Eurographics Workshop on 3D Object Retrieval (2015)
I. Pratikakis, M. Spagnuolo, T. Theoharis, L. Van Gool, and R. Veltkamp (Editors)

Morphological analysis of 3D faces for weight gain assessment

D.Giorgi and M. A. Pascali and G. Raccichini and S. Colantonio and O. Salvetti

Institute of Information Science and Technologies, National Research Council, Pisa, Italy

Abstract

In this paper we analyse patterns in face shape variation due to weight gain. We propose the use of persistent homology descriptors to get geometric and topological information about the configuration of anthropometric 3D face landmarks. In this way, evaluating face changes boils down to comparing the descriptors computed on 3D face scans taken at different times. By applying dimensionality reduction techniques to the dissimilarity matrix of descriptors, we get a shape space in which each face is a point and face shape variations are encoded as trajectories in that space. Our first results show that persistent homology is able to identify features which are well-related to overweight, and may help assessing individual weight trends. The research is carried out in the context of the European project SEMEOTICONS, which is developing a multisensory platform which detects and monitors over time facial signs of cardio-metabolic risk.

Categories and Subject Descriptors (according to ACM CCS): I.4.7 [IMAGE PROCESSING AND COMPUTER VISION]: Feature Measurement—Feature representation, Size and shape

1. Introduction

Back in 1942, D'Arcy Wentworth Thompson expressed the importance of investigating biological form in a fully quantitative manner [TW*42]:

The study of form may be descriptive merely, or it may become analytical. We begin by describing the shape of an object in the simple words of common speech: we end by defining it in the precise language of mathematics; and the one method tends to follow the other in strict scientific order and historical continuity.

We may say that D'Arcy Thompson's vision has come true: in the last century, *morphometrics* came of age, as the discipline dealing with the quantative study of form [Rey96]. This was mainly accomplished by applying univariate and multivariate statistics to measures such as linear distances, angles and ratios. In the 1980's, it became clear that a more complex approach to the study of shape was needed, which had to be able to capture the *geometry* of the morphological structures under study and retain its information through the analysis. It was the birth of *Geometric Morphometrics* [Cor93], which quantifies the variation in the shape of anatomical objects using the Cartesian coordinates of anatomical landmarks, after the effects of non-shape variations (translation, rotation, scale) have been factored out. A rich statistical theory for shape analysis supported the analysis of shape variation [DM98, Boo96].

On the other hand, D'Arcy Thompson's dream of a quantitative investigation of shape has been realized also in the discipline of 3D shape analysis and description: over the last decades, computer vision and computer graphics scientists brought fundamental methodological advances for shape quantification, in the form of shape descriptors and similarity measures [BFGS14]. In a sense, we could say that morphometrics belongs to a much broader family of computational methods for quantitative morphological investigation. Therefore, we believe that morphometrics could greatly benefit from the developments in the field of shape analysis and description, especially for 3D data, whose usage was one of the major changes over the last decade.

In this paper, we are interested in quantifying patterns in face shape variation due to weight gain. We take advantage of computational topology, an emerging yet well-established field of research in Computer Graphics [BDFF*08], which gives accurate descriptors of 3D data. This research is carried out in the context of the European project SEMEOTICONS [sem], which is developing a multisensory platform in the form of a mirror. The platform detects and moni-

DOI: 10.2312/3dor.20151055

tors over time facial signs of cardio-metabolic risk, cardiovascular diseases being one of the leading causes of mortality worldwide. The mirror is designed to fit into daily-life settings (the home, the gym, the pharmacy). Therefore, it requires contact-less data aquisition and non-invasive sign detection and analysis. According to a semeiotic model of the face for cardio-metabolic risk [CFG*14], the face signs include 3D morphological face descriptors of overweight and obesity, to be computed on a 3D face model reconstructed from range data acquired by a 3D scanner.

Though it is well known that the face is involved in the process of fat accumulation, there is no consensus in the literature about which are the facial morphological correlates of body fat. An increase in some facial dimensions was observed in a study about the face morphology of obese adolescents [FDT*04]. Moreover, there are studies that show that some geometrical facial features may be related to Body Mass Index and Waist Circumference [LK14]. Those features, measured on 2D images, include Euclidean distances, angles and face areas defined by selected soft-tissue landmarks. We argue that shape changes cannot be attributed to single, local variations in the position of individual landmarks. Also, shape changes involve shifts in the position of landmarks relative to other landmarks. Therefore, we need techniques that enable one to globally analyse the landmark configuration on a face.

We propose the use of Persistent Homology [ELZ02], a technique which grows a space incrementally and analyses the placement of topological events within the history of this growth. We compute Persistent Homology on the configuration of soft-tissue face landmarks, namely on a filtered complex whose nodes correspond to anthropometric landmarks, and edge lengths to their Euclidean or geodesic distance. The output is a shape descriptor (a persistence interval) giving information on the geometry and topology of the landmark structure. As persistence intervals can be efficiently compared using suitable distances, evaluating face changes boils down to comparing the persistence intervals computed on 3D face scans taken at different times.

By applying dimensionality reduction techniques to the matrices of dissimilarities between persistent intervals, we get a shape space in which each face is a point, and face shape variations are trajectories in that space. We experimented on a dataset of synthetic 3D faces simulating weight changes, generated using a parametric morphable model [PKA*09]. Though the research is in its initial phase, our results are promising. By analysing the position of thin and fat people in the shape space, we show that persistent homology is able to identify features which are well-related to overweight. Also, by analysing the shape patterns of single individuals as trajectories in shape space, we show that our technique helps assessing trends in weight change on individuals.

2. Basics in persistent homolgy

Persistent homology analyses the placement of topological events in a growing space: for example, the birth of a connected component and its death when it merges into another component. The lifespan of topological attributes is encoded in a simple descriptor called persistence interval. The aim is to furnish a scale to assess the relevance of topological attributes, under the assumption that longevity is equivalent to significance [ELZ02]. Persistent homology can be useful in analysing 3D shape data and shape changes in particular.

The first concept we need is that of a filtered complex. A complex K is filtered by a filtration $\{K^i\}_{i=0,\dots,n}$ if $K^n = K$ and K^i is a subcomplex of K^{i+1} for each $i = 0,\dots,n-1$. One can think of K as a complex that grows from an initial state K^0 to a final state $K^n = K$. An example is the Rips filtration: if a space X is known through a finite number of samples, for a real number $\varepsilon > 0$ the *Rips complex* $R_\varepsilon(X)$ is the complex whose k-simplexes are the subsets $\{x_0, x_1, \dots, x_k\}$ of X such that $d(x_i, x_j) \leq \varepsilon$ for all pairs x_i, x_j with $0 \leq i, j \leq k$. Whenever $\varepsilon < \varepsilon'$, there is an inclusion $R_\varepsilon(X) \to R_{\varepsilon'}(X)$ that reveals a growing complex.

Given a filtered complex, its topological attributes change through the filtration, since new components appear or connect to the old ones, tunnels are created and closed off, voids are enclosed and filled in, etc. In particular, as for 0-homology, each homology class corresponds to a connected component, and a homology class is born when a point is added, forming a new connected component, thus being a 0-cycle. A homology class dies when two points belonging to different connected components, i.e. they belongs to two different 0-cycles, are connected by a 1-chain, thus becoming a boundary. More formally, given a filtered simplicial complex $\{K^i\}_{i=0,\dots,n}$, the *j- persistent k-th homology group* of K^i can be defined as a group isomorphic to the image of the homomorphism $\eta_k^{i,j} : H_k(K^i) \to H_k(K^{i+j})$ induced by the inclusion of K^i into K^{i+j}. In other words, the *j-persistent homology group* of K^i counts how many homology classes of K^i still survive in K^{i+j}. Persistence represents the life-time of cycles in the growing filtration. The persistent homology of a filtered complex can be represented by a set of intervals, called *persistence intervals*: a persistence interval is a pair (i, j), with $i, j \in \mathbf{Z} \cup \{+\infty\}$ and $0 \leq i < j$, such that there exists a cycle that is completed at level i of the filtration and becomes a boundary at level j.

3. Face description and comparison

Our aim is to define a metric in the space of faces that provides information about face shape changes due to weight gain. We proceed in three steps:

- Represent faces using 23 landmarks $\{l_1, \dots, l_{23}\}$, $l_i \in \mathbf{R}^3$. The landmarks are a subset of Farkas' landmarks, picked up according to the findings in [LK14]. Figure 1 shows the set of landmarks on a template face model [PKA*09].

- Describe the landmark configuration via persistent homology. This requires choosing the filtration (Section 3.1) and computing persistence intervals (Section 3.2).
- Build a metric space on which to analyze patterns of shape variation (Section 4).

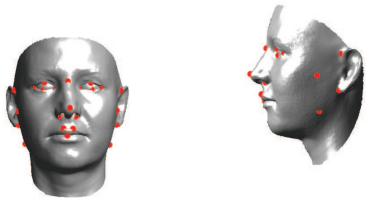

Figure 1: *The 23 landmarks used to represent faces.*

3.1. Building the filtered complex

Building a filtered complex for faces requires defining the simplicial complex and the filtration strategy.

A first choice for the complex could be the flag complex obtained from the landmarks restricted to three dimensions, in which every pair of landmarks has an edge between them, and triangles and terahedra are consequently included. As suggested in [HGK12], the drawback is that the flag complex would not generally have a geometric realization, whereas we do know that our data are inherently three-dimensional, since the landmarks come from human faces which live in the 3D Euclidean space. To preserve the Euclidean nature of data, we preferred a geometrically realizable subcomplex of the flag complex, namely the Delaunay triangulation of the 23 landmarks. We decided to compute the Delaunay triangulation on a template face, namely the Basel Face Model [PKA*09], rather than computing it on individual faces. We do this so that the abstract complex is fixed across all possible subjects, since we want to compare the absolute structure of the landmark configuration, common to all faces. Indeed, if the Dealunay triangulation was calculated separately for each face landmark configuration, the Delaunay triangulations could be slightly different, due to differences in the distances between landmarks for each subject.

Once the complex is fixed, what varies on individual faces are the properties used for filtering the complex. We experimented with different filtrations. The first choice is a Rips filtration using the Euclidean distance between landmarks, similar to what has been done in [HGK12] to study the outcome of clinical orthodontic procedures. All vertices enter at time t_0; the edge between landmarks i and j enters at time $t_{i,j} = M - d_E(l_i, l_j)$ where $d_E(l_i, l_j)$ is the Euclidean inter-landmark distance between landmarks i and j and $M = \max_{q,r} d_E(l_q, l_r)$; triangles and terahedra join the filtration when all of their faces have. Differently from [HGK12], we chage sign to the Euclidean distances. This is done so that

Figure 3: *Euclidean (left) and geodesic (right) distance between two landmarks.*

landmarks that are far apart will have a smaller entry time. Figure 2 shows the process of growth of the complex. Also, differently from [HGK12], we do not consider any normalization of the function across different individuals, as we are not interested in comparing inter-landmark distances within an individual with those same inter-landmark distances in other individuals, but rather on evaluating changes on the same individual in the process of gaining weight.

Moreover, we experimented with another filtration, which is similar to the one above but with geodesic distances instead of Euclidean ones: $t_{i,j} = M - d_G(l_i, l_j)$ where $d_G(l_i, l_j)$ is the geodesic distance between l_i and l_j and $M = \max_{q,r} d_G(l_q, l_r)$ Geodesic distances take into account the intrinsic properties of faces, as they are bound to *walks* on surfaces. Geodesic distances encode different shape features than Euclidean distances: for example, the geodesic distance between the two landmarks in Figure 3 measures the length of the path passing below the chin, whereas the Euclidean distance measures the horizontal distance between the points.

We believe the filtered complexes above are a sensible choice to study the structure of a landmark configuration. At any rate, we underline that different choices for the complex are possible, as discussed in Section 5.

3.2. Computing and comparing persistence diagrams

We computed persistent homology on each filtered complex, with homolgy of dimension 0, 1, and 2. That is, we computed 0-, 1-, 2-dimensional persistence intervals for each face. Figure 4 shows a face and its persistent intervals. Persistence diagrams were compared via the Bottleneck distance [dFL06]. The computations were performed in MATLAB, with code adapted from the program JavaPlex [TVJA11].

4. Analysis of patterns of shape variation

As a longitudinal study on real subjects to monitor weight and face changes was not available, a dataset of synthetic 3D faces simulating weight changes was generated using a parametric morphable model [PKA*09] and used for our first experiments. The morphable model provides specific

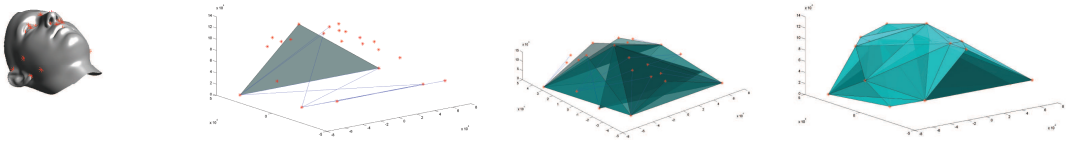

(a) (b) (a) (b)

Figure 2: *From left to right, the growth of the filtered complex, from single vertices associated with landmark points (superimposed on a template face model) to the complete complex.*

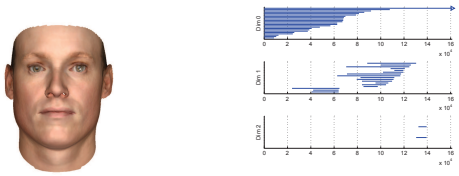

Figure 4: *A face and its persistence intervals.*

Figure 5: *A seed face in our dataset, with fattening simulated in 10 steps.*

parameters for simulating fattening. 25 faces are generated as seeds, and each face is morphed to simulate the process of gaining weight, with 10 equally spaced intervals (Figure 5). This gives a dataset of 250 faces, divided into 10 groups ordered according to increasing fatness. Computing persistent homology on this dataset gives 250×3 persistent intervals, one for each dimension (0-, 1-, 2-degree homology).

4.1. Metric shape space from shape dissimilarities

In statistical shape analysis, the analysis of shape variation is usually carried out in a feature space. Since persistent diagrams are not feature vectors, we took a different route and worked in the space of shape dissimilarities between persistence intervals. Indeed, given a set of faces, the input to our analysis was the dissimilarity matrix whose entries are pairwise distances between face descriptors, namely the Bottleneck distances between persistence intervals. This approach is general and flexible, in that it can be adapted to arbitrary descriptors, other than feature-based (e.g., graphs).

To build the dissimilarity shape space, we computed the Bottleneck distance on the faces of our dataset. This yielded three different dissimilarity matrices, one for each homology dimension (0, 1, 2); each matrix is of dimension 250×250. Then, a dimensionality reduction technique was applied to the matrices of dissimilarities between persistence intervals. We experimented with MultiDimesional Scaling, and got three $250 \times p$ matrices, with $p < 250$. In the analysis that follows, we set $p = 2$. Each row in the matrices represents the coordinates of a face in the lower-dimensional embedding space. Note that different embedding techniques could be used, possibly including non-linear dimensionality reduction techniques such as isometric feature mapping [TDSL00] and Laplacian eigenmaps [NSW08].

As expected, given the intrinsic characteristic of faces, we found that homology of dimension 2 was not significant, whereas homology of degree 0 and 1 proved to be more informative. This can be seen from the observation of the dissimilarity matrices in Figure 6. Therefore, in what follows we will only analyse data pertaining to homology of degree 0 and 1. We first analyse visually the results separately, that is, for 0- and 1-homolgy, and for Euclidean- and geodesic-based filtration. This is done to study the different information they provide (Section 4.2). Then, we analyse quantitatively the results of integrated distances, given by the sums of matrices, in terms of classification rate (Section 4.3).

4.2. Qualitative analysis

The analysis of scatterplots in the embedding space seems to confirm that the proposed technique is able to identify 3D features which are well-related to overweight and obesity. Figure 7 (left) and 8 (left) show the first two embedded MDS coordinates, labeled by fatness level, from 1 to 10, for the filtrations based on Euclidean and geodesic distance, respectively, in dimension 0 and 1. It can be seen that in both cases the subjects are well distributed in the space according to their fatness level. In other words, our technique seems to be able to separate faces of people in different groups. This can be better appreciated in Figure 7 (right) and 8 (right), which show the first two embedded MDS coordinates for a subset of faces, namely thinner people (red), medium people (green) and fatter people (blue) in our dataset.

Since our essential task is the description of morphological change over time on a subject, we must check whether

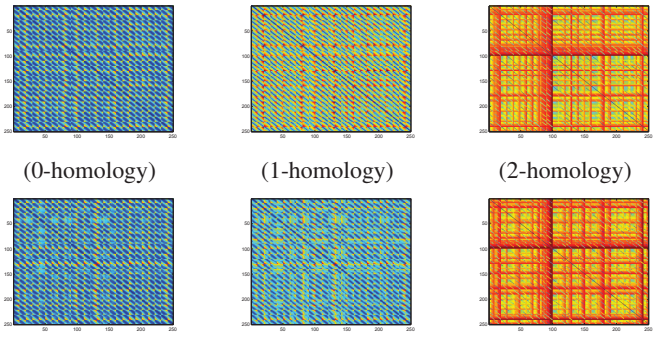

Figure 6: *Dissimilarity matrices between persistence intervals, for the Euclidean (top) and geodesic filtration (bottom).*

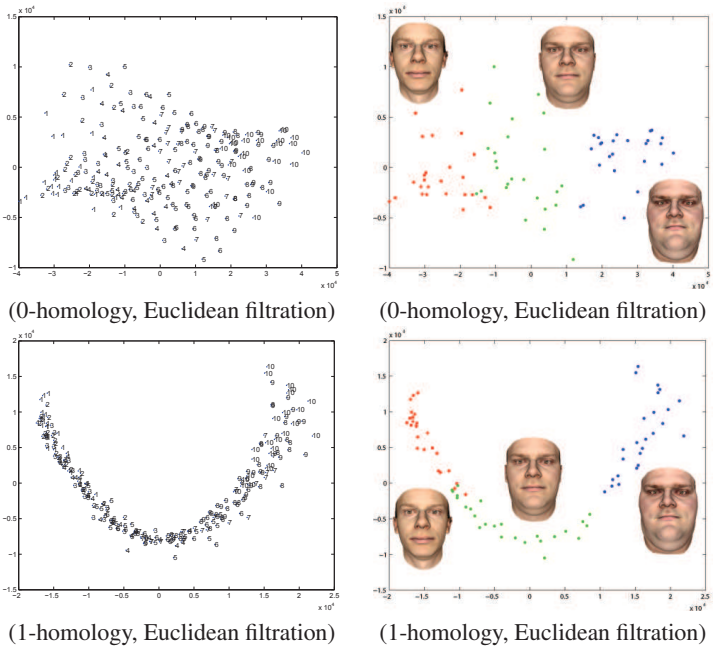

Figure 7: *Scatterplot of the first two embedded MDS coordinates, labeled by fatness level, from 1 to 10, for the Euclidean-based filtrations (left). The scatterplot with only a subset of faces shown (right).*

our technique enables us to discover a trend in a longitudinal study. A way to do this is visualizing the shape patterns of individuals as trajectories [CA13] in the dissimilarity space. Each individual has a trajectoty made of ten consecutive points. For a given trajectory, we can analyse four attributes, namely *location* (the starting and ending points); *size* (the magnitude of the vector between the endpoints); *orientation* (the direction of the vector between the endpoints); and *shape*. In our context, the location depends on the specific, initial traits of each individual. The size is a measure of the difference in shape between the thinnest and the fattest morphing of the individual. The orientation is crucial: a consistent orientation would indicate that our technique is able to detect and encode the process of getting weight. Figure 9 shows the trajectories of three sample faces in our dataset in the embedding space given by the first two coordinates. It is clear that the orientation is consistent, from left to right in accordance with weight gain, but the shape of trajectories do differ, especially for 0-homology. Figure 10 shows the trajectories if only the first embedding coordinate is taken into account. As the trajectories are more homogeneous, it seems that the first coordinate alone is able to identify the trend in fat variation better than the first two coordinates.

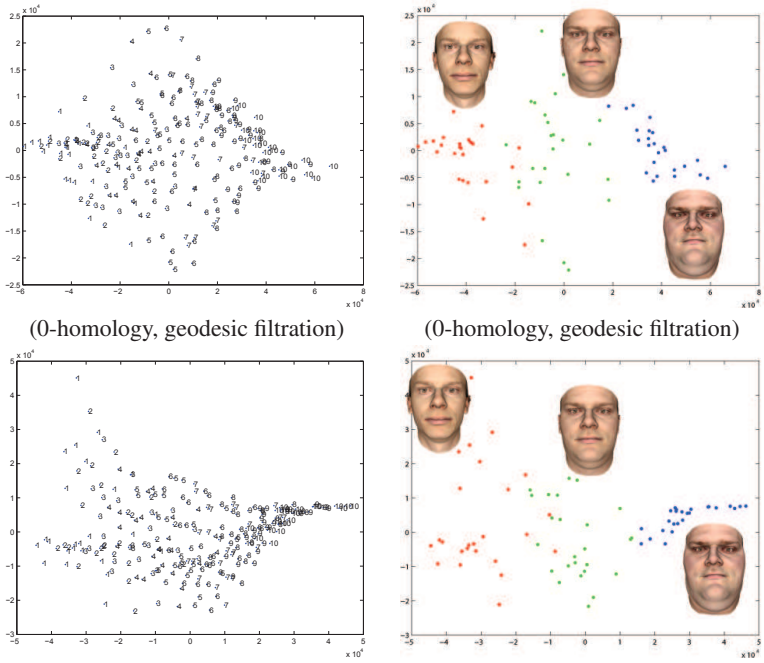

(0-homology, geodesic filtration) (0-homology, geodesic filtration)

(1-homology, geodesic filtration) (1-homology, geodesic filtration)

Figure 8: *Scatterplot of the first two embedded MDS coordinates, labeled by fatness level, from 1 to 10, for the geodesics-based filtrations (left). The scatterplot with only a subset of faces shown (right).*

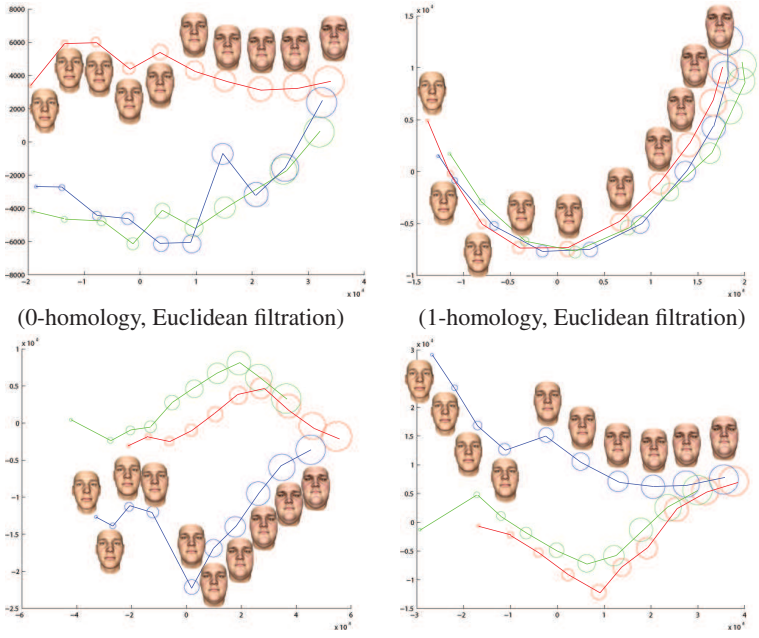

(0-homology, Euclidean filtration) (1-homology, Euclidean filtration)

(0-homology, geodesic filtration) (1-homology, geodesic filtration)

Figure 9: *Trajectories of three sample faces in the space given by the first two embedding coordinates.*

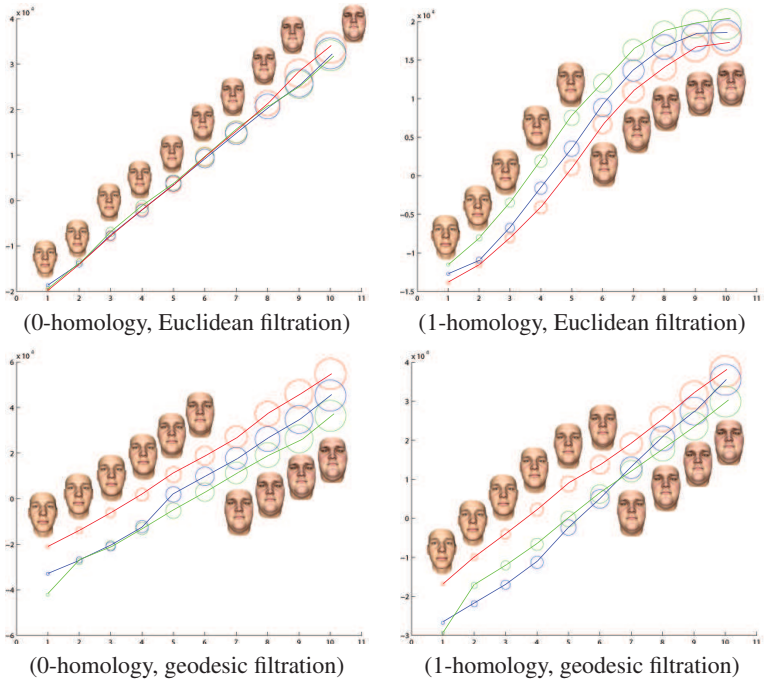

(0-homology, Euclidean filtration) (1-homology, Euclidean filtration)

(0-homology, geodesic filtration) (1-homology, geodesic filtration)

Figure 10: *Trajectories of three sample faces in the space given by the first embedding coordinate.*

4.3. Quantitative analysis

We have seen that persistence intervals based on both the Euclidean distance and the geodesic distance between landmarks are able to quantify shape variation, and that they take different properties into account. This suggests that summing up their information could be beneficial. We can define the distance between two faces as the sum of the Bottleneck distances between their persistence intervals in different dimensions, or with different filtrations. Let us denote $d^0_{E,G}$ ($d^1_{E,G}$) the sum of distances in dimension 0 (dimension 1) obtained with the Euclidean and the geodesic filtration, and $d^{0,1}_E$ ($d^{0,1}_E$) the sum of distances with the Euclidean (geodesic) filtration in both dimensions 0 and 1. Table 1 reports the classification rate on our dataset for the integrated distances above defined. We also evaluated two popular state-of-the-art shape descriptors, namely Shape Distributions (SD) [OFCD02] and Spherical Harmonics (SH) [KFR03]. The classification rate refers to the number of subjects correctly attributed to their group, out of the 10 groups in the dataset, in a leave-one-out experiment. Notice that the classification task is very challenging, since the variation among consecutive groups, in terms of fat gain, is rather small. Therefore, we considered three classification rates, for correct prediction within the first, second and third choice. It can be seen that integrated distances perform better than or comparably to Spherical Harmonics, and significantly better than Shape Distributions. The best performance seems to be provided by the filtration based on geodesic distances, with distances summed over homology dimensions.

	1st	2nd	3rd
$d^0_{E,G}$	42%	63%	**88.4%**
$d^1_{E,G}$	41.2%	64%	83.2%
$d^{0,1}_E$	30.8%	54.4%	77.2%
$d^{0,1}_G$	**42.4%**	**72.2%**	83.6%
SD	22.4%	44.8%	65.6%
SH	38.4%	65.6%	88.0%

Table 1: *Classification rates for integrated distances and competitors. The best rates are marked as bold.*

5. Conclusions

We described an ongoing work in the European project SEMEOTICONS, which is developing a multisensory platform which detects and monitors over time facial signs correlated with cardio-metabolic risk, and gives personalized guidance to individuals to improve their habits. Our contribution to the project is the automatic assessment of weight gain via 3D shape analysis, being obesity and overweight one of the main factors of cardio-metabolic risk. We used persistent homology, which offered a tunable framework for face description and comparison. We qualitatively and quantitatively described the behaviour of our descriptors on a syn-

thetic dataset of 3D faces, and the results seem to be promising. Additional experiments on real face data acquired from volunteers are planned in the future.

A first direction of future research is the study of different filtrations, defined on the whole face rather than on the landmark configuration, for example using Morse filtrations based on curvature or distances from reference points. A landmark-free approach would also give the advantage of not having to pre-compute face landmarks. Although landmarks can be identified with different strategies [OSB13], locating landmarks with the desired accuracy could be difficult on bad-quality face scans.

The future work also include the use of hybrid descriptors which analyse both shape and texture, as for example the hybrid geodesic distance in [BCGS].

Finally, an advantage of our framework is that it is flexible, meaning that it can be adapted to the study of face properties other than weight accumulation, by defining a different, *ad hoc* filtered complex. Therefore, we plan to investigate on the study of other signs, including facial asymmetry.

Acknowledgements

This work is partially supported by the EU FP7-ICT-2013.5.1-611516 Project SEMEOTICONS, started in November 2013.

References

[BCGS] BIASOTTI S., CERRI A., GIORGI D., SPAGNUOLO M.: Phog: Photometric and geometric functions for textured shape retrieval. *Computer Graphics Forum 32*. 8

[BDFF*08] BIASOTTI S., DE FLORIANI L., FALCIDIENO B., FROSINI P., GIORGI D., LANDI C., PAPALEO L., SPAGNUOLO M.: Describing shapes by geometrical-topological properties of real functions. *ACM Computing Surveys 40*, 4 (2008), 12. 1

[BFGS14] BIASOTTI S., FALCIDIENO B., GIORGI D., SPAGNUOLO M.: Mathematical tools for shape analysis and description. *Synthesis Lectures on Computer Graphics and Animation 6*, 2 (2014), 1–138. 1

[Boo96] BOOKSTEIN F. L.: Biometrics, biomathematics and the morphometric synthesis. *Bulletin of mathematical biology 58*, 2 (1996), 313–365. 1

[CA13] COLLYER M. L., ADAMS D. C.: Phenotypic trajectory analysis: comparison of shape change patterns in evolution and ecology. *Hystrix 24*, 1 (2013), 75–83. 5

[CFG*14] COPPINI G., FAVILLA R., GASTALDELLI M., COLANTONIO S., MARRACCINI P.: Moving medical semeiotics to the digitalrealm. semeoticons approach to face signs of cardiometabolic risk. In *HEALTHINF 2014* (2014). 2

[Cor93] CORTI M.: Geometric morphometrics: An extension of the revolution. *Trends in Ecology & Evolution 8*, 8 (1993). 1

[dFL06] D'AMICO M., FROSINI P., LANDI C.: Using matching distance in size theory: A survey. *International Journal of Imaging Systems and Technology 16*, 5 (2006), 154–161. 3

[DM98] DRYDEN I. L., MARDIA K. V.: *Statistical shape analysis*, vol. 4. Wiley Chichester, 1998. 1

[ELZ02] EDELSBRUNNER H., LETSCHER D., ZOMORODIAN A.: Topological persistence and simplification. *Discrete and Computational Geometry 28*, 4 (2002), 511–533. 2

[FDT*04] FERRARIO V. F., DELLAVIA C., TARTAGLIA G. M., TURCI M., SFORZA C.: Soft tissue facial morphology in obese adolescents: a three-dimensional noninvasive assessment. *The Angle orthodontist 74*, 1 (2004), 37–42. 2

[HGK12] HEO G., GAMBLE J., KIM P. T.: Topological analysis of variance and the maxillary complex. *Journal of the American Statistical Association 107*, 498 (2012), 477–492. 3

[KFR03] KAZHDAN M., FUNKHOUSER T., RUSINKIEWICZ S.: Rotation invariant spherical harmonic representation of 3 d shape descriptors. In *SGP* (2003), vol. 6. 7

[LK14] LEE B. J., KIM J. Y.: Predicting visceral obesity based on facial characteristics. *BMC complementary and alternative medicine 14*, 1 (2014), 248. 2

[NSW08] NIYOGI P., SMALE S., WEINBERGER S.: Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry 39*, 1-3 (2008). 4

[OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *ACM Transactions on Graphics (TOG) 21*, 4 (2002), 807–832. 7

[OSB13] O. C., S. U., B. S.: A comparative study of face landmarking techniques. *EURASIP Journal on Image and Video Processing 13* (2013). 8

[PKA*09] PAYSAN P., KNOTHE R., AMBERG B., ROMDHANI S., VETTER T.: A 3d face model for pose and illumination invariant face recognition. IEEE. 2, 3

[Rey96] REYMENT R. A.: An idiosyncratic history of early morphometrics. In *Advances in morphometrics*. Springer, 1996, pp. 15–22. 1

[sem] http://www.semeoticons.eu/. 1

[TDSL00] TENENBAUM J. B., DE SILVA V., LANGFORD J. C.: A global geometric framework for nonlinear dimensionality reduction. *Science 290*, 5500 (2000), 2319–2323. 4

[TVJA11] TAUSZ A., VEJDEMO-JOHANSSON M., ADAMS H.: Javaplex: A research software package for persistent (co)homology. Software available at http://code.google.com/javaplex, 2011. 3

[TW*42] THOMPSON, WENTWORTH D., ET AL.: On growth and form. *On growth and form.* (1942). 1