CAG-WM: A Synthetic-Data-Driven Coronary World Model for Autonomous Guidewire Navigation

Yue Cao[1] [0009–0003–4482–8440], Zhuo Zhang[1, ✉] [0000–0002–3946–0720], Xiaoshan Peng[1] [0009–0002–2074–6374], Guipeng Lan[1] [0000–0001–7321–7460], Shuai Xiao[1] [0000–0003–4058–8120], and Jiachen Yang[1] [0000–0003–2558–552X]

Tianjin University, Tianjin, China
z_zhuo@tju.edu.cn

Abstract. Automated guidewire navigation in percutaneous coronary intervention can reduce radiation exposure and clinician workload. However, existing methods suffer from low sample efficiency and struggle with long-horizon guidewire navigation. World models offer a promising solution, but scarce paired angiographic-3D data leaves end-to-end world model approaches unexplored. To address these challenges, we propose CAG-WM, the first synthetic-data-driven world model for autonomous coronary guidewire navigation. Specifically, we construct dynamic synthetic coronary angiography–3D data and train a world model within a SOFA-based simulation environment to learn vessel–guidewire dynamics and perform internal multi-step rollouts for long-horizon navigation. CAG-WM is trained on 100 synthetic cases with full-view angiographic renderings and corresponding 3D voxelized coronary models, and evaluated on 20 multi-center clinical cases and 10 public datasets. Experimental results show that CAG-WM achieves a success rate exceeding 80% in clinical cases, significantly outperforming competing methods. Even in unsuccessful cases, the guidewire remains close to target lesion. Despite slightly longer per-step inference time, CAG-WM achieves the shortest total time (about 2 min) thanks to its higher success rate, which reduces premature termination at the maximum step limit. Code and data is available at https://github.com/CaoYue2004/CAG-WM.

Keywords: Coronary Angiography · Guidewire Navigation · World Model.

1 Introduction

Autonomous guidewire navigation holds significant clinical value in percutaneous coronary intervention (PCI) [2], which analyzes real-time angiographic images to generate optimal control commands, enabling automated guidewire advancement to stenotic lesions. It effectively reduces cumulative radiation exposure to physicians, mitigates steep learning curve of manual catheterization, and decreases inter-operator variability [15, 22]. Furthermore, this task represents a compelling frontier for embodied AI, where an agent must perceive angiographic scenes, reason about vascular anatomy, and execute precise guidewire control, spurring growing interest in AI-driven endovascular intervention [17, 21].

Recently, reinforcement learning-based (RL) [5, 14, 19, 20, 25] and planning-based approaches [14, 20] have been explored for autonomous coronary guidewire navigation, demonstrating promising results. However, these methods lack explicit modeling of guidewire-vessel dynamics, resulting in poor sample efficiency and limited generalization across diverse coronary anatomies. Without the ability to anticipate the consequences of actions, agents cannot reason beyond immediate observations, making it difficult to plan ahead or adapt to unseen anatomical variations. World models [8] have emerged as a promising framework to tackle these limitations. By learning environment dynamics, they enable agents to plan through imagined trajectories rather than costly trial-and-error interactions.

Nevertheless, world models remain rarely underexplored in vascular domain. Robert-shaw et al. [18] introduced the first world model for autonomous guidewire navigation in aortic arch. Despite this pioneering effort, several limitations remain. The study relies on limited angiography and corresponding 3D datasets, potentially restricting its broader clinical value, and does not explicitly model vascular dynamics. Furthermore, it lacks an end-to-end mapping from angiographic images to navigation actions, and the derivation of observational states, such as guidewire position, is not clearly specified. Consequently, the development of world models for coronary intervention remains largely open.

To address these challenges, we propose **C**oronary **A**n**g**iography **W**orld **M**odel (**CAG-WM**), the first synthetic-data-driven world model for autonomous coronary guidewire navigation. CAG-WM learns guidewire-vessel interaction dynamics from CAG observations and generates control actions to reach stenotic targets efficiently. To mitigate data scarcity, we construct a **diverse synthetic dataset** of full-view CAG paired with corresponding 3D voxelized models, incorporating dynamic vascular properties. We develop a **world model**-based framework specifically tailored for the continuous control demands of guidewire navigation, coupled with a SOFA-based [7] simulation environment that enables large-scale policy optimization and facilitates learning of environment dynamics under diverse anatomical variations. CAG-WM is trained on 100 synthetic cases and evaluated on 20 multi-center clinical cases and 10 public datasets [23]. Experimental results demonstrate a success rate exceeding 80% in real coronary navigation, significantly outperforming both RL baselines and existing world model methods. Even in unsuccessful cases, guidewire navigated by CAG-WM reaches significantly closer to target lesion compared to other methods. Although per-step inference time is slightly longer than RL methods, CAG-WM achieves the shortest total navigation time (about 2 min), owing to its higher success rate which reduces episodes terminated by reaching maximum step limit.

2 Method

We formulate coronary guidewire navigation as a sequential decision problem where the agent receives partial observations o_t (CAG images) and outputs

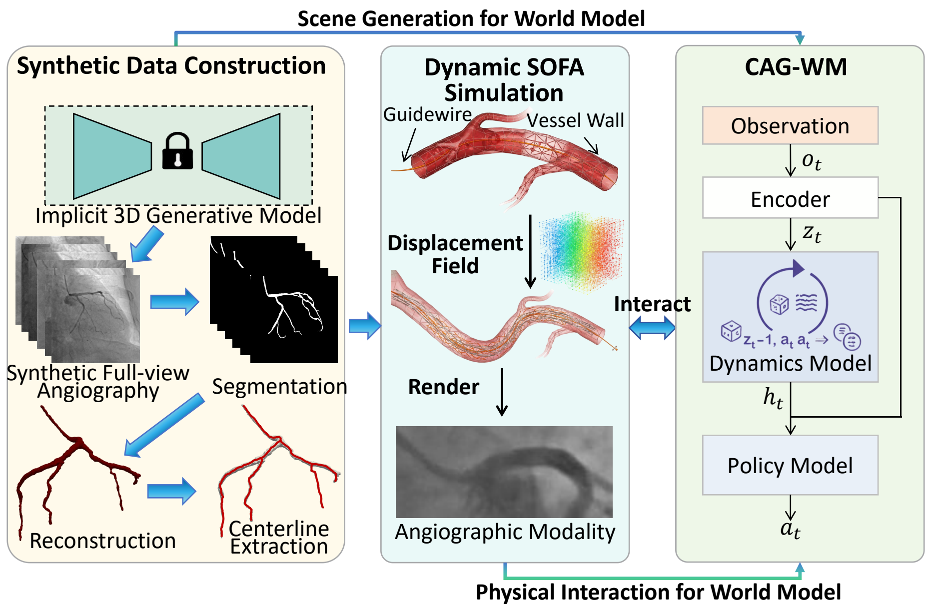

Fig. 1. Overview of the proposed CAG-WM framework.

guidewire control actions. Our objective is to learn a policy that navigates guidewire from its initial position to target while avoiding vessel wall contact and minimizing procedural time. To this end, we propose CAG-WM, a coronary world model as illustrated in Fig. 1: (1) synthetic data construction, (2) dynamic SOFA simulation, and (3) world model training for latent dynamics and optimal control inference. Each of these components is described in detail below.

2.1 Synthetic Data Construction

To address the scarcity of paired clinical data, we propose a synthetic framework for generating consistent coronary environments for world model training.

Given synthesized CAG images $\{I_i\}_{i=1}^N$ rendered from implicit 3D coronary models [4], we extract binary vessel masks $\{M_i\}_{i=1}^N$ via segmentation. Under known projection geometries, we apply multi-view geometric consistency constraints to reconstruct a coarse 3D voxelized volume $\mathcal{V} \subset \mathbb{R}^3$. A skeleton graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ is then extracted from $\mathcal{V}$ and globally connected using a minimum spanning tree. For each centerline node $n \in \mathcal{N}$, the vessel radius r_n is estimated by back-projecting n onto all views and averaging the distances to the corresponding vessel boundaries. The final 3D coronary model is defined as

$$\mathcal{M} = \bigcup_{n \in \mathcal{N}} B(n, r_n), \tag{1}$$

where $B(n, r_n)$ denotes a 3D ball centered at n with radius r_n. The resulting synthetic dataset provides full-view CAG, 3D coronary structures, and centerline representations, forming a comprehensive foundation for training CAG-WM.

2.2 Dynamic SOFA Simulation

The training environment for CAG-WM is based on a physics-driven simulation of guidewire-vessel interaction in the SOFA framework [7]. Guidewire is modeled as a deformable Timoshenko beam, while vessel wall is treated as a dynamically deforming boundary governed by a diffeomorphic displacement field that provides time-varying wall positions and velocities. The guidewire dynamics with contact constraints are solved via constraint-based formulation,

$$\begin{aligned} \left(\mathbf{M} + \Delta t\,\mathbf{D} + \Delta t^2\mathbf{K}\right)\Delta\mathbf{v} &= -\Delta t\left(\mathbf{f} + \Delta t\,\mathbf{K}\mathbf{v}\right) + \Delta t\,\mathbf{H}^\top\boldsymbol{\lambda}, \\ \boldsymbol{\lambda} &= \arg\min_{\lambda\geq 0}\left\|\mathbf{H}(\mathbf{v}+\Delta\mathbf{v}) + \frac{\mathbf{g}(t)}{\Delta t} - \mathbf{H}_w\mathbf{v}_w(t)\right\| \end{aligned} \tag{2}$$

where $\mathbf{M}$, $\mathbf{D}$, and $\mathbf{K}$ denote the beam mass, damping, and stiffness matrices, $\mathbf{H}$ encodes contact constraints, and $\boldsymbol{\lambda}$ are Lagrange multipliers. The bias terms $\mathbf{g}(t)$ and $\mathbf{v}_w(t)$ account for vessel wall deformation, and $\boldsymbol{\lambda}$ is solved using a Gauss–Seidel scheme for stable guidewire–vessel interaction.

The 3D coronary models from Sec. 2.2 are imported into SOFA as vessel walls, with guidewire modeled using real physical parameters. CAG images with the guidewire overlaid are rendered to provide clinically realistic observations.

2.3 CAG-WM

We develop CAG-WM, a TD-MPC2-based [11] coronary world model to learn vessel–guidewire dynamics for autonomous navigation. Fig. 2 illustrates our framework, where observations extracted from CAG images are encoded into a latent representation and processed by the world model to generate control actions, with a reward function guiding policy optimization.

Coronary-Specific Observation Modeling At step t, given a synthesized angiography image I_t, we randomly sample a target point $\mathbf{p}^*$ from the vessel centerline $\mathcal{C}$ [26], and obtain the guidewire position $\{\mathbf{g}_i\}_{i=1}^N$ from the SOFA simulation. The guidewire tip $\mathbf{g}_{\text{tip}}$ and its M nearest centerline points $\{\mathbf{c}_j\}_{j=1}^M$ are used to describe the local vessel structure for observation encoding.

Two TCN [1] encoders are used to encode $\{\mathbf{g}_i\}_{i=1}^N$ and $\{\mathbf{c}_j\}_{j=1}^M$, capturing spatial and temporal information. The resulting features are concatenated with $\mathbf{p}^*$ and fused by an MLP to form final multimodal observation representation.

Reward Design To promote reliable and efficient guidewire navigation, we design a task-specific reward. The reward at step t is defined as

$$R_t = \lambda_{\text{tar}} r_t^{\text{tar}} + \lambda_{\text{path}} r_t^{\text{path}} - \lambda_{\text{err}}\mathbf{1}[\text{error}_t] - \lambda_{\text{stuck}} r_t^{\text{stuck}}, \tag{3}$$

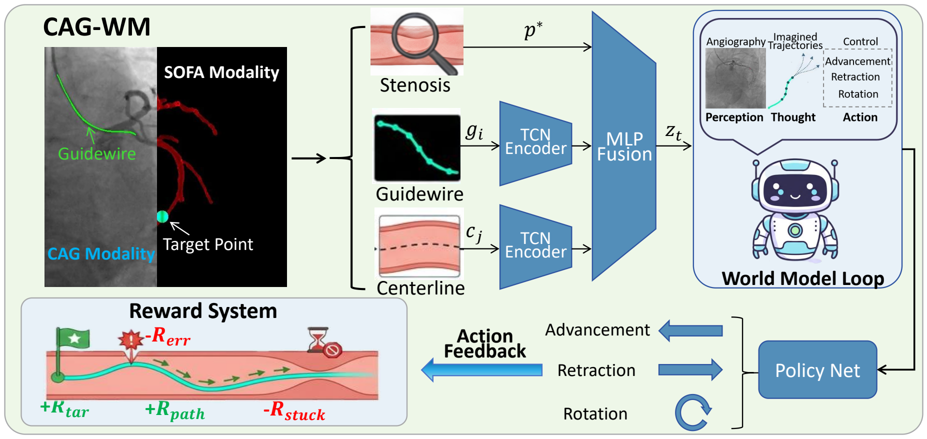

Fig. 2. Details of our CAG-WM. Observations from CAG images are encoded into a latent representation z_t via TCN, which is fed into the world model loop for imagination-based planning and action generation, guided by our reward system.

where r_t^{tar} rewards reaching the final target and intermediate waypoints along the planned path, r_t^{path} encourages forward progress by reducing the remaining path length, $\mathbf{1}[\text{error}_t]$ is an indicator of simulation failure (e.g., vessel wall damage), and r_t^{stuck} penalizes stagnation when the guidewire exhibits negligible displacement. The nonnegative weights $\lambda_{\text{tar}}, \lambda_{\text{path}}, \lambda_{\text{err}}, \lambda_{\text{stuck}}$ control the trade-off between efficiency, safety, and stability.

3 Experiments

3.1 Dataset and Metrics

Training Set The training set consists of 100 synthetic coronary cases, comprising CAG images, 3D vessel models, and centerlines, generated using a trained implicit 3D generative model [4, 3, 13]. Corresponding 3D vessel models and centerlines are reconstructed following Sec 2.2. All perceptual information required by observation is pre-extracted prior to training. Dataset is available at github.

Test Set The test set comprises 20 real clinical cases collected from five centers (InCAG) and 10 cases from public dataset, VMR [23]. For each coronary centerline, 15 target points are sampled to serve as navigation goals. For each target point, 10 repeated trials are conducted to evaluate the robustness of the methods, and the results are reported as mean ± standard deviation across trials.

Evaluation Metrics To comprehensively evaluate the performance, efficiency, and stability of CAG-WM, we design several evaluation metrics. Success Rate

(SR) is defined as the percentage of episodes that successfully reach the target. Relative Distance (RD) measures the proportion of remaining centerline distance to the target at episode termination, normalized by the total initial centerline distance from the starting point to the target, evaluated only on failed episodes. Time (T) represents the duration required to navigate from the starting point to the target. Statistical comparisons are conducted using two-tailed paired Student's t-tests, with the significance level set to $p < 0.05$.

3.2 Implementation Details

In data preparation stage, the implicit 3D generative model is trained on 30k real CAG images, achieving an FID score of 8.91. The segmentation network UNet++ [27] are trained on ARCADE [16] and achieves a Dice score of 0.761. During the training of CAG-WM, experiments are conducted on a single NVIDIA L20 GPU. Each model is trained for 10^7 steps with a learning rate of 3×10^{-4}, and the maximum number of steps per episode is set to 400. In Eq. (3), the coefficients are set as $\lambda_{\text{tar}} = 5.0$, $\lambda_{\text{path}} = 0.1$, $\lambda_{\text{err}} = 1.0$, and $\lambda_{\text{stuck}} = 0.1$.

4 Results

4.1 Comparative Experiments

To validate the superiority of CAG-WM, we conduct comprehensive comparative experiments. The baseline methods are categorized into three groups: (1) model-free RL methods, including SAC [9] and TD3 [24]; (2) planning-based methods that directly use simulator dynamics, namely CEM-MPC [6]; (3) world-model-based methods, including DreamerV3 [10], TD-MPC [12], and TD-MPC2 [11].

As shown in Table 1, CAG-WM achieves the best SR and RD on both the clinical and public datasets, reaching 84.7%/17.6% and 83.1%/18.9% respectively, significantly outperforming model-free RL methods and the planning-based CEM-MPC ($p < 0.001$). Among world model baselines, CAG-WM also demonstrates clear advantages. DreamerV3 is better suited for raw image inputs rather than extracted positional features, which limits its performance in our setting. TD-MPC struggles with continuous action spaces compared to our method. Against our backbone TD-MPC2, CAG-WM achieves approximately 10% and 8% improvements in SR and RD respectively, validating the effectiveness of our encoder modifications. Regarding the time metric, although our per-step inference time is slower than SAC and TD3, the higher SR reduces episodes terminated by reaching the maximum step limit, resulting in the shortest overall navigation time and highlighting the efficiency of our approach.

Figure 3.A illustrates navigation performance across three clinical casesin both CAG and SOFA modalities, where the guidewire successfully selects correct vessel branch and traverses tortuous structures, demonstrating robust comprehension of complex coronary anatomy. Figure 3.B shows guidewire tip trajectory heatmaps alongside corresponding 3D coronary voxels, revealing comprehensive

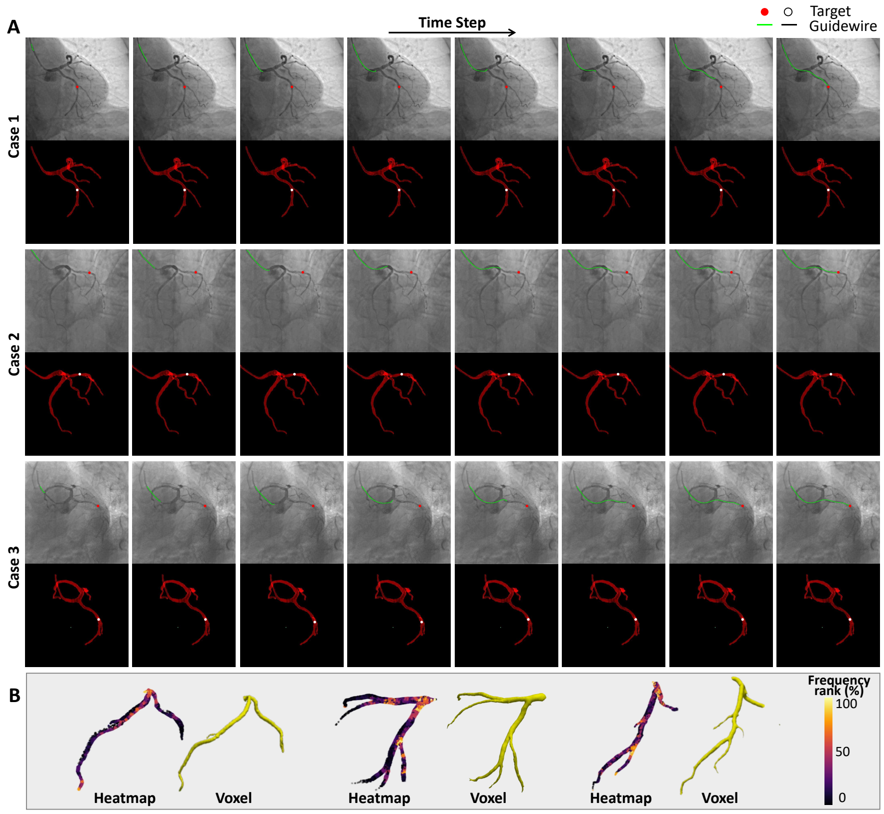

Fig. 3. A. Visualization of autonomous guidewire navigation in coronary arteries by CAG-WM. For each case, the first row shows the CAG modality (green line: guidewire; red dot: target) and the second row shows the SOFA modality (black line: guidewire; white dot: target). B. Heatmap visualization of guidewire tip positions during training and corresponding 3D coronary voxel models.

exploration across entire vascular tree and confirming the guidewire's capacity to access all reachable vessel positions. More experimental videos are provided in the supplementary materials.

4.2 Ablation Study

To justify the design of our reward function, we conduct an ablation study on individual reward terms. As shown in Table 2, the full reward achieves the highest SR on both datasets, confirming that each component plays a complementary role in stabilizing learning and improving navigation performance. Without R_{s}, the guidewire frequently hesitates at vascular bifurcations due to the absence of stagnation penalty, leading to redundant movements and failed navigation. Without R_{p}, the agent lacks dense centerline guidance throughout the naviga-

Table 1. Quantitative comparison of CAG-WM with competing methods.

Method	InCAG			VMR		
	SR (%) ↑	RD (%) ↓	T (s) ↓	SR (%) ↑	RD (%) ↓	T (s) ↓
SAC [9]	$47.3_{\pm 17.1}$	$49.5_{\pm 21.9}$	$256.67_{\pm 24.34}$	$46.0_{\pm 18.3}$	$51.2_{\pm 20.4}$	$261.43_{\pm 26.17}$
TD3 [24]	$50.7_{\pm 15.4}$	$42.1_{\pm 23.8}$	$239.13_{\pm 23.48}$	$49.3_{\pm 16.2}$	$43.7_{\pm 22.5}$	$244.82_{\pm 21.93}$
CEM-MPC [6]	$52.7_{\pm 15.9}$	$40.8_{\pm 16.5}$	$279.63_{\pm 29.23}$	$52.0_{\pm 14.7}$	$41.9_{\pm 0.178}$	$283.17_{\pm 31.45}$
DreamerV3 [10]	$61.3_{\pm 12.8}$	$37.6_{\pm 17.7}$	$186.79_{\pm 20.30}$	$59.3_{\pm 13.5}$	$38.9_{\pm 16.8}$	$191.24_{\pm 22.56}$
TD-MPC [12]	$43.0_{\pm 16.4}$	$53.2_{\pm 27.1}$	$267.13_{\pm 23.10}$	$42.7_{\pm 17.1}$	$54.8_{\pm 25.9}$	$271.36_{\pm 25.83}$
TD-MPC2 [11]	$75.3_{\pm 8.3}$	$25.1_{\pm 10.4}$	$156.47_{\pm 17.66}$	$74.0_{\pm 9.1}$	$26.3_{\pm 11.2}$	$160.83_{\pm 19.24}$
CAG-WM	$\mathbf{84.7}_{\pm 7.7}$	$\mathbf{17.6}_{\pm 16.7}$	$\mathbf{127.64}_{\pm 14.34}$	$\mathbf{83.3}_{\pm 8.2}$	$\mathbf{18.9}_{\pm 15.4}$	$\mathbf{132.57}_{\pm 15.89}$

Table 2. Ablation study on reward components, where R_t: reward for reaching target point; R_p: reward for reducing centerline path distance to the target; R_e: penalty for mechanical interaction with vessel wall; R_s: penalty for guidewire being stuck.

R_t	R_p	R_e	R_s	InCAG			VMR		
				SR (%) ↑	RD (%) ↓	T (s) ↓	SR (%) ↑	RD (%) ↓	T (s) ↓
✓				$58.7_{\pm 14.2}$	$43.1_{\pm 19.8}$	$103.71_{\pm 18.23}$	$57.3_{\pm 15.6}$	$44.7_{\pm 21.1}$	$107.34_{\pm 19.87}$
✓	✓			$60.3_{\pm 13.1}$	$41.2_{\pm 18.3}$	$\mathbf{99.86}_{\pm 17.54}$	$58.7_{\pm 14.3}$	$42.8_{\pm 19.6}$	$\mathbf{103.52}_{\pm 18.63}$
✓	✓	✓		$62.9_{\pm 10.3}$	$39.6_{\pm 15.4}$	$129.33_{\pm 16.87}$	$61.3_{\pm 12.1}$	$41.5_{\pm 17.6}$	$133.17_{\pm 17.92}$
✓	✓	✓	✓	$\mathbf{84.7}_{\pm 7.7}$	$\mathbf{17.6}_{\pm 16.7}$	$127.64_{\pm 14.34}$	$\mathbf{83.3}_{\pm 0.082}$	$\mathbf{18.9}_{\pm 15.4}$	$132.57_{\pm 15.89}$

tion process, and relying solely on sparse target rewards proves insufficient for effective long-horizon navigation.

Although our full method exhibits longer navigation time, this is attributed to R_{e} enforcing mechanical constraints that prevent unrealistic vessel wall penetration, avoiding premature episode termination and ensuring physically valid trajectories that better reflect real clinical conditions.

5 Conclusion

In this study, we propose CAG-WM, the first coronary world model enabling end-to-end autonomous guidewire navigation from angiographic observations to control actions, providing a new pathway to reduce operator radiation exposure and workload. We construct a dynamic full-view coronary angiography–3D synthetic dataset and integrate a world model-based framework with a SOFA-based simulation environment to model vessel–guidewire interactions and optimize long-horizon decision-making. Validation on multi-center clinical cases and public datasets demonstrates that CAG-WM significantly outperforms RL and existing world model approaches in navigation success rate. Meanwhile, even in failed cases, our method achieves a closer proximity to the target point than competing methods. Furthermore, CAG-WM achieves the shortest navigation time, as its higher success rate reduces episodes terminated by the maximum step limit. This work represents the first systematic application of world modeling to dynamic coronary autonomous navigation and lays the foundation for AI-driven robotic coronary intervention. Future work will expand validation cohorts and advance evaluation on physical vascular phantoms and in vivo experiments. Beyond coronary guidewire navigation, the world model framework holds potential for broader PIC tasks, including balloon dilation and stent deployment, as well as pre-operative planning and intraoperative decision support.

Disclosure of Interests. The authors have no competing interests in the paper.

References

1. Bai, S., Kolter, J.Z., Koltun, V.: An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv:1803.01271 (2018)
2. Bhatt, D.L.: Percutaneous coronary intervention in 2018. JAMA **319**(20), 2127–2128 (05 2018). https://doi.org/10.1001/jama.2018.5281, https://doi.org/10.1001/jama.2018.5281
3. Cao, Y., Zhang, Z., Xiao, S., Li, J., Lan, G., Wen, J., Yang, J.: Automatic translational correction of multi-view coronary angiography based on auto-annotation data generation. Proceedings of the AAAI Conference on Artificial Intelligence **40**(4), 2652–2660 (Mar 2026). https://doi.org/10.1609/aaai.v40i4.37253, https://ojs.aaai.org/index.php/AAAI/article/view/37253
4. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., Mello, S.D., Gallo, O., Guibas, L., Tremblay, J., Khamis, S., Karras, T., Wetzstein, G.: Efficient geometry-aware 3D generative adversarial networks. In: arXiv (2021)
5. Cho, Y., Park, J.H., Choi, J., Chang, D.E.: Sim-to-real transfer of image-based autonomous guidewire navigation trained by deep deterministic policy gradient with behavior cloning for fast learning. In: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 3468–3475 (2022)
6. Chua, K., Calandra, R., McAllister, R., Levine, S.: Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In: Proceedings of the 32nd International Conference on Neural Information Processing Systems. p. 4759–4770. NIPS'18, Curran Associates Inc., Red Hook, NY, USA (2018)
7. Faure, F., Duriez, C., Delingette, H., Allard, J., Gilles, B., Marchesseau, S., Talbot, H., Courtecuisse, H., Bousquet, G., Peterlik, I., Cotin, S.: SOFA: A Multi-Model Framework for Interactive Physical Simulation, pp. 283–321. Springer Berlin Heidelberg, Berlin, Heidelberg (2012). https://doi.org/10.1007/8415_2012_125
8. Ha, D., Schmidhuber, J.: Recurrent world models facilitate policy evolution. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 31. Curran Associates, Inc. (2018)
9. Haarnoja, T., Zhou, A., Abbeel, P., Levine, S.: Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. arXiv e-prints arXiv:1801.01290 (Jan 2018). https://doi.org/10.48550/arXiv.1801.01290
10. Hafner, D., Pasukonis, J., Ba, J., Lillicrap, T.: Mastering diverse control tasks through world models. Nature pp. 1–7 (2025)
11. Hansen, N., Su, H., Wang, X.: Td-mpc2: Scalable, robust world models for continuous control (2024)
12. Hansen, N., Wang, X., Su, H.: Temporal difference learning for model predictive control (2022)

13. Li, J., Zhang, Z., Cao, Y., Lan, G., Wen, J., Xiao, S., Yang, J.: Geometrically constrained stenosis editing in coronary angiography via entropic optimal transport. arXiv preprint arXiv:2605.08851 (2026)
14. Li, Z., Lambranzi, C., Wu, D., Segato, A., De Marco, F., Poorten, E.V., Dankelman, J., De Momi, E.: Robust path planning via learning from demonstrations for robotic catheters in deformable environments. IEEE Transactions on Biomedical Engineering **72**(1), 324–336 (2025). https://doi.org/10.1109/TBME.2024.3452034
15. Madder, R.D., VanOosterhout, S., Mulder, A., Elmore, M., Campbell, J., Borgman, A., Parker, J., Wohns, D.: Impact of robotics and a suspended lead suit on physician radiation exposure during percutaneous coronary intervention. Cardiovascular Revascularization Medicine **18**(3), 190–196 (2017). https://doi.org/https://doi.org/10.1016/j.carrev.2016.12.011
16. Popov, M., Amanturdieva, A., Zhaksylyk, N., Alkanov, A., Saniyazbekov, A., Aimyshev, T., Ismailov, E., Bulegenov, A., Kuzhukeyev, A., Kulanbayeva, A., Kalzhanov, A., Temenov, N., Kolesnikov, A., Sakhov, O., Fazli, S.: Dataset for automatic region-based coronary artery disease diagnostics using x-ray angiography images. Scientific Data **11**(1), 20 (Jan 2024). https://doi.org/10.1038/s41597-023-02871-z, https://doi.org/10.1038/s41597-023-02871-z
17. Robertshaw, H., Karstensen, L., Jackson, B., Sadati, S., Rhode, K., Ourselin, S., Granados, A., Booth, T.: Artificial intelligence in the autonomous navigation of endovascular interventions: a systematic review. Frontiers in Human Neuroscience **17** (08 2023). https://doi.org/10.3389/fnhum.2023.1239374
18. Robertshaw, H., Wu, H.R., Granados, A., Booth, T.C.: World model for ai autonomous navigation in mechanical thrombectomy. In: Gee, J.C., Alexander, D.C., Hong, J., Iglesias, J.E., Sudre, C.H., Venkataraman, A., Golland, P., Kim, J.H., Park, J. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2025. pp. 680–690. Springer Nature Switzerland, Cham (2026)
19. Scarponi, V., Lecomte, F., Duprez, M., Nageotte, F., Cotin, S.: Autonomous guidewire navigation in dynamic environments. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 6931–6938 (2024)
20. Schegg, P., Dequidt, J., Coevoet, E., Leurent, E., Sabatier, R., Preux, P., Duriez, C.: Automated planning for robotic guidewire navigation in the coronary arteries. In: 2022 IEEE 5th International Conference on Soft Robotics (RoboSoft). pp. 239–246 (2022). https://doi.org/10.1109/RoboSoft54090.2022.9762096
21. Weisz, G., Metzger, D.C., Caputo, R.P., Delgado, J.A., Marshall, J.J., Vetrovec, G.W., Reisman, M., Waksman, R., Granada, J.F., Novack, V., Moses, J.W., Carrozza, J.P.: Safety and feasibility of robotic percutaneous coronary intervention: Precise (percutaneous robotically-enhanced coronary intervention) study. Journal of the American College of Cardiology **61**(15), 1596–1600 (2013). https://doi.org/https://doi.org/10.1016/j.jacc.2012.12.045
22. Weyland, C.S., Şeker, F., Potreck, A., Hametner, C., Ringleb, P.A., Möhlenbruch, M.A., Bendszus, M., Pfaff, J.A.R.: Radiation exposure per thrombectomy attempt in modern endovascular stroke treatment in the anterior circulation. European Radiology **30**, 5039 – 5047 (2020)
23. Wilson, N., Ortiz, A., Johnson, A.: The vascular model repository: A public resource of medical imaging data and blood flow simulation results. Journal of medical devices **7**, 0409231–409231 (12 2013). https://doi.org/10.1115/1.4025983
24. Wu, J., Wu, Q.M.J., Chen, S., Pourpanah, F., Huang, D.: A-td3: An adaptive asynchronous twin delayed deep deterministic for continuous action spaces. IEEE Access **10**, 128077–128089 (2022). https://doi.org/10.1109/ACCESS.2022.3226446

25. Yoo, S., Kwon, H., Choi, J., Chang, D.E.: A robust deep reinforcement learning framework for image-based autonomous guidewire navigation. In: 2025 IEEE International Conference on Robotics and Automation (ICRA). pp. 468–474 (2025). https://doi.org/10.1109/ICRA55743.2025.11127741
26. Zhao, Z., Zhang, Z., Liu, Y., Zhang, Z., Yu, H., Wang, D., Wang, L.: Deformcl: Learning deformable centerline representation for vessel extraction in 3d medical image. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 30896–30905 (2025)
27. Zhou, Z., Siddiquee, M.M.R., Tajbakhsh, N., Liang, J.: Unet++: Redesigning skip connections to exploit multiscale features in image segmentation. IEEE Transactions on Medical Imaging (2019)