Dynamic Cage-Driven 3D Range-Scan Alignment

L. M. Seversky[1,2] and L. Yin[2]

[1] Air Force Research Laboratory, Rome, NY
[2] Department of Computer Science, State University of New York at Binghamton, Binghamton NY

Abstract

This paper presents a novel and automatic approach for aligning range-scan data of objects exhibiting non-rigid, articulated motion using a cage-driven reduced deformable model. Reduced deformable models have previously been used for non-rigid registration. However, these approaches usually assume a model apriori or determine one in step with the registration which adds complexity. We choose a cage-based space deformation mapping as the reduced deformable model and formulate the scan alignment problem as a space deformation problem. This cage-based deformation mapping provides a compact deformation model which is inherently geometric. We seek the deformation of a source cage (and embedded geometry) that results in the best alignment of the source and target scans. The main advantage of our approach is that the reduced deformable model is constructed automatically from the underlying object geometry and is independent of the alignment procedure as it does not require explicit partitioning of the object into parts or the establishment of joints. Our alignment algorithm is completely automatic and does not require initial correspondences between the surfaces to be aligned.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric Algorithms

1. Introduction

Acquired time-varying 3D geometry is becoming more and more prevalent as 3D acquisition and reconstruction technologies continue to advance. Solving the 3D registration problem is a fundamental first step towards the processing and higher level analysis of this emerging data type. This work addresses the registration problem for range-scan data of non-rigid articulated real world objects.

Non-rigid registration methods can be categorized by how the underlying object's motion is modeled. A number of approaches seek to model or track the motion of each point [ARV07]. Alternatively, global approaches solve for a registration that conforms to a single global deformation model [BM92, JV05]. The thin-plate spline is a popular global model which forces the deformation to be globally smooth. However, due to the global smoothness constraint it is unable to effectively capture large or piecewise rigid deformations which are common in many real world objects.

Recently, reduced deformable models (RDM) have been used in the context of non-rigid registration. RDM approaches model motion using a few deformation parameters and have been shown to be well suited for modeling deforming articulated shapes [CZ09]. However, automatically determining an appropriate RDM for registration is not a trivial task. Recent approaches either assume a deformable model

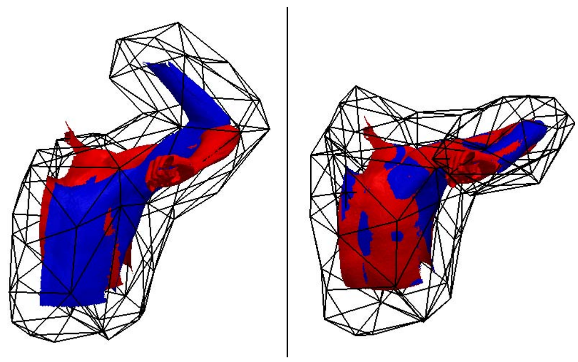

Figure 1: *Cage-Driven Alignment. Left: Source (blue), target (red) and initial cage (black) enclosing source. Right: Resulting alignment and corresponding deformed cage.*

apriori or find a model in step with the registration which adds additional complexity to the already difficult registration problem.

In this paper, we present a non-rigid registration approach using a *cage-based* space deformation mapping as the reduced deformable model, which has not been previously considered. We leverage recent space deformation techniques [JMD*07, BCWG09] and formulate the registration problem as a cage-based deformation problem in which the deformation parameters of the chosen RDM is a function of the geometry of a closed 3D mesh enclosing the object in space. Figure 1 illustrates an example cage and alignment using our framework. This geometric interpretation of a RDM allows us to automatically construct an *appropriate* model directly from the underlying object geometry, simplifying both model construction and the alignment problem formulation.

1.1. Related Work

Recently, reduced deformable models (RDM) have become popular for non-rigid registration of real world objects due to their ability to model object motion using a compact set of deformation parameters. In general, a RDM provides a low parameter model for modeling object motion by partitioning an object into parts where each part moves together as a unit and the set of parts compose an object's deformation. RDMs have been extensively used to model surface and free-form deformation [JT05, DSP06], which is not the same problem addressed in this paper, since in this work no correspondence or motion information is known.

Work by Huang et al. [HAWG08] obtains a deformation by clustering the motion of the surface into parts which can be described by a single rigid transformation. In Li et al. [LSP08] a rigid transformation is assigned to each node in a deformation graph serving as the RDM. Most recently Chang et al. [CZ09] use a linear blend skinning model and explicitly solve for skinning weights to determine which points move together. A grid based approach allows the weights to be dynamically determined as part of the optimization formulation resulting in more realistic deformations.

Different from previous approaches, we consider a cage-based space deformation mapping as our RDM which enables us to automatically construct the model from the underlying object geometry independent of registration. The resulting model takes on a simple geometric form which is able to express natural as-rigid-as-possible shape deformations. As shown in the next section, the cage-based RDM can be easily incorporated into a non-rigid registration framework and can also be used to provide an initial alignment in place of other methods which require a sparse initial set of correspondences.

2. Cage-Driven Alignment

Before we present our algorithm for cage-driven registration, a brief overview of space deformation techniques is provided. Space deformation methods deform the ambient space in which an object is embedded. Specifically, we use a cage-based approach which deforms an object by positioning a coarse closed triangular mesh (i.e. *a cage*) around the object. The object is then represented in terms of the cage vertices and face normals by computing a weight at each cage vertex and face at the position of every object point, thus forming an embedding of the object with respect to the cage. As the cage vertices are moved to new locations, new object point positions can also be determined. These embedding and deformation properties can be expressed as the linear combination

$$p = F(p;V;N) = \sum_{i \in I_V} \phi_i(p) v_i + \sum_{j \in I_T} \psi_j(p) n(t_j) \qquad (1)$$

where p is a point on the object and $\phi(p)$ and $\psi(p)$ are the coordinates for p which are a function of the point location and cage vertices and normals, respectively. The cage mesh vertices and triangles are denoted by $V = \{v_i\}_{i \in I_v}$ and $T = \{t_j\}_{j \in I_T}$, and N is the set of of triangle normals where $n(t_j)$ denotes the outward normal of the triangle t_j. Similarly, the deformation as a result of a deformed cage C' is expressed as

$$p' = F(p;V';N') = \sum_{i \in I_V} \phi_i(p) v'_i + \sum_{j \in I_T} \psi_j(p) n(t'_j) \qquad (2)$$

where p' is the new point location. As shown in Ben-Chen et al. [BCWG09], we use Variational Harmonic Maps for coordinate construction over other methods [JMD*07, LLCO08], due to their ability to produce well behaved as-rigid-as-possible deformations for complex shapes as well as the existence of closed form expressions.

Using a cage-based space deformation mapping as a RDM for non-rigid registration has many benefits due to the properties of space deformation models. First, they are general as the deformation model is independent of the underlying object representation. Second, the computational complexity of deforming the cage is independent of the complexity of the underlying object geometry. Third, the expressiveness of the deformation model can be controlled by changing the detail and resolution of the cage. As a result, relatively simple cages can be used to model the deformation of complex shapes regardless of the shapes genus or surface geometry representation. Given this background on space deformation methods we present our cage-driven alignment algorithm in detail.

2.1. Alignment Algorithm

We propose to align 3D range scans of a moving object by modeling the motion of the object using a reduced deformable model (RDM). Due to the many attractive properties of space deformation techniques we choose a cage-based space deformation mapping for the RDM where the cage describes the deformation implicitly without the need for partitioning, clustering, or transformation and joint assignment.

Let X and Y denote the source and target 3D point set data and let C denote a source closed polygonal cage mesh positioned around point set X. The deformation of the source cage, C' is sought which aligns the source point set X with the target point set Y. Unlike the user-driven cage deformation problem formulation in [BCWG09], we do not have any source and target point position correspondences and instead must solve for these as part of our optimization problem. Therefore, two terms are incorporated into the objective function which measure the local and global accuracy of the alignment. Thus, the goal is to find the deformed cage C' defined by vertices V and normals N, which are denoted for convenience as the stacked matrix Z, which minimizes the cost function given by

$$E(Z) = \alpha||F(X,Z) - Y||^2 + \beta||f_z - g||^2 + \lambda||\hat{H}Z||_F^2 \quad (3)$$

where the first term $||F(X,Z) - Y||^2$ measures the accuracy of the point set alignment at a fine scale. This term evaluates the closest point distance between the deformed point set $F(X,Z)$ and the target set Y.

The second term measures the accuracy of the registration at a more global level by considering the registration as the alignment between two Gaussian Mixture Models (GMM), where each of the point sets are the GMM centroids given by $f(x) = \sum_{i=1}^{m} \alpha_i \phi(x|u_i, \Sigma_i)$ and $g(x) = \sum_{j=1}^{n} \beta_j \phi(x|v_j, \Gamma_j)$, respectively [JV05]. The l_2 distance between the transformed source distribution and target distribution is given by $f_Z(x) = \sum_{i=1}^{m} \alpha_i \phi(x|F(u_i, Z))$, where $F(u_i, Z)$ deforms the object point p_i associated with centroid u_i with respect to the cage with vertices and normals equal to Z.

The last term, $||\hat{H}Z||_F^2$ ensures that the cage deforms in an as-rigid-as-possible fashion by enforcing that nearby points sampled on the cage boundary undergo similar transformations [BCWG09]. Finally, the scalars α, β, λ are weight coefficients for each term respectively.

Implementation Details: In order to minimize the cost function in (3), an unconstrained optimization problem is solved. An initial deformation is obtained by matching source and target cages using the robust non-rigid point registration method described in [CR03]. We compute the coordinates (ϕ, ψ) for the points belonging to X and Y and the Hessians for the boundary sample points once. Additionally, not all the points need to be considered if only a sparse correspondence is sought. The distributions f and g are constructed from points uniformly sampled from X and Y, respectively and are assumed to be spherical GMMs with uniform scale.

Cage Construction: Cages for interactive deformation applications have traditionally been constructed manually. As a result, we developed a simple automatic cage construction method for use in our alignment procedure. We construct a watertight cage mesh by sampling a set of balls uniformly from the range-scan data and using the union-of-balls reconstruction method of [ACK01]. We manually choose a fixed radius for all balls which determines the tightness of the reconstructed mesh. Simplification is then performed on the reconstructed mesh to obtain the cage at the desired resolution.

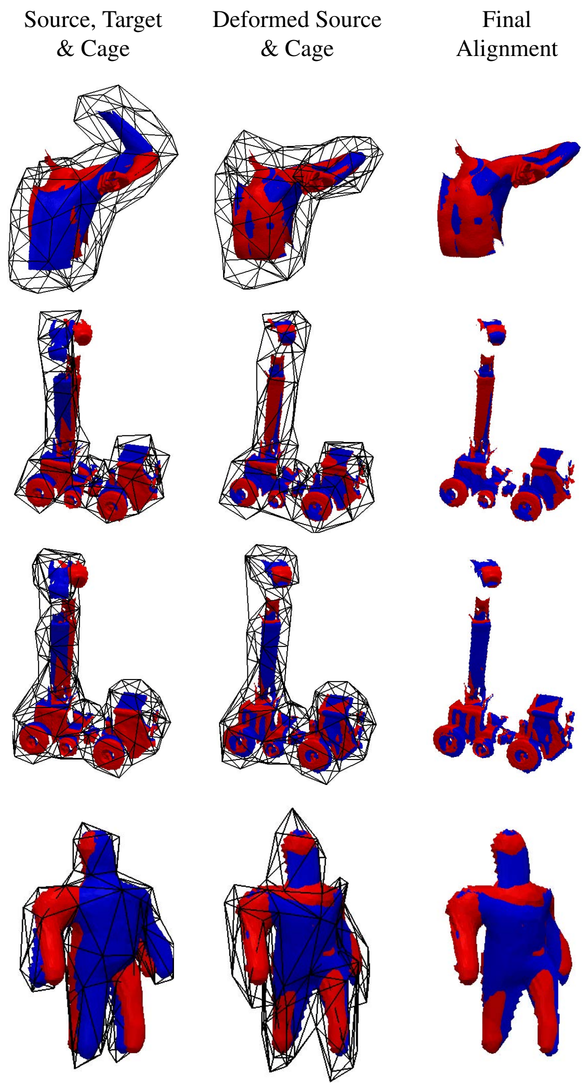

Figure 2: *Alignment of frames from the shoulder, car, and walk datasets. The source object is denoted in blue, the target in red, and cage denoted by the black wireframe.*

3. Results & Discussion

Examples of aligned scans from the shoulder, car, and synthetic walk sequences are shown in Figure 2. Alignment of both adjacent and non-adjacent frames were evaluated. In these examples, each frame averaged 80K, 5.7K, and 5.4K points for the shoulder, car, and man sequences. Similarly, the source cage meshes averaged 150, 200, and 150 vertices, respectively. All cages were constructed automatically from the range-scan data. However, for some frames in the car sequence a single unified cage was not generated due to the large amounts of missing data. We manually joined the separate pieces into a single cage in these instances as shown in Figure 2. Other less severe instances of missing data and occlusions such as in the torso and hand of the shoulder scans did not present a problem for our cage construction and alignment algorithms. However, the synthetic walking sequence proved to be more challenging as shown in Figure 3, due to the stretching of the extremities. In this case, cages which more closely mirrors the articulated joints of the underlying object should produce better alignments.

The tests were performed on an Intel Core Duo 2.4GHz laptop with 4GB of RAM. The alignment time is dominated by the time for coordinate construction which is a function of the cage resolution and number of points in the range-scan. However, if a sparse correspondence is sought, the total time can be reduced by constructing coordinates for a smaller subset of the points in the source scan. Further, when aligning a sequence of shapes it may be possible to reuse coordinates over subsequent frames. Finally, our approach assumes that the source and target cages enclose similar shapes. As a result, partial matching of objects due to missing parts or large occlusions is a challenge for our approach.

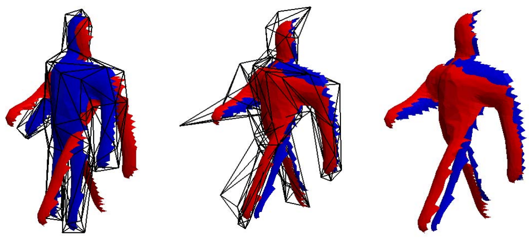

Figure 3: *An example of a misalignment from the synthetic walk dataset due to the stretching in the extremities. Left: Source, Target & Cage. Middle: Deformed Source & Cage. Right: Final Alignment.*

4. Conclusion

This work introduces a method for automatically aligning range scans of deforming objects by using a space deformation mapping as the reduced deformable model. We show that a cage-based deformation mapping provides a compact deformation model that is inherently geometric and as a result can be easily manipulated in terms of its geometry. We demonstrate that our alignment framework can match articulated shapes with significant occlusions and missing data automatically without requiring an initial correspondence. Finally, we present a simple approach for constructing cage-based RDMs automatically from the underlying object geometry which does not depend on explicit object partitioning or joint modeling. Investigating automatic cage building techniques which provide consistent multi-resolution cages and support for partial registration will be considered in future work. Further, we believe our cage-driven approach holds large potential for use in existing robust registration methods.

Acknowledgments

We wish to thank Y. Pekelny and C. Gotsman for creating the car dataset, I. Baran for the walking dataset, and B. Allen for the shoulder dataset and also W. Chang for providing us this data.

References

[ACK01] Amenta N., Choi S., Kolluri R.: The power crust, unions of balls, and the medial axis transform. *Computational Geom.: Theory and Applications 19*, 2-3 (2001).

[ARV07] Amberg B., Romdhani S., Vetter T.: Optimal step nonrigid icp algorithms for surface registration. *IEEE Conference on Computer Vision and Pattern Recognition* (2007), 1–8.

[BCWG09] Ben-Chen M., Weber O., Gotsman C.: Variational harmonic maps for space deformation. *ACM Trans. Graphics 28*, 3 (2009).

[BM92] Besl P., McKay N.: A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence 14*, 2 (1992), 239–256.

[CR03] Chui H., Rangarajan A.: A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding 89*, 2-3 (2003), 114–141.

[CZ09] Chang W., Zwicker M.: Range scan registration using reduced deformable models. *Computer Graphics Forum 28*, 2 (2009), 447–456.

[DSP06] Der K., Sumner R., Popović J.: Inverse kinematics for reduced deformable models. *ACM Trans. Graphics 25*, 3 (2006), 1179.

[HAWG08] Huang Q., Adams B., Wicke M., Guibas L.: Non-rigid registration under isometric deformations. *Computer Graphics Forum 27*, 5 (2008), 1449–1457.

[JMD*07] Joshi P., Meyer M., DeRose T., Green B., Sanocki T.: Harmonic coordinates for character articulation. In *SIGGRAPH* (2007), p. 71.

[JT05] James D., Twigg C.: Skinning mesh animations. In *SIGGRAPH* (2005), p. 407.

[JV05] Jian B., Vemuri B.: A robust algorithm for point set registration using mixture of gaussians. *ICCV* (2005), 1246–1251.

[LLCO08] Lipman Y., Levin D., Cohen-Or D.: Green coordinates. In *SIGGRAPH* (2008), p. 78.

[LSP08] Li H., Sumner R., Pauly M.: Global correspondence optimization for non-rigid registration of depth scans. *Computer Graphics Forum 27*, 5 (2008), 1421–1430.