

MINIMUM DESCRIPTION LENGTH WITH LOCAL GEOMETRY

Martin Styner^{1,2}, Ipek Oguz¹, Tobias Heimann³, Guido Gerig⁴

Departments of ¹Computer Science and ²Psychiatry, University of North Carolina, Chapel Hill, USA

³Div. Medical and Biological Informatics, DKFZ, Heidelberg, Germany

⁴Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, USA

ABSTRACT

Establishing optimal correspondence across object populations is essential to statistical shape analysis. Minimizing the description length (MDL) is a popular method for finding correspondence. In this work, we extend the MDL method by incorporating various local curvature metrics. Using local curvature can improve performance by ensuring that corresponding points exhibit similar local geometric characteristics that can't always be captured by mere point locations. We illustrate results on a variety of anatomical structures. The MDL method with a combination of point locations and curvature outperforms all the other methods we analyzed, including traditional MDL and spherical harmonics (SPHARM) correspondence, when the analyzed object population exhibits complex structure. When the objects are of simple nature, however, there's no added benefit to using the local curvature. In our experiments, we did not observe a significant difference in the correspondence quality when different curvature metrics (e.g. principal curvatures, mean curvature, Gaussian curvature) were used.

Index Terms— Correspondence, Image Shape Analysis, Modeling, Statistics, Image Registration

1. INTRODUCTION

Since its publication by Cootes et al. [1], Active Shape Models have become very popular in medical imaging. However, this method introduces the problem of correspondence: how to choose points on a surface such that the chosen points correspond across a population of objects. In 2D, this can be done by manually chosen landmarks [2]. However, this becomes a complicated problem in 3D, because manually choosing points is a very cumbersome process given the high number of landmarks needed. This process also is more than likely to lead to inconsistent results among different experts. Automated methods are therefore necessary for establishing correspondence. Christensen [3], Szeliski [4] and Rueckert [5] describe methods for warping the space in which the shapes are embedded, such that models can then be built from the resulting deformation field. Meier [6] uses shape warping using spherical harmonics, but also takes into

This research is supported by the NIH Roadmap for Medical Research Grant U54 EB005149-01, UNC Neurodevelopmental Disorders Research Center HD 03110, and the NIH NIBIB grant P01 EB002779. The caudate and striata segmentations used in this manuscript were provided by J.Levitt and M.E.Shenton (Harvard/BWH). The femur segmentations were provided by M.Fleute (Praxim-Medivision). The lateral ventricle segmentations were provided by D.Weinberger and D.Jones (NIMH). We are thankful to J.Miller (GE Research) and S.Barre (Kitware) for their valuable discussions regarding the development of our tools.

account differences in normals and curvature. Brett [7], Rangarajan [8] and Tagare [9] propose shape features to establish point correspondences. Currently, the Minimum Description Length (MDL) method, introduced by Davies et al. [10, 11, 12], is the most commonly used technique to solve this problem in an automated fashion. However, MDL performs rather poorly on objects with complicated geometry, when the surface is convoluted and therefore the vertex locations exhibit sharp changes that cause 'problem areas' for the optimizer. In this work, we are exploring ways to enhance the MDL method to overcome such issues.

The main idea of our method is to use the data matrix of the MDL method to encode information other than the vertex locations of the surface meshes. Since we would like to explore the local geometry of the objects, a natural source of additional information is local curvature measurements. We also experiment with using the curvature metrics combined with the vertex locations. We apply all these techniques to a wide variety of anatomical correspondence problems.

It should be noted that this technique can be used to incorporate more than local geometry information to the correspondence computation. One can use various types of knowledge about the objects, such as measurements obtained through fMRI or DTI, to improve correspondence quality.

2. METHODOLOGY

In this work, we are comparing results from many different techniques. These can be summarized as 1) Traditional MDL, 2) MDL with curvature only (CurvMDL), 3) MDL with location and curvature (CombinationMDL), and 4) Spherical harmonics (SPHARM) correspondence. We describe each of these techniques in this section. Note that the results from SPHARM correspondence were used as an initialization point for all MDL-based methods.

2.1. Traditional MDL

The MDL method [12] is an information theoretic approach to the correspondence problem. The main idea is that the 'simplest' description of a population is the best; simplicity is measured in terms of the length of the code to transmit the data as well as the model parameters. In this paper, we are using the gradient descent optimization implementation as described in [13, 14], which is based on a slightly modified version of the MDL objective function for computational simplicity, as suggested by Thodberg et al. [15].

2.2. MDL with Curvature (*CurvMDL*)

To use local curvature in the MDL computation rather than vertex locations, we simply substitute the data matrix encoding the spatial locations in traditional MDL. In this alternative matrix, the columns are the local curvature measurements of the objects instead of the spatial locations. Therefore, the cost function for MDL is modified to be based on the eigenvalues of the curvature matrix instead of the eigenvalues of the location matrix.

The specific metric to be chosen for representing curvature is one of the issues in integrating local geometry into the MDL cost function. We have experimented with three pairs of complementary curvature metrics. All curvature metrics have been computed analytically based on the SPHARM coefficients.

1. Principal Curvatures (*Kappa*): The sorted eigenvalues of the second fundamental form are called the principal curvatures κ_1 and κ_2 . These values measure the maximum and minimum values of bending of a surface, and they bound the local values of the normal curvature.
2. Mean Curvature and Gaussian Curvature (*HK*): The mean of the principal curvatures, defined as $H = \frac{1}{2}(\kappa_1 + \kappa_2)$, is called the mean curvature. Koenderink [16] describes it as "the nose dive averaged over all directions", where the 'nose dive' refers to the amount of twist-free turning of the principal frame field. The Gaussian curvature, defined as $K = (\kappa_1 * \kappa_2)$, is a measure of the spherical spread of the surface normals [16].
3. Curvedness and Shape Index (*CS*): An intuitive representation of the local curvature is given by the curvedness C and the shape index S , introduced by Koenderink [16]. The shape index S , taking values in the interval $[-1..1]$, describes the local shape in terms of concavity and convexity. The curvedness C , taking values in $(-\infty.. \infty)$, represents how curved the surface is. C and S basically correspond to a polar representation of the principal curvatures, in an attempt to decouple measurements for the size and shape of the curved surface. This can be a useful property when comparing objects of different sizes (and thus different curvature ranges), as trying to match variables with different ranges could adversely affect the optimization process.

2.3. MDL with Curvature and Location (*CombinationMDL*)

Meier [6] uses a combination of curvature and location in solving pairwise correspondence. Following his example, we explored using the curvature and location together in a population-wise correspondence setting. Once again, this is done by simply substituting the data matrix encoding the spatial locations in traditional MDL. However, we are weighting the curvature measurements in order to normalize them with respect to spatial location values, or, more specifically, to ensure that the two metrics have similar local variances (as the variance, rather than absolute values, is what the MDL cost function depends upon). The specific weights are defined as the inverse of the average local variance, computed separately for spatial location and each curvature metric.

2.4. SPHARM Correspondence

Introduced by Brechbühler et al. [17], SPHARM is a parametric surface description computed via optimizing an equal area mapping of a 3D quadrilateral voxel mesh onto a unit sphere, while minimizing angular distortions. The basis functions of the parameterized

surface are spherical harmonics. The SPHARM correspondence is determined by aligning the parameterization so that the ridges of the first order ellipsoid coincide. Note that SPHARM can only represent objects of spherical topology.

3. EVALUATION OF CORRESPONDENCE QUALITY

To compare the results of the various correspondence methods, evaluation metrics are needed. The choice of evaluation metrics to be used is important since the definition of a "good" correspondence can greatly vary among different applications. In this work, we are using the well established generalization and specificity metrics [12, 18], which favor structurally feasible correspondence. This is a relevant correspondence evaluation method for anatomically defined correspondence as presented in this work.

Given a statistical shape model, generalization is a measure of how well the model can describe unseen objects of the same class. Generalization ability is especially important when the model should learn the characteristics of the object class from a limited training set. In practice, the generalization ability is computed by leave-one-out reconstruction from the set of objects and averaging the reconstruction error. Therefore, a good model should exhibit low generalization values. Generalization is denoted by $G(M)$, where M is the number of shape eigenmodes used in reconstruction.

Specificity is a metric of how well the model fits the object class, in that it measures the distance between objects in the training set and new objects generated using the model. A specific model should only generate objects similar to those in the training set. The specificity $S(M)$, again a function of the number of shape eigenmodes M , is computed via generating a large number of random objects from the model shape space and comparing them to the objects in the training set.

In Figs. 2- 5, the generalization and specificity measures for various correspondence methods are plotted as a function of the number of shape eigenmodes M . Note that the case $M = 0$ corresponds to the population average. $G(0)$, therefore, measures how far, on the average, the individual shapes are from the population average. $S(0)$, on the other hand, measures the distance between the population average and the individual shape that is closest to that average.

4. RESULTS AND DISCUSSION

We have applied the presented correspondence methods to four different object populations: lateral ventricles, caudate, striata (which is the union of caudate, putamen and nucleus accumbens), and left femoral heads. All populations of brain structures (lateral ventricles, caudate and striata) included healthy subjects as well as patients with various disorders. This setting leads to a higher variability in the populations.

For all populations, traditional MDL performed better than SPHARM, which agrees with findings reported in [18]. SPHARM correspondence does not handle well rotational variations in the populations, since it is based on first order ellipsoid alignment. It is surprising, however, that SPHARM nonetheless performed better than pure curvature MDL (*CurvMDL*) for all populations. We believe this is due to the noisy nature of curvature measurements. The natural solution is to combine location and curvature for MDL computation. The performance of *CombinationMDL* is satisfactory for all populations. However, for some populations, there seems to be little reason to use *CombinationMDL* instead of the traditional MDL. This finding is expected since the traditional MDL is quite

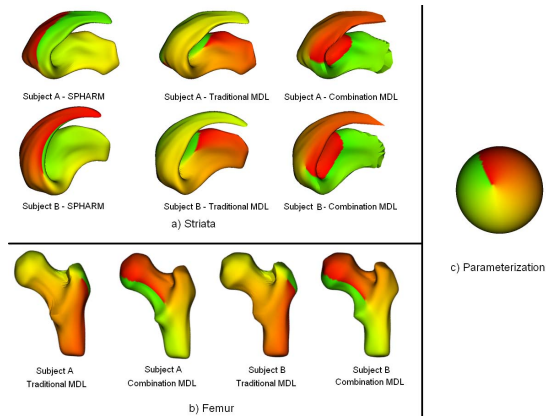


Fig. 1. The results of the SPHARM, traditional MDL and CombinationMDL demonstrated on two striata and two femoral heads. The coloring shows the ϕ coordinates on the spherical parameterization of the objects, with each (x, y, z) point on the object surface mapped to a (ϕ, θ) point on the unit sphere. Locations with same coloring on the two subjects correspond to each other for a given method. As shown in (c), the line separating green and red regions corresponds to a longitude line on the unit sphere. Note that the SPHARM correspondence for the striata is very poor, and the CombinationMDL method gives the best visual correspondence for both populations.

adequate for simpler shapes like caudates. However, for objects with complex curvature patterns, such as the striata and the femur illustrated in Fig. 1, the incorporation of curvature improves the correspondence.

For all populations, the choice of particular curvature metric proved not to be critical, since any of the 3 pair of metrics we used resulted in performances in the same range. This result is not surprising, since given any pair of metrics we presented, one can easily compute the other 2 pairs: each pair encodes the same information about the surface shape. The limited differences between the results from the various metrics is mainly due to the different numerical properties of the metrics, as these influence the performance of the optimizer. To improve graph readability, we only show results using the C-S pair of metrics on most of the figures in this paper.

The first population consists of 116 lateral ventricle segmentations. For this population, the traditional MDL performs very well, and introducing additional curvature information does not have any benefits. In fact, since curvature measurements are more prone to sampling noise than locations, using the curvature data is similar to adding noise to the dataset. However, the performances of the two methods are still in the same range. Figure 2 summarizes these results.

The next dataset includes 56 caudate segmentations. The traditional MDL and the CombinationMDL perform similarly well on this population as expected, given the highly simple shape of the objects in the population. In fact, Figure 3 demonstrates that even SPHARM correspondence performs reasonably well on this dataset, since there are no twists in the structure.

The third dataset we used consists of 18 femoral heads. Here, due to the complex structure of the objects, the correspondence improves when curvature information is included. Note that, for the generalization metric, traditional MDL catches up with the performance of CombinationMDL when a high number of shape eigenmodes (M) are used. The CombinationMDL has superior specificity values independent of M . Figure 4 captures these results.

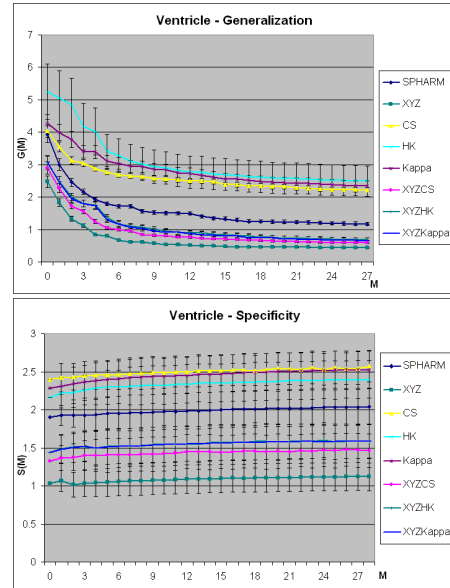


Fig. 2. The generalization and specificity comparison on a population of lateral ventricles. In all figures, XYZ refers to the traditional MDL method; CS, HK and Kappa refer to CurvMDL method using the specified curvature metrics; XYZCS, XYZHK and XYZKappa refer to CombinationMDL method. Note that the choice of particular curvature metric has very little effect on the results for both CurvMDL and CombinationMDL methods. Therefore, only results using the C-S metrics will be shown in the subsequent figures to improve graph readability. For the lateral ventricle population, SPHARM and CurvMDL both perform poorly. The performance of traditional MDL and CombinationMDL is within the same range.

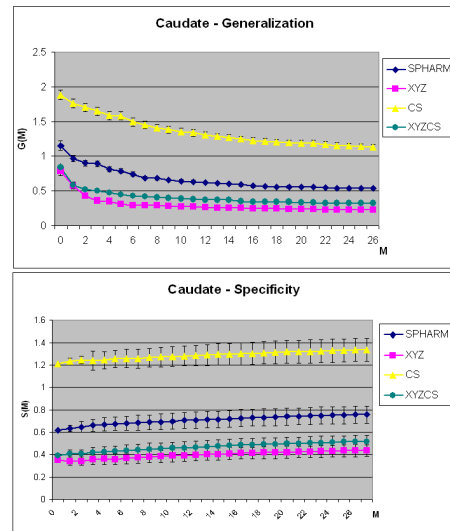


Fig. 3. The generalization and specificity comparison on a population of caudates. Given the very simple structure of the objects, there is no extra benefit in adding curvature information to MDL computation. Even SPHARM performs well on this dataset, since the first order ellipsoid alignment is satisfactory. The pure curvature method (CurvMDL) performs poorly.

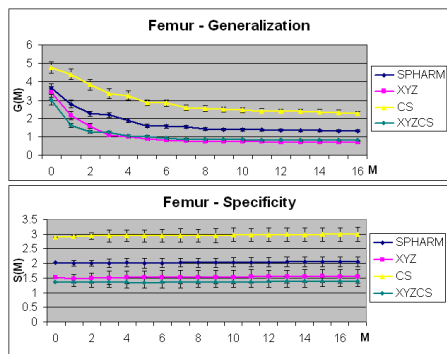


Fig. 4. The generalization and specificity comparison on a population of femurs. CombinationMDL provides an improved correspondence compared to traditional MDL, even though the improvement is negligible when a higher number of shape eigenmodes (M) is used. SPHARM and CurvMDL both perform poorly.

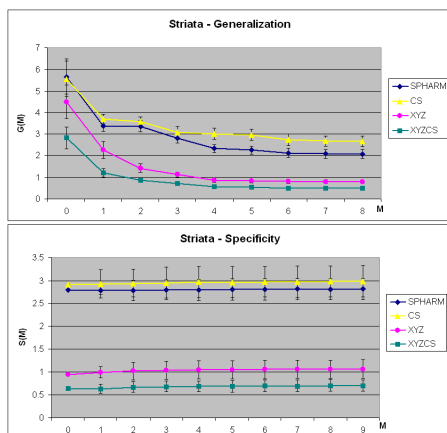


Fig. 5. The generalization and specificity comparison on a population of striata. CombinationMDL clearly improves correspondence. SPHARM and CurvMDL both perform poorly.

Finally, we have compared the methods on a population of 20 striata. The striatum best demonstrates the type of object where using curvature for establishing correspondence is beneficial. A typical striatum has both highly convex and highly concave surface patches, and the curvature pattern quickly changes along the surface. Traditional MDL does not perform optimally on these convoluted objects because the vertices with similar spatial locations across subjects do not necessarily correspond together. In this case, the local geometry becomes varied enough that it provides a means of identifying corresponding points, and thus it helps to use curvature as a stabilizer on striatum correspondence. The results are shown in Figure 5.

5. CONCLUSIONS

In this paper, we presented a framework for integrating curvature measurements into MDL computation. We compared results of SPHARM correspondence, traditional MDL, as well as pure-curvature (CurvMDL) and combination MDL, on 4 populations of different anatomical structures. Our method invariably produces

results at least comparable to traditional MDL, and handles objects with highly complex curvature structure much better than traditional MDL.

Incorporating local measurements into MDL correspondence computation allows many new research paths. Other than geometric information such as curvature, any other measurements can also be used to improve correspondence. In future work, we plan to explore possibilities of using connectivity (DTI), functional (fMRI) and angiographic (MRA) data. This would allow the correspondence to be a function of not only object shape, but also of brain connectivity or function.

6. REFERENCES

- [1] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models – their training and application," *Computer Vision and Image Understanding*, vol. 61, no. 1, pp. 38–59, 1995.
- [2] F. L. Bookstein, *Morphometric Tools for Landmark Data: Geometry and Biology*, Cambridge University Press, 1991.
- [3] G. Christensen, S. Joshi, and M. Miller, "Volumetric transformation of brain anatomy," *IEEE Trans. Med. Imag.*, pp. 864–877, 1997.
- [4] R. Szeliski and S. Lavallee, "Matching 3-d anatomical surfaces with non rigid deformations using octree-splines," *Int. J. Computer Vision*, 1996.
- [5] D. Rueckert, A.F. Frangi, and J.A. Schnabel, "Automatic construction of 3d statistical deformation models using non-rigid registration," in *MICCAI*, 2001, pp. 77–84.
- [6] D. Meier and E. Fisher, "Parameter space warping: Shape-based correspondence between morphologically different objects," *Trans. Med. Imag.*, vol. 21, no. 1, pp. 31–47, 2002.
- [7] A.D. Brett and C.J. Taylor, "Construction of 3d shape models of femoral articular cartilage using harmonic maps," in *MICCAI*, 2000, pp. 1205–1214.
- [8] A.Rangarajan, H. Chui, and F.L. Bookstein, "The softassign procrustes matching algorithm," in *IPMI*, 1997, pp. 29–42.
- [9] H. Tagare, "Shape-based nonrigid correspondence with application to heart motion analysis," in *IEEE Trans. Med. Imag.*, 1999, pp. 570–580.
- [10] Rhodri H. Davies, Carole J. Twining, Tim F. Cootes, John C. Waterton, and Chris J. Taylor, "A minimum description length approach to statistical shape modelling," *IEEE trans. Medical Imaging*, vol. 21, no. 5, pp. 525–537, 2002.
- [11] Rhodri H. Davies, Carole J. Twining, Timothy F. Cootes, John C. Waterton, and Christopher J. Taylor, "3D statistical shape models using direct optimisation of description length," in *Proc. European Conference on Computer Vision, Part III*. 2002, pp. 3–20, Springer.
- [12] Rhodri H. Davies, *Learning Shape: Optimal Models for Analysing Shape Variability*, Ph.D. thesis, University of Manchester, Manchester, UK, 2002.
- [13] Tobias Heimann, Ivo Wolf, Tomos G. Williams, and Hans-Peter Meinzer, "3D active shape models using gradient descent optimization of description length," in *Proc. IPMI*. 2005, pp. 566–577, Springer.
- [14] Tobias Heimann, Ipek Oguz, Ivo Wolf, Martin Styner, and Hans-Peter Meinzer, "Implementing the automatic generation of 3d statistical shape models with itk," *ISC/NA-MIC Workshop on Open Science at MICCAI 2006*, 2006.
- [15] Hans Henrik Thodberg, "Minimum description length shape and appearance models," in *Proc. IPMI*, 2003, pp. 51–62.
- [16] J. J. Koenderink, *Solid Shape*, MIT Press, 1990.
- [17] C. Brechbühler, G. Gerig, and O. Kübler, "Parametrization of closed surfaces for 3-D shape description," *Computer Vision and Image Understanding*, vol. 61, pp. 154–170, March 1995.
- [18] Martin Styner, Kumar T. Rajamani, Lutz-Peter Nolte, Gabriel Zsemlye, Gábor Székely, Christopher J. Taylor, and Rhodri H. Davies, "Evaluation of 3D correspondence methods for model building," in *Proc. IPMI*, 2003, pp. 63–75.