Finding Mirror Symmetry via Registration and Optimal Symmetric Pairwise Assignment of Curves

Marcelo Cicconet
Harvard Medical School
Boston, MA, USA
cicconet@gmail.com

David G. C. Hildebrand
Harvard Medical School
Boston, MA, USA
david@hildebrand.name

Hunter Elliott
Harvard Medical School
Boston, MA, USA
elliott.hunter@gmail.com

Abstract

We demonstrate that the problem of fitting a plane of mirror symmetry to data in any Euclidian space can be reduced to the problem of registering two datasets. The exactness of the resulting solution depends entirely on the registration accuracy. This new Mirror Symmetry via Registration (MSR) framework involves (1) data reflection with respect to an arbitrary plane, (2) registration of original and reflected datasets, and (3) calculation of the eigenvector of eigenvalue -1 for the transformation matrix representing the reflection and registration mappings. To support MSR, we also introduce a novel 2D registration method based on random sample consensus of an ensemble of normalized cross-correlation matches. With this as its registration back-end, MSR achieves state-of-the-art performance for symmetry line detection in two independent 2D testing databases. We further demonstrate the generality of MSR by testing it on a database of 3D shapes with an iterative closest point registration back-end. We finally explore its applicability to examining symmetry in natural systems by assessing the degree of symmetry present in myelinated axon reconstructions from a larval zebrafish. Using the MSR-computed plane of symmetry, we introduce techniques for the optimal symmetric pairwise assignment between axon reconstructions and provide visualizations illustrating how neighborhood relationships between nearby axon pairs compare with the relationships between their mirror-reflected counterparts along the anteroposterior axis.

1. Introduction

Symmetry is frequently found in nature and man-made designs, and the human visual system exploits this fact to facilitate object recognition [38]. Similarly, computational tools can take advantage of symmetry for simplified data representation, since it implies a great degree of information redundancy [11]. Therefore, symmetry detection has great potential utility in practical computer vision applications including object recognition and image compression.

In natural images, any present symmetric objects are often surrounded by clutter or partially occluded. This makes symmetry detection challenging, forcing methods to be robust to outliers. Perhaps as a consequence, the symmetry detection approaches that currently perform best ([22], [7]) are partially or entirely based on sampling or voting schemes. However, despite increased resiliency through such schemes, the current state-of-the-art methods still leave substantial room for improvement.

In this work, we show that the problem of finding mirror symmetry—also known as reflection symmetry or bilateral symmetry—in $\mathbb{R}^n$ can be reduced to a registration problem using a new method that we refer to as Mirror Symmetry via Registration (MSR). This is accomplished by computing the eigenvector of eigenvalue -1 for a transformation matrix computed from reflection and registration mappings. We provide straightforward theoretical deductions to support this claim and demonstrate its utility through examples.

To enhance symmetry detection with MSR, we also present a registration algorithm of the random sample consensus (RANSAC) class for two-dimensional (2D) images. This algorithm infers optimal parameters from a collection of patch-to-image registrations computed via Normalized Cross-Correlation (NCC) [16]. By combining these approaches, we achieve state-of-the-art performance for 2D symmetry line and segment detection on two independent testing databases: the CVPR 2013 Symmetry Detection from Real World Images competition [19] and the NYU Symmetry database [7].

To highlight the MSR procedure's generality, we also applied it to symmetric three-dimensional (3D) objects from the McGill 3D Shape Benchmark [34]. These tests achieved 86% accuracy when using an Iterative Closest Point (ICP) algorithm for the underlying registration [6, 4].

Finally, we recognize that symmetry detection algorithms can be helpful for analyzing natural systems. Toward this goal, we apply MSR to start examining questions

2473-9944/17 $31.00 © 2017 IEEE
DOI 10.1109/ICCVW.2017.206

in the field of neuroscience that are concerned with the degree of symmetry in neuronal morphology and patterning. We develop techniques to optimally assign pairs of symmetric axons (represented as 3D curves), and to visualize local symmetries and relationships between pairs of axons in slices perpendicular to the symmetry plane.

The approach we present here exposes symmetry detection to a new line of attack and introduces registration-technique developers to a new type of data on which to test novel methods. The paper is organized as follows: Section 2 reviews existing literature to put MSR in perspective. Section 3 contains mathematical preliminaries and a description of MSR. Section 4 details quantitative experiments on 2D and 3D testing databases. Section 5 analyzes mirror symmetry between pairs of curves in 3D. Finally, section 6 discusses MSR advantages, limitations, and possible future directions.

2. Previous Work

Listed here is a review of some relevant previous work, though necessarily incomplete due to space restrictions. For a more comprehensive review, we refer the reader to [21].

2.1. Mirror Symmetry Detection on 2D Data

1993. [23]: The image is reflected w.r.t. quantized candidate lines, and the correlation between the resulting images and the original image is calculated. **2007.** [20]: Develop a symmetry-based method to identify dihedral and frieze symmetry as well as asymmetric sub-patterns to generate a fold-then-cut plan that can be used to recreate the input papercut pattern and synthesize new papercut patterns. **2006.** [22]: Scale-invariant feature transform (SIFT) features were grouped into "symmetric constellations" by a voting scheme. Dominant symmetries present in the image emerged as local maxima. [9]: Proposes an artificial neural network that extracts axes of symmetry from visual patterns. **2012.** [15]: Generalized mirror symmetry detection to a curved transflection (glide reflection) symmetry detection problem. Estimated symmetry via a set of contiguous local straight reflection axes. **2013.** [14]: Developed a 3-step algorithm, wherein (1) SIFT correlation measures are computed along discrete directions, (2) symmetrical regions are identified from matches in the directions characterized by maximum correlations, and then steps (1) and (2) are repeated at different scales. [31]: Created a 2-step algorithm, wherein candidates for mirror-symmetric patches are identified using a Hough-like voting scheme and then validated using a principled statistical procedure inspired from *a contrario* theory. [24]: Introduced a combinatorial gestalt algebra technique to be used on top of SIFT descriptors. [19]: Evaluated the performance of various symmetry detection methods on a common database, with [22] emerging as overall winner. **2014.** [8]: Described a pairwise voting-scheme based on tangents computed via wavelet filtering. [5]: Presented an adaptive feature point detection algorithm to overcome susceptibility to clutter in feature-based methods. **2015.** [39]: Exhibited use of traditional edge detectors and a voting process, respectively, before and after a novel edge description and matching step based on locally affine-invariant features. **2016.** [7]: Introduced a pairwise convolutional approach to mirror symmetry detection similar to [8]. The method outperformed [22] by a small margin and its authors released a new database, which we use here for testing. [10]: Exploited ambiguities and challenges in symmetry detection to propose a method for producing re-CAPTCHA solutions based on symmetry.

2.2. Mirror Symmetry Detection on 3D Data

1992. [25] Shows that for a body which exhibits planar symmetry its plane of symmetry is perpendicular to a principal axis and contains the object's center of mass. **1997.** [37]: Converted the symmetry detection problem to the correlation of the Gaussian image. **2002.** [3]: Presented an approach similar to MSR in the reflection and registration steps. However, the symmetry plane was fit on the set of midpoints, not obtained as the eigenvalue solution. Mathematical proofs for the results were not provided and tests were only conducted in 3D for symmetry in human faces. **2006.** [32]: Described a planar reflective symmetry transform that captures a continuous measure of a shape's degree of mirror symmetry with respect to all possible planes. [30]: Presented a more robust Gaussian image-based approach. [35]: Uses [25] to "solve for the current plane of maximum symmetry in a closed form manner by considering the center of mass [...] and weighted covariance matrix [...] relative to the weights". **2007.** [26]: Present a symmetrization algorithm for geometric objects, whereby optimal displacement vectors are used to drive a constrained deformation model that pulls the shape towards symmetry. **2011.** [40]: Introduced a 2-step method, wherein landmark-related region detection is followed by a learning stage that computes a symmetry plane from the landmarks based on training input consisting of standard symmetry planes identified by medical experts. [29]: Reviewed [3] and ICP variations. **2013.** [27]: Discussed applications in computer graphics and geometry that can utilize symmetry information for more effective processing. [13]: Described bilateral symmetry plane estimation for 3D shapes that is carried out in the spherical harmonic domain. **2014.** [36]: Presented an algorithm that generates a set of candidate symmetries by matching local maxima of a surface function based on heat diffusion in local domains, with a global optimum obtained by a voting scheme. **2015.** [41]: Developed a skeleton-intrinsic symmetrization method for recovering the aesthetics of mirror symmetry from asymmetric shapes while preserving their general pose. This was accomplished by mea-

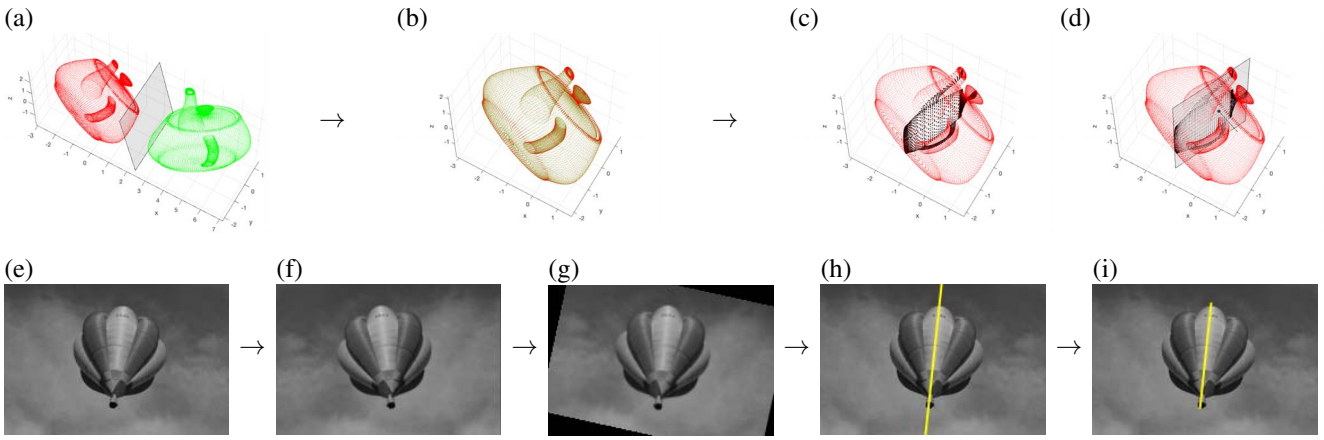

Figure 1. **3D Teapot Point Clouds:** Symmetry plane detection in 3D using MSR. (a) The original data (red) is reflected (green) with respect to an arbitrary plane. (b) Registration between original and reflected point clouds using an Iterative Closest Point (ICP) algorithm. (c) Visualization of every midpoint (black) between a point in the original set and the corresponding point in the reflected set. (d) Computing the symmetry plane in this case involves either fitting a plane to the midpoints in (c) or analytically solving an eigenvalue problem on a function of the transformation matrices corresponding to the reflection and ICP registration mappings. **2D Baloon Images:** Symmetry line and segment detection in 2D using MSR. (e) Input. (f) Mirror reflection of the original image with respect to a vertical line through the center of the image. (g) Registration of (b) with respect to (a). (h) The symmetry line computed by the proposed algorithm (MSR) is shown in yellow. (i) The symmetry segment can be computed from (d) via a post-processing phase, as in [7].

suring intrinsic distances over a curve skeleton backbone for symmetry analysis, symmetrizing about the skeleton, and propagating the symmetrization from skeleton to shape. [2]: Propose a scale invariant structure feature which describes points on extremum curvature along edges for detecting visually salient, structure based symmetry patterns. **2016.** [17]: Achieved symmetry plane detection by generating a candidate plane based on a matching pair of sample views and then verifying whether the number of remaining matching pairs fell within a preset minimum number.

3. Method

Definition 1 (of Mirror Symmetry)**.** *A set of points $P \subset \mathbb{R}^n$ is said to present* mirror, reflection*, or* bilateral *symmetry if there exists a hyperplane $H \subset \mathbb{R}^n$ of dimension $n-1$ such that the* mirror reflection *of P with respect to H produces a set of points Q such that $P = Q$.*

Definition 2 (of Mirror Reflection)**.** *Let $H \subset \mathbb{R}^n$ be a $(n-1)$-dimensional hyperplane, v a unit vector perpendicular to H, and p a fixed point in H, so that $H = \{q \in \mathbb{R}^n : \langle q-p, v\rangle = 0\}$. The* mirror reflection *of a set of points P with respect to H is the set $\{q - 2\langle q-p, v\rangle v : q \in P\}$.*

Notice that mirror symmetry is a *property* of a set of points present (as in "the set *is* mirror symmetric"), whereas mirror reflection is a mathematical transform (e.g. "we took the mirror symmetric of the set with respect to an arbitrary plane").

The mirror reflection of a point x with respect to a plane through the origin and with normal vector v is given by $x \mapsto S_v x$, where $S_v = I - 2vv^\top$, where I is the identity matrix. The reflection with respect to a plane through an arbitrary point p and with normal vector v is given by:

$$x \mapsto S_{p,v}(x) = S_v x + 2dv\,, \tag{1}$$

where $d = \langle p, v\rangle$ is the "signed" distance between the plane and the origin. For simplicity of notation, we will henceforth denote $S_{p,v}(x)$ as $S_{p,v}x$.

The symmetry plane in $\mathbb{R}^n$ can be computed in 3 steps, as illustrated in Figure 1:

1. Reflect original data with respect to an arbitrary plane.
2. Register original and reflected sets.
3. Infer optimal symmetry plane from the parameters of the reflection and registration mappings.

Remarks:

- Depending on the registration algorithm used, it can help to start in MSR step 1 with an arbitrary plane that is near—or a good guess for—the actual symmetry plane. We employ this strategy for the application described in Section 5. Alternatively, several runs with different initial planes can be attempted and the one for which the registration algorithm returns the most confident result chosen. We use this second strategy for the 3D experiments in Section 4. On 2D experiments, the initial reflection is with respect to the vertical line passing through the center of the image.

- All steps in the MSR framework are exact (when factoring out numerical errors) except for registration. If the data is not perfectly mirror symmetric, then the registration will not be precise. The MSR approach reduces mirror symmetry detection to a registration problem, with the caveat that its robustness depends entirely on the robustness of the underlying registration method.

- MSR step 3 can be performed in one of two ways: either by fitting a plane through the midpoints of all corresponding original-transformed point pairs, or by solving an eigenvalue problem related to the global (reflection and rigid) transformation that was applied to the original data during registration. We adopt the latter approach here.

We now mathematically demonstrate why the MSR approach works for detecting mirror symmetry.

Let $P = \{p_1, ..., p_N\}$ be a point cloud and $Q = \{q_1, ..., q_N\}$ the reflection of P given by $S_{p,v}$, that is, $q_i = S_{p,v} p_i \, \forall i$.

Proposition 1. *Let $m_i = \frac{1}{2}(p_i + q_i)$, so that $M = \{m_1, ..., m_N\}$ is the set of midpoints between corresponding points in P and Q. Then the set M is contained in the plane with normal vector v passing through dv.*

Proof. For $x \in P$, the reflection by $S_{p,v}$ is $S_v x + 2dv$, so

$$\langle \frac{1}{2}(x + S_v x + 2dv) - dv, v \rangle = \tag{2}$$

$$\langle \frac{1}{2} x, v \rangle + \langle \frac{1}{2} S_v x, v \rangle + \langle dv, v \rangle - \langle dv, v \rangle. \tag{3}$$

But S_v is symmetric, so $\langle S_v x, v \rangle = \langle x, S_v v \rangle$. Further, $S_v v = -v$ because S_v is the reflection with respect to the plane through the origin with normal vector v, so $\langle x, S_v v \rangle = -\langle x, v \rangle$. Therefore, (3) is equal to 0. □

Let now R be the rigid transformation defined by $R(x) = R_0 x + t$, where R_0 is a rotation matrix and t a translation vector. If we reflect a point $x \in P$ by $S_{p,v}$ and then transform it through R, the result is $R_0(S_v x + 2dv) + t$.

Proposition 2. *Let $T = S_v R_0^\top$ and w equal the unit eigenvector of T corresponding to the eigenvalue -1. That is, $Tw = -w$. We will show in* Proposition 3 *that such a w exists. Let $r = \frac{1}{2}(R_0(2dv) + t)$, with d as previously defined. Then the midpoints $\frac{1}{2}(x + R_0(S_v x + 2dv) + t)$ lie in the plane with normal vector w passing through r.*

Proof.

$$\langle \frac{1}{2}(x + R_0(S_v x + 2dv) + t) - r, w \rangle = \tag{4}$$

$$\langle \frac{1}{2}(x + R_0(S_v x + 2dv) + t) - \frac{1}{2}(R_0(2dv) + t), w \rangle = \tag{5}$$

$$\frac{1}{2}\langle x + R_0(S_v x), w \rangle = \tag{6}$$

$$\frac{1}{2}(\langle x, w \rangle + \langle R_0 S_v x, w \rangle) = \tag{7}$$

$$\frac{1}{2}(\langle x, w \rangle + \langle x, S_v R_0^\top w \rangle) = \tag{8}$$

$$\frac{1}{2}(\langle x, w \rangle + \langle x, -w \rangle) = \tag{9}$$

$$0. \tag{10}$$

□

Proposition 3. *If S is a reflection and R is a rotation, then SR is a reflection. As a consequence, SR necessarily has -1 as an eigenvalue.*

Proof. Since S and R are orthogonal, so is SR (this follows immediately from the definition of orthogonality and the fact that $(SR)^\top = R^\top S^\top$). Further, since the determinant of the product equals the product of the determinants, the determinant of SR is -1. Considering the reflection SR, let H be the reflection hyperplane with normal vector v. In this case, $SRv = -v$, so v is the eigenvector of SR corresponding to the eigenvalue -1. □

3.1. Symmetry Detection Pipeline

Given these results, the precise MSR pipeline for finding the mirror symmetry plane for a set of points P is:

1. Choose an initial reflection plane, given by a point p and a perpendicular vector v. For example, p can be the average (center of mass) of the points in P and v set as the vector $(1, 0, ..., 0)$.

2. Reflect all the points $x \in P$:

$$d = \langle p, v \rangle \, , \tag{11}$$

$$x \mapsto S_{p,v} x = S_v x + 2dv \, , \tag{12}$$

   to obtain a new set of points Q.

3. Register Q to P through a rigid transformation, obtaining a rotation matrix R_0 and a translation vector t. That is, the registration transformation is given by $x \mapsto R_0 x + t$.

4. Compute the eigenvector $\bar{v}$ of the matrix $S_v R_0^\top$ corresponding to the eigenvalue -1, where $\bar{v}$ is the vector perpendicular to the symmetry plane.

5. Completely define the symmetry plane by computing a point $\bar{p}$ through which it passes:

$$\bar{p} = \frac{1}{2}(R_0(2dv) + t) \, . \tag{13}$$

3.2. Consensus of Patch-to-Image Registrations

We initially tested MSR with a number of off-the-shelf registration algorithms in our 2D experiments: ICP on edge maps, intensity-based on images or edge maps, and speeded up robust feature (SURF)-based on images or edge maps. However, none of these options produced results with comparable precision/recall numbers obtained using the previous state-of-the-art symmetry detection algorithms.

Therefore, we designed a new registration method based on a consensus over an ensemble of patch-to-image registration outputs (i.e., a RANSAC approach). First, we assume that the transformation is rigid (rotation and/or translation), which is sufficient for MSR. Then, for every angle $\alpha = 0, \frac{360}{N}, 2\frac{360}{N}, ..., (N-1)\frac{360}{N}$, we sample hundreds of square patches from the moving image and register each with respect to the target image using NCC [16], selecting only those registrations for which the maximum in correlation space is above an empirically chosen threshold ($\frac{1}{4}$). For this process, we found that $N = 60$ typically provides good results. Finally, we look for the K best local maxima in the space of registration parameters found via NCC. For precision/recall evaluations, we chose $K = 10$.

This NCC-based registration approach performed better in the MSR framework than the others we tested. In a one-shot[1] symmetry line detection experiment on the NYU Symmetry database (176 images), our NCC-based registration achieved 95% accuracy, while other methods accomplished accuracies near 73% (Figure 3 (a)).

4. Quantitative Experiments

4.1. Accuracy Metric

For 2D cases, we examined MSR accuracy using established metrics. When detecting symmetry segments, the metric described in [19] was used. When detecting symmetry lines, an extension of the metric for segments [7] was used. In brief, the correctness criteria for segments was based on both angle and center proximity between the prediction result and the ground truth. The correctness criteria for lines was similar, except that center proximity was replaced with the distance between the center of the ground truth and the prediction line because the prediction line has no defined center. For thorough evaluation across approaches, precision/recall curves were generated as in [7] for each method from up to the top ten results.

[1]Only the first guess for symmetry line was used for evaluation.

For 3D cases, we evaluated MSR accuracy by visual inspection of projections of the data along three mutually perpendicular directions, one of which was orthogonal to the estimated symmetry plane.

4.2. Results

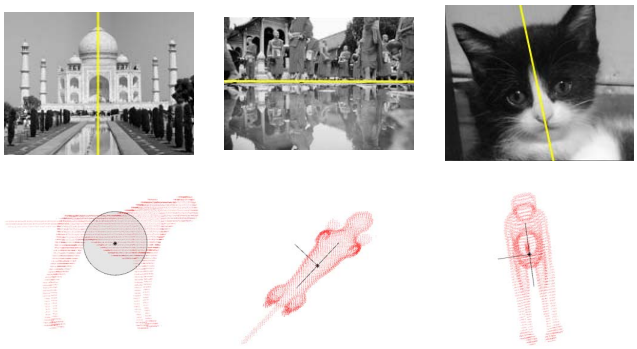

Figure 2. **Images:** Sample of MSR results for 2D images from the NYU Symmetry database. **Point Clouds:** Sample of MSR results for 3D shapes from the McGill 3D Shape Benchmark [34]. The columns show mutually perpendicular views. The left-most view is orthogonal to the MSR-computed symmetry plane.

2D. The previous state-of-the-art for single-symmetry line or segment detection in 2D was a pairwise convolutional method [7], referred to here as Convolutional Approach to Reflection Symmetry (CARS). Released with its description was the database used for testing (the NYU Symmetry database), with which we tested the MSR approach. However, given that CARS does not appear to be peer-reviewed yet, we additionally conducted testing with the CVPR 2013 database, for which Loy's method [22] reported best results.

In accordance with the registration method comparison depicted in Figure 3 (a), we adopted the NCC-based registration described in Subsection 3.2 to compute precision/recall curves for evaluation. Though the MSR method outputs only *lines* for 2D cases, not *segments* (limited subsets of lines), the latter is required for proper comparison with CARS. For this reason, we post-processed the symmetry line resulting from MSR into segments using a previously reported algorithm [7]. Evaluation results are shown in Figure 3 (b,c) and are accompanied by examples of symmetry line detection outputs in Figure 2.

For 2D symmetry detection, our MSR approach outperforms the previous state-of-the-art peer-reviewed method [22] for single-symmetry *segment* detection on the CVPR 2013 database, while reaching similar performance as CARS (pre-print [7]). Additionally, MSR outperforms CARS for single-symmetry *line* detection on this database. Note that line detection results for [22] were not reported and are therefore not available for comparison. MSR also

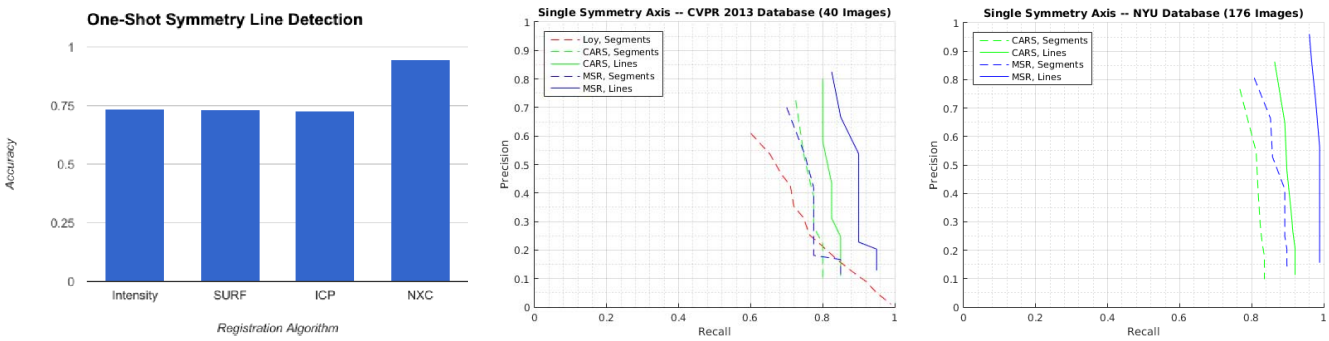

(a) (b) (c)

Figure 3. (Left) Accuracy of one-shot symmetry line detection achieved by combining the MSR approach described in Section 3 with various registration algorithms. Intensity- and SURF-based registrations were performed on the gradients of the images rather than images themselves, as we could not detect a significant difference between the two. ICP was applied on top of Ultrametric Contour Maps [1]. NCC indicates results from using the NCC-based registration described in Subsection 3.2. For this experiment, the NCC-based registration was applied on image gradients for 40×40 patches with a maximum side length of 200. (Middle, Right) Single-symmetry detection on the CVPR 2013 and NYU Symmetry databases. Loy refers to [22], which performed best in the CVPR 2013 competition. CARS refers to [7], which yielded the previous state-of-the-art. The method we describe here is referred to as MSR (Mirror Symmetry via Registration). Comparison with all existing algorithms would be impractical, so we are comparing only with the state-of-the-art on each dataset.

outperforms CARS on both segment and line detection on the NYU Symmetry database.

3D. To the best of our knowledge, there are no general-purpose databases or accuracy metrics for 3D mirror symmetry detection tests. We thus created a testing database consisting of hand-picked 203 symmetric 3D shapes from the McGill 3D Shape Benchmark [34]. The included shapes consist of surface points corresponding to objects such as cups, airplanes, and insects. Because each shape was represented by a set of points, we chose the ICP algorithm [6, 4] as the registration back-end for all 3D testing. We should point out that due to the lack of a common dataset for comparison, we decided not to test all possible registration methods as back-end to MSR. ICP has some well known shortcomings (e.g. sensitivity to local minima and initialization), but we adopted it because it was the most conveniently available in Matlab.

We ran the MSR method three times per object, each with a different initial reflection hyperplane. These hyperplanes were always selected as passing through the object's center of mass and with perpendicular vectors given by the canonical basis $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$. As the final solution, we chose the result whose registration confidence was highest.

By visual inspection of projections along the three mutually perpendicular vectors, one of which was always orthogonal to the symmetry plane, we found that MSR achieved 87% accuracy. (From the set of 203 shapes, symmetry was correctly detected in 177.) One example is shown in Figure 2.

5. Symmetry of Curves in 3D

This project was largely driven by a practical application in the field of neuroscience. The presence of bilateral symmetry in the morphologies of neurons is an indicator that specialized genetic programs, rather than experience and neuronal activity, may be responsible for the way that they develop. Previous work in larval zebrafish examined symmetry in neuronal circuitry locally with the goal of understanding the role that left-right asymmetries play in lateralized behaviors (such as an observed bias in turn direction during swimming) [reviewed in [33]]. We sought to examine symmetry more globally throughout the entire brain of a larval zebrafish by analyzing the precise shapes and positions of myelinated axon projections. We further aimed to determine whether the spatial relationships between the projections of neurons on one side are also present for neurons in the contralateral hemisphere. Maintenance of such relationships at a fine scale would suggest that the developmental programs for each side are hard-coded and should provide insights into the strategies employed by neurons to correctly reach their downstream. Thus, we developed methods to analyze the degree of bilateral symmetry in myelinated axons projections reconstructed from a whole larval zebrafish brain [12].

Myelinated axon reconstructions were manually extracted from serial-section electron micrographs. The resulting data consisted of curves represented as sequences of points in 3D, which we refer to as *skeletons*. We first sought to find the plane of bilateral symmetry given that the projections appeared nearly mirror symmetric. A visually acceptable result was obtained by application of the MSR

approach with the ICP algorithm as the registration back-end and a manually selected initial reflection plane corresponding to the vector given by the canonical basis $(1, 0, 0)$.

Figure 4 (a) illustrates the myelinated axon reconstructions and the MSR-computed symmetry plane using ICP as the registration back-end. Figure 4 (b) shows projections of the data along three mutually perpendicular directions, where the side projection is orthogonal to the symmetry plane.

We next sought to find the optimally symmetric left-right pairing of skeletons. To do so, we first defined a metric of similarity between skeletons.

5.1. Similarity Between Curves

A skeleton s is a discrete curve in $\mathbb{R}^n$:

$$s = \{s_i : i = 1, ..., n_s\} . \tag{14}$$

Given two skeletons s and t, their similarity can be computed via Dynamic Time Warping (DTW), a variation of Dynamic Programming that is widely used for sequence matching [18].

DTW admits a parameter for the cost of matching a point in one sequence with a gap in another. We set this gap cost to 0 (zero), since our data is sampled at a nearly constant rate and we want to find the optimal subsequence match in case one sequence is shorter than or offset with respect to the other. For optimal skeleton pairing after the hyperplane of symmetry is found, however, we add to the default DTW cost a penalty proportional to the portions of the two sequences that remained unmatched.

Let l_s^m and l_t^m be the lengths of the matched portions of the sequences s and t, respectively. Let l_s and l_t be the total lengths of s and t. We define penalties for unmatched points as

$$c_s = \frac{l_s}{l_s^m} \text{ and } c_t = \frac{l_t}{l_t^m} . \tag{15}$$

Given the default DTW matching cost $c_0(s, t)$, between s and t, the penalized cost $c(s, t)$ is given by

$$c(s, t) = c_0(s, t) \cdot c_s \cdot c_t . \tag{16}$$

Notice that the larger the unmatched portions, the larger the factors c_s and c_t, and therefore the larger $c(s, t)$.

5.2. Optimal Pairwise Assignment

Given a reference symmetry plane H, a pairwise symmetry measure is computed by comparing one skeleton with the reflection of the other with respect to H (Figure 4).

Given a matrix of pairwise costs C, where $C(i, j) = C(j, i)$ is the symmetry measure between skeletons of indexes i and j, we apply the Munkres assignment algorithm [28] (also known as the Hungarian method) to compute the globally optimal pairwise assignment between skeletons.

5.3. Slice Visualization

Besides analysis on a skeleton level, we studied the neighbor relationship between pairs of skeletons (in a subset of skeletons) at a z-slice level. We were interested in the symmetry they display in terms of relative displacements, as well as in visualizing more closely how the skeletons are arranged as z varies. Figure 5 shows the output of this analysis for a particular z-slice.

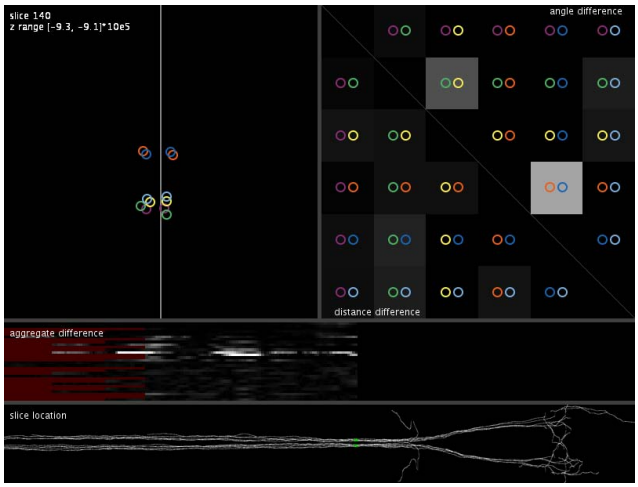

Figure 5. One frame of a visualization movie for a slice of data – i.e., a portion of data in a particular z-interval – for a subset of 12 skeletons. Top-left panel: skeleton points (in coordinates related to the computed symmetry plane) that fall on a given z range have their x and y coordinates averaged and plotted as circles, where identical colors correspond to pairings given by the Munkres algorithm. the vertical line is the projection of the found symmetry plane. Top-right panel: angle- and distance-based relative differences between pairs of skeletons from "one side" of the symmetry plane. Second panel from bottom: linearized version of the difference matrix; every column aggregates the difference matrix of a particular slice and the horizontal dimension relates to the z coordinate. Bottom panel: "geographical" location of the slice (shown in green) with respect to the subset of skeletons with the z coordinate is represented horizontally.

First the coordinates of the points are altered so that the symmetry plane is the plane $\{x = 0\}$, as illustrated in Figure 4 (d). Given z_0, z_1, and a skeleton S containing points $s = (s_x, s_y, s_z)$, let $S_{z_0,z_1} = \{s \in S : z_0 \leq s_z < z_1\}$. That is, S_{z_0,z_1} is the subset of points of S such that the coordinates s_z are in the interval $[z_0, z_1)$. We refer to the subset of $\mathbb{R}^3$ such that $z \in [z_0, z_1)$ as the *slice* $[z_0, z_1)$, as in Figure 4 (d).

For each slice $[z_0, z_1)$ and skeleton S, we define $\langle S_{z_0,z_1} \rangle$ as the mean of the elements in S_{z_0,z_1}. This is the representative of the skeleton S in slice $[z_0, z_1)$ and is used for plotting (top-left panel in Figure 5) and analysis (top-right panel).

Let now $s_1, ..., s_n$ be a set of representative points in a

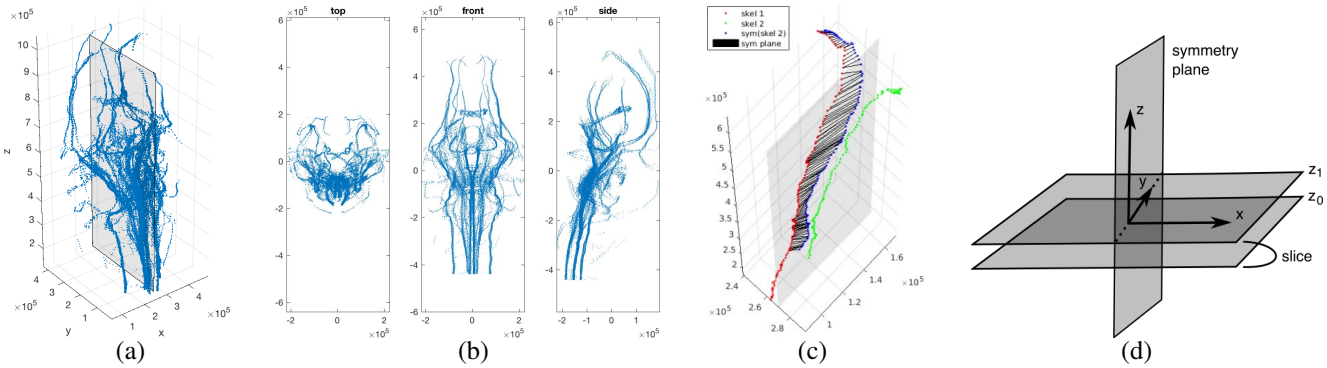

Figure 4. (a) Symmetry plane computed using the MSR approach for myelinated axon reconstructions (blue) obtained from a larval zebrafish using serial-section electron microscopy. (b) Projections of myelinated axon reconstructions along three mutually perpendicular directions. The side projection is orthogonal to the MSR-computed symmetry plane. (c) Pairwise skeleton symmetry. Skeleton 2 is reflected with respect to the symmetry plane and the result matched to skeleton 1 via DTW. The DTW matching cost is the measure of symmetry between the two skeletons. (d) For within-slice pairwise symmetry analysis, point coordinates were altered so that the symmetry plane is the plane $\{x = 0\}$. A slice $[z_0, z_1)$ is the subset of $\mathbb{R}^3$ such that $z \in [z_0, z_1)$.

fixed slice for skeletons $S_1, ..., S_n$ and $t_1, ..., t_n$ the representative points (for the same slice) of the respective skeletons $T_1, ..., T_n$ that were paired to $S_1, ..., S_n$ by the Munkres algorithm.

We visualized the local (at the slice level) symmetry between the sets s_i and t_i (for $i = 1, ..., n$) by plotting them across different slices. Furthermore, we devised a measure of relative displacements between the two sets as follows.

Let $\{t_i^s\}$ be the reflections of $\{t_i\}$ with respect to the computed plane of symmetry. Given two indexes i and j, the *angle difference* between pairs s_i, s_j and t_i^s, t_j^s is defined as

$$a_{i,j} = \frac{1}{2}\left(1 - \frac{\langle s_j - s_i, t_j^s - t_i^s\rangle}{\|s_j - s_i\|\|t_j^s - t_i^s\|}\right), \qquad (17)$$

and the *distance difference* as

$$d_{i,j} = \frac{|\|s_j - s_i\| - \|t_j^s - t_i^s\||}{M}, \qquad (18)$$

where M is the maximum of $|\|s_j - s_i\| - \|t_j^s - t_i^s\||$ across all pairs i, j and stacks. Notice that $a_{i,j}$ and $d_{i,j}$ vary from 0 (no difference) to 1 (maximum difference) and that if points s_i and s_j are perfectly symmetric with respect to points t_i and t_j, then $a_{i,j} = 0$ and $d_{i,j} = 0$.

A difference matrix D can then be defined for each slice by setting $D(i, j) = a_{i,j}$ if $j > i$ and $D(i, j) = d_{j,i}$ if $j < i$. An example is shown at the top-right panel of Figure 5. Notice that the angle difference for the blue and orange pairs is high because the angle between the left blue-orange line (connecting the blue and orange points on the left) and the reflected right blue-orange line is large.

The difference matrices can be vectorized (linearized) and plotted vertically for every slice in order to highlight slices for which the representative points most deviate from a symmetric displacement. This is shown in the panel labeled "aggregate difference" in Figure 5.

6. Conclusion

In this paper, we introduced Mirror Symmetry via Registration (MSR), a new framework for mirror symmetry detection that is based on registration and invariant to dimension. For all but the registration phase, this approach is mathematically exact. That is, mirror symmetry detection in $\mathbb{R}^n$ is as good as the best available registration method. In addition, we described a new 2D image registration algorithm based on RANSAC over a set of patch-to-image registrations.

To illustrate MSR performance, we provided experimental results from testing on 2D and 3D databases. To show the utility of MSR in analyses of natural systems, we described its application to 3D symmetry detection in the myelinated axons of a larval zebrafish. We further analyzed symmetry in zebrafish axons by introducing techniques for the optimal symmetric pairwise assignment between axons, and to visualize how the relationship between pairs of axons and their symmetrics varies across the anteroposterior axis. For more details on our biological findings, we refer the reader to [12].

One limitation of MSR is that it does not output the intersection of the computed symmetry hyperplane with the symmetric object. In 2D, for example, it only outputs the symmetry line, not the symmetry segment.

Potential improvements to MSR include its extension to enable detection of multiple symmetry axes and, on the theoretical side, a metric for quantifying plane similarity in $\mathbb{R}^n$ for $n > 2$ should be developed to better measure accuracy.

References

[1] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5):898–916, May 2011. 6

[2] I. Atadjanov and S. Lee. Bilateral symmetry detection based on scale invariant structure feature. In *Image Processing (ICIP), 2015 IEEE International Conference on*, pages 3447–3451. IEEE, 2015. 3

[3] M. Benz, X. Laboureux, T. Maier, E. Nkenke, S. Seeger, F. W. Neukam, and G. Häusler. The symmetry of faces. In *Proceedings of the Vision, Modeling, and Visualization Conference 2002 (VMV 2002), Erlangen, Germany, November 20-22, 2002*, pages 43–50, 2002. 2

[4] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(2):239–256, 1992. 1, 6

[5] D. Cai, P. Li, F. Su, and Z. Zhao. An adaptive symmetry detection algorithm based on local features. In *Visual Communications and Image Processing Conference, 2014 IEEE*, pages 478–481, Dec 2014. 2

[6] Y. Chen and G. G. Medioni. Object modelling by registration of multiple range images. *Image Vision Comput.*, 10(3):145–155, 1992. 1, 6

[7] M. Cicconet, V. Birodkar, M. Lund, M. Werman, and D. Geiger. A convolutional approach to reflection symmetry. *Pattern Recognition Letters*, 95:44 – 50, 2017. 1, 2, 3, 5, 6

[8] M. Cicconet, K. Gunsalus, D. Geiger, and M. Werman. Mirror symmetry histograms for capturing geometric properties in images. CVPR, 2014. Columbus, Ohio. 2

[9] K. Fukushima and M. Kikuchi. Symmetry axis extraction by a neural network. *Neurocomputing*, 69(16):1827–1836, 2006. 2

[10] C. Funk and Y. Liu. Symmetry recaptcha. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2

[11] R. Gens and P. M. Domingos. Deep symmetry networks. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2537–2545. Curran Associates, Inc., 2014. 1

[12] D. G. C. Hildebrand, M. Cicconet, R. M. Torres, W. Choi, T. M. Quan, J. Moon, A. W. Wetzel, A. S. Champion, B. J. Graham, O. Randlett, et al. Whole-brain serial-section electron microscopy in larval zebrafish. *Nature*, 545(7654):345–349, 2017 May 18. 6, 8

[13] R. Kakarala, P. Kaliamoorthi, and V. Premachandran. Three-dimensional bilateral symmetry plane estimation in the phase domain. In *2013 IEEE Conference on Computer Vision and Pattern Recognition, Portland, OR, USA, June 23-28, 2013*, pages 249–256, 2013. 2

[14] S. Kondra, A. Petrosino, and S. Iodice. Multi-scale kernel operators for reflection and rotation symmetry: Further achievements. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on*, pages 217–222, June 2013. 2

[15] S. Lee and Y. Liu. Curved glide-reflection symmetry detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(2):266–278, Feb 2012. 2

[16] J. Lewis. Fast template matching. In *Canadian Image Processing and Pattern Recognition Society, Quebec City, Canada*, pages 120–123, 1995. 1, 5

[17] B. Li, H. Johan, Y. Ye, and Y. Lu. Efficient 3d reflection symmetry detection: A view-based approach. *Graphical Models*, 83:2–14, 2016. 3

[18] S. Z. Li and A. Jain, editors. *Dynamic Time Warping (DTW)*, pages 231–231. Springer US, Boston, MA, 2009. 7

[19] J. Liu, G. Slota, G. Zheng, Z. Wu, M. Park, S. Lee, I. Rauschert, and Y. Liu. Symmetry detection from real world images competition 2013: Summary and results. CVPR Workshop, 2013. Portland, Oregon. 1, 2, 5

[20] Y. Liu, J. Hays, Y.-Q. Xu, and H.-Y. Shum. Digital papercutting. In *ACM SIGGRAPH 2005 Sketches*, page 99. ACM, 2005. 2

[21] Y. Liu, H. Hel-Or, C. S. Kaplan, L. Van Gool, et al. Computational symmetry in computer vision and computer graphics. *Foundations and Trends® in Computer Graphics and Vision*, 5(1–2):1–195, 2010. 2

[22] G. Loy and J.-O. Eklundh. Detecting symmetry and symmetric constellations of features. In *Proceedings of the 9th European Conference on Computer Vision - Volume Part II*, ECCV'06, pages 508–521, Berlin, Heidelberg, 2006. Springer-Verlag. 1, 2, 5, 6

[23] T. Masuda, K. Yamamoto, and H. Yamada. Extraction of symmetry properties using correlation with rotated and reflected images. *Electronics and Communications in Japan (Part III: Fundamental Electronic Science)*, 76(1):8–19, 1993. 2

[24] E. Michaelsen, D. Muench, and M. Arens. Recognition of symmetry structure by use of gestalt algebra. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on*, pages 206–210, June 2013. 2

[25] P. MINOVIC, S. ISHIKAWA, and K. KATO. Three-dimensional symmetry identification part i: Theory. *Memoirs of the Kyushu Institute of Technology. Engineering*, 21:1–16, 1992. 2

[26] N. J. Mitra, L. J. Guibas, and M. Pauly. Symmetrization. *ACM Transactions on Graphics (TOG)*, 26(3):63, 2007. 2

[27] N. J. Mitra, M. Pauly, M. Wand, and D. Ceylan. Symmetry in 3d geometry: Extraction and applications. *Comput. Graph. Forum*, 32(6):1–23, 2013. 2

[28] J. Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32-38, 1957. 7

[29] C. Padia and N. Pears. A review and characterization of icp-based symmetry plane localisation in 3d face data. Technical report, Department of Computer Science, University of York, UK, 2011. 2

[30] G. Pan, Y. Wang, Y. Qi, and Z. Wu. Finding symmetry plane of 3d face shape. In *18th International Conference on Pattern Recognition (ICPR 2006), 20-24 August 2006, Hong Kong, China*, pages 1143–1146, 2006. 2

[31] V. Patraucean, R. von Gioi, and M. Ovsjanikov. Detection of mirror-symmetric image patches. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on*, pages 211–216, June 2013. 2

[32] J. Podolak, P. Shilane, A. Golovinskiy, S. Rusinkiewicz, and T. A. Funkhouser. A planar-reflective symmetry transform for 3d shapes. *ACM Trans. Graph.*, 25(3):549–559, 2006. 2

[33] M. Roussigne, P. Blader, and S. W. Wilson. Breaking symmetry: The zebrafish as a model for understanding left-right asymmetry in the developing brain. *Developmental neurobiology*, 72(3):269–281, 2012. 6

[34] K. Siddiqi, J. Zhang, D. Macrini, A. Shokoufandeh, S. Bouix, and S. Dickinson. Retrieving articulated 3-d models using medial surfaces. *Mach. Vision Appl.*, 19(4):261–275, May 2008. 1, 5, 6

[35] P. Simari, E. Kalogerakis, and K. Singh. Folding meshes: Hierarchical mesh segmentation based on planar symmetry. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing*, SGP '06, pages 111–119, Aire-la-Ville, Switzerland, Switzerland, 2006. Eurographics Association. 2

[36] I. Sipiran, R. Gregor, and T. Schreck. Approximate symmetry detection in partial 3d meshes. *Comput. Graph. Forum*, 33(7):131–140, 2014. 2

[37] C. Sun and J. Sherrah. 3d symmetry detection using the extended gaussian image. *IEEE Trans. Pattern Anal. Mach. Intell.*, 19(2):164–168, 1997. 2

[38] T. Vetter, T. Poggio, , and H. Bulthoff. The importance of symmetry and virtual views in three-dimensional object recognition. *Current Biology*, 4(1):18–23, 1994. 1

[39] Z. Wang, Z. Tang, and X. Zhang. Reflection symmetry detection using locally affine invariant edge correspondence. *Image Processing, IEEE Transactions on*, 24(4):1297–1301, April 2015. 2

[40] J. Wu, R. Tse, C. L. Heike, and L. G. Shapiro. Learning to compute the plane of symmetry for human faces. In *Proceedings of the 2Nd ACM Conference on Bioinformatics, Computational Biology and Biomedicine*, BCB '11, pages 471–474, New York, NY, USA, 2011. ACM. 2

[41] Q. Zheng, Z. Hao, H. Huang, K. Xu, H. Zhang, D. Cohen-Or, and B. Chen. Skeleton-intrinsic symmetrization of shapes. *Comput. Graph. Forum*, 34(2):275–286, 2015. 2