

Intuitive Alignment of Point-clouds with Painting-based Feature Correspondence

Shane Transue
Computer Science and Engineering
University of Colorado Denver
shane.transue@ucdenver.edu,

Min-Hyung Choi
Computer Science and Engineering
University of Colorado Denver
min.choi@ucdenver.edu

Abstract. Throughout the course of several years, significant progress has been made with regard to the accuracy and performance of pair-wise alignment techniques; however when considering low-resolution scans with minimal pair-wise overlap, and scans with high levels of symmetry, the process of successfully performing sequential alignments in the object reconstruction process remains a challenging task. Even with the improvements in surface point sampling and surface feature correspondence estimation, existing techniques do not guarantee an alignment between arbitrary point-cloud pairs due to statistically-driven estimation models. In this paper we define a robust and intuitive painting-based feature correspondence selection methodology that can refine input sets for these existing techniques to ensure alignment convergence. Additionally, we consolidate this painting process into a semi-automated alignment compilation technique that can be used to ensure the proper reconstruction of scanned models.

1 Introduction

Research related to the process of reconstructing a three dimensional virtual object from a set of independent scans or range-images has progressed to the point where pairwise alignment between image pairs is a mature process. Considerable effort has been invested in the development of techniques that derive the proper alignment between two individual point-clouds and these techniques continue to evolve and become more accurate. However, there still remain several critical factors that contribute to successfully obtaining a proper alignment between point-cloud pairs to completely reconstruct a virtual object that closely mirrors the object in reality. These factors include scan resolution, noise, and pair-wise scan alignment convergence.

Several alignment algorithms have been proposed to perform pair-wise alignments between point-clouds. Each of these proposed algorithms utilizes different techniques for finding correspondence between the features of the surface of the scanned object in each

point-cloud. Although existing techniques can estimate an initial alignment of the two scanned surfaces, and minimize the distance between the points within the point-clouds for a refined alignment, nothing guarantees these algorithms will converge. For an implementation that ensures a proper alignment and exhibits a high level of flexibility with respect to noise, erroneous surface property characteristics, and background information, we must utilize additional methods for finding the corresponding features within a point-cloud pair. We propose the introduction of a complete object reconstruction methodology that addresses the lack of cohesion between existing pair-wise alignment techniques and introduces an intuitive method for effectively reconstructing virtual models from point-cloud sets.

In this paper we introduce an intuitive painting-based selection technique for feature correspondence that can be used to aid in the flexibility and robustness of the alignment between low-overlap scan pairs. We then consider how this contribution can aid in the process of ensuring that alignments are accurate enough for estimate techniques to fall within the convergence threshold existing refined alignment algorithms as part of an alignment compilation process. This process allows for the rapid and robust reconstruction of a virtual object model utilizing several point-clouds without physical markers, scanner position tracking or viewing restrictions.

2 Pairwise Point-Cloud Alignment

Over the course of several years there has been widespread research [3] into the process of aligning two point-clouds collected from 3D scanning devices. A significant amount of this work is originally based off of the work by Besl and McKay [2] with the introduction of the iterative closest point (ICP) algorithm. The concept behind the ICP algorithm, along with its many variants [7], is based on the notion that the distance between the two point-clouds should be minimized over the course of several iterations. This process is required to perform a refined pair-wise alignment between two point-clouds and has been extensively utilized for 3D object reconstruction.

Based on the process of aligning point-clouds utilizing the distance minimization approach introduced by the ICP algorithm, several critical problems can be identified when the algorithm or its variants are utilized for object reconstruction. Instances of symmetry and a lack of ideal scanning conditions greatly degrade the success and accuracy of the alignment process. We can also identify the erroneous contribution of all background information that does not directly define the surface of the scanned object. All of this information must be removed and discarded prior to the alignment process and cannot always be represented as statistical outliers. Similarly, the introduction of point-cloud measurement errors and scanned material properties all contribute to the misalignment of the images and must be properly handled before an attempted alignment.

Prior variants of the ICP algorithm aim to selectively address some of these concerns. Yet, several additional parameters must be considered prior to the application of this

alignment technique. These include the positional and rotational offsets of the scanning device, perspective projection distortions, and the notion of pair-wise scan overlap. In general, the technique utilized by the ICP algorithm is unsuited to address the alignment between scan pairs that exhibit vast differences in these additional parameters. However, since the ICP algorithm works well within a set of provided constraints (correct initial orientation, close proximity, proper sampling) and provides an accurate final alignment of two point-clouds, we categorize this process as a refined pair-wise alignment technique.

Based on the limitations that have been imposed by utilizing the ICP algorithm for general point-cloud alignment and the notion of having the point-clouds initially closely aligned, alternative approaches have been introduced to address these concerns. The introduction of persistent feature histograms [6] and point feature histograms (FPH) for three dimensional alignments by Rusu et al. [1] for Same Consensus Initial Alignment (SAC-IA) have been formulated to provide the estimate alignment required by the ICP algorithm. These techniques fall within what we defined as the classification of estimate alignment techniques. The key notion behind these contributions is that they utilize specific orientation independent surface features to estimate an initial (guess) alignment between the scan pairs using feature correspondence. The objective of these approaches is to orient the scan pairs such that they will fall within the convergence threshold of the ICP algorithm. In the context of constructing a virtual model from a set of point-clouds, the alignment process [5] is defined through an iterative process of first estimating the correspondence between two scans and then obtaining a highly accurate alignment through the utilization of a refined alignment technique.

3 Point-cloud Acquisition

The process that defines the collection of independent point-clouds from a scanning device dictates the requirements for the alignment techniques utilized in the construction of a virtual model. Based on the available information about the scan (position, rotation, visual markers, etc), the alignment technique may vary considerably.

In our approach we assume that when a point-cloud is generated or collected from a 3D scanning device, none of this additional information is present. The point-cloud only contains the explicit data that describes the surface of the object being scanned. Furthermore, we assume that the point-cloud may contain all of the original background information, erroneous data introduced by reflective and transparent materials, and statistical outliers. From these assertions we can drastically reduce the requirements related to the point-cloud acquisition process. The image in Figure 1 (left) provides an illustration of an individual point-cloud that contains the point-cloud data of the scanners view along with the RGB value associated with each point. The two right images in Figure 1 illustrate two edited point-clouds of the same vehicle from two different viewing angles, where the background data has been eliminated. This process is explicitly required for most existing alignment techniques; however in our case this is not required.



Fig 1. An example of an individual point-cloud collected with background information present in the time of flight (TOF) scan (left). The front right corner (left) and the rear left corner (right) contain a correspondence overlap of approximately 60%

4 Painting-based Feature Correspondence

Computationally, the process of identifying similar features between two images is an ongoing research topic in the field of computer graphics and object recognition [4, 11]. When considering the process that we mentally follow when trying to identify an object in two images, we look for edges and unique features that visually stand out from the surrounding surfaces. If we can identify a unique object in an image, it is common to be able to identify the same object within another image from a different perspective. In this section we introduce an intuitive method for identifying objects and features that exist within both images of a point-cloud pair and allow for the dynamic selection of these features for pairwise alignment based on this intuition.

The ability to identify objects and surface features that exist in a point-cloud pair is directly related to feature correspondence which has been utilized in recent estimate alignment techniques (such as SAC-IA). However instead of using statistical analysis, we look at human intervention where objects within a scene can be visually identified by unique features, even from multiple angles. For example, in Figure 1, the wheels of the vehicle can be visually identified from both viewing angles. Intuitively we can identify that the wheels exist in both scans and correspond to the wheels of the scanned vehicle. This is an example that represents an orientation independent feature that is common to both scans. When we attempt to align scans that contain uncommon features, it can be difficult for pair-wise estimate alignment techniques to converge. Therefore our objective is to interject human intervention into this process when these estimate alignment techniques fail to obtain an alignment within the required ICP convergence threshold. Additionally, we extend this approach to handle even more complex alignment scenarios where features are hard to objectively identify as shown in Figure 2.



Fig. 2. Illustration of two point-clouds defining a dirt pile surface that lacks unique point-features that would be used for existing key-point selection techniques to generate an estimate alignment.

Based on the observation that similar features can be identified from various viewing angles, we introduce the notion of selecting these features as a subset of points that are prime candidates for alignment. This is where we introduce the notion of painting-based selection. Visually the term paint means to highlight or select the points that will be used in an alignment. These points are then made visually distinct in size and color. Similar techniques feature point-selection for identifying exact key-point correspondences between scans; however in an instance where the geometry does not present a clear feature that can be identified by a point between each scan, this is extremely challenging. The image in Figure 2 depicts two scans of a dirt pile from two perspectives where using a feature-point selection technique would be challenging due to the lack of a unique feature in each scan. Therefore with our technique we can simply paint over similar curvatures in the surface common to both of these scans to identify a rough feature correspondence.

In our proposed approach it is vital that the surface curvature and image data is readily visible. To aid in the visual clarity of the displayed point-clouds we provide both point-based normal shading with estimated surface normals and include support for an RGB color channel. This allows for the curvature of the scanned surface to be readily visible to the user. The image in Figure 3 provides an illustration of the selected corresponding vehicle features (windows and wheels) that overlap between the two point-clouds. While the painted regions do not have to exactly correspond to the features within each scan, the ability to identify these features is a critical aspect of our approach.



Fig 3. The prominent features in common between these two scans have been selected (windows and wheels) for alignment. These are the features common to both point-clouds and the selection ensures that the pairwise alignment process is less prone to becoming trapped in a local minima.

The ability to explicitly identify unique characteristics of a scanned objects surface from multiple point-clouds requires additional developments in usability engineering to make the process intuitive. This is due to the integration of human intervention into the process of selecting key features (as opposed to key-points) as a pre-processing step for pair-wise alignment. This means that only the selected data-points will be utilized for the alignment process and all other points will simply be ignored as if they were not part of the point-cloud (user directed sub-sampling).

To aid in the accessibility of selecting the features or regions that should be utilized in a pairwise alignment, we introduce an interactive painting tool that dynamically selects the data-points of a point-cloud in 3D space based on user selection. This is accomplished with through the use of a circular paint-brush that is used to define a cylindrical volume in 3D space. As the brush is moved throughout the point-cloud, the data-points encapsulated in the swept volume will be nominated for alignment. Considering the flexibility of this

tool we also include the ability to modify the radius and length of this cylindrical volume to account for fine details and also provide a similar means for deselecting data-points to refine the selected feature regions. The circle within the image represents the selection region that defines what points will be added to the candidate point set for alignment.

The images in Figure 4 illustrate how a quick brush-stroke can be used to both select and deselect features from two example point-clouds collected with 3D scanning devices. The user can simply click and drag the mouse to produce these painting effects.

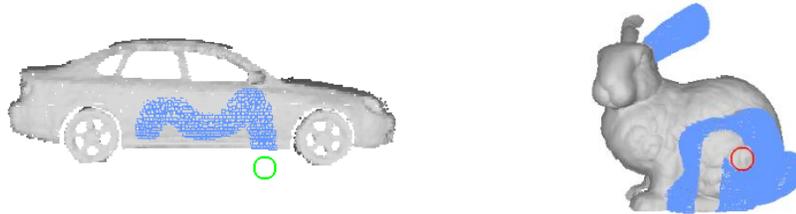


Fig 4. Brush-based painting tool that selects the data-points within the radius of the brush-selection circle as it is moved across the surface of the scanned object. The unique brush trail left on the surface of the scan illustrates the points selected for alignment (left). The deselection brush allows for points to be removed from the alignment (right).

The ability to control which features are selected for alignment provides an extensive level of flexibility for the input of an estimate alignment algorithm. In contrast to existing alignment techniques [10], this provides alignment parameter modification at a higher level. If we consider scan data that contains erroneous measurements or extensive background information, we can simply exclude those samples from the alignment process. This will completely remove the error in the alignment that these additional data points generate. This is due to an attempt to align features that only exist within one scan. This will only increase the error metric of the resulting alignment.

The images in Figure 5 illustrate the selection of the common features of the focus object in two point-clouds. The painted data indicates the points that will be utilized for the pairwise alignment between the two scans containing extensive background information.

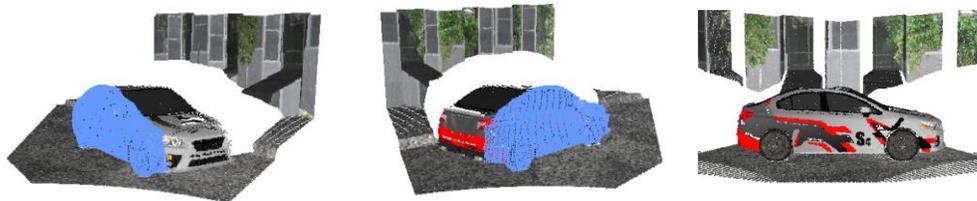


Fig 5. Two scans (left, center) from different perspectives that contain extensive background information. The ability to paint the corresponding features that are common to each scan and the focus object allows for an alignment (right) to be performed with directly corresponding points.

5 Scan Alignment Compilation

The separation of an alignment methodology to include two separate classifications of pairwise alignments (estimate/refined), does not inherently guarantee the success of the resulting alignment. However due to performance requirements, most existing alignment techniques utilize some form of sampling to reduce the number of data-points directly involved in a pairwise alignment. Generally uniform sampling is utilized to sample the entire point-cloud for the operation; however, other approaches [9] attempt to identify regions with steep gradients to improve alignment accuracy. Yet these more sophisticated approaches sample the entire point-cloud which may have non-corresponding features.

The problem that statistically-based techniques introduce is that when they are used consecutively, the output of the estimate alignment technique must be used as input for the refined alignment. Therefore if the estimate alignment diverges or encounters a local minima that does not properly align the provided point-clouds, the initial alignment constraint identified for the refined alignment process has been violated. This will result in an invalid refined alignment and thus make the entire process depend on a successful result of the estimate alignment. Since this process includes several pair-wise alignments we must ensure that each is successful before continuing. This introduces the concept of an alignment checkpoint that allows for verification techniques to be used to ensure that the result of the estimate alignment was successful before a refined alignment is performed. If the estimate technique cannot derive the optimal alignment between the two point-clouds, this is where we will utilize human intervention to modify the candidate data-points to

Given a scan batch that contains several point-clouds, where each scan has pairwise overlap with at least one other scan, each pairwise alignment is performed consecutively with verification checkpoints. The compilation process consists of three main steps: data acquisition, pairwise alignment, and an application of the transformations provided by each pairwise alignment identified as the alignment compilation process.

1. Acquire Scan Batch (point-clouds with pairwise overlap)
2. For each scan pair with overlap:
 - a. Paint common features between scans in current pair
 - b. Estimate Alignment (ensure alignment, adjust painting)
 - c. Perform Refined Alignment
 - d. Checkpoint: Ensure alignment, continue to next pair
3. Compile Alignment

The structure of this approach is defined to reconcile the fragmentation between potentially noisy and error prone raw point-cloud data, estimate pairwise alignment and the use of refined alignment techniques to ensure that the virtual object construction process is robust. When a statistically-based algorithm fails to provide a valid result, the alternative is to manually alter the numerical parameters of the alignment algorithm. This is inherently a complex process and the modifications are not readily apparent or intuitive.

Here we allow for alignment verification and provide an intuitive painting-based selection method for manual intervention to ensure the alignment is successful based on a more accurate selection of corresponding features. The final process in the alignment compilation is defined by the collection and application of the transformations provided by each pairwise alignment. After the estimate and refined alignments have been applied to a scan pair, the resulting output will be a transformation matrix that represents the translation and rotation required to align the scan pair. Once this transformation matrix has been obtained for all pairs, they can be sequentially applied to construct a final model.

The critical feature imposed by this method is that if we verify the each pairwise alignment was successful (within the error metric of the refined alignment technique); we can ensure the validity of the resulting object constructed from the aligned scans after the alignment is compiled. While this approach can provide a viable result it is however sensitive to the accumulative error introduced by each refined pairwise alignment (due to the consecutive application of the refined alignment transformations). However with the flexibility of this compilation process, we can choose to employ a global error minimization alignment technique between all scans to further reduce this accumulative error after the batch has been compiled.

6 Evaluation and Discussion

The evaluation of this approach is defined by the demonstration of this methodology for the efficient construction of virtual object models using existing alignment techniques in instances where they would otherwise fail to converge. We illustrate how painting-based feature selection and alignment compilation can be utilized to ensure alignment convergence in the process of semi-automated object reconstruction. Specifically we look at instances where the employed estimate alignment technique (SAC-IA) will fail to provide an estimate alignment within the convergence threshold of the ICP algorithm, and how this can be addressed using our featuring painting technique.

One particularly challenging alignment is a pair of point-clouds that contain a high level of symmetry and non-corresponding features. This instance is not generally handled well by estimate alignment techniques due to the similarities between each side of the vehicle and features present in one scan but not in the other (non-correspondence). We initially consider all points within these scans to try to establish an estimate alignment using the SAC-IA approach. The result of this failed alignment is illustrated in Figure 6.



Fig 6. SAC-IA fails to provide an accurate alignment (right) of the two input point-clouds (left and center) due to the influence of the incorrect correspondence between the trunk and hood.

The painting technique we have introduced allows for a subset of features of each scan to be properly selected so that the resulting alignment is geometrically accurate and within the ICP convergence threshold. In Figure 7 we see an illustration of the scan-pair from Figure 6; however instead of sampling all points, we selectively paint the features common to both scans (in this instance, the side of vehicle). This selection refines the set of candidate points used during this estimate alignment. The improved alignment result is derived from the elimination of large non-corresponding features (such as the hood and trunk of the vehicle).

Table 1. SAC-IA (PCL Implementation) Estimate Alignment Parameters

Voxel Grid Size: 0.05	Maximum Iterations: 200
Max Correspondence Distance: 0.1	Normal k-Neighbors: 20
Minimum Sample Distance: 0.05	Feature k-Neighbors: 20

The elimination of these features has such a profound impact on the estimate alignment that the natural variance between individual painting patterns has a minimal negative effect. These alignments were performed using painted patterns illustrated in Figure 7 and the SAC-IA implementation provided in the Point Cloud Library (PCL) with the parameters identified in Table 1.

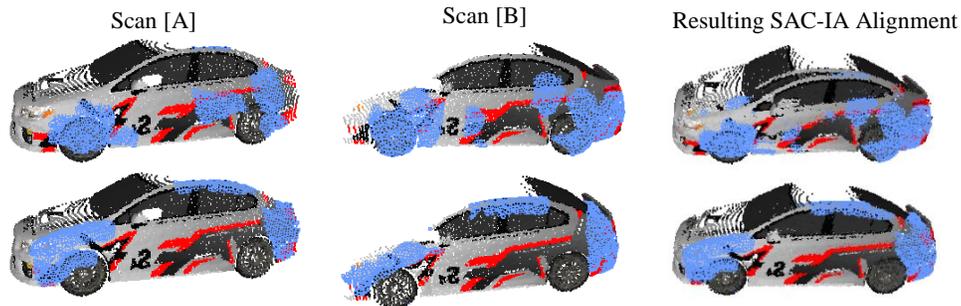


Fig 7. Estimate alignment convergence using SAC-IA algorithm for two distinct painting patterns. Top row: painted correspondence variant 1, bottom row: correspondence variant 2. The results of each SAC-IA estimate alignment between the scans in columns 1 and 2 are shown in the third column. In both instances, the resulting alignment is within the convergence threshold of ICP.

Based on the alignment provided by the painting assisted SAC-IA estimate alignment, we can now apply a refined alignment technique such as the ICP algorithm to finalize this pair-wise alignment. Using the same subsets of selected points or an expanded selection, we can ensure convergence of the ICP algorithm [12] with the following parameters:

Table 2. ICP (Scanalyze Implementation) Refined Alignment Parameters:

Iterations: 60	Culling %: 40
Sample Rate: 0.5	Method: Point-to-Plane
Norm-space Sampling: True	Threshold Type/Value: Relative/3

As introduced in section four, we can also utilize this painting technique to address instances where key-point selection techniques fail to be intuitive and easy to use. Specifically considering the alignment of point-clouds that contain uniform surfaces without unique features, the selection of individual points that will accurately provide a correct estimate alignment is challenging. Applying our painting technique solves this problem. We present a higher level objective to the user: identify general regions within the surface of the scanned object where the curvature appears similar. The result of painting a general region of these two dirt scans provides an accurate alignment utilizing the SAC-IA, ICP alignment sequence defined above. The painted region and result of this alignment sequence is illustrated in Figure 8.



Fig 8. Identifying corresponding regions is more intuitive for uniform surfaces that lack unique key-points. The successful alignment of these scans is presented using our painting technique with the SAC-IA estimate alignment algorithm.

Given that the user can interactively select general regions, correspondence between point-cloud features are often easier to identify than key-points. This is especially true for low resolution point-clouds that contain uniform surfaces with significant noise. However when identifying feature regions, we note that the incorrect selection of these corresponding features can significantly degrade the quality of the estimate alignment. In general the requirement is that similar features (based on geometric composition, color, etc.) can accurately be identified by the user. Including both selection and deselection paint-brushes provides the user with the ability to iteratively refine their selected features. If an estimate alignment is still unsuccessful with the current select point subset, this selection can easily be modified to more accurately define the identified feature. This provides the user with our higher level objective of region selection.

Using both the painting-assisted SACI-IA and ICP algorithms, the images in Figure 9 illustrate the application of our iterative alignment compilation technique to successfully align low resolution scan pairs with minimal overlap to construct full virtual models. Each model illustrated below is composed of only four scans with minimal overlap $\sim 30\%$.



Fig 9. Alignment results of three individual scan batches: Stanford bunny (left), WRX (center), and the Legacy dataset (right). Through the use of the painting-assisted alignment, highly symmetrical scans can be effectively aligned with minimal human intervention.

7 Conclusion

In this paper we presented a new and intuitive painting-based methodology for selecting corresponding features between scan pairs. Upon doing so we also greatly reduced the alignment error exhibited by estimate alignment techniques for scans with low overlap, symmetry, and uniform surfaces. Additionally, we introduced an alignment compilation process that ensures that through the process of applying subsequent pairwise alignments, the resulting transformations will be valid and produce an accurate representation of the scanned object. This work provides a step towards delivering a flexible and robust alignment methodology that can be utilized aid in the process of semi-automated object reconstruction when pair-wise alignment techniques fail to converge.

References

1. R. B. Rusu, N Blodow, M. Beetz. "Fast Point Feature Histograms (FPFH) for 3D registration", International Conference on Robotics and Automation - ICRA , pp. 3212-3217, 2009.
2. P. J. Besl and N. D. McKay, "A Method for Registration of 3-D Shapes," IEEE Trans. Pattern Anal. Mach. Intell., vol. 14, 1992.
3. ter Haar, F.B.; Veltkamp, R.C., "Automatic multiview quadruple alignment of unordered range scans" Shape Modeling and Applications, 2007. SMI '07. IEEE International Conference on, vol., no., pp.137,146, 13-15 June 2007.
4. Minsu Cho, Jungmin Lee, and Kyoung Mu Lee, "Feature Correspondence and deformable Object Matching via Agglomerative Correspondence Clustering", IEEE International Conference of Computer Vision (ICCV), 2009.
5. Shane Transue and Min-Hyung Choi, "Enhanced Pre-conditioning Algorithm for the Accurate Alignment of 3D Range Scans", International Conference on Image Processing, Computer vision, and Pattern Recognition. 2013.
6. R. B. Rusu, N. Blodow, Z. C. Marton, and M. Beetz, "Aligning Point Cloud Views using Persistent Feature Histograms," in Proceedings of the 21st IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Nice, France, September 22-26, 2008.
7. Szymon Rusinkiewicz and Marc Levoy. "Efficient Variants of the ICP Algorithm". Proceedings of the International Conference on 3-D Digital Imaging and Modeling, pp. 145–152, 2001.
8. Yang Chen, and Gerard Medioni. "Object Modeling by Registration of Multiple Range Images". International Journal of Image and Vision Computing, 10(3), pp. 145–155, 1992.
9. Torsello, A.; Rodola, E.; Albarelli, A., "Sampling Relevant Points for Surface Registration," 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2011 International Conference on , vol., no., pp.290,295, 16-19 May 2011.
10. Avishek Chatterjee, Suraj Jain, and Venu Madhav Govindu. "A pipeline for building 3D models using depth cameras". ICVGIP '12 Proceedings of the Eighth Indian Conference on Computer Vision, Graphics and Image Processing. Article No. 38, 2012.
11. David G. Lowe. "Distinctive Image Features from Scale-Invariant Keypoints". vol. 60 of Int. J. Comput. Vision, pp. 91-110. 2004.
12. Scanalyze: A system for aligning and merging range data. Stanford University, 1997-2004.