Symmetry Restoration by Stretching

Misha Kazhdan* Nina Amenta† Shengyin Gu† David F. Wiley† Bernd Hamann†

Abstract

We consider restoring the bilateral symmetry of an computer model of an object which has been deformed by compression. This problem arises in paleontology, where symmetric bones are typically deformed in the process of fossilization. Our input is a user-selected set of point-pairs on the deformed object, which are assumed to be mirror-images on the original symmetric object, with some added noise. We carefully consider the formulation of the problem, and give a closed-form solution.

1 Introduction

Much of what we know about evolution comes from the study of fossils. From the shapes of the bones of extinct animals we form hypotheses about how they moved, what they ate, how they are related to each other, and so on. Yet these shapes are usually deformed by the geological processes which occur during fossilization, for example the skull in Figure 1. For some fossils, for example skulls and vertebrae, we can assume that the original shape was roughly bilaterally symmetric. We can use this assumption to reverse the deformation, or at least limit the family of possible reconstructions. This process is sometimes called retrodeformation.

Usually the input for retrodeformation is a set of point-pairs, chosen by the paleontologist on the deformed specimen. We assume the point-pairs are stored in a $3 \times 2n$ matrix P with the assumption that point p_{2i} was the mirror image of p_{2i+1}, on the original object, before it was deformed by geological forces. The point-pairs are chosen using the expert's experience of the biological shape. We use this approach, although we note that developing automatic methods for finding point-pairs or other useful descriptions of the input data is another, quite interesting, research question (see [5]).

We consider the problem of finding a deformation which restores bilateral symmetry to the point-pairs. To get a reasonable mathematical formulation, we need to consider a limited set of transformations. We consider *single axis stretches*. A single axis stretch is produced by choosing a direction vector and scaling only in that direction; it is represented by a symmetric matrix A for which two of its eigenvalues are one and the third is greater than one. Single-axis stretches are important, since the simplest hypothesis for how a fossil is deformed is that it is compressed in a single direction. We want to find a single-axis stretch A such that AP is as nearly symmetric as possible.

*Department of Computer Science, Johns Hopkins University

†Department of Computer Science, University of California at Davis

Figure 1: A deformed dinosaur skull in the Carnegie Museum of Science.

Problem 1 *Let P be a set of point-pairs. Find the single-axis stretch A, a translation vector t, and a plane of reflection, such that the mean-squared error*

$$E = \sum_{i=1}^{n} ||A(p_{2i} + t) - \mathit{Refl}_{v,c}(A(p_{2i+1} + t)||^2 \qquad (1)$$

is minimized. Here $\mathit{Refl}_{v,c}$ is the affine transformation reflecting space across the plane with normal v passing through point c.

The choice of mean-squared error is natural, and consistent with usual practice in paleontology.

But there is a problem with this formulation: there is not a unique solution in the absence of noise. Instead, there is a one-dimensional set of single-axis stretches that produce different perfectly symmetric shapes. As an analogy, think of fitting a plane to set of points that lie on a line; there is no unique solution. If noise is added, there is a unique solution, but it provides information only about the noise, not about the unknown plane that contains the points. Similarly, when P is noisy the unique minimum error solution selects one of the possible symmetrizing single-axis stretches, but based on the noise rather than on any information about the original shape. Instead of returning a single solution, our computation returns a description of the entire set of possible symmetrizing single-axis stretches. We also consider choosing, as a canonical solution, the one requiring minimum deformation. If there are other criteria - comparison with other fossils, for instance - these could be used instead to select a solution.

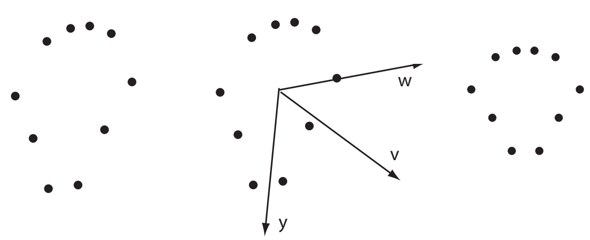

Figure 2: Two ideas for retrodeformation. On the left, a perfectly symmetric set of point-pairs, deformed by compression along a single axis. Center, it seems intuitively clear that stretching in direction v is the most efficient way to make w and y perpendicular. Right, making the entire point set isotropic also makes it symmetric.

Our results: In this paper we tie together two ideas for restoring symmetry, illustrated in Figure 2. The first is a procedure for retrodeformation suggested by the physical anthropolgist Christoph Zollikofer [12] (Appendix E): given a vector w estimating the average direction of the vectors $p_{2i} - p_{2i+1}$, and an estimate y of the projection of that vector on the sagittal plane of reflection, stretch in the direction v bisecting $\angle w, y$ until they become perpendicular. This method is presented without a proof of optimality. We connect it with a "well known" idea in the area of symmetry detection: if you transform P to an isotropic matrix $\tilde{P}$ (that is, the principal components of $\tilde{P}$ are all vectors of length one), then it becomes as symmetric as possible. We prove that the least-squares best-fit plane T to the mid-points of the segments $\tilde{p}_{2i}, \tilde{p}_{2i+1}$ is the plane of symmetry that minimizes the error E above for $\tilde{P}$. This nice property remains true for all linear transformations which maintain symmetry across T (the *symmetry preserving transformations*). including the symmetrizing single-axis stretches. To produce a cannonical solution within this set of symmetrizing single-axis stretches, we choose the vectors w and y using T, and we prove that applying Zollikofer's formula we produce the symmetrizing single-axis stretch which minimizes the deformation of P. **Related work:** In paleontology, this problem has been approached in different ways. An article by Motatni [7] gave a closed-form minimum-error solution in two dimensions, using a somewhat different set-up. A similar method has recently been used to study trilobites [1]. Our work is more closely related to that of Zollikofer [12].

More research in computer science and morphometrics has focused on detecting symmetry; see prior work by the first author [2] and references therein, [9], [10], [11], and, in morphometrics [4] and [3]. A notable exception is [6], where detected approximate symmetries were grouped and aligned to restore the symmetry of bent objects (ie, straightening out a snake).

2 Isotropy and symmetry preserving transformations

We assume throughout that P is not co-planar and, without loss of generality, that the center of mass of P is at the origin.

We say a set of points P is *isotropic* if its 3×3 covariance matrix $PP^t = I$ (all of its principle components are one). We say a set P of point-pairs is *symmetric* if there exists a plane T through the origin with normal v such that P is symmetric across T. We say P is *perfectly symmetrizable* if there exists any matrix A such that AP is symmetric. We build on the following key idea.

Observation 1 *If P is perfectly symmetrizable and isotropic, then P is symmetric.*

This observation follows from the work of [8]. We use this idea as follows: we first apply a linear transformation, to produce an isotropic set $\tilde{P}$, and then we find an optimal plane of symmetry. The transformation $M^{-1/2}$ taking P to $\tilde{P}$ will most likely not be a single-axis stretch, but finding the optimal plane of symmetry is a good first step. Details of the definition of $M^{-1/2}$ can be found in Appendix A.

Still assuming that $\tilde{P}$ is symmetric, we consider the optimal plane of symmetry T of $\tilde{P}$ and the transformations that preserve symmetry across it. Let R be any rotation matrix which takes T into the plane $x = 0$. The symmetry of $\tilde{P}$ is preserved by the multiplication $SFR\tilde{P}$ where S is any rotation and F is any matrix of the form

$$F = \begin{bmatrix} a & 0 & 0 \\ 0 & b & c \\ 0 & d & e \end{bmatrix} \quad (2)$$

This gives us a set of transformations $V = SFR$, such that any $V\tilde{P}$ is symmetric. We call these the *symmetry preserving transformations* of $\tilde{P}$. The set of symmetry preserving transformations is seven dimensional; although there are five degrees of freedom in choosing F and three degrees of freedom in choosing S, the fact that rotating one choice of F about the x-axis produces some other choice of F reduces the dimensionality to seven.

3 Cross-covariance

We define the 3×3 *cross-covariance* matrix C_Q of a set of point pairs Q as:

$$C_Q = \sum_i q_{2i} q_{2i+1}^t + q_{2i+1} q_{2i}^t$$

This matrix expresses the covariance of the points on one side with respect to the points on the other. Since C_Q is symmetric, it can be decomposed into $C = R^t \Lambda R$, with smallest eigenvalue λ_1.

Lemma 2 *Let Q be any set of point-pairs, not necessarily symmetric or isotropic. Let u be the eigenvector with minimum eigenvalue of the cross-correspondence matrix C_U of Q. The plane T_Q through the origin with normal u is the least-squares best-fit plane to the set of midpoints $\tilde{m}_i = (q_{2i} + q_{2i+1})/2$ of Q.*

Proof. We consider the sum of the squared distances of all of the midpoints $\tilde{m}_i$ to the plane through the origin with normal u:

$$\begin{gathered} \sum_i \langle (q_{2i} + q_{2i+1})/2, u \rangle^2 = \\ 1/4 \left(u^t \left(\sum_i (q_{2i} + q_{2i+1})(q_{2i} + q_{2i+1})^t \right) u \right) = \\ 1/4 (u^t (QQ^t) u + u^t C u) = 1/4 (1 + u^t C_Q u) \end{gathered}$$

Since C_Q is symmetric, the error is minimized when u is the eigenvector of C_Q corresponding to the smallest eigenvalue λ_1. □

Notice that when Q *is* symmetric, T_Q passes through the midpoints and is exactly the plane of symmetry of Q.

Here is a useful property of the cross-covariance matrix of a symmetric set of point pairs.

Lemma 3 *Let P be any symmetric set of point-pairs. The cross-covariance matrix C has exactly one negative eigenvalue.*

Intuitively, this eigenvalue corresponds to the reflection; the proof is found in Appendix B. A general P, not necessarily symmetric, might not have exactly one negative eigenvalue, in which case it would not much resemble a set of symmetric point-pairs. We say that P is *approximately symmetrizable* if C does indeed have exactly one negative eigenvalue.

4 Noise and optimality

Even when the input P is not perfectly symmetrizable, we can use the cross-covariance matrix of $\tilde{P}$ to find an approximate plane of symmetry T. We let R be the rotation in the decomposition $C_{\tilde{P}} = R^t \Lambda R$ of the cross-covariance matrix of $\tilde{P}$. Then we define $T = R^t T_0$, where T_0 is the plane $x = 0$. By Lemma 2, T is the least-squares best fit to the midpoints of $\tilde{P}$.

We now show that T is also the optimal plane of symmetry, not only for $\tilde{P}$ but for any symmetry-preserving transformation of $\tilde{P}$. For convenience, we will first consider only symmetry-preserving transformations of the form $V = FR$, that is, the ones that transform P so that T goes to $x = 0$.

Lemma 4 *Let P be an approximately symmetrizable set of point-pairs. Then T is the plane minimizing the symmetry error of (Equation 1) for the transformed set $V\tilde{P}$, where $V = FR$.*

The proof can be found in the Appendix.

Theorem 5 *Let P be a set of approximate symmetrizable point-pairs. Then ST is the plane minimizing the symmetry error of (Equation 1) for the transformed set $V\tilde{P}$, where $V = SFR$ is any symmety preserving transformation of $\tilde{P}$. Also, ST is the least-squares best-fit plane to the midpoints of the transformed points of $V\tilde{P}$.*

This theorem follows from Lemmas 4 and 2.

5 Single-axis stretches

Single-axis stretches have the form:

$$A = (\alpha - 1) u u^t + I$$

where u is the unit vector in the direction of stretching, and the stretching factor is α, which we define to be greater than one. The entire set of single-axis stretches is three-dimensional, but not all single-axis stretches are symmetrizing transformations ($A = VM^{-1/2}$). Consider the three eigenvectors v_1, v_2, v_3 of C, the cross-covariance matrix of $\tilde{P}$. Let w_1, w_2, w_3 be the pre-images of these vectors with respect to $M^{-1/2}$; eg. $u_1 = M^{-1/2} w_1$. The single-axis stretches which are also symmetrizing transformations are exactly those which make Aw_1 perpendicular to the plane ST spanned by

Aw_2, Aw_3. This introduces two additional constraints, so that the dimension of the set of symmetrizing single-axis stretches is only one.

6 Optimal single-axis stretch

Single-axis stretches are symmetric matrices, so the condition $\langle Aw_1, Aw_2 \rangle = 0$ can be rewritten as $w_1^t A^t A w_2 = \langle A^2 w_1, w_2 \rangle = 0$, and we rewrite the condition $\langle Aw_1, Aw_3 \rangle = 0$ similarly. Observe that $A^2 = (\alpha^2 - 1)vv^t + I$, that is, we stretch twice in the same direction. Our condition that Aw_1 should be perpendicular to both Aw_2, Aw_3 means that $A^2 w_1$ should be some multiple of $w_2 \times w_3$. We construct a 2D coordinate system for the plane through w_1 and $w_2 \times w_3$:

$$X = \frac{w_2 \times w_3}{||w_2 \times w_3||} \quad Y = \frac{w_1 - \langle w_1, X \rangle X}{||w_1 - \langle w_1, X \rangle X||}$$

Here X is the normal of the plane spanned by w_2, w_3 and Y is the projection of w_1 onto that plane. Let θ be the (known) angle such that

$$w_1 = \cos(\theta) X + \sin(\theta) Y$$

Any single-axis stretch which aligns w_1 and X does so by stretching along a vector v that lies in the plane containing them. (We prove this formally in the Appendix). Vector v must also lie in this plane, so we can write v as:

$$v = \cos(\phi) X + \sin(\phi) Y$$

So we have the situation in Figure 3. We can re-write

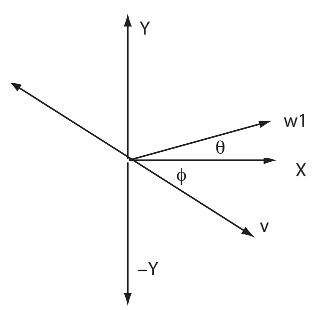

Figure 3: Both the vector v along which stretching occurs and w_1 lie in the same plane. X is the normal to the plane spanned by w_2, w_3 and Y is the projection of w_1 onto that plane. Intuitively, the optimal choice for v is the vector half-way between w_1 and $-Y$, and indeed this is what the proof indicates.

$$\begin{aligned} A^2 w_1 &= (\alpha^2 - 1)\langle v, w_1 \rangle v + w_1 \\ &= (\alpha^2 - 1)[\cos\phi\cos\theta + \sin\phi\sin\theta][\cos\phi X + \sin\phi Y] \\ &\quad +(\cos\theta X + \sin\theta Y) \\ &= \left[(\alpha^2 - 1)[\cos^2\phi\cos\theta + \cos\phi\sin\phi\sin\theta + \cos\theta]\right] X \\ &\quad + \left[(\alpha^2 - 1)[\cos\phi\sin\phi\cos\theta + \sin^2\phi\sin\theta + \sin\theta]\right] Y \end{aligned}$$

Since $A^2 w_1$ should be perpendicular to Y, this implies that

$$0 = (\alpha^2 - 1)[\cos(\phi)\sin(\phi)\cos(\theta) + \sin^2(\phi)\sin(\theta)] + \sin(\theta)$$

Solving for α^2, we get

$$\begin{aligned} \alpha^2 &= 1 - \frac{\sin(\theta)}{\cos(\phi)\cos(\theta)\sin(\phi) + \sin^2(\phi)\sin(\theta)} \\ &= 1 - \frac{2\sin(\theta)}{\sin(2\phi - \theta) + \sin(\theta)} \end{aligned}$$

There are two solutions for $\alpha > 1$ for any ϕ in the range

$$0 > \phi > \theta - \pi/2$$

and α goes to infinity on the boundaries of this region. The minimum value of α in this region occurs when $\sin(2\phi - \theta)$ is maximized, that is at

$$\phi = \frac{\theta - \pi/2}{2}$$

At this point the angle $\phi - \pi/2 = \theta - \phi = \theta/2 + \pi/4 = \beta$, and

$$\alpha^2 = \frac{1 + \sin(\theta)}{1 - \sin(\theta)} = \frac{2\sin^2(\theta/2 + \pi/4)}{2\cos^2(\theta/2 + \pi/4)} = \tan^2(\beta)$$

as observed by Zollikofer.

7 Conclusion

We have shown that given any approximately symmetrizable set P of point-pairs, transforming it to be isomorphic and then fitting a plane to the set of midpoints using least-squares gives us a plane T which minimizes the symmetry error under any symmetry-preserving transformation. Using T to define vectors w_1, w_2, w_3 for P, we can find the single-axis stretch which minimizes the deformation of P.

This is only a cannonical choice of the best restoration of symmetry from an infinite set of possible transformations. If one had two compressed fossils of the same object, assuming they were compressed from two sufficiently different directions, it should be possible to estimate the true original symmetric shape almost perfectly. It would be interesting to find an opportunity to apply this idea in paleontology.

References

[1] C. Crônier, J.-C. Auffray, and P. Courville. A quantitative comparison of the ontogeny of two closely-related upper devonian phacopid trilobites. *Lethia*, 38:123–135, 2005.

[2] M. Kazhdan, B. Chazelle, D. Dobkin, T. Funkhouser, and S. Rusinkiewicz. A reflective symmetry descriptor for 3d models. *Algorithmica*, pages 201–225, October 2003.

[3] J. Kent and K. Mardia. Shape, procrustes tangent projections and bilateral symmetry. *Biometrika*, 88:469–485, 2001.

[4] K. V. Mardia, F. L. Bookstein, and I. J. Moreton. Statistical assessment of bilateral symmetry of shapes. *Biometrika*, 87:285–300, 2000.

[5] N. J. Mitra, L. Guibas, and M. Pauly. Partial and approximate symmetry detection for 3d geometry. *ACM TOG, Proceedings of SIGGRAPH*, 2006. to appear.

[6] N. J. Mitra, L. J. Guibas, and M. Pauly. Symmetrization. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, page 63, New York, NY, USA, 2007. ACM.

[7] R. Motani. New technique for retrodeforming tectonically deformed fossils, with an example for ichthyosaurian specimens. *Lethaia*, 30:221–228, 1997.

[8] D. O'Mara and R. Owens. Measuring bilateral symmetry in digital images. In *Procs. of TENCON '96*, pages 151–156, November 1996.

[9] M. Pauly, N. J. Mitra, J. Wallner, H. Pottmann, and L. J. Guibas. Discovering structural regularity in 3d geometry. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, pages 1–11, New York, NY, USA, 2008. ACM.

[10] J. Podolak, P. Shilane, A. Golovinskiy, S. Rusinkiewicz, and T. Funkhouser. A planar-reflective symmetry transform for 3d shapes. *ACM Trans. Graph.*, 25(3):549–559, 2006.

[11] S. Thrun and B. Wegbreit. Shape from symmetry. In *Proceedings of the International Conference on Computer Vision (ICCV)*, Bejing, China, 2005. IEEE.

[12] C. P. E. Zollikofer and M. S. Ponce de Len. *Virtual Reconstruction: A Primer in Computer-assisted Paleontology and Biomedicine*. New York: Wiley, 2006.

A Making a Point Set Isotropic

For completeness, we describe how to find a matrix taking any point set P into an isotropic set $\tilde{P}$ (this is in no way novel, but if one wanted to implement our ideas it is important to be clear about this). We assume that the center of mass of P is the origin. If not, we pre-process it by subtracting the center of mass t from each point, where

$$t = \frac{1}{2n}\sum_{i=1}^{2n} p_i$$

We first compute the covariance matrix of P:

$$M = \sum_i p_i p_i^t$$

Since this is a symmetric, positive semi-definite matrix, M can be expressed as:

$$M = Q^t \begin{bmatrix} a & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & c \end{bmatrix} Q$$

where Q is a rotation and the eigenvalues a, b, c are all positive, so that $M^{-1/2}$ is equal to:

$$M^{-1/2} = Q^t \begin{bmatrix} 1/\sqrt{a} & 0 & 0 \\ 0 & 1/\sqrt{b} & 0 \\ 0 & 0 & 1/\sqrt{c} \end{bmatrix} Q.$$

Then the isotropic set of point-pairs is $\tilde{P} = M^{-1/2}P$; and indeed we can observe that $M^{-1/2}PP^tM^{-1/2} = I$. This transformation is uniquely defined up to an orthogonal transformation (adding any rotation or reflection still produces an isotropic point set).

B Technical Lemmas

Lemma 4 *Let P be any symmetric set of point-pairs. The cross-covariance matrix C has exactly one negative eigenvalue.*

Proof. When P is symmetric, there exists a vector v such that:

$$\langle p_{2i}, v\rangle = -\langle p_{2i+1}, v\rangle$$

for all i. Thus, for every i, there exists a vector w_i with $\langle w_i, v\rangle = 0$ and a scalar α_i such that:

$$p_{2i} = w_i + \alpha_i v \qquad \text{and} \qquad p_{2i+1} = w_i - \alpha_i v.$$

Computing the cross-covariance matrix of P we get:

$$C = -\sum_i \alpha_i^2 vv^t + \sum_i w_i w_i^t.$$

Thus, v is an eigenvalue of C with negative eigenvalue $-\sum \alpha_i^2$. Furthermore, if w is any vector perpendicular to v, then $w^tCw = \sum\langle w_i, w\rangle^2$ which cannot be negative. Thus, C has only one negative eigenvalue. □

Lemma 5 *Let P be an approximately symmetrizable set of point-pairs. Then T is the plane minimizing the symmetry error of (Equation 1) for the transformed set $V\tilde{P}$, where $V = FR$.*

Proof. We expand Equation 1 giving the reflective error as a function of the linear transformation A and the candidate plane of reflection's unit normal w:

$$\begin{aligned} E(A,w) &= \sum_{i=1}^{n} ||A(p_{2i}) - \mathrm{Refl}_w(A(p_{2i+1})||^2 \\ &= \sum_{i=1}^{n} ||V(\tilde{p}_{2i}) - \mathrm{Refl}_w(V(\tilde{p}_{2i+1})||^2 \\ &= \sum_{i=1}^{n} ||V(\tilde{p}_{2i} - \tilde{p}_{2i+1}) + 2\langle V(\tilde{p}_{2i+1}), w\rangle w||^2 \\ &= \sum_{i=1}^{n} \|V(\tilde{p}_{2i} - \tilde{p}_{2i+1})\|^2 + 4\langle V(\tilde{p}_{2i+1}), w\rangle^2 \\ &\qquad + 4\langle V(\tilde{p}_{2i} - \tilde{p}_{2i+1}), w\rangle\langle V(\tilde{p}_{2i+1}), w\rangle \\ &= \sum_{i=1}^{n} \|V(\tilde{p}_{2i} - \tilde{p}_{2i+1})\|^2 + 2w^tVCV^tw. \end{aligned}$$

Thus, the plane minimizing the symmetry error is the plane whose normal is the eigenvector of VCV^t with smallest corresponding eigenvalue. We have $F^t(1,0,0) = (a,0,0)$, and since R^t maps the vector $v = (1,0,0)$ to the only eigenvector of C with negative eigenvalue, it follows that $(1,0,0)$ is an eigenvector of VCV^t with negative eigenvalue. Additionally, VCV^t has only one eigenvector with negative eigenvalue. Thus $(1,0,0)$ must be the vector minimizing the symmetry error and T is the plane of reflection minimizing the error. □

Lemma 6 *Given two distinct lines spanned by the vectors X and Y, any single-axis stretch that aligns the two lines must do so by stretching along a vector contained in the plane $V = Span\{X, Y\}$.*

Proof. Let w be a unit vector and α be a scalar defining the stretching transformation, so that:

$$\lambda X - Y = (\alpha - 1)\langle Y, w\rangle w$$

for some λ. Since the left-hand side is contained in V and since w is not, the only way for equality to hold is if $0 = (\alpha - 1)\langle Y, w\rangle$. But this must imply that X is just a scalar multiply of Y, contradicting the initial assumption that the two lines are distinct. □