On the Quantitative Analysis of Craniofacial Asymmetry in 3D

Federico M. Sukno[1], Mario A. Rojas[2,3], John L. Waddington[3] and Paul F. Whelan[2]

[1]Department of Information and Communication Technologies, Pompeu Fabra University, Barcelona, Spain
[2]Centre for Image Processing & Analysis, Dublin City University, Dublin, Ireland
[3]Molecular & Cellular Therapeutics, Royal College of Surgeons in Ireland, Dublin, Ireland

Abstract— We address a systematic evaluation of facial asymmetry from a population of 100 high-quality laser scans, which are first symmetrized and then manipulated to introduce 25 synthetic patterns with a variety of asymmetries. A quantitative evaluation is performed by comparing these known asymmetries with those estimated by different automatic algorithms. Estimation of the actual asymmetries present in the original surface was also addressed. We find that widely used methods based on least-squares minimization not only fail to produce accurate estimates but, in some cases, recover asymmetry patterns that are radically different from the actual asymmetry of the input surfaces, with low or even negative correlation coefficients. A number of alternative algorithms are tested, including landmark-, midline- and surface-based approaches. Among these, we find that the best performance is obtained by a hybrid approach combining surface and midline points, framed within a least median of squares algorithm with weights that decay exponentially with the distance from the midline and an additional term to ensure that the recovered pattern of asymmetry is itself symmetric.

I. INTRODUCTION

Facial symmetry is a recurrent topic in computer vision, with several and diverse applications including face recognition [10], [19], [29], gender classification [30], landmark detection [23], [26], assessment of orthodontic surgery [1], [5] or facial attractiveness [25], to name a few. A particularly interesting application arises from the relation between craniofacial dysmorphology and neuropsychiatric disorders of developmental origin. Among the most evident examples of this relation are the distinctive facial characteristics of patients with Down syndrome [6], but these have also been identified in autism [9], [22], schizophrenia [13], bipolar disorder [11], fetal alcohol syndrome [21], [16], etc. When looking for dysmorphology patterns, facial asymmetry has been highlighted as an especially relevant feature [9], [13], [16]. However, in contrast to the evident dysmorphology in diseases like Down syndrome, dysmorphology in other disorders such as schizophrenia, autism or bipolar disorder can be very subtle to the extent that it cannot be identified by the human eye.

Thus, there is an interest in highly accurate methods to analyze facial asymmetry. All spatial components of the face (right-left, cranio-caudal, anterior-posterior) can be asymmetric [25], hence the analysis needs to be performed in 3D with the help of automatic algorithms. The latter arises as a consequence of the complex asymmetry patterns resulting from a full-3D analysis of the facial surface, but also due to the fact that human perception of asymmetry is orientation-dependant, with a bias toward the vertical direction [27].

Furthermore, one of the added benefits of 3D analysis is that surface information is expressed directly in physical units (e.g. mm), which allows for extracting quantitatively meaningful information. Thus, there exist several methods [1], [2], [4], [5], [9], [17], [20], [28] that provide an estimate of the magnitude of facial asymmetry for each surface point, calculated as the deviation with respect to a perfectly symmetric one, directly expressed in mm. On the other hand, validation of the estimated asymmetry has been largely overlooked and the few attempts in this direction have either focused on qualitative assessment [20] or reported results on single cases [4], [17], which limits their significance.

In this paper we address a systematic evaluation of facial asymmetry from a population of 100 high-quality laser scans, which are first symmetrized and then manipulated to introduce 25 synthetic patterns with a variety of asymmetries. A quantitative evaluation is performed by comparing these known asymmetries with those estimated by different automatic algorithms. We find that widely used methods based on least-squares minimization not only fail to produce accurate estimates but, in some cases, they recover asymmetry patterns that correlate poorly with the actual asymmetries present in the facial surfaces. We also present alternative cost functions and show that these can perform considerably better in terms of estimation accuracy.

The next section provides some basic concepts and notation, together with a review of the most related literature. The different alternatives explored for automatic asymmetry estimation are described in Section III, followed by experiments with synthetic and real asymmetries in Sections IV and V, respectively. Concluding remarks are provided in Section VI.

II. PROBLEM STATEMENT AND RELATED WORK

Estimation of facial symmetry and, more generally, shape symmetry has been extensively studied and an in-depth review of the field is beyond the scope of this work. We will focus on 3D methods that allow obtaining point-wise estimates of facial asymmetry directly in physical units, as these methods make possible to quantitatively assess the accuracy of the provided estimates.

Given a facial surface **S**, the most popular approach to analyze its symmetry is to divide it into left and right hemispheres, so that asymmetry is computed by comparing paired points from both sides [5], [8], [9], [16], [17], [20], [28], [29]. This is typically done as follows:

1. Given a facial surface $\mathcal{S}$, its mirrored version $\mathcal{S}_M$ is obtained by reflection into some arbitrary plane [15];
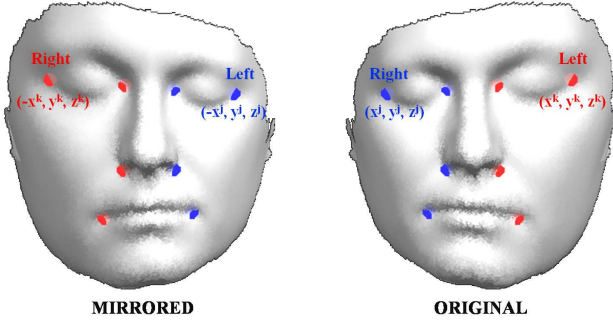
Fig. 1. Example of a facial surface and a mirrored version of it, obtained by inverting the sign of its x-coordinates. Eight left/right landmarks are also displayed, with red and blue colors indicating the side to which they belong to highlight that, after reflection, left and right sides are swapped.

conceptually, this would ideally be the mid-sagittal plane, but in practice that is not strictly required. In this paper we use the plane $x = 0$.

2. Both surfaces $\mathcal{S}$ and $\mathcal{S}_M$ are put into correspondence. Without loss of generality this can be indicated by a transformation T and a reflection of vertex indices, $(\cdot)^{ref}$, applied to $\mathcal{S}_M$ so that $T(\mathcal{S}_M^{ref})$ is as close as possible to $\mathcal{S}$ under the appropriate metric.

3. Asymmetry is computed by (point-wise) comparison between $\mathcal{S}$ and $T(\mathcal{S}_M^{ref})$.

Reflection of vertex indices is necessary to bring $\mathcal{S}_M$ back into point-wise correspondence with $\mathcal{S}$: e.g. let $\mathbf{x}^j = (x^j, y^j, z^j)^T$ be the right outer-eye corner of $\mathcal{S}$; after reflection it becomes $\mathbf{x}_M^j = (-x^j, y^j, z^j)^T$ which shall now correspond to the **left** outer-eye corner of $\mathcal{S}_M$, as illustrated in Fig. 1. The operator $(\cdot)^{ref}$ is a function $f : \mathbb{N} \to \mathbb{N}$ that swaps left and right vertex indices to deal with this issue. We will use $\bar{j}$ to denote the *complement* (paired symmetric index) of the j-th index, i.e.

$$(j)^{ref} = \bar{j}, \qquad (\bar{j})^{ref} = j$$

Some simplification to the above procedure is possible by ensuring that the reflection plane of step 1 coincides with the axis of bilateral symmetry or mid-sagittal plane $\mathbf{p}_S$ [3]. In such case, the registration in step 2 to put $\mathcal{S}$ and $\mathcal{S}_M$ in correspondence is not necessary. For example, in [1], [7], [14] asymmetry is quantified from the difference of landmark distances to the mid-sagittal plane. However, the correct estimation of $\mathbf{p}_S$ is actually analogous to solving the registration in step 2 and pair-wise comparison between $\mathcal{S}$ and $T(\mathcal{S}_M^{ref})$ is a generalization of earlier works based on simplified calculations that can be equivalently obtained from step 3.

Therefore, estimation of $\mathbf{p}_S$ or, more generally, alignment of $\mathcal{S}$ and $\mathcal{S}_M$ becomes a crucial aspect for facial symmetry analysis. An overwhelming majority of studies have addressed this from least-squares cost functions (LMS), typically applying Procrustes methods for shape alignment (in this case constrained to rigid transformations). While least-squared alignment has solid theoretical foundations within

shape theory, it does not necessarily produce the correct alignment for identification of asymmetry. To illustrate this, consider the following example:

- Let $\mathcal{S}_{symm}$ be a facial surface that is perfectly symmetric about the plane $x = 0$. Thus, symmetric paired points (e.g. eye-, nose-, mouth-corners) have the same y and z coordinates and x coordinates with the same magnitude but inverted sign.
- Let us generate $\mathcal{S}_1$ by expanding the left side of $\mathcal{S}_{symm}$ linearly with respect to the coordinates of the x-axis. If we represent surface $\mathcal{S}_{symm}$ by a set of n_V vertices $\{\mathbf{v}_s^i\}_{i=1}^{n_V}$ with $\mathbf{v}_s^i = (x_s^i, y_s^i, z_s^i)^T$ then:

$$\forall (\mathbf{v}_s^i \in \mathcal{S}_{symm}, \mathbf{v}_1^i \in \mathcal{S}_1):$$
$$\mathbf{v}_1^i = \begin{cases} \mathbf{v}_s^i + (\alpha \, x_s^i, 0, 0)^T & \text{if } x_s^i > 0 \\ \mathbf{v}_s^i & \text{otherwise} \end{cases}$$

where α is a constant that controls the degree of expansion.

Surface $\mathcal{S}_1$ is asymmetric, with asymmetry pattern $\mathcal{A}_1$. Note that $\mathcal{A}_1$ is not simply the displacement that we have introduced, as that would imply that points on the left side are asymmetric while points on the right side are not, since $\mathbf{v}_1^i - \mathbf{v}_s^i = 0$ for any $\mathbf{v}_1^i$ on the right side of the face. In other words, the asymmetry pattern must be symmetric for paired points. In essence $\mathcal{A}_1$ is the deviation of $\mathcal{S}_1$ from its closest possible symmetric surface $\mathcal{S}_{1s}$. Thus, we first symmetrize $\mathcal{S}_1$, as follows[1]:

$$\mathcal{S}_{1s} = \frac{\mathcal{S}_1 + M_x^{ref}(\mathcal{S}_1)}{2} \qquad (1)$$

where the operator $M_x^{ref}(\cdot)$ denotes mirror reflection of both coordinates (which is done by means of a sign inversion of the x-coordinates) and vertex indices. Thus:

$$\forall (\mathbf{v}_1^i \in \mathcal{S}_1, \mathbf{v}_{1s}^i \in \mathcal{S}_{1s}) : \mathbf{v}_{1s}^i = \frac{\mathbf{v}_1^i + M_x^{ref}(\mathbf{v}_1^i)}{2}$$

Given $\mathbf{v}_s^i = (x_s^i, y_s^i, z_s^i)^T$, $x_s^i > 0 \Leftrightarrow x_s^{\bar{i}} < 0$, therefore:

$$\mathbf{v}_{1s}^i = \begin{cases} \frac{1}{2}(x_s^i + \alpha \, x_s^i, y_s^i, z_s^i)^T + \frac{1}{2}(-x_s^{\bar{i}}, y_s^{\bar{i}}, z_s^{\bar{i}})^T & \text{if } x_s^i > 0 \\ \frac{1}{2}(x_s^i, y_s^i, z_s^i)^T + \frac{1}{2}(-x_s^{\bar{i}} - \alpha x_s^{\bar{i}}, y_s^{\bar{i}}, z_s^{\bar{i}})^T & \text{otherwise} \end{cases}$$

But $\mathcal{S}_{symm}$ is symmetric with respect to $x = 0$, with $x_s^{\bar{i}} = -x_s^i$, $y_s^{\bar{i}} = y_s^i$ and $z_s^{\bar{i}} = z_s^i$, hence:

$$\mathbf{v}_{1s}^i = \left(x_s^i + \frac{\alpha}{2} x_s^i, \, y_s^i, \, z_s^i\right)^T \qquad \forall i$$

Now, our asymmetry pattern can be calculated as the deviation of $\mathcal{S}_1$ from its symmetrization $\mathcal{S}_{1s}$, that is:

$$\mathcal{A}_1 = \mathcal{S}_1 - \frac{\mathcal{S}_1 + M_x^{ref}(\mathcal{S}_1)}{2} = \frac{\mathcal{S}_1 - M_x^{ref}(\mathcal{S}_1)}{2}$$

which matches the usual definition of asymmetry, as *the difference between a face and a mirrored version of itself*, except for the $\frac{1}{2}$ scaling factor, which is uniform for all points on the surface and does not alter the spatial pattern

---

[1]This is possible because we know by construction that the symmetry plane of $\mathcal{S}_1$ is at $x = 0$.
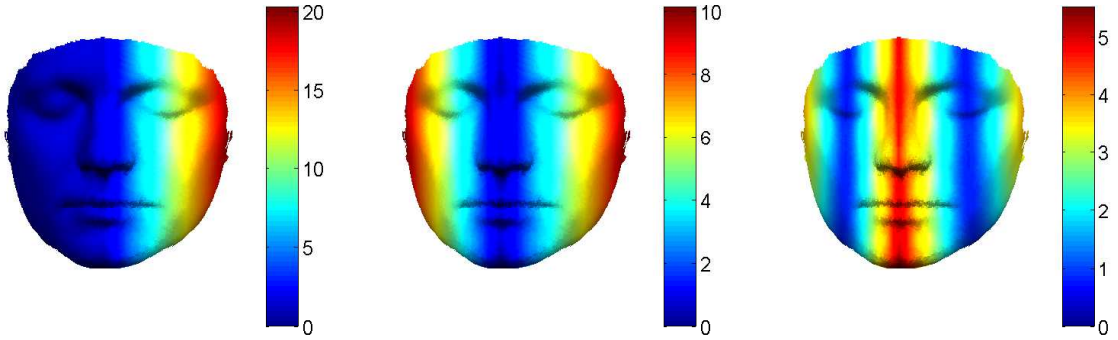
Fig. 2. Example of asymmetric surface $\mathcal{S}_1$ generated by linear expansion of the x-axis on the left facial side. Pattern magnitudes are color-coded on the surfaces according to the respective scales indicated by the color bars, expressed in mm. Left: expansion applied to the left side; Middle: the resulting pattern of asymmetry; Right: the pattern of asymmetry estimated by the LMS approach. Correlation between the actual (middle) and recovered (right) patterns is -0.158 and their average difference over the whole surface is 4.02 mm.

of asymmetry. We can also see that the resulting asymmetry pattern is itself symmetric, as required.

Because we have constructed $\mathcal{S}_1$ by synthetically deforming $\mathcal{S}_{symm}$, we know its symmetry plane (at $x = 0$) and its asymmetry pattern (i.e. $\mathcal{A}_1$ is our ground truth). Now, let us try to estimate the asymmetry of $\mathcal{S}_1$ automatically, under the simplified assumption of having a dense correspondences between left and right paired points[2]. Our estimated asymmetry pattern will be:

$$\hat{\mathcal{A}}_1 = \frac{\mathcal{S}_1 - T(M_x^{ref}(\mathcal{S}_1))}{2}$$

where $T(\cdot)$ is the rigid transformation that aligns $\mathcal{S}_1$ and its reflection $M_x^{ref}(\mathcal{S}_1)$, which in this case should be simply the identity. It is easy to see that the LMS approach will not produce the correct solution, as it estimates T as follows:

$$\arg\min_{\hat{T}} \sum_{\forall \mathbf{v}_S} \|\mathbf{v}_S - \hat{T}(\mathbf{v}_M^{ref})\|^2, \quad \mathbf{v}_S \in \mathcal{S}, \ \mathbf{v}_M^{ref} \in \mathcal{S}_M^{ref}$$

for a generic surface $\mathcal{S}$ and its mirrored-reflection $\mathcal{S}_M^{ref}$. Fig. 2 illustrates the difference between the resulting LMS estimate $\hat{\mathcal{A}}_1$ and the actual asymmetry $\mathcal{A}_1$. Some robust alternatives have been explored to alleviate this shortcoming [2] but we will show in Section IV that they are often insufficient.

Finally, it should be noted that we adopt a template-based representation where left-right correspondences are known, thus allowing to constrain the analysis to rigid transformations. While some methods address superimposition of $\mathcal{S}$ and $\mathcal{S}_M^{ref}$ within non-rigid alignment frameworks [2], [4], [9], this is necessary only when paired correspondences are not available (i.e. the registration algorithm needs both determine the transformation and the correspondences between $\mathcal{S}$ and $\mathcal{S}_M^{ref}$). Under ideal conditions, the non-rigid registration would produce the same correspondences as those provided by the template and the deformation it recovers could be decomposed into rigid and non-rigid parts: the first one is

equivalent to $\hat{T}$ and the second one to the difference that remains after rigid alignment, $\mathcal{S} - \hat{T}(\mathcal{S}_M^{ref})$.

III. ASYMMETRY ESTIMATION

As discussed in the previous section, a key element for the estimation of facial asymmetry is alignment of its surface with a mirrored-reflected version of itself. Different choices of the objective functions can lead to quite different results. Because asymmetry will be estimated as the difference after alignment (on a point-wise basis), we need to determine a suitable strategy for it. With this goal, we address a quantitative comparison to determine the accuracy of asymmetry estimation for different types of alignment.

There are two main elements that need to be chosen to determine the objective function that is used: 1) the type of input data, 2) the error metric. Input data can be a) a sparse set of anatomically corresponding points (landmarks), b) a dense set of surface points (which are also semi-landmarks in our case, since they are in correspondence), c) a specific subset of corresponding points, such as all surface-points corresponding to the facial midline. There are many options for error metrics, among which averages or medians of squares are the most common.

Apart from the LMS approach already described, below we provide a list of the alignment options that will be evaluated. Unless explicitly indicated otherwise, we provide the formulation for the case of using surface points that are in correspondence (hence behaving as quasi-landmarks). Adaptations for the case of using landmarks or midline points are rather trivial and therefore not indicated.

1) Least mean of squares (LMS): This is the most common choice used in shape analysis. We can re-write its cost function from Section II in terms of corresponding indices in the original and mirrored-reflected surfaces, respectively i and $\bar{i}$, as follows:

$$\arg\min_{\hat{T}} \sum_{\forall i} \|\mathbf{v}_S^i - \hat{T}(\mathbf{v}_M^{\bar{i}})\|^2 \tag{2}$$

This minimization can be solved in closed-form but is very sensitive to outliers. Two popular ways to tackle this well

---

[2]In our case this is known by construction; in general it must be estimated, e.g. by template matching or iterative closest point registration.

known weakness are $i)$ to explicitly detect (and discard) outliers or $ii)$ replace the mean with the median, implicitly tolerant to outliers.

2) Confidence-weighted LMS (CW-LMS): Following the work by Van Leemput et al. [18], the alignment between original and mirrored faces is computed by Claes et al [2] using a weighted objective function:

$$\arg\min_{\hat{T}} \sum_{\forall i} w_i \|\mathbf{v}_S^i - \hat{T}(\mathbf{v}_M^{\bar{i}})\|^2 \qquad (3)$$

The objective is to assign higher weights to points that correctly follow the expected distribution of the error and lower weights to points that behave atypically (outliers). Thus, the weights w_i are calculated as the probability of the i-th point to be an inlier divided by the total probability (i.e. including that of being an outlier). A Gaussian distribution of the alignment error $\mathcal{N}(e, 0, \sigma)$ is assumed for inliers and a uniform distribution $1/\lambda$ for outliers. Starting from an initial estimate of $\hat{T}$ using equal (positive) weights for all points, σ and λ are estimated from the resulting errors:

$$e_i = \|\mathbf{v}_S^i - \hat{T}(\mathbf{v}_M^{\bar{i}})\| \qquad (4)$$

$$\sigma = \sqrt{\frac{\sum_{\forall i} w_i e_i}{\sum_{\forall i} e_i}} \qquad (5)$$

$$\lambda = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}K^2\right) \qquad (6)$$

where K is a constant that controls the significance level, which we set to $K = 2$ following [2]. Then we can compute weights w_i for each point and re-estimate $\hat{T}$ from eq. (3):

$$w_i = \frac{\mathcal{N}(e_i, 0, \sigma)}{\mathcal{N}(e_i, 0, \sigma) + \lambda} \qquad (7)$$

Equations (3) to (7) are iterated until convergence of $\hat{T}$.

3) Least median of squares (LMedS): It consists of a slight modification of the objective function in eq. (2), as follows:

$$\arg\min_{\hat{T}} \left(\text{median}\left\{\|\mathbf{v}_S^i - \hat{T}(\mathbf{v}_M^{\bar{i}})\|^2\right\}_{i=1}^{n_V} \right), \qquad (8)$$

where n_V is the number of vertices of surface $\mathcal{S}$. This apparently minor change has a considerable impact [24]. Firstly, the objective function is now strongly tolerant to outliers. Secondly, the minimization cannot be solved in closed-form as before, but needs to be estimated numerically.

4) Distance-weighted LMedS (DW-LMedS): We can also introduce weights in the LMedS scheme. However, since LMedS is intrinsically tolerant to outliers, confidence weights do not make much sense. Instead, we introduce the following modifications: 1) weights that decrease with the distance from the facial midline, based on the observation that facial asymmetry tends to grow as we depart from the midline; 2) symmetry enforcement of the recovered asymmetry patterns.

Intuitively, given a global transformation relating $\mathbf{S}$ and $\mathbf{S}_M^{ref}$, the difference between corresponding points would be amplified proportionally to the distance to the midline. The

inverse of the distance, however, is unbounded. Thus, we use the negative exponential to obtain weights with a controlled behavior that resembles to the inverse distance, leading to the cost function below:

$$\arg\min_{\hat{T}} \left(\text{median}\left\{\| e_{DW}^i \|^2\right\}_{i=1}^{n_V} + \sum_{i=1}^{n_V} \frac{|e_i^2 - e_{\bar{i}}^2|}{n_V} \right) \qquad (9)$$

$$e_{DW}^i = \exp\left(\frac{-1}{\tau} \frac{d_{ML}(\mathbf{v}_S^i)}{\max\left(d_{ML}(\mathbf{v}_S)\right)}\right)\left(\mathbf{v}_S^i - \hat{T}(\mathbf{v}_M^{\bar{i}})\right)$$

$$e_i = \|\mathbf{v}_S^i - \hat{T}(\mathbf{v}_M^{\bar{i}})\|, \qquad e_{\bar{i}} = \|\mathbf{v}_S^{\bar{i}} - \hat{T}(\mathbf{v}_M^i)\|$$

where $d_{ML}(\mathbf{v}_S^i)$ is the distance of the i-th vertex to the facial mid-line and the argument of the exponential is normalized with respect to the maximum value of d_{ML} in the surface. The constant τ controls the rate of decay of the exponential, which nearly vanishes when d_{ML} reaches

$$\frac{3}{2}\tau \max\left(d_{ML}(\mathbf{v}_S)\right)$$

using the standard criterion of 95% decay of the exponential (the division by 2 is because the weights are squared in the cost function). Since the maximum distance is an estimate of half the face width, we can conclude that this weighting scheme considers points within $(3/4)\tau$ times the facial width.

As noted by Klingenberg et al. [15], a drawback of using robust methods to estimate the alignment $T(\cdot)$ is that we may end up with asymmetry patterns that are not symmetric themselves. Thus, the second term in the cost function of eq. (9) penalizes the mean difference of squared errors between points that must have a symmetric pattern. Since both terms in eq. (9) are of comparable magnitude there is no need for an additional constant to weight their contribution. The median in the first term provides robustness to outliers, which is not desired in the second term because the asymmetry pattern should be strictly symmetric.

5) Hemispheres + midline (HM-LMedS and HM-DW-LMedS): As we shall see in the results, alignments based on the facial midline provide an interesting complementarity to those based on all surface points. Indeed, the distance-based weights of DW-LMedS give more relevance to errors near the midline, which results in more accurate estimates of the asymmetry patterns. However, there is a huge imbalance between the number of points in the midline and in the rest of the surface. A simple way to make sure we capture the effects of both the midline and both facial hemispheres is to consider their contribution in separate terms of the cost function. Hence, we update eq. (9) into:

$$\arg\min_{\hat{T}} \left(\text{median}\left\{\| e_{DW}^i \|^2\right\}_{\forall i \notin \mathcal{I}_{ML}} + \right.$$
$$\left. + \frac{1}{2} \text{median}\left\{e_i^2\right\}_{\forall i \in \mathcal{I}_{ML}} + \sum_{i=1}^{n_V} \frac{|e_i^2 - e_{\bar{i}}^2|}{n_V} \right) \qquad (10)$$

where the first term corresponds to the hemispheres and the second term to the midline, defined as the points with no

paired correspondence (or, analogously, those that are the paired points of themselves):

$$\mathcal{I}_{ML} = \left\{ i \in [1; n_V] \,|\, \bar{\bar{i}} = i \right\}$$

Eq. (10) corresponds to HM-DW-LMedS as it allows to introduce distance-based weights, similarly to eq. (9). We can also fix all weights to 1, so that eq. (10) is simply a symmetry-enforced LMedS that considers separately the hemispheres and the midline. We denote the latter as HM-LMedS.

IV. EVALUATION OF ASYMMETRY ESTIMATES

A. Ground truth generation

In order to evaluate how the different alignment options discussed in Section III perform, we need to know the asymmetry patterns contained in our test surfaces. For a perfectly symmetric surface, we can proceed as described in the example from Section II and impose synthetic patterns of asymmetry which will serve as our ground truth.

Obtention of facial surfaces with perfect bilateral symmetry about $x = 0$ can be done in a rather straight-forward manner. Starting from a non-symmetric surface $\mathcal{S}_1$ we simply apply eq. (1) to obtain $\mathcal{S}_{1s}$. All we need for this is a rough estimate of correspondences of the paired points (to do the index-reflection), which can be obtained by template matching or registration. An advantage of proceeding this way is that, although we will introduce a synthetic pattern of asymmetry, we start from the symmetrization of real faces, which allows including the actual variability of facial geometry in our experiments.

We generated 25 synthetic asymmetry patterns, which included linear expansion of x, y and z axes (as in Section II), module expansions (with respect to specific facial landmarks taken as the origin, e.g. nose tip, eye- and mouth-corners), vertical (y axis) displacements of one or both facial sides, horizontal (x axis) displacements of specific regions of the face (eyes, nose, mouth), quadratic distortion and combinations of two of these effects at once (see details in Supplementary Materials). Some examples are provided in Fig. 3.

Note that while application of eq. (1) results in a symmetric surface $\mathcal{S}_{1s}$, the latter is not necessarily the closest possible symmetric surface to $\mathcal{S}_1$. This happens only when the symmetry axis of $\mathcal{S}_1$ coincides with the plane chosen to do the mirroring (in our case, $x = 0$). Thus, in general $\mathcal{S}_{1s}$ could not be used to compute the asymmetry pattern of the original surface $\mathcal{S}_1$, but it still suffices for the purpose of starting from a perfectly symmetric face.

The above highlights an important issue: when generating synthetic patterns of asymmetry, we must avoid introducing any translation of rotation bias, since that would imply changes in the transformation $T(\cdot)$ relating the original and mirror-reflected surfaces (which we assume to be the identity), hence invalidating our asymmetry ground truth.

B. Experimental evaluation

We performed tests on a dataset of 100 symmetrized surfaces, each of them deformed with the 25 synthetic patterns of asymmetry described above. To this end, we started from 100 facial scans acquired by means of a hand-held laser scanner (Polhemus FastSCAN^{TM}, Colchester, VT, USA), corresponding exclusively to healthy volunteers who acted as controls in the context of craniofacial dysmorphology research [12]. Each surface was mapped into a symmetric template with left-right paired correspondences. The resulting template-based representation of each surface was then averaged with its mirrored-reflected version to obtain the symmetrized surface, as in eq. (1).

We evaluated the resulting 25×100 faces having synthetic asymmetry, with the objective of quantitatively assessing the asymmetry estimation accuracy of the different alignment options listed in Section III. Let $\mathcal{S}$ be a test surface with synthesized asymmetry $\mathcal{A}$ and $\mathcal{S}_M^{ref}$ its mirrored and reflected version. Asymmetry is estimated by (point-wise) comparison between $\mathcal{S}$ and $T(\mathcal{S}_M^{ref})$. That is, we obtain a function $\hat{A}$ that assigns a vector in $\mathbb{R}^3$ (the estimated asymmetry) to each vertex of $\mathcal{S}$. We can measure how accurate is this estimation by comparing $\hat{A}$ to the actual asymmetry $\mathcal{A}$ (i.e. the ground truth). While the error per vertex can be displayed color-coded on the surface for a qualitative assessment, more compressed representations are needed for a quantitative comparison. Thus, we used the correlation coefficient ρ between $\hat{A}$ and $\mathcal{A}$, which gives a normalized coefficient of agreement between the estimated and actual patterns of asymmetry. An advantage of using correlation is its independence with respect to the magnitude of asymmetry, as opposed to measures such as the average error over the surface used in [4].

Table I provides the average correlation coefficients obtained for each asymmetry pattern by the different alignment options, subdivided into landmark-, midline- and surface-based approaches. When using landmarks for alignment, we relied on a subset of 22 points distributed at corners, midpoints or salient features of the eyes, nose, mouth and chin, as detailed in [26]. Facial surfaces were represented as triangulated meshes with 16363 vertices each: 8124 for each face-side or hemisphere and 115 for the midline.

DW-HM-LMedS produced the best results, especially if looking at its worse-performing pattern, which showed a reasonably high correlation coefficient. In this sense, apart from DW-HML-LMedS, only landmark-based LMedS and HML-LMedS showed positive correlation coefficients for all tested patterns. This means that, except for these three options, all other alignments tested may lead to the extraction of asymmetry patterns that are radically different from the actual patterns of asymmetry of the input surfaces, with low or even negative correlation coefficients.

Results reported for DW-HM-LMedS in Table I correspond to $\tau = 3/4$, but the influence of this parameter was also explored for $2 \geq \tau \geq 2/3$ (i.e. so that weights vanish for distances from 0.5 to 1.5 times the estimated face width).
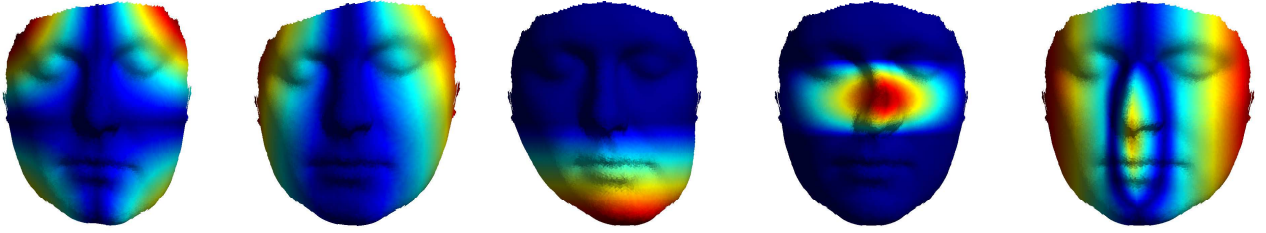
Fig. 3. Examples of a facial surface with 5 asymmetry patterns synthetically generated (corresponding to patterns number 2, 5, 11, 15 and 24 in Table I). The magnitude of asymmetry is color coded on the surface from blue (lowest) to red (highest).

TABLE I

AVERAGE CORRELATION COEFFICIENTS FOR EACH OF THE ASYMMETRY PATTERNS TESTED (FULLY DETAILED IN SUPPLEMENTARY MATERIALS) UNDER DIFFERENT ALIGNMENT OPTIONS TO EXTRACT ASYMMETRY. IN EACH CELL, COLORS WERE USED TO HIGHLIGHT NEARLY PERFECT CORRELATION IF $\rho > 0.9$ (BLUE), WEAK CORRELATION IF $\rho < 0.5$ (ORANGE) OR INVERTED CORRELATION IF $\rho < 0$ (RED).

Asymm Patterns		Landmarks			Midline		Whole surface				Mixed	
		LMS	LMedS	DW-LMedS	LMS	LMedS	LMS	CW-LMS	LMedS	DW-LMedS	HM-LMedS	HM-DW-LMedS
Linear expansions of individual axes (with the nose tip at the origin)	1	0.92	0.95	0.94	1.00	1.00	−0.15	−0.17	−0.93	−0.92	0.82	0.87
	2	0.87	1.00	1.00	1.00	1.00	0.90	0.90	0.97	0.98	0.98	0.99
	3	0.73	0.99	1.00	1.00	1.00	0.70	0.77	0.95	0.95	0.96	0.96
Module expansions with respect to a given landmark	4	0.95	1.00	1.00	1.00	1.00	0.67	0.70	0.93	0.94	0.95	0.97
	5	0.95	0.99	1.00	1.00	1.00	0.48	0.50	0.87	0.86	0.93	0.94
	6	0.75	0.96	0.95	0.99	0.99	0.64	0.67	0.85	0.88	0.90	0.93
	7	0.93	0.98	0.98	0.99	0.99	0.62	0.64	0.84	0.88	0.92	0.94
Combinations of vertical shift (y axis) and expansion of x or z axis	8	0.54	0.99	1.00	1.00	1.00	0.61	0.65	0.94	0.95	0.95	0.96
	9	0.94	0.99	0.99	1.00	1.00	0.46	0.47	0.86	0.43	0.92	0.94
	10	0.89	0.94	0.95	1.00	1.00	0.51	0.51	−0.05	−0.05	0.88	0.90
Linear or rotational horizontal shifts of the upper or lower face	11	0.65	0.97	0.97	−0.13	1.00	0.84	0.92	1.00	1.00	1.00	1.00
	12	0.95	1.00	1.00	−0.33	1.00	0.65	0.71	1.00	1.00	1.00	1.00
	13	0.84	1.00	1.00	−0.01	1.00	0.92	0.97	1.00	1.00	1.00	1.00
	14	0.67	1.00	1.00	−0.14	1.00	0.60	0.64	1.00	1.00	1.00	1.00
Progressive horizontal shifts of specific regions (e.g. mouth, nose, eyes)	15	0.76	1.00	1.00	−0.18	1.00	0.88	0.94	1.00	1.00	1.00	1.00
	16	0.61	1.00	1.00	−0.04	1.00	0.92	0.97	1.00	1.00	1.00	1.00
	17	0.83	1.00	1.00	−0.47	1.00	0.76	0.84	1.00	1.00	1.00	1.00
Quadratic distortions limited both vertically and horizontally from specific landmarks	18	0.12	0.98	0.97	−0.25	1.00	0.82	0.90	1.00	1.00	1.00	1.00
	19	0.37	1.00	1.00	0.07	1.00	0.91	0.97	1.00	1.00	1.00	1.00
	20	−0.53	0.27	0.09	−0.44	−0.42	0.97	1.00	1.00	1.00	0.98	0.97
	21	0.87	0.85	0.96	0.83	0.60	0.53	0.56	0.78	0.65	0.93	0.94
	22	0.72	1.00	1.00	−0.51	−0.33	0.84	0.91	1.00	1.00	0.94	0.94
Combinations of quadratic distortion and axes expansion	23	0.70	0.62	0.75	0.27	0.28	0.92	0.92	0.97	0.97	0.99	0.99
	24	0.95	0.85	0.86	0.96	0.95	−0.07	−0.18	−0.77	−0.76	0.67	0.84
	25	−0.08	0.78	−0.38	0.92	0.92	0.27	0.29	0.89	0.88	0.90	0.91

In this range, correlation coefficients showed only small variations, bounded within 10% of the values indicated in Table I.

V. TESTS WITH REAL ASYMMETRIES

In previous sections we showed that different choices for aligning a shape with its mirrored version may result in significantly different asymmetry patterns, even though we are analyzing the same facial surface and, therefore, there should be a unique pattern of asymmetry. As our tests were carried out on surfaces with synthesized asymmetry, we were able to evaluate the success of the different options. In this section we address the following question: will we also recover significantly different patterns of asymmetry in real facial scans? To this end, we used the same population of 100 scans as in the previous section, but without modifications to the acquired surfaces, so that they retain their original asymmetry.

It is evident that in this situation we do not know the actual asymmetry, hence some qualitative assessment might be helpful. With this in mind we start by providing an example where different alignment methods produce radically different results, as shown in the first row of Fig. 4: from the three alignment options tested, the first one on the left recovers a pattern of asymmetry that is rather scattered over the whole face, the middle-one finds a dominant asymmetry at the outer part of the cheeks and the last one (on the right) suggests that most of the asymmetry is concentrated on the lower face. Because symmetry perception is known to be orientation dependant [27], we have also re-aligned the recovered symmetry plane with the vertical direction of the image, so that the displayed snapshots enforce the
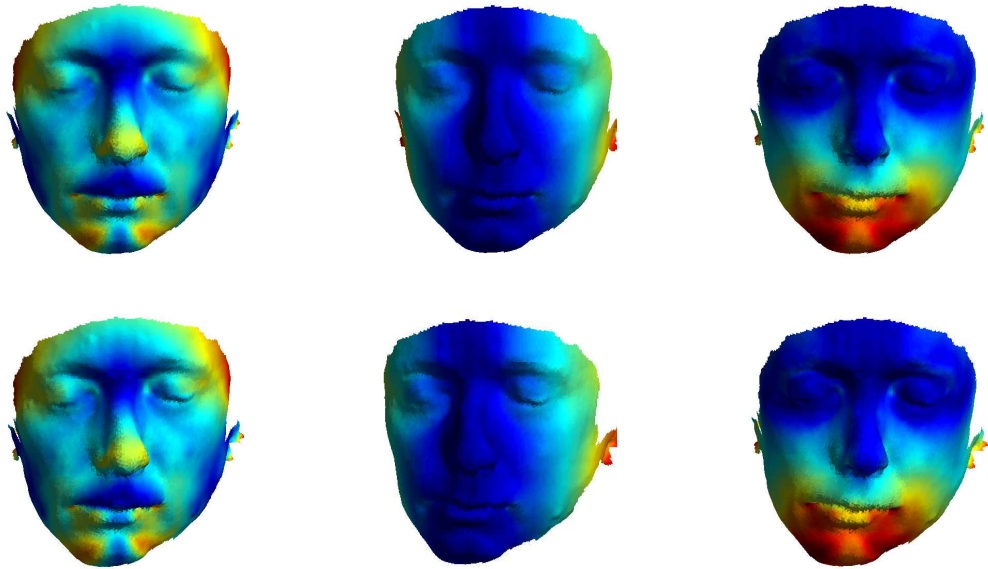
Fig. 4. Three different patterns of asymmetry recovered from the same input face: Landmark-based LMS (left), midline-based LMS (middle) and DW-HM-LMedS (right). The top-row shows the test surface with the recovered asymmetry color-coded on the surface. The bottom-row shows the recovered asymmetry magnified by a factor of 2.

(perceived) validity of each of the recovered patterns (i.e. if the shapes were rotated differently, we may perceive that a different pattern should have been recovered).

Note that each of the patterns recovered in Fig. 4 could have been used as a synthetic pattern; it would then become the ground truth against which we test, evidently favoring one of the alignment methods in particular. Therefore, none of the 3 patterns can be regarded as incorrect at this point: each of them could be correct with respect to a different symmetric surface, namely, the symmetric surface to which we have to apply the recovered asymmetry pattern to obtain the test shape in Fig. 4.

As mentioned in Section IV-A, such ambiguities can be explained due to a translation or rotation bias. This can be clearly seen in the lower part of Fig. 4: the second row of the figure shows a magnification of the recovered patterns of asymmetry by a factor of 2. With this magnification, we clearly see that the pattern in the central column, recovered by LMS alignment of the facial midline, contains a rotation bias, i.e. the head seems to be rotated to its right-hand side (viewers left) making the snapshot depart from a frontal shot toward a half-profile view.

From the discussion above we are able to discard the middle pattern as incorrect due to rotation bias. Unfortunately, deciding whether the left-most or right-most pattern in Fig. 4 is the correct one is not trivial, at least from visual inspection. On the other hand, our experiments with synthetic patterns, on which we were especially careful of avoiding bias, indicate that DW-HM-LMedS is typically more accurate than LMS, suggesting that the pattern on the right is the best asymmetry estimate.

We conclude this section with some quantitative results. Although we cannot compare the obtained patterns to ground truth, we can measure how similar are the recovered patterns: this was done by computing the correlation coefficients between each pair of patterns recovered by different alignment methods. Averages over 100 facial scans showed correlations between 0.3 and 1.0; for approximately half of the scans there were at least two of the recovered patterns showing a negative correlation.

VI. CONCLUSIONS

This paper presents a quantitative evaluation of different methods to estimate the bilateral asymmetry of 3D facial scans. To this end, a variety of synthetic asymmetry patterns are used to deform symmetrized surfaces obtained from 100 facial scans. We find that different estimation methods can produce large variations in the recovered patterns and that the widespread use of least-squares approaches might produce patterns that correlate poorly with the actual asymmetries present in the facial surfaces. A number of solutions were tested, including both existing and novel alternatives. Among the latter, an LMedS approach with exponential decay from the midline and symmetry enforcement (HM-DW-LMedS) was shown to produce estimates highly correlated with the actual asymmetry for all tested patterns.

Another interesting result is that landmark- or surface-based estimates (within the same algorithmic choice) can produce significantly different results. In contrast to current trends that favor surface-based methods, these produced slightly worse results than landmark-based methods. The reason seems related to the excessive weight of peripheral facial regions such as the cheeks, due to their comparatively larger area, as well as the fact that asymmetry tends to increase as we depart from the midline. On the other hand, landmarks are typically placed rather centrally on the facial surface, attached to its major features (e.g. eyes, nose,

mouth), but the exact choice of what points are considered is often arbitrary, which can impact on the stability of the estimated asymmetry.

The above provides some intuition about our design of HM-DW-LMedS. The exponential decay on moving away from the midline provides an interesting trade-off, assigning more weight near the facial features typically covered by landmarks but still including the whole surface in the cost function. The exponential decay can also be beneficial with respect to the exact area that is considered to be within the facial region: in general, its boundaries are difficult to determine and prone to higher errors than the central features of the face; hence, assigning them lower weights can improve the robustness of the algorithms.

Finally, we also performed experiments to estimate the asymmetry of the original facial scans. While in this case we cannot measure the estimation error (as the ground truth is unknown), we did confirm important differences in the asymmetry patterns that were extracted by the different methods.

ACKNOWLEDGMENTS

REFERENCES

[1] M. Alqattan, J. Djordevic, A.I. Zhurov, and S. Richmond. Comparison between landmark and surface-based three-dimensional analysis of facial asymmetry in adults. *European Journal of Orthodontics*, doi: 10.1093/ejo/cjt075, 2013.

[2] P. Claes, M. Walters, D. Vandermeulen, and J.G. Clement. Spatially-dense 3D facial asymmetry assessment in both typical and disordered growth. *Journal of Anatomy*, 219(4):444–455, 2011.

[3] B. Combes, R.H. Hennessy amd J.L. Waddington, N. Roberts, and S. Prima. Automatic symmetry plane estimation of bilateral objects in point clouds. In *Proc. IEEE Int. Conf. on Computer Vision and Pattern Recognition, Anchorage, AK, USA*, pages 1–8, 2008.

[4] B. Combes and S. Prima. New algorithms to map asymmetries of 3D surfaces. In *Proc. 11th Int. Conf. Medical Image Computing and Computer Assisted Intervention, New York, NY, USA*, pages 17–25, 2008.

[5] J. Djordjevic, A.M. Toma, A.I. Zhurov, and S. Richmond. Three-dimensional quantification of facial symmetry in adolescents using laser surface scanning. *European Journal of Orthodontics*, 36(2):125–132, 2014.

[6] V.F. Ferrario, C. Dellavia, G. Serrano, and C. Sforza. Soft tissue facial angles in Downs syndrome subjects: a three-dimensional non-invasive study. *European Journal of Orthodontics*, 27(4):355–362, 2005.

[7] V.F. Ferrario, C. Sforza, C. Dellavia, G.M. Tartaglia, A. Colombo, and A. Caru. A quantitative three-dimensional assessment of soft tissue facia asymmetry of cleft lip and palate adult patients. *Journal of Craniofacial Surgery*, 14(5):739–746, 2003.

[8] M.H. Hajeer, A.F. Ayoub, and D.T. Millet. Three-dimensional assessment of facial soft-tissue asymmetry before and after orthognatic surgery. *British Journal of Oral and Maxilofacial Surgery*, 42:396–404, 2004.

[9] P. Hammond, C. Forster-Gibson, A.E. Chudley, T.J. Hutton, S.A. Farreal, J. McKenzie, J.J.A. Holden, and M.E.S. Lewis. Face-brain asymmetry in autism spectrum disorders. *Molecular Psychiatry*, 13:614–623, 2008.

[10] J. Harguess and J.K. Aggarwal. A case for the average-half-face in 2D and 3D for face recognition. In *Proc. IEEE Int. Conf. on Computer Vision and Pattern Recognition, Miami, FL, USA*, pages 7–12, 2009.

[11] R.H. Hennessy, P.B. Baldwin, D.J. Browne, A. Kinsella, and J.L. Waddington. Frontonasal dysmorphology in bipolar disorder by 3D laser surface imaging and geometric morphometrics: Comparison with schizophrenia. *Schizophrenia Research*, 122(1-3):63–71, 2010.

[12] R.H. Hennessy, A. Kinsella, and J.L. Waddington. 3D laser surface scanning and geometric morphometric analysis of craniofacial shape as an index of cerebro-craniofacial morphogenesis: initial application to sexual dimorphism. *Biological Psychiatry*, 51(6):507–514, 2002.

[13] R.H. Hennessy, A. Lane, A. Kinsella, C. Larkin, E. O′Callaghan, and J.L. Waddington. 3D morphometrics of craniofacial dysmorphology reveals sex-specific asymmetries in schizophrenia. *Schizophrenia Research*, 67:261–268, 2004.

[14] C.S. Huang, X.Q. Liu, and Y.R. Chen. Facial asymmetry index in normal young adults. *Orthodontics & Craniofacial Research*, 16(2):97–104, 2013.

[15] C.P. Klingenberg, M. Barluenga, and A. Meyer. Shape analysis of symmetric structures: quantifying variation among individuals and asymmetry. *Evolution*, 56(10):1909–1920, 2002.

[16] C.P. Klingenberg, L. Wetherill, J. Rogers, E. Moore, R. Ward, I. Autti-Ramo, A. Fagerlund, S.W. Jacobson, L.K. Robinson, H.E. Hoyme, S.N. Mattson, T.K. Li, E.P. Riley, T. Foroud, and CIFASD Consortium. Parental alcohol exposure alters the patterns of facial asymmetry. *Alcohol*, 44(7-8):649–657, 2010.

[17] V. Krajicek, J. Dupej, J. Veleminska, and J. Pelikan. Morphometric analysis of mesh asymmetry. *Journal of WSCG*, 20:65–72, 2012.

[18] K. Van Leemput, F. Maes, D. Vandermeulen, A. Colchester, and P. Suetens. Automated segmentation of multiple sclerosis lesions by model outlier detection. *IEEE Transactions on Medical Imaging*, 20(8):677–688, 2001.

[19] Y. Liu, K.L. Schmidt, J.F. Cohn, and S. Mitra. Facial asymmetry quantification for expression invariant human identification. *Computer Vision and Image Understanding*, 91:138–159, 2003.

[20] P. Meyer-Marcotty, A. Stellzig-Eisenhauer, U. Bareis, J. Hartmann, and J. Kochel. Three-dimensional perception of facial asymmetry. *European Journal of Orthodontics*, 33:647–653, 2011.

[21] T.E.M. Mutsvangwa, E.M. Meintjes, D.L. Viljoen, and T.S. Douglas. Morphometric analysis and classification of the facial phenotype associated with fetal alcohol syndrome in 5- and 12-year-old children. *American Journal of Medical Genetics*, 152A(1):32–41, 2010.

[22] H.M. Ozgen, J.W. Hop, J.J. Hox, F.A. Beemer, and H. van England. Minor physical anomalies in autism: a meta-analysis. *Molecular Psychiatry*, 15:300–307, 2010.

[23] N. Perakis, G. Passalis, T. Theoharis, and I.A. Kakadiaris. 3D facial landmark detection under large yaw and expression variations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(7):1552–1564, 2013.

[24] P.J. Rousseeuw. Least median of squares regression. *Journal of the American Statistical Association*, 79(388):871–880, 1984.

[25] C. Sforza, A. Laino, G. Grandi, L. Pisoni, and V.F. Ferrario. Three-dimensional facial asymmetry in attractive and normal people from childhood to young adulthood. *Symmetry*, 2010(2):1925–1944, 2010.

[26] F.M. Sukno, J.L. Waddington, and P.F. Whelan. Asymmetry patterns shape contexts to describe the 3D geometry of craniofacial landmarks. *Computer Vision, Imaging and Computer Graphics - Theory and Applications. Communications in Computer and Information Science*, 458:19–35, 2014.

[27] M.S. Treder. Behind the looking-glass: a review on human symmetry perception. *Symmetry*, 2(3):1510–1543, 2010.

[28] T.J. Verhoeven, J.W. Nolte, T.J.J. Maal, S.J. Berge, and A.G. Becking. Unilateral condylar hyperplasia: a 3-Dimensional quantification of asymmetry. *PLOS ONE*, 8(3):e59391, 2013.

[29] Y. Wang, J. Liu, and X. Tang. Robust 3D face recognition by local shape difference boosting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(10):1858–1870, 2010.

[30] B. Xia, B. Ben Amor, H. Drira, M. Doudi, and L. Ballihi. Gender and 3D facial symmetry: whats the relationship? In *Proc. 10th IEEE Int. Conf. on Automatic Face and Gesture Recognition, Shanghai, China*, 2013.