

Ear Cartilage Inference for Reconstructive Surgery with Convolutional Mesh Autoencoders

Eimear O’ Sullivan¹, Lara van de Lande², Antonia Osolos², Silvia Schievano²,
David J. Dunaway², Neil Bulstrode^{2,3}, and Stefanos Zafeiriou¹

¹ Department of Computing, Imperial College London, London, UK

² UCL Great Ormond Street Institute of Child Health, London, UK

³ Craniofacial Unit, Great Ormond Street Hospital for Children, London, UK

Abstract. Many children born with ear microtia undergo reconstructive surgery for both aesthetic and functional purposes. This surgery is a delicate procedure that requires the surgeon to carve a “scaffold” for a new ear, typically from the patient’s own rib cartilage. This is an unnecessarily invasive procedure, and reconstruction relies on the skill of the surgeon to accurately construct a scaffold that best suits the patient based on limited data. Work in stem-cell technologies and bioprinting present an opportunity to change this procedure by providing the opportunity to “bioprint” a personalised cartilage scaffold in a lab. To do so, however, a 3D model of the desired cartilage shape is first required. In this paper we optimise the standard convolutional mesh autoencoder framework such that, given only the soft tissue surface of an unaffected ear, it can accurately predict the shape of the underlying cartilage. To prevent predicted cartilage meshes from intersecting with, and protruding through, the soft tissue ear mesh, we develop a novel intersection-based loss function. These combined efforts present a means of designing personalised ear cartilage scaffold for use in reconstructive ear surgery.

Keywords: Mesh Autoencoders · Ear reconstruction · Cartilage.

1 Introduction

Microtia is a congenital condition characterized by underdevelopment of the external ear and affects 1 in 2,000 births [9]. Microtia can be unilateral or bilateral, and occurs with varying degrees of severity, from mild abnormalities in the shape of the external ear to a complete absence of both the ear and the auricular canal [3]. Typically, microtia is treated between the ages of 9 and 12, for both functional and aesthetic purposes.

The current gold standard surgical intervention for ear reconstruction in microtia cases is a two-stage process. The first stage requires the implantation of a “scaffold” to provide the frame for a new external ear. This scaffold is constructed from the patients own rib cartilage and placed under the skin where the underdeveloped ear used to be. It remains here for a number of weeks to accommodate the stretch of the surrounding skin. A second procedure is then

performed to release the scaffold from the skin to form a normal ear shape [14]. The drawbacks of this approach include the increased co-morbidity and potential for chest wall depression where the cartilage is harvested. Tissue engineered scaffolds, and non-absorbable auricular frameworks have also been used [8, 33]. These scaffolds are made by molding biodegradable polymers into the mirror image shape of the intact contralateral ear and are then seeded with cartilage excised from the patient to encourage biocompatibility prior to subcutaneous implantation [24]. Ongoing work with stem cell tissue engineering and biofabrication indicate that it could soon be possible to bioprint 3D ear scaffolds or human patients, reducing the requirements for invasive cartilage harvesting and complex tissue engineered solutions [17, 20]. While it is unclear when these advances will be available for human adoption, promising results have been seen in nude mice [20]. With this in mind, we aim to address the following question; given the ideal shape of a reconstructed ear, can we predict the underlying scaffold required that would need to be bioprinted to produce this shape?

Recent works in the field of geometric deep learning [6] have given rise to a wide variety of approaches that can be used for analysis and learning on 3D mesh structures. Mesh autoencoders [28] are of particular interest, as they allow 3D shapes to be encoded to a single latent vector representation. Highly accurate reconstructions can then be obtained by decoding these latent representations. Most, if not all, recent works on mesh autoencoders focus on compressing and then reconstructing a given shape, typically the human body, face and head [5, 16, 28]. Here, we aim to adapt the standard autoencoder framework to enable the encoding of one shape class and the decoding of another, thereby providing a means of inferring the structure of the cartilage lying beneath the soft tissue of the human ear. To enforce the intuition that the reconstructed cartilage mesh should lie beneath the surface of the skin, or soft tissue, of the ear, we propose a new loss function, which we refer to as an *intersection loss*.

In summary, we present a method to determine the optimal scaffold shape required for high quality ear reconstruction by adapting the standard mesh autoencoder framework such that it can be used to encode a specific structure (the ear soft tissue), and decode an alternative structure (the ear cartilage). We introduce a combined ear-and-cartilage dataset. In addition, we propose a novel intersection loss function that aims to ensure that all cartilage is within the outer surface of the soft tissue mesh. Given a lack of ground truth data for the ear scaffolds that correspond to certain ears, we present our method using actual ear cartilage and demonstrate that it is highly accurate, even with relatively small number of samples. Going forward, this makes a strong case for the use of such a method to predict the required scaffold to construct a desired ear shape.

2 Related Work

Geometric Deep Learning: Geometric deep learning is an umbrella term for the emerging techniques that facilitate the generalisation of deep neural networks and machine learning approaches to non-Euclidean domains, such as

high-dimensional graphs and manifolds, and 3D mesh structures [6]. Common applications include node classification [22], molecule prediction [30], and determining mesh correspondences [4]. Multiple approaches have been presented to generalise standard 2D convolutional operators for higher-dimensional domains and irregular graph structures, [7, 15, 23, 29]. New methods for mesh pooling, or graph coarsening, have also been developed [10, 12, 32].

Mesh Autoencoders: Convolutional mesh autoencoders have proven to be a highly effective means of modelling 3D data. Among the first of these was the COMA model proposed by Ranjan *et.al.* [28]. By storing the barycentric coordinates of vertices removed during mesh pooling, or encoding, and using these during the unpooling, or decoding, stage, the pooled mesh could be restored to its original topology. Truncated Chebyshev spectral convolutions were used [10]. Bouritsas *et.al.* [5] later improved on this by replacing the initial spectral convolutions with anisotropic spiral convolutions to achieve more accurate mesh reconstructions. The spiral convolution was further adapted by [16], providing a mesh convolution operator that enables fast and highly efficient training. The combination of convolutional mesh autoencoders with the spiral convolution operator has been shown to be well suited to small datasets; this was demonstrated both [5] and [16], where high reconstruction accuracies were achieved when only 80 samples were used for training.

3D Loss Functions: There exist a variety of loss functions that are typically used for 3D mesh reconstruction in geometric deep learning. The standard L1 loss function can be applied to mesh vertices, and is common when processing meshes that are in correspondence [28]. A normal consistency loss can be applied to encourage reconstructed normals to match those of the original mesh. Where exact correspondences are unavailable, the chamfer distance is often used to encourage similarity between two distinct point clouds [13]. A Laplacian objective can be used to encourage the reconstruction of smooth meshes [11, 25], while the purpose of a mesh edge loss is to encourage edge length regularization. Where the reconstruction of sharp features, such as edges and corners, is desirable, a quadric loss has been shown to be beneficial [1].

While these losses have proven effective in many applications, they were not designed to enforce a particular relationship between two given surfaces. We aim to address this by proposing an intersection loss. This loss aims to ensure that points in the reconstructed sample lie on the desired side of a given plane. In the case of ear cartilage reconstruction, this helps to enforce a fundamental human understanding; that the ear cartilage should lie *within* the surface of the soft tissue of the ear.

3 Dataset

While a number of datasets for the human ear soft tissue exist [18, 27, 31], there are, to the best of our knowledge, none that capture the paired relationship between ear cartilage and soft tissue. To address this, we construct a database of these paired samples. Data was collected from patients aged between 9 and

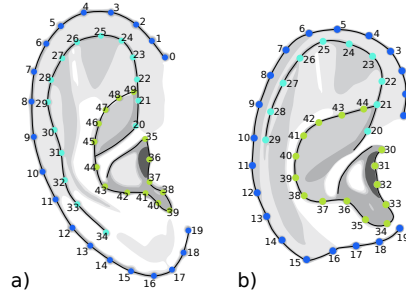


Fig. 1. a) Soft tissue (adapted from [34]) and b) Cartilage landmark sets.

12 referred for a Computed Tomography (CT) scan and were reported without any abnormalities. The mean CT slice thickness and pixel size was $1mm$. The soft tissue and ear cartilage segmentations were acquired using Mimics InPrint⁴. The range used for soft tissue segmentation was -200 to $+200$ Hounsfield units. The ear cartilage, which has a density between that of bone and soft tissue and which can vary depending on the age of the patient, was segmented by an expert clinician. The best range for ear cartilage segmentation was found to be between -30 and -10 Hounsfield units. Although difficulties regarding the segmentation of cartilage from a CT scan limited the number of available samples, a dataset of 100 paired ear-and-cartilage instances from 68 individuals was collected.

All left meshes were mirrored so that they lay in the plane of the right ear samples. Landmarks were manually applied to both soft tissue and cartilage meshes using a set of 50 and 45 landmarks respectively, as shown in Fig. 1. The ear landmarks from [34] were used for the soft tissue, and adapted to create the cartilage landmark set. Soft tissue meshes were rigidly aligned with a soft tissue ear template mesh. The template mesh used was a non-watertight triangular mesh, consisting of 2800 vertices, and was the same as that used in [27]. The same rigid transformation was applied to the cartilage, to maintain the spatial relationship between paired cartilage and soft tissue meshes. A cartilage template mesh, containing the same number of vertices and an identical mesh topology to that of the soft tissue, was then rigidly aligned with the raw cartilage mesh. A Non-rigid Iterative Closest Point (NICP) procedure [2] was used to obtain mesh correspondences for both cartilage and soft tissue. Throughout the data processing, care was taken to ensure the spatial relationship between the soft tissue and cartilage meshes was maintained.

4 Network Architecture

Our mesh autoencoder architecture is based on that of Gong *et.al* [16]. The encoder is comprised of three convolutional layers, followed by mesh pooling.

⁴ Mimics Inprint, Materialise, Leuven, Belgium




Fig. 2. Network architecture. \mathbb{R}^{l_v} is the size of the latent vector.

The decoder architecture is identical to that of the encoder, but in reverse. 16 latent vectors were used for the encoding. Spiral convolutions with a length of 9 are used in all layers [16]. An overview of the architecture is shown in Fig. 2.

Mesh autoencoder frameworks typically use the barycentric coordinates of mesh vertices removed during the pooling stage to restore these vertices during the unpooling stage [28]. For computational efficiency, these values are calculated once prior to training, and then used throughout the training process. The autoencoder output is therefore a mesh with an identical topology to the input. When using an autoencoder to predict the shape of something other than the original input, however, the assumptions of identical mesh topology and barycentric coordinate locations do not necessarily hold true. Therefore, prior to training, we pool the cartilage mesh template independent of the soft tissue ear mesh and acquire the barycentric vertex locations specific to the cartilage. These are then used during the unpooling stage in place of those of the soft tissue ear mesh. This approach could be used to allow the decoder to return an arbitrary mesh topology, however, as the cartilage and soft tissue ear template have identical topologies, we make no such changes here.

4.1 Loss Functions

An L1 loss was applied to all reconstructed cartilage vertices. In addition to this, we also apply a loss to penalise the reconstruction if cartilage vertices are seen to protrude through the skin. This loss can be calculated using the normal direction of the vertices in cartilage and soft tissue meshes. For each vertex in the reconstructed cartilage, we find the closest vertex (nearest neighbour) in the corresponding soft tissue mesh. As the position of the predicted cartilage vertices with respect to their closest soft tissue vertex is likely to change many times during the training process, this cannot be pre-computed. For each point in the set of cartilage vertices, S_C , the closest vertex, x_s , in the set of soft tissue vertices, S_S , is instead calculated.

To determine whether the cartilage vertex intersects the surface of the soft tissue, the following formula can be applied:

$$d = n_s \cdot (x_c - x_s) \quad (1)$$

where d is the distance from the point to the soft tissue surface, x_c is a vertex in the cartilage mesh, x_s is the closest vertex in the soft tissue mesh, and n_s is the

corresponding normal. The sign of d gives additional insight into the cartilage point location with respect to the soft tissue; if $d = 0$, the cartilage point lies on the plane of the soft tissue vertex. When $d < 0$, the vertex lies on the opposite side of the plane into which the normal is directed and the vertex is within the soft tissue surface. If $d > 0$, the vertex lies on the side of the plane into which the normal is directed and an intersection has occurred. In this case, the value of d can be used to determine the extent of the intersection.

The L1 loss and Intersection loss were combined as in Equation 2, where λ_i and λ_{int} are the weighting values applied to the L1 and intersection losses respectively. Optimal performance was observed for $\lambda_i = 1$ and $\lambda_{int} = 10$.

$$l_{total} = \lambda_i l_1 + \lambda_{int} l_{int} \quad (2)$$

As the nearest neighbour search for each vertex in the input point cloud is independent, the calculations can be trivially parallelised and efficiently computed on a GPU.

4.2 Training

The autoencoder framework is implemented in PyTorch [26]. Models were trained with a batch size of 4. The Adam optimiser, initialised with a learning rate of 0.001, and learning rate decay of 0.99, was used [21]. Mesh downsampling was achieved using the QSlim approach, while mesh upsampling in the decoder was achieved as in [28]. Early stopping was used to prevent the model from overfitting to the training data. Spiral convolutions, with a fixed length of 9, are used in each layer of the encoder-decoder framework [16]. With 80 samples in the training database, and running for up to 300 epochs, training takes 91 seconds on average, equating to 0.327 seconds per epoch on an Nvidia Titan Xp GPU.

5 Experiments

A five-fold cross-validation procedure was used. If both the right and left ear data from a given subject were used, care was taken to ensure that both sides were in the same dataset to prevent data leakage between the two datasets. Beyond this, the data was split randomly between training and testing sets. Each training and test set therefore consisted of 80 and 20 samples respectively. Mean accuracies over all data splits are reported below. Ground truth in this study was considered to be the registered cartilage meshes. Errors were calculated as the Euclidean distance between corresponding vertices in the ground truth and predicted cartilage meshes.

5.1 Ablation Study

In this section we analyse the effects of the proposed loss functions on the reconstruction quality. Three cases were tested; using just the L1 loss, using a




Fig. 3. Predicted cartilage for a given soft tissue sample. Rows show the soft tissue ear mesh, ground truth cartilage, and predicted cartilage meshes. The final column displays the heatmap for the Euclidean error in millimeters (mm).

combination of the L1 loss and chamfer distance, and using the L1 loss with the proposed intersection loss. For a baseline comparison, we compare a scaled version of the soft tissue to the ground truth cartilage mesh. All soft tissue meshes were scaled by approximately 93%, equivalent to reducing the mesh by $1.5mm$ to allow for skin thickness [19]. The mean vertex error, and mean number of vertex intersections per mesh, are given in Table 1. Due to the small size of the human ear, results are given in micrometers (μm). Simply scaling the soft tissue does not give a close estimate of the ground truth cartilage mesh, and results in the greatest number of intersections of the cartilage though the surface of the soft tissue. As shown in the table, the introduction of the intersection loss plays a role in reducing the mean vertex error, in addition to reducing the number of cartilage vertices that intersect the soft tissue ear mesh. Interestingly, the worst results are observed when using the L1 loss in conjunction with the chamfer distance. Examples of the predicted cartilage meshes for a given soft tissue mesh are shown Fig. 3.

	Scaling		L1		L1 + Ch.		L1 + Int.	
	mean	ints.	mean	ints.	mean	ints.	mean	ints.
Test	324.6 ± 42.7	1683	135.1 ± 122.8	27	135.3 ± 123.0	28	134.7 ± 121.9	17
Train	—	—	578.4 ± 578.4	212	583.8 ± 413.2	215	571.5 ± 419.9	202

Table 1. Ablation Study: The effect of different loss functions on reconstruction accuracy and number of intersections. Errors are in micrometers (μm). The mean error, standard deviation, and mean number of intersections are on a per sample basis. *Chamfer Loss* is abbreviated as *Ch.*, *Intersection Loss* is abbreviated as *Int.*

	L1	L1 + Ch.	L1 + Int.
Test	210.3 ± 154.3	211.5 ± 154.7	208.9 ± 155.7
Train	712.9 ± 522.9	716.6 ± 528.4	702.6 ± 530.3

Table 2. Mean cartilage landmark error. A latent vector of size 16 was used. All values are reported in micrometers (μm).

5.2 Landmark Error

Table 2 reports the landmark error for the inferred cartilage meshes for each of the loss combinations. Minor discrepancies are noted between mean loss values for each of the implemented loss combinations, however using the intersection loss again leads to an increase in the overall accuracy. This indicates that the intersection loss discourages the predicted cartilage from protruding through the ear soft tissue, without negatively impacting the result of the reconstruction.

6 Discussion and Conclusion

Future work would look to expanding the combined ear-cartilage database to include a greater number of samples from more subjects. Neural networks tend to perform better with large quantities of data, and one of the main limitations of the system is the relatively small number of samples that can be used for training purposes; this is highlighted by the discrepancy between the reconstruction accuracies for the training and testing datasets. Though we have chosen a network architecture that has been shown to perform well on small datasets and have endeavoured to keep the number of parameters in the autoencoder small to reduce overfitting, this remains an area where further improvements could be achieved. In particular, the network has difficulties in accurately reconstructing the sharper ridges and external regions of the cartilage. This can be seen most clearly in the rows 1 and 3 of Fig. 3.

Determining the shape of the cartilage scaffold required for ear reconstruction is by no means trivial and it takes a trained and experienced surgeon to carve a cartilage scaffold from autologous rib cartilage. The ability to predict the shape of the required cartilage scaffold, paired with the continued advancements in bioprinting technologies presents an alternative approach to creating this scaffold. With the reduced requirements for cartilage to be harvested from the ribs of the patient, morbidity and operating times would likely decrease, and the risk to chest wall deformity would be reduced.

We have presented an autoencoder framework that can be used to infer the ear cartilage shape given that of the ear soft tissue. We adapt the conventional autoencoder framework so that the encoder framework can be used to encode and decode meshes with different topologies. Additionally, we present a novel loss function, the intersection loss, and show how this can be used to reduce the number of cartilage vertices that intersect, or protrude through, the soft tissue ear surface, thereby encouraging the inference of more realistic ear cartilage samples. The loss functions can be easily optimized for GPU, facilitating

fast training speeds. This approach can also allow for patient specific cartilage design; in the case of unilateral microtia, the shape of the soft tissue from the contralateral ear can be used for cartilage inference.

Acknowledgements. Stefanos Zafeiriou acknowledges support from EPSRC Fellowship DEFORM (EP/S010203/1).

References

1. Agarwal, N., Yoon, S.e., Gopi, M.: Learning embedding of 3D models with quadric loss (Jul 24, 2019)
2. Amberg, B., Romdhani, S., Vetter, T.: Optimal step nonrigid ICP algorithms for surface registration. pp. 1–8. IEEE (2007)
3. Bly, R., Bhrany, A., Murakami, C., Sie, K.: Microtia reconstruction. *Facial Plast. Surg. Clin. North Am.* **24**(4), 577–591 (Nov 2016)
4. Boscaini, D., Masci, J., Rodolà, E., Bronstein, M.M.: Learning shape correspondence with anisotropic convolutional neural networks (May 20, 2016), <https://arxiv.org/abs/1605.06437>
5. Bouritsas, G., Bokhnyak, S., Ploumpis, S., Zafeiriou, S., Bronstein, M.: Neural 3d morphable models: Spiral convolutional networks for 3d shape representation learning and generation. pp. 7212–7221. IEEE (Oct 2019)
6. Bronstein, M.M., Bruna, J., LeCun, Y., Szlam, A., Vandergheynst, P.: Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine* **34**(4), 18–42 (2017)
7. Bruna, J., Zaremba, W., Szlam, A., LeCun, Y.: Spectral networks and locally connected networks on graphs. *ICLR 2014* (Dec 20, 2013), <https://arxiv.org/abs/1312.6203>, first proposal of convolution operations on graphs in the spectral domain
8. Cao, Y., Vacanti, J.P., Paige, K.T., Upton, J., Vacanti, C.A.: Transplantation of chondrocytes utilizing a polymer-cell construct to produce tissue-engineered cartilage in the shape of a human ear. *Plastic and reconstructive surgery* **100**(2), 297–302 (Aug 1997)
9. Cubitt, J.J., Chang, L., Liang, D., Vandervord, J., Marucci, D.D.: Auricular reconstruction. *Journal of Paediatrics and Child Health* **55**(5), 512–517 (May 2019)
10. Defferrard, M., Bresson, X., Vandergheynst, P.: (Jun 30, 2016), <http://arxiv.org/abs/1606.09375>
11. Desbrun, M., Meyer, M., Schröder, P., Barr, A.: Implicit fairing of irregular meshes using diffusion and curvature flow. pp. 317–324. *SIGGRAPH ’99*, ACM Press/Addison-Wesley Publishing Co (Jul 1, 1999)
12. Diehl, F., Brunner, T., Truong Le, M., Knoll, A.: Towards graph pooling by edge contraction (2019), <https://graphreason.github.io/papers/17.pdf>
13. Fan, H., Su, H., Guibas, L.: A point set generation network for 3d object reconstruction from a single image. pp. 2463–2471. IEEE (Jul 2017)
14. Fattah, A., Sebire, N.J., Bulstrode, N.W.: Donor site reconstitution for ear reconstruction. *Journal of Plastic, Reconstructive & Aesthetic Surgery* **63**(9), 1459–1465 (2009)
15. Fey, M., Lenssen, J.E., Weichert, F., Muller, H.: Splinecnn: Fast geometric deep learning with continuous b-spline kernels. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition pp. 869–877 (Jun 2018)

16. Gong, S., Chen, L., Bronstein, M., Zafeiriou, S.: (Nov 13, 2019), <https://arxiv.org/abs/1911.05856>
17. Griffin, M.F., Ibrahim, A., Seifalian, A.M., Butler, P.E.M., Kalaskar, D.M., Ferretti, P.: Argon plasma modification promotes adipose derived stem cells osteogenic and chondrogenic differentiation on nanocomposite polyurethane scaffolds; implications for skeletal tissue engineering. *Materials Science & Engineering C* **105**, 110085 (Dec 2019)
18. Jin, C.T., Guillon, P., Epain, N., Zolfaghari, R., van Schaik, A., Tew, A.I., Hetherington, C., Thorpe, J.: Creating the sydney york morphological and acoustic recordings of ears database. *IEEE Transactions on Multimedia* **16**(1), 37–46 (Jan 2014)
19. Jung, B.K., Kim, J.Y., Kim, Y.S., Roh, T.S., Seo, A., Park, K.H., Shim, J.H., Yun, I.S.: Ideal scaffold design for total ear reconstruction using a three-dimensional printing technique. *Journal of biomedical materials research. Part B, Applied biomaterials* **107**(4), 1295–1303 (May 2019)
20. Kang, H.W., Lee, S.J., Ko, I.K., Kengla, C., Yoo, J.J., Atala, A.: A 3d bioprinting system to produce human-scale tissue constructs with structural integrity. *Nature biotechnology* **34**(3), 312–319 (Mar 2016)
21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (Dec 22, 2014), <https://arxiv.org/abs/1412.6980>
22. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations (ICLR)* (Sep 9, 2016), <https://arxiv.org/abs/1609.02907>
23. Levie, R., Monti, F., Bresson, X., Bronstein, M.M.: Cayleynets: Graph convolutional neural networks with complex rational spectral filters. *IEEE Transactions on Signal Processing* **67**(1), 97–109 (Jan 1, 2019)
24. Liu, Y., Zhang, L., Zhou, G., Li, Q., Liu, W., Yu, Z., Luo, X., Jiang, T., Zhang, W., Cao, Y.: In vitro engineering of human ear-shaped cartilage assisted with CAD/CAM technology. *Biomaterials* **31**(8), 2176–2183 (2009)
25. Nealen, A., Igarashi, T., Sorkine, O., Alexa, M.: Laplacian mesh optimization. pp. 381–389. *GRAPHITE '06, ACM* (Nov 29, 2006)
26. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in PyTorch (2017)
27. Ploumpis, S., Ververas, E., O' Sullivan, E., Moschoglou, S., Wang, H., Pears, N., Smith, W.A.P., Gecer, B., Zafeiriou, S.: Towards a complete 3d morphable model of the human head (Nov 18, 2019), <https://arxiv.org/abs/1911.08008>
28. Ranjan, A., Bolkart, T., Sanyal, S., Black, M.J.: Generating 3d faces using convolutional mesh autoencoders (Jul 26, 2018), <http://arxiv.org/abs/1807.10267>
29. Tang, S., Li, B., Yu, H.: Chebnet: Efficient and stable constructions of deep neural networks with rectified power units using chebyshev approximations (Nov 7, 2019), <https://arxiv.org/abs/1911.05467>
30. Veselkov, K., Gonzalez, G., Aljifri, S., Galea, D., Mirnezami, R., Youssef, J., Bronstein, M., Laponogov, I.: Hyperfoods: Machine intelligent mapping of cancer-beating molecules in foods. *Scientific reports* **9**(1), 9237–12 (Jul 3, 2019)
31. Yan, P., Bowyer, K.W.: Biometric recognition using 3d ear shape. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29**(8), 1297–1308 (Aug 2007)
32. Ying, R., You, J., Morris, C., Ren, X., Hamilton, W.L., Leskovec, J.: Hierarchical graph representation learning with differentiable pooling (Jun 22, 2018), <https://arxiv.org/abs/1806.08804>

33. Zhou, G., Jiang, H., Yin, Z., Liu, Y., Zhang, Q., Zhang, C., Pan, B., Zhou, J., Zhou, X., Sun, H., Li, D., He, A., Zhang, Z., Zhang, W., Liu, W., Cao, Y.: In vitro regeneration of patient-specific ear-shaped cartilage and its first clinical application for auricular reconstruction. *EBioMedicine* **28**(C), 287–302 (Feb 2018)
34. Zhou, Y., Zaferiou, S.: Deformable models of ears in-the-wild for alignment and recognition. pp. 626–633. *IEEE* (May, 2017)