An automated approach for point cloud alignment based on density histograms

Jan Martens, Jörg Blankenbach
Geodetic Institute and Chair for Computing in Civil Engineering & Geo Information Systems, RWTH Aachen University, Germany
jan,martens@gia.rwth-aachen.de, blankenbach@gia.rwth-aachen.de

Abstract. Building Information Modeling is growing more relevant as digital models are not only used during the construction phase but also throughout the building's life cycle. The digital representation of geometric, physical and functional properties enables new methods for planning, execution and operation. Digital models of existing buildings are commonly derived from surveying data such as laser scanning which needs to be processed either manually or automatically throughout various steps. Aligning point clouds along a coordinate system's main axes is a common manual task benefitting subsequent manual and automated processing steps. With the intention of automating the alignment task, we hereby present an enhanced approach based on point density histograms. Our results show that this approach is computationally cheap, robust towards non-uniform point cloud resolutions and clutter as well as easy to integrate into existing BIM modelling software.

1. Introduction

The rise of Building Information Modeling (BIM) in recent years has led to a gradual paradigm shift in the AECO industry. With BIM providing a digital representation of all geometrical, physical and functional characteristics of buildings and constructions, it offers the opportunity of sharing information and tracking object states throughout the building life cycle. Therefore, object-oriented models, so called BIM models, are used. Up to now BIM has been predominantly used during the planning phase. However, solutions for using BIM during the execution and operation phase gain an increasing relevance. The benefits of using them e.g. for facility management (FM) or building preservation has earned BIM methods wide-spread interest. Problematically, creating BIM models from existing buildings without the existence of digital data or even digital models from planning is a costly procedure and requires trained staff. For this purpose, recent surveying techniques such as laser scanning or photogrammetry promise a lot of potential for as-is/as-built data capturing since they allow reality capturing with a high spatial resolution in a short period of time. However, laser scanning results in huge amounts of data in the form of point clouds which have to be analyzed for BIM modeling in a time-consuming process that is still characterized by much manual work. This has led to increased interest in techniques which either fully or partially automate the modeling process from point clouds.

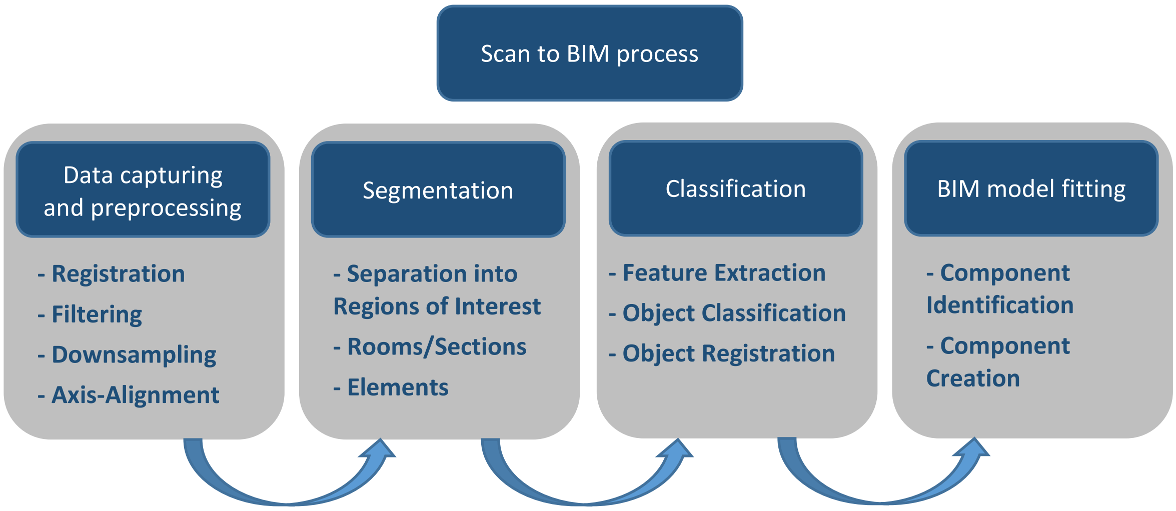

Figure 1: Stages of the “Scan to BIM“-process.

The process of automatic BIM model generation from point clouds (scan to BIM) can roughly be subdivided into the phases preprocessing, segmentation, classification and BIM model fitting as outlined in the adapted Figure 1 following Loges & Blankenbach (2017). Data preprocessing generally represents the point cloud processing step performed immediately after data capturing. The general goal of preprocessing is to denoise the point cloud by removing outliers, smoothing out noise and possibly downsampling and transforming it into a suitable format for later computations. Furthermore, this step can also be beneficial for manual modelling and visualization, as preprocessed data eases the modeling of structures, reduces the risk of potential errors and is generally more pleasant to deal with.

In this work, we focus on point cloud alignment, as it is most often the first operation done in manual modeling and carries great importance in automated reconstruction. Aligning a point cloud or pose normalization refer to rotating the point cloud such that at least one of the dominant building axes captured with it, is oriented along one of the global coordinate system axes. In the field of automation, the benefits of the alignment process lie in the fact that most data structures which ease point neighborhood lookups during point cloud processing such as voxel grids and octrees use boundary volumes oriented along the global coordinate system axes. Therefore, axis-aligned point clouds ensure optimal compactness and therefore usage of these structures. An example on how proper alignment reduces aliasing artifacts which might otherwise occur for voxel-based feature extraction is illustrated in Figure 2. Whereas in the aligned point cloud edges derived from the voxel-wise scatter values are clearly visible, edges in the unaligned point cloud are harder to identify and suffer from visible noise. This problem is commonly caused by planar surfaces which are not properly aligned with the axis-aligned voxels of the used grid.

Aside from benefits for data structures and some feature extraction techniques, many successful methods for segmentation and analysis make use of the Manhattan World assumption, which assumes structures contained in the point cloud to be oriented along a rectangular grid. Determining the global orientation of the grid and rotating the point cloud to match this grid thus represents an important precondition for such algorithms to work properly.

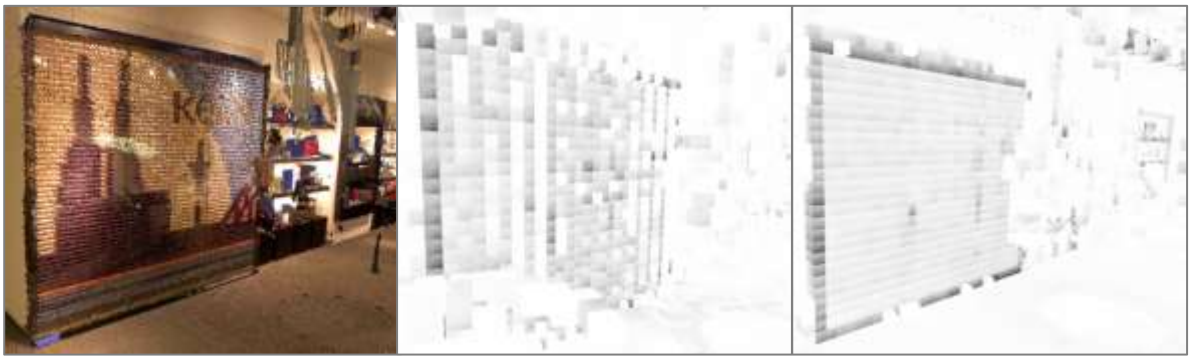

Figure 2: Impact of axis-alignment on voxel-based calculations. Left: Original point cloud. Center: Scatter values per voxel of un-aligned point cloud (contrast enhanced for visibility). Right: Scatter values per voxel of axis-aligned point cloud (contrast enhanced for visibility).

2. Related Work

The problem of pose normalization for 3D shapes and meshes has been extensively discussed in other works. However, publications in this field for point clouds are either sparse or concerned with solving different, related problems. Many works in the field of pose normalization commonly deal with alignment as a preprocessing step for shape matching or discuss it in context of feature extraction. Lian et al. (2010) defined a metric for 3D mesh rectilinearity which was dependent on the mesh's pose. Since their metric would measure the projected area of a mesh in relation to the total area, it could also be used to find suitable mesh orientations. An adaptation of their algorithm for point clouds does not exist though. Works done by Podolak et al. (2006), Paquet et al. (2000), Chaouch and Verroust-Blondet (2008) and Kazhdan et al. (2003, 2007) indicate that both symmetry and the presence of reflection planes help estimate correct rotation and object poses, however, laser scans of entire buildings rarely follow these assumptions.

Many of aforementioned works compared the pose normalization problem together with the principal component analysis (PCA). Most often the PCA is being used for feature extraction, but also represents a common way for calculating an object's principal axis. This is achieved by constructing the covariance matrix of a point set, followed by an eigenvalue decomposition of this matrix resulting in an orthogonal basis of principal axes. These principal axes, however, are often sensitive to noise and local sampling such that only the most dominant principal axis ends up being useful for shape alignment. In their publications, Kazhdan et al. (2007) compared multiple feature-based 3D pose retrieval methods for meshes using PCA as the baseline for evaluating their experiments. As pointed out and thoroughly discussed in their works, PCA does not represent a viable way of aligning objects along the axes of a global coordinate system. In context of point cloud processing, Okorn et al. (2010) used point density histograms for the segmentation of floor and ceiling planes and to estimate dominant orientations of cross sections of single rooms rather than entire point clouds. Aside from the limitation to single rooms which requires said rooms to be presegmented, this method requires the selected cross sections to be free of clutter. The method thus requires careful selection of a suitable cross section, limiting it even further to a handful of scenarios.

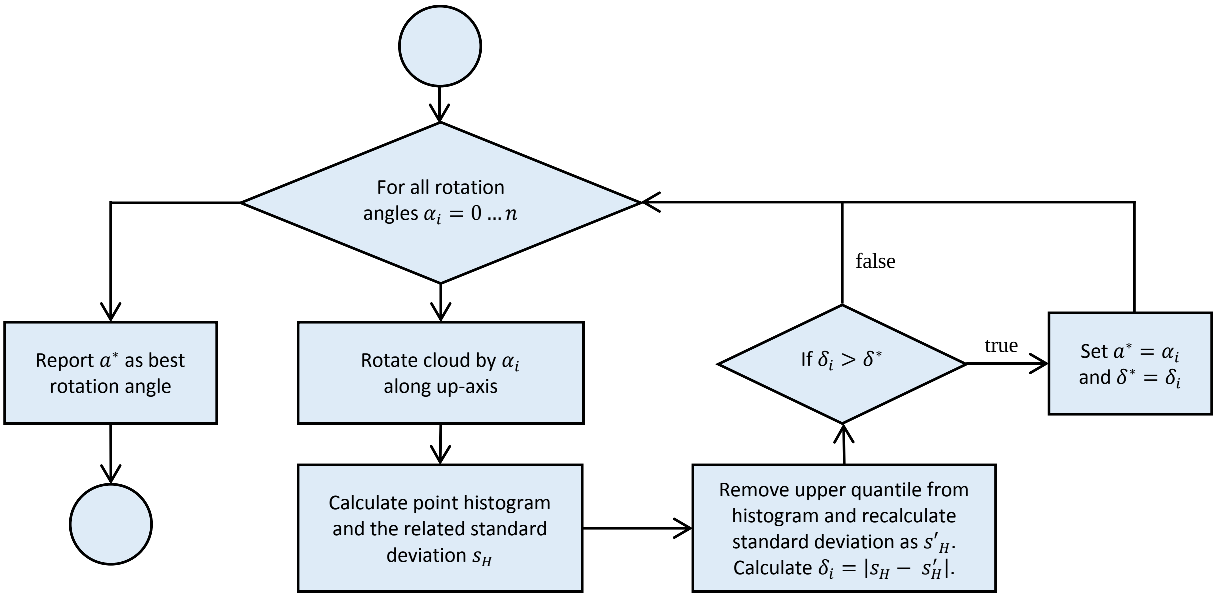

Figure 3: Workflow diagram of our point cloud alignment method.

Combining the ideas of both, symmetrical properties and dominant planes, Fu et al. (2008) presented an approach which would construct a convex hull from an input mesh and identify the dominant plane of the hull polygon. Due to the reliance on a convex hull, this method could be adapted for use with point clouds and furthermore indicates that the detection of planes in man-made structures plays an integral role in object alignment.

3. Methods

Our approach to point cloud alignment makes use of point density histogram to identify dominant wall segments, but contains multiple analysis steps shown in Figure 3. We initially start by defining a vector pointing “upwards” as our rotation axis and selecting an axis for alignment (which is oriented orthogonally to the upwards vector). This vector is typically chosen to be the global z-axis which is usually fixed in case of point cloud data captured with a terrestrial laser scanner in static mode. Afterwards, we incrementally rotate the point cloud and construct the density histograms. These histograms are created by counting the number of points falling into a specific range along the alignment axis. The shape of these histograms strongly depends on the chosen rotation, a behaviour explored further in Figure 3. As point densities are larger along planar surfaces such as walls, clearly identifiable peaks will emerge if a viable rotation angle for alignment has been chosen. Contrary, an unsuitable rotation angle results in a rather flat distribution in the density histogram. Based on this observation, the alignment problem is reduced to finding dominant peaks in the density histogram as their presence indicates that walls are aligned with the main axes. As a fast and robust metric for detecting histograms related to the desired rotation angle, the per-bin standard deviation s_H of the point distribution histogram is being calculated. Afterwards, the histogram bins falling into an upper quantile are being removed. In case of histograms related to suitable angles, this operation removes peaks and in consequence greatly reduces the standard deviation of the point density distributions. The difference δ_s between the original standard deviation s_H and the new standard s'_H deviation is used as an indicator for the desired histogram and therefore the desired rotation angle. The desired rotation angle in the search space thus maximizes the calculated difference δ_s of the standard deviations:

$$\delta_s = \max(|s_H - s'_H|)$$

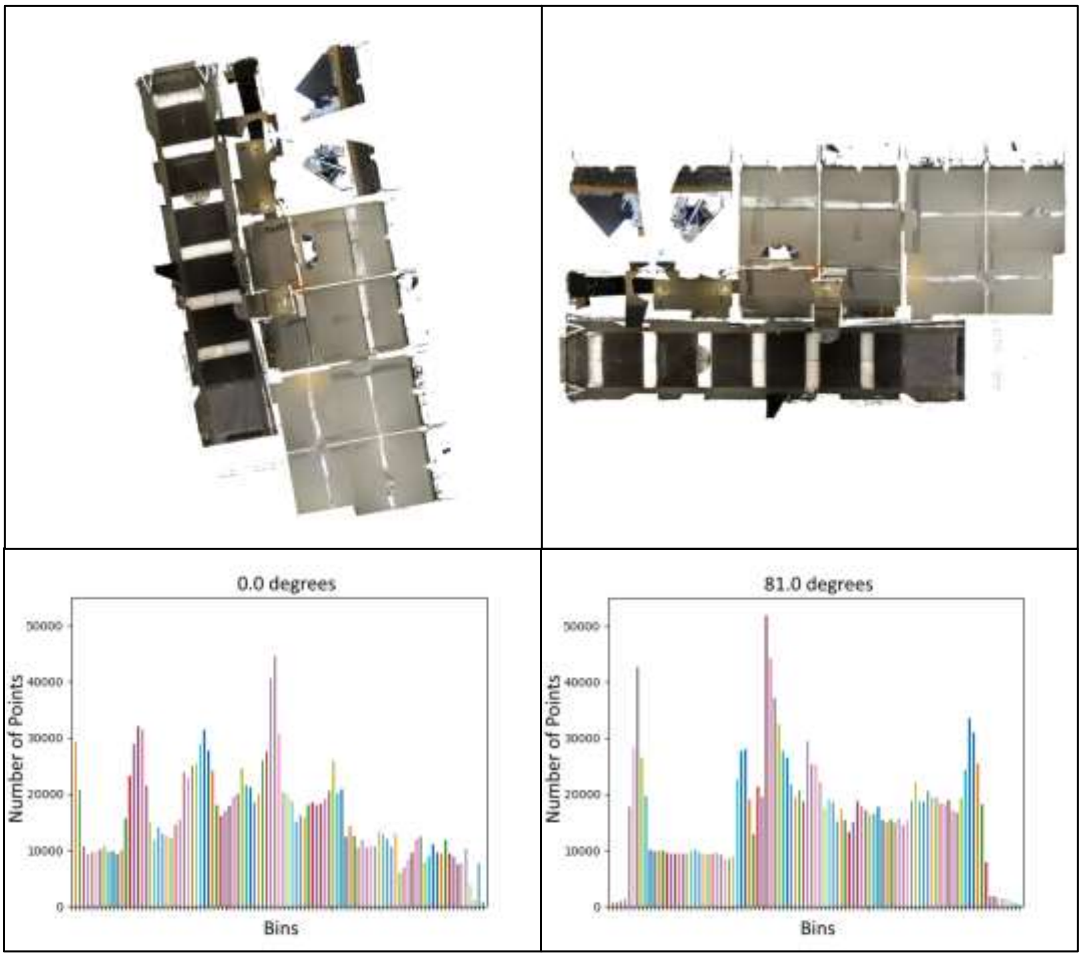

Figure 4: Comparison of density histograms for point cloud with 1,629,143 points. Left column: Top view of unaligned point cloud and correspong density histogram. Right column: Top view of aligned point cloud and correspong density histogram. Note how histogram peaks are more dominant in the aligned point cloud, with noticable flat regions.

Even though searching for the largest peak in the density histograms seems like a simpler option, we decided to use aforementioned method as it has proven to be more robust towards local point density variations and clutter. If desired, this alignment step may be repeated for the remaining axes. To ensure that the point cloud is still roughly located at the same position, all rotations are performed around the point cloud's centre of gravity.
The described algorithm depends on three parameters which are the resolution of the rotation angle, the number of histogram bins and upper quantile size.

In case of the number of histogram bins, it is advisable to choose their amount such that the resulting bin sizes roughly correspond to a wall's thickness. The size of the rotation angle increments on the other hand has an impact on the size of the search space and therefore on the algorithm's running time. In case of the point cloud with 1,629,143 points in Figure 4, we observed an overall execution time of around 8.408s for an angular resolution of 1 degree, but a significantly reduced execution time of 1.762s for an angular resolution of 5 degrees. The accuracy also depends on the upper histogram bin quantile, for which we found a value of 90% to be adequate for our experiments. In our example, removing this upper quantile from the histogram results reduces the overall standard deviation of the remaining histogram bins by 6,313.613 points.

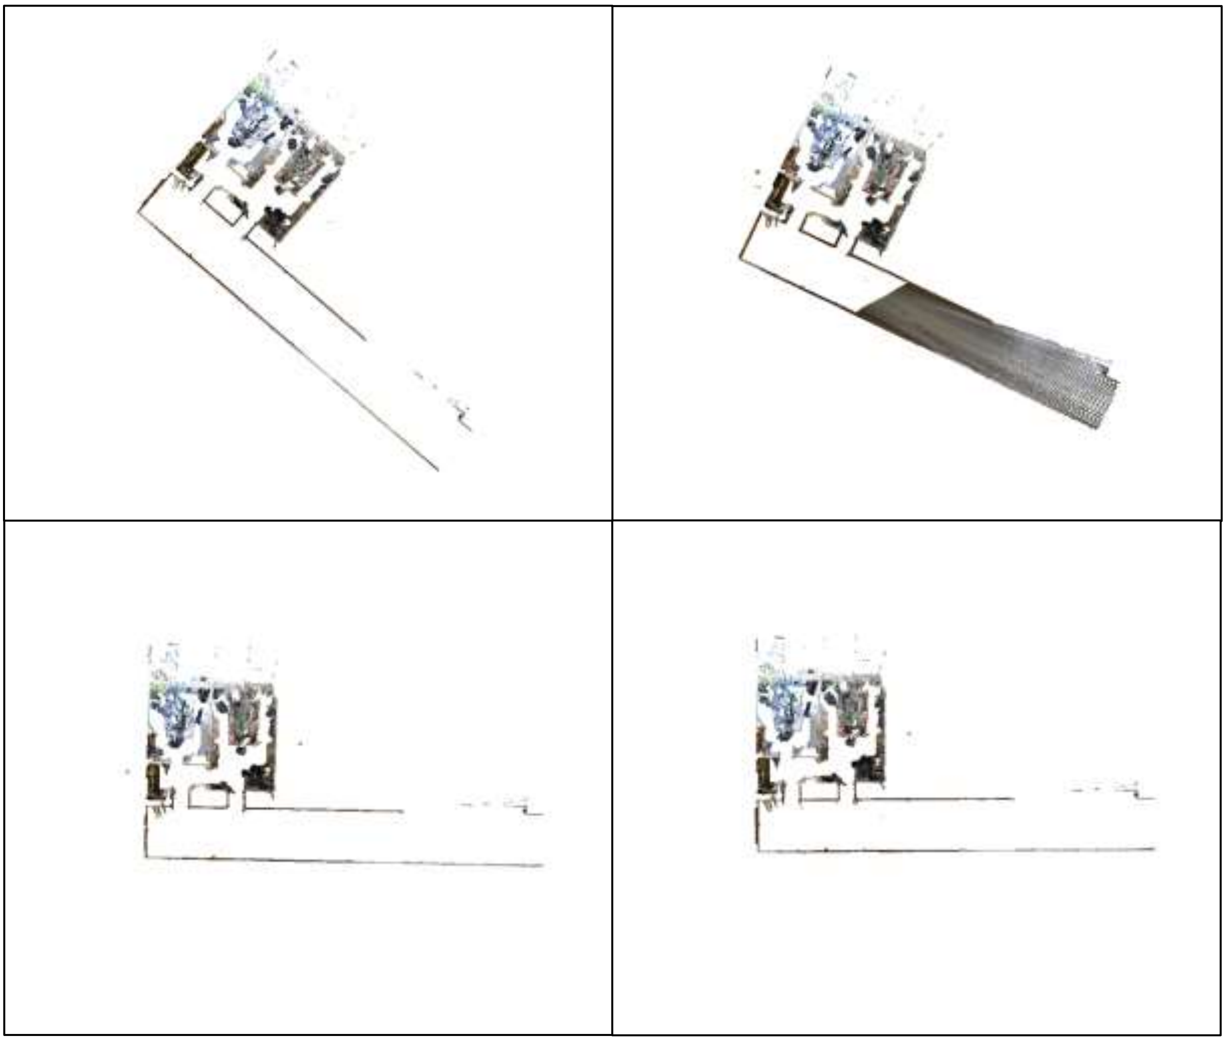

Figure 5: Top view of cross sections of original unaligned and axis-aligned point clouds. Top left: Original point cloud. Top right: The PCA-based alignment method is fast (approx. 0.116s), but fails to find the desired rotation angle. Alignment may even tilt the cloud, as apparent by the floor points which have been moved into the displayed cross section plane. Bottom left: Alignment with the histogram-based method. Using 10 bins results in a minor error. Bottom right: Using 100 bins for the histogram-based method provides excellent accuracy.

4. Results

We evaluated our method on various data sets (further examples shown in figure 6) and found that it performs desirably for evenly sampled point clouds. In direct comparison, it performs considerably better than an alignment based on the principal component analysis (PCA). For evaluation purposes, we implemented our algorithm and a PCA-based alignment algorithm as command line tools and subsequently embedded them both into Autodesk Revit as plugins, making them available inside a BIM-enabled modeling environment. It is important to point out that our experiments in Autodesk Revit were limited to 1,000,000 points due to API restrictions, while the command line tool has no such point count limits. Figure 5 displays the result of our algorithm in Autodesk Revit and compares it to the results of PCA-based alignment. As displayed in the figure, our method is capable of identifying the desired rotation angle even for point clouds where the PCA-based method struggles. Not only are point clouds being rotated as intended, but the algorithm is also capable of dealing with the uneven sampling mentioned earlier. Aligning subsampled point clouds also gives satisfactory results, meaning that accelerating the alignment by using subsampled data sets of the original point cloud is a viable option. Moreover, we found the algorithm capable of dealing with scans spanning along multiple floors. We also evaluated our method on point clouds acquired from different sensors with results being listed in Table 1. All show an uneven resolution with locally high point densities.

The PCA-based alignment method on the other hand merely calculates a local frame for a point cloud and aligns it with the global frame. This method is significantly faster in comparison, often finishing in less than a second, but delivers poor results as regions with high point densities will have a higher impact on the rotation. Another problem arises from laser scans where only one section of the building has been captured. In these cases, the scanned rooms have a bigger influence on the calculation of the principal axis than sparsely scanned corridors.

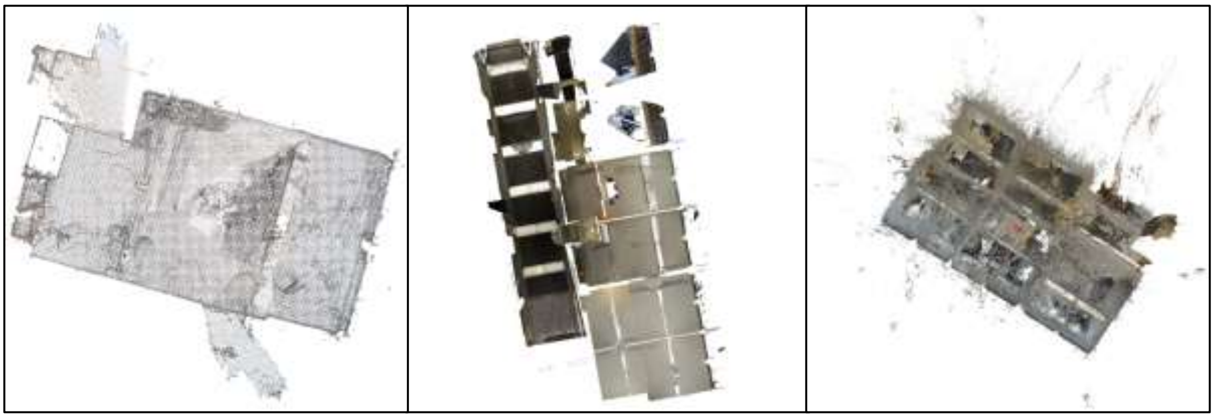

Figure 6: Visual comparison of point clouds used for evaluation (high-resolution versions of these images are shown in the appendix). Left: Single office captured with Google Tango device. Notice the low spatial resolution and rather even point distribution. Center: Multiple offices captured with Riegl VZ400 laser scanner. Overall scan resolution is high, and very dense close to the scanning device. Right: Point cloud reconstructed from image matching. Characteristics are the strong noise and fluctuation point densities.

Table 1: Execution times for PCA- and histogram-based algorithms with an angle resolution of 1° for various point clouds. The first cloud was acquired with a Google Tango device, while “Riegl1” and “Riegl2” have been captured with a geodetic laser scanner (Riegl VZ400); both show an uneven point distribution characteristic for TLS. The final cloud has been captured via image matching.

Point Cloud	Number of Points	Time (PCA) [s]	Time (Histogram) [s]
Tango	1,629,143	0.031	0.874
Riegl1	5,702,432	0.203	8.408
Riegl2 (original)	17,934,986	3.558	93.393
Riegl2 (Revit)	1,000,000	0.116	5.326
Image Matching	4,579,656	0.515	23.743

5. Conclusion and Outlook

We presented an enhanced method for the automated alignment of point clouds to a global coordinate system. Our method has proven to be fast and robust towards uneven sampling and delivers more precise results than the PCA-based method. Although we have already carried out extensive test runs with various point clouds of building, we intend to evaluate our method also with point clouds of engineering structures (e.g. bridges). With the algorithm being based on point density histograms, our brute-force algorithm could be optimized further by means of using more sophisticated optimization methods such as gradient descent, Monte Carlo methods or genetic algorithms. Large point clouds in particular would benefit from this extension. Modifying the algorithm to use histograms based on voxels rather than points might be an option to make it more robust towards densely sampled regions.

Interestingly, first trials of our command-line tool indicated that pose normalization with respect to multiple axes seems possible as well, with the alignment being performed successively for each axis. However, further evaluations need to be performed to verify this assumption. Aside from these technical aspects, an option for automated parameter selection would be useful to simplify user input and manual parameter selection.

6. References

Loges, S., Blankenbach, J. (2017), As-built Dokumentation für BIM – Ableitung von bauteilorientierten Modellen aus Punktwolken, Photogrammetrie – Laserscanning – optische 3D-Messtechnik: Beiträge der Oldenburger 3D-Tage 2017.

Okorn, B. E., Xiong, X., Akinci, B., Huber, D. (2010), Toward Automated Modeling of Flor Plans, Proceedings of the Symposium on 3D Data Processing, Visualization and Transmission.

Lian, Z., Rosin, P., Sun, X. (2010), Rectilinearity of 3D meshes, International Journal of Computer Vision.

Chaouch, M., Verroust-Blondet, A. (2008), A Novel Method for Alignment of 3D Models, IEEE International Conference on Shape Modeling and Applications, Proceedings.

Fu, H., Cohen-Or, D., Dror, G., Sheffer, A. (2008), Upright Orientation of Man-Made Objects, ACM Transaction on Graphics, Proceedings of SIGGRAPH.

Podolak, J., Shilane, P., Golovinskiy, A., Rusinkiewicz, S., Funkhouser, T., (2006), A Planar-reflective Symmetry Tranform for 3D Shapes, ACM Transactions on Graphics (TOG)

Kazhdan, M., Funkhouser, T., Rusinkievicz, S. (2003), Rotation Invariant Spherical Harmonic Representation of 3D Shape Descriptors, Symposium on Geometry Processing.

Kazhdan, M., (2007), An Approximate and Efficient Method for Optimal Rotation of 3D Models, IEEE Transaction on Pattern Analysis and Machine Intelligence.

Paquet, E., Rioux, M., Murching, A., Naveen, T., Tabatabai, A., (2000), Description of Shape Information for 2D and 3D Objects, Signal Processing: Image Communication.

7. Appendix

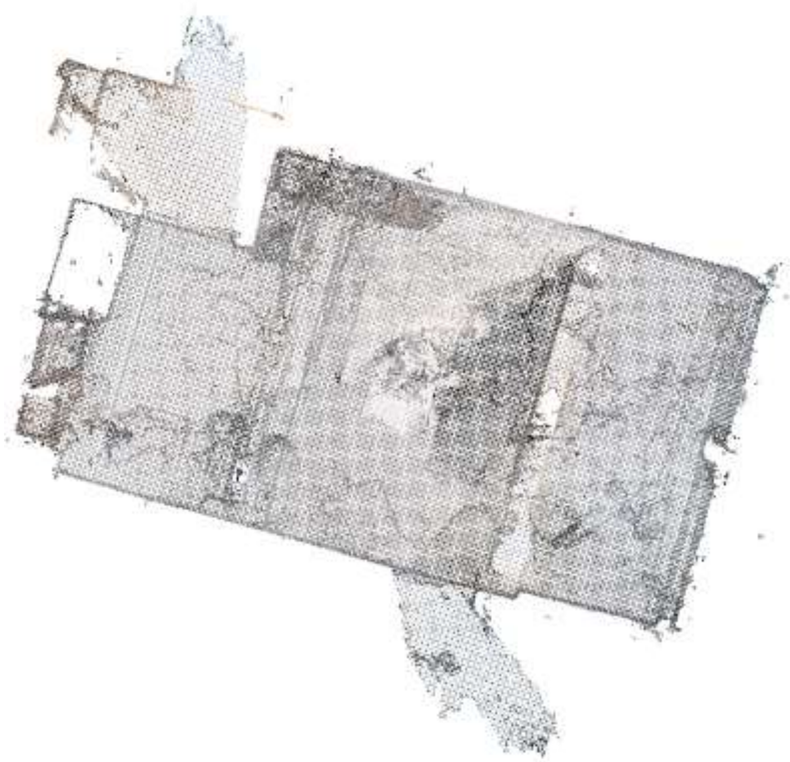

Figure 7: Point cloud acquired with Time-of-Flight camera of Google Tango device. Point distributions are even, but very coarse.

Figure 8: Point cloud reconstructed from multiple registered scans captured with a Riegl VZ400 laser scanner. The overall resolutions is very fine, but suffers from density fluctuations close to the scan positions.

Figure 9: Reconstructed point cloud created through image matching. Point densities vary greatly with points scattered along their respective surfaces. Heavy noise and low point densities in regions such as the ceiling are quite characteristic.