

3D Landmark Localisation

Luke Gahan, Supervised by Prof. Paul F. Whelan

Abstract—In this work we examine in detail the landmark localisation algorithm proposed by Gupta et al. [9]. This algorithm automatically localises 10 facial fiducial points using both texture and range images, in conjunction with anthropometric information. For six of the landmarks detected a modified version of the EBGM technique developed by Wisott et al. is used [8]. The landmark localisation performance of this system is determined as the original authors only provide 2D standard deviation results. While the nose tip localisation performs poorly with a 3D mean error of 6.15mm, the remaining landmarks are all localised with an error of under 3.35mm, with the outer eye corner and mouth corner detections performing particularly well. The influence of the inclusion of texture and/or range information when localising landmarks is examined. It is determined that best performance is achieved when both texture and range information is used. Finally, the relative localisation performance improvement achieved by using the modified EBGM technique determined. For all landmarks examined the use of EBGM improves localisation.

I. INTRODUCTION

Facial landmark localisation is the primary step in a number of computer vision systems including, facial recognition, facial pose estimation, medical diagnostics and multimedia applications. Historically most landmark localisation algorithms have used standard 2D images. Such systems, no matter how accurate, are always going to be limited by the fact that they are operating on dimensionally reduced representations of 3D objects. A significant amount of extra information about the human face is contained in the 3D spatial dimension.

A number of different approaches have been taken with regard to localising facial landmarks in 3D images. Geometry based techniques have received a good deal of attention. Segundo et al. present an effective system which uses surface classification techniques in order to localise landmarks [1]. The authors record a 3D localisation error of under 10mm for 90% of images in their test set. Creusot et al. combine machine learning and a large number of geometric techniques in their system [2]. The authors note that while this system does not out perform others in terms of precision it does perform quite well in terms of robustness. Since the algorithm used is not sequential in nature, a failure to detect certain landmarks does not influence the localisation of subsequent landmarks. This system provides a framework for landmark localisation and leaves potential for future improvement.

Zhao et al. present a statistical model based approach in [3]. This system works well in challenging situations where there is facial occlusion and/or very expressive faces. This system learns the spatial relationships between different landmarks and uses this in conjunction with local texture and range information. The authors use PCA to create a statistical facial feature map. This is essentially a combination of individual geometry (landmark coordinates), shape (range images) and

texture (texture images) models. The authors report a mean 3D error rate of below 4.65mm for all 15 facial landmarks.

Perakis et al. use local shape descriptors to localise facial landmarks [4], [5]. These local shape descriptors characterise the shape profile at a given landmark. By evaluating the shape index at a landmark in a number of training images a model can be constructed. These descriptors are generated by examining the principal curvature and spin image at a landmark. A facial landmark model is then created. This is used to constrain the relative locations of detected landmarks. Models are also created for the left and right hand side of the face. These are used to deal with profile or semi-profile faces. The systems achieves relatively good results with a mean 3D error of below 5.58mm for all 8 landmarks.

One particular approach which has received increased attention in recent years is the use of Gabor filters for facial landmark localisation. While an in depth discussion of Gabor filter theory is outside the scope of this work, an excellent tutorial is provided in [6]. Jahanbin et al. use Gabor filter banks for landmark localisation in [7]. This technique implements the same landmark localisation procedure as Wiscott et al. used in their Elastic Bunch Graph Match system (without the elastic constraint) [8]. While the authors do not present in depth results in this particular paper, it does serve as a basis for later work carried out by the same research group [9].

This particular system combines curvature detection, Gabor filter and expert knowledge of the human face to localise landmarks. Anthropometric information based on the work carried out by Farkas et al. in the medical field [10]. This information plays a vital role in establishing a sensible search region which is then examined to further improve the accuracy of localisation. For the detection of inner eye corners and mouth corners the authors used a technique which they term 2D + 3D EBGM which involves using the techniques developed by Wiscott et al. on both range and portrait images. For the outer eye corners the 2D EBGM is used rather than 2D + 3D. The authors state the reason for this is that "the outer corners of the eyes do not have distinct surface curvature characteristics". The authors of this paper only publish the 2D standard deviation error results detected. Even though these are impressive it does mean that the algorithm cannot be benchmarked against state of the art techniques.

In this work the method developed by Gupta et al. is examined in depth. The mean 3D error is evaluated with a view to determining the systems performance in terms of actual localisation rather than just standard deviation. The influence of the inclusion of texture information on landmark localisation is also examined. For the six landmarks where EBGM is used for localisation, the results are compared for 2D (texture), 3D (range) and 2D + 3D (range and texture) EBGM. Finally the relative performance improvement obtained by



Fig. 1: 10 facial landmarks

using EBGGM is evaluated.

II. DATABASE

The Texas 3DFR database is used for the development and testing of this system [9], [11], [12]. This is the same database used by the original authors. It contains 1149 high resolution (751 x 501 pixels, 0.32 mm per pixel) pairs of portrait and range images. The database contains 118 healthy adult subjects. 25 facial landmarks have been manually located. Both range and portrait images were acquired simultaneously using a regularly calibrated stereo vision system. Median filtering and bi-cubic interpolation are used to remove impulse noise and large holes. Finally Gaussian smoothing ($\sigma = 1$ pixel) is applied to reduce noise levels.

III. AUTOMATIC LANDMARK LOCALISATION USING ANTHROPOMETRIC INFORMATION

The landmark localisation procedure carried out remains as faithful as possible to the method developed by Gupta et al. [9]. Generally speaking the algorithm first uses curvature information to detect an approximate location for a particular landmark. Using anthropometric information a search region is defined around this approximation. The search region is then evaluated using some other technique, as described below. The 10 landmarks localised are shown in figure 1.

For a more detailed description of the development and implementation of this algorithm please see Appendix B.

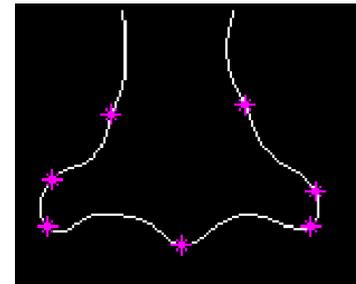
A. Nose Tip (prn)

The Iterative Closest Point (ICP) algorithm is used to register each face in the database to a frontal template face. These aligned images are used in all subsequent steps. Once all images have been aligned the manually localised tip of the template face is taken as an approximate location for tip of the nose in all images. A window of 96 mm x 96mm is then defined around this approximated nose tip. Since all faces have been frontally aligned, the actual nose tip is present in this large window for all cases. This means that the method is not fully automated since it relies on the manually localised tip of the template face.

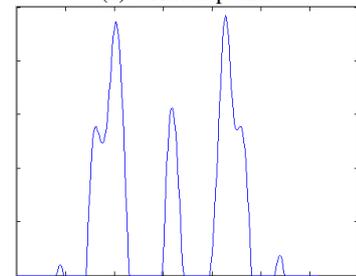
It has been observed that the Gaussian surface curvature of the tip of the nose is distinctly elliptical ($K > 0$), [1], [2], [13]. For this reason the Gaussian surface curvature ($\sigma = 15$ pixels) is evaluated within the search region about the nose tip approximation. The maximum Gaussian curvature within the region is taken as final location of the nose tip (prn).

B. Nose Width Points ($al-al$)

The nose width points are localised by first defining a search region around the detected nose tip. The size of this window (42 mm x 50 mm) is defined based on the mean and standard deviation values published by Farkas [14]. Further detail is provided in [9]. A Laplacian of Gaussian edge detector ($\sigma = 7$ pixels) is then used within this region. Moving in a horizontal direction from the nose tip, the first edge encountered is considered to be the nose contour and is retained. A curvature detection technique developed by Rodriguez and Aggarwal is used to detect the nose width points which are considered to be points of negative curvature (as per notation used) [15]. This is done by generating an unwrapped chain code for the nose contour. A derivative of Gaussian filter is used on this one dimensional signal to detect points of critical curvature (Fig. 2).



(a) Critical points



(b) Derivative of chain code

Fig. 2: Detection of nose width points

The choice of which critical points are selected as nose width points is based on selecting the first critical points immediately above and below the vertical coordinate of the nose tip. The widest of these are selected as nose width points.

C. Inner Eye Corner ($en-en$) & Center of Nose Root (m')

A search region for the left and right inner eye corners is defined using the location of the detected nose tip and nose width points. The vertical limit defined based on the fact that for the average adult, the distance between inner eye corners

and the tip of the nose in the vertical direction is 0.3803 times the distance between the tip of the nose and the top point of the head [9], [10]. Gupta et al. allow for variations in the measure by setting the upper vertical limit at $(prn_y + 0.3803 \times 1.5|prn_y - V_y)$, where V_y is the Y coordinate of the highest vertical point in the 3D model. The horizontal limit is obtained by using the locations of the nose width points and the nose tip. Specifically horizontal limits are defined from the nose tip to $al_{x,left/right} \pm 0.5|al_{x,left} - al_{x,right}|$ for the left and right inner eye corners.

The Gaussian curvature within this region is evaluated and the location of maximum curvature is used as an approximation for the location of the inner eye corner ($\sigma = 15$ pixels, as in III-A). Finally a region of 20mm x 20mm is defined around this peak of Gaussian curvature.

In order to further improve the localisation of the inner eye corner a modified version of the EBGm technique is used [7], [8]. In brief, this technique involves comparing the Gabor coefficients generated for each pixel in the search region with the coefficients for the landmarks of 89 training images. These 89 images consist of neutral and expressive faces. The images are selected in an attempt to cover as much feature variance as possible (i.e. closed/open mouth and eyes). Since Gupta et al. do not state specifically (the database partitions are provided) which 89 images from the database they use as example images, we selected the example images based on the information provided in [9]. 80 Gabor coefficients (known as a Gabor jet) are generated at each landmark for each of the example images. A filter bank of 40 Gabor filters is used (5 scales x 8 orientations). 40 coefficients are generated for both range (3D) and texture (2D) images. While the specific parameters of these filters are not provided in [9], we used the filter bank outlined in by Wiscott et al. [8]. It is important to note that all images should be scaled by $\frac{1}{3}$ when Gabor filtering is applied. The final location of the inner eye corner is obtained by finding the pixel which has a Gabor jet most similar to that of any training landmark. The similarity score is given in equation 1, where J and J' are the jets to be compared. The jets contain either 40 or 80 coefficients depending on which form of EBGm is to be used. Gupta et al. chose to use 2D and 3D Gabor coefficients. In this work 2D, 3D and 2D+3D results are compared.

$$S(\vec{J}, \vec{J}') = \frac{\sum_{i=1}^{40/80} a_i a'_i \cos(\Phi_i - \Phi'_i)}{\sqrt{\sum_{i=1}^{40/80} a_i^2 \sum_{i=1}^{40/80} a'_i{}^2}} \quad (1)$$

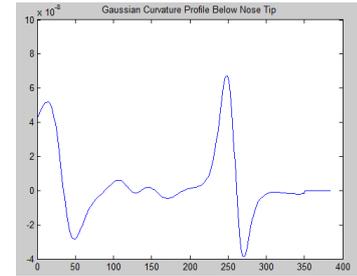
The center of the nose root is determined by finding the mid-point between the two inner eye corners.

D. Outer Eye Corners (ex-ex)

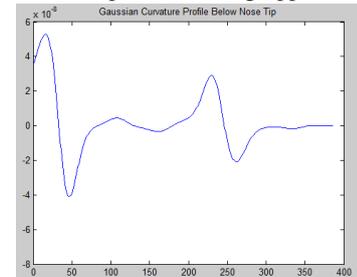
A search region for the outer eye corners is defined based on the location of the detected inner eye corners as per [9]. This 20 x 34 mm region is evaluated using the same search procedure as used for the inner eye corners. Gupta et al. chose to use 2D EBGm search as the outer eye corner region does not have distinct enough curvature characteristics. In this work all three EBGm techniques are evaluated.

E. Mouth Width Corners (ch-ch)

The lip curvature is examined in order to determine a search region for the mouth width corners. The Gaussian curvature of both the upper and lower lips is elliptical in nature. The regions immediately above the upper lip and below the lower lip are hyperbolic ($K < 0$). As shown in figure 3a where the peaks correspond to the nose tip, upper lip, lower lip and chin, moving from left to right. These properties can be used to define upper and lower search limits for the mouth corners. The horizontal limits are defined by $[(al_{x,left} - 0.7|al_{x,left} - al_{x,right}|), (al_{x,left})]$ for ch_{left} and $[(al_{x,right} + 0.7|al_{x,left} - al_{x,right}|), (al_{x,right})]$ for ch_{right} . In order to remove noise a certain amount of smoothing must be carried out when calculating Gaussian curvature. In some cases the Gaussian curvature of the upper or lower lip is too weak and cannot be localised, as shown in figure 3b. In such cases the troughs in Gaussian curvature immediately above and below the lip region are used as limits. While these are usually stronger features than the lips, errors can arise when searching for peak mean curvature in the next stage of the algorithm as there is a high mean curvature along the jaw line.



(a) Lip curvature profile - strong upper & lower lip



(b) Lip curvature profile - weak lower lip

Fig. 3: Comparison of lip curvature profiles

The mean curvature ($\sigma = 2$ pixels) is then calculated for the defined search region. Since the mouth corners are regions of high mean curvature the peak curvature value in this region is taken as an estimate for of the mouth corner. A 30mm x 11mm search region is defined around these mouth corner estimates. The same EBGm procedure used in sections III-C and III-D is used to precisely localise the mouth corners. Gupta et al. chose to use 2D+3D EBGm. In this work 2D, 3D and 2D+3D EBGm results are compared.

IV. EXPERIMENTAL RESULTS & DISCUSSION

Provided below is a summary of results obtained. For more comprehensive results see appendix C.

A. Test Data

The performance of the landmark localisation algorithm is evaluated using the Texas 3DFR database. As mentioned in section II the database contains 1149 pairs of portrait and range images. 89 of these pairs are used in the EBGM search. The remaining 1060 pairs are used as test data.

B. Landmark Localisation Results

The landmark localisation results obtained for the Texas 3DFR database are given in table II. All results are given in millimetres.

As mentioned previously Gupta et al. do not provide 3D error results [9]. In order to ensure that this method faithfully reproduced the original method a comparison with the 2D results obtained by Gupta et al. is provided in XXXX.

Landmark	Author	X std. dev	Y std. dev	2D std. dev
Prn	Gupta	1.045	1.68	1.978
	This method	0.766	1.714	1.705
Al Left	Gupta	0.721	1.655	1.805
	This method	0.647	0.710	0.739
Al Right	Gupta	0.798	1.646	1.829
	This method	0.546	0.814	0.818
En Left	Gupta	1.488	1.245	1.940
	This method	1.249	0.908	1.363
En Right	Gupta	1.354	1.344	1.908
	This method	1.378	0.792	1.417
M'	Gupta	1.355	1.811	2.261
	This method	1.415	1.010	1.417
Ex Left	Gupta	1.795	1.285	2.208
	This method	1.727	1.047	1.850
Ex Right	Gupta	2.126	1.384	2.537
	This method	1.940	1.248	2.149
Ch Left	Gupta	1.948	0.933	2.160
	This method	1.749	1.692	2.321
Ch Right	Gupta	1.976	1.045	2.235
	This method	1.429	0.844	1.460

TABLE I: Error standard deviation comparison with Gupta et al. [9]

The mean error result of the nose tip is noticeably larger than the localisation of the other landmarks. On closer examination it appears that in all cases the detected nose tip is above the manually localised nose tip (in the Y direction). This can clearly be seen in the boxplot in figure 4. This figure shows clearly that the median value for the X error is 0mm as expected in a normal error distribution. The Y distribution is extremely skewed to one side of the manually localised nose tip (a negative Y error is above the manual location for an upright face). Since the standard deviation of the Y error is relatively small it seems that the issue is that the peak of Gaussian curvature does not correspond to the same location the manual annotators have identified as the nose tip.

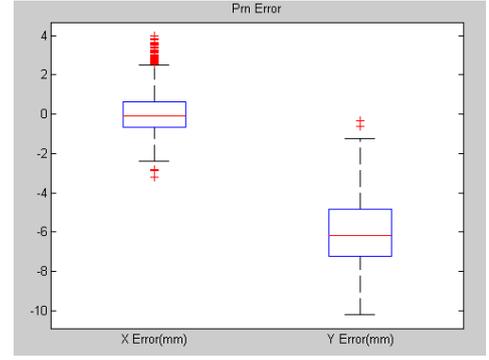


Fig. 4: Boxplot of nose tip localisation error

The mean error results obtained for the nose width points are reasonable while the standard deviations are impressive. The most impressive results are those obtained using the modified EBGM technique. A 3D mean error of under 2mm is recorded for both inner eye corners. The outer eye corners which are slightly more difficult to localise are detected with a mean error of under 2.6mm. A mean error of below 2.16mm is achieved for both mouth corners. The algorithm does have particular difficulty with faces where facial hair is present. This is as expected when using Gabor filters as there is a significantly different response to a Gabor filter when facial hair is present.

One interesting point to note is that the three worst results obtained are for the three landmarks localised using techniques which do not involve training. The training stage of EBGM uses manual landmark locations. This means that when EBGM is used, the algorithm searches for a location on an unknown image which is most similar to the training data, which is based on manual locations. For the nose tip and width points the algorithm searches for a particular image feature (e.g. maximum Gaussian curvature) which is said to be present at that landmark. Perhaps using EBGM for all landmarks might yield better performance. Another possible issue could be marker bias. No details are provided about how many annotators are used but using separate annotators for test and training data could be a possible solution.

Landmark	3D mean (mm)	3D std. dev (mm)
Prn	6.147	1.746
Al Left	3.354	1.647
Al Right	3.310	1.877
En Left	1.821	1.499
En Right	1.747	1.520
M'	2.760	1.591
Ex Left	2.477	2.576
Ex Right	2.590	2.990
Ch Left	2.159	3.039
Ch Right	2.016	2.153

TABLE II: Landmark localisation error

C. Texture & Range Comparison

The inner eyes and outer mouth corners are detected using 2D + 3D EBGM while 2D EBGM is used for the outer eye

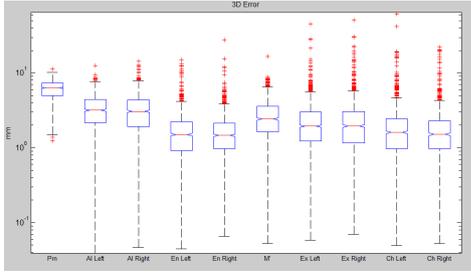


Fig. 5: 3D error boxplot

corners. The same similarity metric is used in each case (1) with the only difference being the coefficients examined.

Landmark	Method	3D mean	3D std. dev
En Left	2D	1.833	1.533
	3D	2.181	1.695
	2D + 3D	1.821	1.499
En Right	2D	1.753	1.546
	3D	1.986	1.581
	2D + 3D	1.747	1.520
Ex Left	2D	2.477	2.576
	3D	5.103	5.280
	2D + 3D	2.388	2.116
Ex Right	2D	2.590	2.990
	3D	8.908	7.215
	2D + 3D	2.489	2.273
Ch Left	2D	2.197	2.830
	3D	2.537	2.886
	2D + 3D	2.159	3.039
Ch Right	2D	2.151	2.442
	3D	2.202	1.613
	2D + 3D	2.016	2.153

TABLE III: 2D, 3D & 2D+3D EBGM comparison

Interestingly, table III shows that for the inner and outer eye corners the inclusion of range coefficients improves localisation results. Gupta et al. use 2D + 3D for the inner eye corner while they choose to use just 2D for the outer eye corners. The results obtained here suggest that a similar improvement in localisation could be achieved with the inclusion of range information. The original authors state that only 2D is used as the outer eye corners do not have any distinct curvature characteristics. While it is clear that just using 3D information results in poor localisation performance it should be noted that the 3D information only influences the result of localisation when a 3D coefficient is more similar to one of the training image coefficients than any of the 2D coefficients. This means that in some individual cases the inclusion of 3D information may adversely affect localisation but for the entire database there is a decrease in mean error.

With regard to the mouth corners the use of texture and range information results in the best mean error performance. This is the same as the behaviour for the other landmarks. Once again the worst mean error is recorded when just range information is used.

It is clear that in all cases examined the inclusion of more information (texture & range) in the EBGM stage results in better overall localisation. This suggests that the similarity

score and the procedure Gupta et al. use for choosing the landmark location works quite well. It suggests that in the majority of cases the inclusion of extra information leads to enhanced localisation performance. Obviously there is a computational overhead to be considered when including this extra information but in cases where speed isn't an issue it seems that the inclusion of 2D and 3D information leads to the best localisation performance.

Since the 2D and 3D EBGM techniques are directly comparable, table III shows that for all landmarks examined texture information yields better results. Though for the inner eye corners and mouth corners this difference is quite small.

D. Improvement Achieved Using EBGM

For each of the landmarks localised using EBGM a similar approach is taken. An approximation of the landmark location is first obtained and a search window is then defined around this approximation. In this section the accuracy of this approximation is evaluated and the relative improvement achieved by the EBGM stage is examined. The inner eye and mouth corners approximations are obtained by examining the curvature of the local area. The outer eye corner approximation is based on the location of inner eye corners.

Landmark	Method	3D mean (mm)	3D std. dev (mm)
En Left	Approximation	2.417	1.555
	EBGM	1.821	1.499
En Right	Approximation	2.767	1.819
	EBGM	1.747	1.520
Ex Left	Approximation	10.231	7.099
	EBGM	1.821	1.499
Ex Right	Approximation	10.430	7.736
	EBGM	2.590	2.990
Ch Left	Approximation	5.359	4.345
	EBGM	2.159	3.039
Ch Right	Approximation	5.714	4.749
	EBGM	2.016	2.153

TABLE IV: Comparison of approximations & EBGM results

The results in table IV clearly show that an improvement in localisation is achieved using EBGM. In the case of the inner eye corners a relatively small improvement is gained by using EBGM. In certain cases where speed is an issue, the approximation of the inner eye corner could be used as the final location of that landmark. This would not result in a major decrease in localisation accuracy. The remaining four landmarks show a significant improvement when EBGM is used, with the outer eye corner localisation improving significantly. As a whole the results show that the increased computation overhead is warranted in terms of localisation performance.

V. CONCLUSION

We have shown that the method developed by Gupta et al. performs well in terms of landmark localisation. The one weak point is the localisation of the nose tip which is quite poor. Even though the localisation of the tip is poor it does

not appear to adversely affect the localisation of subsequent landmarks where the location of the nose tip is used to define a search region. Another better performing method, such as that used by Segundo et al., could perhaps be used for the localisation of the nose tip [1].

It was determined that for the EBGM stage, the inclusion of both texture and range information yields the best results. Interestingly, for the inner eye corners and mouth corners the error results recorded are similar for each of the EBGM methods. For the outer eye corner 3D EBGM performed quite poorly, with 2D and 2D+3D obtaining similar results. This suggests that for outer eye corner detection, 2D EBGM could be used without a significant (0.3mm) decrease in mean error.

Finally the improvement in localisation achieved by using EBGM was examined. A noticeable improvement was observed for the outer eye corner and mouth corners localisation. The distinct curvature features of the inner eye region results in it being easier to localise and as a consequence the approximation obtained was reasonably good.

REFERENCES

- [1] M. P. Segundo, L. Silva, O. R. P. Bellon, and C. C. Queirolo, "Automatic face segmentation and facial landmark detection in range images," *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, vol. 40, no. 5, pp. 1319–1330, 2010.
- [2] C. Creusot, N. Pears, and J. Austin, "A machine-learning approach to keypoint detection and landmarking on 3d meshes," *International Journal of Computer Vision*, vol. 102, no. 1-3, pp. 146–179, 2013.
- [3] X. Zhao, E. Dellandrea, L. Chen, and I. A. Kakadiaris, "Accurate landmarking of three-dimensional facial data in the presence of facial expressions and occlusions using a three-dimensional statistical facial feature model," *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, vol. 41, no. 5, pp. 1417–1428, 2011.
- [4] P. Perakis, G. Passalis, T. Theoharis, and I. Kakadiaris, "3d facial landmark detection under large yaw and expression variations," 2012.
- [5] G. Passalis, P. Perakis, T. Theoharis, and I. A. Kakadiaris, "Using facial symmetry to handle pose variations in real-world 3d face recognition," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 33, no. 10, pp. 1938–1951, 2011.
- [6] J. R. Movellan, "Tutorial on gabor filters," *Open Source Document*, 2002.
- [7] S. Jahanbin, A. C. Bovik, and H. Choi, "Automated facial feature detection from portrait and range images," in *Image analysis and interpretation, 2008. SSIAI 2008. IEEE southwest symposium on*. IEEE, 2008, pp. 25–28.
- [8] L. Wiskott, J.-M. Fellous, N. Kuiger, and C. Von Der Malsburg, "Face recognition by elastic bunch graph matching," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 19, no. 7, pp. 775–779, 1997.
- [9] S. Gupta, M. K. Markey, and A. C. Bovik, "Anthropometric 3d face recognition," *International Journal of Computer Vision*, vol. 90, no. 3, pp. 331–349, 2010.
- [10] L. G. Farkas and I. R. Munro, *Anthropometric facial proportions in medicine*. Charles C. Thomas Publisher, 1987.
- [11] S. Gupta, K. R. Castleman, M. K. Markey, and A. C. Bovik, "Texas 3d face recognition database," in *Image Analysis & Interpretation (SSIAI), 2010 IEEE Southwest Symposium on*. IEEE, 2010, pp. 97–100.
- [12] S. Gupta, M. K. Markey, and A. C. Bovik. (2010) Texas 3d face recognition database. [Online]. Available: <http://live.ece.utexas.edu/research/texas3dfr/index.htm>
- [13] A. B. Moreno, A. Sánchez, J. F. Vélez, and F. J. Díaz, "Face recognition using 3d surface-extracted descriptors," in *Irish Machine Vision and Image Processing Conference*, vol. 2003. Citeseer, 2003.
- [14] L. G. Farkas and S. A. Schendel, "Anthropometry of the head and face," *American Journal of Orthodontics and Dentofacial Orthopedics*, vol. 107, no. 1, pp. 112–112, 1995.
- [15] J. J. Rodriguez and J. Aggarwal, "Matching aerial images to 3-d terrain maps," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 12, no. 12, pp. 1138–1149, 1990.
- [16] Paul F Whelan (2011), "VSG Image Processing and Analysis (VSG IPA) Toolbox", Technical Report, Centre for Image Processing & Analysis, Dublin City University
- [17] Paul F Whelan (2011), "VSG Image Processing and Analysis (VSG IPA) Matlab Toolbox - Software", <http://www.cipa.dcu.ie/code.html>, Centre for Image Processing & Analysis, Dublin City University.
- [18] Paul F. Whelan and D. Molloy (2000), "Machine Vision Algorithms in Java: Techniques and Implementation", Springer (London), 298 Pages. ISBN 1-85233-218-2.
- [19] Paul F Whelan (2011), "EE544: Computer Vision Course Notes", Centre for Image Processing & Analysis, Dublin City University.
- [20] Paul F Whelan (2011), "EE425/EE453: Image Processing and Analysis Course Notes", Centre for Image Processing & Analysis, Dublin City University.