Fast Few-Shot Embolization Simulations in Ischemic Stroke using Equivariant Neural Fields

Thijs P. Kuipers[1,2,3], Mahsa Mojtahedi[3], Emiel D. den Haan[3], Praneeta R. Konduri[4], Henk A. Marquering[1,2], and Erik J. Bekkers[5]

[1] Dept. of Biomedical Engineering and Physics, Amsterdam UMC, The Netherlands
[2] Dept. of Radiology and Nuclear Medicine, Amsterdam UMC, The Netherlands
[3] inSteps B.V., the Netherlands
[4] Stanford University, United States of America
[5] Informatics Institute, University of Amsterdam, The Netherlands
t.p.kuipers@amsterdamumc.nl

Abstract. Ischemic stroke is a leading cause of death and disability, in which rapid restoration of blood flow is critical to patient outcome. During embolization, the lodging location of the occluding clot within the cerebral vasculature strongly influences treatment difficulty, as clot stiffness, deformability, and vessel anatomy determine the clot retrievability during endovascular thrombectomy. High-fidelity finite-element method (FEM) simulations can simulate embolization, but their computational cost makes them impractical for real-time clinical use. We propose a fast and data-efficient meshless neural surrogate for simulating Lagrangian embolization dynamics. Our model predicts particle acceleration updates using SE(3)-equivariant neural fields, allowing training from sparsely sampled high-resolution simulations while supporting inference at full spatial resolution. By combining neural dynamics with rotation and translation equivariance, the surrogate generalizes across patient-specific vascular geometries using only 20 training simulations. On held-out test simulations, our approach achieves stable rollouts with a median error of 1.76 mm (Q1–Q3: 0.93–2.82), a median lodging error of 1.27 mm (0.77–1.42), corresponding to a median relative localization error of 1.8%, and a 200–400× speedup over FEM simulations, demonstrating both the accuracy and efficiency of our approach.

Keywords: neural dynamics · finite element method · implicit neural representation · digital twin · ischemic stroke

1 Introduction

Treating acute ischemic stroke (AIS) is a race against time: the brain loses an estimated 1.9 million neurons for every minute of untreated ischemia [19]. AIS occurs when a blood clot occludes an artery in the brain, preventing blood flow. In large-vessel occlusions, endovascular thrombectomy (EVT), in which the occluding clot is mechanically removed, is the standard of care. Clinical outcome is strongly time dependent, with the best recovery when the clot is

removed in a single pass [8, 23]. In addition, embolization, where the clot enters and subsequently lodges in the cerebral vasculature, plays an important role in determining downstream treatment difficulty and outcome. This phase determines factors including the clot's location, geometry, and mechanical state, all of which influence retrievability during EVT [7, 22, 16]. These factors vary substantially between patients, which has motivated the development of digital twins for patient-specific, in-silico modeling of endovascular treatment strategies and stroke progression [5, 2, 9]. Accurately capturing embolization, therefore, provides meaningful, patient-specific insight and establishes realistic initial conditions for subsequent EVT treatment simulations.

High-fidelity finite-element modeling (FEM) is capable of simulating EVT, including embolization, and can provide informative visualizations to support treatment planning [14, 11]. However, the computational cost of FEM simulations makes them impractical for time-critical clinical use. Neural surrogate models offer a potential alternative by learning system dynamics directly from simulation data. Eulerian formulations, where physical quantities evolve on a fixed spatial mesh, have previously been used to model artery wall shear stress prediction [21] and cardiovascular flow simulations [17]. However, embolization is inherently a Lagrangian process, characterized by particles moving through space. Lagrangian neural models have been explored for particle and mesh-based dynamics [18, 1], but these methods typically require large numbers of high-fidelity training simulations.

Embolization dynamics are equivariant to rigid body transformations, as rotating or translating the system does not alter clot–vessel interactions. We therefore propose a fast, meshless neural surrogate based on SE(3)-equivariant neural field representations, allowing for learning from sparsely sampled high-resolution simulations while supporting full-resolution inference. Our approach achieves strong generalization, outperforming state-of-the-art neural physics solvers X-MeshGraphNet [15] and Transolver++[12], across patient-specific vascular geometries while being trained on only 20 reference simulations. Code is available.[1]

2 Method

The equivariant neural dynamics model follows an encoder-decoder design built from a single shared equivariant cross-attention block: the encoder maps the particle state to a compact latent representation, and the decoder queries this latent state to compute clot accelerations. An overview is shown in Figure 1.

2.1 Problem formulation

We consider a discrete-time Lagrangian particle system for simulating clot dynamics within a blood vessel. At each time step t, the system state is represented as a point cloud:

$$\mathcal{S}^t = \{(\mathbf{x}_i^t, \mathbf{v}_i^t, \mathbf{u}_i^t))\}_{i=1}^N, \tag{1}$$

[1] `https://github.com/ThijsKuipers1995/neural_clot_sim`

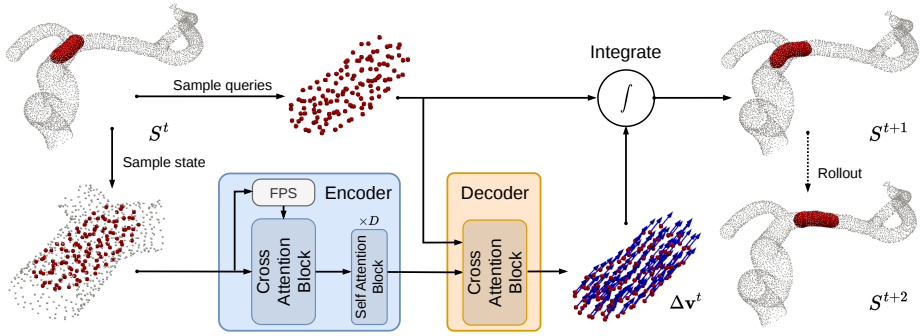

Fig. 1. Given the full system state S^t, we first perform spatial subsampling and discard vessel nodes located more than 2 mm away from the clot. The reduced state is encoded via cross-attention and farthest point sampling (FPS) to obtain a latent state. Subsequently, a decoder predicts velocity increments $\Delta \mathbf{v}^t$ for query points sampled on the clot volume and integrates them to advance the system to the next state S^{t+1}.

where $\mathbf{x}_i^t \in \mathbb{R}^3$ and $\mathbf{v}_i^t \in \mathbb{R}^3$ are the particle positions and velocities, and $\mathbf{u}_i^t = P(\mathbf{x}_i^t)$ is a precomputed pressure-driven flow field. Particles are divided into static nodes and dynamic nodes that evolve. The system evolves with a fixed timestep ($\Delta t = 1$). Our objective is to learn a function:

$$f_\theta : S_t \mapsto \{\Delta \mathbf{v}_i^t\}_{i=1}^N, \tag{2}$$

that predicts per-particle velocity increments given the current state. State updates are performed using a first-order Euler integration scheme:

$$\mathbf{v}^{t+1} = \mathbf{v}^t + \Delta \mathbf{v}^t \quad \text{and} \quad \mathbf{x}^{t+1} = \mathbf{x}^t + \mathbf{v}^{t+1}. \tag{3}$$

2.2 Embolization Simulations

We construct two embolization simulation datasets: a high-fidelity dataset with complex vascular geometries, and a simplified dataset with tapering tube geometries featuring a single 180-degree bend.

Complex Vessel Dataset We performed high-fidelity FEM simulations in Abaqus (SIMULIA, Providence, RI) of large vessel occlusion with varying cerebral vessel geometries. A 10 mm clot enters the internal carotid artery (ICA) and lodges in the proximal or distal segment of the first branch of the middle cerebral artery (MCA-M1). Clots are modeled using a Hyperfoam material formulation with parameters selected following prior computational thrombectomy studies [13, 4]. Clot entry and lodging are driven by a constant pressure load of 0.002 MPa. Vessel geometries are rigid and generated as described in [10]. Each simulation consists of 195 states and approximately 27000 nodes, and requires 30–60 minutes on an AMD Ryzen 9 9950X3D 16-core CPU.

Simple Bent Tube Dataset We additionally construct a simplified tube dataset consisting of 20 simulations of a bent cylindrical vessel in which the clot lodges at a single curvature. Each simulation contains 80 timesteps and 1000 nodes. This controlled setting serves as a sanity check to verify baseline behavior. It is not used for model selection or tuning.

2.3 SE(3)-Equivariant Neural Dynamics

The core of the proposed model is a G-equivariant kernelized cross-attention operator, which we adapt from [6]. For G, we consider the platonic subgroups of SO(3). For an input sequence of length N and feature dimension d, feature ($\mathbf{H}$) and coordinate ($\mathbf{X}$) tensors are defined over the group:

$$\mathbf{H} \in \mathbb{R}^{|G| \times N \times d}, \quad \mathbf{X} \in \mathbb{R}^{|G| \times N \times 3}. \tag{4}$$

Given query and context tensors $(\mathbf{H}_q, \mathbf{X}_q)$ and $(\mathbf{H}_c, \mathbf{X}_c)$, cross-attention is computed independently for each group element $g \in G$ in parallel using kernelized attention [3], which makes computation time scale linearly with the context length:

$$\mathrm{Attn}(\mathbf{H}_q, \mathbf{X}_q; \mathbf{H}_c, \mathbf{X}_c) = \frac{1}{N_c} \bigoplus_{g \in G} \mathbf{Q}^g ((\mathbf{K}^g)^\top V^g), \tag{5}$$

where $\bigoplus$ denotes concatenation along the group dimension and $(\cdot)^g \in \mathbb{R}^{N \times d}$ denotes tensors in the reference frame corresponding to group element g. The query, key, and value tensors are defined as:

$$\mathbf{Q}^g = \mathrm{RoPE}(\phi_q(\mathbf{H}_q)^g, \mathbf{X}_q^g), \quad \mathbf{K}^g = \mathrm{RoPE}(\mathbf{1}, \mathbf{X}_c^g), \quad \mathbf{V}^g = \phi_c(\mathbf{H}_c)^g. \tag{6}$$

Here, ϕ_q and ϕ_c are learnable equivariant linear projections and $\mathbf{1}$ denotes a tensor of ones. Rotary positional embeddings (RoPE) [20] inject relative geometric information, ensuring translation equivariance. Because the keys depend only on coordinates, information propagation is guided by the system's geometry.

Importantly, the attention operator retains the original kernelized attention form. Equivariance is induced by evaluating shared embeddings across transformed group elements. This weight-sharing constrains the operator to be G-equivariant, preserving the original computation graph, and inducing equivariance without additional computational overhead. When $q = c$, Eq. 5 reduces to the self-attention. We construct attention blocks containing one equivariant attention layer followed by an equivariant two-layer MLP with GeLU activation using pre-layer normalization.

Encoder At time t, the input state $\mathcal{S}^t$ consists of a sequence of N node coordinates ($\mathbf{x}_i^t$), scalar features (one-hot encoded node types distinguishing clot and vessel nodes), and vector features (velocity $\mathbf{v}_i^t$, flow $\mathbf{u}_i^t$, and zero-mean mesh-space coordinates $\mathbf{x}_i^{t=0}$). Vector features transform under rotations, whereas scalars remain invariant. Coordinates and vector features are lifted by rotating them for each $g \in G$. Scalars are copied across G. The input state is first

downsampled to $M < N$ latent points using farthest point sampling, ensuring uniform geometric coverage. The latent state is obtained using a single cross-attention block (Eq. 5) with the downsampled points as queries and the full-resolution state as context. The latent state is then updated using a sequence of D self-attention blocks.

Decoder and Neural Field Dynamics To predict per-particle velocity increments $\Delta\mathbf{v}^t$, query points are sampled from the clot domain at time t. Query features consist of the same scalar and vector quantities as in the encoder and are lifted to G accordingly. A single cross-attention block updates these queries using the latent dynamics state as context, defining a neural field representation of the system dynamics. The resulting equivariant features are projected back to $\mathbb{R}^3$ to obtain $\Delta\mathbf{v}^t$, which are integrated using Eq. 3.

2.4 Implementation and Training Details

The model uses the tetrahedral subgroup of SO(3) with $|G| = 12$ rotations and two attention heads per group element. We use 768 channels (32 channels per head). The encoder consists of $D = 1$ equivariant self-attention blocks. Input states are first sampled from the full simulation state by uniformly sampling 512 vessel nodes and 1536 clot nodes. Latent states consist of $M = 512$ points. During training, 2048 query points are sampled to predict velocity increments. Clot-vessel interactions are restricted to vessel nodes within radius $r = 2$ mm of clot nodes to introduce locality and improve computational efficiency. During training, both input states and queries are augmented with Gaussian noise with a standard deviation of 1×10^{-4}. The data is augmented with random rotations during training for all models, including baselines. Models are trained for 1000 epochs with batch size 8 using AdamW, an initial learning rate of 5×10^{-6}, cosine decay to 1×10^{-6}, and linear warmup during the first 10 epochs. We use mean squared error loss. Input features are scaled by a factor of 0.5. The 30 vessel simulations are split into 20 development cases and 10 held-out test cases. Five-fold cross-validation is performed on the development set using a 16/4 train/validation split. The final model is trained on all 20 development simulations and evaluated on the test set. The simple tube dataset uses a 13/1/6 train/val/test split. Models are trained on an NVIDIA RTX 5070 GPU.

3 Experiments and Results

We evaluate the predictive accuracy, rollout stability, computational efficiency, and generalization across vessel geometries. Model selection is performed via five-fold cross-validation, followed by comparison to established neural physics solver baselines on a held-out test set and additional analyses of robustness and out-of-domain behavior. During inference, rollouts stop when the clot's 10-step average velocity falls below 0.05 mm/Δt, indicating that it has lodged.

Table 1. Cross-validation results comparing SE(3)-equivariant and translation-equivariant models. Values are median (Q1-Q3) in mm across five folds, with RMSEs averaged over three seeds per trajectory. Best results in **bold**.

Equivariance	Per-step RMSE	Rollout RMSE	Lodging Error
Translation	**0.12** (0.11–0.13)	3.39 (3.18–4.29)	6.70 (3.05–7.32)
SE(3)	0.14 (0.12–0.14)	**2.33** (2.11–2.36)	**3.69** (2.35–5.45)

Table 2. Comparison of baseline models on the test set. We report the median as well as the interquartile range (Q1-Q3) in mm calculated over all test trajectories and three different seeds. Best results in **bold**. (MGN: X-MeshGraphNet).

Dataset	Model	Per-step RMSE	Rollout RMSE	Lodging Error
Tube	MGN	0.12 (0.11–0.14)	**0.24** (0.23–0.25)	0.40 (0.35–0.45)
	Transolver++	**0.05** (0.05–0.07)	0.44 (0.41–0.45)	0.48 (0.26–0.61)
	Ours	0.24 (0.24–0.25)	1.09 (1.09–1.10)	**0.39** (0.35–0.45)
Vessel	MGN	0.10 (0.08–0.13)	5.12(3.87–5.65)	18.0 (16.5–23.8)
	Transolver++	**0.07** (0.06–0.9)	3.34 (2.59–5.15)	17.7 (11.7–23.1)
	Ours	0.14 (0.13–0.16)	**1.76** (0.93–2.82)	**1.27** (0.77–1.42)

Model Validation: 5-Fold Cross-Validation We assess the effect of rotational equivariance by comparing the proposed SE(3)-equivariant model to a translation-equivariant variant. Performance is evaluated using three metrics on the vessel dataset. The per-step root mean squared error (RMSE) measures one-step prediction accuracy by comparing predictions computed from the reference state at each timestep of the simulation. The rollout RMSE quantifies error accumulation when the model is iteratively applied to generate the full trajectory and assesses simulation stability. Lodging error is the Euclidean distance between the final predicted and reference clot positions. RMSEs are averaged over three seeds per trajectory and across trajectories within each fold.

The final results are reported as median (Q1–Q3) across folds in Table 1. Both models achieve comparable per-step RMSE, while the SE(3)-equivariant model yields substantially lower rollout and lodging errors, indicating improved rollout stability. Our model achieves a full-resolution rollout in approximately 8 seconds, corresponding to a 200–400× speedup when compared to the high-fidelity FEM simulations.

Test-Set Comparison to Baselines We evaluate on held-out simulations from both the simple tube and complex vessel datasets. We compare against two state-of-the-art neural physics solvers: X-MeshGraphNet [15] (MGN) and Transolver++ [12]. Baselines and training details are adapted from their original implementations and tuned on the validation sets. MGN uses a single scale 5-nearest-neighbor graph augmented with explicit clot-vessel interaction edges, 12

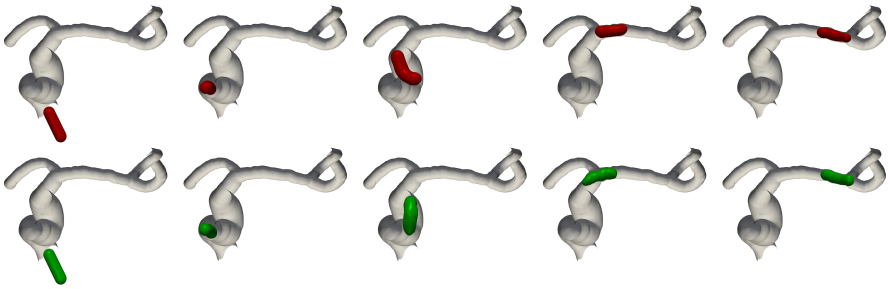

Fig. 2. Rollout trajectory of the clot (length: 10 mm) within the vessel. The predicted rollout is shown in red, while the reference (ground truth) trajectory is shown in green.

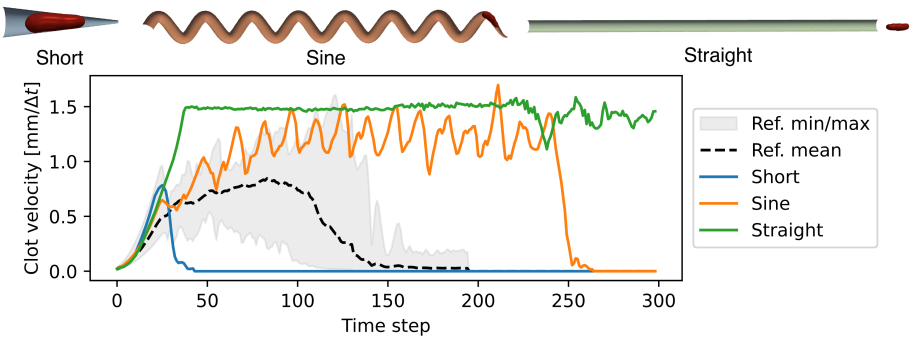

Fig. 3. Out-of-domain geometries (short, sine, and straight) with predicted clot rollouts that induce significantly different velocity profiles than seen in the training set.

message-passing layers with 128 channels. Transolver++ uses 5 attention blocks with 8 heads with 256 channels.

We report our results in Table 2, with a representative rollout shown in Fig. 2. On the tube dataset, all methods perform well. On the vessel dataset, Transolver++ maintains the lowest per-step error, while our method achieves substantially lower rollout and lodging errors than both baselines. The median lodging error of 1.27 mm corresponds to approximately 1.8% of the total clot trajectory length. We note that the baseline models exhibit instabilities during all rollouts on the vessel dataset, including clipping through the vessel wall and partial clot fragmentation. In contrast, these instabilities are absent in our model, although one out of the ten trajectories showed consistent early lodging. Because the model depends on stochastic input states, we assess stability across three seeds using the mean pairwise Euclidean distance. Rollouts with different seeds differ on average by 1.02 mm (SD 1.14 mm).

Out-of-Domain Geometries We assess generalization to unseen vessel geometries using three out-of-domain tubes: short (20 mm), sine (200 mm), and straight (160 mm). The short and sine tubes taper from 2.5 mm to 0.2 mm, while the straight tube maintains a constant radius of 2.5 mm. We evaluate whether the model lodges the clot earlier, later, or not at all, and whether lodging is driven by geometric narrowing rather than implicit velocity thresholds.

Results are shown in Fig. 3. In the short and sine tubes, the model successfully lodges the clot (velocity goes to zero), occurring notably earlier and later than in the training set. No lodging occurs in the straight tube, as expected. Notably, for the sine tube, the number of velocity oscillations matches the number of geometric extrema, indicating a direct coupling between vessel curvature and clot dynamics. The sine and straight tubes require substantially longer rollouts (250–300 steps) than the reference simulations (130–160 steps), demonstrating stable long-horizon behavior, as well as generalization to out-of-domain geometries.

4 Discussion and Conclusion

We have presented a data-efficient, meshless neural surrogate for simulating Lagrangian dynamics of embolization in ischemic stroke. Our validation experiments demonstrate that explicitly incorporating SE(3) equivariance substantially improves long-horizon stability. While per-step errors are similar to a translation-equivariant variant, the equivariant model better captures the coupled interactions between clot and vessel, which are highly sensitive to small perturbations in position and velocity. By enforcing rotational and translational symmetry, the model generalizes across different orientations and positions of the system, resulting in more robust rollouts and more accurate clot lodging even with only 20 training simulations.

Comparisons to state-of-the-art neural physics solvers further highlight the importance of equivariance for long-horizon stability in a few-shot setting. Both MGN and Transolver++ achieved low per-step errors on both datasets. Despite successfully simulating the simple dataset, the baselines do not reliably lodge the clot on the complex vessel dataset. It seems that their lack of explicit constraints on rotational symmetry prevents these models from accurately capturing the intricate interactions between the clot and vessel. This leads to error accumulation over time and unstable rollouts. In contrast, our model demonstrates physically plausible behavior across diverse and out-of-distribution vessel geometries, including long trajectories and correct clot lodging driven by geometric constraints rather than implicit velocity thresholds.

While our approach demonstrates strong performance in a few-shot regime, we do not assess how well it scales with larger training datasets. Additionally, although the model can perform rollouts at full resolution, it still relies on a fixed-resolution encoding of the input state. The current simulation setup uses a single set of clot mechanical parameters and a constant, clot-independent pressure field. Addressing these limitations and extending our model to include addi-

tional stages of the EVT treatment are important next steps toward full in-silico treatment simulation.

Overall, our results indicate that meshless, SE(3)-equivariant neural field dynamics can efficiently capture the complex clot–vessel interactions that determine lodging and trajectory, providing accurate, robust predictions in a few-shot regime. Our approach enables fast, patient-specific simulations of embolization in ischemic stroke and lays the groundwork for future extensions to full neural treatment modeling.

Acknowledgments. This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101136438 (GEMINI project). This publication reflects only the author's view and the European Commission is not responsible for any use that may be made of the information it contains.

Disclosure of Interests. Authors Thijs P. Kuipers, Mahsa Mojtahedi, and Emiel D. den Haan were employed by inSteps B.V. during the conduct of this work. Henk A. Marquering owns stock in inSteps B.V. and Nicolab B.V. Praneeta R. Konduri owns stock in inSteps B.V.

Bibliography

[1] Alkin, B., Fürst, A., Schmid, S., Gruber, L., Holzleitner, M., Brandstetter, J.: Universal physics transformers: A framework for efficiently scaling neural operators. Advances in Neural Information Processing Systems **37**, 25152–25194 (2024)

[2] Chaparro-Cárdenas, S.L., Ramirez-Bautista, J.A., Terven, J., Córdova-Esparza, D.M., Romero-Gonzalez, J.A., Ramírez-Pedraza, A., Chavez-Urbiola, E.A.: A technological review of digital twins and artificial intelligence for personalized and predictive healthcare. In: Healthcare. vol. 13, p. 1763. MDPI (2025)

[3] Choromanski, K., Likhosherstov, V., Dohan, D., Song, X., Gane, A., Sarlos, T., Hawkins, P., Davis, J., Mohiuddin, A., Kaiser, L., et al.: Rethinking attention with performers. arXiv preprint arXiv:2009.14794 (2020)

[4] Fregona, V., Luraghi, G., Fereidoonnezhad, B., Gijsen, F.J., Majoie, C.B., Matas, J.F.R., Migliavacca, F.: Impact of thrombus composition on virtual thrombectomy procedures using human clot analogues mechanical data. Journal of the mechanical behavior of biomedical materials **163**, 106886 (2025)

[5] Hoekstra, A.G., Marquering, H., consortium https://dth-gemini. eu/team/, G.: Towards a generation of digital twins in healthcare of ischaemic and haemorrhagic stroke. In: International Conference on Computational Science. pp. 239–245. Springer (2024)

[6] Islam, M.M., Anand, R., Wessels, D.R., de Kruiff, F., Kuipers, T.P., Ying, R., Sánchez, C.I., Vadgama, S., Bökman, G., Bekkers, E.J.: Platonic transformers: A solid choice for equivariance. arXiv preprint arXiv:2510.03511 (2025)

[7] Johnson, S., Chueh, J., Gounis, M.J., McCarthy, R., McGarry, J.P., McHugh, P.E., Gilvarry, M.: Mechanical behavior of in vitro blood clots and the implications for acute ischemic stroke treatment. Journal of neurointerventional surgery **12**(9), 853–857 (2020)

[8] Kang, D.H., Kim, B.M., Heo, J.H., Nam, H.S., Kim, Y.D., Hwang, Y.H., Kim, Y.W., Kim, D.J., Kim, J.W., Baek, J.H., et al.: Effects of first pass recanalization on outcomes of contact aspiration thrombectomy. Journal of NeuroInterventional Surgery **12**(5), 466–470 (2020)

[9] Katsoulakis, E., Wang, Q., Wu, H., Shahriyari, L., Fletcher, R., Liu, J., Achenie, L., Liu, H., Jackson, P., Xiao, Y., et al.: Digital twins for health: a scoping review. NPJ digital medicine **7**(1), 77 (2024)

[10] Kuipers, T.P., Konduri, P.R., Marquering, H., Bekkers, E.J.: Generating cerebral vessel trees of acute ischemic stroke patients using conditional set-diffusion. In: Medical Imaging with Deep Learning (2024)

[11] Liu, R., Jin, C., Wang, L., Yang, Y., Fan, Y., Wang, W.: Simulation of stent retriever thrombectomy in acute ischemic stroke by finite element analysis.

Computer Methods in Biomechanics and Biomedical Engineering **25**(7), 740–749 (2022)
[12] Luo, H., Wu, H., Zhou, H., Xing, L., Di, Y., Wang, J., Long, M.: Transolver++: An accurate neural solver for pdes on million-scale geometries. arXiv preprint arXiv:2502.02414 (2025)
[13] Luraghi, G., Bridio, S., Lissoni, V., Dubini, G., Dwivedi, A., McCarthy, R., Fereidoonnezhad, B., McGarry, P., Gijsen, F.J., Matas, J.F.R., et al.: Combined stent-retriever and aspiration intra-arterial thrombectomy performance for fragmentable blood clots: A proof-of-concept computational study. Journal of the mechanical behavior of biomedical materials **135**, 105462 (2022)
[14] Luraghi, G., Bridio, S., Matas, J.F.R., Dubini, G., Boodt, N., Gijsen, F.J., van der Lugt, A., Fereidoonnezhad, B., Moerman, K.M., McGarry, P., et al.: The first virtual patient-specific thrombectomy procedure. Journal of Biomechanics **126**, 110622 (2021)
[15] Nabian, M.A., Liu, C., Ranade, R., Choudhry, S.: X-meshgraphnet: Scalable multi-scale graph neural networks for physics simulation. arXiv preprint arXiv:2411.17164 (2024)
[16] Patki, P., Simon, S., Manning, K.B., Costanzo, F.: Computational analysis of effects of clot length on acute ischemic stroke recanalization under different cyclic aspiration loading conditions. International journal for numerical methods in biomedical engineering **39**(2), e3667 (2023)
[17] Pegolotti, L., Pfaller, M.R., Rubio, N.L., Ding, K., Brufau, R.B., Darve, E., Marsden, A.L.: Learning reduced-order models for cardiovascular simulations with graph neural networks. Computers in Biology and Medicine **168**, 107676 (2024)
[18] Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., Battaglia, P.: Learning mesh-based simulation with graph networks. In: International conference on learning representations (2020)
[19] Saver, J.L.: Time is brain—quantified. Stroke **37**(1), 263–266 (2006)
[20] Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., Liu, Y.: Roformer: Enhanced transformer with rotary position embedding. Neurocomputing **568**, 127063 (2024)
[21] 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. Computers in Biology and Medicine **173**, 108328 (2024)
[22] Vandelanotte, S., Staessens, S., François, O., De Wilde, M., Desender, L., De Sloovere, A.S., Dewaele, T., Tersteeg, C., Vanhoorelbeke, K., Vanacker, P., et al.: Association between thrombus composition and first-pass recanalization after thrombectomy in acute ischemic stroke. Journal of Thrombosis and Haemostasis **22**(9), 2555–2561 (2024)
[23] Zaidat, O.O., Castonguay, A.C., Linfante, I., Gupta, R., Martin, C.O., Holloway, W.E., Mueller-Kronast, N., English, J.D., Dabus, G., Malisch, T.W., et al.: First pass effect: a new measure for stroke thrombectomy devices. Stroke **49**(3), 660–666 (2018)