

World Model for AI Autonomous Navigation in Mechanical Thrombectomy

Harry Robertshaw¹[0000–0002–5486–8062], Han-Ru Wu²[0000–0001–6717–6615],
Alejandro Granados¹[0000–0002–2866–1324], and Thomas C
Booth^{1,3}[0000–0003–0984–3998]

¹ Surgical & Interventional Engineering, School of Biomedical Engineering & Imaging Sciences, King’s College London, UK

² Nuffield Department of Primary Care Health Sciences, University of Oxford, UK

³ Department of Neuroradiology, King’s College Hospital, UK
thomas.booth@kcl.ac.uk

Abstract. Autonomous navigation for mechanical thrombectomy (MT) remains a critical challenge due to the complexity of vascular anatomy and the need for precise, real-time decision-making. Reinforcement learning (RL)-based approaches have demonstrated potential in automating endovascular navigation, but current methods often struggle with generalization across multiple patient vasculatures and long-horizon tasks. We propose a world model for autonomous endovascular navigation using TD-MPC2, a model-based RL algorithm. We trained a single RL agent across multiple endovascular navigation tasks in ten real patient vasculatures, comparing performance against the state-of-the-art Soft Actor-Critic (SAC) method. Results indicate that TD-MPC2 significantly outperforms SAC in multi-task learning, achieving a 65% mean success rate compared to SAC’s 37% ($p < 0.001$), with notable improvements in path ratio. TD-MPC2 exhibited increased procedure times, suggesting a trade-off between success rate and execution speed. These findings highlight the potential of world models for improving autonomous endovascular navigation and lay the foundation for future research in generalizable AI-driven robotic interventions.

Keywords: World model · Reinforcement learning · Mechanical Thrombectomy · Autonomous Navigation · AI-driven Robotics.

1 Introduction

Stroke is a leading cause of death and disability worldwide, imposing significant strain on healthcare systems and families [18]. Mechanical thrombectomy (MT) has improved ischemic stroke outcomes, reducing mortality and disability compared to medical therapy alone [4,21]. However, its effectiveness declines with delayed treatment, emphasizing the need for rapid intervention [21,1]. Despite its efficacy, only 3.1% of stroke admissions in, for example, the UK receive MT, far below the estimated 15% eligibility rate, due to limited access to MT-capable centers and long transfer times [26,19,20]. Additionally, operators face

radiation exposure, increasing cancer and cataract risks, while protective gear contributes to orthopedic strain [16,28,17]. Robotic surgical systems offer a potential solution by improving accessibility and reducing operator dependency [25]. Tele-operated robotic MT could (1) allow specialists in centralized centers to perform procedures remotely, while (2) AI assistance robots may enable less experienced operators (interventional radiologists in peripheral hospitals) to perform MT effectively. Integrating AI into robotic systems could further enhance efficiency and safety, with autonomous surgical robots already demonstrating superior performance over manual techniques in some scenarios [23].

Autonomous endovascular navigation research has primarily focused on the aortic arch [13,15,25], with recent work exploring micro-guidewire and micro-catheter navigation in MT’s second phase [3]. These studies use multiple test vasculatures, but examine simple tasks with short episodes, limiting clinical applicability. Additionally, autonomous two-device navigation for phase one of MT has been investigated using a larger episode size but was limited to testing on a single vasculature [24]. These studies use reinforcement learning (RL) algorithms, the majority of which are sensitive to architecture, hyperparameters, and are unable to perform over long time horizons, while often being designed for single-task learning only [11,7]. For complex or multiple tasks, this limits RL to computationally expensive models or tasks where tuning is prohibitive [9].

To move towards realizing the benefits of autonomous MT navigation, it is necessary to develop policies capable of performing long navigation tasks across multiple patients vasculatures that can be adapted efficiently to new robots, tasks, and environments [22]. One promising technique may be world models, a learned representation of the environment that simulates its dynamics, enabling a single agent to optimize actions in a virtual setting without relying solely on real-world data [8]. Using this model, large-scale learning from diverse datasets could create AI navigation systems that can understand, predict, and adapt to real-world complexities [12]. Recent work has shown that single configurations of RL algorithms (DreamerV3 and TD-MPC2) with no hyperparameter tuning can outperform specialized methods across diverse benchmark tasks [9,10]. They can also complete farsighted tasks such as collecting diamonds in Minecraft without human data or curricula and capturing expectations of future events during autonomous driving [9,12]. Although the translation of these models to real-world applications is limited, they hold the potential for creating agents capable of performing long navigation tasks on a diverse range of patient anatomies. TD-MPC2 has shown significant improvements upon both Soft Actor-Critic (SAC) (a configuration of which is the current state-of-the-art for autonomous endovascular interventions compared to benchmarks [15,2]), and DreamerV3 when examining multi-task environments with continuous action spaces. It is also able to utilize multiple data types, such as human demonstrator, RL collected, and its own online interaction data (from a single task or across multiple tasks).

The aim of this study was to propose a framework for a world model for autonomous endovascular interventions, and more specifically, MT. The primary objective was to demonstrate that a singular RL agent could be used to perform

multiple endovascular navigation tasks across multiple patient vasculatures. Our contributions are as follows: 1) we implemented a world model capable of performing multiple endovascular navigation tasks across multiple patient vasculatures, for the first time, 2) we proposed a framework for endovascular intervention multi-task learning, 3) we compared our results to the current state-of-the-art RL algorithms for autonomous endovascular interventions, demonstrating performance across the largest dataset examined for autonomous MT.

2 Methods

2.1 Navigation tasks

MT typically involves navigating a guide catheter from the femoral or radial artery to the internal carotid artery (ICA). An ‘access catheter’ is usually placed within the guide catheter and taken ahead of the guide catheter tip during navigation. Once the access catheter is within the carotid artery, the guide catheter is advanced to make a stable platform. At this point, the guidewire and access catheter are retracted, and a micro-guidewire within a micro-catheter is passed through the stable guide catheter and navigated to the target thrombus site. Frequently, the final step is to remove the micro-guidewire and exchange it for a stent retriever to remove the thrombus, restoring blood flow to the brain.

The first phase of MT (navigation of guidewire and guide catheter from the femoral to the ICA) has been split into five tasks for multi-task RL training. These are based on previously defined phases of MT [5] and are described in Figure 1. A target was randomly sampled from the set of points within the target vessel, and guide catheter and guidewire navigated to it from a randomly sampled point in the starting vessel.

The device’s distal position was described by three points equally spaced 2 mm apart along the guidewire, denoted as $(x'_g, y'_g)_{i=1,2,3}$, with $(x'_g, y'_g)_1$ representing the instrument tip. The target location was specified by the coordinates (x'_t, y'_t) . Observations comprised current and previous device positions, target location, and the previous action taken.

2.2 Data collection

The *in silico* environment for the navigation task utilized the stEVE framework (Available at: <https://github.com/lkarstensen/stEVE> [15]). The BeamAdapter plugin for Simulation Open Framework Architecture (SOFA, v23.06) was used to model the Neuron MAX 0.088" guide catheter (Penumbra, California, USA) (160 vertices with Young’s modulus of 47 MPa) and 0.035" guidewire (Terumo, Tokyo, Japan) (120 vertices with Young’s modulus of 43 MPa) used for navigation [27,6]. The simulation assumed rigid vessel walls (with an empty lumen). The simulation’s real-world device behavior was determined by utilizing a tensile testing machine to measure the tensile strength of the devices, which facilitated the calculation of their stiffness. Friction between the wall and




Fig. 1. Full MT vasculature with each navigation task labeled. A_1 : Common iliac artery to top of descending aorta, A_{2L} : top of descending aorta to left common carotid artery (CCA), A_{2R} : top of descending aorta to right CCA, A_{3L} : left CCA to left ICA, A_{3R} : right CCA to right ICA.

guidewire was iteratively tuned to mimic guidewire behavior in a test-bench set-up. Such methodology previously allowed the translation of autonomous endovascular navigation agents from *in silico* to *in vitro* aortic arch models and from *in silico* to *ex vivo* porcine liver vasculatures [15,14].

A computed tomography angiography (CTA) scan which encompassed the aortic arch and extended to the cerebral vessels, and a CTA scan encompassing the abdominal and thoracic regions, including the femoral arteries, descending aorta, and the aortic arch (obtained with UK Research Ethics Committee AN/ON/YMOUS), were processed into surface meshes using segmentation tools in 3D Slicer (v5.8.0). The centerlines with radii for all scans were then extracted. The centerlines of the abdominal and thoracic regions were then scaled appropriately so that the radius at the last point of the descending aorta would match the first point on the aortic arch centerline. These two sets of centerlines were then joined, and the radius at each centreline point was used to generate a surface mesh, which was loaded into SOFA. This was repeated for ten aortic arch and cerebral vessel CTAs, with the same CTA of the abdominal and thoracic regions being scaled to fit each one. The mean right and left-hand-side vessel tortuosity were 1.19 ± 0.07 mm and 1.14 ± 0.05 mm, respectively. Type-I aortic arches were found in 80% (8/10) of vasculatures, while the remaining were Type-II. Additional augmentation of the entire mesh was applied during training via random scaling (0.7 to 1.3 for height and width) to enhance generalization.

Input parameters, including device rotation and translation speed, were applied at the proximal device end. Rotation and translation speed was constrained

to a maximum of 180°s^{-1} and 40 mm s^{-1} , respectively. Similar to a clinical scenario with fluoroscopy, feedback during the navigation was given as two-dimensional (x', y') tracking coordinates of three points along each device’s tip; no visual information showing the geometry of the patient vasculature was given.

2.3 Controller architecture

Two types of RL agents were trained in this study. One type used an implementation of the model-free RL algorithm, SAC from the stEVE framework [15], whereby the critic learns the value and the actor optimizes the critic directly to maximize cumulative rewards [7]. This is useful for continuous action spaces, leveraging experience replay for data efficiency, but requires tuning and struggles under high-dimensional inputs [9]. The architecture includes a Long Short-Term Memory (LSTM) layer for learning trajectory-dependent state representations and feedforward layers for controlling the devices. The controller takes observations as input, and a Gaussian policy network outputs mean (μ) and standard deviation (σ) for expected actions, representing the micro-catheter and micro-guidewire rotations and translations. During training, actions are sampled from the σ , but for evaluation, μ is used directly for deterministic behavior.

The other type of agent was trained using TD-MPC2, a model-based RL algorithm designed for sample-efficient learning and effective planning in continuous action spaces (Available at: <https://github.com/nicklashansen/tdmpc2> [10]). TD-MPC2 combines temporal difference learning with model-predictive control, leveraging a learned dynamics model to simulate environment transitions and plan action sequences over a short horizon. This approach enables the agent to optimize actions based on both predicted rewards and task constraints while using limited real-world interaction. By incorporating a latent dynamics model and a cross-entropy planning method, TD-MPC2 is particularly well-suited for complex tasks with longer episode lengths. This study used the base TD-MPC2 configuration but increased the length of the replay buffer from 1×10^6 to 1×10^7 . Additionally, we implemented an LSTM layer as an observation embedder, allowing storage of the environment state at each step so that the vessel structure could be estimated by the path of the device tip.

The dense reward function used during training in both models is shown in Equation 1. *Pathlength* is defined as the distance between the guidewire tip and the target along the centerlines of the arteries, with $\Delta\text{pathlength}$ representing the change in pathlength at time t from the previous step at time $t = -1$.

$$R = -0.00015 - 0.001 \cdot \Delta\text{pathlength} + \begin{cases} 1 & \text{if target reached} \\ 0 & \text{else} \end{cases} \quad (1)$$

2.4 Training and evaluation

Initially, five agents were trained using SAC (one for each defined phase in Figure 1). Each of these agents was then evaluated for 250 episodes while the trajectory data was recorded. This data was used to fill the replay buffer of the

SAC world model and the TD-MPC2 world model before initiating training. All models were trained for 1×10^7 exploration steps, with each navigation task (or episode) considered complete when the target was reached within a threshold equal to the entire cross-section of the vessel. An episode termination of 200 steps (approximately 27 s) was set for computational efficiency. Training was performed on an NVIDIA DGX A100 node (Santa Clara, California, USA).

We train and test on 10 patient-specific vascular anatomies, with evaluations conducted every 2.5×10^5 exploration steps for 250 episodes, recording the success rate, procedure time, and path ratio. Comparative statistical analyses were conducted using two-tailed paired Student’s t-tests with a predetermined significance threshold set at $p = 0.05$.

3 Results

3.1 Single task training

Results from single-task training for each of the five MT phases can be seen in Table 1. Success rates of 100% were recorded for phases A_1 and A_{2R} , while success rates and high path ratios were exhibited in phase A_{2L} (SR: 82%, PR: 90%), A_{2R} (SR: 90%, PR: 90%), and A_{3L} (SR: 92%, PR: 93%). This data was used to fill replay buffers before training for the world models in Section 3.2.

Table 1. Single task training using SAC. *Success Rate*: percentage of evaluation episodes where the target is reached. *Path Ratio*: remaining distance to the target point in unsuccessful episodes, calculated by dividing the remaining distance by the initial distance. *Procedure Time*: time from the start of navigation to the target location for successful episodes. *Exploration Steps*: number of training steps taken to reach the point at which the results are provided. The reported values are mean \pm standard deviation (standard deviation values may exceed logical bounds (0–100%) due to the statistical calculation).

Task	Success Rate (%)	Procedure Time (s)	Path Ratio (%)	Exploration Steps
A_1	100 \pm 0	9.8 \pm 1.5	100 \pm 0	0.25×10^5
A_{2L}	82 \pm 39	12.3 \pm 9.9	90 \pm 18	8.5×10^6
A_{2R}	100 \pm 0	10.7 \pm 7.4	100 \pm 0	9.5×10^6
A_{3L}	90 \pm 30	8.7 \pm 6.1	90 \pm 26	2.5×10^6
A_{3R}	92 \pm 27	10.9 \pm 6.9	93 \pm 17	7.0×10^6

3.2 World model

Results from training on multiple vasculatures and navigation tasks showed a significant increase in mean success rate across all tasks for TD-MPC2, 65% compared to 37% for the current state-of-the-art (SAC) ($p < 0.001$). These results can be seen in Table 2. Furthermore, significant increases in success rate

and path ratio were seen across tasks when moving from SAC to TD-MPC2: A_{2R} (SR: 8% to 66% [$p < 0.001$], PR: 29% to 77% [$p < 0.001$]), A_{3L} (SR: 16% to 50% [$p < 0.001$], PR: 26% to 6% [$p < 0.001$]), and A_{3R} (SR: 56% to 90% [$p < 0.001$], PR: 63% to 90% [$p = 0.001$]). Additionally, an increase in procedure time was observed across all tasks when moving from SAC to TD-MPC2. This was significant in A_1 (PT: 8.9% to 16.7% [$p < 0.001$]), A_{3L} (PT: 7.2% to 13.7% [$p = 0.021$]), and A_{3R} (PT: 5.6% to 12.9% [$p < 0.001$]). Exploration steps recorded for SAC and TD-MPC2 were 4.5×10^6 (reached after 75 hours) and 0.5×10^6 (reached after 25 hours), respectively.

Table 2. World model results per navigation task, with mean results across all tasks and vasculatures.

Task	SAC			TD-MPC2		
	Success Rate (%)	Procedure Time (s)	Path Ratio (%)	Success Rate (%)	Procedure Time (s)	Path Ratio (%)
A_1	96 ± 20	8.9 ± 1.6	89 ± 16	96 ± 20	16.7 ± 4.1	90 ± 3
A_{2L}	10 ± 30	14.5 ± 7.1	27 ± 30	22 ± 41	17.0 ± 6.1	45 ± 33
A_{2R}	8 ± 27	12.8 ± 9.9	29 ± 27	66 ± 47	16.5 ± 6.2	77 ± 32
A_{3L}	16 ± 37	7.2 ± 7.0	26 ± 34	50 ± 50	13.7 ± 6.3	66 ± 38
A_{3R}	56 ± 50	5.6 ± 4.5	63 ± 43	90 ± 30	12.9 ± 5.8	90 ± 21
Mean	37 ± 48	8.2 ± 4.6	47 ± 40	65 ± 48	15.2 ± 5.8	73 ± 33

4 Discussion

The results of this study demonstrate notable progress in the development of RL models for autonomous MT, addressing multiple navigation tasks across ten real patient vasculatures in a clinically-based simulation using industry-standard devices. This represents the largest dataset examined for autonomous MT, surpassing previous studies that focused on limited anatomical variations and smaller-scale evaluations [24,15]. The results from single-task training indicate that SAC performs well in short tasks, achieving success rates of 100% in phases A_1 and A_{2R} . Other tasks, such as A_{2L} , A_{3L} , and A_{3R} , exhibited moderate to high success rates, with corresponding path ratios. These results allowed suitable data to be used for inputs to the world model, and illustrate SAC’s ability to effectively learn task-specific policies with high levels of accuracy and path efficiency when trained on shorter tasks. The reduced performance on left-sided navigation tasks (A_{2L} and A_{3L}) highlights potential challenges related to increased vasculature complexity. This variation requires further exploration to identify anatomical features that influence performance but reflects clinical scenarios where the left CCA is typically harder to catheterize and may need specialist access catheters.

The world model results demonstrate a clear advantage of TD-MPC2 over SAC. A statistically significant increase in mean success rate was observed, with TD-MPC2 achieving 65% compared to 37% for SAC ($p < 0.001$). Notably, TD-MPC2 significantly outperformed SAC in A_{2R} , A_{3L} , and A_{3R} , while recording the same or non-significant increases in success rate for A_1 and A_{2L} . This

highlights the advantages of TD-MPC2’s model-based framework, leveraging a learned dynamics model and horizon-based planning to optimize actions in diverse and complex environments. SAC’s model-free approach may struggle to generalize effectively to anatomical variations and multiple objectives. However, in this context and other work, TD-MPC2 failed to solve all tasks it was trained on and exhibits limited scalability due to online planning [7]. The findings suggest that TD-MPC2 is better suited for scenarios requiring multi-task generalization and adaptability to unseen environments. In TD-MPC2, the lower success rates seen in A_{2L} (22%) and A_{3L} (50%) compared to other tasks may be attributed to the lower performance of the corresponding single-task models. Despite gains in success rate and path ratio, TD-MPC2 resulted in increased procedure time across all tasks. The longer procedure times suggest that TD-MPC2 may conceivably prioritize exploration and more deliberate navigation strategies, leading to improved performance but at the cost of increased procedure time. Future work could explore optimizing the trade-off between efficiency and exploration in navigation.

These results underscore the importance of scalable and generalizable RL approaches for realistic clinical applications of autonomous MT. While the success rate of TD-MPC2 (65%) highlights the need for further optimization, it represents a significant step toward bridging the gap between experimental performance and clinical feasibility. The comparison of single and multi-task training reveals critical trade-offs between task specialization and generalization. SAC demonstrated high success rates in single-task training, but limited applicability to real-world clinical scenarios, where agents must navigate diverse vasculatures. Conversely, TD-MPC2’s superior performance in multi-task settings indicates its potential as a robust framework for autonomous MT navigation, though its performance remains suboptimal for clinical deployment. For world models, performance relies on pre-existing data to train the world model, which might not always be feasible, especially in novel environments such as this one. Here we propose a balance between number of tasks and a single agent to execute them.

Future work should focus on improving TD-MPC2’s training pipeline to improve generalization to unseen diverse vasculatures. Current world models have demonstrated scalability to 80-150 tasks [9,7,10], underscoring their potential for future work as it looks to incorporate more diverse datasets, more detailed simulation environments that consider contact forces [3], while transitioning to *in vitro* testbeds. Although the vessels navigated in A_1 and A_2 are almost entirely rigid in practice, vessel deformation and motion for A_3 may enhance *in vitro* translation. Using differentiable physics simulators could further enable efficient policy learning using First-order Gradients, while work is needed to provide gradient stability over long horizons. Future work may also compare image-based RL, such as DreamerV3 [9], to the current approach, to evaluate if human-like visual input provides benefit. Additionally, balancing task-specific performance with generalization should remain a priority for advancing the clinical utility of AI-based autonomous endovascular navigation, paving the way for safer and more efficient outcomes in real-world applications.

5 Conclusion

This study demonstrates the feasibility of a world model-based RL approach for autonomous MT, addressing key challenges in generalization and multi-task learning. The proposed world model leverages existing open-source repositories to maximize reproducibility [15,10]. Through evaluations across multiple patient vasculatures, TD-MPC2 achieved superior generalization compared to SAC, with a significant increase in success rate and path ratio. Future research should focus on refining world models for more efficient planning, integrating differentiable physics simulators to enhance gradient stability, and expanding datasets for better generalization across diverse vasculatures. The findings presented here provide a critical step toward the development of AI-driven autonomous robotic endovascular interventions with improved safety and precision.

Acknowledgments. Partial financial support was received from the WELLCOME TRUST (Grant Agreement No 203148/A/16/Z), the Engineering and Physical Sciences Research Council Doctoral Training Partnership (Grant Agreement No EP/R513064/1), and the MRC IAA 2021 Kings College London (MR/X502923/1). For the purpose of Open Access, the Author has applied a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission.

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

References

1. Asdaghi, N., Wang, K., Gardener, H., Jameson, A., Rose, D.Z., Alkhachroum, A., et al.: Impact of time to treatment on endovascular thrombectomy outcomes in the early versus late treatment time windows. *Stroke* **54**, 733–742 (3 2023). <https://doi.org/10.1161/STROKEAHA.122.040352>
2. Author, A.: Benchmarking reinforcement learning algorithms for autonomous mechanical thrombectomy. This paper is currently under review. An anonymized version of this work has been submitted as supplementary material
3. Author, A.: Reinforcement learning for safe autonomous two device navigation of cerebral vessels in mechanical thrombectomy. This paper is currently under review. An anonymized version of this work has been submitted as supplementary material
4. Bendszus, M., Fiehler, J., Subtil, F., Bonekamp, S., Aamodt, A.H., Fuentes, B., et al.: Endovascular thrombectomy for acute ischaemic stroke with established large infarct: multicentre, open-label, randomised trial. *The Lancet* (11 2023). [https://doi.org/10.1016/S0140-6736\(23\)02032-9](https://doi.org/10.1016/S0140-6736(23)02032-9)
5. Crossley, R., Liebig, T., Holtmannspoetter, M., Lindkvist, J., Henn, P., Lonn, L., Gallagher, A.: Validation studies of virtual reality simulation performance metrics for mechanical thrombectomy in ischemic stroke. *Journal of NeuroInterventional Surgery* pp. 775–780 (2019). <https://doi.org/10.1136/neurintsurg-2018-014510>
6. Faure, F., Duriez, C., Delingette, H., Allard, J., Gilles, B., Marchesseau, S., et al.: SOFA, a Multi-Model Framework for Interactive Physical Simulation (2012). https://doi.org/https://doi.org/10.1007/8415_2012_125

7. Georgiev, I., Giridhar, V., Hansen, N., Garg, A.: Pwm: Policy learning with large world models (7 2024), <http://arxiv.org/abs/2407.02466>
8. Ha, D., Schmidhuber, J.: Recurrent world models facilitate policy evolution. In: Advances in Neural Information Processing Systems (2018), <https://worldmodels.github.io>
9. Hafner, D., Pasukonis, J., Ba, J., Lillicrap, T.: Mastering diverse domains through world models (1 2023), <http://arxiv.org/abs/2301.04104>
10. Hansen, N., Su, H., Wang, X.: Td-mpc2: Scalable, robust world models for continuous control. In: The Twelfth International Conference on Learning Representations (2024), <http://arxiv.org/abs/2310.16828>
11. Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., Meger, D.: Deep reinforcement learning that matters. In: Thirty-Second AAAI Conference on AI and Thirtieth Innovative Applications of AI Conference and Eighth AAAI Symposium on Educational Advances in AI (2 2018), www.aaai.org
12. Hu, A., Russell, L., Yeo, H., Murez, Z., Fedoseev, G., Kendall, A., Shotton, J., Corrado, G.: Gaia-1: A generative world model for autonomous driving (9 2023), <http://arxiv.org/abs/2309.17080>
13. Jianu, T., Huang, B., Vu, M.N., Abdelaziz, M.E., Fichera, S., Lee, C.Y., et al.: Cathsim: An open-source simulator for endovascular intervention. IEEE Transactions on Medical Robotics and Bionics **6**, 971–979 (2024). <https://doi.org/10.1109/TMRB.2024.3421256>
14. Karstensen, L., Ritter, J., Hatzl, J., Pätz, T., Langejürgen, J., Uhl, C., Mathis-Ullrich, F.: Learning-based autonomous vascular guidewire navigation without human demonstration in the venous system of a porcine liver. International Journal of Computer Assisted Radiology and Surgery **17**, 2033–2040 (11 2022). <https://doi.org/10.1007/s11548-022-02646-8>
15. Karstensen, L., Robertshaw, H., Hatzl, J., Jackson, B., Langejürgen, J., Breininger, K., et al.: Learning-based autonomous navigation, benchmark environments and simulation framework for endovascular interventions (10 2024), <http://arxiv.org/abs/2410.01956>
16. Klein, L.W., Miller, D.L., Balter, S., Laskey, W., Haines, D., Norbash, A., et al.: Occupational health hazards in the interventional laboratory: Time for a safer environment. Society of Interventional Radiology **250**, 538–544 (2 2009)
17. Madder, R.D., VanOosterhout, S., Mulder, A., Elmore, M., Campbell, J., Borgman, A., et al.: Impact of robotics and a suspended lead suit on physician radiation exposure during percutaneous coronary intervention. Cardiovascular Revascularization Medicine **18**, 190–196 (4 2017). <https://doi.org/10.1016/j.carrev.2016.12.011>
18. Martin, S.S., Aday, A.W., Almarzooq, Z.I., Anderson, C.A., Arora, P., Avery, C.L., et al.: 2024 heart disease and stroke statistics: A report of us and global data from the american heart association. Circulation **149**, E347–E913 (2 2024). <https://doi.org/10.1161/CIR.0000000000001209>
19. McMeekin, P., James, M., Price, C.I., Ford, G.A., White, P.: The impact of large core and late treatment trials: An update on the modelled annual thrombectomy eligibility of uk stroke patients. European Stroke Journal **9**, 566–574 (9 2024). <https://doi.org/10.1177/23969873241232820>
20. Ngankou, E.H., Gory, B., Marnat, G., Richard, S., Bourcier, R., Sibon, I., et al.: Thrombectomy complications in large vessel occlusions: Incidence, predictors, and clinical impact in the etis registry. Stroke **52**, E764–E768 (12 2021). <https://doi.org/10.1161/STROKEAHA.121.034865>

21. Nogueira, R.G., Jadhav, A.P., Haussen, D.C., Bonafe, A., Budzik, R.F., Bhuva, P., et al.: Thrombectomy 6 to 24 hours after stroke with a mismatch between deficit and infarct. *New England Journal of Medicine* **378**, 11–21 (1 2018). <https://doi.org/10.1056/nejmoa1706442>
22. O’Neill, A., Rehman, A., Maddukuri, A., Gupta, A., Padalkar, A., Lee, A., et al.: Open x-embodiment: Robotic learning datasets and rt-x models. In: *IEEE International Conference on Robotics and Automation*. pp. 6892–6903 (2024). <https://doi.org/10.1109/ICRA57147.2024.10611477>
23. Rivero-Moreno, Y., Rodriguez, M., Losada-Muñoz, P., Redden, S., Lopez-Lezama, S., Vidal-Gallardo, A., et al.: Autonomous robotic surgery: Has the future arrived? *Cureus* (1 2024). <https://doi.org/10.7759/cureus.52243>
24. Robertshaw, H., Karstensen, L., Jackson, B., Granados, A., Booth, T.C.: Autonomous navigation of catheters and guidewires in mechanical thrombectomy using inverse reinforcement learning. *International Journal of Computer Assisted Radiology and Surgery* (6 2024). <https://doi.org/10.1007/s11548-024-03208-w>
25. Robertshaw, H., Karstensen, L., Jackson, B., Sadati, H., Rhode, K., Ourselin, S., et al.: Artificial intelligence in the autonomous navigation of endovascular interventions: a systematic review. *Frontiers in Human Neuroscience* **17** (8 2023). <https://doi.org/10.3389/fnhum.2023.1239374>
26. SSNAP: Ssnap annual report 2023. Tech. rep., Sentinel Stroke National Audit Programme (2023), www.hqip.org.uk/national-programmes
27. Wei, Y., Cotin, S., Dequidt, J., Duriez, C., Allard, J., Kerrie, E.: A (Near) Real-Time Simulation Method of Aneurysm Coil Embolization (8 2012). <https://doi.org/10.5772/48635>
28. Weyland, C.S., Seker, F., Potreck, A., Hametner, C., Ringleb, P.A., Möhlenbruch, M.A., et al.: Radiation exposure per thrombectomy attempt in modern endovascular stroke treatment in the anterior circulation. *European Radiology* (4 2020). <https://doi.org/10.1007/s00330-020-06837-2>