ORIGINAL ARTICLE

Learning-based autonomous vascular guidewire navigation without human demonstration in the venous system of a porcine liver

Lennart Karstensen[1,4] · Jacqueline Ritter[1] · Johannes Hatzl[2] · Torben Pätz[3] · Jens Langejürgen[1] · Christian Uhl[2] · Franziska Mathis-Ullrich[4]

Abstract

Purpose The navigation of endovascular guidewires is a dexterous task where physicians and patients can benefit from automation. Machine learning-based controllers are promising to help master this task. However, human-generated training data are scarce and resource-intensive to generate. We investigate if a neural network-based controller trained without human-generated data can learn human-like behaviors.

Methods We trained and evaluated a neural network-based controller via deep reinforcement learning in a finite element simulation to navigate the venous system of a porcine liver without human-generated data. The behavior is compared to manual expert navigation, and real-world transferability is evaluated.

Results The controller achieves a success rate of 100% in simulation. The controller applies a wiggling behavior, where the guidewire tip is continuously rotated alternately clockwise and counterclockwise like the human expert applies. In the ex vivo porcine liver, the success rate drops to 30%, because either the wrong branch is probed, or the guidewire becomes entangled.

Conclusion In this work, we prove that a learning-based controller is capable of learning human-like guidewire navigation behavior without human-generated data, therefore, mitigating the requirement to produce resource-intensive human-generated training data. Limitations are the restriction to one vessel geometry, the neglected safeness of navigation, and the reduced transferability to the real world.

Keywords Endovascular intervention · Guidewire navigation · Deep reinforcement learning · Autonomous · Learning from scratch

✉ Lennart Karstensen
Lennart.karstensen@ipa.fraunhofer.de

Torben Pätz
Torben.Paetz@mevis.fraunhofer.de

Franziska Mathis-Ullrich
Franziska.ullrich@kit.edu

1 Fraunhofer IPA, Theodor-Kutzer-Ufer 1-3, 68167 Mannheim, Germany

2 Department of Vascular and Endovascular Surgery, University Hospital Heidelberg, Im Neuenheimer Feld 420, 69120 Heidelberg, Germany

3 Fraunhofer MEVIS, Max-von-Laue-Str. 2, 28359 Bremen, Germany

4 Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Engler-Bunte-Ring 8, 76131 Karlsruhe, Germany

Introduction

On a global scale, the main cause of death is cardiovascular diseases (31.5%), in particular, ischemic heart (14.8%) and cerebrovascular (11.7%) diseases [1]. Endovascular, catheter-based interventions are the gold standard for treatment of numerous cardiovascular diseases [2]. In these interventions, a combination of a guidewire and catheter is navigated from an insertion point to the region of interest. The devices are inserted into the artery, typically in the groin or at the wrist, and manually navigated by twisting and pushing under image guidance with intraoperative fluoroscopy. The angled tip of the guidewire allows probing the desired branch at bifurcations. Typically, a bifurcation is mastered by rotating the guidewire, such that the angled tip enters the targeted vessel, and the remainder of the guidewire follows. When the guidewire sits firmly in the targeted vessel, the catheter is pushed over it. Additionally, the guidewire tip geometry

can be modified by altering the overlap of guidewire and catheter. When advancing the guidewire through a vessel certain movement pattern can ease the navigation process, e.g., constantly wiggling the guidewire clockwise and counterclockwise can prevent the guidewire from entangling if the vessel meanders or smaller side vessels are present. After a successful navigation, the subsequent treatment through the catheter may be performed.

Improvements in treatment quality and speed can have a high impact on patients' lives. Automation of catheter navigation has the potential to reduce radiation exposure of physicians when they must no longer stand near the patient. Even though radiation exposure can be reduced with measures of best practice, further reduction of radiation exposure is a sensible goal for the health of physicians [3, 4]. Furthermore, rural areas have reduced access to endovascular treatments compared to urban areas. A reliable autonomous system takes over the basic surgery and dexterity tasks, and the physician can focus on diagnosis and decision making, thus further enabling telesurgery and mitigating the supply gap [5, 6].

Endovascular robotics already transitioned from research to the industry in the form of telemanipulated manipulators, e.g., with the discontinued Magellan System (Hansen Medical, USA) or the currently available Corpath System (Corindus, USA) [7]. Among others, the main advantages are reduced radiation for the physicians [8] or telemedicine without direct patient visibility [9]. The Corpath System already incorporates automation of some movement patterns utilized by physicians in manual operations, e.g., the "wiggle" pattern [10], but no fully automated navigation is possible.

Current efforts to automate catheter and guidewire navigation can be divided into three categories: magnetically actuated, active, and passive catheters.

Control of a magnetic catheter tip in a closed-loop approach for eye surgery has been presented in [11]. The same principle is adapted to endovascular catheters in [12]. A recent development regarding autonomous magnetically actuated catheter navigation is the Advanced Robotics for Magnetic Manipulation System [13] that successfully positioned a magnetic catheter in a gelatinous phantom of a human torso with an embedded aorta using ultrasound images as feedback.

Fagogenis et al. [14] demonstrate autonomous navigation of a catheter with an active tip inside an in-vivo beating heart. Research regarding steerable catheters improves the feasibility and use cases of these devices, e.g., Gopesh et al. [15] fabricated and animal tested a steerable microcatheter for the endovascular treatment of cerebral disorders.

Passive catheters and guidewires are used much more frequently than the active or magnetic ones because there are a larger number of variants, they are less expensive and can be significantly smaller due to their simpler design. Furthermore, utilizing commercially available passive catheters enables faster clinical applicability as the invasive devices are already certified. Rafii-Tari et al. [16] propose an approach to learn optimum motion trajectories from expert demonstration and present their system in a phantom using electromagnetic tracking as feedback. Tercero et al. [17] demonstrate a system that navigates a catheter by following a pre-computed software map of movement commands. This software map may be adapted to a patient based on pre-interventional images.

Artificial intelligence embedded into robotics exhibits a high potential to improve the quality of surgical procedures. Diseases may be treated either through surgeons cooperating with intelligent robots or, in the future, by fully or semi-autonomous robotic systems [6, 18]. Therefore, approaches for automated navigation of passive guidewires or catheters emerge, which utilize deep learning. Chi et al. [19] investigate learning catheterization of the brachycephalic artery using expert demonstration with a generative adversarial imitation learning agent. Subsequently, a proximal policy optimization agent learns to catheterize the left common carotid artery. Both agents are trained in a type-1 aortic arch and evaluated in the type-1 and a type-2 aortic arch phantom using electromagnetic tracking as feedback. Zhao et al. [20] propose a convolutional neural network to estimate suitable manipulation actions. Additionally, an operating force mode estimator is trained to ensure safe operation. The estimated most suitable action is only executed if the operating force is normal. If the force becomes abnormal, predefined avoiding actions are executed. The neural networks are trained using demonstration data from expert surgeons. The controller is evaluated in a phantom with a gray-scale camera simulating X-ray imaging.

The combination of endovascular robotics and artificial intelligence has the potential to greatly advance autonomous navigation of endovascular catheters. The aforementioned approaches utilize human-generated data. However, availability of large-scale, high-quality, and correctly labeled data is limited and resource-intensive to generate. A controller trained from scratch without the necessity of data created by humans would circumvent the need for resource-intense data generation.

This work investigates the training of a neural network-based controller for endovascular navigation without human-generated data. The controller design builds upon prior work [21]. We investigate whether the controller can learn behaviors comparable to manual navigation. For this purpose, the controller learns the autonomous navigation of a guidewire from the insertion point to arbitrary targets within a vessel system. Exemplary the venous system of a porcine liver is utilized, as the porcine and human vessel system is similar, and pigs are often used for endovascular animal trials. The controller is trained and evaluated in a finite element
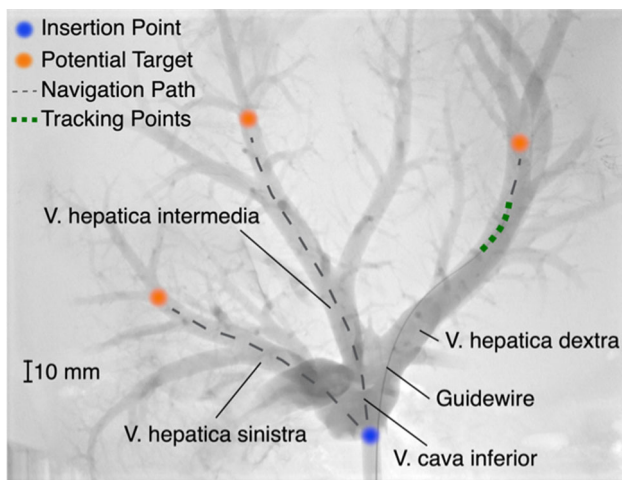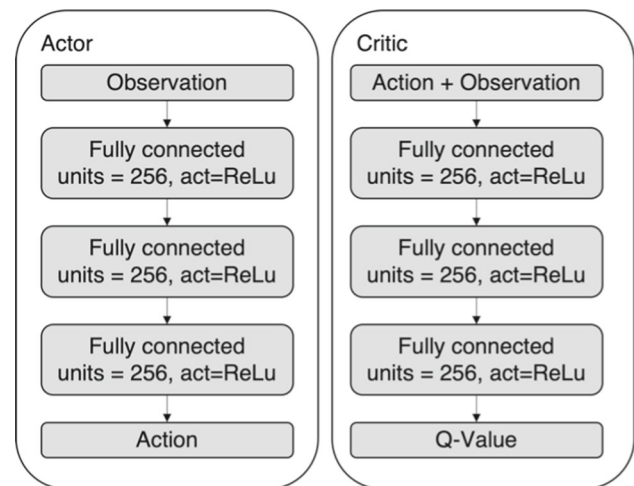
Fig. 1 X-Ray of an ex-vivo porcine liver with contrast agent, a guidewire inserted in the vena hepatica dextra, possible target points including their navigation paths and the tracking points used as feedback

Fig. 2 The structure of the actor and critic neural networks of the controller

simulation of one ex-vivo porcine liver specimen. For qualitative comparison, the same navigation is performed by a medical expert in a phantom modeled after the same specimen. The resulting successfully trained controller shows that the wiggle movement pattern has been learned as an optimal behavior without previous human demonstration. The main scope of this work is the investigation of the learned behavior patterns and maneuvers of the controller. The transferability of the simulation-trained controller to the real world is evaluated by performing autonomous navigation in the same ex vivo porcine liver specimen.

Methods and materials

Endovascular navigation task

This work aims to autonomously navigate an endovascular guidewire through the veins of a porcine liver. The same specimen is utilized for in silico training and evaluation, manual phantom navigation, and ex vivo transferability testing. The guidewire is inserted in the vena cava inferior and enters the liver coming from the direction of the heart. The anatomy and nomenclature of the liver veins, the insertion point and possible targets with their navigation paths are illustrated in Fig. 1. The vena cava inferior separates into three main liver veins, i.e., vena hepatica dextra, vena hepatica intermedia, and vena hepatica sinistra. All three main hepatic veins bifurcate in several large and small side branches depending on the individual's specific geometry. Here, we consider all larger branches, which are reachable by manual navigation. Smaller side branches are neglected due to mechanical limitations of the guidewire dexterity.

The guidewire (Terumo Radifocus Guide Wire M, angled, Stiff Type, 0.035″ Diameter, 3 cm flexible tip) is manipulated at its base outside the liver. During the navigation procedure, the guidewire position is used as feedback in the form of discrete tracking points, as illustrated in Fig. 1.

A navigation task starts at the vena cava inferior shortly before the trifurcation to the hepatic veins. The navigation target is a pre-selected point on the centerline of any branch. A navigation task is considered successful if the guidewire tip reaches the selected target within 40 s. In a manual preliminary experiment, 20 s was determined to be sufficient time for a navigation task which was doubled to allow exploration during training. The target is considered reached if the distance between the guidewire tip and the target is less than 15 mm.

The control task is characterized by a partially observable Markov decision process [22]. The Markov property requires that a single input state can describe the complete state of the given problem. However, the guidewire control is only partially observable as the two-dimensional X-ray images used as feedback during navigation provide the position of the guidewire in the image plane. The depth is only estimated from knowledge of the vessel anatomy. Information about the absolute rotation of the guidewire tip is not perceivable from the image as two images of the guidewire tip rotated at a specific angle upwards or downwards are ambiguous.

Controller architecture and training procedure

The controller is a neural network, trained using Deep Deterministic Policy Gradients [23] with Hindsight Experience Replay [24]. Actor and a critic both have three fully connected layers of 256 Neurons, as illustrated in Fig. 2. The observation is defined as the current and the last two

guidewire positions, the target position and the last two actions leading to the current state, thus creating a memory to make the rotation of the guidewire tip observable. By observing the movement direction of the tip, it can be estimated if it is turned upwards or downwards. The guidewire position is given as the tracking coordinates of five points on the guidewire described as $(x, z)_i$ $i = \{1, 2, 3, 4, 5\}$ relative to the insertion point where all points are spaced evenly and 1 mm apart and $(x, z)_1$ is coincident with the guidewire tip. The target position is given as the target coordinates (x, z) of the current target. The action is given as the guidewire rotation and translation speed.

For all coordinate-based inputs, only the x and z coordinates are considered, representing the image plane. They are normalized between -1 and 1 with $(0, 0)$ indicating the geometric center of the porcine liver and -1 and 1 being the lowest and largest coordinate on the liver, respectively. The action output of the actor network is the guidewire rotation and translation speed. Rotation and translation speeds are continuous values between -6.28 rad/s to 6.28 rad/s and -10 mm/s to 10 mm/s, respectively.

Hindsight Experience Replay provides best results with sparse rewards, reducing the required task-specific knowledge to set up the reinforcement learning environment, i.e., it is only necessary to detect whether the target is reached. Thus, the agent receives -1 reward for every step and an additional $+1$ reward for reaching the target. One step is defined as applying an action to the environment, receiving the observation, and subsequently calculating the next action. The episode is considered finished if the target is reached or 40 s have passed. One episode is defined as performing one navigation task. The control frequency is 7.5 Hz. For training and evaluation, targets are selected randomly on the branch centerlines and an episode always starts at the insertion point. The necessary insertion depth, depending on the distance between insertion point and target, has no influence on the maximum duration.

Every 750 training episodes, the performance is measured for 150 consecutive evaluation episodes. Performance is defined as the percentage of evaluation episodes where the controller successfully reaches the target. During training episodes, stochastic noise is added to the action calculation for exploration, but not during evaluation episodes. Training and evaluation are performed with 15 digital agents in parallel. Each agent periodically performs 50 training and 10 evaluation episodes.

Training and evaluation environment

Training of the controller is solely performed in the simulation environment. The controller is evaluated in the same simulation environment. To examine real world transferability, the controller is additionally tested on a testbench in the ex vivo porcine liver used as basis for the simulation.

The simulation is modeled within the SOFA framework [25] using the BeamAdapter plugin [26]. Phantom walls are assumed rigid and the lumen empty. Friction between wall and guidewire as well as guidewire stiffness has been iteratively tuned to mimic guidewire behavior in the testbench. The simulation receives guidewire rotation and translation speed as input and the output is the guidewire position as coordinate points evenly spaced at 1 mm starting from the guidewire tip. The tracking position is received with a delay, which can be described by a normal distribution with a mean delay of 0.13 s and variance of 0.01 s. This mimics the processing delay of the testbench tracking system. A rendering of the simulation model can be seen in Fig. 3a.

To transfer the geometry of the porcine liver into the simulation environment, a CT scan (Artis Zeego, Siemens Healthineers) with contrast agent (potassium iodide) is performed. The veins are segmented in the CT image semiautonomously [27]. The resulting segmentation is smoothed and transformed into a surface mesh.

The phantom testbench can be seen in Fig. 3b. The phantom consists of a rigid structure with lumens in the shape of the extracted vessel geometry, manufactured by stereolithography. A medical expert manually navigates the guidewire using continuous fluoroscopy (Artis Pheno, Siemens Healthineers) as feedback for comparison to the autonomous controller.

The ex vivo testbench is presented in Fig. 3c. The porcine liver is obtained from a local abattoir and prepared for guidewire navigation. A custom tube connector attached to the vena cava inferior allows guidewire insertion and the attachment of the constant flow pump. All other openings are closed by stitching or gluing.

The liver is placed inside a water-filled specimen container. A constant flow pumps water into the vena cava inferior to ensure that the veins are inflated. A customized guidewire manipulator translates and rotates the guidewire at its base. It allows for continuous rotation and a translation of up to 300 mm. The guidewire manipulator receives control commands by the trained controllers.

During the navigation task, continuous fluoroscopy (Artis Zeego, Siemens Healthineers) with a framerate of 7.5 images per second allows tracking of the guidewire. A tracking system evaluates each fluoroscopy image to extract the guidewire position. A recurrent convolutional neural network is utilized to segment the guidewire from the real-time fluoroscopy images, and a backprojection reconstructs the 3D shape of the guidewire. The output of the tracking system is the guidewire position as coordinate points evenly spaced
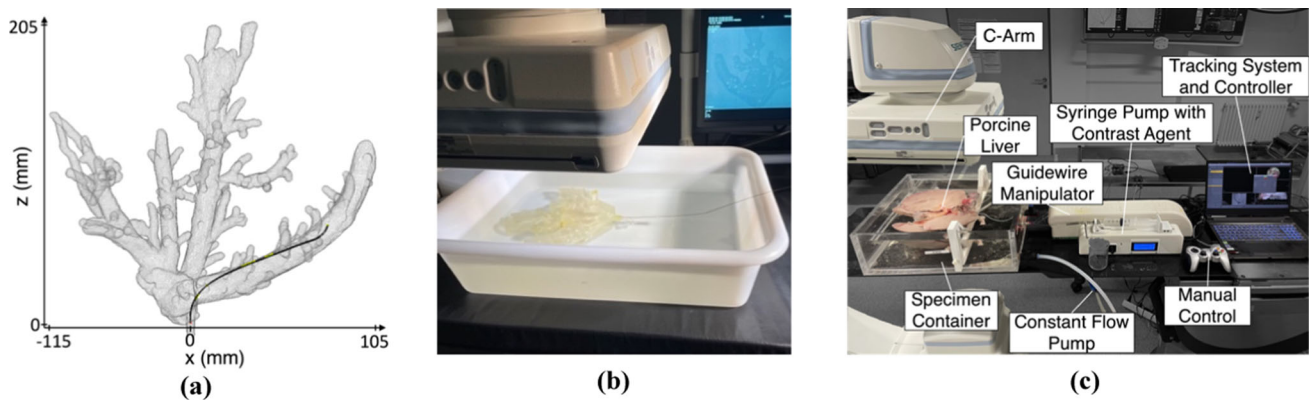
Fig. 3 **a** Simulation model, **b** Phantom testbench and **c** Ex-vivo testbench for guidewire navigation
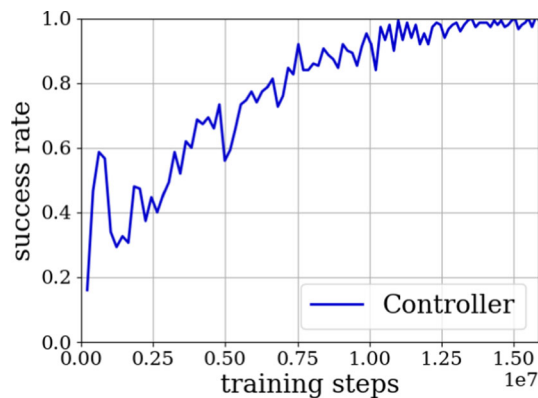
Fig. 4 Success rate during the simulation training of the controller

and 1 mm apart from each other starting from the guidewire tip with a tracking error of approximately 1.3 mm.

Results

Evaluation results during training of the controller are graphed in Fig. 4. As shown in the figure, the controller steadily improves. A success rate of 100%, meaning 150 out of 150 evaluation episodes with randomly sampled targets have been successful, is reached for the first time after 13.5×10^6 training steps but continues to fluctuate between 90 and 100% for further evaluation steps.

Figure 5 and Online Resource 1 illustrate the trajectory of the guidewire tip from the insertion point to the target by autonomous controller in silico and by a medical expert in the phantom. The autonomous controller maneuvers directly to the target, while the guidewire is continuously rotated alternately clockwise and counterclockwise, resulting in a wiggling behavior of the guidewire. In comparison, the medical expert utilizes the wiggle motion only where necessary, e.g., navigation in the vena hepatica intermedia shows very little rotation, while navigation in the vena hepatica dextra

utilizes the wiggle motion frequently. Duration of the navigation procedure is similar. For the vena hepatica sinistra, intermedia and dextra, the autonomous and the manual navigation takes 14 s and 12 s, 17 s and 16 s, and 18 s and 17 s, respectively.

The controller, solely trained in the simulation, was evaluated in the ex vivo porcine liver specimen used to extract the vessel geometry for the simulation and phantom. 10 targets are specified in each of the three main liver veins: left (*sinistra*), center (*intermedia*) and right (*dextra*). The resulting number of successful navigation (*success*), failure due to probing the wrong branch (*wrong branch*) and failure due to entanglement of the guidewire tip while advancing the guidewire in the correct branch (*entanglement*) are presented in Table 1. The controller achieves a success rate of 30%, the failure rate of probing the wrong branch is 33.3%, and the failure rate of entanglement is 36.6%. Success rate and failure types vary between the branches.

Discussion and conclusion

The navigation of guidewires in endovascular interventions is a dexterous task, and physicians and patients can benefit from automation. Machine learning-based controllers are promising to help master this task. However, training data provided by human operators are scarce and resource-intensive to generate. As it reduces the need for high-quality human-labeled data, learning with self-generated data is desirable. In this work, we have shown that a learning-based controller is capable of learning human-like guidewire navigation behavior without the usage of human-generated data. Even if no or limited data are available, a learning-based controller can perform high-quality guidewire navigation. Therefore, mitigating the requirement to use human-generated training data when training a reinforcement learning-based controller.

We have trained a neural network-based controller to navigate the venous system of a porcine liver from scratch without
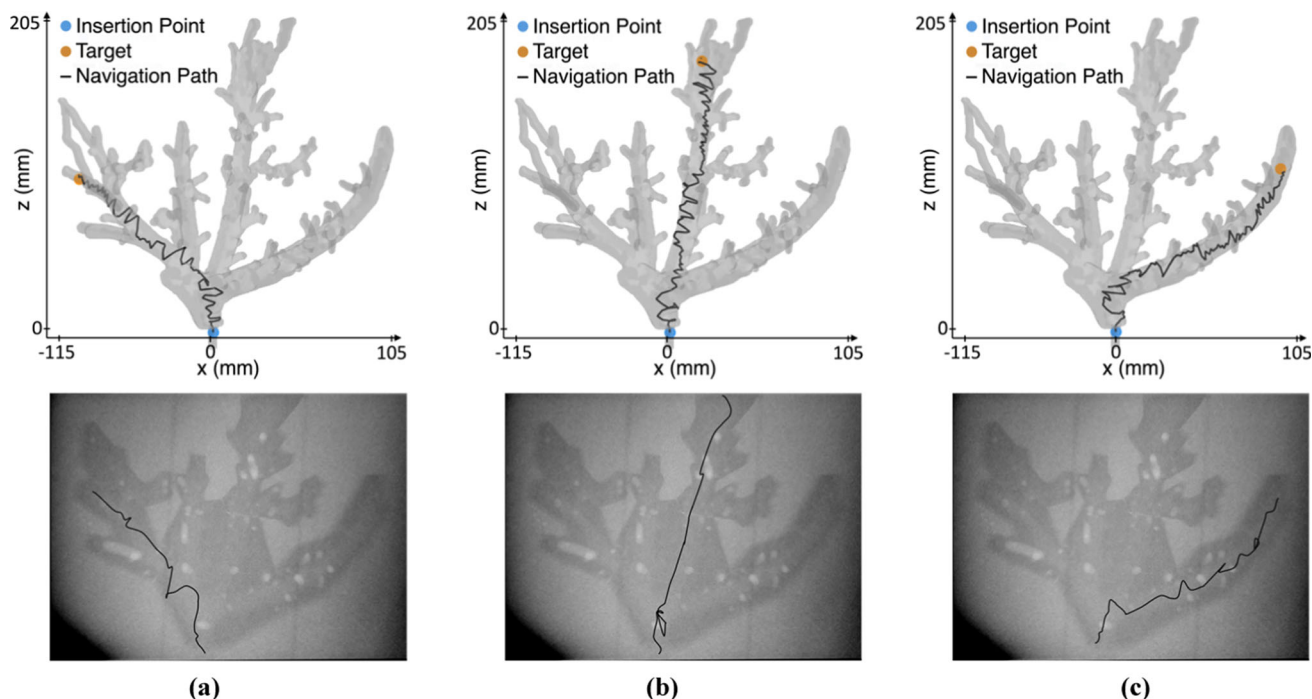
Fig. 5 Trajectories of the guidewire tip for the navigation in the vena hepatica **a** sinistra, **b** intermedia and **c** dextra. Autonomous in-silico navigation is shown in the top row and manual phantom navigation in the bottom row

Table 1 Evaluation results of the controller on the testbench stating the amount of successful navigation attempts and failures due to wrong branch navigation and entanglement

Branch	Sinistra	Intermedia	Dextra	Sum
Success	1	4	4	9
Wrong branch	6	0	4	10
Entanglement	3	6	2	11

human demonstration. The controller was trained and evaluated in a simulation environment, compared to manual expert navigation in a phantom, and transferability to real-world applications was tested by applying the controller to the ex vivo porcine liver, which was the model for the simulation.

Concluding from the simulation results one may deduce that the controller design is capable of learning to navigate an endovascular guidewire through liver veins for a static setup. The controller manages to navigate to the target in 100% of the attempts. The reason for the fluctuation of success rate at the end of training is presumably due to the controller being updated further leading to worsening of the performance. The applied high learning rate allows fast learning at the beginning but can lead to suboptimal updates when the resulting neural network controller is close to its optimum. A solution may be utilizing an adjustable learning rate, which decreases during training. The controller proves to have learned the behavior of wiggling the guidewire alternately clockwise and counterclockwise, like the maneuver a physician applies

in this situation. Interestingly, the maneuver was learned by the controller without being provided training data recorded from a physician that would indicate or teach this motion pattern. This shows that learning-based controllers are capable of learning motion patterns, which are inherently successful.

When transferred to the testbench with an ex vivo liver, the success rate of the controller drops to 30% of 30 given navigation tasks. Navigation failures by probing the wrong branch are most likely caused by registration inaccuracies, as registration for soft, deformable tissue without any visible landmarks remains a difficult task. The controller is not able to detect the failure from the guidewire position feedback. From the entanglement failures, it can be deduced that the controller does not learn general detection of entanglements and subsequent reaction to them.

There are limitations if one wants to transfer the results into clinical practice. In this study, only one vessel geometry was used, and the controller learns the movements to navigate this specific geometry instead of generalized guidewire navigation. In real-world scenarios, each patient has a unique

vessel geometry and training a specific controller for each patient is impractical. A controller with a neural network structure and observations as presented is presumably not able to adapt to general geometries. Instead, either the geometry can be given as observation or neural network elements that can recognize the geometry during navigation need to be added, e.g., recurrent layers. We currently investigate generalization across different patient geometries with promising initial results. Additionally, the safeness of the guidewire navigation regarding wall contacts and applied forces is not considered in this study. Instead, only the navigation behavior to robustly reach the target is evaluated, as the sparse reward of the training process only values reaching the target. For a clinical transition, an analysis if the behavior is safe for the patient is necessary. Applicability of this simulation-trained controller in the real world is limited. Inaccuracies of the guidewire stiffness, the segmented vessel geometry, the interaction between guidewire and vessel walls, and shape variations of the ex vivo vessels, which are not considered in the simulation, lead to a significantly reduced real-world performance. By incorporating variations of the dynamics and observations, a controller may be trained more generally and should increase its real-world performance.

In view of future research fields, the reduced real-world performance indicates the importance of evaluation in the real world when developing endovascular robots and stresses the simulation-to-real gap. The utilization of phantoms can be an intermediate step to ease the transition. Furthermore, the guidewire becoming entangled represents a common problem, which should be a matter of future investigation. The controller should be able to detect entanglements and react appropriately. It could be beneficial to use the fluoroscopy image as controller input instead of guidewire tracking points to detect entanglements. Another area of research should be navigating a two-instrument system, by including a catheter. With the catheter, the shape and stiffness of the guidewire tip can be modulated, like a concentric tube robot. This additional degree of freedom may improve the navigation capabilities of the controller.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Ethics approval This article does not contain any studies with human participants or animals performed by any of the authors. Specimen from porcine offal, for which ethics approval was not required, was used for this study.

Informed consent This article does not contain patient data.

References

1. GBD 2015 Risk Factors Collaborators (2016) Global, regional, and national comparative risk assessment of 79 behavioural, environmental and occupational, and metabolic risks or clusters of risks, 1990–2015: a systematic analysis for the Global Burden of Disease Study 2015. Lancet 388(10053):1659–1724. https://doi.org/10.1016/S0140-6736(16)31679-8

2. Lanzer P (2018) Textbook of catheter-based cardiovascular interventions: a knowledge-based approach. Springer, Berlin

3. Ho P, Cheng SWK, Wu PM, Ting ACW, Poon JTC, Cheng CKM, Mok JHM, Tsang MS (2007) Ionizing radiation absorption of vascular surgeons during endovascular procedures. J Vasc Surg 46(3):455–459. https://doi.org/10.1016/j.jvs.2007.04.034

4. Böckler D (2020) Praktische Tipps für den persönlichen Strahlenschutz bei endovaskulären Eingriffen im Hybrid-Operationssaal. Gefässchirurgie 25(1):19–30. https://doi.org/10.1007/s00772-020-00620-9

5. Yan Y, Hu K, Alcock S, Ghrooda E, Trivedi A, McEachern J, Kaderali Z, Shankar J (2022) Access to endovascular thrombectomy for stroke in rural versus urban regions. Can J Neurol Sci J Can Sci Neurol 49(1):70–75. https://doi.org/10.1017/cjn.2021.35

6. Zhou XY, Guo Y, Shen M, Yang GZ (2020) Application of artificial intelligence in surgery. Front Med 14(4):417–430. https://doi.org/10.1007/s11684-020-0770-0

7. Ghamraoui AK, Ricotta JJ (2018) Current and future perspectives in robotic endovascular surgery. Curr Surg Rep 6(12):21. https://doi.org/10.1007/s40137-018-0218-5

8. Weisz G, Metzger DC, Caputo RP, Delgado JA, Marshall JJ, Vetrovec GW, Reisman M, Waksman R, Granada JF, Novack V, Moses JW, Carrozza JP (2013) Safety and feasibility of robotic percutaneous coronary intervention. J Am Coll Cardiol 61(15):1596–1600. https://doi.org/10.1016/j.jacc.2012.12.045

9. Madder R, VanOosterhout S, Jacoby M, Collins JS, Borgman A, Mulder A, Elmore M, Campbell J, McNamara R, Wohns D (2017) Percutaneous coronary intervention using a combination of robotics and telecommunications by an operator in a separate physical location from the patient: an early exploration into the feasibility of telestenting (the REMOTE-PCI study). EuroIntervention 12(13):1569–1576. https://doi.org/10.4244/EIJ-D-16-00363

10. Berczeli M, Legeza P, Lumsden A (2021) Catheter robots in the cardiovascular system. In: Küçük S (ed) Latest developments in medical robotics systems. IntechOpen. https://doi.org/10.5772/intechopen.97314

11. Ullrich F, Schuerle S, Pieters R, Dishy A, Michels S, Nelson BJ (2014) Automated capsulorhexis based on a hybrid magnetic-mechanical actuation system. Proceedings of—IEEE international conference on robotics and automation, pp 4387–4392. https://doi.org/10.1109/ICRA.2014.6907498

12. Boskma KJ, Scheggi S, Misra S (2016) Closed-loop control of a magnetically-actuated catheter using two-dimensional ultrasound images. In: Proceedings of IEEE/RAS-EMBS international conference on biomedical robotics and biomechatronics, 2016-July, pp 61–66. https://doi.org/10.1109/BIOROB.2016.7523599

13. Heunis CM, Wotte YP, Sikorski J, Furtado GP, Misra S (2020) The ARMM system—autonomous steering of magnetically-actuated catheters: towards endovascular applications. IEEE Robot Autom Lett 5(2):705–712. https://doi.org/10.1109/LRA.2020.2965077

14. Fagogenis G, Mencattelli M, Machaidze Z, Rosa B, Price K, Wu F, Weixler V, Saeed M, Mayer JE, Dupont PE (2019) Autonomous robotic intracardiac catheter navigation using haptic vision. Sci Robot 4(29). https://doi.org/10.1126/scirobotics.aaw1977

15. Gopesh T, Wen JH, Santiago-Dieppa D, Yan B, Pannell JS, Khalessi A, Norbash A, Friend J (2021) Soft robotic steerable microcatheter for the endovascular treatment of cerebral disorders. Sci Robot 6(57):eabf0601. https://doi.org/10.1126/scirobotics.abf0601

16. Rafii-Tari H, Bicknell C, Yang G (2013) Learning-based modeling of endovascular catheterization. In: International conference on medical image computing and computer-assisted intervention, pp 369–377. https://doi.org/10.1007/978-3-642-40763-5_46

17. Tercero C, Ikeda S, Uchiyama T, Fukuda T, Arai F, Okada Y, Ono Y, Hattori R, Yamamoto T, Negoro M, Takahashi I (2007) Autonomous catheter insertion system using magnetic motion capture sensor for endovascular surgery. Int J Med Robot 3(1):52–58. https://doi.org/10.1002/rcs.116

18. Mirnezami R, Ahmed A (2018) Surgery 3.0, artificial intelligence and the next-generation surgeon. Br J Surg 105(5):463–465. https://doi.org/10.1002/bjs.10860

19. Chi W, Dagnino G, Kwok TMY, Nguyen A, Kundrat D, Abdelaziz MEMK, Riga C, Bicknell C, Yang GZ (2020) Collaborative robot-assisted endovascular catheterization with generative adversarial imitation learning. In: Proceedings of IEEE international conference on robotics and automation, pp 2414–2420. https://doi.org/10.1109/ICRA40945.2020.9196912

20. Zhao Y, Guo S, Wang Y, Cui J, Ma Y, Zeng Y, Liu X, Jiang Y, Li Y, Shi L, Xiao N (2019) A CNN-based prototype method of unstructured surgical state perception and navigation for an endovascular surgery robot. Med Biol Eng Comput 57(9):1875–1887. https://doi.org/10.1007/s11517-019-02002-0

21. Karstensen L, Behr T, Pusch TP, Mathis-Ullrich F, Stallkamp J (2020) Autonomous guidewire navigation in a two dimensional vascular phantom. Curr Dir Biomed Eng 6(1):20200007. https://doi.org/10.1515/cdbme-2020-0007

22. Åström KJ (1965) Optimal control of Markov processes with incomplete state information. J Math Anal Appl 10(1):174–205. https://doi.org/10.1016/0022-247X(65)90154-X

23. Lillicrap TP, Hunt JJ, Pritzel A, Heess N, Erez T, Tassa Y, Silver D, Wierstra D (2016) Continuous control with deep reinforcement learning. In: 4th International conference on learning representations, ICLR 2016—conference track proceedings

24. Andrychowicz M, Wolski F, Ray A, Schneider J, Fong R, Welinder P, McGrew B, Tobin J, Abbeel P, Zaremba W (2017) Hindsight experience replay. In: Advances in neural information processing systems, 2017-December, pp 5049–5059

25. Faure F, Duriez C, Delingette H, Allard J, Gilles B, Marchesseau S, Talbot H, Courtecuisse H, Bousquet G, Peterlik I, Cotin S (2012) SOFA: a multi-model framework for interactive physical simulation. In: Payan Y (ed) Studies in mechanobiology, tissue engineering and biomaterials, vol 11. Springer, Berlin, pp 283–321. https://doi.org/10.1007/8415_2012_125

26. Duriez C, Cotin S, Lenoir J, Neumann P (2006) New approaches to catheter navigation for interventional radiology simulation. Comput Aided Surg 11(6):300–308. https://doi.org/10.3109/10929080601090623

27. Zidowitz S, Drexl J, Kröger T, Preusser T, Ritter F, Weihusen A, Peitgen H-O (2007) Bayesian vessel extraction for planning of radiofrequency-ablation. In Bildverarbeitung für die Medizin 2007. Informatik aktuell., München, pp 187–191. https://doi.org/10.1007/978-3-540-71091-2_38