Optimized knot placement for B-splines in deformable image registration

Travis J. Jacobson^{a)} and Martin J. Murphy

Department of Radiation Oncology, Virginia Commonwealth University, Richmond, Virginia 23298

(Received 15 February 2011; revised 13 June 2011; accepted for publication 15 June 2011; published 22 July 2011)

Purpose: To develop an automatic knot placement algorithm to enable the use of NonUniform Rational B-Splines (NURBS) in deformable image registration.

Methods: The authors developed a two-step approach to fit a known displacement vector field (DVF). An initial fit was made with uniform knot spacing. The error generated by this fit was then assigned as an attractive force pulling on the knots, acting against a resistive spring force in an iterative equilibration scheme. To demonstrate the accuracy gain of knot optimization over uniform knot placement, we compared the sum of the squared errors and the frequency of large errors.

Results: Fits were made to a one-dimensional DVF using 1–20 free knots. Given the same number of free knots, the optimized, nonuniform B-spline fit produced a smaller error than the uniform B-spline fit. The accuracy was improved by a mean factor of 4.02. The optimized B-spline was found to greatly reduce the number of errors more than 1 standard deviation from the mean error of the uniform fit. The uniform B-spline had 15 such errors, while the optimized B-spline had only two. The algorithm was extended to fit a two-dimensional DVF using control point grid sizes ranging from 8×8 to 15×15 . Compared with uniform fits, the optimized B-spline fits were again found to reduce the sum of squared errors (mean ratio = 2.61) and number of large errors (mean ratio = 4.50).

Conclusions: Nonuniform B-splines offer an attractive alternative to uniform B-splines in modeling the DVF. They carry forward the mathematical compactness of B-splines while simultaneously introducing new degrees of freedom. The increased adaptability of knot placement gained from the generalization to NURBS offers increased local control as well as the ability to explicitly represent topological discontinuities. © 2011 American Association of Physicists in Medicine. [DOI: 10.1118/1.3609416]

Key words: Deformable image registration, B-splines, NURBS

I. PURPOSE

Deformable image registration (DIR) in medicine seeks to map the movement of anatomy from one imaging moment to another. The movement is typically represented by a free-form displacement vector field (DVF), which relates the position of each anatomical element in one image to its position in the other. Steep gradients in the DVF and discontinuous motion between organs present significant problems to DIR.^{1–3} This is observed, for example, with lung tissue sliding along the boundary with the lung wall.⁴ The difficulties arise because deformable registration algorithms must regularize and smooth the displacement vector fields, in order to have a well-posed problem with a tractable number of degrees of freedom. This works against the need to model sharp features in the DVF.

Parametric B-spline basis functions are commonly used to model a free-form deformation field.^{5,6} The basis functions are defined to have continuous first derivatives across the knot boundaries connecting the B-spline segments, which forces the DVF to be smooth and continuous everywhere. In the conventional representation of a DVF with B-splines, the spline knots are distributed at equal spacing throughout the registered volume,⁷ resulting in a uniform spatial resolution. To resolve locally sharp features in the DVF, one must distribute a large number of control points throughout the image. This is inefficient and partially defeats the B-spline advantage of local support, and thus has moti-

vated research into B-spline configurations that can be adapted to local structure in the DVF. We note the adaptive-basis approach of Rohde *et al.*⁸ as well as hierarchical B-spline models.^{9,10} In hierarchical models, the knot spacing remains periodic and uniform, but the image volume is subdivided into regions where the local density of knots can be increased. This allows higher resolution near sharp features, but the knot placement is still not optimal.

B-splines have the useful property that one can change the differentiability of the spline tensor at a knot location by placing multiple knots there. For example, when using quadratic B-splines, the placement of two knots at a particular point will allow the DVF to have a discontinuous first derivative there. This would, in principle, allow one to model discontinuous organ motion. The difficulty is that the convention of uniform knot placement is only coincidentally likely to locate a knot at a point of discontinuity. What we really need is a B-spline model that allows the knots to be clustered where they can be most useful in describing local structure, and spread out where the DVF is relatively homogeneous. Ideally, this would be an automatic process.

Nonuniform Rational B-Splines (NURBS) have unequally spaced knots and weighting parameters, in addition to the control points. As a consequence, they are more flexible than uniform B-splines and are widely used in computer graphics to create complex free-form surfaces.^{11–13} They are less frequently used to fit predefined curves and

surfaces,¹⁴ which is the problem posed by DIR. We have begun testing NURBS as a more general basis representation for the DVF.

When fitting a function with NURBS, it is common practice to separate the search for the optimal knot vector from the search for the optimal control points and weights.¹⁵ We have adopted that strategy for now and present here a novel approach for optimizing the knot placement automatically. It is our hope that the method can be generalized to place multiple knots at (or near) points of discontinuous anatomical motion.

II. METHODS

Knot segments define the local support of control points. Increasing the density of knots in a given area of the B-spline curve increases the number of control points active in that area, thereby increasing the ability of the B-spline to emulate sharply varying features of the target DVF. The previous methods of knot placement optimization have used iterative samplings of dominant features in the target function, such as areas of maximum curvature.¹⁶ Instead, we have developed a two-step approach that allows knots to migrate to the areas in which a previous fit, computed with uniform knot spacing, produced the largest errors. This strategy more readily lends itself to the DIR process than the previous methods. While we do not have adequate *a priori* knowledge of the underlying DVF, we can easily obtain intensity difference maps of the two images. This will allow us to anticipate the areas in which we expect to observe the greatest change in the DVF, hence where a high concentration of knots will be most beneficial. For didactic purposes, we demonstrate the algorithm by applying it directly to a known DVF $d(x)$ in one dimension.

An initial fit is made with a uniform one-dimensional B-spline $f(x) = \sum_i c_i B_i(x)$. Distance to the target function $d(x)$ is computed at each data point, and these values are squared to form the error function $\varepsilon^2(x) = [f(x) - d(x)]^2$. Simply moving the knots directly to the points of largest error produced unsatisfactory results. This was due in part to the number of peaks in the error function and the frequency with which multiple peaks would be present along one knot segment. To compensate for these issues, we implemented a force equilibration scheme. The error $\varepsilon^2(x)$ at any particular position x exerts a force to pull the knots toward that point (or, conversely, resisting efforts to move the knot further away). The force is a simple springlike resistance proportional to the product of the distance of the knot from a particular error magnitude $\varepsilon^2(x)$ and the error magnitude itself. By summing over all of the restoring forces on a particular knot i over the range of points x bounded by knot $i - 1$ and $i + 1$, one obtains the knot position at which the forces are equilibrated. For each knot i , solve for the knot position u_i that minimizes the sum of the contributions of each point of the error function contained in the knot segment belonging to u_i

$$\arg \min \int_{x=u_{i-1}}^{x=u_{i+1}} |x - u_i| \cdot \varepsilon^2(x) dx$$

Starting with the first knot, each knot position was solved using the Nelder–Mead Downhill Simplex technique for minimization.¹⁷ The knot positions were found in sequence, such that the updated knot position u_i was used in the calculation of knot position u_{i+1} .

One can see intuitively that if a knot starts out somewhere between two symmetric error peaks, it will settle down halfway between them. If one of the two error peaks is larger, the knot will gravitate closer to it, but will still be restrained by the other errors, causing it to settle at a point that neutralizes the forces on it. The sum total of the restoring forces prevents the knots from piling up at the location of maximum error.

To mimic the hierarchical process of current B-spline DIR implementations, and also to underline the consistent advantage of the knot optimization method, the target DVF is fit iteratively. The initial fit is made with one free knot, and each iteration contributes an additional free knot. During each iteration, the knot placement is optimized by the force equilibration scheme, the target is refit with the optimized knot vector, the error function is updated, a new knot is placed among the optimized knots at the point of the largest error, and the equilibration process is repeated. Iteration proceeds until a predetermined error tolerance criterion is met.

To provide a simple demonstration of our knot placement algorithm, we have fitted a one dimensional profile of a slice of the DVF taken from the POPI-model dataset, as computed by the Demons algorithm.¹⁸ This real-world target data possesses many of the characteristic traits that can illustrate the advantage of nonuniform B-splines over their uniform counterparts: namely, areas of slowly varying displacements contrasted with areas of locally abrupt variations and discontinuities.

We then extended the algorithm to two dimensions, which requires that separate knot vectors be optimized in u and v . This means that the knot positions are restricted to a rectilinear grid rather than being completely free-form. In essence, the optimization moves lines of knots instead of individual knot positions. This is performed by first computing initial fit errors at each point on the surface DVF and

FIG. 1. Sum of squared errors versus number of free knots.

FIG. 2. (Top panel) Target DVF fit with a uniform B-spline. (Bottom panel) Target DVF fit with an optimized nonuniform B-spline. The absolute difference between fit and target at each data point is shown magnified by a factor of 10.

then projecting them onto the x and y axis as one-dimensional error profiles to which we can separately apply the force equilibration scheme described above. This strategy was used to fit a surface DVF taken from a slice of the POPI-model.

III. RESULTS

For the 1D DVF fit, 21 iterations were completed for a total of 22 free knots. Figure 1 plots the sum of the squared errors for the two methods for the first 20 iterations. It can be seen that the optimized nonuniform B-spline fit produced a smaller error at each step than the uniform B-spline fit. The improvement in accuracy ranged from a factor of 1.32 to a factor of 10.57, with a mean factor of 4.02.

Figure 2 shows the uniform and optimized 1D B-spline fits after the final iteration. In the optimized fit, the knots, depicted on the curve by green circles, are concentrated in areas of locally abrupt deformations and away from

FIG. 3. Histogram of errors in the two B-spline fits.

FIG. 4. Contour plot of surface DVF overlaid with the knot positions resulting from optimization.

smoothly varying features of the target DVF. This arrangement of knots allows for more efficient utilization of free parameters in the fitting routine, resulting in smaller error values.

When deformable image registration methods are evaluated for accuracy, they can show an acceptable overall error while having difficulty in regions of high DVF gradients or discontinuities. This can be seen when one makes a histogram of the errors at each voxel or evaluation landmark. In our study, we have been specifically interested in improving the performance of DIR in these difficult regions. We used the error histogram to track improvements by observing changes in the frequency of errors that are more than one standard deviation from the mean.

A histogram of the errors of the Fig. 2 1D fits is included in Fig. 3. The optimized B-spline was found to greatly reduce the number of errors more than 1 standard deviation from the mean error of the uniform fit. The uniform B-spline had 15 such errors, while the optimized B-spline had only two.

In two dimensions, the algorithm separately optimizes two knot vectors, which form an interlocking mesh across the surface. Figure 4 shows a contour plot of the 235-by-141 surface DVF, overlaid with the knot positions resulting from optimization of a 14×14 control point grid. It can be seen that the knots are drawn to areas of sharply varying gradients, and as a result they tend to frame the lungs. This preferential placement of knots results in a smaller error than the uniform B-spline fit because of the more efficient use of the fitting parameters (control points). Error comparisons between the uniform and optimized B-spline fits for different control point grid sizes are contained in Table I.

TABLE I. Ratio of the sum of squared errors and the number of large errors between uniform and optimized B-spline fits.

Control point grid size	8 × 8	10 × 10	12 × 12	14 × 14	15 × 15
Ratio of sum of squared errors	4.39	3.92	1.97	1.52	1.25
Ratio of number of large errors	9.39	7.54	2.47	1.82	1.29

IV. CONCLUSIONS

We have developed a novel approach for knot placement in nonuniform B-spline fitting routines in 1D and showed how it can be extended into higher dimensions. The algorithm is automatic and robust. We have demonstrated that with all else being equal, optimization of the knot spacing results in increased fitting accuracy. By optimizing knot spacing, we are introducing an increase in computation time, but there is not a one-to-one correspondence between control point computation time and knot value computation time. Furthermore, the number of free knots in a three-dimensional grid is much smaller than the number of control points. We are confident that the gain in accuracy will outweigh the tradeoff in computation time, but this is certainly an important issue that we will address when the technique has been fully developed and incorporated into the deformable image registration process.

^{a)}Author to whom correspondence should be addressed. Electronic mail: jacobsonjtj@vcu.edu; Telephone: 804-628-7775.

¹R. Kashani, M. Hub, J. M. Balter, M. L. Kessler, L. Dong, L. Zhang, L. Xing, Y. Xie, D. Hawkes, J. A. Schnabel, J. McClelland, S. Joshi, Q. Chen, and W. Lu, "Objective assessment of deformable image registration in radiotherapy: A multi-institution study," *Med. Phys.* **35**(12), 5944–5953 (2008).

²D. Demirovic, A. Serifovic-Trbalic, E. Skejic, and N. Prljaca, "Evaluation of the diffusion-like accelerated Demons algorithm," in *XXII International Symposium on Information, Communication and Automation Technologies (ICAT 2009)* (2009), pp. 1–4.

³Z. Wu, E. Rietzel, V. Boldea, D. Sarrut, and G. Sharp, "Evaluation of deformable registration of patient lung 4DCT with sub-anatomical region segmentations," *Med. Phys.* **35**(2), 775–781 (2008).

⁴D. Yang, W. Lu, D. A. Low, J. O. Deasy, A. J. Hope, and I. El Naqa, "4D-CT motion estimation using deformable image registration and 5D respiratory motion modeling," *Med. Phys.* **35**(10), 4577–4590 (2008).

⁵D. Rueckert, "Non-rigid registration: Concepts, algorithms, and applications," in *Medical Image Registration*, edited by J. V. Hajnal, D. C. G. Hill, and D. J. Hawkes (CRC, Boca Raton, 2001).

⁶K. K. Brock *et al.*, "Results of a multi-institution deformable registration accuracy study (MIDRAS)," *Int. J. Radiat. Oncol., Biol., Phys.* **76**(2), 583–596 (2010).

⁷J. Kybic and M. Unser, "Fast parametric elastic image registration," *IEEE Trans. Image Process.* **12**(11), 1427–1442 (2003).

⁸G. K. Rohde, A. Aldroubi, and B. M. Dawant, "The adaptive bases algorithm for nonrigid image registration," *IEEE Trans. Med. Imaging* **22**, 1470–1479 (2003).

⁹D. R. Forsey and R. H. Bartels, "Surface fitting with hierarchical splines," *ACM Trans. Graphics* **14**, 134–161 (1995).

¹⁰J. A. Schnabel *et al.*, "A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations," in *Proceedings of the Medical Image Computing and Computer-assisted Intervention (MICCAI) 2208*, edited by W. J. Niessen and M. A. Viergever (Springer-Verlag, London, UK, 2001), pp. 573–581.

¹¹J. D. Foley, *Introduction to Computer Graphics* (Addison-Wesley, Reading, Massachusetts, 1994).

¹²A. Watt and M. Watt, *Advanced Animation and Rendering Techniques* (AMC, Addison-Wesley, Reading, Massachusetts, 1992).

¹³H. J. Lamoussin, "NURBS-based free-form deformations," *IEEE Comput. Graphics Appl.* **14**(6), 59–65 (1994).

¹⁴W. Ma and J. P. Kruth, "NURBS curve and surface fitting for reverse engineering," *Int. J. Adv. Manuf. Technol.* **14**(12), 918–927 (1998).

¹⁵L. Piegel and W. Tiller, *The NURBS Book* (Springer-Verlag, Berlin, 1997).

¹⁶H. Park and J.-H. Lee, "B-spline curve fitting based on adaptive curve refinement using dominant points," *CAD* **39**(6), 439–451 (2007).

¹⁷J. A. Nelder and R. Mead, "A simplex method for function minimization," *Comput. J.* **7**, 308–313 (1965).

¹⁸J. Vandemeulebroucke, D. Sarrut, and P. Clarysse, "The POPI-model, a point-validated pixel-based breathing thorax model," *XVth International Conference on the Use of Computers in Radiation Therapy (ICCR)* (Toronto, Canada, 2007).