



UNIVERSITY OF LEEDS

This is a repository copy of *Image-Derived Phenotype Extraction for Genetic Discovery via Unsupervised Deep Learning in CMR Images*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/187260/>

Version: Accepted Version

Proceedings Paper:

Bonazzola, R, Ravikumar, N orcid.org/0000-0003-0134-107X, Attar, R et al. (3 more authors) (2021) Image-Derived Phenotype Extraction for Genetic Discovery via Unsupervised Deep Learning in CMR Images. In: MICCAI 2021: Medical Image Computing and Computer Assisted Intervention. MICCAI 2021, 27 Sep - 01 Oct 2021, Strasbourg, France. Springer International Publishing , pp. 699-708. ISBN 9783030872397

https://doi.org/10.1007/978-3-030-87240-3_67

© 2021 Springer Nature Switzerland AG. This is an author produced version of a conference paper, published in MICCAI 2021: Medical Image Computing and Computer Assisted Intervention. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Image-derived phenotype extraction for genetic discovery via unsupervised deep learning in CMR images

Rodrigo Bonazzola^{1,2}, Nishant Ravikumar^{1,2}, Rahman Attar³, Enzo Ferrante⁴,
Tanveer Syeda-Mahmood⁵, and Alejandro F Frangi^{1,2}

¹ Centre for Computational Imaging and Simulation Technologies in Biomedicine (CISTIB), School of Computing, University of Leeds, Leeds, UK

`r.bonazzola1@leeds.ac.uk`

² Department of Bioengineering, Faculty of Engineering, Imperial College London, London, UK

³ Center for Computational Imaging and Simulation Technologies in Biomedicine, School of Medicine, University of Leeds, Leeds, UK

⁴ Research Institute for Signals, Systems and Computational Intelligence, sinc(i), FICH-UNL / CONICET, Santa Fe, Argentina

⁵ IBM Almaden Research Center, San Jose, USA

Abstract. Prospective studies with linked image and genetic data, such as the UK Biobank (UKB), provide an unprecedented opportunity to extract knowledge on the genetic basis of image-derived phenotypes. However, the extent of phenotypes tested within so-called genome-wide association studies (GWAS) is usually limited to handcrafted features, where the main limitation to proceed otherwise is the high dimensionality of both the imaging and genetic data. Here, we propose an approach where the phenotyping is performed in an unsupervised manner, via autoencoders that operate on image-derived 3D meshes. Therefore, the latent variables produced by the encoder condense the information related to the geometry of the biologic structure of interest. The network’s training proceeds in two steps: the first is genotype-agnostic and the second enforces an association with a set of genetic markers selected via GWAS on the intermediate latent representation. This genotype-dependent optimisation procedure allows the refinement of the phenotypes produced by the autoencoder to better understand the effect of the genetic markers encountered. We tested and validated our proposed method on left-ventricular meshes derived from cardiovascular magnetic resonance images from the UKB, leading to the discovery of novel genetic associations that, to the best of our knowledge, had not been yet reported in the literature on cardiac phenotypes.

1 Introduction

The emergence of population-scale prospective studies with linked imaging and genetic data, such as the UK Biobank [7], has enabled research into the genetics of image-derived phenotypes, a field called imaging genetics. One of the

main challenges of this field is the high dimensionality of both the imaging and genetic datasets. This problem is usually addressed on the imaging side by deriving handcrafted features from the images, based on prior expert knowledge supporting their clinical relevance. For example, in the case of cardiac images, these phenotypes could be the volumes of the different cardiac chambers, the myocardial mass, or functional parameters such as the ejection fraction.

In this work, we propose a different approach, based on unsupervised learning, to extract phenotypes from image-derived 3D meshes which describe the organs of interest’s geometry. This approach is outlined in figure 1. Its input consists of

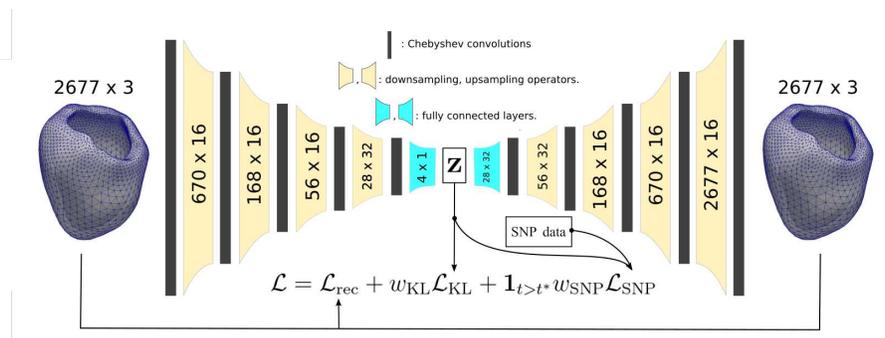


Fig. 1. Scheme of the encoder-decoder network implemented in this work. $\mathbf{1}$ is the indicator function. The numbers $V \times C$ within the downsampling and upsampling operators represent the dimension of their output, where V is the number of vertices and C is the number of channels.

a set of 3D meshes representing the geometry of the organs of interest, obtained by previous segmentation of the images. A graph-convolutional autoencoder is trained for t^* epochs to perform dimensionality reduction on the 3D meshes, without any genetic data being input to the network. Each component of the latent representation found by the autoencoder, which will encode the modes of variation of the set of meshes, are employed as a phenotype in a genome-wide association study (GWAS), where a reduced set of candidate loci are chosen based on a significance criterion. Then, the autoencoder is further trained adding to the cost function a term that enforces an association between each of the chosen loci and the corresponding latent variables; this step aims to produce a fine-tuning of the phenotype. Finally, the latent representation is tested in a GWAS using an independent set. We applied this approach to cardiovascular magnetic resonance (CMR) images from the UK Biobank, where the left ventricle (LV) at end-diastole was our object of study.

Our contributions are, therefore, two-fold: on the one hand, we propose an approach to study the genetic basis of image-derived phenotypes, based on unsupervised deep learning; on the other hand, we discover novel genetic associations

that have not been previously reported in the literature on cardiac phenotypes, to the best of our knowledge.

1.1 Related work

Other studies using unsupervised approaches to derive phenotypes have been performed. Studies on the genetic basis of 3D-mesh-derived facial features have been published [8], which used hierarchical clustering to obtain distinctive regions of the face. Each of these regions was projected onto a linear space using principal component analysis (PCA), and canonical correlation analysis (CCA) was used to find linear combinations of region-specific PCs that are maximally correlated with each of the single nucleotide polymorphisms (SNPs) tested.

On the other hand, several studies have been published in the recent years in the field of cardiac imaging genetics using CMR data, which focused on hand-crafted phenotypes. In [3], the authors investigate left-ventricular wall thickness at end-diastole, performing an association test with a set of genetic variants in a vertex-by-vertex fashion. Other studies ([2] and [12]) perform GWAS on global LV phenotypes: chamber volume at end-diastole and end-systole, stroke volume, ejection fraction and myocardial mass.

Also, work has been performed on CMR images to extract biomarkers using unsupervised dimensionality reduction approaches [4]. Still, this work has not relied on an intermediate 3D mesh representation of the cardiac chambers, nor has been used for the purpose of genetic discovery.

2 Methods

2.1 Description of the data

All the data used for this work comes from the UK Biobank project, data accession number 11350.

Cardiovascular Magnetic Resonance (CMR) data. The CMR imaging protocol used to obtain the raw imaging data is described elsewhere [11]. For a given individual and time point, this data consists of a stack of 10-12 short-axis view slices (SAX) along with three long-axis view (LAX) slices. The cardiac segmentation algorithm utilised is described in detail in [1]. This algorithm produces as output a set of registered meshes, i.e. meshes with the same number of vertices and the same connectivity. For LV, the meshes encompass the endocardial and epicardial surfaces. As mentioned before, in this work we studied only end-diastole.

The LV mesh for subject i , $i = 1, \dots, N$, can then be represented as pairs (\mathbf{S}_i, A) , where $\mathbf{S}_i = [x_{i1} \ y_{i1} \ z_{i1} \ | \ \dots \ | \ x_{iM} \ y_{iM} \ z_{iM}] \in \mathbb{R}^{M \times 3}$ is the shape and A is the adjacency matrix of the mesh. The number of individuals and mesh vertices are $N = 29051$ and $M = 2677$, respectively.

Genotype data SNP microarray data is available for all the individuals in the UK Biobank cohort. This microarray covers $\sim 800\text{k}$ genetic variants including SNPs and short indels. The design of this microarray has been described in detail in [6]. An augmented set of $\sim 9.5\text{M}$ variants was obtained from these genotyped markers through imputation, after filtering by a minor allele frequency (MAF) threshold of 1%, a Hardy-Weinberg equilibrium p -value threshold of 10^{-5} and an imputation info score of 0.3. Also, only the autosomes (chromosomes 1 through 22) were used.

2.2 Graph-convolutional autoencoder

To perform dimensionality reduction, we propose using an encoder-decoder approach. The encoder E consists of convolutional and pooling layers, whereas D consists of unpooling layers. To leverage the topology of the mesh, we utilise graph-convolutional layers. Since the vertices are not in a rectangular grid, the usual convolution, pooling and unpooling operations defined for such geometry are not adequate for this task and need to be suitably generalised. There are several methods to achieve this, but they all can be classified into two large groups: spatial or spectral [14]. In this work we applied a method belonging to the latter category, which relies on expressing the features in the Fourier basis of the graph, as will be explained below after providing some background of spectral graph theory.

The Laplace-Beltrami operator of a graph with adjacency matrix A is defined as $L = D - A$, where D is the degree matrix, i.e. a diagonal matrix where $D_{ii} = \sum_j A_{ij}$ is the number of edges connected to vertex i . The Fourier basis of the graph can be obtained by diagonalising the Laplace operator, $L = U^t A U$. The columns of U constitute the Fourier basis, and the operation of convolution \star for a graph can be defined in the following manner

$$x \star y = U(U^t x \odot U^t y), \quad (1)$$

where \odot is the element-wise product (also known as Hadamard product).

All spectral methods for convolution rely on this definition, and differ from one another in the form of the kernel. In this work, a parameterisation proposed in [9] was used. The said method is based on the Chebyshev family of polynomials $\{T_i\}$. The kernel g_ξ is defined as:

$$g_\xi(L) = \sum_{i=1}^K \xi_i T_i(L), \quad (2)$$

where K is the highest degree of the polynomials considered (in this work $K = 6$). Chebyshev polynomials have the advantage of being computable recursively through the relation $T_i(x) = xT_{i-1}(x) - T_{i-2}(x)$ and the base cases $T_1(x) = 1$

and $T_2(x) = x$. It is also worth mentioning that the filter described by equation 2, despite its spectral formulation, has the characteristic of being local.

Following [13], each of the three spatial coordinates of each vertex are input as a separate channel of the autoencoder. Downsampling and upsampling operations used in this study are based on a surface simplification algorithm proposed in [10]. These operations are defined before training each layer, using a single template shape. Here we utilise the mean shape $\bar{\mathbf{S}} = (1/N) \sum_{i=1}^N \mathbf{S}_i$ as a template.

2.3 GWAS

According to the traditional GWAS scheme [5], we tested each genetic variant l , with dosage $X_l \in [0, 2]$, for association with each of the LV latent features z_j through a univariate linear model $z_k = \beta_{lk}X_l + \epsilon_{lk}$, where ϵ_{lk} is the component not explained by the genotype, which we model as a normal random variable. The null hypothesis tested is that $\beta_{lk} = 0$. From linear regression, one obtains an estimate $\hat{\beta}_{lk}$ of the effect size β_{lk} , along with the standard error of this estimate, $se(\hat{\beta}_{lk})$. Finally, the p -value for the association can be computed from these values.

Before GWAS, the phenotypes (i.e. latent variables) were adjusted for a set of covariates: height, BMI, age, sex, diastolic and systolic blood pressure. For this, a multilinear model is used, and the new phenotypes are the residues obtained from there. Also, rank-based inverse normalisation is performed on these phenotypes so that the usual closed-form formulas for hypothesis testing can be utilised.

To avoid issues related to population stratification, only individuals with British ancestry were utilised, leading to the aforementioned sample size of $N = 29051$. No filtering was performed based on pathologies.

2.4 Proposed algorithm

The proposed method is described in Algorithm 1.

Data: 3D meshes \mathbf{S}_i and linked genotype dosages X_{il}
Result: Network weights, GWAS summary statistics.
Hyperparameters: network architecture, $w_{\text{KL}} > 0$, $w_{\text{SNP}} > 0$.
ProcrustesAlignment(\mathbf{S}_i);
PartitionDataset(\mathbf{S}_i, X_{il});
InitialiseWeights();
while *Stop criterion is not met* **do**
| Perform optimisation step with loss \mathcal{L}_1 ;
end
Select the best epoch within the validation set;
for *each latent variable k* **do**
| Perform GWAS within the training set;
| Extract set of genetic markers S_k based on a significance criterion;
| **for** *each SNP l_k in S_k* **do**
| | **if** $\text{Corr}(X_{l_k}, z_k) < 0$ **then**
| | | $X_{l_k} \mapsto 2 - X_{l_k}$
| | **end**
| **end**
end
while *Stop criterion is not met* **do**
| Perform optimisation step with loss \mathcal{L}_2 ;
end
Select the best epoch within the validation set;
Perform GWAS within the test set;
Algorithm 1: Workflow of the proposed method.

Loss function The loss function \mathcal{L}_1 for the first training stage consists of two terms:

$$\mathcal{L}_1 = \mathcal{L}_{\text{rec}} + w_{\text{KL}}\mathcal{L}_{\text{KL}}, \quad (3)$$

where \mathcal{L}_{rec} is the reconstruction loss and \mathcal{L}_{KL} is the variational regularisation term, computed as the Kullback-Leibler divergence of the latent representation \mathbf{z} with an isotropic normal distribution. For the second training stage, another term is added, \mathcal{L}_{SNP} , which encourages a stronger association between each of the latent variables k and a set of SNPs S_k :

$$\mathcal{L}_2 = \mathcal{L}_1 + w_{\text{SNP}}\mathcal{L}_{\text{SNP}} \quad (4)$$

$$\mathcal{L}_{\text{SNP}} = - \sum_k \sum_{l_k \in S_k} \text{Corr}(X_{l_k}, z_k) \quad (5)$$

During training, the correlation in equation 5 is computed as the sample correlation for each batch. The reference and alternative alleles are chosen so that the correlation is positive within the training set⁶.

In this work, our criterion for SNP selection was passing the Bonferroni significance threshold (taken as the usual genome-wide threshold of 5×10^{-8} divided by the number of latent variables tested). As we explain below, this criterion led to a single term in the summation.

2.5 Implementation details

The architecture of the network is detailed in figure 1. After each convolutional layer, a ReLU activation function was applied. Importantly, the latent representation’s size $\dim(\mathbf{z})$ was chosen as 4. We found this number enough to capture the most salient global features explaining the variability in LV shape. The reconstruction loss employed was the vertex-wise mean squared error (MSE), averaged across the vertices of each mesh. An ablation study was performed to assess the impact of different parameters: learning rate γ , w_{KL} and w_{SNP} . γ was chosen as 10^{-3} , whereas for the first training stage $w_{\text{KL}} = 10^{-3}$. Batch size was 100. Also, in the second training stage and for each parameter configuration, experiments with different seeds were performed. These seeds controlled the partition of the full dataset into training, validation and test sets. The sizes of these sets were 5000, 1000 and 23051, respectively.

3 Results and discussion

After the first training stage, GWAS was performed on each of the 4 components of \mathbf{z} . Only one of them, z_1 , yielded a Bonferroni-significant association. Figure 2 shows both the morphological impact of changes in this variable, and the Manhattan plot displaying the GWAS p -values. The association lies on chromosome 6, and we mapped it to gene PLN with high confidence, based on the literature on the genetics of cardiac phenotypes.

As can also be determined visually, latent variable z_1 was found to correlate with LV sphericity index s . Spearman correlation between s and z_1 is 0.432 (computed across real meshes). The index s was calculated as $s = A_{\text{sph}}(V_{\text{CH}})/A_{\text{CH}}$, where $A_{\text{sph}}(V) = (36\pi V^2)^{\frac{1}{3}}$ is the surface area of a sphere of volume V , whereas A_{CH} and V_{CH} are the surface area of and the volume enclosed by the convex hull of the LV, respectively. To the best of our knowledge, this association between PLN and LV sphericity had not been previously reported.

For the second training stage, the effect of the \mathcal{L}_{SNP} term on the strength of the genetic association was studied. \mathcal{L}_{SNP} consists of a single term, corresponding to the leading SNP in the PLN locus, rs11153730. Results are displayed in figure 3. Each observation corresponds to a different experiment. In each experiment, the sets of training and testing samples varies according to the random seed,

⁶ Swapping the alleles corresponds to performing the transformation $X_l \mapsto 2 - X_l$.

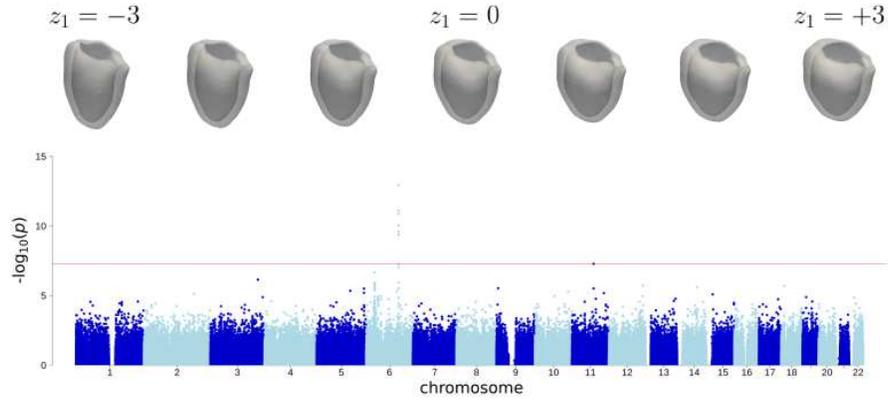


Fig. 2. Effect of latent variable z_1 (before re-training) and corresponding Manhattan plot.

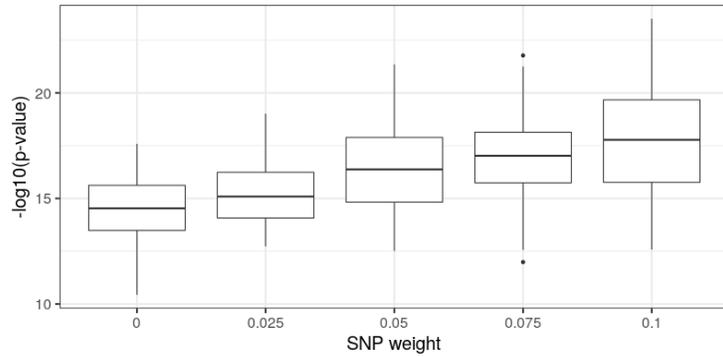


Fig. 3. Distribution of p -values for the z_1 -rs11153730 associations after re-training, for experiments with different values of w_{SNP} with $w_{\text{KL}} = 0.1$. Each box contains 60 ± 10 experiments.

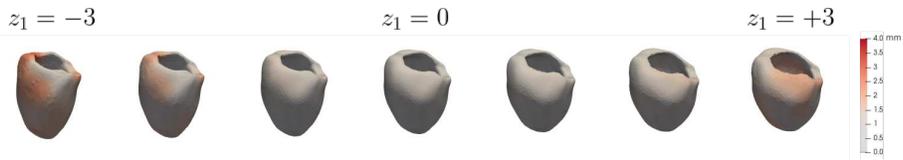


Fig. 4. Morphologic effect of latent variable z_1 , after re-training. Color represents the deviations in shape with respect to the meshes from figure 2.

but their size is fixed. We confirmed that the additional training is beneficial for obtaining a better phenotype, i.e. a phenotype that yields a stronger association on the GWAS. It is worth noting that the effectiveness of this procedure was not obvious in advance, since individual SNPs’ effect sizes on complex traits are generally small. We studied the change in shape modeled by the z_1 latent variable after retraining with $w_{\text{SNP}} = w_{\text{KL}} = 0.1$, for one particular experiment. This change is shown in figure 4. The w_{KL} coefficient was increased with respect to the first training stage to avoid overfitting of the \mathcal{L}_{SNP} term.

4 Conclusions

We have proposed and validated an approach based on graph-convolutional autoencoders to extract phenotypes from 3D meshes of biological structures, for genetic discovery through GWAS; and applied it to CMR-derived left-ventricular meshes.

In particular, a genetic association was found between a genetic locus linked to PLN gene and a latent variable that correlates with LV sphericity. To the best of our knowledge, this had not been reported before, even though mutations PLN have been linked to dilated cardiomyopathy, a disease characterised by an increase in LV sphericity. Furthermore, we have shown that an additional stage of training that aims to refine the phenotype effectively improves the genetic association, and is therefore recommended.

On the other hand, we acknowledge that our method as presented here results in a loss of information as compared to the original CMR images, since it only captures the LV myocardial surface’s geometry. However, this is not a drawback of the method itself, since texture information can be added seamlessly as node features (in addition to spatial coordinates). Besides, diseases could be investigated in this context by encouraging the latent representation to be discriminative of disease status, whenever diagnosis information is available. Finally, UKB provides images for the whole cardiac cycle, from which dynamic patterns could be extracted. These ideas are left as future work.

Acknowledgments

This work was funded by the following institutions: The Royal Academy of Engineering [grant number CiET1819\19] and EPSRC [TUSCA EP/V04799X/1] (R.B., N.R. and A.F.F.), The Royal Society through the International Exchanges 2020 Round 2 scheme [grant number IES\R2\202165] (R.B., E.F. and A.F.F). E.F. was also funded by ANPCyT [grant number PICT2018-3907] and UNL [grant numbers CAI+D 50220140100-084LI, 50620190100-145LI].

References

1. Attar, R., Pereañez, M., Gooya, A., Albà, X., Zhang, L., de Vila, M.H., Lee, A.M., Aung, N., Lukaschuk, E., Sanghvi, M.M., Fung, K., Paiva, J.M., Piechnik, S.K., Neubauer, S., Petersen, S.E., Frangi, A.F.: Quantitative CMR population imaging on 20,000 subjects of the UK Biobank imaging study: LV/RV quantification pipeline and its evaluation. *Medical Image Analysis* **56**, 26–42 (Aug 2019). <https://doi.org/10.1016/j.media.2019.05.006>
2. Aung, N., Vargas, J.D., Yang, C., Cabrera, C.P., Warren, H.R., Fung, K., Tzanis, E., Barnes, M.R., Rotter, J.I., Taylor, K.D., Manichaikul, A.W., Lima, J.A., Bluemke, D.A., Piechnik, S.K., Neubauer, S., Munroe, P.B., Petersen, S.E.: Genome-Wide Analysis of Left Ventricular Image-Derived Phenotypes Identifies Fourteen Loci Associated With Cardiac Morphogenesis and Heart Failure Development. *Circulation* **140**(16), 1318–1330 (Oct 2019). <https://doi.org/10.1161/CIRCULATIONAHA.119.041161>
3. Biffi, C., de Marvao, A., Attard, M.I., Dawes, T.J.W., Whiffin, N., Bai, W., Shi, W., Francis, C., Meyer, H., Buchan, R., Cook, S.A., Rueckert, D., O’Regan, D.P.: Three-dimensional cardiovascular imaging-genetics: a mass univariate framework. *Bioinformatics* **34**(1), 97–103 (Jan 2018). <https://doi.org/10.1093/bioinformatics/btx552>
4. Biffi, C., Oktay, O., Tarroni, G., Bai, W., De Marvao, A., Doumou, G., Rajchl, M., Bedair, R., Prasad, S., Cook, S., et al.: Learning interpretable anatomical features through deep generative models: Application to cardiac remodeling. In: *International conference on medical image computing and computer-assisted intervention*. pp. 464–471. Springer (2018)
5. Bush, W.S., Moore, J.H.: Chapter 11: Genome-wide association studies. *PLOS Computational Biology* **8**(12), 1–11 (12 2012). <https://doi.org/10.1371/journal.pcbi.1002822>, <https://doi.org/10.1371/journal.pcbi.1002822>
6. Bycroft, C., Freeman, C., Petkova, D., Band, G., Elliott, L.T., Sharp, K., Motyer, A., Vukcevic, D., Delaneau, O., O’Connell, J., Cortes, A., Welsh, S., McVean, G., Leslie, S., Donnelly, P., Marchini, J.: Genome-wide genetic data on ~500,000 UK Biobank participants. *bioRxiv* (2017). <https://doi.org/10.1101/166298>
7. Bycroft, C., Freeman, C., Petkova, D., Band, G., Elliott, L.T., Sharp, K., Motyer, A., Vukcevic, D., Delaneau, O., O’Connell, J., Cortes, A., Welsh, S., Young, A., Effingham, M., McVean, G., Leslie, S., Allen, N., Donnelly, P., Marchini, J.: The UK Biobank resource with deep phenotyping and genomic data. *Nature* **562**(7726), 203–209 (Oct 2018). <https://doi.org/10.1038/s41586-018-0579-z>

8. Claes, P., Roosenboom, J., White, J.D., Swigut, T., Sero, D., Li, J., Lee, M.K., Zaidi, A., Mattern, B.C., Liebowitz, C., et al.: Genome-wide mapping of global-to-local genetic effects on human facial shape. *Nature genetics* **50**(3), 414–423 (2018)
9. Defferrard, M., Bresson, X., Vandergheynst, P.: Convolutional neural networks on graphs with fast localized spectral filtering. In: *Advances in Neural Information Processing Systems* (2016), <https://arxiv.org/abs/1606.09375>
10. Garland, M., Heckbert, P.S.: Surface simplification using quadric error metrics. In: *Proceedings of the 24th annual conference on Computer graphics and interactive techniques - SIGGRAPH '97*. pp. 209–216. ACM Press (1997). <https://doi.org/10.1145/258734.258849>
11. Petersen, S.E., Matthews, P.M., Bamberg, F., Bluemke, D.A., Francis, J.M., Friedrich, M.G., Leeson, P., Nagel, E., Plein, S., Rademakers, F.E., Young, A.A., Garratt, S., Peakman, T., Sellors, J., Collins, R., Neubauer, S.: Imaging in population science: cardiovascular magnetic resonance in 100,000 participants of UK Biobank - rationale, challenges and approaches. *Journal of Cardiovascular Magnetic Resonance* **15**(1), 46 (Dec 2013). <https://doi.org/10.1186/1532-429X-15-46>
12. Pirruccello, J.P., Bick, A., Wang, M., Chaffin, M., Friedman, S., Yao, J., Guo, X., Venkatesh, B.A., Taylor, K.D., Post, W.S., Rich, S., Lima, J.A.C., Rotter, J.L., Philippakis, A., Lubitz, S.A., Ellinor, P.T., Khera, A.V., Kathiresan, S., Aragam, K.G.: Analysis of cardiac magnetic resonance imaging in 36,000 individuals yields genetic insights into dilated cardiomyopathy. *Nature Communications* **11**(1) (Dec 2020). <https://doi.org/10.1038/s41467-020-15823-7>
13. Ranjan, A., Bolkart, T., Sanyal, S., Black, M.J.: Generating 3D Faces Using Convolutional Mesh Autoencoders. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *Computer Vision – ECCV 2018*, vol. 11207, pp. 725–741. Springer International Publishing, Cham (2018). https://doi.org/10.1007/978-3-030-01219-9_43
14. Zhou, J., Cui, G., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., Sun, M.: Graph neural networks: A review of methods and applications. *arXiv preprint arXiv:1812.08434* (2018)