

LaB-GATr: geometric algebra transformers for large biomedical surface and volume meshes

Julian Suk^{*[0000–0003–0729–047X]}, Baris Imre^{*[0009–0005–3724–699X]}, and
Jelmer M. Wolterink^[0000–0001–5505–475X]

Department of Applied Mathematics, Technical Medical Centre, University of
Twente, Enschede, The Netherlands
{j.m.suk, j.m.wolterink}@utwente.nl

Abstract. Many anatomical structures can be described by surface or volume meshes. Machine learning is a promising tool to extract information from these 3D models. However, high-fidelity meshes often contain hundreds of thousands of vertices, which creates unique challenges in building deep neural network architectures. Furthermore, patient-specific meshes may not be canonically aligned which limits the generalisation of machine learning algorithms. We propose LaB-GATr, a transformer neural network with geometric tokenisation that can effectively learn with large-scale (bio-)medical surface and volume meshes through sequence compression and interpolation. Our method extends the recently proposed geometric algebra transformer (GATr) and thus respects all Euclidean symmetries, i.e. rotation, translation and reflection, effectively mitigating the problem of canonical alignment between patients. LaB-GATr achieves state-of-the-art results on three tasks in cardiovascular hemodynamics modelling and neurodevelopmental phenotype prediction, featuring meshes of up to 200,000 vertices. Our results demonstrate that LaB-GATr is a powerful architecture for learning with high-fidelity meshes which has the potential to enable interesting downstream applications. Our implementation is publicly available. [†]

Keywords: Deep learning · Attention models · Cardiovascular hemodynamics · Neuroimaging · Geometric algebra.

1 Introduction

Deep neural networks can leverage biomedical data to uncover previously unknown cause-and-effect relations [28] and enable novel ways of medical diagnosis and treatment [15,21]. There has been active research into using deep neural networks for biomedical modelling, such as cardiovascular biomechanics estimation [2] and neuroimage analysis based on the cortical surface [29]. These applications feature 3D mesh representations of patient anatomy. Depending on the downstream application, 3D meshes either discretise the surface of an organ

^{*}Equal contribution.

[†]github.com/sukjulian/lab-gatr

or vessel with, e.g., triangles, or the interior with, e.g., tetrahedra. Graph neural networks (GNN) have gained traction in the context of learning with 3D meshes due to their direct applicability, flexibility regarding mesh size, and good performance on several benchmarks [9,23,25]. However, GNNs are known to suffer from over-squashing, i.e. the loss of accuracy when compressing exponentially growing information into fixed-sized channels when propagating messages on long paths [1]. This makes them inefficient at accumulating large enough receptive fields, capable of learning global interactions across meshes. PointNet++ [17] can circumvent this issue to some extent via pooling to coarser surrogate graphs, but this requires problem-specific hyperparameter setup and might still fail around bottlenecks. In contrast, the transformer architecture [27], treats its input as a sequence of tokens (or patches) and models all pair-wise interactions, thus aggregating global context after a single layer. In medical imaging, transformers have been successfully applied in the context of, e.g., semantic segmentation of 2D histopathology images [12] and 3D brain tumor magnetic resonance images (MRI) [10]. Nevertheless, applications to 3D biomedical meshes remain scarce due to the difficulty of finding a unified framework that addresses their sheer size, commonly in the hundreds of thousands of vertices. Dahan et al. [6,5] addressed this by aligning 3D cortical surface meshes across subjects via morphing to an icosphere, which allows for segmentation into coarser triangular patches.

The recently proposed geometric algebra transformer (GATr) [4] has achieved state-of-the-art accuracy on several geometric tasks, largely due to the incorporation of task-specific symmetries (SE(3)-equivariance) and modelling of interactions via geometric algebra.* Even though GATr uses memory-efficient attention [18] with linear complexity, large 3D meshes still cause it to exceed GPU memory during training. In the context of graph transformers, the common bottleneck of GPU memory has been addressed by (1) sparse, local attention mechanisms together with expander graph rewiring [7,22], as well as (2) vertex clustering together with learned graph pooling and upsampling [11,14].

In this work, we scale GATr to **large-scale biomedical (LaB-GATr)** surface and volume meshes. Since each of these meshes discretises a continuous shape, mesh connectivity should be treated as an artefact, not a feature. Thus, we opt for a vertex clustering approach to make self-attention tractable. We derive a general-purpose tokenisation algorithm and interpolation method with learned feature representations in geometric algebra space, which retains all symmetries of GATr as well as the original mesh resolution, while decreasing the number of active tokens used in self-attention. We demonstrate the efficacy of our method on three tasks involving biomedical meshes. LaB-GATr sets a new state-of-the-art in prediction of blood velocity in high-fidelity, synthetic coronary artery meshes [23] which would exceed 48 GB of VRAM when using GATr with the same number of parameters. Furthermore, Lab-GATr excels at phenotype prediction of cortical surfaces from the Developing Human Connectome Project (dHCP) [8], setting a new state-of-the-art in postmenstrual age (PMA) estimation without morphing the cortical surface mesh to an icosphere. We provide a

*We refer the interested reader to [3,20,19].




Fig. 1. LaB-GATr takes input features in the form of multivectors, which are constructed by embedding, e.g., mesh vertices as points and surface normal vectors as planes (see Section 2.2). In the tokenisation module, the features are pooled to a coarse subset of mesh vertices via message passing (see Section 2.4). The tokenisation allows control over the number of tokens which are processed by the GATr module. Downstream, the interpolation module lifts the tokenisation back to original mesh resolution (see Section 2.4). An optional (*) class token is appended to the token sequence for mesh-level output. Subsequently, scalar or vector-valued output features are extracted.

modular implementation of Lab-GATr and an interface with which the geometric algebra back-end can be treated as black box.

2 Methods

Fig. 1 provides an overview of our method. We extend GATr by geometric tokenisation for scaling to high-fidelity meshes. In the following, we discuss background on geometric algebra and transformers before introducing our contribution.

2.1 Geometric algebra

GATr is built on the geometric algebra $\mathbf{G}(3, 0, 1)$ which uses projective geometry: a fourth coordinate is appended to 3D space, so translation can be expressed as linear maps. At the core of geometric algebra lies the introduction of an associative (but not commutative) geometric product of vectors y, z , simply denoted as yz . Given a 4D orthogonal basis $\{e_i\}_i$, it holds that $e_0e_0 = 0$, $e_ie_i = 1$ ($i \neq 0$), and $e_ie_j = -e_je_i$ ($i \neq j$). All possible, linearly independent geometric products span a 16-dimensional vector space of multivectors $x \in \mathbf{G}(3, 0, 1)$:

$$x = (x_s, \underbrace{x_0, x_1, x_2, x_3}_{\text{vectors}}, \underbrace{x_{01}, x_{02}, x_{03}, x_{12}, x_{13}, x_{23}}_{\text{bivectors}}, \underbrace{x_{012}, x_{013}, x_{023}, x_{123}}_{\text{trivectors}}, x_{0123}) \quad (1)$$

which can represent geometric quantities like points and planes.

Table 1. Embedding of some common geometric objects as $x \in \mathbf{G}(3, 0, 1)$. See (1) for the 16 multivector components. Other multivector components remain zero. In geometric algebra, geometric operations can be multivectors just like geometric objects.

Geometric object / operation	Multivector mapping
Scalar $s \in \mathbb{R}$	$x_s = s$
Plane with normal $\nu \in \mathbb{R}^3$ and offset $\delta \in \mathbb{R}$	$(x_0, x_1, x_2, x_3) = (\delta, \nu)$
Point $\rho \in \mathbb{R}^3$	$(x_{012}, x_{013}, x_{023}, x_{123}) = (\rho, 1)$
Translation $\tau \in \mathbb{R}^3$	$(x_s, x_{01}, x_{02}, x_{03}) = (1, \frac{1}{2}\tau)$

2.2 Embedding meshes in $\mathbf{G}(3, 0, 1)$

Consider an arbitrary surface or volume mesh consisting of n vertices. For each vertex, we construct a d -dimensional positional encoding which describes its unique geometric properties within this mesh. Each positional encoding is composed of a set of c geometric objects, e.g. the surface normal vector (for surface meshes) or the (scalar) distance to the surface (for volume meshes). Table 1 provides a look-up on how to embed relevant geometric objects, such as points and planes, as multivectors $x \in \mathbf{G}(3, 0, 1)$ (see Fig. 1a). Consequently, we describe each mesh by a tensor $X^{(0)} \in \mathbb{R}^{n \times d}$ with $d = c \cdot 16$.

2.3 Geometric algebra transformers

Given a tensor of input features $X^{(l)}$, we define transformer blocks as follows:

$$A^{(l)} = X^{(l)} + \xi \left(\text{Concat}_h \text{ Softmax} \left(\frac{q_h(X^{(l)})k_h(X^{(l)})^\top}{\sqrt{d}} \right) v_h(X^{(l)}) \right)$$

$$X^{(l+1)} = A^{(l)} + \text{MLP}(A^{(l)}).$$

As is common practice, $q_h, k_h, v_h: \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times d}$ consist of layer normalisation composed with learned linear maps. Multi-head self-attention [27] over heads indexed by h is implemented via concatenation followed by a learned linear map ξ . GATr [4] introduced layer normalisation, linear and nonlinear maps $\mathbf{G}(3, 0, 1)^{n \times c} \rightarrow \mathbf{G}(3, 0, 1)^{n \times c}$, and gated activation functions which can be used to construct ξ, q_h, k_h, v_h , and MLP in geometric algebra. GATr blocks are equivariant under rotations, translations and reflections $\rho \in \text{E}(3)$ of the input geometry encoded in $X^{(l)}$, i.e. they map $\rho X^{(l)} \mapsto \rho X^{(l+1)}$.

2.4 Learned, geometric tokenisation

In the following, we introduce tokenisation layers for large surface and volume meshes, allowing us to compress the token sequence for the transformer blocks.

Pooling From the point cloud P_{fine} consisting of the n mesh vertices, we compute a subset of points P_{coarse} via farthest point sampling. We partition the mesh vertices into tokens by assigning each $p \in P_{\text{coarse}}$ the (disjoint) cluster of closest points in P_{fine} (see Fig. 1b):

$$C(p) = \{v \in P_{\text{fine}} \mid p = \arg \min_{q \in P_{\text{coarse}}} \|v - q\|_2\}.$$

This means that each point in P_{fine} is clustered with the point in P_{coarse} to which it is the closest. Define $n_{\text{coarse}} = |P_{\text{coarse}}|$. Given a tensor of input features $X \in \mathbb{R}^{n \times d}$, we perform learned message passing within these clusters (see Fig. 1c) as follows:

$$\begin{aligned} m_{v \rightarrow p} &= \text{MLP}(X^{(0)}|_v, p - v) && (\text{message from } P_{\text{fine}} \text{ to } P_{\text{coarse}}) \\ X^{(1)}|_p &= \frac{1}{|C(p)|} \sum_{v \in C(p)} m_{v \rightarrow p} && (\text{aggregation of } X^{(1)} \in \mathbb{R}^{n_{\text{coarse}} \times d}) \end{aligned}$$

where $X|_q$ denotes the row of X corresponding to point q . This layer maps $\mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n_{\text{coarse}} \times d}$ with $n_{\text{coarse}} < n$. Within this framework, we can embed $p - v$ as translation (see Table 1) and use the multilayer perceptron (MLP) introduced by [4] to reduce the number of tokens in a way that is fully compatible with GATr. In particular, this layer respects all symmetries of $\mathbf{G}(3, 0, 1)$.

Interpolation Given a tensor $X^{(l)} \in \mathbb{R}^{n_{\text{coarse}} \times d}$ we define learned interpolation to the original mesh resolution $Y \in \mathbb{R}^{n \times d}$ as follows:

$$\begin{aligned} X^{(l+1)}|_v &= \frac{\sum_p \lambda_{p,v} X^{(l)}|_p}{\sum_p \lambda_{p,v}}, \quad \lambda_{p,v} := \frac{1}{\|p - v\|_2^2 + \epsilon}, \\ Y &= \text{MLP}(X^{(l+1)}, X^{(0)}) \end{aligned}$$

where for each $v \in P_{\text{fine}}$ we sum over the three (four) closest points $p \in P_{\text{coarse}}$ for surface (volume) meshes (see Fig. 1d) and ϵ is a small constant. This layer lifts the tokenisation and ensures neural-network output in original mesh resolution. The interpolation provably behaves as expected in $\mathbf{G}(3, 0, 1)$ (see appendix) and by using the MLP introduced by [4] this layer respects all symmetries of $\mathbf{G}(3, 0, 1)$.

2.5 Neural network architecture

LaB-GATr is composed of geometric algebra embedding, tokenisation module, GATr module, and interpolation module followed by feature extraction (see Fig. 1). We embed the input mesh in $\mathbf{G}(3, 0, 1)$ (see Section 2.2) based on an application-specific set of geometric descriptors. The embedding $X^{(0)} \in \mathbb{R}^{n \times d}$ is pooled (see Section 2.4) and the resulting n_{coarse} tokens are fed into a GATr module. For mesh-level tasks, we append a global class token which we embed




Fig. 2. Qualitative results of cardiovascular hemodynamics estimation on test-split arteries. Left: (top) surface and (bottom) volume mesh. Middle: (top) wall shear stress and (bottom) velocity field, via computational fluid dynamics (CFD). Right: LaB-GATr prediction.

as the mean of the n_{coarse} tokens. For vertex-level tasks, the m output tokens of the transformer are interpolated (see Section 2.4) back to the original mesh resolution. For classification tasks (like segmentation), Softmax can be used over the channel dimension at the mesh- or vertex-level output.

3 Experiments and results

We evaluate LaB-GATr on three tasks previously explored with GNNs and transformers. Since all of these are regression tasks, we train LaB-GATr under L^1 loss using the Adam optimiser [13] and learning rate $3e-4$ with exponential decay on NVIDIA L40 (48 GB) GPUs. We use the same number of channels and attention heads throughout all experiments, leading to around 320k trainable parameters.

3.1 Coronary artery models

Surface-based WSS estimation The publicly available dataset [25] consists of 2,000 synthetic, coronary artery surface meshes (around 7,000 vertices) with simulated, steady-state wall shear stress (WSS) vectors. We chose $n_{\text{coarse}} = 0.1 \cdot n$ in the tokenisation (see Section 2.4) and embed vertex positions as points (see Table 1), surface normal vectors as oriented planes, and geodesic distances to the artery inlet as scalars. The GATr back-end (see Section 2.5) was set up identical to [4]. We trained LaB-GATr on a single GPU for 4,000 epochs ($58 \frac{\text{s}}{\text{epoch}}$) with batch size 8 on an 1600:200:200 split of the dataset. Fig. 2 shows an example of

Table 2. Comparison of LaB-GATr with state-of-the-art baselines. For WSS and hemodynamics estimation, we report mean ε (lower is better) across the test set. For PMA estimation, some of the referenced works report lowest and some average MAE across three training runs. We report lowest MAE in brackets where applicable.

Domain	Dataset	Model	Disparity (average) (lowest)		Metric
Cardiovascular	WSS [25]	GEM-CNN [24]	7.8		ε [%]
		GATr [4]	5.5		
		<i>LaB-GATr</i>	5.5*		
		<i>LaB-GATr</i> ¹	7.0		
		<i>LaB-GATr</i> ²	12.3		
Neuroimaging	PMA [8] (native)	SEGNN [23]	7.4		MAE [weeks]
		<i>LaB-GATr</i>	3.3		
		SiT [6]	–	(0.68)	
		MS-SiT [5]	0.59	–	
		[26]	–	(0.54)	
		<i>LaB-GATr</i>	0.54	(0.52)	

*with 10-fold compression ¹static pooling module ²static interpolation module

LaB-GATr prediction on a test-split mesh. Table 2 shows approximation error ε [24,4] of LaB-GATr compared to the baselines. LaB-GATr matches GATr’s accuracy $\varepsilon = 5.5\%$ (standard deviation was $\pm 2.0\%$ over 200 test cases) despite its 10-fold compression of the token sequence. Ablation of the learnable parameters of the pooling and interpolation module reveals that learned interpolation plays a bigger role in performance than learned pooling.

Volume-based velocity field estimation The publicly available dataset [23] consists of 2,000 synthetic, bifurcating coronary artery volume meshes (around 175,000 vertices) with simulated, steady-state velocity vectors. We chose $n_{\text{coarse}} = 0.01 \cdot n$ for the tokenisation and embed vertex positions as points, directions to the artery inlet, outlets, and wall as oriented planes, and distances to the artery inlet, outlets, and wall as scalars. The size of these volume meshes caused GATr to exceed memory of an NVIDIA L40 (48 GB) GPU. Leveraging our tokenisation enabled training LaB-GATr with batch size 1. We trained LaB-GATr on four GPUs in parallel for 300 epochs ($10:24 \frac{\text{min}}{\text{epoch}}$) on an 1600:200:200 split of the dataset. Fig. 2 shows an example of LaB-GATr prediction on a test-split mesh. Table 2 compares LaB-GATr to [23]. LaB-GATr sets a new state-of-the-art for this dataset with $\varepsilon = 3.3\%$ (standard deviation was $\pm 4.4\%$ over 200 test cases) compared to the previous 7.4%.

3.2 Postmenstrual age prediction from the cortical surface

The publicly available third release of dHCP [8] consists of 530 newborns’ cortical surface meshes (81,924 vertices each) which are symmetric across hemispheres.




Fig. 3. Postmenstrual age (PMA) prediction (values in weeks) on test-split subjects. LaB-GATr predictions based on the newborn’s cortical surface correlated well with the reference values.

We estimate the subjects’ postmenstrual age (PMA) at the time of scan. We chose $n_{\text{coarse}} = 0.024 \cdot n$ for the tokenisation and embed vertex positions as points, surface normal vectors as oriented planes, the reflection planes between symmetric vertices as oriented planes, and myelination, curvature, cortical thickness, and sulcal depth as scalars. Since we observed quick convergence on the validation split, we trained LaB-GATr on a single GPU for only 200 epochs ($1:38 \frac{\text{min}}{\text{epoch}}$) with batch size 4 on the 423:53:54 splits used in [5]. Figure 3 shows a Bland-Altman plot and the cortical surface of two subjects, indicating good accuracy. Table 2 shows mean absolute error (MAE) and comparison to the baselines. In contrast to all baselines, LaB-GATr runs directly on the cortical surface mesh without morphing to a sphere. LaB-GATr sets a new state-of-the-art in PMA estimation on "native space" [9] with average MAE = 0.54 weeks (standard deviation was ± 0.39 weeks over 54 test cases times three runs).

4 Discussion and conclusion

In this work, we propose LaB-GATr, a general-purpose geometric algebra transformer for large-scale surface and volume meshes, often found in biomedical engineering. We extend GATr by learned tokenisation and interpolation in geometric algebra. Our method can be understood as a thin PointNet++ [17] wrapper which is adapted to projective geometric algebra. Through the self-attention mechanism, LaB-GATr models global interactions within the mesh, while avoiding the over-squashing phenomenon exhibited by GNNs. Notably, LaB-GATr is equivariant to rotations, translations, and reflections of the input mesh, thus circumventing the problem of cortical surface alignment. In our experiments, LaB-GATr matched the performance of GATr [4] on the same dataset, suggesting near loss-less compression. Even though LaB-GATr introduces additional, trainable parameters over GATr, we found that computing self-attention between less tokens outweighed the parameter overhead and led to favourable training times. Beside estimating PMA of subjects from dHCP, we also attempted gestational age (GA) estimation with LaB-GATr. However, we observed considerably lower

accuracy. We believe that GA can be best explained by vertex-specific biomarkers such as myelination, which was provided on sphericalised and subsequently sub-sampled brains and back-projection to the cortical surface erased some of their spatial context.

In theory, our geometric pooling scheme introduces all necessary building blocks to define patch merging and to build a geometric version of sliding window (Swin) attention [16], which is an interesting direction for future work. We believe that Lab-GATr has potential as general-purpose model for learning with large (bio-)medical surface and volume meshes, enabling interesting downstream applications. In particular, we are interested to explore brain parcellation and attention-map-based analysis of biomedical pathology. Geometric algebra introduces an inductive bias to our learning system. In future work, we aim to investigate to what extent this affects LaB-GATr predictions and derive theoretical guarantees.

Acknowledgments. This work is funded in part by the 4TU Precision Medicine programme supported by High Tech for a Sustainable Future, a framework commissioned by the four Universities of Technology of the Netherlands. Jelmer M. Wolterink was supported by the NWO domain Applied and Engineering Sciences Veni grant (18192).

We would like to thank Simon Dahan for his efforts in providing the dHCP dataset as well as Pim de Haan and Johann Brehmer for the fruitful discussions about geometric algebra.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Alon, U., Yahav, E.: On the bottleneck of graph neural networks and its practical implications. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021 (2021)
2. Arzani, A., Wang, J.X., Sacks, M., Shadden, S.: Machine learning for cardiovascular biomechanics modeling: Challenges and beyond. *Annals of Biomedical Engineering* **50**, 1–13 (04 2022)
3. Brandstetter, J., van den Berg, R., Welling, M., Gupta, J.K.: Clifford neural layers for PDE modeling. In: The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023 (2023)
4. Brehmer, J., de Haan, P., Behrends, S., Cohen, T.: Geometric algebra transformer. In: *Advances in Neural Information Processing Systems*. vol. 37 (2023)
5. Dahan, S., Fawaz, A., Suliman, M.A., da Silva, M., Williams, L.Z.J., Rueckert, D., Robinson, E.C.: The multiscale surface vision transformer. *ArXiv* (2023)
6. Dahan, S., Fawaz, A., Williams, L.Z.J., Yang, C., Coalson, T.S., Glasser, M.F., Edwards, A.D., Rueckert, D., Robinson, E.C.: Surface vision transformers: Attention-based modelling applied to cortical analysis. In: *International Conference on Medical Imaging with Deep Learning, MIDL*, 6-8 July 2022, Zurich, Switzerland (2022)
7. Deac, A., Lackenby, M., Velickovic, P.: Expander graph propagation. In: Rieck, B., Pascanu, R. (eds.) *Learning on Graphs Conference, LoG*, 9-12 December 2022, Virtual Event (2022)

8. Edwards, A.D., al.: The developing human connectome project neonatal data release. *Frontiers in Neuroscience* **16** (2022)
9. Fawaz, A., Williams, L.Z.J., Alansary, A., Bass, C., Gopinath, K., da Silva, M., Dahan, S., Adamson, C., Alexander, B., Thompson, D., Ball, G., Desrosiers, C., Lombaert, H., Rueckert, D., Edwards, A.D., Robinson, E.C.: Benchmarking geometric deep learning for cortical segmentation and neurodevelopmental phenotype prediction. *bioRxiv* (2021)
10. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D.: Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In: *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*. pp. 272–284. Springer International Publishing, Cham (2022)
11. Janny, S., Bénateau, A., Nadri, M., Digne, J., Thome, N., Wolf, C.: EAGLE: large-scale learning of turbulent fluid dynamics with mesh transformers. In: *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023* (2023)
12. Ji, Y., Zhang, R., Wang, H., Li, Z., Wu, L., Zhang, S., Luo, P.: Multi-compound transformer for accurate biomedical image segmentation. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*. pp. 326–336 (2021)
13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (2015)
14. Kong, K., Chen, J., Kirchenbauer, J., Ni, R., Bruss, C.B., Goldstein, T.: GOAT: A global transformer on large-scale graphs. In: *International Conference on Machine Learning, ICML, 23-29 July 2023, Honolulu, Hawaii, USA* (2023)
15. Li, G., Wang, H., Zhang, M., Tupin, S., Qiao, A., Liu, Y., Ohta, M., Anzai, H.: Prediction of 3d cardiovascular hemodynamics before and after coronary artery bypass surgery via deep learning. *Communications Biology* **4** (01 2021)
16. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: *2021 IEEE/CVF International Conference on Computer Vision, ICCV, Montreal, QC, Canada, October 10-17, 2021* (2021)
17. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: *Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA* (2017)
18. Rabe, M.N., Staats, C.: Self-attention does not need $O(n^2)$ memory. In: n/a (2021)
19. Ruhe, D., Brandstetter, J., Forré, P.: Clifford group equivariant neural networks. In: *Thirty-seventh Conference on Neural Information Processing Systems* (2023)
20. Ruhe, D., Gupta, J.K., Keninck, S.D., Welling, M., Brandstetter, J.: Geometric clifford algebra networks. In: *International Conference on Machine Learning, ICML, 23-29 July 2023, Honolulu, Hawaii, USA* (2023)
21. Sarasua, I., Pölsterl, S., Wachinger, C.: Transformesh: A transformer network for longitudinal modeling of anatomical meshes. In: *Machine Learning in Medical Imaging*. pp. 209–218. Springer International Publishing, Cham (2021)
22. Shirzad, H., Velingker, A., Venkatachalam, B., Sutherland, D.J., Sinop, A.K.: Exphormer: Sparse transformers for graphs. In: *International Conference on Machine Learning, ICML, 23-29 July 2023, Honolulu, Hawaii, USA* (2023)
23. Suk, J., Brune, C., Wolterink, J.M.: Se(3) symmetry lets graph neural networks learn arterial velocity estimation from small datasets. In: Bernard, O., Clarysse, P., Duchateau, N., Ohayon, J., Viallon, M. (eds.) *Functional Imaging and Modeling of the Heart*. pp. 445–454. Springer Nature Switzerland, Cham (2023)

24. Suk, J., de Haan, P., Lippe, P., Brune, C., Wolterink, J.M.: Mesh neural networks for $se(3)$ -equivariant hemodynamics estimation on the artery wall. ArXiv (2022)
25. Suk, J., Haan, P.d., Lippe, P., Brune, C., Wolterink, J.M.: Mesh convolutional neural networks for wall shear stress estimation in 3d artery models. In: Statistical Atlases and Computational Models of the Heart. Multi-Disease, Multi-View, and Multi-Center Right Ventricular Segmentation in Cardiac MRI Challenge. pp. 93–102. Springer International Publishing, Cham (2022)
26. Unyi, D., Gyires-Tóth, B.: Neurodevelopmental phenotype prediction: A state-of-the-art deep learning model. In: Machine Learning for Health, ML4H, 28 November 2022, New Orleans, Louisiana, USA & Virtual. Proceedings of Machine Learning Research, vol. 193, pp. 279–289. PMLR (2022)
27. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Annual Conference on Neural Information Processing Systems, December 4-9, 2017, Long Beach, CA, USA (2017)
28. Vosylus, V., Wang, A., Waters, C., Zakharov, A., Ward, F., Le Folgoc, L., Cupitt, J., Makropoulos, A., Schuh, A., Rueckert, D., Alansary, A.: Geometric deep learning for post-menstrual age prediction based on the neonatal white matter cortical surface. In: Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Graphs in Biomedical Image Analysis. pp. 174–186. Springer International Publishing, Cham (2020)
29. Zhao, F., Wu, Z., Li, G.: Deep learning in cortical surface-based neuroimage analysis: a systematic review. *Intelligent Medicine* **3**(1), 46–58 (2023)