Towards an observer-oriented theory of shape comparison

Patrizio Frosini

Department of Mathematics and ARCES, University of Bologna
patrizio.frosini@unibo.it

9^{th} Eurographics Workshop on 3D Object Retrieval
Lisbon, 8 May 2016

Outline

Our basic questions

Assumptions in our model

Mathematical setting and theoretical results

A first step towards the application of our model: GIPHOD

Our basic questions

We are interested in these questions:

- Is there a general metric model for shape comparison?
- What should be the role of the observer in such a model?
- How could we approximate the metric used in that model?

Our talk will be devoted to illustrate these questions and to propose some answers by means of a mathematical approach based on persistent homology and group invariant non-expansive operators.

Assumptions in our model

Truth often depends on the observer's perspective:

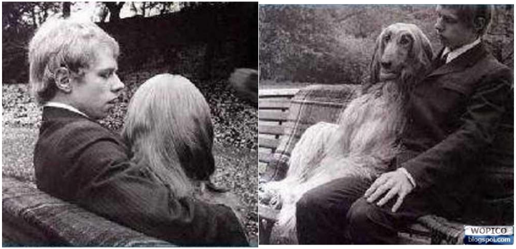

Multiple perspectives are usually unavoidable!

In the past this observation was mostly confined to the philosophical debate, but nowadays it starts to be quite relevant also in several scientific applications involving Information Technology.

Assumptions in our model

We will make these assumptions:

1. No object can be studied in a direct and absolute way. Any object is only knowable through acts of measurement made by an observer.
2. Any act of measurement can be represented as a function defined on a topological space.
3. The observer usually acquires measurement data by applying operators to the functions describing these data. These operators are frequently endowed with some invariances that are relevant for the observer.
4. Only the observer is entitled to decide about shape similarity.

An example of measurement

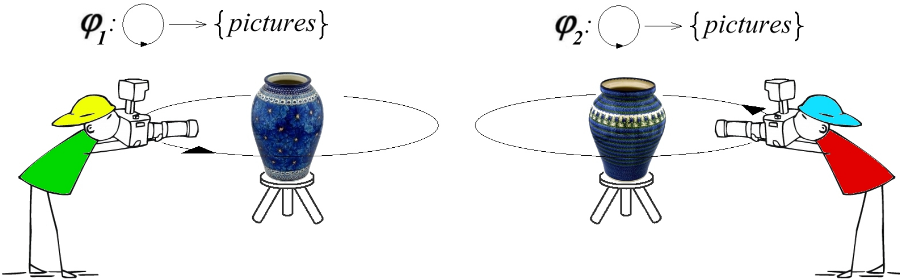

Another example of measurement

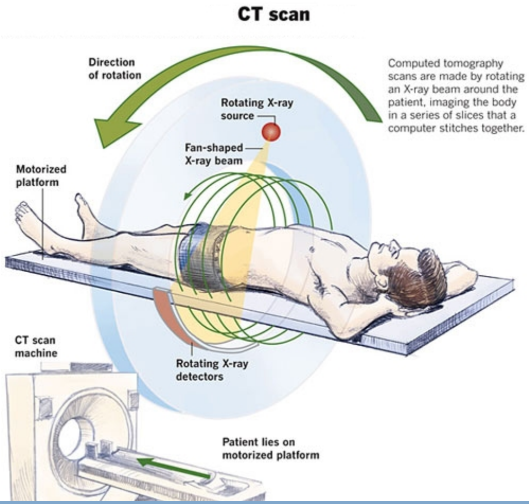

Another example of measurement

An example of operator

GRAYSCALE IMAGE

$\varphi : [0, a] \times [0, b] \longrightarrow \mathbb{R}$ $\quad T(\varphi)$ = CONVOLUTION OF φ

Choice of the operators

- The observer cannot usually choose the functions representing the measurement data, but can often choose the operators that will be applied to those functions.
- The choice of the operators reflects the invariances that are relevant for the observer.
- In some sense we could state that the observer can be represented as a collection of (suitable) operators, endowed with the invariance he/she has chosen.

In this talk we will confine ourselves to examine the case of operators that act on a space Φ of continuous functions and take Φ to itself.

From comparing sets in $\mathbb{R}^n$ to comparing functions

Instead of directly focusing on the objects we are interested in, we focus on the functions describing the measurements we make on them, and on the "glasses" that we use "to observe" the functions. In our approach, these "glasses" are G-operators which act on the functions.
These operators represent the observer's perspective.

In some sense, the family of operators defines the observer.

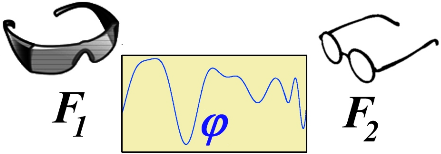

Observers are usually interested in invariances

The observer usually takes some invariance into account. We suggest that this invariance could be represented by a group of homeomorphisms. The reason is that, if measurement data are described by functions from a topological space X to $\mathbb{R}^k$, a natural way of stating the equivalence between two functions $\varphi_1, \varphi_2 : X \to \mathbb{R}^k$ consists in saying that $\varphi_1 \equiv \varphi_2 \circ g$ for a suitable homeomorphism g chosen in a given group G of self-homeomorphisms of X. The composition of φ_2 with g to obtain φ_1 can be seen as a kind of alignment of data, as happens in image registration. The choice of the group G corresponds to the selection of the alignments of data that are judged admissible by the observer.

These observations justify the introduction of the G-invariant (pseudo)metric that will be defined in the next slide.

Natural pseudo-distance associated with a group G

In our model data are compared by the following pseudo-metric.
(pseudo-metric=metric without the property $d(x,y)=0 \implies x=y$).

Definition

Let X be a compact space. Let G be a subgroup of the group $\mathrm{Homeo}(X)$ of all homeomorphisms $f: X \to X$. The pseudo-distance $d_G : C^0(X,\mathbb{R}^k) \times C^0(X,\mathbb{R}^k) \to \mathbb{R}$ defined by setting

$$d_G(\varphi,\psi) = \inf_{g\in G} \max_{x\in X} \|\varphi(x) - \psi(g(x))\|_\infty$$

is called the natural pseudo-distance associated with the group G.

In plain words, the definition of d_G is based on the attempt of finding the best correspondence between the functions φ,ψ by means of homeomorphisms belonging to the chosen group G.

Difficulty in computing d_G

The natural pseudo-distance d_G represents our ground truth.

Unfortunately, d_G is difficult to compute.

Nevertheless, in this talk we will show that d_G can be approximated with arbitrary precision by means of a **DUAL** approach based on persistent homology and G-invariant non-expansive operators.

Reference: P. Frosini, G. Jabłoński, *Combining persistent homology and invariance groups for shape comparison*, Discrete & Computational Geometry, vol. 55 (2016), n. 2, pages 373–409.

G-invariant non-expansive operators

Let us consider the following objects:

- A triangulable space X with nontrivial homology in degree k.
- A set Φ of continuous functions from X to $\mathbb{R}$, that contains the set of all constant functions.
- A topological subgroup G of $\mathrm{Homeo}(X)$ that acts on Φ by composition on the right.
- A subset $\mathscr{F}$ of the set $\mathscr{F}^{\mathrm{all}}(\Phi, G)$ of all G-invariant non-expansive operators from Φ to Φ (GINOs).

The operator space $\mathscr{F}^{\mathrm{all}}(\Phi, G)$

In plain words, $F \in \mathscr{F}^{\mathrm{all}}(\Phi, G)$ means that

1. $F : \Phi \to \Phi$
2. $F(\varphi \circ g) = F(\varphi) \circ g$. ($F$ is a G-operator)
3. $\|F(\varphi_1) - F(\varphi_2)\|_\infty \leq \|\varphi_1 - \varphi_2\|_\infty$. ($F$ is non-expansive)

The operator F is not required to be linear.

Some simple examples of F, taking Φ equal to the set of all continuous functions $\varphi : \mathbf{S}^1 \to \mathbb{R}$ and G equal to the group of all rotations of $\mathbf{S}^1$:

- $F(\varphi) :=$ the constant function $\psi : \mathbf{S}^1 \to \mathbb{R}$ taking the value $\max \varphi$;
- $F(\varphi)$ defined by setting $F(\varphi)(x) := \max\left\{\varphi\left(x - \frac{\pi}{8}\right), \varphi\left(x + \frac{\pi}{8}\right)\right\}$;
- $F(\varphi)$ defined by setting $F(\varphi)(x) := \frac{1}{2}\left(\varphi\left(x - \frac{\pi}{8}\right) + \varphi\left(x + \frac{\pi}{8}\right)\right)$.

Persistent homology

We recall that persistent homology is a theory describing the m-dimensional holes (components, tunnels, voids, ...) of the sublevel sets of a topological space X endowed with a continuous function $\varphi : X \to \mathbb{R}^k$. In the case $k = 1$, persistent homology is described by suitable collections of points called persistence diagrams or, equivalently, by particular functions called persistent Betti number functions. Two such diagrams (or functions) can be compared by a suitable metric d_{match}, called bottleneck (or matching) distance.

The research concerning k-dimensional persistent homology is still at an early stage of development for $k > 1$. Because of this fact, in the rest of this talk we will confine ourselves to consider the case $k = 1$, for which well-established results and algorithms are available.

The pseudo-metric $D^{\mathscr{F}}_{\text{match}}$

For every $\mathscr{F} \subseteq \mathscr{F}^{\text{all}}(\Phi, G)$ and every $\varphi_1, \varphi_2 \in \Phi$ we set

$$D^{\mathscr{F}}_{\text{match}}(\varphi_1, \varphi_2) := \sup_{F \in \mathscr{F}} d_{match}(\rho_k(F(\varphi_1)), \rho_k(F(\varphi_2)))$$

where $\rho_k(\psi)$ denotes the persistent Betti number function of ψ in degree k, while d_{match} denotes the usual bottleneck distance that is used to compare the persistence diagrams associated with $\rho_k(F(\varphi_1))$ and $\rho_k(F(\varphi_2))$.

Proposition

$D^{\mathscr{F}}_{match}$ is a G-invariant and stable pseudo-metric on Φ.

The G-invariance of $D^{\mathscr{F}}_{match}$ means that for every $\varphi_1, \varphi_2 \in \Phi$ and every $g \in G$ the equality $D^{\mathscr{F}}_{\text{match}}(\varphi_1, \varphi_2 \circ g) = D^{\mathscr{F}}_{\text{match}}(\varphi_1, \varphi_2)$ holds.

An equivalence result

We observe that the pseudo-distance $D^{\mathscr{F}}_{\mathrm{match}}$ and the natural pseudo-distance d_G are defined in quite different ways.

In particular, the definition of $D^{\mathscr{F}}_{\mathrm{match}}$ is based on persistent homology, while the natural pseudo-distance d_G is based on the group of homeomorphisms G.

In spite of this, the following statement holds:

Theorem

If $\mathscr{F} = \mathscr{F}^{\mathrm{all}}(\Phi, G)$, then the pseudo-distance $D^{\mathscr{F}}_{\mathrm{match}}$ coincides with the natural pseudo-distance d_G on Φ.

Our main idea

The previous theorem suggests to study $D_{\text{match}}^{\mathscr{F}}$ instead of d_G.

To this end, let us choose a finite subset $\mathscr{F}^*$ of $\mathscr{F}$, and consider the pseudo-metric $D_{\text{match}}^{\mathscr{F}^*}$.

Obviously, $D_{\text{match}}^{\mathscr{F}^*} \leq D_{\text{match}}^{\mathscr{F}}$.

We observe that if $\mathscr{F}^*$ is dense enough in $\mathscr{F}$, then the new pseudo-distance $D_{\text{match}}^{\mathscr{F}^*}$ is close to $D_{\text{match}}^{\mathscr{F}}$.

In order to make this point clear, we need the next theoretical result.

Compactness of $\mathscr{F}^{\mathrm{all}}(\Phi, G)$

The following result holds:

Theorem

If Φ is a compact metric space with respect to the sup-norm, then $\mathscr{F}^{\mathrm{all}}(\Phi, G)$ is a compact metric space with respect to the distance d defined by setting

$$d(F_1, F_2) := \max_{\varphi \in \Phi} \|F_1(\varphi) - F_2(\varphi)\|_\infty$$

for every $F_1, F_2 \in \mathscr{F}$.

Approximation of $\mathscr{F}$

This statement follows:

Corollary

Assume that the metric space Φ is compact with respect to the sup-norm. Let $\mathscr{F}$ be a subset of $\mathscr{F}^{\mathrm{all}}(\Phi, G)$. For every $\varepsilon > 0$, a finite subset $\mathscr{F}^*$ of $\mathscr{F}$ exists, such that

$$\left| D_{match}^{\mathscr{F}^*}(\varphi_1, \varphi_2) - D_{match}^{\mathscr{F}}(\varphi_1, \varphi_2) \right| \leq \varepsilon$$

for every $\varphi_1, \varphi_2 \in \Phi$.

This corollary implies that the pseudo-distance $D_{match}^{\mathscr{F}}$ can be approximated computationally, at least in the compact case.

Our idea in a nutshell

- The natural pseudo-metric d_G can be approximated with arbitrary precision by the pseudo-metric $D_{\mathrm{match}}^{\mathscr{F}^*}$.
- While d_G is usually difficult to compute, $D_{\mathrm{match}}^{\mathscr{F}^*}$ can be efficiently computed by algorithms developed for persistent homology.
- The set $\mathscr{F}^*$ of G-invariant non-expansive operators (GINOs) represents the observer.

We highlight that the set $\mathscr{F}^*$ is not chosen once and for all. To change $\mathscr{F}^*$ means to change the observer and the invariances he/she is interested in. This fact depends on our assumption that only the observer is entitled to decide about shape similarity.

Our basic questions

Assumptions in our model

Mathematical setting and theoretical results

A first step towards the application of our model: GIPHOD

GIPHOD (Group Invariant Persistent Homology On-line Demonstrator)

GIPHOD is an on-line demonstrator, allowing the user to choose an image and an invariance group. GIPHOD searches for the most similar images in the dataset, with respect to the chosen invariance group.
Purpose: to show the use of our theoretical approach for image comparison.
Dataset: 10.000 quite simple grey-level synthetic images obtained by adding randomly chosen bell-shaped functions. The images are coded as functions from $\mathbb{R}^2 \to [0,1]$.

GIPHOD can be tested at `http://giphod.ii.uj.edu.pl`.

Thanks to everyone that will give suggestions for improvement (please send them to `grzegorz.jablonski@uj.edu.pl`*)*

GIPHOD

Joint project with Grzegorz Jabłoński (Jagiellonian University and IST Austria) and Marc Ethier (Université de Saint-Boniface - Canada)

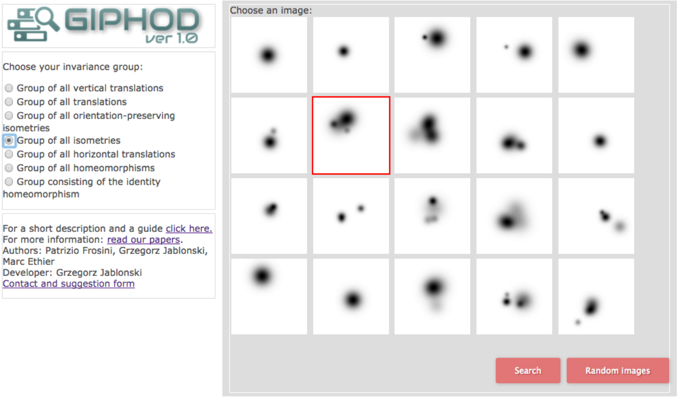

GIPHOD: Examples for the group of isometries

We will now show some results obtained by GIPHOD when the invariance group G is the group of isometries:

Some data about the pseudo-metric $D^{\mathscr{F}}_{match}$ in this case:

- Mean distance between images: 0.2984;
- Standard deviation of distance between images: 0.1377;
- Number of GINOs that have been used: 5.

GIPHOD: Examples for the group of isometries

Here are the five GINOs that we used for the invariance group of isometries:

- $F(\varphi) = \varphi$.
- $F(\varphi) :=$ constant function taking each point to the value $\int_{\mathbb{R}^2} \varphi(\mathbf{x}) \, d\mathbf{x}$.
- $F(\varphi)$ defined by setting

$$F(\varphi)(\mathbf{x}) := \int_{\mathbb{R}^2} \varphi(\mathbf{x}-\mathbf{y}) \cdot \beta\left(\|\mathbf{y}\|_2\right) \, d\mathbf{y}$$

  where $\beta : \mathbb{R} \to \mathbb{R}$ is an integrable function with $\int_{\mathbb{R}^2} |\beta(\|\mathbf{y}\|_2)| \, d\mathbf{y} \leq 1$
  (we have used three operators of this kind).

GIPHOD: Examples for the group of isometries

GIPHOD ver 1.0

For a short description and a guide click here.
For more information: read our papers.
Authors: Patrizio Frosini, Grzegorz Jablonski, Marc Ethier
Developer: Grzegorz Jablonski
Contact and suggestion form

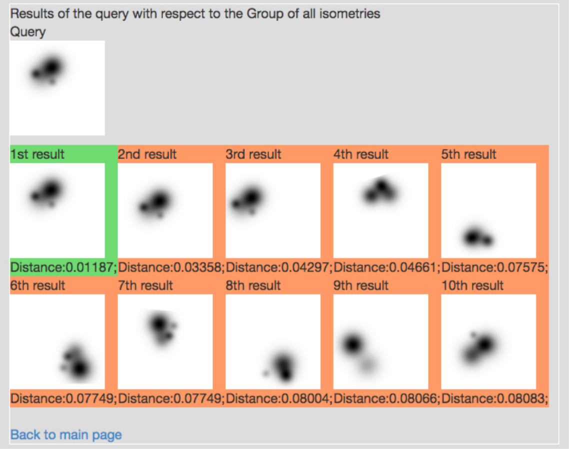

Mean distance: 0.2984; Standard deviation of distance: 0.1377; Number of GINOs: 5.

GIPHOD: Examples for the group of isometries

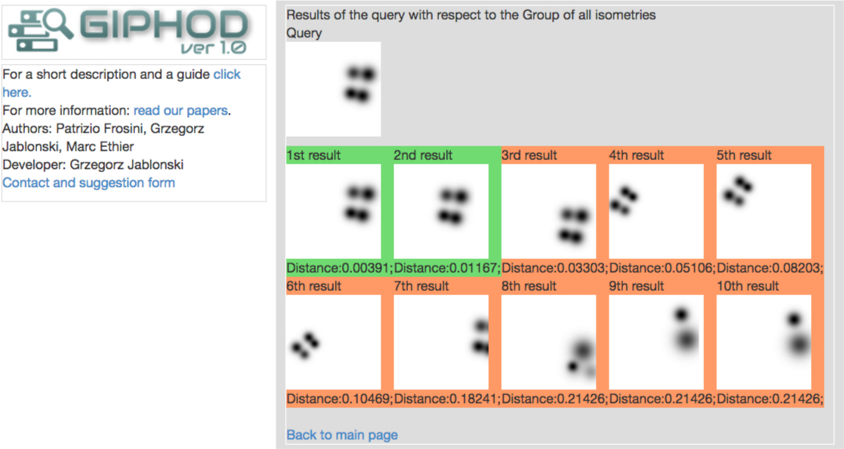

Mean distance: 0.2984; Standard deviation of distance: 0.1377; Number of GINOs: 5.

GIPHOD: Examples for the group of isometries

GIPHOD ver 1.0

For a short description and a guide click here.
For more information: read our papers.
Authors: Patrizio Frosini, Grzegorz Jablonski, Marc Ethier
Developer: Grzegorz Jablonski
Contact and suggestion form

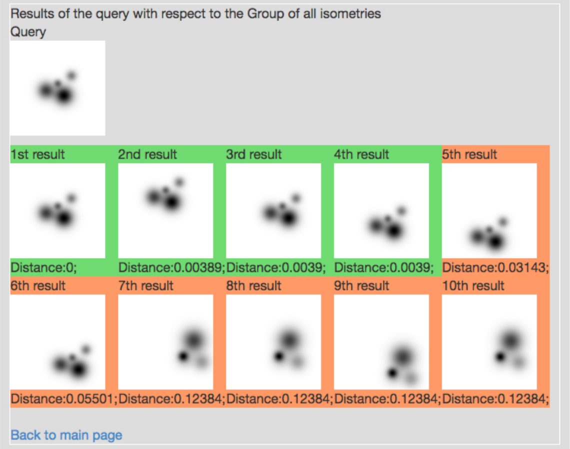

Mean distance: 0.2984; Standard deviation of distance: 0.1377; Number of GINOs: 5.

Conclusions

The model that we have proposed to study is based on the idea that, from the mathematical point of view, a shape should not be considered as a subset or a submanifold of a Euclidean space, but as a quotient of the space Φ of the signals that can be perceived by the chosen observer with respect to the action of a given invariance group G. According to this model, each observer should be represented by a collection of group invariant non-expansive operators acting on Φ. This idea is supported by some formal results showing how the emerging theory of persistent homology could be used to study the approach to shape comparison that we have proposed in this talk. We suggest that this approach could possibly contribute to bridge the semantic gap by means of the framework of topological data analysis.

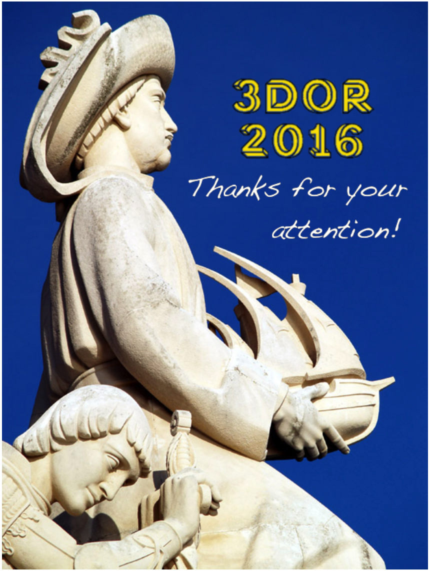
3DOR
2016
Thanks for your
attention!