Articulated Rigid Registration for Serial Lower-Limb Mouse Imaging

Xenophon Papademetris[1,2], Donald P. Dione[3], Lawrence W. Dobrucki[3], Lawrence H. Staib[1,2], and Albert J. Sinusas[2,3]

[1] Departments of Biomedical Engineering
[2] Diag. Radiology
[3] Medicine, Yale University New Haven, CT 06520-8042
xenophon.papademetris@yale.edu

Abstract. This paper describes a new piecewise rotational transformation model for capturing the articulation of joints such as the hip and the knee. While a simple piecewise rigid model can be applied, such models suffer from discontinuities at the motion boundary leading to both folding and stretching. Our model avoids both of these problems by constructing a provably continuous transformation along the motion interface. We embed this transformation model within the robust point matching framework and demonstrate its successful application to both synthetic data, and to serial x-ray CT mouse images. In the later case, our model captures the articulation of six joints, namely the left/right hip, the left/right knee and the left/right ankle. In the future such a model could be used to initialize non-rigid registrations of images from different subjects, as well as, be embedded in intensity-based and integrated registration algorithms. It could also be applied to human data in cases where articulated motion is an issue (e.g. image guided prostate radiotherapy, lower extremity CT angiography).

1 Introduction

While non rigid image registration has been extensively applied to brain image analysis [7,2,4,5] (e.g. for comparing shape and function between individuals or groups, developing probabilistic models and atlases, and measuring change within an individual) it has not been extensively applied to other parts of the body to date. Unlike the brain which is a single organ enclosed in the skull with no articulated joints, the abdominal/pelvic cavities and especially regions close to limb joints contain many organs/glands whose relative position/orientation vary substantially from subject to subject. This is of particular importance for non-rigid registration, as the process typically relies on a good initialization often performed by estimating the global linear transformations between the two images. Given the relatively high degrees of freedom (DOF) available in most non-rigid registration methods, the final estimate of the nonlinear transformation is critically dependent on this early step to bring the nonlinear optimization process to a position to enable it to converge to the appropriate local minimum. While the estimation of the initial linear transformation is relatively straightforward in the case of brain images (using both intensity and/or feature methods),

J. Duncan and G. Gerig (Eds.): MICCAI 2005, LNCS 3750, pp. 919–926, 2005.

such simple transformations are inadequate (even for the purpose of initialization) in regions where there are articulated joints. Here the relative orientation of, for example, the proximal and the distal leg is highly dependent on the state of the knee joint and can vary substantially between subjects even when extra care is taken to implement a standardized imaging protocol. This is particularly true in our application of serial hybrid 3-dimensional imaging for the purpose of quantifying the remodeling of existing collateral arteries (arteriogenesis) and increased microvascular density (angiogenesis) associated with peripheral arterial disease of the lower extremities. In this application, x-ray computed tomographic (CT) angiography is used to evaluate lower extremity arteriogenesis, while registered single photon emission computed tomographic (SPECT) images provide a quantitative index of changes in either tissue perfusion or critical radio-labeled molecular markers of the underlying biological process.

While the problem of modeling articulated joints has received extensive interest in computer graphics, it has received little attention to our knowledge, in the medical image analysis literature, the recent early work of Martin-Fernandez et al. [6] being one exception. In their work, the problem of estimating 2D hand motion tracking is modeled using a weighted sum of multiple rigid transformation. While this model is adequate for simple motions, it does not explicitly address the key problem of folding (i.e. ensuring that the resulting transformation remains invertible). A more interesting set of work is the polyrigid/polyaffine model proposed by Arsigny et al. [1], where an elegant method for constructing an invertible transformation as a weighted sum of piecewise rigid/affine transformations is described. The weights are obtained by the solution of an ordinary differential equation. Arsigny et al. present 2D results in the matching of histochemical slices to anatomical data, and they note that this model would also be applicable to articulated registration issues such as the one of concern in this paper. A weakness of their method is the fact that there is no closed form solution to the resulting transformation, rather a numerical solution of the differential equation is needed to generate the final transformation, which makes the model highly computationally intensive.

In this work, we present both a theoretical model with a closed form solution for modeling the piecewise rotational motion of both the hip, the knee, and the ankle joints in mice (and by extension man) in a manner that ensures that the overall map is smooth and invertible and apply this model to the problem of registration of serial CT images acquired from mouse models of angiogenesis. The transformation model is used within the robust point matching registration framework to estimate the piecewise rigid transformation, which can then be used as the initialization to a full nonrigid registration.

2 Methods

2.1 A Model for Non-folding Piecewise Rotations

Our model for computing a non-folding mapping appropriately blends piecewise rotations making the following assumptions: (i) at each joint the motion is a

rotation about a rotation axis passing through the joint origin (e.g. the knee). (ii) The two "limbs" linked at the joint can be described without loss of generality as a stationary limb and a moving limb. (iii) A surface can be found which separates the two limbs in the original position (i.e. in the reference image); we label this the joint or motion interface, and (iv) both the joint origin and the axis of rotation are restricted to lie outside the moving limb – this last assumption effectively reduces each joint to having only two degrees of freedom.

Given these assumptions, we now describe our blending piecewise rotation model of articulated joints which provably results in no folding (i.e. the determinant of the Jacobian matrix of the resulting transformations is positive everywhere [2]). Consider the example shown in Figure 1. Here the top part of the cylinder (the moving limb) rotates counter-clockwise, with angle q degrees, with respect to the bottom part of the cylinder (the stationary limb). The axis of rotation shown in yellow (coming out of the paper). If a "naive" piecewise rotation model is used, we observe (Figure 1 middle-top) both folding of the surface where the moving limb moves towards the interface (near point A) and stretching where it moves away from the interface (near point B). Our proposed model (Figure 1 middle-bottom) corrects for both of these problems by appropriately manipulating the angle of rotation. The key to our model is the fact that any complicated 3D rotation can be described as a single rotation about a single axis (the so-called angle-axis representation). Without loss of generality, let the rotation axis pass through the local coordinate origin and be aligned to the $z-$axis. (This can easily be achieved by a global rigid transformation). We next employ a cylindrical polar coordinate system (r, z, θ), where r is the radial distance from the origin and θ the rotational coordinate. In this polar coordinate system, any rotation about the $z-$axis can be expressed as a translation in θ . For example the "naive" piecewise rotation model can be written as:

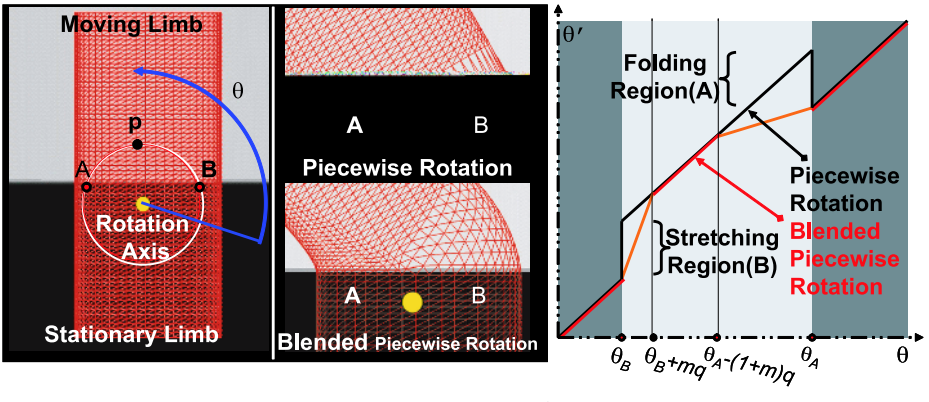

Fig. 1. The blending piecewise model. **Left:** A simple case of articulated motion where the top part of a cylinder is rotating independently of the bottom part. **Middle Top:** Simple piecewise rotation model exhibiting folding (A) and stretching (B) behavior. **Middle Bottom:** Proposed blending piecewise rotation model which is free from folding and stretching. **Right:** Following transformation to a cylindrical polar coordinate system we plot the original (θ) vs the transformed (θ') polar angle for both the piecewise and the blending models. Note that the blending piecewise rotation model results in an continuous invertible mapping unlike the "naive" piecewise rotation model.

$$\text{Static Limb: } (r,z,\theta) \mapsto (r,z,\theta), \quad \text{Moving Limb: } (r,z,\theta) \mapsto (r,z,\theta+q) \tag{1}$$

$$\text{Static Limb: } \theta' = \theta, \quad \text{Moving Limb: } \theta' = \theta + q \tag{2}$$

where in equation 2 we explicitly reduce this mapping to a one dimensional problem of finding $\theta \mapsto \theta'$, since (r, z) remain constant. Using the same notation, we express our blended piecewise rotation model as:

$$\theta' = \theta + q\phi_{r,z}(\theta) \tag{3}$$

where $\phi_{r,z}(\theta)$ is the to-be constructed continuous function on each circle of constant r, z. Since there is no change in the r, z coordinates, the key to constructing an invertible, continuous mapping is reduced to simply constructing an invertible continuous function $\phi_{r,z}(\theta)$ which is constrained to be close to 0 in the static limb and 1 in the moving limb so as to adhere as much as possible to the overall piecewise rotation. One possible solution to this is demonstrated in figure 1(right) where we plot θ against θ'. For a given circle at constant (r, z), let A and B be the points where the circle intersects the motion interface, with angular coordinates θ_A and θ_B respectively. Our blending model sets $\phi_{r,z} = 0$ in the static limb, and performs all blending in the moving limb – this enables easy hierarchical updating of multi-joint structures such as the leg. Next, in the moving limb, in the region immediately after B (i.e. $\theta_B < \theta < \theta_B + mq$, where $m > 0$ is the dimensionless "extent of blending" parameter typically set to 1) $\phi_{r,z}$ ramps up from zero to one to correct for the stretching problem. In the region $\theta_B + mq < \theta < \theta_A - (1+m)q$ we have $\phi = 1$ resulting in unconstrained rotation i.e. $\theta' = \theta + q$. Finally as we approach the moving interface $\theta_A - (1+m)q < \theta < \theta_A$ we ramp down the rotation angle back to zero to ensure continuity at the interface point A. This map is both continuous and invertible; a cubic blending method can also be used to make it C^1 continuous if desired. An easily handled special case occurs when the part of the circle that lies in the moving limb is too short (in an angular sense) such that $\theta_A - (1+m)q < \theta_B + mq$. In this case $\phi_{r,z}$ never reaches one and there is no middle portion to the line, rather the ramping functions intersect. We apply folding correction to the right of the intersection and stretching correction to the left.

The practical application of this blending model consists of finding for any given point $p = (r, z, \theta)$, the θ-coordinates of the appropriate intersection points A, B via a bisection search strategy along the circle centered at $(0, z)$. Potential singularities can exist close to the rotation axis $r = 0$; these are avoided by design (see assumption iv above) by ensuring that the rotation axis lies outside the moving limb. A sufficient condition for this is to ensure that the rotation axis lies on a plane (*the local xy-plane*) that is outside the moving limb – this effectively reduces each joint to having only two degrees of freedom.

2.2 Practical Implementation

First, a label image is constructed where each voxel contains a value equal to the index of each limb (i.e. air=-1, main body=0, left hip=1, left knee=2 etc.). Next, for each joint we identify both the position of the joint and a local *xy-plane*

that stays completely outside the moving limb on which the rotation axis for the joint is constrained to lie. The joint or motion interface is set to be, for each joint, the bounding surface of the static limb. In the case of the left knee joint, for example, the motion interface is the boundary of the static "limb" which is the union of the main body trunk, the whole right leg, and the left proximal leg, whereas the moving "limb" is the union of the left distal leg and the left foot. The overall transformation model consists in applying a series of transformation in a hierarchical manner, i.e. in the case of a point x in the left foot the overall mapping is:

$$x' = T_{\text{global}} . T_{\text{left-hip}} . T_{\text{left-knee}} . T_{\text{left-ankle}}(x) \tag{4}$$

Here T_{global} is an arbitrary global transformation, whereas $T_{\text{left-hip}}$, $T_{\text{left-knee}}$, $T_{\text{left-ankle}}$ are modeled using our blended piecewise rotation model described above. As an example, the 2-DOF (α, q) transformation $T_{\text{left-ankle}}$ is applied to x as follows:

1. Transform the global coordinate system to the local coordinate system such that the joint origin for the ankle is mapped to the coordinate origin, and the normal of the local *xy-plane* to the local z-axis. (Transformation T_1).
2. Initialize the the rotation axis to be the local x-axis and rotate it by an angle α about the local z-axis to obtain the final rotation axis $\mathbf{l}$.
3. Convert to a cylindrical polar system such that $\mathbf{l}$ maps to the cylindrical polar z-axis. (Transformation T_2).
4. For any point (r, z, θ) compute $\phi_{r,z}(, \theta)$ and map to $(r, z, \theta + q\phi_{r,z}(\theta))$, using the construction discussed in Section 2.1, below equation 3.
5. Transform by $T_1^{-1} . T_2^{-1}$ to get back to the original coordinate system.
6. Consecutively apply $T_{\text{left-knee}}$, $T_{\text{left-hip}}$ and T_{global} in order to get the final transformed point x'.

2.3 Articulated Rigid Robust Point Matching

We embed this new transformation model within the robust point matching (RPM) framework [3]. The registration procedure consists of two alternative steps: (i) the correspondence estimation step – which allows for the handling of outliers both in the reference and target point sets and (ii) the transformation estimation step. The correspondence estimation step is identical to the standard RPM implementation [3] and will not be described here. Given two sets of corresponding points we estimate the $N = 6+2n$ parameters of the articulated model, where n is the number of the articulated joints, by optimizing a functional of the form:

$$G = \underset{g}{\arg\min} \sum_i w_i (g(X_i) - V_i)^2 \tag{5}$$

where g is the articulated rigid transformation, X_i is point i in the reference point set, V_i is its corresponding point in the target data set and the weights w_i are set to give less weight to points termed as outliers by the correspondence process. This functional is optimized using a Conjugate Gradient method.

3 Results

3.1 Synthetic Data

To test both the utility of the model as well as the convergence of our registration algorithm, we constructed a three piece (two-joint) synthetic model using the left leg from a real micro-CT mouse image, as shown in Figure 2. The leg is divided into three parts (proximal leg or femur, distal leg and foot) as shown in Figure 2(left) and has two joints (knee,ankle).

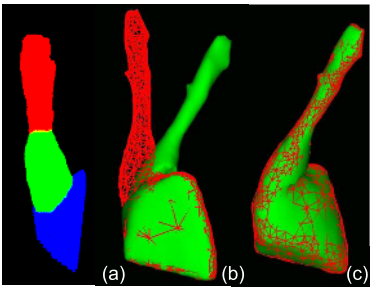

Fig. 2. A Synthetic Mouse Leg Model

A simulated motion of this model using our articulated blended transformation is shown in Figure 2(b) where the original model is shown in red and the transformed model in green. The result of the registration algorithm which consists of the RPM matching strategy and the hierarchical articulated rigid model presented in this paper is shown in Figure 2(c). Note that for this example, the joint rotations where of the order of 25° and the error in the parameter estimates was less than 10% for all parameters.

3.2 Serial Mouse Image Registration

Imaging: We tested the initial utility of our algorithm on three different pairs of mouse micro CT images (resolution $100 \times 100 \times 100 \mu m^3$. In two of the pairs the images were acquired 3 weeks apart as part of an angiogenesis imaging protocol, whereas the third case was specially planned to test this algorithm. In this last example, the mouse was positioned in the scanner, imaged and then removed from the scanner. Then the technologist attempted to position the mouse approximately in the same position as in the first acquisition, thereby simulating the actual situation in the case of serial imaging. The images consisted of the lower half of the mouse (roughly from below the lungs to the feet). All figures are based on this last example.

Articulated Model Construction: Prior to registration, the articulated model with six joints was constructed which is shown in Figure 3 (left). The mouse was divided into seven parts namely the main body trunk, the left and right femurs (or proximal legs), the left and right distal legs and the left and right feet, by manually partitioning the image by interactively placing cutting planes to separate the parts. Whereas a single plane was sufficient for the knee and ankle joints, two planes where needed to delineate the boundary between the hips and the main body. In addition, at each joint, the joint origin was located (at the intersection of the bone and the cutting planes) and the local x- and y-axes were defined such that the local x-axis was roughly aligned to the most likely rotation axis. The joint origin was subsequently translated along the local z-axis away from the cutting plane to ensure that the local xy-plane did not intersect the moving limb at each joint, as required by our model.

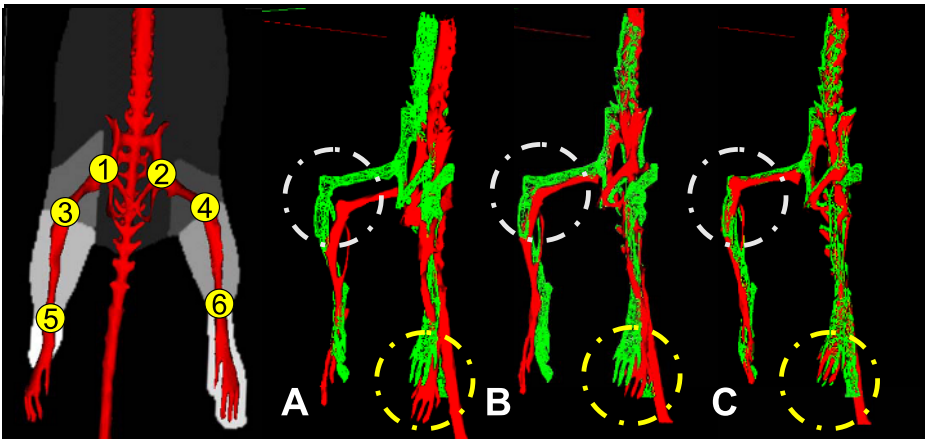

Fig. 3. Left: Schematic of the articulated model with the six joints overlaid on the mouse bony anatomy. (1,2=left/right hip, 3,4=left/right knee and 5,6=left/right ankle) **Right 3 views:** (A) starting position for the registration (red=reference, green=target), (B) result after global rigid alignment, (C) result after articulated rigid alignment. For this result we used points extracted from the bone surfaces of the two images. In particular note the improvement in the registration when using the articulated model at the knee (highlighted by a white circle) and the foot (yellow circle).

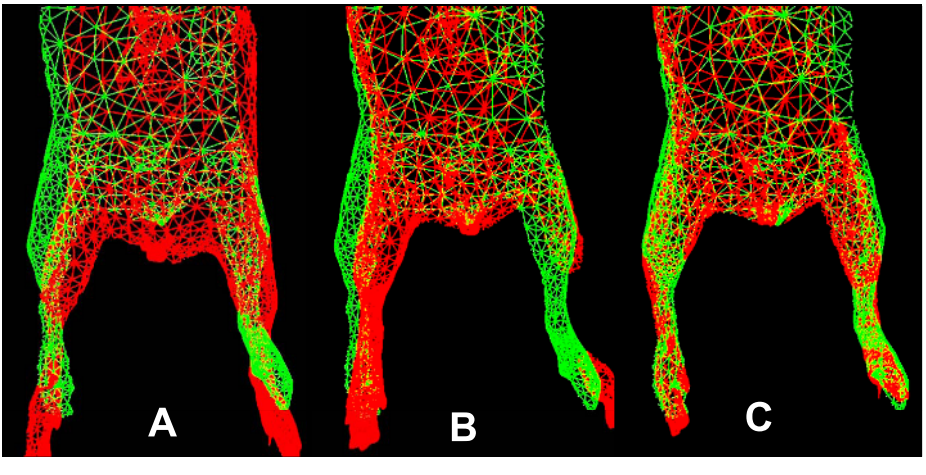

Fig. 4. Surface based registration. In this example, using the same images as in Figure 3 above, we performed the registrations using points sampled from the outer skin surfaces (red=reference, green=target). (A) starting position, (B) after global rigid alignment, (C) after articulated rigid alignment.

Registration: We estimated two different variants of the registration, namely (i) using bone surface points as shown in Figure 3, and (ii) using skin surface points, as shown in Figure 4. While in the case of CT data, the use of the bone

surfaces is optimal for estimating the articulation such surfaces are not easily extractable from MRI-data, hence the use of skin surfaces to test for the more general applicability of the method. Visually, at least, our model performs as expected and successfully captured the articulation at the joints, as shown in the figures. In particular, the bone point version of the algorithm was tested on all three datasets and the algorithm successfully recovered joint rotations in the range of 10 to 40°. For this application we represented the bony anatomy with approximately 600 points. RPM was run with a temperature range 10.0 : 2.0 mm and the annealing factor was set to 0.93 [3].

4 Conclusions

In this paper we presented, the first to our knowledge implementation of an articulated rigid registration which embeds a blending piecewise model of the articulated joint that provably results in a continuous and smooth transformation across motion interfaces. The ultimate goal of this work is to use the output of this algorithm to optimally initialize a non-rigid registration algorithm for capturing the deformation of the soft tissue in addition to the overall limb motion, which will in turn provide accurate registrations of both serial intra-mouse images as well as inter-mouse images. This model of articulation is also applicable to human image data, as for example, in the case of the registration of pre-therapy and intra-therapy images in image guided prostate radiotherapy, where there is a significant articulated motion component.

References

1. V. Arsigny, X. Pennec, and N. Ayache. Polyrigid and polyaffine transformations: A new class of diffeomorphisms for locally rigid or affine registration. In Randy E. Ellis and Terry M. Peters, editors, *Proc. of MICCAI'03, Part II*, volume 2879 of *LNCS*, pages 829–837, Montreal, November 2003. Springer Verlag.
2. G. E. Christensen, M. I. Miller, and M. W. Vannier. Individualizing neuroanatomical atlases using a massively parallel computer. *Computer*, pages 32–38, January 1996.
3. H. Chui, L. Win, R. T. Schultz, J. S. Duncan, and A. Rangarajan. A unified non-rigid feature registration method for brain mapping. *Medical Image Analysis*, 7(2):113–130, 2003.
4. D. L. Collins, G. Le Goualher, and A. C. Evans. Non-linear cerebral registration with sulcal constraints. In *Medical Image Computing and Computer Assisted Intervention*, pages 974–984. Springer, Berlin, 1998.
5. C. Davatzikos. Spatial transformation and registration of brain images using elastically deformable models. *Comp. Vision and Image Understanding*, 66(2):207–222, 1997.
6. M. A. Martin-Fernandez, E. Munoz-Moreno, M. Martin-Fernandez, and C. Alberola-Lopez. Articulated registration: Elastic registration based on a wire-model. In *Medical Imaging 2005: Image Processing*, Proceedings of the SPIE 5747.
7. A. W. Toga. *Brain Warping*. Academic Press, San Diego, 1999.