

On the foundations of vision modeling. II. Mining of mirror symmetry of 2-D shapes

Jianhong Shen*

Abstract

Vision can be considered as a feature mining problem. Visually meaningful features are often geometrical, e.g., boundaries (or edges), corners, T-junctions, and symmetries. Mirror symmetry or near mirror symmetry is very common and useful in image and vision analysis. The current paper proposes several different approaches to extract the symmetry mirrors of 2-dimensional (2-D) mirror symmetric shapes. Proper mirror symmetry metrics are introduced based on Lebesgue measures, Hausdorff distance, and lower-dimensional feature sets. Theory and computation of these approaches and measures are studied.

1 Introduction and Motivation

Vision, a multi-level process of perception generation from observed 2-D images or image sequences, is first of all, a data mining problem. At the most fundamental level, it is to mine the basic geometric, topological, material, and chromatic (or painting) features from the input (intensity or range) image data. These basic features or patterns are then further encoded, correlated, and re-organized (typically in a tree-code form [5]), leading to the perception of our evolving but stationary, generic, and spatially well ordered 3-D world, which is certainly not merely a collection of molecular or atomic clutters.

The current paper is intended to make contributions to the mining of *mirror symmetry* of 2-D shapes, perhaps the most common and accessible symmetry type in image and vision analysis. Figure 1 demonstrates three familiar examples which are mirror symmetric or close to be.




Figure 1: Examples of (nearly)mirror symmetric shapes.

In manufacturing industries (e.g., automobiles), mirror symmetry offers great convenience in merchandise design and production. In life science, it is so astonishing that genetic DNA codes at the molecular level could lead to such universal (near) mirror symmetry in the global appearance of human beings and most animals. Finally, in the psychology of architecture design, the ancient mystery still remains why mirror symmetry looks so pleasing to human beings (see for example, the Eiffel Tower, Paris, and the Big Dome of MIT, Boston, USA).

In this paper, we shall call any compact set $A \subset R^m$ with positive Lebesgue measure

$$|A| = \int_A 1 \, dx_1 \cdots dx_m > 0,$$

a (Lebesgue) compact positive shape. In image and vision analysis, most shapes are Lipschitz domains (except fractal coastlines). Let l denote any affine hyperplane of co-dimension 1 in R^m , \mathbf{n} one of its two unit normals, and $\mathbf{q} \in l$ a fixed reference point. The reflection R_l about mirror l is defined to be, for any $\mathbf{p} \in R^m$,

$$R_l \mathbf{p} = \mathbf{p} - 2\langle \mathbf{p} - \mathbf{q}, \mathbf{n} \rangle \mathbf{n}.$$

*Jianhong (Jackie) Shen is with the School of Mathematics, University of Minnesota, Minneapolis, MN 55455, USA. Telephone: (612) 625-3570. Email: jhshen@math.umn.edu. URL: <http://www.math.umn.edu/~jhshen>. This project is partially supported by NSF's Program of Applied Mathematics under grant number DMS-0202565.

A shape A is said to be mirror symmetric about l , or l -symmetric, if

$$R_l A = A, \quad a.e.,$$

and then l is called the symmetry mirror of A .

We shall work with $m = 2$ only, as motivated by image and vision analysis. *All* the results are valid for higher dimensions, except where the opposite is explicitly stated.

The current paper strives to make both theoretical and computational contributions to the mining of mirror symmetry. We intend to answer the following two key questions:

- 1) Suppose that a *generic* shape A is mirror symmetric or close to be (as common in image and vision analysis). How to extract the symmetry mirror l in a computationally efficient way, which is also stable and insensitive to the inaccuracy of shape extraction algorithms and various types of perturbations?
- 2) Instead of a *binary* decision like “ A is either mirror symmetric or not,” can we develop meaningful metrics to be able to say that “ A is 0.618 mirror symmetric in a scale of $[0, 1]$?”

The main content is almost explicit from the titles of the five major sections:

- §2. Mirror identification and computation: methods of cut-and-weight, and PCA (principle component analysis).
- §3. Mirror symmetry via lower-dimensional features.
- §4. Metrics of mirror symmetry: the symmetry index, and Hausdorff distance for feature sets.
- §5. Mirror symmetry of *convex* shapes (via support functions).
- §6. Mirror symmetry of C^1 shapes via *normal pairs*.

Finally, we want to emphasize that, in order to clearly understand which mechanisms human observers truly employ, all the models and algorithms developed here should be further validated by psychology experiments in human vision research. For the opposite direction, i.e., mathematical modeling directly motivated by laws from vision psychology and psychophysics, please see the author’s work in [8].

Generous tolerance from our readers on the clumsy English is secretly but deeply appreciated.

2 Mirror Identification and Computation

In this section, we propose two approaches to identify the mirror of a shape which is mirror symmetric or close to be. (§ 4 clarifies such closeness.) The algorithms must be stable, meaning that when a perfectly mirror symmetric shape A is perturbed to A^ϵ , the identified new mirror l_ϵ should undergo a minor change as well. In image and vision analysis, such perturbation is very common since shapes are often approximately segmented and extracted from intensity images by some other numerical algorithms (e.g., such as the implementation of the celebrated Mumford-Shah model [6]).

2.1 Method of cut-and-weight Denote by $B_r(\mathbf{a})$ the unit disk of radius r , centered at \mathbf{a} . Define the mass center $[A] \in \mathbb{R}^2$ of a compact shape A by

$$(2.1) \quad [A] = \frac{1}{|A|} \int_A \mathbf{p} \, dx dy, \quad \mathbf{p} = (x, y).$$

For each $r \geq 0$: $|A - B_r([A])| > 0$, similarly define

$$\mathbf{m}_r = [A - B_r([A])],$$

which represents the mass center of A with a disk hole cut out. It is in this sense that we have named our approach after “cut-and-weight.”

PROPOSITION 2.1. *Let A be a positive compact shape. Then there exists a unique $r_A > 0$, so that*

$$|A - B_r([A])| > 0, \quad r \in [0, r_A); \quad |A - B_{r_A}([A])| = 0.$$

Proof. Define $m(r) = |A - B_r([A])|$. It can be easily established that

- a) $m(0) > 0$ and $m(r) \geq 0$ is non-increasing;
- b) $m(r)$ is a continuous function; and
- c) $m(r) \equiv 0$ when r is large enough.

Then it follows from the three properties that we must have $r_A = \min m^{-1}(0) > 0$. □

Therefore, for any $r < r_A$, $\mathbf{m}_r \in R^2$ is well-defined. \mathbf{m}_r gives a one-parameter continuous curve in R^2 . (Moreover, examples can be constructed in which $\lim_{r \rightarrow r_A^-} \mathbf{m}_r$ does not exist.)

THEOREM 2.1. *If A is mirror symmetric about mirror l : $R_l A = A$. Then, for any $r < r_A$, $\mathbf{m}_r \in l$.*

Proof. It suffices to show that $R_l \mathbf{m}_r = \mathbf{m}_r$. First

$$(2.2) \quad R_l[A] = \frac{1}{|A|} \int_A R_l \mathbf{p} \, dx dy = \frac{1}{|A|} \int_{R_l A} \mathbf{p} \, dx dy = [A].$$

By Euclidean invariance,

$$R_l B_r([A]) = B_r(R_l[A]) = B_r([A]).$$

Therefore,

$$R_l(A - B_r([A])) = A - B_r([A]).$$

In (2.2), having A replaced by $A - B_r([A])$, we obtain

$$R_l \mathbf{m}_r = R_l[A - B_r([A])] = [A - B_r([A])] = \mathbf{m}_r.$$

This completes the proof. □

Notice that for generic shapes, $\mathbf{m}_r (r < r_A)$ is a continuous curve containing more than one points. We shall call it the mass center curve. The above theorem implies when a generic shape A is l -symmetric, the mirror l is completely determined by the mass center line. This nourishes our first algorithm for mirror identification.

In practice, due to both imperfect shape extraction and boundary perturbation, A can only be *nearly* mirror symmetric. As a result, the mass center curve is only *nearly* straight. Thus, instead of the ideal algorithm based on the Euclidean principle “two points determine a straight line,” we turn to the technique of *weighted least square fitting* [9] to extract the mirror l .

Definition. (Generalized Mirror) Let A be a generic compact shape with positive mass $|A|$. A straight line l_0 is said to be its mirror if

$$(2.3) \quad l_0 = \operatorname{argmin}_l \int_0^{r_A} d^2(\mathbf{m}_r, l) m(r) dr = \operatorname{argmin}_l e[l],$$

where $d(\mathbf{p}, l)$ denotes the Euclidean point-line distance, and $m(r)$ is defined as in Proposition 2.1.

This notion of generalized mirrors applies meaningfully to nearly mirror symmetric shapes, meaning that it well approximates human perceptual estimation. Later on in § 4, a more general definition for generic shapes will be introduced.

In this weighted fitting model, the weight $m(r)$ has been motivated by the following consideration: larger mass $m(r) = |A - B_r([A])|$ makes the mass center $\mathbf{m}_r = [A - B_r([A])]$ more reliable, since from the statistical point of view, averaging is taken over a larger pool of samples.

We now discuss how to computationally search for the mirror l_0 of a nearly mirror symmetric compact shape A . Fix any reference point \mathbf{q} , for example, $\mathbf{q} = (0, 0)$. For any straight line l , let \mathbf{p} denote the orthogonal projection of \mathbf{q} on l . Define

$$s = d(\mathbf{q}, l) = |\mathbf{p} - \mathbf{q}|, \quad \mathbf{n} = (\mathbf{p} - \mathbf{q})/s.$$

Notice that generically $s > 0$ and \mathbf{n} is indeed well-defined. Moreover, $l \leftrightarrow (s, \mathbf{n}) \in \mathbb{R}_+^1 \times S^1$ is a complete representation. Thus we shall simply identify l with the pair (s, \mathbf{n}) , and the weighted squared fitting error in (2.3) is now given by

$$e[l] = e[s, \mathbf{n}] = \int_0^{r_A} ((\mathbf{m}_r - \mathbf{q}) \cdot \mathbf{n} - s)^2 m(r) dr.$$

To minimize it, we solve the equilibrium equations in $s \in \mathbb{R}_+^1$ and $\mathbf{n} \in S^1$. First for s ,

$$0 = \frac{\partial e[s, \mathbf{n}]}{\partial s} = -2 \int_0^{r_A} ((\mathbf{m}_r - \mathbf{q}) \cdot \mathbf{n} - s) m(r) dr,$$

which gives

$$s = \frac{\mathbf{n} \cdot \int_0^{r_A} \mathbf{m}_r m(r) dr - \mathbf{n} \cdot \mathbf{q}}{M}, \quad M = \int_0^{r_A} m(r) dr.$$

It shows that in equilibrium, s is completely determined by the normal \mathbf{n} of the mirror. Now for any $\mathbf{n} \in S^1$, define the tangential projection $\pi_{\mathbf{n}}(\mathbf{u}) = \mathbf{u} - (\mathbf{u} \cdot \mathbf{n})\mathbf{n}$. Taking the partial derivative $\partial e[s, \mathbf{n}] / \partial \mathbf{n}$ along the unique circle S^1 (i.e., manifold differentiation) leads to the nonlinear equilibrium equation for \mathbf{n} :

$$0 = \int_0^{r_A} ((\mathbf{m}_r - \mathbf{q}) \cdot \mathbf{n} - s) \pi_{\mathbf{n}}(\mathbf{m}_r - \mathbf{q}) m(r) dr.$$

Generically, these two equations on (s, \mathbf{n}) can be solved by suitable iterative schemes $(s_k, \mathbf{n}_k) \rightarrow (s_{k+1}, \mathbf{n}_{k+1})$ (see, for example, Chan and Shen [2] for similar energy minimization problems involving unit normals).

We must point out that in practice, the continuous integral $\int_0^{r_A}$ is computed by a Monte-Carlo type scheme. Both \mathbf{m}_r and $m(r)$ are randomly sampled and computed at some r_1, r_2, \dots, r_K . The discrete version for the weighted squared error becomes

$$e_d[s, \mathbf{n}] = \sum_{k=1}^K ((\mathbf{m}_{r_k} - \mathbf{q}) \cdot \mathbf{n} - s)^2 m(r_k).$$

Note that even the integration weight $|r_{k+1} - r_k|$ has been dropped since our intrinsic focus is merely on the optimal mirror l . K should be large enough so that $\{\mathbf{m}_{r_1}, \dots, \mathbf{m}_{r_K}\}$ contains at least two distinguishable (according to some threshold) points. Larger K also helps stabilize the solution by averaging out the effects of noise and perturbations.

2.2 Method of principle component analysis (PCA) The second approach for mirror identification is based on PCA or singular value decompositions, and thus carries a more genuine taste of data mining. It is fundamentally stochastic in treating a shape as a collection of random samples from a Gaussian distribution:

$$p(\mathbf{p}) = \frac{1}{(2\pi \det \Sigma)^{n/2}} \exp\left(-\frac{1}{2}(\mathbf{p} - \mathbf{p}_0)^t \Sigma^{-1}(\mathbf{p} - \mathbf{p}_0)\right).$$

Then mirror identification is simply to compute the principle directions of the distribution, i.e., the eigenvectors of the correlation matrix Σ .

Let A be a given compact shape with positive Lebesgue mass $|A|$. The expectation type functional E_A is defined for any locally square integrable function (scalar or vectorial) $f(\mathbf{p}) = f(x, y)$:

$$E_A f = \frac{1}{|A|} \int_A f(\mathbf{p}) dx dy.$$

Mass center $[A] = (x_0, y_0)$ in the previous section is simply $E_A[\mathbf{p}]$. Furthermore, define the auto-correlation matrix to be

$$\Sigma_A = \begin{bmatrix} E_A(x - x_0)^2 & E_A(x - x_0)(y - y_0) \\ E_A(x - x_0)(y - y_0) & E_A(y - y_0)^2 \end{bmatrix}.$$

The foundation of this pseudo-probabilistic approach is solidified by the following theorem.

THEOREM 2.2. *If A is l -symmetric, then the mirror's normal \mathbf{n} must be the eigenvector of Σ_A .*

Proof. Notice that the auto-correlation matrix Σ_A is both rotation and translation invariant with respect to A (in the sense of orthonormal similarity). Thus, without loss of generality, assume that $[A] = (x_0, y_0) = (0, 0)$ and l is the x -axis. Then the off-diagonal in Σ_A is

$$\frac{1}{|A|} \int_A xy dx dy = \frac{1}{2|A|} \int_A (xy + x(-y)) dx dy = 0.$$

Thus Σ_A is diagonal and the x and y axes are precisely its principle directions. In particular, the y axis is the normal direction of the mirror. \square

The proof has been greatly simplified by rotational invariance. In practice, this is also exactly the main virtue of the PCA approach since the unknown rotation (i.e. the mirror direction) is found by solving a spectral analysis problem.

But the correlation matrix Σ_A has two principle directions: one is normal to the mirror and the other tangential. The magnitude of the associated eigenvalues does not provide valuable information for identifying which one is the normal (since the shape can be either “fat” or “slim” with respect to the mirror). Thus an extra validation step is required to further single out the normal from the two independent eigenvectors. The computation is straightforward by comparing the symmetry indices (§ 4) of the two hypotheses.

We conclude the section by a stability theorem.

THEOREM 2.3. (LEBESGUE STABILITY) *Suppose that A is a positive compact shape symmetric with respect to the mirror $l = ([A], \mathbf{n})$, and the two eigenvalues λ_1 and λ_2 of Σ_A are distinct. Let A^ε be a perturbed compact shape so that $|A^\varepsilon \ominus A| = O(\varepsilon)$ as $\varepsilon \rightarrow 0$, and $l^\varepsilon = ([A^\varepsilon], \mathbf{n}^\varepsilon)$ the identified mirror from the PCA algorithm. Then*

$$|[A] - [A^\varepsilon]| + |\mathbf{n} - \mathbf{n}^\varepsilon| = O(\varepsilon).$$

Proof. We only sketch the proof by pointing out that when $\lambda_1 \neq \lambda_2$, the spectral decomposition of symmetric matrices Σ

$$\Sigma = \lambda_1 \mathbf{n}_1 \otimes \mathbf{n}_1 + \lambda_2 \mathbf{n}_2 \otimes \mathbf{n}_2$$

is a local diffeomorphism between symmetric matrices and the product space $O(R^2) \times R^2$. Here $O(R^2)$ represents the orthogonal group and R^2 for the eigenvalue pair. It is also straightforward to show that

$$|[A] - [A^\varepsilon]| = O(\varepsilon), \quad \Sigma_{A^\varepsilon} = \Sigma_A + O(\varepsilon).$$

\square

3 Mirror Symmetry via Lower-Dimensional Features

In image analysis and visual perception, many deterministic shape features are invariant under diffeomorphisms. Let $F(A) \subset R^2$ denote the spatial locations of such a class of features (e.g., edges, corners, T-junctions), and $H : R^2 \rightarrow R^2$ a given diffeomorphism. Then,

$$F(H(A)) = H(F(A)).$$

Two visually important classes of lower-dimensional features are edges and corners. We shall denote them by ∂A and $\angle A$ separately. Their Hausdorff dimensions are one and zero for generic Lipschitz shapes. (In image and vision analysis, corners mostly belong to compact shapes with piecewise C^1 boundaries. A corner is where the tangent space cannot be defined.) It is easy to see that indeed they are invariant under diffeomorphisms:

$$\partial H(A) = H(\partial A), \quad \angle H(A) = H(\angle A).$$

Euclidian transforms (i.e. translation, rotation, and reflection) are special examples of (linear) diffeomorphisms. Especially, for mirror reflection R_l , we have $F(R_l A) = R_l F(A)$. Thus,

PROPOSITION 3.1. *If A is l -symmetric, so must be $F(A) : R_l F(A) = F(A)$.*

This provides the foundation for mirror detection via feature sets. The advantage of going from shapes to feature sets is the dimensionality reduction and the resulting efficient computational approach. It is believed in cognitive science that human visual perception often cleverly relies on feature and pattern extraction and processing. Wavelets theory is the most well known example for nicely encoding and processing singular features such as edges [1, 4].

Such dimensionality reduction does come at a potential cost. Features are often only necessary but not sufficient, meaning that if $A = B$, then $F(A) = F(B)$, but the opposite is generally unguaranteed. A possible way to improve is to further quantify the features. Take the class of corner features for example, one may enrich a corner position \mathbf{p} by adding the associated interior angle θ : $\mathbf{p} \rightarrow (\mathbf{p}, \theta)$ (reference to Figure 2).

Let $d\mu$ be a suitable measure on $F(A)$, for example, the 1-D Hausdorff measure dH^1 for ∂A , and the atomic (or counting) measure for $\angle A$. Then both the cut-and-weight and the PCA approaches can be similarly developed on $(F(A), d\mu)$ for mirror detection, with integration $\int_A dx dy$ substituted by $\int_{F(A)} d\mu$. Figure 2 illustrates the cut-and-weight approach for the corner feature set $\angle A$.




Figure 2: Method of cut-and-weight for the corner feature set.

4 Metrics of Mirror Symmetry

The previous section deals with mirror symmetric or nearly mirror symmetric shapes for which the extracted mirrors well approximate human perception. In this section, we intend to quantitatively measure the degree of mirror symmetry, and develop proper metrics which allow us to say that “the degree of mirror symmetry of shape A is 0.618 from a scale of $[0,1]$.” We call 0.618 the symmetry index of A .

4.1 Symmetry index Let $L_+^1 = L_+^1(\mathbb{R}^2)$ denote all non-negative Lebesgue integrable functions. The space of 2-D compact shapes can be naturally embedded into L_+^1 by the characteristic map:

$$A \rightarrow 1_A(\mathbf{p}) = 1, \quad \mathbf{p} \in A; 0, \quad \text{otherwise.}$$

Then it is easy to see that

$$|A \ominus B| = \int_{\mathbb{R}^2} |1_A - 1_B| dx dy = \|1_A - 1_B\|_{L^1}.$$

The notion of mirror symmetry applies to L_+^1 as well. $f \in L_+^1$ is said to be mirror symmetric about a straight line l , if $f = f \circ R_l$. It is obvious that a compact shape A is mirror symmetric about l if and only if $1_A(\mathbf{p})$ is.

For any straight line l , and $f \in L_+^1(\mathbb{R}^2)$, define the l -symmetric component $(f|l)$ and l -asymmetric component $[f|l]$ of f by

$$(f|l)(\mathbf{p}) = \frac{f(\mathbf{p}) + f(R_l \mathbf{p})}{2}, \quad [f|l](\mathbf{p}) = \frac{f(\mathbf{p}) - f(R_l \mathbf{p})}{2}.$$

Then

PROPOSITION 4.1. f is l -symmetric if and only if $(f|l) = f$, or equivalently, $[f|l] = 0$.

Definition. (Symmetry Index) Assume that $f \in L_+^1$, and l is a given straight line.

1) The symmetry index of f with respect to l is defined to be

$$s_l[f] = 1 - \frac{\|[f|l]\|}{\|(f|l)\|}.$$

Here the norms are both in L^1 . Also define $s_l[0] = 1$.

2) The (overall) *symmetry index* of f is defined to be

$$s[f] = \sup_l s_l[f].$$

We define the symmetry indices of a compact shape by setting $s_l[A] = s_l[1_A]$ and $s[A] = s[1_A]$. The proof of the next list of properties is left to the readers.

PROPOSITION 4.2. *Suppose A is a compact shape, $f \in L_+^1(\mathbb{R}^2)$, and l a mirror. Then,*

(a) $0 \leq s_l[f] \leq 1$, and both equalities are achievable.

(b) f is l -symmetric $\Leftrightarrow s_l[f] = 1$.

(c) $s_l[f] = 0 \Leftrightarrow f \times f \circ R_l = 0$, a.e.; $s_l[1_A] = 0 \Leftrightarrow |R_l A \cap A| = 0$.

(d) For a compact shape A ,

$$s_l[A] = \frac{|A \cap R_l A|}{|A \cap R_l A| + |A \ominus R_l A|/2}.$$

Thus the notion of symmetry index does coincide with our common intuition in visual perception. What seems more interesting is the following theorem which claims, roughly speaking, that if you add more symmetry, the symmetry index does increase.

THEOREM 4.1. *Suppose $g \in L_+^1(\mathbb{R}^2)$ is l -symmetric (i.e. $s_l[g] = 1$). Then*

$$s_l[f + g] \geq s_l[f], \quad \forall f \in L_+^1.$$

The equality holds only when $g = 0$ or $s_l[f] = 1$.

Proof. By linearity and $[g|l] = 0$, $(g|l) = g$,

$$[f + g|l] = [f|l] + [g|l] = [f|l]; \quad (f + g|l) = (f|l) + g.$$

Therefore

$$s_l[f + g] = 1 - \frac{\|[f|l]\|}{\|(f|l)\| + \|g\|} \geq 1 - \frac{\|[f|l]\|}{\|(f|l)\|}.$$

The equality holds if and only if $\|g\| = 0$ or $\|[f|l]\| = 0$, which completes the proof. \square

As for compact shapes, a similar but independent theorem can be established.

THEOREM 4.2. *If a compact shape B is l -symmetric, then for any compact shape A ,*

$$s_l[A \cup B] \geq s_l[A].$$

The technical independence of the two theorems is due to that generally,

$$1_{A \cup B} \neq 1_A + 1_B.$$

Proof. First if $A \subset C$, then $1_C \geq 1_A$, and $(1_C|l) \geq (1_A|l)$. In particular, $(1_{A \cup B}|l) \geq (1_A|l)$. Now that $1_{A \cup B} = 1_B + 1_{A-B}$ and B is l -symmetric, we must have

$$[1_{A \cup B}|l] = [1_{A-B}|l].$$

Thus the proof will be complete if we are able to show

$$|[1_{A-B}|l]| \leq |[1_A|l]|.$$

Suppose $|[1_{A-B}|l](\mathbf{p})| \neq 0$. By definition we must have $[1_{A-B}|l] = \pm 1/2$, and one and only one from $\{\mathbf{p}, R_l \mathbf{p}\}$ belongs to $A - B$. Then $\{\mathbf{p}, R_l \mathbf{p}\} \cap B$ must be empty since B is l -symmetric, which implies that one and only one of $\{\mathbf{p}, R_l \mathbf{p}\}$ is in A . Thus, $[1_A|l](\mathbf{p}) = \pm 1/2$. This concludes the proof. \square

The concept of symmetry index brings two new degrees of freedom:

- 1) any compact shape can have certain degree of mirror symmetry; and,
- 2) mirror identification becomes an optimization problem: $\sup_l s_l[A]$.

For a general compact shape A , it seems to the author that there is no direct efficient optimization algorithm. When A is a smooth shape, $S_l[A]$ is smooth in l as well, and gradient descent marching in l works. Still uniqueness is not guaranteed, but existence can be established.

THEOREM 4.3. (EXISTENCE OF OPTIMAL MIRRORS) *There exists a mirror l_0 so that $s_{l_0}[A] = \max_l s_l[A]$.*

Proof. Let $l_k = (\mathbf{p}_k, \mathbf{n}_k)$ be a maximizing mirror sequence. ($l = (\mathbf{p}, \mathbf{n})$ means that l goes through \mathbf{p} with normal \mathbf{n} .) Without loss of generality, assume that $s_{l_k}[A] > 0, k = 1, 2, \dots$. Let $h(A)$ denote the convex hull of A . Then $l_k \cap A$ cannot be empty. Otherwise, the set of inner products $\langle h(A) - \mathbf{p}_k, \mathbf{n}_k \rangle$ must be either simultaneously non-positive or non-negative. It implies that $h(A)$ and thus A must lie on one single side of l . In particular, $R_{l_k}A \cap A$ is empty, and by Proposition 4.2, $s_{l_k}[A] = 0$, contradicting to our assumption. Therefore, we can assume that $\mathbf{p}_k \in h(A), k = 1, 2, \dots$. We finally wrap up the proof by applying the compactness of both $h(A)$ and the unit circle S^1 to the sequence $(\mathbf{p}_k, \mathbf{n}_k)$. \square

4.2 Hausdorff distance for feature sets The above approach only applies to 2-D shapes and functions. For lower-dimensional feature sets $F(A)$ such as ∂A and $\angle A$ (see § 3), we turn to Hausdorff distance.

For any compact set $F \subset R^2$, its δ -neighborhood is defined as

$$F_\delta = \{\mathbf{p} | d(\mathbf{p}, F) \leq \delta\}.$$

Then the Hausdorff distance between two compact sets E and F is defined by

$$d_H(E, F) = \inf\{\delta | E \subset F_\delta \quad \text{and} \quad F \subset E_\delta\}.$$

Let $F = F(A)$ denote a feature set of A . For any mirror l , we define the symmetry index of F with respect to l by

$$s_l[F] = \exp(-d_H(F, R_l F)).$$

As in the previous section, the overall mirror symmetry index is defined to be

$$s[F] = \sup_l s_l[F].$$

Notice that Hausdorff does weigh each individual feature point, which is different from Lebesgue measures.

Like the symmetry index for shapes and functions, we have

- (a) $0 < s_l[F(A)] \leq 1$ for any feature set $F(A)$.
- (b) A feature set $F(A)$ is l -symmetric if and only if $s_l[F(A)] = 1$.
- (c) The optimal mirror for $F[A]$ exists.

Many properties in the previous section can be similarly studied.

5 Mirror Symmetry of Convex Shapes

In this section, we propose a new approach for studying mirror symmetry of compact convex shapes, based on the tool of *support functions*.

Let A be a compact convex shape in R^2 . Its support function $h_A(\mathbf{u})$ is defined for any vector $\mathbf{u} \in R^2$ by:

$$h_A(\mathbf{u}) = \max\langle A, \mathbf{u} \rangle = \max_{\mathbf{p} \in A} \langle \mathbf{p}, \mathbf{u} \rangle.$$

Recall the basic properties of the support function (e.g., [7]):

- (a) For any scalar $\lambda > 0$, $h_A(\lambda \mathbf{u}) = \lambda h_A(\mathbf{u})$.

(b) For any $\mathbf{u}, \mathbf{v} \in R^2$, $h_A(\mathbf{u} + \mathbf{v}) \leq h_A(\mathbf{u}) + h_A(\mathbf{v})$.

(c) For any two convex compact shapes A and B , let $d_H(A, B)$ be their Hausdorff distance, and restrict h_A on the unit circle S^1 . Then

$$d_H(A, B) = \|h_A - h_B\|_{L^\infty(S^1)}.$$

In fact, (a) and (b) *completely* characterize a support function [7]. That is, a function $f(\mathbf{u})$ is the support function for some convex shape if and only if it carries the properties (a) and (b). Moreover, the correspondence $A \leftrightarrow h_A$ is one-to-one and even isomorphic in some suitable sense based on the last property.

The support function is not translation invariant in A since

$$h_{A+\mathbf{p}}(\mathbf{u}) = h_A(\mathbf{u}) + \langle \mathbf{p}, \mathbf{u} \rangle.$$

We propose to homogenize h_A by

$$H_A(\mathbf{u}) = h_A(\mathbf{u}) + h_A(-\mathbf{u}).$$

We shall call H_A the homogenized support function. By design it is translation invariant in A , and

$$H_A(-\mathbf{u}) = H_A(\mathbf{u}),$$

which makes $H_A(\mathbf{u})$ well defined on the real projective line RP^1 . Let $\mathbf{u} = \exp(i\theta)$ be restricted on the unit circle S^1 and define $H_A(\theta) = H_A(\mathbf{u})$. Then $H_A(\theta)$ is π -periodic:

$$H_A(\theta + \pi) = H_A(\theta).$$

Now let l be any given mirror. It is easy to see that

$$H_{R_l A}(\mathbf{u}) = H_A(R_l \mathbf{u}).$$

In particular, if A is l -symmetric, so is $H_A(\cdot)$. Since H_A is translation invariant in A , we can assume that the mass center of A is the origin: $[A] = (0, 0)$. Then l passes through the origin. Let $z_0 = \exp(i\theta_0)$ be the trace of l on the unit circle (any one from the two trace points). Then the mirror reflection $\mathbf{u} = \exp(i\theta) \rightarrow R_l \mathbf{u}$ is equivalent to

$$\theta \rightarrow 2\theta_0 - \theta.$$

Therefore, we have proved that

PROPOSITION 5.1. *If a convex compact shape A with mass center $[A] = 0$ is l -mirror symmetric, then*

$$H_A(\theta) = H_A(2\theta_0 - \theta).$$

Here $\exp(i\theta_0)$ is the trace of l on the unit circle.

This theorem provides the foundation for mirror mining based on the homogenized support function H_A restricted on the unit circle. We shall see that mirror identification now becomes a phase identification problem.

Since $H_A(\theta)$ is π -periodic, its Fourier series are in the form of

$$H_A(\theta) = \sum_{-\infty}^{\infty} a_n e^{-2in\theta}.$$

$a_{-n} = \bar{a}_n$ because H_A is real. Thus

$$H_A(2\theta_0 - \theta) = \sum_{-\infty}^{\infty} (\bar{a}_n e^{4in\theta_0}) e^{-2in\theta},$$

and we must have $a_n = \bar{a}_n e^{4in\theta_0}$, $n = 1, 2, \dots$. Define $z_0 = e^{4i\theta_0}$ and $z_n = a_n / \bar{a}_n \in S^1$ (possibly having those vanishing a_n 's dropped). Then

$$(5.4) \quad z_0^n = z_n, \quad n = 1, 2, \dots$$

In practice, for near symmetric convex shapes, (5.4) only holds approximately. Thus we apply the least square fitting to extract z_0 . First truncate the Fourier series to finite N terms, and then solve the minimization problem

$$\min_{z_0: |z_0|=1} \sum_{n=1}^N w_n |z_0^n - z_n|^2,$$

with a suitable set of positive weights (w_n), often larger for lower frequency modes. Finally, notice that a single optimal z_0 can produce four angles $\theta_0, \theta_0 + \pi/2, \theta_0 + \pi, \theta_0 + 3\pi/2$, corresponding to the optimal mirror and the dual mirror l^\perp perpendicular to it.

6 Mirror Symmetry of C^1 Shapes via Normal Pairs

In this section, we give a new characterization for mirror symmetry of 2-D compact C^1 shapes, i.e., shapes with C^1 boundary manifolds.

Definition. (Normal Pairs) Let A be a compact C^1 shape with boundary ∂A . A pair of distinct points $\mathbf{p}, \mathbf{q} \in \partial A$ is said to be a *normal pair* if

$$\mathbf{p} - \mathbf{q} \perp T_{\mathbf{p}} \partial A \quad \text{and} \quad \mathbf{p} - \mathbf{q} \perp T_{\mathbf{q}} \partial A,$$

Where $T_{\mathbf{p}} \partial A$ denotes the tangent space at \mathbf{p} .

PROPOSITION 6.1. (2-D ONLY) *Suppose A is l -symmetric. Then any pair of distinct points $\mathbf{p}, \mathbf{q} \in \partial A \cap l$ is a normal pair.*

Proof. Simply notice that if $\mathbf{p} \in \partial A \cap l$, then $T_{\mathbf{p}} \partial A$ must be perpendicular to l .

This proposition is generally false in higher dimensions.

If all normal pairs are along the mirror, then mirror identification certainly becomes the detection of normal pairs. But the truth is that generically, there indeed exists at least one normal pair off the mirror.

THEOREM 6.1. (2-D ONLY) *Let A be a compact C^1 shape that is l -symmetric. If, in addition, A is connected, then A has at least two distinct normal pairs.*

Proof. Since A is connected and l -symmetric, it is not difficult to show that $\partial A \cap l$ is non-empty and contains at least two points, which contribute at least one normal pair according to the previous proposition. (Pay attention to the assumption that ∂A is a C^1 submanifold of R^2 , which prohibits singular contacting boundary points shaped like a soft \mathbf{x} .)

Now we look for another normal pair which is perpendicular to the mirror. For simplicity, assume that l goes through the origin $(0, 0)$. Let \mathbf{n} denote any one of the two unit normals of l . Define the normal projection function

$$g(\mathbf{p}) = \langle \mathbf{p}, \mathbf{n} \rangle, \quad \mathbf{p} \in R^2.$$

Since g is continuous and A is compact, there exists some $\mathbf{q} \in A$ so that

$$\mathbf{q} = \operatorname{argmax}_{\mathbf{p} \in A} g(\mathbf{p}).$$

By mirror symmetry, we must have

$$R_l \mathbf{q} = \operatorname{argmin}_{\mathbf{p} \in A} g(\mathbf{p}).$$

It is then not difficult to show that (1) $\mathbf{q}, R_l \mathbf{q} \in \partial A$; (2) $T_{\mathbf{q}} \partial A, T_{R_l \mathbf{q}} \partial A \perp \mathbf{n}$; and (3) $\mathbf{q} - R_l \mathbf{q} \perp l$. Thus \mathbf{q} and $R_l \mathbf{q}$ is a new normal pair. This concludes the proof. \square

For an ideal disk, there are infinitely many normal pairs. For a generic ellipse, there are two and only two normal pairs corresponding exactly to the two symmetry mirrors. For generic shapes, normal pair is still a powerful tool to extract either the tangential or normal directions of the symmetry mirrors, even though success is not 100% guaranteed.

We conclude this section by stating an algorithm for computationally identifying normal pairs.

Normal Shooting Algorithm (NSA). Start with any initial guess $\mathbf{p}_0 \in \partial A$, and do the following iteration $\mathbf{p}_{k+1} = \text{NSA}(\mathbf{p}_k)$. Let N_k denote the normal line of the current tangent space $T_{\mathbf{p}_k} \partial A$. Then take any $\mathbf{p}_{k+1} \in (N_k - \{\mathbf{p}_k\}) \cap \partial A$ (which must be non-empty).

It is not difficult to establish the following theorem.

THEOREM 6.2. *Suppose that (\mathbf{p}_{2k}) converges to \mathbf{p}^* , and (\mathbf{p}_{2k+1}) to \mathbf{q}^* , as $k \rightarrow \infty$. The \mathbf{p}^* and \mathbf{q}^* is a normal pair.*




Figure 3: The Normal Shooting Algorithm (NSA) and its convergence.

7 Conclusion

The current paper proposes several different approaches to extract symmetry mirrors from 2-D compact mirror symmetric shapes. Visually meaningful symmetry measures are introduced based on Lebesgue measure, Hausdorff distance, and lower-dimensional feature sets. Theory and computation of these approaches and measures are developed.

Important applications can be found for human and object detection, template based object tracking, data compression, and high-level inpainting algorithms based on global features [3].

The present work also sets an example for studying more complex symmetry types such as affine and projective symmetries, which are also very important in mathematical image and vision analysis.

To understand which mechanisms human observers truly employ, however, all the models and algorithms developed here must be further validated by psychology experiments in human vision research.

Acknowledgments

The author would like to thank Tony Chan for his integrity and dignity as a scholar, advisor, and dear friend. Thanks also must go to the Program of Applied Mathematics, National Science Foundation, for her generous support to the author's research.

References

- [1] E. J. Candés and D. L. Donoho. Curvelets and reconstruction of images from noisy radon data. *Wavelet Applications in Signal and Image Processing VIII*, A. Aldroubi, A. F. Laine, M. A. Unser eds., Proc. SPIE 4119, 2000.
- [2] T. F. Chan and J. Shen. Variational restoration of non-flat image features: models and algorithms. *SIAM J. Appl. Math.*, 61(4):1338–1361, 2000.
- [3] T. F. Chan and J. Shen. Mathematical models for local nontexture inpaintings. *SIAM J. Appl. Math.*, 62(3):1019–1043, 2001.
- [4] I. Daubechies. *Ten lectures on wavelets*. SIAM, Philadelphia, 1992.
- [5] U. Grenander. *Lectures in Pattern Theory. I. II. and III*. Springer, 1976-1981.
- [6] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Applied. Math.*, 42:577–685, 1989.
- [7] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1970.
- [8] J. Shen. On the foundations of vision modeling. I. Weber's law and Weberized TV restoration. *Physica D*, to appear, 2002.
- [9] G. Strang. *Introduction to Applied Mathematics*. Wellesley-Cambridge Press, MA, 1993.