Mirror-symmetry in images and 3D shapes

Viorica Pătrăucean
Geometrica, Inria Saclay

Journées de Géométrie Algorithmique – CIRM Luminy 2013

Symmetry is all around...

In nature: mirror, rotational, helical, scale (fractal)

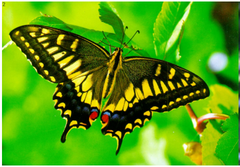

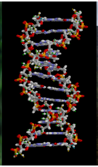


Man-made objects: mirror, rotational, quadrilateral



Music: translation, glide reflection

Beethoven's *Moonlight* Sonata

Chopin's Waltz

Symmetry in *Image and 3D Shape Analysis*

- Object recognition in human perception is greatly enhanced by the presence of symmetries [BR79]
- Automatic *object recognition in images* and *3D shape matching* work better in the presence of symmetries?

Symmetry in *Image and 3D Shape Analysis*

- Object recognition in human perception is greatly enhanced by the presence of symmetries [BR79]
- Automatic *object recognition in images* and *3D shape matching* work better in the presence of symmetries?

Outline:

II Symmetry in images
III Symmetry in 3D shapes
IV Putting all together

PART II: Symmetry in images

Object recognition in images

- Recognition rate improves when symmetry information is available [PLC^{+}12]
- Popular state-of-the-art methods use bag-of-words approaches based on SIFT-like descriptors
- Only few recent works try to include symmetry descriptors [HS12]

PART II: Symmetry in images

Object recognition in images

- Recognition rate improves when symmetry information is available [PLC+12]
- Popular state-of-the-art methods use bag-of-words approaches based on SIFT-like descriptors
- Only few recent works try to include symmetry descriptors [HS12]
- Reliable symmetry detection in images is difficult
  - Background clutter
  - Partial (not perfect) symmetries
  - Perspective effect, occlusions

PART II: Symmetry in images

Object recognition in images

- Recognition rate improves when symmetry information is available [PLC+12]
- Popular state-of-the-art methods use bag-of-words approaches based on SIFT-like descriptors
- Only few recent works try to include symmetry descriptors [HS12]
- Reliable symmetry detection in images is difficult
  - Background clutter
  - Partial (not perfect) symmetries
  - Perspective effect, occlusions

Contribution: *Symmetry detection in images with controlled number of false positives*

Symmetry detection in images with controlled number of false positives

Joint work with R. Grompone von Gioi (ENS Cachan) and M. Ovsjanikov (École Polytechnique) – Symmetry Competition Workshop CVPR2013

Proposed method

- Mirror symmetry detection
- Orthogonal view of objects, no perspective skew
- Detection = two-stage process: *candidate selection* + ***validation***
- Main concern: diminish false detections → *a contrario* approach

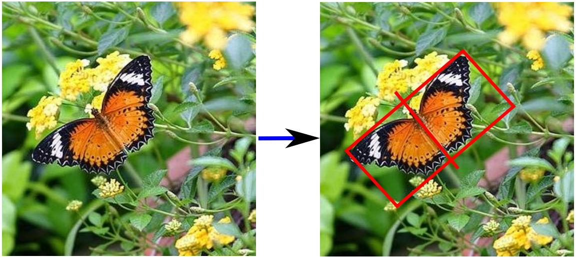

Candidate selection

Goal: No false negatives!

Candidate: image patch $(x_1, y_1, x_2, y_2, width) \sim (\rho, \theta, width, height, offset)$

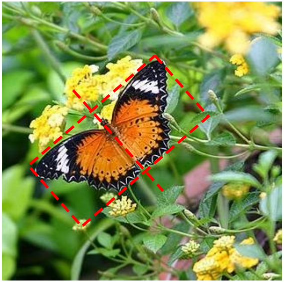

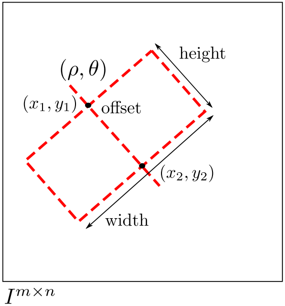

5 degrees of freedom → exhaustive search not feasible

Candidate selection

Goal: No false negatives!

Candidate: image patch $(x_1, y_1, x_2, y_2, width) \sim (\rho, \theta, width, height, offset)$

A. Reisfeld voting using SIFT features $\rightarrow (\rho, \theta)$ [LE06]

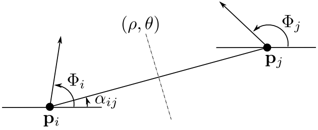

$$\mathbf{w}_{ij} = 1 - \cos(\Phi_i + \Phi_j - 2\alpha_{ij})$$

B. Exhaustive search along (ρ, θ) axis using *integral images* $\rightarrow (width, height, offset)$

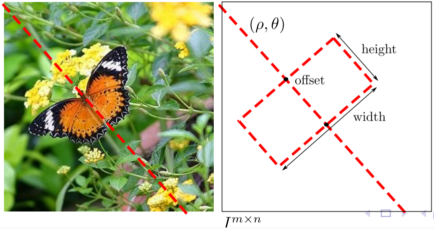

Validation

Goal: No false positives!

Candidate: image patch $s(x_1, y_1, x_2, y_2, width) \sim s(\rho, \theta, width, height, offset)$

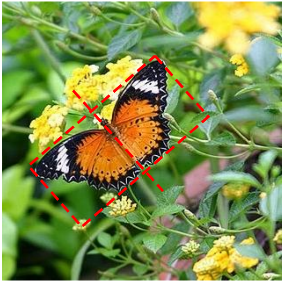

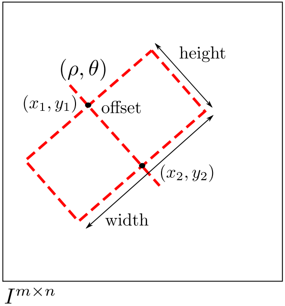

Is the given patch a meaningful symmetry?

Validation

Goal: No false positives!

Candidate: image patch $s(x_1, y_1, x_2, y_2, width) \sim s(\rho, \theta, width, height, offset)$

Measure for the **degree of symmetry**: *gradient orientation error*

for all the η *pairs of pixels* in *patch*
accumulate normalised angular error

$$k(s) = \sum_{i}^{\eta} \frac{|\delta_i|}{\pi}$$

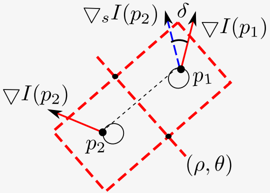

$$k(s) \in [0, \eta] \begin{cases} 0, & \text{perfect symmetry} \\ \eta, & \text{worst symmetry} \end{cases}$$

Need a detection threshold on $k(s)$

Validation

Goal: No false positives!

***A contrario* theory**: formalises the *non-accidentalness principle* *("no perception in noise")* using a *multiple hypothesis testing* approach

- Random model: *null hypothesis* $\mathcal{H}_0$
- Candidate s: *outlier* w.r.t. $\mathcal{H}_0$? $\xrightarrow{yes}$ s – *meaningful* detection

Validation

Goal: No false positives!

***A contrario* theory**: formalises the *non-accidentalness principle* (*"no perception in noise"*) using a *multiple hypothesis testing* approach

- Random model: *null hypothesis* $\mathcal{H}_0$
- Candidate s: *outlier* w.r.t. $\mathcal{H}_0$? $\xrightarrow{yes}$ s – *meaningful* detection

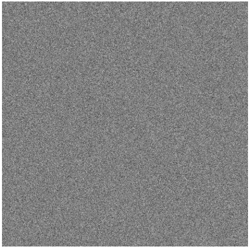

Validation

Goal: No false positives!

***A contrario* theory**: formalises the *non-accidentalness principle* (*"no perception in noise"*) using a *multiple hypothesis testing* approach

- Random model: *null hypothesis* $\mathcal{H}_0$
- Candidate s: *outlier* w.r.t. $\mathcal{H}_0$? $\xrightarrow{yes}$ s – *meaningful* detection

Random model $\mathcal{H}_0$: Gaussian white noise

p-value: $\mathbb{P}_{\mathcal{H}_0}[k_X(s) \leq k(s)]$

Number of candidates: $N_I = (mn)^{\frac{5}{2}}$

$\varepsilon = 1 \rightarrow$ *one* false positive per image

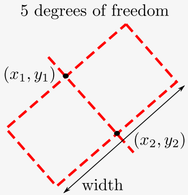

Meaningfulness test: $N_I \cdot \text{p-value}(s) \leq \varepsilon$

Results

- No detection in noise images
- Satisfactory results in general

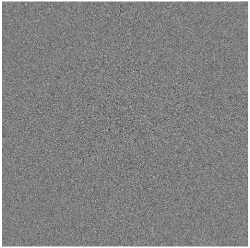

Human labels

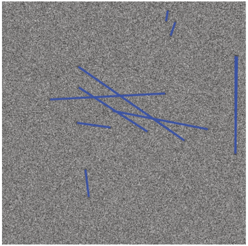

State-of-the-art

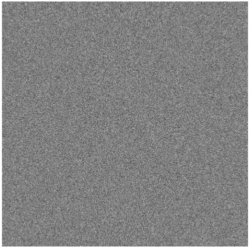

Proposed

Results

- No detection in noise images
- Satisfactory results in general

Human labels

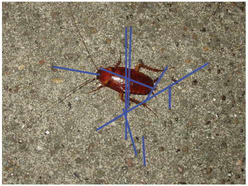

State-of-the-art

Proposed

Results

- No detection in noise images
- Satisfactory results in general

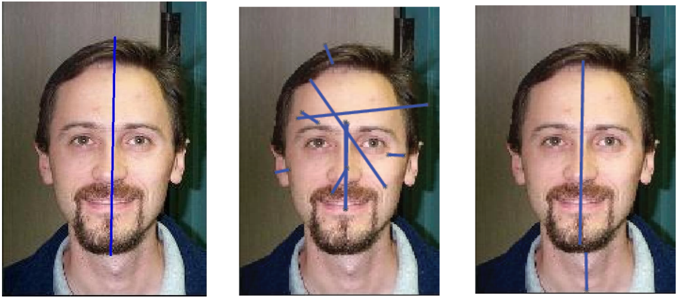

Human labels State-of-the-art Proposed

Results

- No detection in noise images
- Satisfactory results in general

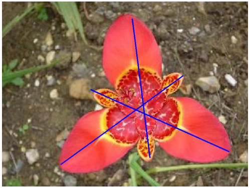

Human labels

State-of-the-art

Proposed

Summary

Is symmetry information **useful** and **used** in object recognition in images?

- **Conclusion:** Using symmetry information improves object recognition in images
- **Contribution:** Parameterless mirror-symmetry detection with controlled number of false positives
- **Limitations:** No perspective skew
- **Future work:** Design symmetry descriptors for object recognition tasks

PART III: Symmetry in 3D shapes

3D shape matching: Find *isometric correspondences* between shapes with *intrinsic symmetries*

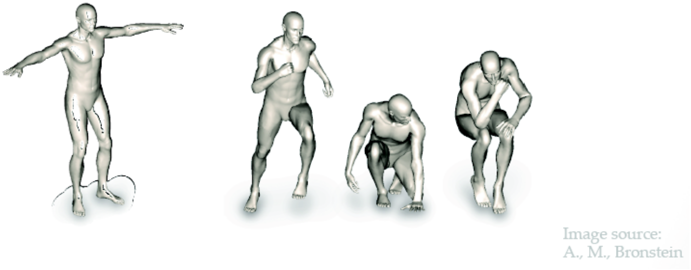

Isometries approximate well articulated motion of humans and animals

Approximately preserve geodesic distances between pairs of points

PART III: Symmetry in 3D shapes

3D shape matching: Find *isometric correspondences* between shapes with *intrinsic symmetries*

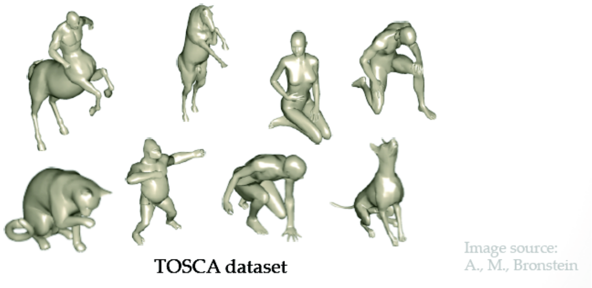

Most approximately isometric shapes contain intrinsic symmetries

Approximately preserve geodesic distances between pairs of points

PART III: Symmetry in 3D shapes

Given a pair of shapes with intrinsic symmetries, find isometric correspondences

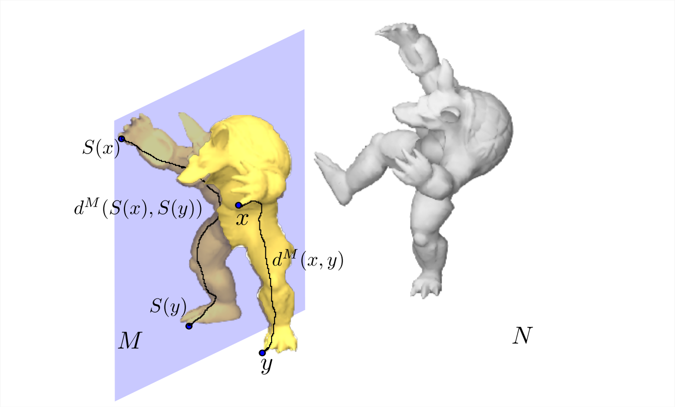

Approximately preserve geodesic distances between pairs of points

PART III: Symmetry in 3D shapes

Given a pair of shapes with intrinsic symmetries, find isometric correspondences

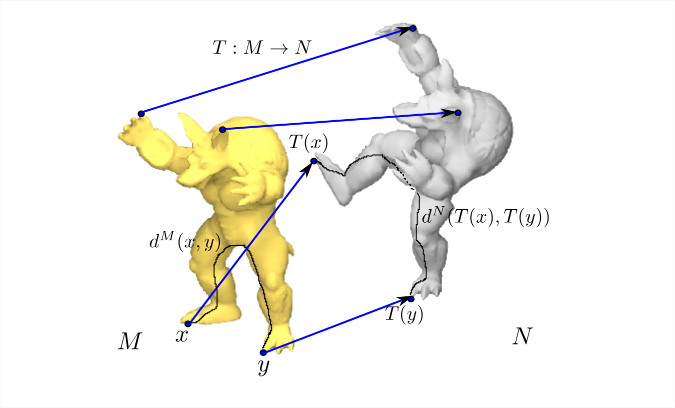

$$\hat{T} = \arg\min_{T} \sum_{x,y} |d^M(x,y) - d^N(T(x),T(y))|$$ **difficult**

PART III: Symmetry in 3D shapes

Given a pair of shapes with intrinsic symmetries, find isometric correspondences

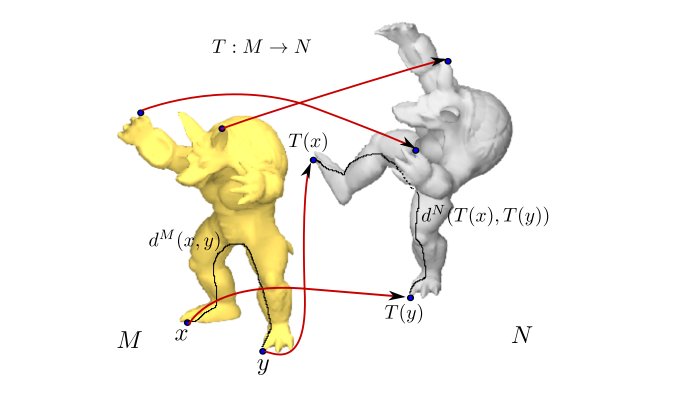

At least two equally-good solutions; non-convex problem $\leftarrow$ symmetry ambiguity

PART III: Symmetry in 3D shapes

Given a pair of shapes with intrinsic symmetries, find isometric correspondences

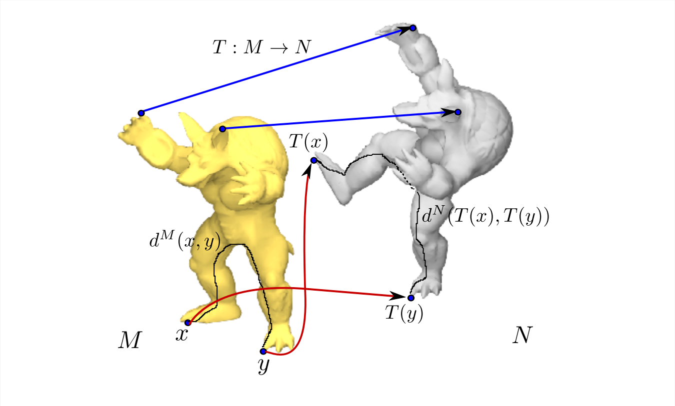

Symmetry flipping $\to$ continuity issue; **symmetry makes the problem harder**

PART III: Symmetry in 3D shapes

Given a pair of shapes with intrinsic symmetries, find isometric correspondences

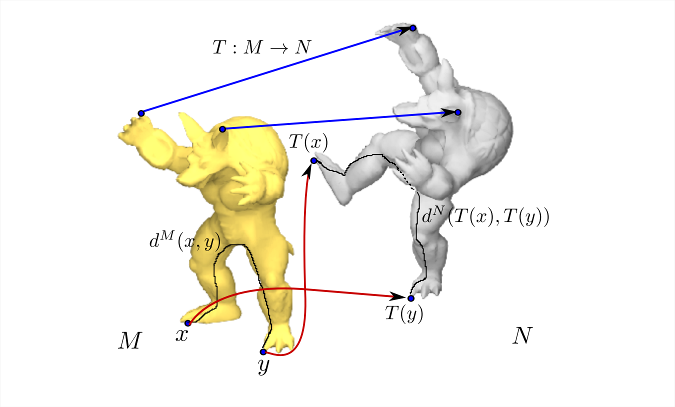

Existing work: coarse-to-fine approaches, based on some initial point correspondences

PART III: Symmetry in 3D shapes

Given a pair of shapes with intrinsic symmetries, find isometric correspondences

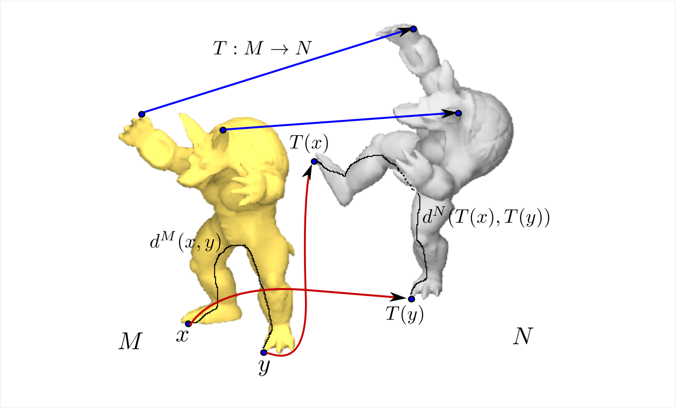

Existing work: coarse-to-fine approaches, based on some initial point correspondences → **complex and error-prone** [GYF11]

PART III: Symmetry in 3D shapes

Given a pair of shapes with intrinsic symmetries, find isometric correspondences

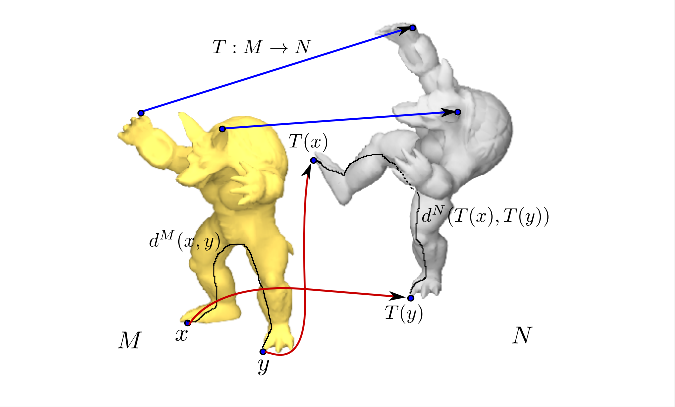

Contribution: *Matching of 3D shapes with intrinsic symmetries without point correspondences*

Matching of 3D shapes with intrinsic symmetries

Joint work with M. Ovsjanikov (École Polytechnique), Q. Mérigot (CNRS), L. Guibas (Stanford University) – SGP2013

Proposed method

- Uses *functional map* framework [OBCS$^+$12] for shape matching
- Matching is done between halves of the shapes
- Two equally good solutions are returned

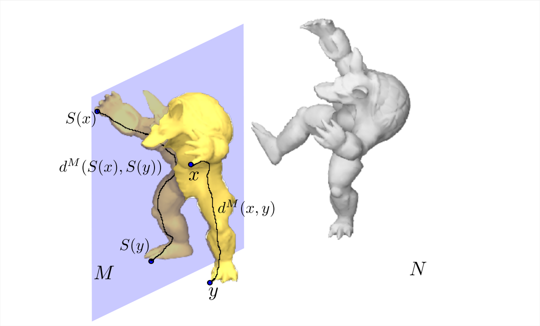

Functional map representation

New concept of matching between shapes

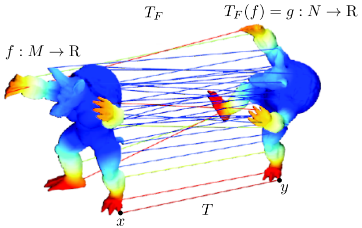

- $T : N \to M$, pointwise map $T(y) = x$
- $T_F : L^2(M) \to L^2(N)$, function-wise map $T_F(f) = g$, where $g = f \circ T$
- $T \Rightarrow T_F$ and $T_F \Rightarrow T$

Functional map representation

New concept of matching between shapes

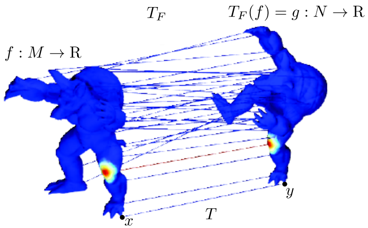

- $T : N \to M$, pointwise map $T(y) = x$
- $T_F : L^2(M) \to L^2(N)$, function-wise map $T_F(f) = g$, where $g = f \circ T$
- $T \Rightarrow T_F$ and $T_F \Rightarrow T$

Functional map representation

How to compute T_F?

- T_F – linear map between vector spaces
- Choose *convenient* bases (LB eigenfunctions) for the two vector spaces: $f = \sum_i a_i \Phi_i^M$ and $g = \sum_i b_i \Phi_i^N$
- T_F – matrix representation $Ca = b$
- Given enough pairs of functions defined on the two shapes, C can be recovered through a least squares system
- Function constraints: descriptor preservation (HKS, WKS), landmark correspondences
- Functional map: state-of-the-art results in isometric shape matching
- Drawback: symmetry flipping – needs some point correspondences

Quotient space matching

Given a pair of shapes with **known** intrinsic symmetries, solve the symmetry ambiguity by *matching between halves of the shapes*

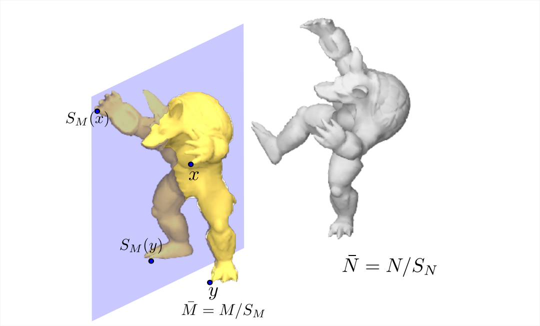

- Quotienting the shape corresponds to splitting the function space

Quotient space matching

Given a pair of shapes with known symmetries:

- Split the space of functions into symmetric and antisymmetric subspaces

$$L^2_+(M) = \{f \in L^2(M) \mid f \circ S_M = f\}; \quad L^2(M) = L^2_+(M) \oplus L^2_-(M)$$

- There exists an orthogonal basis of $L^2_+(M)$ formed by LB eigenfunctions
- Functional map decomposed into parts $C = C_+ \oplus C_-$, estimated independently

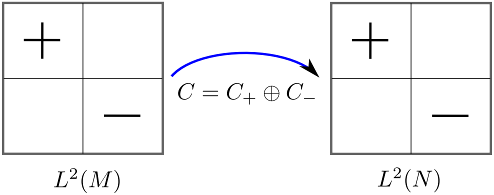

Quotient space matching

Given a pair of shapes with known symmetries:

- Solve for C_+ as before
- C_+ is easier to compute than C
  - C_+ is unique
  - The descriptors (HKS, WKS) are usually symmetric functions
- Use C_+ to recover a point-to-orbit map
- Use the known symmetries to compute two equally-good point-to-point maps

Semi-quotient space matching

More practical scenario: the symmetry is known on only one shape

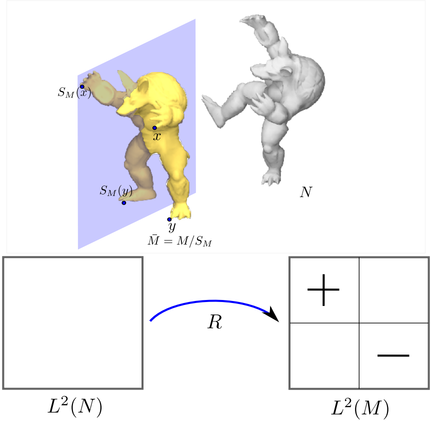

- For the map $C : L^2(N) \to L^2(M)$, compute its symmetric part R
- Transfer the symmetry, and proceed as before

Results

Symmetry transfer accuracy

State-of-the-art

Proposed method

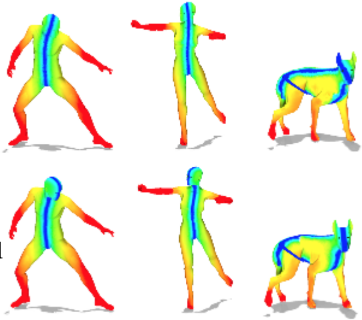

Results

Quotient space matching

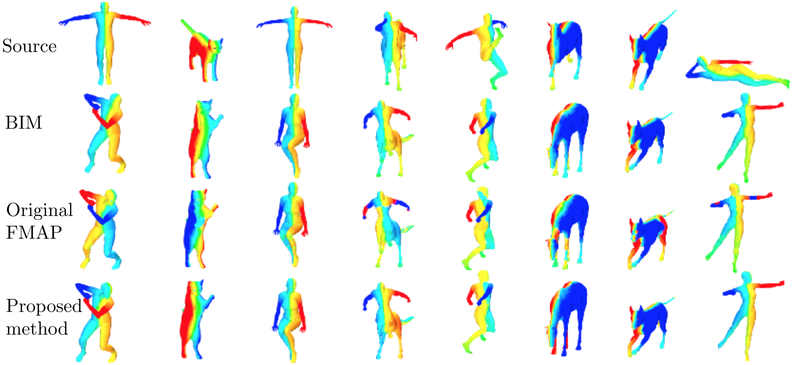

Summary

Is symmetry information **useful** and **used** in 3D shape matching?

- **Conclusion:** Shape matching apparently more difficult for shapes with symmetries
- **Contribution:** Matching of shapes with intrinsic symmetry without point correspondences
- **Bonus:** Accurate symmetry estimation on an unknown shape through symmetry transfer

PART IV: Putting all together

- Large public repositories of images (flickr) and 3D shapes (Google Warehouse)
- Increasing interest for joint image and 3D shape analysis
- Possible application: automatic texture mapping from natural images

PART IV: Putting all together

- Large public repositories of images (flickr) and 3D shapes (Google Warehouse)
- Increasing interest for joint image and 3D shape analysis
- Possible application: automatic texture mapping from natural images

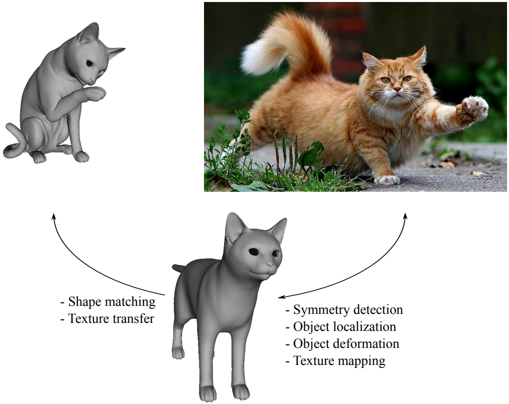

References I

- H. B. Barlow and B. C. Reeves, *The versatility and absolute efficiency of detecting mirror symmetry in random dot displays.*, Vision research **19** (1979), no. 7, 783–793.
- Kim V. G., Lipman Y., and Fun, *Blended intrinsic maps*, ACM TOG (2011).
- D. C. Hauagge and N. Snavely, *Image matching using local symmetry features*, Proc. of CVPR, 2012.
- G. Loy and J.-O. Eklundh, *Detecting symmetry and symmetric constellations of features*, Proc. of ECCV, 2006.
- Ovsjanikov, Ben-Chen, Solomon, Butscher, and Guibas, *Functional maps: A flexible representation of maps between shapes*, SIGGRAPH, 2012.
- M. Park, S. Lee, P.-C. Chen, S. Kashyap, A. A. Butt, and Y. Liu, *Performance evaluation of state-of-the-art discrete symmetry detection algorithms*, Proc. of ECCV, 2012.