTarget (cloud_target);
est.setFocalLengths (f_x, f_y);
est.setDepthThreshold (epsilon);
CorrespondencesPtr corresps (new
Correspondences ());
est.determineCorrespondences (*corresps);
```



**Figure 3.** An example of correspondence rejection. Good correspondence pairs (green) are kept while outliers (red) are sorted out to improve convergence. (a) The rejection of pairs based on the distance between the points, (b) the rejection of pairs based on normal compatibility, (c) the rejection of pairs with duplicate target matches, and (d) the rejection of pairs that contain boundary points.

This approach is fast but imprecise for point clouds with large depth discontinuities or for frames that are far away from each other. That is why this method is recommended to be used only after the two point clouds have been brought close together, making it good for aligning consecutive point clouds in a stream recorded at a high frame rate. A code snippet for projection-based correspondence estimation is presented in Listing 4.

#### Rejecting and Filtering Correspondences

Since invalid correspondences can negatively affect the registration results, most registration pipelines feature a rejection step. This step consists of filtering the point pairs matched in the previous stage to facilitate the transformation estimation algorithm toward convergence to the global minimum. This step can take advantage of auxiliary information available from the input point clouds, such as local surface normals or statistics about the correspondences. Referring to Figure 3, the following correspondence rejection methods are commonly deployed and available in the PCL.

- 1) *Correspondence Rejection Based on Distance:* This method filters out point pairs with a distance larger than a given threshold [see Figure 3(a)]. It is used in [1] and was already mentioned in the original formulation of the ICP algorithm [2] ( $\rightarrow$  CorrespondenceRejectorDistance).
- 2) *Rejection Based on Median Distance:* Unlike the previous rejector, this one does not use a fixed threshold but computes it as the median of all point-to-point distances in the input set of correspondences. Hence, it considers the distribution of the distances between the points and adapts to it, becoming smaller as the two point clouds get closer during ICP iterations. Compared with an adaptive threshold based on the mean value, the median is often more effective in reducing the influence of outliers ( $\rightarrow$  CorrespondenceRejectorMedianDistance).
- 3) *Rejecting Pairs with Duplicate Target Matches:* Usually, each sampled point in the source cloud is assigned to a correspondence in the target cloud. Hence, a point in the target cloud could be assigned multiple corresponding source points [see Figure 3(c)]. This rejector only keeps a single pair  $(\mathbf{p}_{i\min}, \mathbf{q}_j)$ , the one with the minimum distance out of all the pairs  $\{(\mathbf{p}_i, \mathbf{q}_j)\}$  ( $\rightarrow$  CorrespondenceRejectorOneToOne).
- 4) *RANSAC-Based Rejection:* This method applies random sample consensus (RANSAC) [6] to estimate a transformation for subsets of the given set of correspondences and eliminates the outlier correspondences based on the Euclidean distance between the points after the computed transformation is applied to the source point cloud. It is effective in keeping the ICP algorithm from converging into local minima, as it always produces slightly different correspondences and is good at filtering outliers. In addition, it provides good initial parameters for the transformation estimation with all inlier correspondence that follows ( $\rightarrow$  CorrespondenceRejectorSampleConsensus).

For the registration of point clouds coming from projective sensors, correspondence rejectors that exploit the image-like

structure of the input data exist. Most notably, a special two-dimensional (2-D) variant (`→ CorrespondenceRejectionSampleConsensus2D`) of the RANSAC-based correspondence rejector discards pairs based on their pixel distance in the image plane after the source point is transformed and projected into the target camera plane.

- 5) *Rejection Based on Normal Compatibility*: This filter uses the normal information about the points and rejects those pairs that have inconsistent normals, i.e., the angle between their normals is larger than a given threshold. It can reject erroneous pairs that seem correct when judged only by the distance between the points, such as the case depicted in Figure 3(b) (`→ CorrespondenceRejectionSurfaceNormal`).
- 6) *Rejection Based on Surface Boundaries*: Furthermore, when two point clouds captured by a projective sensor represent surfaces that have a partial overlap, accepting correspondences containing surface boundary points can introduce errors [Figure 3(d)]. To detect such points, we can exploit the organized nature of the depth maps and eliminate the correspondences that contain points on depth discontinuities by moving a window across the depth image and checking if there are enough points in the window that are within a threshold depth from the center point (`→ CorrespondenceRejectionBoundaryPoints`).
- 7) *Rejector Pipelines*: Most often, not only is a single rejector applied, but several correspondence rejectors are queued to implement a filtering pipeline. For example, Diebel et al. [7] apply a combination of different solutions for rejecting point pairs in their active stereo point clouds, emphasizing the importance of rejecting points on mesh boundaries (see the RGB-D example in the “Transformation Estimation and Weighting” section). They also reject pairs based on their normal compatibility and based on their point-to-point distance using a distance threshold based on the median distance. In the PCL, multiple rejectors can be easily concatenated using the output correspondences of the previous rejector as input to the next one.

We provide a usage example for PCL correspondence rejection in Listing 5. Given the input point clouds and a set of estimated correspondence pairs, `getCorrespondences` runs the rejection and returns the vector of filtered correspondence pairs. Moreover, all correspondence rejection methods in the PCL feature `getRejectedQueryIndices` to retrieve the

#### Listing 5. Correspondence rejection (based on distance).

```
#include <pcl/registration/
correspondence_rejection_distance.h>
(...)
CorrespondenceRejectionDistance rejector;
rejector.setInputSource<PointT> (cloud_src);
rejector.setInputTarget<PointT> (cloud_tgt);
rejector.setInputCorrespondences
(corresps_in);
rejector.setMaximumDistance (max_dist);
rejector.getCorrespondences (corresps_
filtered);
```

indices of query points of rejected correspondence pairs. This information is useful, for example, to determine nonoverlapping parts between the source and the target point cloud.

### Alignment–Error Metrics and Transformation Estimation

Over the years, there have been numerous mathematical approaches for solving for the rigid transformation  $T$  that minimizes the error of the point pairs.  $T$  is composed of a rotation  $R$  and a translation  $t$ . Note that when referring to a transform  $T$  and a point  $p$ , homogeneous coordinates will be used.

There are two main error metrics to be minimized that have been considered in literature: point-to-point (2) and point-to-plane (3), where  $(p_\kappa, q_\kappa)$  is the  $\kappa$ th of the  $N$  pair correspondences from the source cloud to the target cloud:

$$E_{\text{point-to-point}}(T) = \sum_{k=1}^N w_k \|T p_k - q_k\|^2, \text{ and} \quad (2)$$

$$E_{\text{point-to-plane}}(T) = \sum_{k=1}^N w_k ((T p_k - q_k) \cdot n_{q_k})^2. \quad (3)$$

The optional  $w_\kappa$  can be used for weighting the pairs to give them more or less importance in the least squares formulation ( $w_\kappa = 1$  if no weighting is applied).

- 1) *Standard Point-to-Point Error Metric*: The standard error metric used in the ICP algorithm is the point-to-point error metric (2). When it was first mentioned by Arun [8], researchers proposed various ways of minimizing it, which led to the introduction of the ICP algorithm [2]. Eggert et al. [9] evaluated each of these methods in terms of numerical stability and accuracy, reaching the conclusion that they are close performers.


PCL provides an implementation using a closed-form singular value decomposition (SVD) that was first proposed by Horn [10] (`→ TransformationEstimationSVD`).

- 2) *Point-to-Plane Error Metric*: Chen and Medioni [3] introduced the point-to-plane metric (3) and proved that it was more stable and faster to converge than the previous approaches. This algorithm uses the distance between the source point  $p_\kappa$  and the plane described by the target point  $q_\kappa$  and its local surface normal  $n_{q_\kappa}$ . Unlike the point-to-point metric, it does not have a closed-form solution, so the minimization is done with nonlinear solvers (such as Levenberg–Marquadt, as proposed by [11]) or by linearizing it [12] (under the assumption of small rotation angles, i.e.,  $\sin \theta \sim \theta$  and  $\cos \theta \sim 1$ ).

Depending on the underlying surface and the distribution of points, using the point-to-plane error metric can be considerably more robust. A standard procedure for minimizing it relies on the Levenberg–Marquadt nonlinear optimizer [11] (`→ TransformationEstimationPointToPlane`).

- 3) *Linear Least Squares Point-to-Plane*: The PCL also features an alternative method that accumulates the point-to-plane constraints of all the correspondences in a matrix  $A$  and estimates the rotation and translation in the least squares sense by solving a linear system of the form  $A^T A v = A^T b$ , with

$v$  being a six-dimensional parameterized representation of the minimizing transformation, as suggested by [12] ( $\rightarrow$  TransformationEstimationPointToPlaneLLS).  
 4) *Weighted Point-to-Plane Error Metric*: Assigning a different weight to each correspondence can improve convergence. The weighting of the point pairs can be seen as a soft correspondence rejection, adjusting the influence of noisy corresponding points in the minimization process. The weighting



**Figure 4.** The error metrics and transformation estimators. (a) The point-to-point error, (b) the point-to-plane error, and (c) the Generalized-ICP.

**Listing 6. Transformation estimation (weighted point-to-plane).**

```
#include <pcl/registration/
transformation_estimation_point_to_plane_
weighted.h>
(...)
TransformationEstimationPointToPlane
Weighted
<PointXYZ, PointXYZ, double> te;
te.setWeights (correspondence_weights);
te.estimateRigidTransformation (*cloud_src,
*cloud_tgt,
*corresps_filtered, transform);
```

can be a function of the point-to-point or point-to-plane distance between the points, a function of the angle between the normals corresponding to the points, or a function of the noise model of the sensor that has been used. Numerous publications [1], [7], [13], and [14] consider the weighting of the point pairs as a beneficial step for the robustness and convergence rate of the transformation estimation algorithm.

Referring to Figure 4, the PCL offers a multitude of techniques for estimating the transformation between two point clouds given a set of correspondences. As an example for transformation estimators, a code snippet showing how to use the weighted point-to-plane error metric ( $\rightarrow$  TransformationEstimationPointToPlaneWeighted) is given in Listing 6.

- 5) *Weighted Linear Least Squares Point-to-Plane*: Similar to the previous metric, the linear least squares point-to-plane variant can be used with nonconstant weights ( $\rightarrow$  TransformationEstimationPointToPlaneLLSWeighted).  
 6) *Generalized Error Metric and Generalized-ICP*: Segal et al. [15] proposed an error metric that generalizes over point-to-point and point-to-plane. It uses covariances of the local point neighborhoods ( $\Sigma_k^P$  and  $\Sigma_k^Q$ ) to align the underlying surfaces rather than the points themselves:

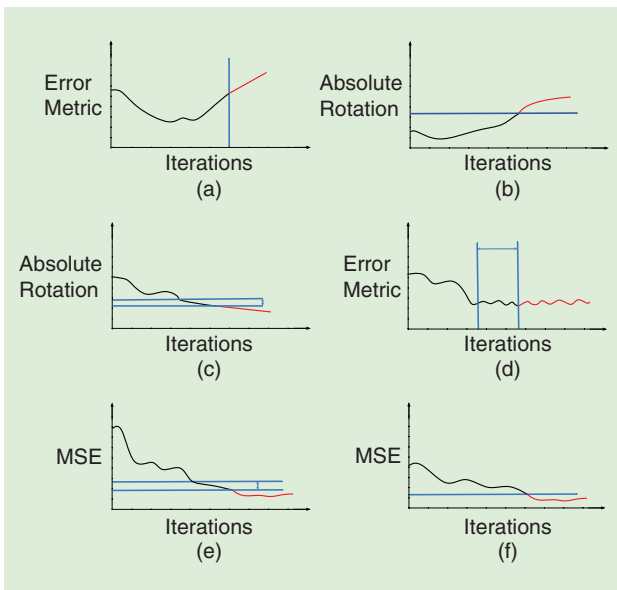
$$E_{\text{Generalized-ICP}}(T) = \sum_{k=1}^N \mathbf{d}_k^{(T)T} (\Sigma_k^Q + T \Sigma_k^P T^T)^{-1} \mathbf{d}_k^{(T)}, \quad (4)$$

where  $\mathbf{d}_k^{(T)} = \mathbf{q}_k - T\mathbf{p}_k$  are the point-to-point distances of the correspondences ( $\mathbf{p}_k, \mathbf{q}_k$ ),  $T$  is the transformation matrix, and  $T^T$  its transpose. In the PCL, Generalized-ICP is implemented as a complete registration component that includes computation of the covariances  $\Sigma_k^P$  and  $\Sigma_k^Q$  for all points in source cloud  $P$  and target cloud  $Q$  ( $\rightarrow$  GeneralizedIterativeClosestPoint).

**Termination Criteria**

For iterative registration algorithms, termination criteria for the transformation refinement process must be defined. In addition to simply setting a maximum number of iterations, many other criteria can be specified. Among others, the following termination criteria are available in the PCL. Using a common interface ( $\rightarrow$  DefaultConvergenceCriteria), they can be efficiently combined and configured, as shown in Listing 7.

- 1) *Maximum Number of Iterations*: Exceeding the number of iterations means that the optimizer diverged. This threshold has to be tuned depending on the complexity of the registration problem. Expect that registering a pair of scans that are far away from each other will require more iterations than two scans that have a good initial alignment [Figure 5(a)].  
 2) *Absolute Transformation Threshold*: The iterations are stopped when the currently estimated transformation is far away from the initial transformation. This is an early termination criteria for registration procedures that diverge. The intuition behind it is that the two clouds to be aligned



**Figure 5.** The ICP termination criteria supported in the PCL. (a) The maximum number of iterations, (b) an absolute transformation threshold, (c) the relative transformation threshold, (d) the maximum number of similar iterations, (e) the relative MSE threshold, and (f) an absolute MSE threshold.

are expected to be within a certain distance range from each other, and transformations that are outside that range need to be rejected, e.g., two consecutive handheld-Kinect scans recorded at 30 Hz cannot be more than 10 cm and 20° apart from each other [Figure 5(b)].

- 3) *Relative Transformation Threshold*: It specifies the minimum transformation difference from one iteration to the next that is considered small enough for the optimizer to have converged [Figure 5(c)].
- 4) *Maximum Number of Similar Iterations*: Under the previous stopping criteria, a minimizer might temporarily seem to have converged, but it is actually oscillating in a local minimum and has the chance to escape from it and converge into the global minimum or a better local minimum. For this reason, we allow the optimizer to spend a certain number of iterations around a minimum point before considering it converged [Figure 5(d)].
- 5) *Relative Mean Square Error (MSE)*: This criterion is similar to the relative transformation threshold, using the MSE metric instead of the rotation/translation increment [Figure 5(e)].
- 6) *Absolute MSE*: It stops the iterations when the error between the two aligned clouds is below a certain value [Figure 5(f)].

All of these techniques represent ways of balancing the quality and the runtime of the registration procedure. They prove to be essential for tuning algorithms for real-time performance.

### Transformation Validation (Optional)

Due to the presence of numerous local minima in the ICP error function, a series of checks can be performed at the end of the registration process to confirm a successful convergence. We suggest checking for the percentage of the overlapping surface of the two point clouds (`→ TransformationValidationEuclidean`). In the case of sensors that obey the pinhole camera model (RGB-D), this can be done efficiently in image space by rendering one point cloud on top of the other with the computed poses. Similarly, one can check if the overlapping region has enough inliers (see Figure 6). Another indicator for the success of the registration is the condition number of the overlapping region [16], which tells if the common surface of the point clouds is stable enough to not allow slippage.

### Coarse Alignment via Descriptor Matching

As discussed, ICP-based algorithms, due to their greedy nature, require a reliable initial alignment to avoid converging to bad local minima. Typical solutions to this problem are represented by running ICP multiple times, each time with a different, random initialization—prohibitively slow for most applications—or by computing an initial transformation from a hand-chosen set of correspondences inserted by a human user (semiautomatic approach). A further approach, which is fast and fully automatic, is determining point-to-point correspondences between 3-D keypoints extracted from both clouds via matching the associated keypoint descriptors. This

set of correspondences is then successively pruned of mismatches by means of an outlier rejection scheme so that it can be deployed to obtain the transformation aligning one cloud to the other. The use of a coarse alignment increases the chance of successful alignment and makes ICP converge faster. In addition, it is a fully automatic approach that does not require user intervention. The PCL includes several state-of-the-art tools for carrying out initial alignment via descriptor matching, which are briefly outlined in the following sections.

---

**A preprocessing sequence is applied to improve the signal-to-noise ratio or to recover information about the surface.**

---

### Estimating Keypoints

Coarse pairwise alignment can be carried out without an explicit keypoint selection stage, i.e., by random selection of points or by uniform sampling of the cloud (as discussed in the “Selection—Sampling Representative Subsets” section). Nevertheless, several choices are currently available in the literature and in the PCL that explicitly aim to extract salient and repeatable keypoints given an input cloud. The advantage of relying on specific 3-D detection schemes is that the repeatability and robustness of the computed 3-D correspondences with respect to nuisances affecting the data are generally improved, thus enhancing the robustness of 3-D registration schemes.

Due to the topic of 3-D keypoint detection being fairly new, a common approach consists of porting the key criteria of successful image interest point detectors from the 2-D domain (i.e., images) to the 3-D domain. This is the case of the popular Harris corner detector [17], which computes the Taylor expansion of the directional intensity variation at each image point. In the 3-D extension of this criterion (`→ Harris3D`), the directional intensity variation is substituted by

#### Listing 7. Combination of termination criteria.

```
#include <pcl/registration/default_
  convergence_criteria.h>
(...)
DefaultConvergenceCriteria<double> conv_crit
  (iteration, transform, *correspondences);
conv_crit.setMaximumIterations (30);
conv_crit.setMaximumIterationsSimilar
  Transforms (3);
conv_crit.setTranslationThreshold (5e-3);
conv_crit.setRotationThreshold (cos (0.5 *
  M_PI / 180.0));
conv_crit.setRelativeMSE (0.01);
do
{ (... ICP iterations ...)
} while (!conv_crit.hasConverged ())
ConvergenceState why_conv =
  conv_crit.getConvergenceState ();
```



computing the angle between the normal of the central point and its neighbors: a point yielding big angle values along several directions will typically denote a corner in 3-D. Two variants of such a detector are also available, one detecting corners in the intensity domain ( $\rightarrow$  `Harris2D`), the other jointly exploiting the intensity as well as the 3-D domain ( $\rightarrow$  `Harris6D`). As for most 3-D keypoint detectors [18], a nonmaxima suppression (NMS) stage is applied to retrieve only local maxima of the saliency measure evaluated at each point by the detector.

Other 3-D keypoint detector approaches were specifically designed for the 3-D or 2.5-D domain. This is the case of intrinsic shape signature (ISS) [19], which is based on the eigenvalue decomposition (EVD) of the  $3 \times 3$  scatter matrix computed between each point and all neighbors within a

spherical support. Nondistinctive points are rejected by thresholding the ratios of the three eigenvalues. The smallest eigenvalue is used as the keypoint saliency, which is subject to the NMS stage. Thus, keypoints retained by this method exhibit a distinct variance of the point coordinates along one direction with respect to the other two. As an example, Listing 8 shows how to compute keypoints using ISS.

Another recent interest point detector is normal aligned radial feature (NARF) [20], a method specifically designed to extract salient keypoints on range images. It focuses on detecting points along object depth borders. In particular, a subset of border points is selected such that they exhibit substantially different dominant directions of the surface in the local neighborhood, so that they can be stable and repeatable with respect to viewpoint changes ( $\rightarrow$  `NarfKeypoint`).

For more details on other relevant proposals in the field, as well as for a recent performance evaluation, we refer the reader to [18].

**Listing 8. Keypoint estimation (intrinsic shape signatures).**


```
ISSKeypoint3D<PointT, PointT> iss_detector;
iss_detector.setSalientRadius (salient_
radius);
iss_detector.setNonMaxRadius (non_max_
radius);
iss_detector.setInputCloud (input_cloud);

PontCloud<PointT>::Ptr keypoints
(PontCloud<PointT>());
iss_detector.compute (*keypoints);
```


**Describing Keypoints—Feature Descriptors**

For each detected keypoint, a descriptor is computed to determine correspondences between the two point clouds. A local descriptor is a compact representation of a point's local neighborhood. In contrast to global descriptors describing a complete object or point cloud, local descriptors try to resemble shape and appearance only in a local neighborhood around a point, and are suitable for representing it in terms of matching. Among the approaches currently available in the literature, we provide a brief outline of the main solutions available in the PCL for the registration of 3-D clouds. Apart from the outlined descriptors, the PCL includes several other methods. We refer the interested reader to [21], a survey focused on the 3-D descriptors available in the PCL (including related code snippets), and to [22], focusing on a performance evaluation of PCL 3-D descriptors.

The spin image descriptor [23] rotates a plane section around the normal of the keypoint. The plane section accumulates the number of points of the cloud intersected over a 2-D histogram while doing a complete spin around the normal [see Figure 7(a)]. One advantage of this approach is



**Figure 6.** An example of transformation validation. If two surfaces match locally, we need to make sure that their empty space also matches and that there is no additional depth information in the overlap region of any of the clouds that does not have correspondences in the other cloud.



**Figure 7.** Functional principles of feature descriptors. (a) The spin image descriptor spins a plane section around the keypoint normal to compute its 2-D histogram. (b) The local reference frame and the three angles,  $\alpha$ ,  $\phi$ , and  $\theta$  computed by FPFH at each point pair. (c) The 3-D spherical grid deployed by 3DSC. Two grid sectors, yielding each a bin in the 3-D histogram, are depicted in green and red, respectively. (d) The 3-D grid deployed by SHOT, which is repeatably oriented by the local reference frame denoted by the blue arrows.

that only a repeatable normal—rather than a full local reference frame (LRF)—is required to compute the descriptor.

The fast point feature histogram (FPFH) descriptor [24] stores the relative orientation of normals and distances between point pairs falling within the spherical neighborhood of a keypoint. Point pairs are formed by the keypoint and its nearest neighbors, as well as by the nearest neighbors of each keypoint's neighbor. For each pair  $(\mathbf{p}_i, \mathbf{p}_j)$ , a repeatable LRF is constructed by means of a Darboux frame coordinate system [24]. Then, three angles  $\alpha, \phi$ , and  $\theta$  are computed out of different comparisons between the unit vectors composing the LRF and the normals of  $(\mathbf{p}_i, \mathbf{p}_j)$ , as shown in Figure 7(b). These three angles are accumulated over three separate histograms for all point pairs, which are in turn juxtaposed so to yield the final descriptor. Listing 9 shows how to compute FPFH feature descriptors for a given set of keypoints.

3-D shape context (3DSC) [25] relies on a spherical grid centered on the keypoint and quantized along its three domains: radius, azimuth, and elevation [see Figure 7(c)]. This quantization defines a 3-D histogram, where each bin is represented by a volume of the grid, and the accumulated value is the number of points falling therein and weighted by the volume size and the local point density. The final descriptor stores all bins of the histogram, by aligning the north pole of the grid with the normal of the central point, and by shifting the grid over the normal plane as many times as the number of azimuth subdivisions. Multiple grid orientations, in turn, yield multiple descriptions for the same point, so to obtain rotation invariance ( $\rightarrow$  ShapeContext3DEstimation).

The unique shape context [26] descriptor has been proposed as an extension of 3DSC where an LRF is associated to each point being described. The LRF provides two repeatable principal directions over the normal plane which uniquely orient the 3-D grid associated with each descriptor, thus yielding a single descriptor at each keypoint rather than multiple ones. This reduces the memory footprint and matching ambiguities during the successive correspondence estimation stage ( $\rightarrow$  UniqueShapeContext).

The signature of histograms of orientations (SHOT) descriptor [27] relies on a 3-D grid centered on the keypoint and oriented according to a repeatable LRF [see Figure 7(d)]. Each unit vector making up this LRF coincides with a principal direction obtained via the EVD of the scatter matrix of the points falling within the grid, and its sign is disambiguated according to a specific procedure relying on the relative density of the cloud [27]. For each grid sector, a histogram is computed by accumulating the angles between the normal of the keypoint and the normal of each point falling within that sector. The final descriptor is formed by orderly juxtaposing all computed histograms ( $\rightarrow$  SHOTEstimation).

### Correspondence Estimation and Filtering

The simplest approach to establishing point-to-point correspondences between the keypoints extracted from the two clouds is

#### Listing 9. Feature estimation (FPFH).

```
FPFHEstimation<PointT, NormalT,
  FPFHSignature33> fpfh;
fpfh.setSearchSurface (cloud);
fpfh.setInputNormals (normals);
fpfh.setRadiusSearch (radius);
fpfh.setInputCloud (keypoints);

PointCloud<FPFHSignature33> features;
fpfh.compute (&features);
```

to associate each descriptor on one cloud to its nearest neighbor in the descriptor space among the descriptor set on the other cloud. Obviously, several correspondences are going to be erroneous due to the presence of nonoverlapping areas between the two clouds, as well as due to spurious matches caused by noise, geometry deformations, and nondistinctive descriptors. Correspondence estimation needs to match only a subset of keypoints and to reject found correspondences, respectively (as discussed in the “Rejecting and Filtering Correspondences” section). One of the most common approaches is to use RANSAC, where the generative model is the 6-degree-of-freedom (6DOF) transformation between the two clouds built over three randomly chosen correspondences (discussed in the “Rejecting and Filtering Correspondences” section).

The alignment converges to the global minimum for optimal alignment.

### Transformation Estimation

Once outlier-free correspondences are reliably estimated, an absolute orientation algorithm is used to estimate the 6DOF transformation between the two clouds based on all remaining correspondences. Among the algorithms available in the literature, the PCL includes an implementation of the method based on the SVD of the matrix representation of the 6DOF transformation proposed in [8] and [10] ( $\rightarrow$  TransformationEstimationSVD).

Note that this is the same type of transformation estimation used in iterative registration where correspondences are not known but instead are iteratively refined by always using the closest points. A code sample showing how to use this transformation estimator has already been presented in Listing 6.

### Experiments and Examples

A whole range of diverse applications benefits from reliable pairwise alignment between 3-D point clouds. These applications differ in terms of 1) characteristics of the acquired data, e.g., high-resolution polygonal mesh as opposed to low-resolution point clouds, 2) working scale, e.g., the large-scale used in most urban and photogrammetry scenarios, as opposed to small-scale, such as what is typically required for reverse engineering and object modeling, and 3) a priori knowledge, which depends on whether estimates

of the acquisition viewpoints are available (in which case the initial estimate only needs to be refined) as opposed to the situation where they are unknown (in which case an initial alignment needs to be computed solely using the acquired data).

In addition to these differences, which often require adopting ad hoc solutions, one can also exploit certain characteristics of the data to improve speed or robustness.

**Publicly available software libraries and data sets are a way to promote such research and to make results comparable.**

Since a fundamental problem of registration is to find the most suitable algorithm with an appropriate parameterization, we discuss three different examples ranging from dense RGB-D image sequences for object modeling to high-resolution and low-resolution laser range

scans for the outdoor mapping of buildings and cities.

### Large-Scale 3-D Laser Scans

Commercially available 3-D laser scanners, such as those used by surveying technicians, acquire high-resolution large-scale scans that usually consist of several million points. With a comparably ample maximum measurement range, these sensors can be used to capture large scenes such as buildings and streets. To create a complete model (e.g., a whole city or an urban district), the acquired scans need to be aligned and aggregated. Frequently, surveying technicians use additional means, such as special reflectors, that assist in obtaining a fine alignment of the data. In our example, we assume that such information is not available and that we want to 1) compute an initial alignment for a pair of large-scale 3-D laser scans, and 2) refine

the initial estimate to obtain a perfect overlay of the acquired data.

### Initial Alignment

As mentioned in the “Coarse Alignment via Descriptor Matching” section, the PCL provides implementations for a variety of 3-D descriptors that can be used for initial alignment. Here, the key issue is not only to find an appropriate descriptor for the problem at hand, but also an appropriate scale, i.e., the support size on which to compute each descriptor. Of particular interest in this context is the distinctiveness of each descriptor that influences the number of possible ambiguities in the matching stage as well as the repeatability, such that the descriptor computed on the same scene point over different scans looks the same. A particularly robust approach is to compute the same class of descriptor at various scales and to focus the registration on points and associated descriptors that prove to be distinctive at the same scale, and are persistent over multiple scales. The PCL provides joint estimation components that finds such keypoints and descriptors ( $\rightarrow$  MultiScaleFeaturePersistence).

Listing 10 shows an example for computing keypoints for an input point cloud and the respective persistent feature descriptors. First, the selected descriptor (in our example we use FPFH [24]) is fetched to an instance of the

**Listing 10. Multiscale keypoint and descriptor computation.**

```
#include <pcl/features/fpfh.h>
#include <pcl/features/multiscale_feature_persistence.h>
(...)
FPFHEstimation<PointXYZ, Normal, FPFHSignature33>::Ptr fest (new (...));
fest->setInputCloud (cloud)
fest->setInputNormals (normals);

MultiscaleFeaturePersistence<PointXYZ, FPFHSignature33> fper;
std::vector<int> keypoints;
std::vector<float> scale_values = { 0.5f, 1.0f, 1.5f };
fper.setScalesVector (scale_values);
fper.setAlpha (1.3f);
fper.setFeatureEstimator (fpfh_estimation);
fper.setDistanceMetric (pcl::CS);
fper.determinePersistentFeatures (*features, keypoints);
```

**Listing 11. Initial alignment using keypoints and descriptors.**

```
#include <pcl/registration/correspondence_estimation.h>
#include <pcl/registration/correspondence_rejection_sample_consensus.h>
#include <pcl/registration/transformation_estimation_svd.h> (...)
Correspondences::Ptr correspondences (new Correspondences);
CorrespondenceEstimation<FPFHSignature33, FPFHSignature33> cest;
cest.setInputSource (source_features);
cest.setInputTarget (target_features);
cest.determineCorrespondences (*correspondences);

Correspondences::Ptr corr_filtered (new Correspondences);
CorrespondenceRejectionSampleConsensus<PointXYZ> rejector;
rejector.setInputSource (source_keypoints);
rejector.setInputTarget (target_keypoints);
rejector.setInlierThreshold (2.5);
rejector.setMaximumIterations (1000000);
rejector.setRefineModel (false);
rejector.setInputCorrespondences (correspondences);
rejector.getCorrespondences (*corr_filtered);

TransformationEstimationSVD<PointXYZ, PointXYZ> trans_est;
trans_est.estimateRigidTransformation (*source_keypoints, *target_keypoints, *corr_filtered, transform);
```



MultiScaleFeaturePersistence class together with a vector of pre-selected scales and a distance metric. The class computes the descriptors with the given estimator for all desired scales and uses the provided distance metric to compute similarity. It returns the indices of all selected interest points that can then be used for an initial alignment.

To compute an initial alignment for a pair of target and source point clouds, we first match the estimated descriptors (CorrespondenceEstimation in Listing 11). We then reject false correspondences by finding a consistent set (in a RANSAC-like fashion) using CorrespondenceRejectionSampleConsensus. It samples from the estimated correspondences, computes a candidate transformation, and then determines all inliers as correspondences where the respective source and target point are sufficiently close after applying the candidate transformation (parameter set using setInlierThreshold). The transformation yielding the highest number of inliers is the most likely initial alignment. All inliers are then used to compute the transformation by means of TransformationEstimationSVD.

### Fine Registration

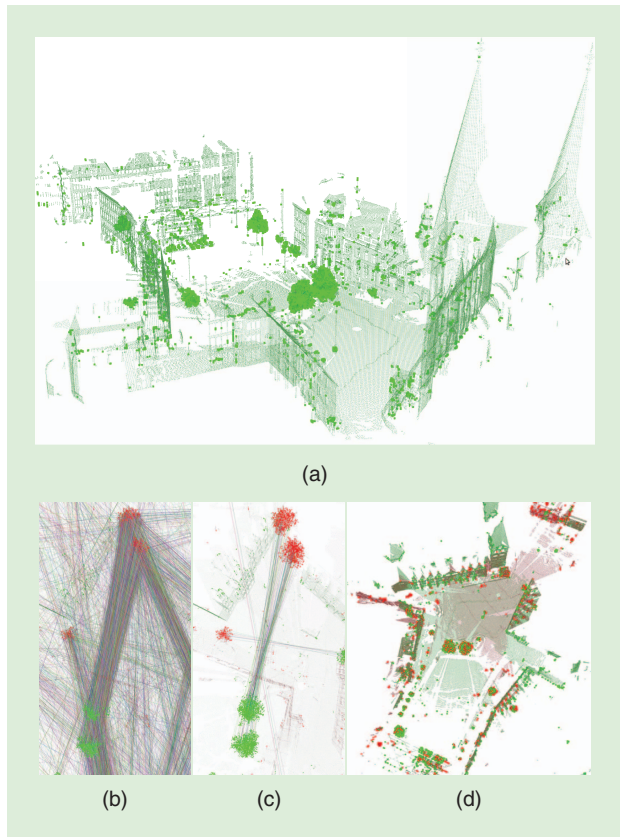
After the initial alignment, the source and target point clouds are already roughly registered. An iterative algorithm is used

to correct the estimated viewpoints and to accurately align the point clouds. In Figure 8, we show the results of the proposed pipeline on typical large-scale 3-D laser scans (from the Bremen data set) in the 3-D Scan Repository (<http://kos.informatik.uni-osnabrueck.de/3-Dscans>).

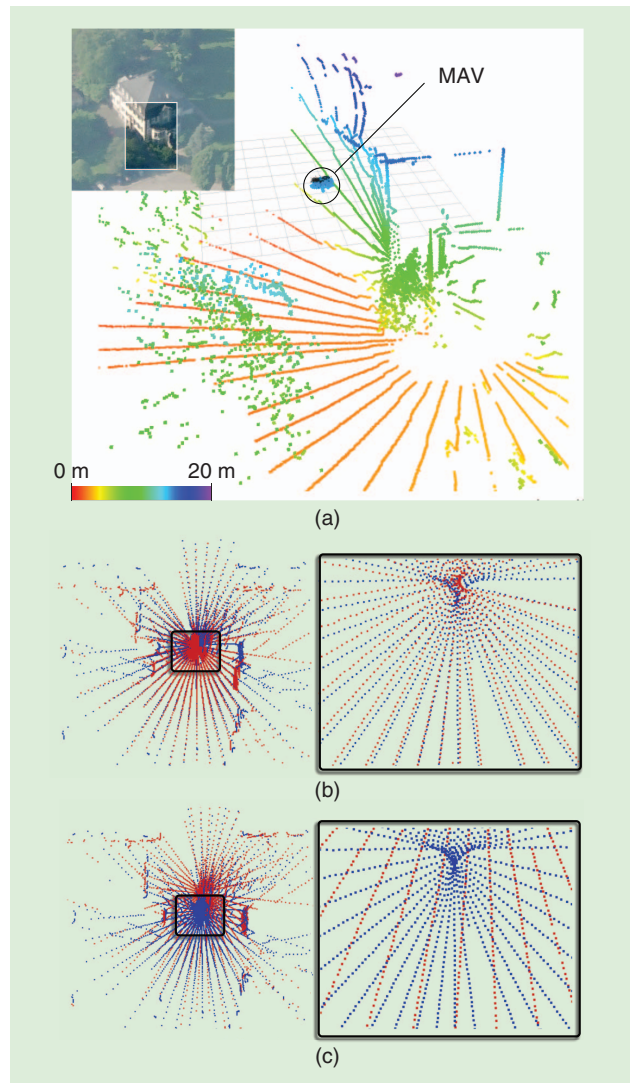
### Sparse Low-Resolution 3-D Laser Scans

In the previous example, the 3-D point clouds were not only accurate, but also dense. We now focus our attention on problems that arise when the data are particularly sparse and how to compensate for the resulting effects.

A particular characteristic of custom-built 3-D scanners composed of a rotating 2-D laser range finder is that the



**Figure 8.** The registration results. (a) The target point cloud with persistent keypoints (visualized with larger radii). (b) A detailed view of the estimated correspondences. (c) A detailed view of the filtered correspondences. (d) The target point clouds after initial and refined alignment.



**Figure 9.** The registration of sparse laser scans. (a) A perspective view of an example scan with MAV and a building. The point colors encode height and a photo highlights the scanned region. (b) The original Generalized-ICP. The standard implementation of Generalized-ICP suffers from inaccurate covariance estimates and incorrectly aligns the two point cloud origins and the individual scan lines. (c) The Generalized-ICP with approximated covariances. When approximating the underlying surface, computing the covariances on the reconstruction, and using these externally computed covariances for Generalized-ICP, we can accurately align the clouds.



acquired point clouds have nonuniform point densities. These densities usually have a high density within each scan line and a certain angle between scan lines, depending on the rotation speed of the scanner. Rotating the scanner quickly

**Pairwise registration is usually carried out by means of one of the several variants of the ICP algorithm.**

increases the frequency with which the surroundings can be perceived (which is important for collision avoidance, e.g., as in our example of an autonomous MAV), but reduces the angular resolution, and thus, the point density in the acquired point clouds (see Figure 9 for an example). The resulting nonuniform point

densities affect classic neighborhood searches in 3-D and cause problems in the local feature estimation and registration. For example, the Generalized-ICP algorithm by Segal et al. [15] uses local surface information in the form of covariance matrices to align surfaces rather than individual points. However, as can be seen in Figure 9, inaccurate estimates (as

caused by the nonuniform point densities) degrade the performance of the algorithm and negatively affect the registration result.

To compensate for these effects, we can approximately reconstruct the underlying surface and then compute the covariances directly on the reconstructed mesh [28]. We thereby exploit the topological information in the organized data and can accurately register such pairs of sparse point clouds. In this example, we show how to approximate the surface reconstruction and the covariances, and how to use custom covariances with Generalized-ICP.

Listing 12 shows the main parts of the implementation using the PCL. We first approximate the surface for both source and target point cloud ( $\rightarrow$  `organizedFastMesh`), and then compute estimates of the local surface normals and the covariances directly on the reconstructed meshes. We then use the standard Generalized-ICP algorithm (discussed in the “Alignment—Error Metrics and Transformation Estimation” section) with the covariances estimated on the reconstructed surface information ( $\rightarrow$  `Generalized ICP`).

### Registering Sequences of RGB-D Images

In the last example, we present a complete pipeline for aligning scans coming from a Microsoft Kinect, which is an example of an RGB-D camera, i.e., a 3-D sensor that simultaneously acquires color and depth images at high frame rates (30 Hz). These sensors have both data sources synchronized and registered, resulting in point clouds where each point has an associated RGB color and lies on a 2-D grid structure. As mentioned

#### Listing 12. Generalized-ICP using approximated covariances.

```
#include <pcl/surface/organized_fast_mesh.h>
#include <pcl/registration/gicp.h>
#include <pcl/features/from_meshes.h>
PointCloud<PointT>::Ptr cloud;
PolygonMesh::Ptr mesh;
(...)
// reconstruct meshes for source and target
OrganizedFastMesh<PointT> fast_mesh;
fast_mesh.setInputCloud(cloud);
fast_mesh.reconstruct(*mesh);
(...)
// compute normals and covariances for source
and target
PointCloud<Normal>::Ptr normals;
boost::shared_ptr<
  std::vector<Eigen::Vector3d> > covs;
pcl::features::computeApproximateNormals(
  *cloud, mesh->polygons, *normals);
pcl::features::computeApproximateCovariances(
  *cloud, normals, *covs);
(...)
// setup Generalized-ICP
GeneralizedIterativeClosestPoint<PointT,
  PointT> gicp;
gicp.setMaxCorrespondenceDistance(1);
gicp.setInputSource(source_cloud);
gicp.setInputTarget(target_cloud);
gicp.setSourceCovariances(source_
  covariances);
gicp.setTargetCovariances(target_
  covariances);

// run registration and get
  transformation PointCloud<PointT> output;
gicp.align(output);
Eigen::Matrix4f transform = gicp.
  getFinalTransformation();
```

#### Listing 13. Filtering (and upsampling) for RGB-D Data.

```
#include <pcl/filters/fast_bilateral.h>
(...)
FastBilateralFilter<PointXYZ> filter;
filter.setInputCloud(cloud_in);
filter.setSigmaS(5);
filter.setSigmaR(0.005f);
bilateral_filter.filter(*cloud_out);
```




Figure 10. A typical result of registering a sequence of RGB-D frames.

earlier, in the PCL we refer to these image-like 3-D representations as organized point clouds.

It is assumed that consecutive clouds acquired from the sensor are temporally (and, thus, spatially) close to each other. For this to hold, not only must the device have a sufficiently high frame rate, but the pipeline must also process frames at a comparable speed. This is the working assumption of cloud registration algorithms deployed within SLAM and Kinect Fusion-like frameworks [29]. Under these conditions the point clouds are close enough to allow ICP to reach a good local minimum in most cases, without the need for the initial alignment procedure, and with just a few iterations. We will briefly walk through the steps of the pipeline.

### Preprocessing—Filtering and Fast Feature Estimation

This type of camera exhibits various forms of noise. Hence, initial filtering becomes necessary when the cameras are used to acquire accurate 3-D models. There are multiple ways in which noise can be efficiently eliminated from generic point clouds, but none of them consider the difficult quantization effects seen in this new type of data. An edge-preserving filtering method, the bilateral filter, has been adapted from the field of 2-D image processing ( $\rightarrow$  `FastBilateralFilter`). It smoothes more when neighboring pixels are similar (flat regions), and less when there are jumps, therefore, it does not affect the edges. Listing 13 provides a usage example.

Accelerating the computation of normals can also be done by using the grid structure of the RGB-D data. The rasterized depth images can be used to speed up the normal estimation process using integral images [30]. The advantage of this method is that it requires only a linear preprocessing stage while allowing for computing mean vectors and normals within a rectangular area of the image in constant time ( $\rightarrow$  `IntegralImageNormalEstimation`).

### Registration Pipeline

For aligning a sequence of RGB-D images, we present a hybrid registration pipeline. Because of the high data rate of these sensors ( $640 \times 480 = 307,200$  points at 30 Hz), the measurements need to be sampled so that the registered point clouds contain only 10% of the original data. As described in the “Selection—Sampling Representative Subsets” section, there are many approaches available for sampling the point cloud.

### Normal Space Subsampling

A particularly robust way to register RGB-D point clouds is to perform normal space sampling ( $\rightarrow$  `NormalSpaceSampling`). This method consists of binning each point into  $M^3$  bins based on the normal angle around each of the  $x$ ,  $y$ , and  $z$ -axis, respectively. Next, a fixed number of samples is extracted randomly from each bin. This ensures that the resulting point cloud has samples at locations from all the differently oriented surfaces in the scene, increasing the probability that the registration will converge to the global minimum.

### Correspondence Estimation and Filtering

Having only a fraction of the points left, using a complex search data structure is not that prohibitively slow anymore. Consequently, we apply the correspondence estimation based on kd-trees (discussed in the “Matching—Closest Points Correspondence Estimation” section) rather than the projection-based estimation.

To reject misleading correspondences, we employ a combination of two correspondence filters (discussed in the “Rejecting and Filtering Correspondences” section): a filter based on normal compatibility ( $\rightarrow$  `CorrespondenceRejectorSurfaceNormal`) to reject point pairs with mismatched normals, and a distance filter using the median distance between correspondence pairs ( $\rightarrow$  `CorrespondenceRejectorMedianDistance`) to eliminate the pairs with outlying distances.

### Transformation Estimation and Weighting

Using the resulting small set of robust correspondences, the transformation between the two point clouds is computed using the weighted point-to-plane distance metric ( $\rightarrow$  `TransformationEstimationPointToPlaneWeighted`). Because our normals have been computed on the original point clouds, they approximate the underlying surface well. As a result, the point-to-plane error metric is preferred as a better estimation of the distance between the source and target surfaces, as opposed to point-to-point measurements. The weighting of the correspondences can be done using a sensor noise model. One such example was proposed by [13] and considers the standard deviation of the depth to increase quadratically with the distance of the scanned surfaces from the image plane. As such, the weight for the correspondence between points  $\mathbf{p}$  and  $\mathbf{q}$  is taken as the maximum of the standard deviation of the two points under the noise model

$$w(\mathbf{p}, \mathbf{q}) = \max(\sigma_z(\mathbf{p}_z), \sigma_z(\mathbf{q}_z)), \quad (5)$$

$$\text{with } \sigma_z(z) = 0.0012 + 0.0019 * (z - 0.4)^2. \quad (6)$$

As convergence criteria for the ICP algorithm (discussed in the “Termination Criteria” section), a good configuration is to use the number of iterations (chosen based on the time we have allocated for registration) and the length of the incremental translation and rotation from the last iteration (chosen to be lower than the expected camera movement between frames).

### Typical Registration Results

For the purpose of the experiment, we recorded sequences of RGB-D frames at 30 Hz. Each frame is registered in a pairwise fashion against the previous one. Results are shown in Figures 1 and 10. Note that if a large number of frames is registered pairwise, individual alignment errors accumulate. There are numerous techniques in the literature to solve this issue, such as loop closing via graph optimization or forms of global registration.

## Conclusion

In this article, we have presented the modular registration framework implemented in the open source PCL. Following the usual scheme of registration algorithms, we explained different modules available in the PCL and presented code snippets showing how to use them. Finally, we presented three examples of registration approaches, each tailored to a specific type of input data to show how the different components in the PCL can be efficiently combined to solve various registration problems.

A sequel to this tutorial focusing on a global multiview registration is planned. The complete example implementations have been made available together with the used data sets at <http://pointclouds.org/media/registration-paper.html>.

## References

- [1] S. Rusinkiewicz and M. Levoy, "Efficient variants of the ICP algorithm," in *Proc. Int. Conf. 3-D Digital Imaging Modeling*, 2001, pp. 145–152.
- [2] P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 2, pp. 239–256, 1992.
- [3] Y. Chen and G. Medioni, "Object modelling by registration of multiple range images," *Image Vis. Comput.*, vol. 10, no. 3, pp. 145–155, 1992.
- [4] M. Muja and D. G. Lowe, "Fast approximate nearest neighbors with automatic algorithm configuration," in *Proc. Int. Conf. Computer Vision Theory Application*, 2009, pp. 331–340.
- [5] J. Elseberg, S. Magnenat, R. Siegwart, and A. Nüchter, "Comparison on nearest-neighbour-search strategies and implementations for efficient shape registration," *J. Softw. Eng. Robot.*, vol. 3, no. 1, pp. 2–12, 2012.
- [6] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [7] J. Diebel, K. Reuterswärd, S. Thrun, J. Davis, and R. Gupta, "Simultaneous localization and mapping with active stereo vision," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Systems*, 2004, pp. 3436–3443.
- [8] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-D point sets," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 9, no. 5, pp. 698–700, 1987.
- [9] D. W. Eggert, A. Lorusso, and R. B. Fisher, "Estimating 3-D rigid body transformations: A comparison of four major algorithms," *Mach. Vis. Appl.*, vol. 9, nos. 5–6, pp. 272–290, 1997.
- [10] B. K. P. Horn, "Closed-form solution of absolute orientation using unit quaternions," *J. Opt. Soc. Amer. A*, vol. 4, no. 4, pp. 629–642, 1987.
- [11] A. W. Fitzgibbon, "Robust registration of 2D and 3D point sets," in *Proc. British Machine Vision Conf.*, 2001, pp. 411–420.
- [12] K.-L. Low, "Linear least-squares optimization for point-to-plane ICP surface registration," Dept. Comput. Sci., Univ. North Carolina, Chapel Hill, NC, Tech. Rep. TR04-004, 2004.
- [13] C. V. Nguyen, S. Izadi, and D. Lovell, "Modeling Kinect sensor noise for improved 3D reconstruction and tracking," in *Proc. 2nd Int. Conf. 3D Imaging, Modeling, Processing, Visualization Transmission*, 2012, pp. 524–530.
- [14] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox, "RGB-D mapping: Using Kinect-style depth cameras for dense 3D modeling of indoor environments," *Int. J. Robot. Res.*, vol. 31, no. 5, pp. 647–663, 2012.
- [15] A. V. Segal, D. Haehnel, and S. Thrun, "Generalized-ICP," in *Proc. Robotics: Science Systems*, 2009.
- [16] N. Gelfand and S. Rusinkiewicz, "Geometrically stable sampling for the ICP algorithm," in *Proc. Int. Conf. 3D Digital Imaging Modeling*, 2003, pp. 260–267.
- [17] C. Harris and M. Stephens, "A combined corner and edge detection," in *Proc. 4th Alvey Vision Conf.*, 1988, pp. 147–151.
- [18] F. Tombari, S. Salti, and L. D. Stefano, "Performance evaluation of 3D key-point detectors," *Int. J. Comput. Vis.*, vol. 102, nos. 1–3, pp. 198–220, 2013.
- [19] Y. Zhong, "Intrinsic shape signatures: A shape descriptor for 3D object recognition," in *Proc. ICCV Workshop 3D Representation Recognition*, 2009, pp. 689–696.
- [20] B. Steder, R. Rusu, K. Konolige, and W. Burgard, "NARF: 3D range image features for object recognition," in *Proc. IROS Workshop Defining Solving Realistic Perception Problems Personal Robotics*, 2010.
- [21] A. Aldoma, Z. Marton, F. Tombari, W. Wohlkinger, C. Potthast, B. Zeisl, R. Rusu, S. Gedikli, and M. Vincze, "Point cloud library: Three-dimensional object recognition and 6 DOF pose estimation," *IEEE Robot. Automat. Mag.*, vol. 19, no. 3, pp. 80–91, 2012.
- [22] L. Alexandre, "3D descriptors for object and category recognition: A comparative evaluation," in *Proc. IROS Workshop Color-Depth Camera Fusion Robotics*, 2012.
- [23] A. Johnson and M. Hebert, "Using spin images for efficient object recognition in cluttered 3D scenes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 21, no. 5, pp. 433–449, 1999.
- [24] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms for 3D registration," in *Proc. IEEE Int. Conf. Robotics Automation*, 2009, pp. 1848–1853.
- [25] A. Frome, D. Huber, R. Kolluri, T. Bülow, and J. Malik, "Recognizing objects in range data using regional point descriptors," in *Proc. European Conf. Computer Vision*, 2004, pp. 224–237.
- [26] F. Tombari, S. Salti, and L. Di Stefano, "Unique shape context for 3D data description," in *Proc. ACM Workshop 3D Object Retrieval*, 2010, pp. 57–62.
- [27] F. Tombari, S. Salti, and L. Di Stefano, "Unique signatures of histograms for local surface description," in *Proc. European Conf. Computer Vision*, 2010, pp. 356–369.
- [28] D. Holz and S. Behnke, "Mapping with micro aerial vehicles by registration of sparse 3D laser scans," in *Proc. Int. Conf. Intelligent Autonomous Systems*, 2014.
- [29] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohli, J. Shotton, S. Hodges, and A. Fitzgibbon, "KinectFusion: Real-time dense surface mapping and tracking," in *Proc. IEEE Int. Symp. Mixed Augmented Reality*, 2011, pp. 127–136.
- [30] S. Holzer, R. B. Rusu, M. Dixon, S. Gedikli, and N. Navab, "Adaptive neighborhood selection for realtime surface normal estimation from organized point cloud data using integral images," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Systems*, 2012, pp. 2684–2689.

**Dirk Holz**, University of Bonn, Bonn, Germany. E-mail: [dirk.holz@ieee.org](mailto:dirk.holz@ieee.org).

**Alexandru-Eugen Ichim**, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland. E-mail: [alexandru.ichim@epfl.ch](mailto:alexandru.ichim@epfl.ch).

**Federico Tombari**, University of Bologna, Bologna, Italy. E-mail: [federico.tombari@unibo.it](mailto:federico.tombari@unibo.it).

**Radu B. Rusu**, Open Perception, Inc., California, United States. E-mail: [radu.b.rusu@gmail.com](mailto:radu.b.rusu@gmail.com).

**Sven Behnke**, University of Bonn, Bonn, Germany. E-mail: [behnke@cs.uni-bonn.de](mailto:behnke@cs.uni-bonn.de).

