

An Automatic Method For Determining Quaternion Field Boundaries for Ball-and-Socket Joint Limits

Lorna HERDA, Raquel URTASUN, Pascal FUA*
Swiss Federal Institute of Technology
CH-1015 Lausanne, Switzerland
{Lorna.Herda, Raquel.Urtasun, Pascal.Fua}@epfl.ch

Andrew HANSON
Indiana University
IN 47405 Bloomington, USA
hanson@cs.indiana.edu

Abstract

To improve the robustness of human motion synthesis and capture algorithms, our goal is to provide an effective framework for imposing joint limits and reducing ambiguities. To this end, we determine these joint limits from measures performed on human subjects and then transpose them to an intuitive joint limits formalism. Here we focus on the case of ball-and-socket joint limits, applied to the specific case of the shoulder complex. The first step is to measure the joint motion range using optical motion capture. We then convert the recorded values to joint poses using a coherent quaternion field representation of the joint orientation space. Finally, we obtain a closed, continuous implicit surface approximation for the quaternion orientation-space boundary delimiting the complete space of valid shoulder orientations, and allowing us to determine the closest valid posture for any given invalid one.

1 Introduction

In the area of Computer Vision, various approaches to tracking and modeling people from video sequences have been proposed [6]. They are not always robust, as image data is inherently noisy and ambiguous. Constraints on the possible poses greatly

*This work is supported in part by OFES. The authors would like to thank Paolo Baerlocher, Ralf Plänkers, Richard Lengagne, Eric Bittar, and especially Amaury Aubel for their valuable contribution to this work.

reduces the ambiguity, restricting the search space of the algorithm to humanly possible configurations. Imposing valid joint limits is a major means of improving performance, and this both for animation and tracking purposes.

Current tools for delimiting valid postures often employ techniques that do not represent joint limits in an intuitively satisfying manner. They tend to be expressed in terms of hard limits on individual rotation angles that do not account for dependencies between those angles [10][12].

By contrast, in this work, we propose a quaternion-based model that allows us to explicitly represent and measure those dependencies: we first measure the joint motion range using optical motion capture. We next convert the recorded values to joint poses encoded by a coherent quaternion field representation of the joint orientation space, a quaternion being the only rotation representation that provides a uniform distance metric allowing the comparison of orientations and thus, meaningful measures of similarity. Finally, we derive a closed, continuous implicit surface approximation for the quaternion orientation-space boundary whose interior represents the complete space of valid orientations, ultimately enabling us to clamp any invalid posture to the closest valid one.

We use the shoulder joint as an example to demonstrate the validity of our approach as this joint, along with the hip, falls under the category of ball-and-socket joints allowing movement in three planes.

In the remainder of the paper, we first briefly review the state of the art. We then introduce our ap-

proach to experimentally sampling the space of valid postures that the joint allows and to representing this space in terms of an implicit surface in quaternion space. Finally, we demonstrate our method’s effectiveness for animation purposes and discuss its applicability in a tracking context.

2 Related approaches

In the field of biomechanics, joint limits have often been measured, using various mechanical and electromagnetic devices, and the resulting statistical data is widely published.

Many of these empirical results are subsequently re-used for modeling human skeletons, when the conversion from the measured data to the representation is readily feasible. The simplest approach is to introduce joint hierarchies formed by independent 1-Degree-Of-Freedom (DOF), described in terms of Euler angles. This formalism, however, does not account for the coupling of the different limits and suffers from an additional weakness known as “Gimbal lock” [13]. The sinus cones, introduced by A.G. Engin in [1] and the spherical polygon, proposed by Korein in [5], improve on this situation. Recently, a new model was suggested by Tolani, Badler and Gallier in [10], where the workspace of angular motion is modeled as a triangular Bézier spline surface.

The 3-DOF motion of a ball-and-socket joint can be decomposed into an angular and an axial component. The angular component, or swing, corresponds to the flexion and abduction DOFs, whereas the axial component, or twist, is the rotation around the axis of the limb itself. In all of these representations, joint limits apply independently to swing and to twist, even though in reality, they are correlated. Wang et al. recorded arm motions and expressed the amount of twist as a function of swing [12]. However, they still do not propose a model that simultaneously provides joint limits for all the DOFs.

In short, the method we propose advances the state-of-the-art because it provides a way to enforce limits on swing and twist while at the same time accounting for their dependencies. Furthermore, the quaternion representation we use is not subject to singularities such as the “Gimbal lock” of Euler angles or mapping rotations of 2π to zero rotation.

3 Measuring and representing shoulder motion

3.1 Approach

The first step towards our objective is to use the Vicon8 motion capture system for collecting 3-D marker data corresponding to extensive shoulder movements along the boundaries. The set of possible joint orientations and positions in space can be considered as a path of referential frames in 3-D space that can be mapped to a sub-space of the space of possible quaternion referential frames of the object, called the Quaternion Gauss map [3]. In the case of the data collected from optical markers, we obtain a cloud of 3-D points, around which we can then wrap a surface, this yielding a volume. In order to determine whether we have exceeded joint boundaries, we need then only perform an inside/outside test on this volume.

We chose quaternions over other formalisms for our work, because only they are free of singularities [9], even if this comes at the cost of an additional parameter and constraint [2]. In practice, rotations are represented by the sub-space of unit quaternions S^3 forming a unit sphere in 4-dimensional space. Any rotation can be obtained by a unit quaternion, so we need to keep in mind that the unitary condition needs to be ensured at all times. Quaternions have a large range of advantages, with respect to rotation representation, interpolation and optimization, that can also be read up in [3]. As we simultaneously measure swing and twist components, and as the quaternion formalism lets us express both within one rotation, we will be able to capture dependencies between swing and twist, one of the major goals of this approach. In this manner, we will have generated joint limits in a semi-automatic way, on the basis of motion capture. We will then be able to make use of these joint limits as constraints for tracking and pose estimation, by eliminating all invalid configurations.

3.2 Motion measurement

In order to capture the entire range of upper arm motion, we used the Vicon Motion Capture System to track four markers placed on the arm of the actor (see Fig. 1a). There is an additional marker at neck-level, serving as fixed reference, given that the actor moves in 3-D space. The motions performed were to cover the entire range of shoulder motions. It has been observed that the inter-subject variance is extremely small [12], therefore measures were performed on

one individual only. The Vicon motion capture system ultimately provides us with reconstructed and labeled markers for every image frame of the sequence.

3.3 Motion representation

Quaternion field technology

Our goal is to create an analyzable topographic space of joint orientations using the experimental 3-D marker coordinates as the basis. To accomplish this, we adopt the method of quaternion fields described in [4]. In the quaternion field method, each joint orientation is first converted to a 3×3 matrix M , where, using Euler’s theorem (not the Euler angles), M may be expressed in terms of its lone real eigenvector \hat{n} and the angle of rotation θ about that axis. This in turn may be expressed as a point in quaternion space, or, equivalently, a point on a three-sphere S^3 embedded in a Euclidean 4D space. To find θ and axis \hat{n} , given *any* rotation matrix or frame M , we need two steps: (1) solve for θ using

$$\text{Tr}M = 1 + 2 \cos \theta \quad (1)$$

and (2) find \hat{n} from the expression of $M - M^t$ that is defined as long as $\theta \neq 0$. In the degenerate case, the matrix and the quaternion are taken to be the identity transformation.

The identification of the corresponding quaternion follows immediately from

$$q(\theta, \hat{n}) = \left(\cos \frac{\theta}{2}, \hat{n} \sin \frac{\theta}{2} \right) \quad (2)$$

up to the sign ambiguity between the two equivalent quaternions q or $-q$, which correspond to the exact same rotation matrix M . A single compact volume containing a dense set of quaternion frames will have two distinct but equivalent realizations due to the $q \Leftrightarrow -q$ ambiguity. We have implemented various procedures to cast out duplicates, or, equivalently, to iteratively check the 4-D dot-product between neighboring quaternions and force all neighbors to have signs that are consistent with positive dot products.

Construction of shoulder coordinate frames

To build our frames from the marker coordinates, we proceed as follows: the first axis of the frame corresponds to the upper arm segment. The second axis is the normal to the plane passing through the shoulder, elbow and wrist markers (this plane representing the axial rotation - or “twist” motion - of the shoulder joint). As for the third axis, it is simply the cross-product of the other axes (see Fig.1b).




Figure 1. (a) Marker positioning for motion capture. (b) Deriving a referential from the markers.

An orthonormal coordinate frame is created in this manner for each image frame of the sequence, along with the rotation relating one image frame to the next. The calculated 3×3 rotation matrices $M(t)$ for each data point are then converted into unit quaternions as described above, producing a dense cloud of points occupying a volumetric portion of the quaternion 3-manifold (the three-sphere).

When transformed to quaternion form, each successive frame from the moving sequence becomes a point on a quaternion curve lying within the three-sphere S^3 embedded in R^4 . Families of frames on a surface become surfaces in quaternion space, and collections or disjoint sets of orientations such as those collected here for the shoulder joint become a cloud of points occupying a volumetric quaternion space with a distinct closed boundary surface.

To visualize the resulting data points, we display them in the 3-D space defined by the spatial coordinates of the quaternion; in other words, these three numbers serve as the coordinates of a quaternion point expressed as projections on three conventional Cartesian axes. Since we are dealing with unit quaternions, the fourth quaternion component q_0 is a dependent variable (up to a sign) and can be deduced from the first three.

3.4 Data analysis

The resulting collections of data points describing the experimentally determined space of possible shoulder orientations show the power of the quaternion field method. Since each possible orientation is a point in a cloud with a clear and unambiguous metric (inherited from the three-sphere), the mutual distances of individual orientations are definable in a mathematically rigorous way, as is the volume giving a measure of the total freedom of motion available to any particular human subject or collection of human subjects. Furthermore, the minimal “shrink-

wrapped" hull surface in quaternion space bounding the volume of experimental points can be computed directly and used as a consistency check on the acceptability of a proposed or measured joint orientation. The quaternion field method thus provides a new and unique method for establishing requirements and constraints for automated or semi-automated joint animations that can be used to reject or readjust proposed orientations to lie within a particular, experimentally verified, joint parameter space.

4 3-D data approximation

In order to make effective use of the collected data, we now need to derive a working representation for it.

Many methods have been studied to solve the problem of shape reconstruction, but we chose the implicit surface representation because it yields a smooth representation of the volume of valid quaternion rotations and allows us to deal with relatively few parameters. At the same time, it provides a convenient metric for estimating proximity, due to the fact that the surface is represented by an equation and not modeled explicitly. Furthermore, when dealing with a value representing a non-valid rotation, the use of implicit surfaces gives us the possibility of clamping this value to the next-best value.

Muraki [7] used implicit surfaces for the purpose of approximating scattered data, introducing the notion of implicit iso-surfaces generated by skeletons. These can be either a point (the center of a sphere or ellipsoid), or a rigid skeleton (a line or a polygon). The principle of this algorithm is to position a surface skeleton at the center of mass of the points and to subdivide it until the process reaches a correct approximation level. This, however, presents several drawbacks, such as the lack of a local criterion and a high computational cost. Furthermore, the algorithm does not perform well on data with "holes", and requires knowledge of surface normals.

Bittar, Tsingos and Gascuel [11] proposed a new method based on Muraki's work in which they outlined a more efficient way of splitting the implicit surfaces without requiring normals, and expressed them in terms of locally defined field functions. This allows the use of an iterative method based on minimizing the distance between the real points and the generated surface. This method forms the basis of our approach, but has been substantially modified to accommodate the fact that (a) our motion capture data forms a volume in quaternion space, so the surface data first needs to be extracted, (b) our data shows variations in data density.

4.1 Implicit surface parameterization

In our implementation, an implicit surface is an iso-surface (or contour surface) of a field defined by one or more spherical primitives. The function f represents our sphere and a point $P(x, y, z)$ belongs to the surface if

$$f(x, y, z) - iso = 0.$$

The field function is defined as follows:

$$f_i(P) = \begin{cases} -k_i r + k_i e_i + 1 & \text{if } r \in [0, e_i] \\ \frac{1}{4} [k_i (r - e_i) - 2]^2 & \text{if } r \in [e_i, R_i] \\ 0 & \text{elsewhere} \end{cases}$$

where $r = d(P, S_i)$ and $R_i = e_i + \frac{2}{k_i}$ [11].

The field functions f_i of the implicit surfaces are parameterized by:

- The radius e_i of the sphere S_i created by a single surface skeleton, so that $f_i(e_i) = iso$.
- The stickiness k_i in e_i that defines the blending properties of the surface.

To speed-up computation and to prevent the approximation from diverging, we fix the stickiness k_i , which means we only need to optimize the remaining four parameters e_i, x_i, y_i, z_i . The chosen value is 5, the iso-value being set to 1. This offers sufficient precision, without letting the primitives become too small, in which case we would need a large number of them to obtain the final fitting solution.

We place the first spherical primitive at the center of mass (x_c, y_c, z_c) , and then iteratively divide it until it approximates as well as possible the cloud of data, by minimizing the energy E :

$$E = \frac{1}{M} \left(\sum_{j=1}^M (f(P_j) - iso)^2 \right) + \alpha + \beta \quad (3)$$

The first term expresses the fact that the surface has to approximate the points. The other two terms α and β respectively prevent the surface primitives from wildly increasing and from moving too far away from the center of mass. Without these constraints, we would obtain large tangent surfaces that tend towards a solution within a small number of iterations.

4.2 Surface primitives subdivision

We have implemented a subdivision-into-6 algorithm based on the local properties of the field function, the child surface primitives being positioned on one axis in a symmetrical manner, so as to yield a




Figure 2. (a) Original quaternion data in 3-D. (b) Surface points. (c) Fitted implicit surface before cleaning. (d) Fitted implicit surface, after a final least-squares fit of the cleaned primitives.

shape as close as possible to the original. The splitting criterion was defined so that the sum is calculated only on the basis of the m_i points that are inside the sphere of influence S_i . The surface primitive with the largest first term in equation 3 is the first candidate for splitting, i.e. the one whose area of influence corresponds to the zone of the surface where reconstruction is worst.

4.3 Results

We have applied all the above algorithms to the collected data for the shoulder joint. Our cloud of data initially being clustered in a volume, we first needed to extract the surface points. To this end, we apply a method based on dividing the 3-D space into overlapping stereoradians placed at the center of mass of the points. The extracted surface points and the resulting fitted implicit surface are shown in Fig.3, where the best solution found was made of 7 surface skeletons. Due to the variation in density of the data and the possible surface gaps this implies, we clean the final fitting result; in other words, we remove all spherical primitives where (a) the radius is larger than 1.0 (as this would obviously include the entire cloud of data of unit quaternions), or (b) the centre of the primitive is further away than 1.0 from the center of mass of the cloud, or (c) where the number of points




Figure 3. Implicit surface joint limits used for animation: an example of a non-valid orientation, and the subsequent corrected orientation.

included in the primitive is small with respect to total density. The initial fit in Fig.3c clearly shows the superfluous overgrown primitives. Fig.3d shows the final result, where after cleaning, we performed a last set of least-squares fitting iterations (without division of the primitives) so as to optimize the fit of the remaining blobs to the surface points.

5 From implicit surfaces to joint limits

In this section, we briefly show how our results can be used for animation purposes and discuss their applicability in a tracking context.

5.1 Implicit surface representation

We are now in a position to directly use the implicit surface representation in order to constrain postures. The animator chooses his keyframe, and due to joint limits enforcement, the position will either be accepted or need to be corrected. Each upper limb orientation is converted to its quaternion equivalent and tested with respect to the implicit surface. Clamping to the closest valid orientation can be readily enforced using the gradient of the total field function. In the example of Fig. 3, we can see a back view of a non-valid orientation, followed by the clamped valid one.

5.2 Video-based tracking

The technique described in Section 5.1 for animation purposes is, potentially, equally applicable to video-based tracking. In earlier work, we have developed an approach to tracking human bodies using




Figure 4. The role of joint limits in a tracking application. In (b) and (c) the left arm is inverted because of missing angular constraint whereas in (d) the model has the correct pose.

video from two or more cameras as input [8]. We extract stereo-data from the images and, at each time step, adjust the position of the body model to fit the data as closely as possible.

Because the model we use has many degrees of freedom, it can in some cases match the data accurately while assuming a position that no real human could. However, by imposing the joint limit constraints, the model switches to a much more plausible pose. An example is shown in Figure 4, where tracking of the arm without constraint enforcing results in a twist at the shoulder level. In this case, the constraints were imposed after the fact but the implicit surface formalism will make it natural to incorporate them directly into the tracking process.

6 Conclusion

We have developed a promising technique for automatically determining ball-and-socket joint limits on the basis of collected data, thus greatly improving our ability to correctly track human upper limbs. We have directly used the implicit surface as a joint limits representation, thus obtaining a tool that provides joint limits in an intuitive manner, without needing to decompose a joint rotation into several components. Furthermore, the method is easily generalized to any joint in the body, and this combination should present advantages in the area of vision as well as animation.

In the future, improvements will be considered with respect to the implicit surface fitting procedure, so as to increase the precision of the method with re-

spect to gaps in the data. These problems arise typically for the areas of non-reachable positions and for under-sampled regions, where it results in holes within the surface points extracted from the data.

References

- [1] A.E. Engin and S.T. Tumer. Three-dimensional kinematic modeling of the human shoulder complex. *Journal of Biomechanical Engineering*, 111:113–121, 1989.
- [2] F.S. Grassia. Practical parameterization of rotations using the exponential map. *Journal of Graphics Tools*, 3(3):29–48, 1998.
- [3] A.J. Hanson. Quaternion gauss maps and optimal framings of curves and surfaces. Indiana University Computer Science Department Technical Report 518 (October, 1998).
- [4] A.J. Hanson. Constrained optimal framings of curves and surfaces using quaternion gauss maps. In *Proceedings of Visualization '98*, pages 375–382. IEEE Computer Society Press, 1998.
- [5] J.U. Korein. *A Geometric Investigation Of Reach*. MIT Press, Cambridge, 1985.
- [6] T.B. Moeslund and E. Granum. A Survey of Computer Vision-Based Human Motion Capture. *Computer Vision and Image Understanding*, 81, March 2001.
- [7] S. Muraki. Volumetric shape description of range data using “Blobby Model”. *Computer Graphics*, 25(4):227–235, 1991.
- [8] R. Plänkers and P. Fua. Articulated soft objects for video-based body modeling. In *ICCV*, Vancouver, Canada, July 2001.
- [9] K. Shoemake. Animating rotation with quaternion curves. In *Computer Graphics*, volume 19, pages 245–254, 1985. Proceedings of SIGGRAPH 1985.
- [10] D. Tolani, N. Badler, and J. Tallier. A kinematic model of the human arm using triangular bézier spline surfaces.
- [11] N. Tsingos, E. Bittar, and M.P. Gascuel. Implicit surfaces for semi-automatic medical organs reconstruction. In *Computer Graphics International '95*, pages 3–15, Leeds, UK, 1995.
- [12] X. Wang, M. Maurin, F. Mazet, N. De Castro Maia, K. Voinot, J.P. Verriest, and M. Fayet. Three-dimensional modelling of the motion range of axial rotation of the upper arm. *Journal of Biomechanics*, 31(10):899–908, 1998.
- [13] A. Watt and M. Watt. *Advanced animation and rendering techniques*, 1992.