Craniofacial Syndrome Identification Using Convolutional Mesh Autoencoders

Lara S. van de Lande (✉ l.lande@ucl.ac.uk)

UCL Great Ormond Street Institute of Child Health & Craniofacial Unit, Great Ormond Street Hospital for Children

Eimear O'Sullivan

Imperial College London

Athanasios Papaioannou

Imperial College London

Richard W. Breakey

UCL Great Ormond Street Institute of Child Health & Craniofacial Unit, Great Ormond Street Hospital for Children

N. Owase Jeelani

UCL Great Ormond Street Institute of Child Health & Craniofacial Unit, Great Ormond Street Hospital for Children

Allan Ponniah

The Royal Free Hospital

Christian Duncan

Alder Hey Children's Hospital

Silvia Schievano

UCL Great Ormond Street Institute of Child Health & Craniofacial Unit, Great Ormond Street Hospital for Children

Roman H. Khonsari

Necker-Enfants Malades Hospital

Stefanos Zafeiriou

Imperial College London

David J. Dunaway

UCL Great Ormond Street Institute of Child Health & Craniofacial Unit, Great Ormond Street Hospital for Children

Research Article

Keywords: Early diagnosis, photographic analysis, Muenke, Crouzon, Apert, CMAs, artificial intelligence (AI), Clinical diagnosis

Loading [MathJax]/jax/output/CommonHTML/jax.js

Posted Date: April 3rd, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-373399/v1>

License: This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Abstract

Clinical diagnosis of craniofacial anomalies requires expert knowledge. Recent studies have shown that artificial intelligence (AI) based facial analysis can match the diagnostic capabilities of expert clinicians in syndrome identification. In general, these systems use 2D images and analyse texture and colour. They are powerful tools for photographic analysis but are not suitable for use with medical imaging modalities such as ultrasound, MRI or CT, and are unable to take shape information into consideration when making a diagnostic prediction. 3D morphable models (3DMMs), and their recently proposed successors, mesh autoencoders, analyse surface topography rather than texture enabling analysis from photography and all common medical imaging modalities and present an alternative to image-based analysis.

We present a craniofacial analysis framework for syndrome identification using Convolutional Mesh Autoencoders (CMAs). The models were trained using 3D photographs of the general population (LSFM and LYHM), computed tomography data (CT) scans from healthy infants and patients with 3 genetically distinct craniofacial syndromes (Muenke, Crouzon, Apert).

Machine diagnosis outperformed expert clinical diagnosis with an accuracy of 99.98%, sensitivity of 99.95% and specificity of 100%. The diagnostic precision of this technique supports its potential inclusion in clinical decision support systems. Its reliance on 3D topography characterisation make it suitable for AI assisted diagnosis in medical imaging as well as photographic analysis in the clinical setting.

Main

Early diagnosis of many genetic disorders improves outcome, but rarity and the variety of possible syndromes make clinical diagnosis problematic. Syndromic craniosynostosis (SC) comprises a group of conditions characterised by premature fusion of the skeletal sutures, compromising normal development of the skull and brain.^{1,2} Delayed diagnosis is common in phenotypically mild SC risking irreversible functional impairment, including visual failure, neurocognitive defects and airway problems that can be avoided by timely diagnosis and treatment. The inadequacy of current screening paradigms makes syndromic craniosynostosis a prime candidate for computer-assisted diagnosis and referral. Muenke, Apert, and Crouzon syndrome are SC variants caused fibroblast growth factor receptor (FGFR) gene mutations and occur between 1:30,000 to 1:65,000 live births³. Phenotypic presentation is variable with some crossover in phenotype between the syndromes.

The application of computer vision and deep learning approaches has already proven effective for the identification and diagnosis of CS from 2-dimensional (2D) images⁴⁻⁷. Among the more prominent examples is DeepGestalt, a deep convolutional neural network (DCNN) trained on tens of thousands of images to identify facial phenotypes for genetic disorders⁸. While such systems demonstrate impressive results, they are unable to take advantage of the rich geometric information in the face and cranium that may give critical insight into the phenotypical variations associations with different syndromes.

Advances in 3-dimensional (3D) modelling and geometric deep learning have resulted in the introduction of a more shape-based approach craniofacial analysis^{9–12}. Previous work has successfully applied statistical shape models in adult populations for diagnosis and surgical simulation in orthognathic patients¹³. The recent introduction of convolutional mesh autoencoder models (CMAs), a deep neural network approach to 3D model construction, offers further potential for the construction of shape-based models^{12,14}. These models learn to extract meaningful shape features from the input data and can consequently be used for classification tasks. Such approaches have not yet been applied for automated diagnosis of craniofacial syndromes using an age matched population.

We report on a geometric deep learning approach to the characterisation and identification of SC. The framework leverages convolutional mesh autoencoders and is trained using 3D data from healthy and syndromic individuals, focused on the identification of three distinct types of SC, namely Apert, Crouzon, and Muenke syndrome. Rather than relying on image data, the proposed model leverages the rich geometric information of the 3-dimensional scans. We demonstrate the power of the model for syndrome characterisation and classification, and illustrate its diagnostic sensitivity with an unusual Crouzon case.

Results

A summary of the demographics for the subjects used in this study is given in Table 1. This includes data from four different databases: CT data from 122 SC patients (Apert, Crouzon, and Muenke, mean age of 5.0 ± 5.1 years, 58% male ($n = 70$)) and 142 healthy infants (mean age, 1.9 ± 1.2 years, 56% male), Stereophotogrammetric data from 196 healthy subjects from the LSFM dataset¹⁵, and 139 healthy subjects from the LYHM database¹¹. As the volunteers in both datasets were typically above the age of four while many cases in the syndromic dataset were quite young, a dataset of CT scans from healthy children was also used. For full details, see Methods.

Table 1

Overview of the face and cranium dataset of the included Syndromic Craniosynostosis and normal samples. All syndromic and infant samples were acquired via CT-scan. The LSFM and LYHM databases were obtained using 3dMD™ photometric stereo capture device set-ups^{10,11}

Type of SC	Number of subjects	Average age, years	Age range at time of scan	Sex (M:F)
<i>Face</i>				
LSFM	196	10.5 ± 4.0	4 years – 17 years	98:98 (50%:50%)
Paediatric	142	1.9 ± 1.2	1 day – 47 months	79:63 (56%:44%)
Apert	47	6.1 ± 6.2	48 days – 20 years	28:19 (60%:40%)
Crouzon	61	5.3 ± 4.4	25 days – 17 years	35:25 (58%:42%)
Muenke	14	1.6 ± 2.1	1 day – 8 years	7:7 (50%:50%)
Total	460		1 day – 20 years	247:213 (54%:46%)
<i>Cranium</i>				
LYHM	139	10.9 ± 3.8	4 years – 18 years	76:63 (55%:45%)
Paediatric	111	1.8 ± 1.1	1 day – 47 months	59:52 (53%:47%)
Apert	39	6.5 ± 6.3	48 days – 20 years	22:17 (56%:44%)
Crouzon	53	5.4 ± 4.4	5 months – 17 years	30:23 (57%:43%)
Muenke	11	1.7 ± 2.3	1 day – 8 years	6:5 (55%:45%)
Total	353		1 day – 20 years	193:160 (55%:45%)

Intrinsic Model Evaluation

Using the available databases, three distinct classes of model were constructed to assess the role of facial and cranial shape in the diagnosis of SC; a face-only model, a head-only model, and a combined head-and-face model (see Methods).

All models reported low reconstruction errors when assessing their reconstruction accuracy and specificity (see Methods). For the face, head and combined models, these values were $1.4 \pm 1.2\text{mm}$, $3.8 \pm 3.1\text{mm}$, and $2.9 \pm 2.5\text{mm}$, respectively. Reconstruction error was higher for models that included the head shape, likely due to the greater degree of variation between subjects in this region. Model specificity was evaluated by randomly synthesising 1000 samples and comparing them to their nearest real neighbour¹⁰. Values of 2.7mm, 4.3mm, and 3.9mm for the face-only, head-only, and combined models respectively indicate that the samples generated are realistic.

Manifold Visualisation

Loading [MathJax]/jax/output/CommonHTML/jax.js

To assess the diagnostic capacity of the model, t-distributed stochastic neighbour embedding (t-SNE) was applied to the high dimensional latent vectors of the patients and volunteers. (Fig. 1) When samples were labelled by syndromic class, clear clusters emerge for the face-only healthy and syndromic groups (Apert, Crouzon, and Muenke). The two clusters that are observed for the healthy individuals are due to age, with samples from the paediatric and LSFM datasets clustering separately. Within the syndromic cluster, we observe further sub-clusters for the included syndromes. The clusters formed for the head-only embeddings are not as distinct as those observed for the head-only cases, however groups for Apert, Crouzon, Muenke, and healthy individuals do emerge. When considering the model constructed using the combined head-and-face template, we again see clear clusters forming for the different subgroups in the dataset.

In all cases, even though the syndromic samples tend to group more tightly together, it is noted that the syndromes themselves seem relatively disentangled. The proximity of the Crouzon cluster to the normal cases in each of the embeddings indicates that this phenotype has a milder manifestation than either Apert or Muenke syndrome. For both the head-only and combined head-and-face embeddings, we note that a number of Muenke and Crouzon samples cluster closer to the group of healthy cases.

Syndrome Classification

Classification was performed with all syndromic and non-syndromic scans. A split of 80%-20% for training and testing data was assessed over 1,000 iterations. The mean sensitivity, specificity, and accuracy over all iterations for each of the assessed regions in the binary classification experiment, and the confusion matrices for both binary and multi-class classification, are shown in Table 2 and Fig. 2, respectively.

Binary classification to identify whether an individual belonged to either the syndromic or non-syndromic groups showed accuracies of greater than 99% in all cases. The high sensitivity of the models indicates that very few syndromic cases were misidentified as normal volunteers (one, sixteen, and two cases per thousand for the face-only, head-only, and combined models respectively). The inverse is also true; as indicated by the high model specificity values, few healthy volunteers were seldom, if ever, misidentified as having a craniofacial syndrome.

The multi-class classification model endeavoured to predict whether a patient belonged to either the non-syndromic, Apert, Crouzon, or Muenke categories. Accuracies of 98.3%, 97.9%, and 98.2% were observed for the face-only, head-only, and combined regions, respectively. When considering the face-only, Muenke patients were the most likely to be mis-diagnosed by our model. As there are the fewest instances of this syndrome in the database, such results are to be expected, and increasing the quantity of Muenke samples would likely lead to increased accuracy for these patients. When the head shape was considered for classification, Crouzon and Apert patients were most likely to be misidentified as each other. As with the binary classification, the poorest performance was seen for the head-only model. These results would indicate that the facial region contains valuable shape information for the correct diagnosis of SC.

Our method has proven as least as sensitive as expert diagnosis in a multidisciplinary clinic and in one case more sensitive. A 2-year-old sibling of a patient with Crouzon syndrome, judged clinically to be unaffected, proved to have Crouzon syndrome on genetic testing. Subsequent three images analysed by our technique demonstrated the ability of the model to diagnose clinically undetectable variations from the norm. (Fig. 3)

Table 2
Classification results for the binary classification experiments.

Model	Sensitivity (%)	Specificity (%)	Accuracy (%)
Face Only	99.95	100.00	99.98
Head Only	98.36	99.41	99.09
Head and Face	99.82	100.00	99.95

Discussion

AI assisted diagnosis is set to play an increasing role in healthcare particularly in relation to rare or difficult to diagnose conditions. In this work, we leverage state-of-the-art geometric deep learning approaches to present a framework for the detection and classification of a subset of syndromic craniosynostoses. Requiring only a 3D input allows potentially deceptive texture information to be omitted, allowing the architecture to focus on the extraction of characteristic shape-based features to return an accurate diagnosis.

Our autoencoder model demonstrates high sensitivity and specificity. The model can be applied for both binary classification (syndromic vs healthy) and multi-class classification (Apert, Crouzon, Muenke, and healthy). The high degree of separation of the different classes in the shape space of the model, as demonstrated by the t-SNE embeddings, intuitively supports these results. The high sensitivity and specificity of the model make suitable as a diagnostic aid in primary and secondary care settings ensuring reliable diagnosis with few false positive results.

Utilising 3D topography rather than the surface texture analysis central to most other facial analysis techniques lends itself to integration with many forms of diagnostic imaging technologies, such as MRI, ultrasound scans, and CT. Our technique facilitates auto-segmentation of surfaces as well as syndrome identification and could therefore be used as a machine learning diagnostic tool to aid radiological diagnosis. Integration with ultrasound imaging is of particular interest, as this presents an opportunity for the foetal detection of genetic disorders¹⁶. The rising availability of 3D scanning applications and cameras on mobile devices presents further possibilities to introduce such a framework in primary and secondary care. In a field where timely diagnosis is necessary for appropriate management, the use of such technologies to detect SC and other conditions will streamline assessment at an earlier stage could be pivotal to improving long-term health outcomes.

While several machine learning approaches for the identification of craniofacial syndromes have been presented [11], to the best of our knowledge, this is the first time the problem of syndrome classification has been approached from a shape-based perspective.

Methods

Institutional review board statement

Patient data for this study were retrospectively retrieved from electronic medical records after receiving approval from the Institutional Review Board: Great Ormond Street Hospital (R&D no. 14DS25). Informed consent was obtained from all participants, or if participants are under 18, from a parent and/or legal guardian. All methods were carried out in accordance with relevant guidelines and regulations.

Data sources

A full summary of the demographics for the databases used is given in Table 1.

All SC patients diagnosed with Apert, Crouzon, or Muenke syndrome at the *Craniofacial Unit of Great Ormond Street Hospital for Children, London, UK* were reviewed retrospectively for preoperative 3D imaging. Computed Tomography (CT) head scans were selected as the most suitable image modality for the assessment of the craniofacial anatomy. CT-scans with insufficient quality for 3D reconstruction due to low number of slices or (movement) artefacts were excluded. Baseline characteristics were collected from corresponding medical charts. We refer to this database as “the syndromic dataset”.

At *Necker Children's Hospital, (Necker-Enfants Malades Hospital), Paris, France*, CT-scans of patients between the age of 0–4 years without a history of craniofacial anomalies were assessed. Patients indicated for a CT-scan between 2011 and 2018 due to headaches, trauma or epilepsy were reviewed for inclusion. The scans were evaluated by two independent reviewers: 1) a paediatric radiologist and 2) a clinical research fellow in craniofacial surgery, to exclude any scans with abnormalities such as fractures, brain tumours, brain damage or craniofacial anomalies. Henceforth, this database of normal craniofacial images will be referred to as ‘the paediatric dataset’. For the paediatric and syndromic datasets, several samples used for the face models were omitted when constructing the head models. Scans with incomplete cranium information or those with previous calvarial surgery, such as posterior vault remodelling, were excluded.

The original data consisting of 10,000 faces of the Large Scale Face Model (LSFM) was used for age-matched diagnostics for cases between the ages of 4 and 17. Of the subjects who met the desired age criteria, 196 samples were randomly selected to provide this age matched reference (7 male, 7 female for each age).

The Liverpool-York head model (LYHM) was constructed using the craniofacial scans from approximately 1,200 individuals aged between 2 and 90 from the Headspace database [11]. The scans from those aged

Loading [MathJax]/jax/output/CommonHTML/jax.js age-matched reference for full head scans from the

Syndromic Craniosynostosis dataset. The database included 139 scans that met the required age criteria. The mean age was 10.9 ± 4.4 years, and 55% were male.

Image pre-processing

CT-scan pre-processing

The DICOM-files of the collected datasets, i.e. the syndromic and paediatric datasets were converted to 3D soft tissue meshes by applying standardised skin setting using Horos, an open-source medical viewer, and exported at stereolithography (STL) files. The STL files were imported in MeshMixer, an open-source software, to undergo a cleaning process where redundant objects, such as draping and gel pads, back of the CT-scanner, pacifiers, and lines and tubes, were removed. The meshes were saved as an Object (OBJ) files. A sparse set of 68 facial landmarks were manually added to the meshes. To encourage good correspondence around the ears, an additional 55 landmarks per ear were also added to the meshes, following the landmark template outlined in ¹⁷. Using the facial landmarks, the raw meshes were subjected to Procrustes analysis, allowing them to be rigidly aligned with the template meshes prior to obtaining dense correspondence.

The construction of a 3D mesh autoencoder model requires that all meshes are re-parametrised to have a consistent topological structure, where each mesh has the same number of vertices connected in consistent triangulation, with similar vertices having the same semantic meaning. Meshes meeting this criterion are said to be in dense correspondence. In this study approach, a non-rigid iterative closest point registration (non-rigid ICP) was used to achieve dense correspondence ¹⁸. This approach allows for the use of landmark points to guide the correspondence process, and the use of data weights to restrict how certain vertices can move during the registration. As this method requires an initial template mesh with the desired final topology, the cropped facial template from ¹⁵ was used for construction of the facial models, while the template from ¹⁹ was used for the craniofacial models. (Fig. 4)

Dense correspondence for the Syndromic Craniosynostosis dataset

The faces and heads of those in the infant and syndromic craniosynostosis datasets differ greatly in shape both from each other, and from the samples used for the construction of previous (cranio)facial 3DMMs. For this reason, the use of landmark guided NICP was not always sufficient to achieve good quality correspondences. As such, Gaussian Processes were applied to the face template mesh to increase deformation flexibility and improve the quality of the correspondences obtained for the syndromic cases.

3D Mesh Autoencoder Construction

Once dense correspondence had been achieved for all meshes, the 3D models were created using mesh autoencoders. Three classes of model were created: face-only, head-only, and combined head and face

models. (Fig. 4) For each of these classes, a model was trained using all samples up to and including the age of three, all samples above the age of three, and a final model constructed from all available data.

The autoencoder architecture applied here is similar to that described in ¹². Four convolutional and down-sampling/up-sampling layers were used for the encoder and decoder for construction of the head-only and combined head-and-face models. Five such layers were used for the face-only model as the template mesh for this model is far denser (28,431 vertices for the face-only model, compared to 7,505 vertices in the head-only model, and 17,039 vertices in the combined model). Encoder convolutional filter sizes of [16, 16, 32, 32, 32] were used for the face-only model, while encoder filter sizes of [16, 16, 32, 32] were used for the head-only and combined models. In both cases, decoder filter sizes were the mirror of the encoder. Each convolutional layer in the encoder was followed by a mesh downsampling layer by a factor of 4 - in the decoder, this was replaced with a layer to upsample the mesh by a factor of 4. An additional convolutional layer with an output dimension of 3 is added to the decoder to allow for the reconstruction of the 3D shape coordinates. An ELU activation function was applied after each convolutional layer. Spiral convolutional with a fixed length of nine were applied in all layers ²⁰, and a latent vector size of 128 was used. Model weights were initialised using Xavier initialisation. Adam optimisation with an initial learning rate of 1×10^{-3} , and a learning rate decay of 0.99 was used. Mesh vertices were used as the autoencoder input and an L1 reconstruction loss was applied to the output. All models were trained for 300 epochs using a batch size of 16.

Models were constructed for the face-only, head-only, and combined head and face regions to gain an insight into how characteristic these regions are of the various craniofacial syndromes. All paediatric healthy cases and syndromic patients with incomplete or poor-quality head CT-scans were omitted from the head-only and combined head and face models. Consequently, these models were constructed using fewer samples. Due to the wide age range of the individuals in the dataset (1 day to 20 years), two additional models were created for each of the regions of interest. The first included all patients and volunteers who were up to and including three years of age. The second consisted of those above the age of three. Nine models were created in total. All results pertaining the age-based models can be found in the supplementary materials.

Error Quantification

To assess reconstruction accuracy, the Euclidean distance, d , between each sample in the real dataset and the corresponding model reconstruction was calculated on a per-vertex basis. For two meshes, A and B , with n vertices, the mean error over all vertices is defined as:

$$d = \frac{\sum_{i=1}^n \sqrt{(x_{i,A} - x_{i,B})^2 + (y_{i,A} - y_{i,B})^2 + (z_{i,A} - z_{i,B})^2}}{n}$$

where x , y , and z correspond to the cartesian coordinates of the mesh vertices. The mean error over all samples was reported.

Loading [MathJax]/jax/output/CommonHTML/jax.js

Model Specificity

Specificity is a metric that is commonly used to evaluate the validity of novel instances created by generative models. To assess this, 1000 samples were randomly synthesised for each of the models and the Euclidean distance to all ground truth samples was calculated. The specificity error was reported as the mean Euclidean distance over all vertices between a synthesised face and the closest ground truth neighbour.

Manifold Visualisation

t-Distributed Stochastic Neighbour Embedding (t-SNE) is a dimensionality reduction technique that allows the high dimensional shape vectors to be embedded in a lower dimensional space and can reveal hidden structures of the data²¹. To assess the diagnostic capacity of the models, t-SNE was applied to the high dimension latent vector encodings for the patients and healthy volunteers, allowing the global manifold of these vectors to be embedded in a 2-dimensional space for visualisation. Samples were then labelled according to their syndromic class (Normal, Apert, Crouzon, or Muenke) with the aim of uncovering distinct groupings, or clusters. All t-SNE embeddings were created using a perplexity of 30 and run for 1,000 iterations.

Classification

Autoencoders are often utilised for their ability to compress data into a much more compact format. This manifests as the latent vectors of the model. These latent vectors provide a natural means by which we can attempt to classify the data and determine its applicability as a diagnostic tool.

Classification was performed using a Support Vector Machine (SVM) with linear kernel and balanced class weighting. A stratified data split with an 80%:20% train:test proportion was used. The scikit-learn SVM implementation with default gamma and regularization parameters ($C = 1.0$) was employed. The mean accuracy, specificity, and sensitivity were calculated following a Monte-Carlo cross-validation system where the training and test sets were randomly selected 10,000 times.

Declarations

Acknowledgements

This work has been funded by Great Ormond Street Hospital for Children Charity (Grant No. 12SG15), the Engineering and Physical Sciences Research Council (EP/N02124X/1) and the European Research Council (ERC-2017-StG-757923). This report incorporates independent research from the National Institute for Health Research Biomedical Research Centre Funding Scheme. The views expressed in this publication are those of the author(s) and not necessarily those of the NHS, the National Institute for Health Research or the Department of Health. The funders had no role in study design, collection, analysis and interpretation of the data, decision to publish, or preparation of the manuscript.

Author contributions

A.P., L.S.v.d.L., and D.J.D., initiated and designed the study. A.P. and E.O.S. developed the algorithms, designed data analysis and performed experiments, and S.Z. validated the methodology, helped design the study and helped with data interpretation. D.J.D., N.U.O.J. and R.K. provided the electronic medical record data, and L.S.v.d.L. and R.W.F.B. extracted the electronic medical record data. L.S.v.d.L., R.W.F.B., and E.O.S. pre-processed the data, L.S.v.d.L., E.O.S., A.P., and wrote the paper in conjunction with R.K., and D.J.D., and S.S., A.P., N.O.U.J. and S.Z. provided manuscript feedback.

Competing interests

A.P. and S.Z. currently work with Huawei Technologies Co., Ltd. They were with Imperial College London and University College London during the experiments, respectively. The other authors declare no competing interests.

References

1. Slater, B. J. *et al.* Cranial sutures: a brief review. *Plast Reconstr Surg.* **121**, 170–178 (2008).
2. Stricker, S. & Mundlos, S. FGF and ROR2 receptor tyrosine kinase signaling in human skeletal development. *Curr Top Dev Biol.* **97**, 179–206 (2011).
3. Reference, G. H. (Accessed August 2020).
4. Basel-Vanagaite, L. *et al.* Recognition of the Cornelia de Lange syndrome phenotype with facial dysmorphology novel analysis. *Clin Genet.* **89**, 557–563 (2016).
5. Cerrolaza, J. J. *et al.* Identification of dysmorphic syndromes using landmark-specific local texture descriptors. in 2016 *IEEE 13th International Symposium on Biomedical Imaging (ISBI)* 1080–1083 (IEEE, 2016).
6. Shukla, P., Gupta, T., Saini, A., Singh, P. & Balasubramanian, R. A deep learning frame-work for recognizing developmental disorders. in 2017 *IEEE Winter Conference on Applications of Computer Vision (WACV)* 705–714 (IEEE, 2017).
7. Valentine, M. *et al.* Computer-aided recognition of facial attributes for fetal alcohol spectrum disorders. *Pediatrics* **140**(2017).
8. Gurovich, Y. *et al.* Identifying facial phenotypes of genetic disorders using deep learning. *Nat Med.* **25**, 60–64 (2019).
9. Blanz, V. & Vetter, T. A morphable model for the synthesis of 3D faces. in Proceedings of the 26th annual conference on Computer graphics and interactive techniques 187–194(1999).
10. Booth, J., Roussos, A., Zafeiriou, S., Ponniah, A. & Dunaway, D. A 3d morphable model learnt from 10,000 faces. in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition 5543–5552(2016).

11. Dai, H., Pears, N., Smith, W. A. & Duncan, C. A 3d morphable model of craniofacial shape and texture variation. in Proceedings of the IEEE International Conference on Computer Vision 3085–3093(2017).
12. Ranjan, A., Bolkart, T., Sanyal, S. & Black, M. J. Generating 3D faces using convolutional mesh autoencoders. in Proceedings of the European Conference on Computer Vision (ECCV) 704–720(2018).
13. Knoops, P. G. *et al.* A machine learning framework for automated diagnosis and computer-assisted planning in plastic and reconstructive surgery. *Scientific reports*. **9**, 1–12 (2019).
14. Bouritsas, G., Bokhnyak, S., Ploumpis, S., Bronstein, M. & Zafeiriou, S. Neural 3d morphable models: Spiral convolutional networks for 3d shape representation learning and generation. in Proceedings of the IEEE International Conference on Computer Vision 7213–7222(2019).
15. Booth, J., Roussos, A., Ponniah, A., Dunaway, D. & Zafeiriou, S. Large scale 3D morphable models. *International Journal of Computer Vision*. **126**, 233–254 (2018).
16. Dall'Asta, A. *et al.* Quantitative analysis of fetal facial morphology using 3D ultrasound and statistical shape modeling: a feasibility study. *American journal of obstetrics and gynecology* 217, 76. e71-76. e78(2017).
17. Zhou, Y. & Zafeiriou, S. Deformable models of ears in-the-wild for alignment and recognition. in 2017 *12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)* 626–633 (IEEE, 2017).
18. Amberg, B., Romdhani, S. & Vetter, T. Optimal step nonrigid ICP algorithms for surface registration. in 2007 *IEEE Conference on Computer Vision and Pattern Recognition* 1–8 (IEEE, 2007).
19. Ploumpis, S., Wang, H., Pears, N., Smith, W. A. & Zafeiriou, S. Combining 3d morphable models: A large scale face-and-head model. in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition 10934–10943(2019).
20. Gong, S., Chen, L., Bronstein, M., Zafeiriou, S. & Spiralnet++ A fast and highly efficient mesh convolution operator. in Proceedings of the IEEE International Conference on Computer Vision Workshops 0–0(2019).
21. Maaten, L. & Hinton, G. Visualizing data using t-SNE. *Journal of machine learning research*. **9**, 2579–2605 (2008).

Figures

Figure 1

t-SNE embeddings of face-only (a), head-only (b), and combined head-and-face (c) models from left to right respectively. In all cases, distinct clusters emerge for healthy and syndromic samples. The clearest disambiguation between samples is observed for the face-only model.

Figure 2

Confusion matrices for the face-only, head-only, and combined head-and-face models in order from left to right. Top row: Binary classification. Bottom row: Multi-class classification.

Figure 3

t-SNE embeddings for all samples, including the atypical Crouzon case, for the face-only (a), head-only (b), and combined head-and-face (c) models.

Figure 4

Mesh templates. a) Face only template, b) Head and Face template, c) Head only template.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SupplementaryMaterials.docx](#)