Implicit Neural Representations for Modeling of Abdominal Aortic Aneurysm Progression

Dieuwertje Alblas[1(✉)] [iD], Marieke Hofman[1], Christoph Brune[1] [iD],
Kak Khee Yeung[2,3] [iD], and Jelmer M. Wolterink[1] [iD]

[1] Department of Applied Mathematics, Technical Medical Centre,
University of Twente, Enschede, The Netherlands
d.alblas@utwente.nl

[2] Department of Surgery, Amsterdam UMC location Vrije Universiteit Amsterdam,
Amsterdam, The Netherlands

[3] Amsterdam Cardiovascular Sciences, Microcirculation,
Amsterdam, The Netherlands

Abstract. Abdominal aortic aneurysms (AAAs) are progressive dilatations of the abdominal aorta that, if left untreated, can rupture with lethal consequences. Imaging-based patient monitoring is required to select patients eligible for surgical repair. In this work, we present a model based on implicit neural representations (INRs) to model AAA progression. We represent the AAA wall over time as the zero-level set of a signed distance function (SDF), estimated by a multilayer perception that operates on space and time. We optimize this INR using automatically extracted segmentation masks in longitudinal CT data. This network is conditioned on spatiotemporal coordinates and represents the AAA surface at any desired resolution at any moment in time. Using regularization on spatial and temporal gradients of the SDF, we ensure proper interpolation of the AAA shape. We demonstrate the network's ability to produce AAA interpolations with average surface distances ranging between 0.72 and 2.52 mm from images acquired at highly irregular intervals. The results indicate that our model can accurately interpolate AAA shapes over time, with potential clinical value for a more personalised assessment of AAA progression.

Keywords: Abdominal aortic aneurysm · Implicit neural representation · Deep learning · Aneurysm progression

1 Introduction

Abdominal aortic aneurysms (AAAs) are progressive local dilatations of the abdominal aorta of at least 30 mm that most frequently occur below the renal

arteries. AAAs are mostly asymptomatic, but rupture of an AAA has a mortality rate of 70–80% [3]. To avert rupture, patients can undergo elective repair via either open surgery or an endovascular procedure. Patients become eligible for surgical repair if the diameter of the AAA exceeds a threshold (5.5 cm in men, 5.0 cm in women) or if the AAA diameter has increased more than 1 cm in a year [16].

Prior to elective repair, patients are monitored via periodic outpatient clinic visits and imaging with ultrasound or CT. Although these longitudinal images are primarily used to measure the diameter of the aneurysm, they contain a wealth of information that may be leveraged to better model AAA progression in individual patients [6]. Detailed insight into personalised AAA progression has the potential to aid the physician in clinical decision-making by filling in the gaps in surveillance data. Previous efforts to model the progression of AAAs based on longitudinal imaging include models based on Gaussian processes that represent an underlying deformation field [4], Markov chains [21], deep belief networks [9], or CNNs operating on the surface of the AAA [10].

Recently, implicit neural representations (INRs) have gained traction as natural representations for signals on a spatial or spatiotemporal domain [20]. INRs are multilayer perceptrons that take continuous coordinates as input and output the value of the signal or function at that point [14]. INRs are attractive representation models as derivatives of the signal can be analytically computed using automatic differentiation. In medical imaging, INRs have been used for, e.g., sparse-view CT reconstruction [13,15] and image registration [19]. Moreover, INRs can be used to accurately represent shapes [12], which has led to applications in cell shape synthesis [18] statistical shape modeling [2,11] or surface fitting based on point cloud annotations [1].

In this work, we propose to use INRs with a time coordinate to represent a longitudinal 3D AAA model of a patient and investigate to what extent such a model can be used to *interpolate* and *extrapolate* the AAA surface in time.

2   Methods

We represent the evolving AAA surface as the zero level set of its temporal signed distance function (SDF). We parametrize this function by a neural network $f(\boldsymbol{x}, t; \theta)$, with weights θ.

2.1   Signed Distance Function

A surface can be implicitly represented by the zero level set of its signed distance function. We consider a manifold evolving over time, that we represent by a temporal SDF: $SDF(\boldsymbol{x}, t) : \mathbb{R}^3 \times \mathbb{R} \mapsto \mathbb{R}$. The value of the $SDF(\boldsymbol{x}, t)$ represents the minimum distance to the surface at location $\boldsymbol{x}$ at time t. The temporal SDF of an evolving 2D manifold $\mathcal{M}$ embedded in $\mathbb{R}^3 \times \mathbb{R}$ is defined as:

$$SDF_{\mathcal{M}}(\boldsymbol{x}, t) = \begin{cases} -d(\boldsymbol{x}, \mathcal{M}) & \boldsymbol{x} \text{ inside } \mathcal{M} \text{ at time } t \\ 0 & \boldsymbol{x} \text{ on } \mathcal{M} \text{ at time } t \\ d(\boldsymbol{x}, \mathcal{M}) & \boldsymbol{x} \text{ outside } \mathcal{M} \text{ at time } t. \end{cases} \tag{1}$$

The zero-level set of this function thus describes the manifold $\mathcal{M}$. Moreover, the signed distance function is a solution to the Eikonal equation at each instance in time: $||\nabla_x SDF(\boldsymbol{x}, t)|| = 1, \forall \boldsymbol{x}, t$.

2.2   Implicit Neural Representations
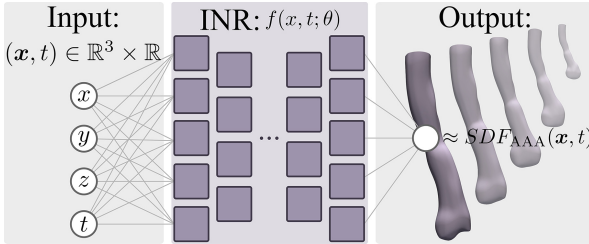
Fig. 1. Schematic representation of our INR, taking spatiotemporal coordinates $(\boldsymbol{x}, t)$ as an input, outputting $SDF(\boldsymbol{x}, t)$ of the AAA surface. Note that a single INR represents the complete evolving AAA of a patient.

In previous work, it has been shown that an SDF of a manifold can be represented by a neural network [1,7,12,18]. Similarly, we embed the remodeling of the AAA over time in an implicit neural representation (INR). We use 4D coordinates from the spatiotemporal domain $\Omega := [-1, 1]^3 \times [-1, 1]$ as input to the network $f(\boldsymbol{x}, t; \theta)$. The output node of our INR approximates the SDF value at the input coordinate. Figure 1 shows a schematic overview of our INR.

We aim to reconstruct $SDF_{\text{AAA}}(\boldsymbol{x}, t)$ given a sequence of point clouds of the AAA surface, representing the aneurysm shape of a single patient over J scans: $\{\mathcal{X}_j\}_{j=1,\ldots,J}$, where $\mathcal{X}_{j,.} \subset [-1, 1]^3$. We denote individual points on the j^{th} AAA surface $\boldsymbol{x}_i^j$. To optimise the INR, we sample points on and off the AAA surface at multiple instances in time.

The loss function we use to optimize the INR consists of two terms: a term $\mathcal{L}_{\text{data}_j}$ at each time point t_j where we have ground-truth scan data, and a term $\mathcal{L}_{\text{reg}}$ that regularises the SDF at times the surface is unknown.

$$\mathcal{L}(\theta) = \sum_{1 \leq j \leq J} \mathcal{L}_{\text{data}_j}(\theta) + \mathcal{L}_{\text{reg}}(\theta), \tag{2}$$

$$\mathcal{L}_{\text{data}_j}(\theta) = \frac{1}{N_j} \sum_{1 \leq i \leq N_j} |f(\boldsymbol{x}_i^j, t_j; \theta)| + \lambda_1 \mathbb{E}(||\nabla_x f(\boldsymbol{x}, t_j; \theta)|| - 1)^2$$

$$+ \lambda_2 \mathbb{E}(|\nabla_t f(\boldsymbol{x}, t_j; \theta)|) \tag{3}$$

$$\mathcal{L}_{\text{reg}} = \lambda_3 \mathbb{E}\left(||\nabla_x f(\boldsymbol{x}, \tilde{t}; \theta)|| - 1\right)^2 + \lambda_4 \mathbb{E}\left(|\nabla_t f(\boldsymbol{x}, t; \theta)|\right). \tag{4}$$

The first term of $\mathcal{L}_{\text{data}_j}$ was introduced in [7]. It ensures $SDF(\boldsymbol{x}_i^j, t_j) = 0$ for all points $\boldsymbol{x}_i^j$ in pointcloud $\mathcal{X}_j$, i.e. that points that are known to be on the AAA
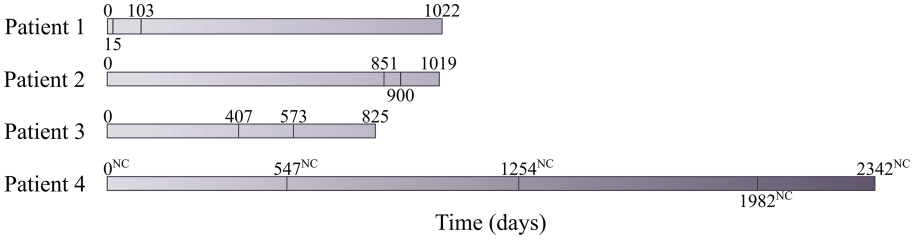
Fig. 2. Timeline of the CT scans of the four patients with longitudinal data, showing scan instances in days. Non-contrast scans are indicated with ^{NC}.

surface are indeed on the zero level set of the SDF. The remaining terms in both parts of the loss function regularise the INR's spatial and temporal gradient. As these terms do not depend on pointcloud data, we evaluate them both at times t_j as well as times data is unavailable. Regularising the norm of the spatial gradient was also introduced in [7] and enforces the INR to be a solution to the Eikonal equation. We evaluate this term at time t_j in $\mathcal{L}_{\text{data}_j}$, and at an arbitrary time point $\tilde{t}$ in $\mathcal{L}_{\text{reg}}$. The temporal regularisation term is introduced in this work to restrict temporal changes of the INR. These are evaluated at time t_j and at multiple arbitrary time points in $\mathcal{L}_{\text{data}_j}$ and $\mathcal{L}_{\text{reg}}$ respectively.

2.3   Data

We retrospectively included longitudinal CT data of four patients scanned at Amsterdam AMC (Amsterdam, The Netherlands) between 2011 and 2020. Three patients were scanned four times, and one patient was scanned five times (Fig. 2). Scan dates were shifted so that the first scan date of each patient became day 0. Patient 1 was scanned three times between day 0 and day 103, followed by a gap of almost three years. The first follow-up image of Patient 2 was after 851 days, after which two additional follow-up images were acquired relatively soon. Patient 3 was scanned more regularly. The follow-up for Patient 4 is the longest, with over 75 months of follow-up. CT scans were a mixture of non-contrast and contrast-enhanced images.

   We obtained automatic segmentations of the AAA and vertebra in each of these patients. All CT scans were processed using TotalSegmentator [17], a Python library based on nn-UNet [8] that segments >100 structures in 3D CT images. This library segmented the vertebra with good accuracy in both non-contrast and contrast-enhanced images and the AAA with high accuracy in all non-contrast images. However, segmentation of the AAA in contrast-enhanced images was unsatisfactory. Instead, we used an in-house dataset of 80 contrast-enhanced CT images of AAA patients with annotations of the AAA ranging between the top of the T12 vertebra and the iliac bifurcation to train an additional nn-UNet model. This model achieved a mean Dice similarity coefficient of 0.90 on a separate test set consisting of 13 contrast-enhanced CT scans.
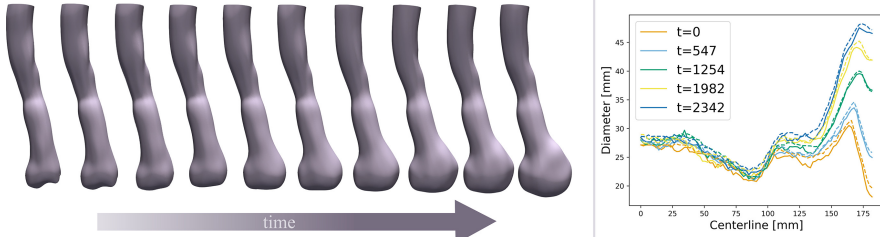
Fig. 3. An optimised INR can be used to extract shape interpolations at an arbitrary number of time points, here we show results for Patient 4. *Left:* We show extracted shapes at ten regularly spaced intervals in time. *Right:* Diameter plots along the centerlines of the aorta, comparing the ground-truth segmentation mask (solid) to the surface fitted by the network at five time points where reference CT scans are available (dashed).

2.4   Preprocessing

In order to evaluate local changes in shape over time, all shapes should be aligned in the same coordinate system. For this, we used rigid registration in ITK on the vertebra segmentations [4]. Subsequently, the surface of each aorta was extracted from the mask and represented as a point set. This resulted in aligned pointcloud representations of the AAA surface for each scan. Finally, before serving as input to the network, the spatial coordinates of the pointclouds of each patient were jointly normalized to the $[-1, 1]^3$ domain. Similarly, the time scale of each patient was normalized to the $[-1, 1]$ interval.

3   Experiments and Results

In all cases, we used an MLP with six fully connected layers containing 256 nodes with Softplus ($\beta = 100$) activations and a final node representing the estimated SDF of the AAA surface. Like [1,7], we used a skip connection, connecting the input to the third hidden layer. The regularization coefficients were set to $\lambda_1 = \lambda_2 = \lambda_3 = \lambda_4 = 0.1$. We used an Adam optimizer with a learning rate of 0.0001 to train our network for 25,000 epochs on an NVIDIA Quadro RTX 6000 GPU. The batch sizes depended on the size of point clouds and ranged between 2877 and 6027.

3.1   Interpolation and Extrapolation

For each patient, we first optimised a single INR (Fig. 1) based on point clouds from all available scans. Because the spatiotemporal input coordinates to the INR are continuous, we can retrieve a shape at any point in time at any resolution. We visualize this in Fig. 3(*left*), where we show ten AAA shapes of Patient 4 at regularly spaced intervals. In Fig. 3(*right*) we compare the diameters along the AAA centerlines of the ground-truth segmentation masks to the
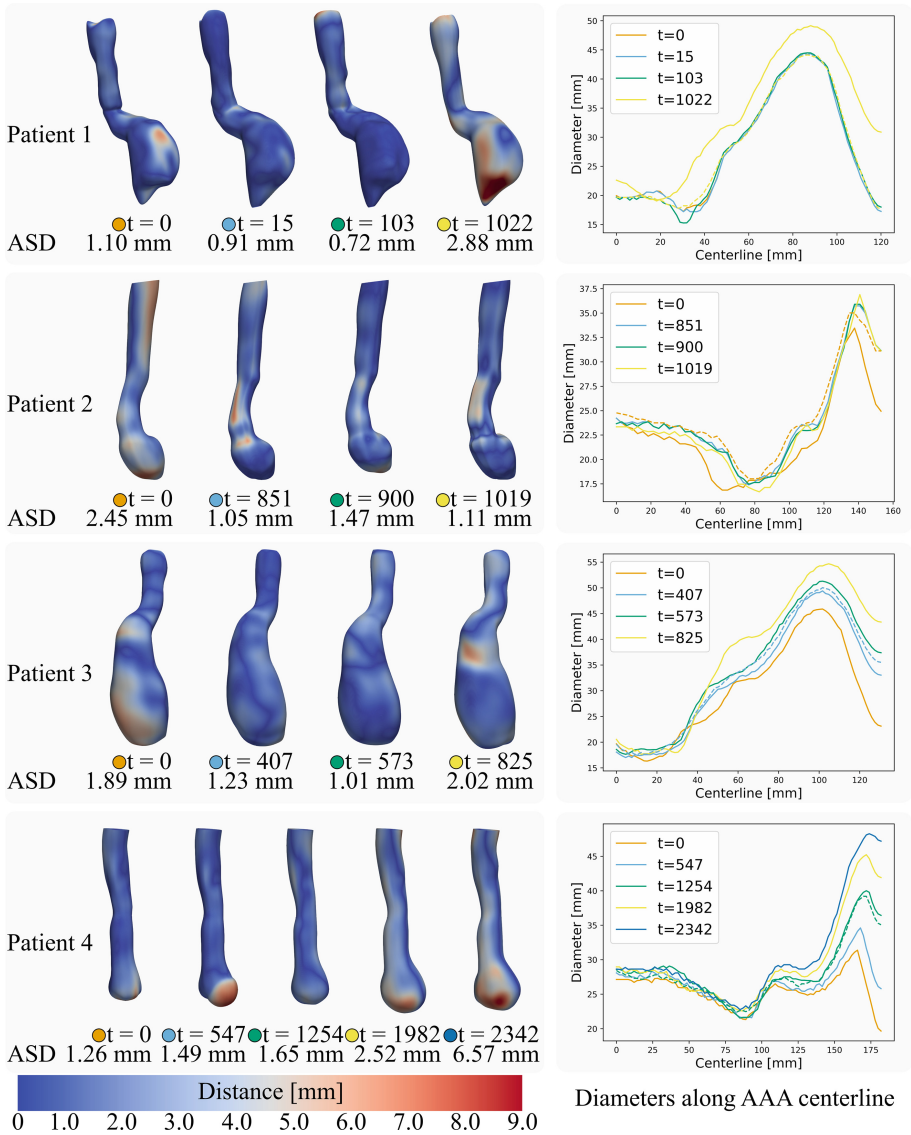
Fig. 4. *Left:* Inter- and extrapolated AAA shapes for each scan. Surface colors indicate distances to reference shapes, averages are indicated below each AAA. Colored dots indicate the corresponding diameter profile in the graph on the right. *Right:* Diameter profiles along each aorta. Solid lines represent reference diameters, dashed lines show interpolated or extrapolated diameters. A complete set of diameter plots can be found in the supplementary materials.

AAA surfaces reconstructed by the network, represented by solid and dashed lines respectively. We observe that the model accurately represents the AAA shapes at scan instances and will thus be used to evaluate the next experiment.

Next, we performed a series of leave-one-out experiments in which we optimised an INR for a patient but left out one of the time points. We used the optimised INR to estimate what the surface would have been at that time point, and compare it to the real reconstruction.

Figure 4 shows the results of these leave-one-out experiments. Results for individual patients are visualized per row. The left column in each row shows the reconstructed AAA shapes at the scan instances when that scan was left out of the training data. Colors indicate the minimal surface distance between the interpolated AAA surface and the reference AAA surface, where lower is better. The right column contains diameter plots for each aorta along its centerline estimated based on an inscribed sphere method [5,9]. Solid lines represent the diameters of the reference AAA surfaces, and the dashed line represents the diameter of the AAA from the scan that was left out. Note that we here show the diameter profile for one leave-one-out experiment per patient and that a full set of diameter profiles can be found in the supplementary materials.

Figure 4 shows that the INR model can *interpolate* AAA shapes to a decent extent. For example, in Patient 3, the interpolated surfaces at $t = 407$ and $t = 573$ had average surface distances of 1.23 and 1.01 mm, respectively, compared to the ground-truth shapes. This is also reflected in the diameter plot for Patient 3, where the interpolated (dashed) line for $t = 407$ days closely follows the reference (solid) line. The results in Fig. 4 also indicate that interpolation might work better in cases where the interval between scans is shorter. For example, interpolation results for Patient 1 at $t = 15$, which is only 15 and 88 days apart from two other scans, have an average surface distance of 0.91 mm. In contrast, interpolation results for Patient 4 at $t = 547$, which is 547 and 707 days apart from two other scans, show relatively large errors on the aneurysm sac. However, this is not consistently the case. For Patient 2, the ASD is 1.47 mm when interpolating at $t = 900$ days, which is larger than the ASD when interpolating for $t = 851$ days. From the diameter plots shown for Patients 3 and 4, we see that interpolations of the model consistently lie between the surrounding two scans and are close to the diameters of the reference shape.

Results also indicate that extrapolation is challenging for the model. The INR particularly struggles to extrapolate over bigger time gaps. For Patient 1, we observe that the extrapolations at $t = 0$ days and $t = 1022$ days have worse results than the interpolations. Moreover, the extrapolation at $t = 1022$ days differs more from the reference shape than at $t = 0$ days due to the difference in time gaps. The diameter profiles for Patient 1 and Patient 2 reveal that the model tends to reconstruct the surface of the last known shape. We hypothesize that this might be due to the temporal regularization term in Eq. 4.

Finally, Fig. 4 indicates that our INR model reacts strongly to small misalignments of the original AAA shapes. Following [4], we register AAA shapes based on segmentation masks of the vertebrae, but this alignment might lead

to small local shifts of the AAA. For example, the result for Patient 2, $t = 0$ in Fig. 4 shows errors on the healthy part of the aorta, an area that, in principle, should not show growth over time.

4   Discussion and Conclusion

In this work, we have obtained a personalised model for AAA progression, based on longitudinal CT data. We combine fully automatic state-of-the-art image segmentation methods, registration, and shape modeling with implicit neural representations and adequate regularisation terms to build personalised models of an evolving anatomical structure. In experiments with four longitudinally scanned AAA patients, we have demonstrated how the model represents the evolving shape of an AAA over time. This may impact patient monitoring and treatment; accurate knowledge about the progression of an AAA allows the physician to personalise surveillance and time intervention better based on AAA diameter and growth rate [16].

One appealing aspect of our approach is the continuity of the implicit neural representations. This allows us to reconstruct an AAA mesh at any point in time, at any desired resolution. We have here modeled shape changes over multiple years with sparse and irregularly spaced shape data. Modeling this change through linear interpolation of alternative surface representations, such as meshes or point clouds, would require point-to-point correspondence, a challenging problem that we here circumvent. Moreover, since our network relies on pointcloud data, it is agnostic to imaging modality. This is important for longitudinal studies of AAAs, where imaging modalities such as MRI and 3D US are increasingly used. All these scans can be incorporated into this framework as long as we can extract AAA surfaces. Furthermore, because we represent an evolving shape in space and time in a differentiable neural network, we can add any gradient-based regularisation term to the loss function. We have here included an Eikonal term and temporal regularization, but this framework could be further extended.

Lastly, we found that our model is sensitive to errors in the initial alignment of AAA shapes. Although we have followed [4] in registering based on the location of the vertebrae, better results can likely be achieved by registering based on other landmarks, such as the renal arteries and iliac bifurcation.

One limitation of the current approach is the relatively limited test set of four longitudinally scanned patients, which we aim to increase in future work. By increasing the data set and combining the here proposed model-driven with a data-driven approach [4, 9, 21], we might address current limitations of our model in extrapolation. Moreover, our current method is based on AAA morphology, but can be extended with biomarkers related to AAA growth and rupture [6]. Furthermore, additional optimization constraints could more properly model the pathophysiology of aneurysms. Whereas our temporal regularization term now aimed to minimise the gradient of the SDF, in future work, we could optimise this gradient within biologically plausible growth rates. This kind of regularization

could also be obtained in a data-driven way, by learning a generalisable model from a larger set of patients with longitudinal data. Finally, there is evidence that intraluminal thrombus shape plays a key role in AAA remodeling [21], and it might be beneficial to explicitly represent thrombus in our INR [1].

In conclusion, we have shown that INRs are promising tools in modeling AAA evolution. In future work, this flexible model could be extended with biologically plausible regularization terms and hemodynamic parameters.

References

1. Alblas, D., Brune, C., Yeung, K.K., Wolterink, J.M.: Going off-grid: continuous implicit neural representations for 3D vascular modeling. In: Statistical Atlases and Computational Models of the Heart. Regular and CMRxMotion Challenge Papers. STACOM 2022. LNCS, vol. 13593, pp. 79–90. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-23443-9_8
2. Amiranashvili, T., Lüdke, D., Li, H.B., Menze, B., Zachow, S.: Learning shape reconstruction from sparse measurements with neural implicit functions. In: International Conference on Medical Imaging with Deep Learning, pp. 22–34. PMLR (2022)
3. Brewster, D.C., Cronenwett, J.L., Hallett, J.W., Jr., Johnston, K.W., Krupski, W.C., Matsumura, J.S.: Guidelines for the treatment of abdominal aortic aneurysms: report of a subcommittee of the joint council of the American association for vascular surgery and society for vascular surgery. J. Vasc. Surg. **37**(5), 1106–1117 (2003)
4. Do, H.N., et al.: Prediction of abdominal aortic aneurysm growth using dynamical Gaussian process implicit surface. IEEE Trans. Biomed. Eng. **66**(3), 609–622 (2018)
5. Gharahi, H., Zambrano, B., Lim, C.Y., Choi, J., Lee, W., Baek, S.: On growth measurements of abdominal aortic aneurysms using maximally inscribed spheres. Med. Eng. Phys. **37**(7), 683–691 (2015)
6. Groeneveld, M.E., et al.: Systematic review of circulating, biomechanical, and genetic markers for the prediction of abdominal aortic aneurysm growth and rupture. J. Am. Heart Assoc. **7**(13), e007791 (2018)
7. Gropp, A., Yariv, L., Haim, N., Atzmon, M., Lipman, Y.: Implicit geometric regularization for learning shapes. In: International Conference on Machine Learning, pp. 3789–3799. PMLR (2020)
8. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nat. Methods **18**(2), 203–211 (2021)
9. Jiang, Z., Do, H.N., Choi, J., Lee, W., Baek, S.: A deep learning approach to predict abdominal aortic aneurysm expansion using longitudinal data. Front. Phys. **7**, 235 (2020)
10. Kim, S., et al.: Deep learning on multiphysical features and hemodynamic modeling for abdominal aortic aneurysm growth prediction. IEEE Trans. Med. Imaging **42**, 196–208 (2023)

11. Lüdke, D., Amiranashvili, T., Ambellan, F., Ezhov, I., Menze, B.H., Zachow, S.: Landmark-free statistical shape modeling via neural flow deformations. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2022. MICCAI 2022. LNCS, vol. 13432, pp. 453–463. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-16434-7_44

12. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: learning continuous signed distance functions for shape representation. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 165–174. IEEE Computer Society (2019)

13. Shen, L., Pauly, J., Xing, L.: NeRP: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. IEEE Transactions on Neural Networks and Learning Systems (2022)

14. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: NeurIPS (2020)

15. Sun, Y., Liu, J., Xie, M., Wohlberg, B., Kamilov, U.S.: Coil: coordinate-based internal learning for tomographic imaging. IEEE Trans. Comput. Imaging **7**, 1400–1412 (2021)

16. Wanhainen, A., et al.: Editor's choice-european society for vascular surgery (ESVS) 2019 clinical practice guidelines on the management of abdominal aorto-iliac artery aneurysms. Eur. J. Vasc. Endovasc. Surg. **57**(1), 8–93 (2019)

17. Wasserthal, J., Meyer, M., Breit, H.C., Cyriac, J., Yang, S., Segeroth, M.: TotalSegmentator: robust segmentation of 104 anatomical structures in CT images. arXiv preprint arXiv:2208.05868 (2022)

18. Wiesner, D., Suk, J., Dummer, S., Svoboda, D., Wolterink, J.M.: Implicit neural representations for generative modeling of living cell shapes. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2022. MICCAI 2022. LNCS, vol. 13434, pp. 58–67. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-16440-8_6

19. Wolterink, J.M., Zwienenberg, J.C., Brune, C.: Implicit neural representations for deformable image registration. In: International Conference on Medical Imaging with Deep Learning, pp. 1349–1359. PMLR (2022)

20. Xie, Y., et al.: Neural fields in visual computing and beyond. Comput. Graph. Forum **41**(2), 641–676 (2022)

21. Zhang, L., Zambrano, B.A., Choi, J., Lee, W., Baek, S., Lim, C.Y.: Intraluminal thrombus effect on the progression of abdominal aortic aneurysms by using a multistate continuous-time Markov chain model. J. Int. Med. Res. **48**(11), 0300060520968449 (2020)