2024 IEEE International Conference on Robotics and Automation (ICRA 2024)
May 13-17, 2024. Yokohama, Japan

Lens Capsule Tearing in Cataract Surgery using Reinforcement Learning

Rebekka Charlotte Peter*[1,2], Steffen Peikert*[2], Ludwig Haide[1], Doan Xuan Viet Pham[1], Tahar Chettaoui[1,3], Eleonora Tagliabue[1], Paul Maria Scheikl[2], Johannes Fauser[1], Matthias Hillenbrand[1], Gerhard Neumann[3], Franziska Mathis-Ullrich[2]

***Abstract*— Cataract is the leading cause of blindness worldwide with an increasing number of patients due to changing demographics, making automation an important part in future surgical treatment. In this work, we focus on a substep of cataract surgery, the Continuous Curvilinear Capsulorhexis (CCC). With a high complexity, this task is an ideal candidate for Reinforcement Learning (RL) in simulation. First, we present an interactive and physically realistic simulation based on the Finite Element Method (FEM) that mimics the tearing behavior of soft tissue during CCC. Then, we train and evaluate RL models in simulation, demonstrating that the trained policies can complete the CCC in 85 % of cases. We also show that applying domain randomization techniques make the policy more robust against changes in geometrical and biomechanical boundary conditions.**

I. Introduction

Cataract is a common eye condition characterized by a clouded lens that results in reduced vision. It is the leading cause of blindness worldwide [1] and is treated by surgery. To restore clear vision, the clouded lens is removed and replaced with an artificial lens through a small incision in the cornea. The lens is located in the lens capsule, a thin, transparent, elastic membrane. To gain access to the lens, a circular opening is made on the anterior side of the capsule (capsulotomy). Once the clouded lens is removed, an artificial lens is inserted.

Age is the most significant risk factor for cataract, with 70 % of individuals over the age of 80 being affected [2]. Due to demographic change and increased life expectancy, the number of individuals aged 80 years and older is anticipated to triple from 2020 to 2050 [3]. To address the increasing number of patients, improved accessibility through effective cataract surgery is required, including computer-assisted and automated surgical techniques. In this work, we focus on the instrument trajectory prediction for the *continuous curvilinear capsulorhexis (CCC)*, a capsulotomy technique which creates the circular opening by grasping and tearing the lens capsule material with a specialized needle or forceps. CCC is considered one of the most difficult steps in cataract surgery. The complex, indirect tearing behavior of the delicate elastic capsule material and the limited working space within the anterior eye chamber make CCC a challenging task to learn. Especially young surgeons with little experience often struggle to perform the tearing under non-optimal conditions [4]. Yet, well-defined success and failure parameters render the procedure a suitable candidate for automation. The main requirements of the lens capsule opening are its circularity with a diameter between 5–6 mm and its location at the center of the pupil.

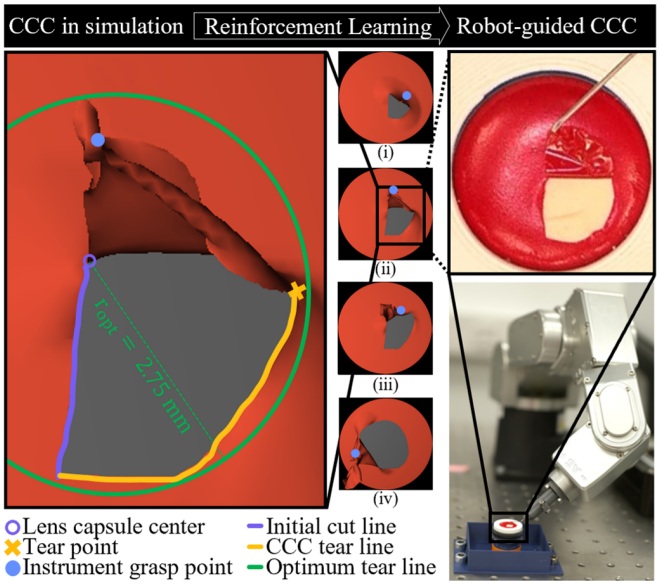

Fig. 1: CCC in simulation and robot-guided CCC on phantom eye. Phases of CCC: (i) Initial state with opening of the capsule bag created with an initial cut from center to periphery; (ii) circular propagation of the capsule bag opening by grasping and tearing the flap; (iii) instrument tip attached to flap closer to the tear point after regrasp; (iv) further propagation and regrasping until an opening of 270° is reached. On the close up, the optimum tear line (green circle) with a radius of 2.75 mm is shown. We define the *tear point* as the last point along the tear line in tearing direction. Experimental setup: With a needle mounted on the endeffector of the Meca500® robotic manipulator, the artificial capsule bag foil (red) is grasped and torn.

*These authors contributed equally to this work.
[1]Carl Zeiss AG, 73447 Oberkochen, Germany: rebekka.peter@zeiss.com
[2]Laboratory for Surgical Planning and Robotic Cognition (SPARC), Dep. Artificial Intelligence in Biomedical Engineering, Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany: {steffen.peikert, paul.m.scheikl, franziska.mathis-ullrich}@fau.de
[3]Institute for Anthropomatics and Robotic (IAR), Karlsruhe Institute of Technology, Germany: gerhard.neumann@kit.edu

To address automation in ophthalmology, various robotic platforms have been proposed. While initial studies show that robotic platforms require more time to perform a given eye surgery task [5], they can provide more safety and accuracy than a human [6]. Typical robotic applications in eye surgery are teleoperation with the daVinci system [7], or custom teleoperation robots [8], [9]. Teleoperating systems can help reduce the surgeons hand tremor due to motion scaling [8], [10], [11], [12], [13]. In one instance, a robot has also been designed to reduce the vibrations in the operation room [14]. Other systems focus on the co-operation between surgeon and robot on retinal endovascular surgery [15] and sub-retinal drug delivery [16].

To automate delicate tasks similar to CCC, reinforcement learning (RL) is considered promising. The main drawback of RL in the surgical setting is its substantial demand for real-world data which is virtually impossible to obtain, necessitating the use of training simulations. These often focus on specific scenarios like soft tissue manipulation [17], tissue retraction [18], [19] and cutting [20], [21], [22], laparoscopic tasks [23] or specific surgical robots like the daVinci system [24].

For cataract surgery, few robotic or automated assistance has been shown. Several learning-based algorithms have been proposed to grade and detect cataract [25], [26]. Notable exceptions are the IRISS system for automated lens extraction [27], [28] and the Femtosecond laser-assisted cataract surgery (FLACS) method for CCC [29], [30]. While initial results for FLACS look promising [29], [31], this method is much more expensive than conventional CCC [32], requires specialized laser-instruments and surgical training [29], [31], and produces weaker cuts [33]. For the automation of the CCC, the definition of an analytical tissue model for the complex tearing behavior of soft tissue is not trivial. Minor variations in the movement of the surgical instrument or the boundary conditions (e.g., the geometry and biomechanical characteristics of the lens capsule) lead to major changes in the tear propagation. This facilitates the need for a machine-learning solution, which can adapt to patient-specific boundary conditions while performing the operation. However, a suitable set of training data covering a large range of tearing behavior is not available. To overcome the lack of data and to address the need for decision-making based on the current state of the scene, RL in simulation is suited. This requires an interactive simulation of the tearing caused by instrument-tissue-interactions during the CCC.

Simulations of the CCC are also used in surgical simulators to support the training of young surgeons. Conventional mass spring models have been investigated to replicate the capsule behavior [34], [35], [36], [37], [38], [39]. Although these models offer fast computation and real-time capabilities, accurate representation of the physical behavior is challenging due to significant simplifications of the represented capsular tissue.

An alternative are Finite-Element Method (FEM) models, which model the underlying differential equations to correctly predict tissue physics. These models are usually computationally expensive and not real-time capable. However, they have been shown useful for haptic simulations [40]. For CCC, a model based on anisotropic tear propagation has been proposed in [41], [42] and a complete FEM 3D eye model has been ported to GPU in [43].

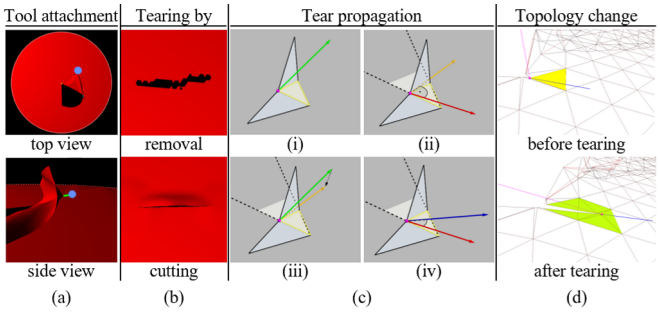

Fig. 2: (a) The tip of the (virtual) instrument is marked with a blue dot. The thin green line visualizes the stiff spring that is used to model the attachment of the instrument to the capsule bag surface. (b) Comparison of tearing by removal and tearing by cutting. (c) Calculation of the next propagation direction. The current end point of the tear is highlighted in magenta, the main FEM triangle to tear in yellow. The maximum principal stress (green arrow) is projected onto the plane of the main triangle (orange vector). Then, the orthogonal direction to this projected vector is found in the neutral plane (red vector). Finally, the orthogonal direction is rotated by a predetermined angle (blue vector). (d) Effect of topology change during tearing: The yellow and its neighboring mesh elements are remeshed after the tearing step.

Fig. 1 gives an overview of our tearing simulation and experimental setup. In this paper, we present a novel RL training scenario for CCC instrument trajectories. We extend the Simulation Open Framework Architecture (SOFA) framework with a soft-tissue tearing simulation, provide a training policy integrated into LapGym [23] and demonstrate that a RL agent can learn and control indirect soft tissue tearing in a simulated environment. We also show that Domain Randomization (DR) is a necessary step for training CCC.

II. Training and Evaluation Environment

To train and evaluate our approach, we rely on a simulation which provides a safe and controlled environment for the RL agent to learn. To ensure that our simulation accurately reflects the physical world, we compare the tearing behavior of the simulation with the tearing behavior of a phantom eye.

A. Phantom Eye

We use the SimuloRhexis® phantom eye (SimulEYE®, Westlake CA) as a physical model on a custom holder for CCC (Fig. 1, right side). It is designed for training surgeons on capsulorhexis techniques and serves as a validation platform for our tearing simulation. The phantom eye is comprised of a circular-shaped foil that simulates the capsule bag and has similar tearing characteristics to human capsular

bags. The foil is placed on top of a flatbed of clay. This acts both as a (semi-) solid surface to prevent gravity from acting on the foil while also mimicking the more dense, but still elastic, lens material. To imitate the forces applied from the zonular fibers to the outer borders of the lens capsule, the foil is clamped with a ring to the holder.

B. Tearing Simulation

1) Lens Capsule Model: We generate a simulated environment for CCC tearing in SOFA [44] which is based on co-rotational FEM to enable large-deformation handling [45]. The anterior lens capsule is modeled with a circular and plane 2D membrane mesh of triangles that can be moved and deformed out-of-plane using linear shape functions for the displacements.

We integrate material parameters of human eyes to provide a realistic tearing behavior in simulation. Due to limited information in literature, we estimate the density of the human lens capsule based on the mass density of collagen, the main component of the lens capsule, which is approximately $1.5\,\mathrm{g/cm^3}$ [46]. With a diameter of $10\,\mathrm{mm}$ and a maximum thickness of $21\,\mu\mathrm{m}$, the mass is approximately $2.5\,\mathrm{mg}$. Due to the piece-wise linear stress-strain relationship of the capsule material [47], we use a linear elastic model in our simulation. The Young's modulus is $5\,\mathrm{N/mm^2}$, which is within the range of values provided in literature [47], [48]. In a real (human) eye, the capsular bag is attached to the zonular fibers, which elastically fixate the lens. In our simulation model, we simplify the capsule fixation by applying positional constraints to the border nodes of our mesh. To replicate the lens material below the capsule bag, a force field of $100\,\mathrm{N/mm}$ is applied, pushing against the mesh and negating gravity. While this does not model a complete lens, it provides a similar elastic force feedback as surgeons would experience within a real human eye. In a human eye, the capsule adheres to the lens cortex and requires a small force to be lifted. We model this adhesion as a planar force field of $5\,\mathrm{N/mm}$ slightly above the mesh of the capsule bag, pointing in the direction of the mesh. Thus, if a node is below the force field, it is pushed downwards and can only move freely once above. As the capsule is modeled by a 2D mesh, two neighboring triangles can be rotated around their common edge without any elastic resistance to such bending. To include a bending stiffness, we add linear springs between the mesh nodes that oppose bending. We also include linear dampers between the nodes to account for the viscoelastic properties of the lens capsule.

2) Surgical Instrument Model: We simulate the surgical instrument as a single point in space, which is attached to a mesh node by a stiff spring as shown in Fig. 2 (a).

3) Tearing: When the instrument is maneuvered during CCC, the attached flap causes the capsule material to undergo deformation resulting in stress within the mesh elements. If the stress exceeds a certain threshold, the material is torn. We employ the von Mises stress criterion, which balances the considerations between tearing from normal stresses and tearing from shear stresses. For our simulation, the von Mises stress threshold is set to $3.2\,\mathrm{N/mm^2}$, similar to values provided in [47]. Instead of removing a whole element from the mesh if the tear threshold is exceeded (*tearing by removal*), *tearing by cutting* is implemented, see Fig. 2 (b). In each forward step of the simulation, it is determined if and how the tear propagates. The propagation is characterized by the direction and the length of the tear in the next time step of the simulation. To determine the direction, the following steps, cf. Fig. 2 (c), are performed: First, the principal stress direction of the triangle with maximum stress is projected onto the plane of the main FEM triangle, which is the mesh element in the direction of the previous propagation vector. Then, we determine the orthogonal direction of this projected vector in the default plane of the mesh to account for normal stresses. Shear stresses are considered by rotating the orthogonal direction by an angle of $40°$ in the same plane. By averaging the resulting vector with the previous propagation direction, we incorporate tear direction continuity. Along the calculated propagation direction, a straight incision is made with a predetermined length. The incision causes a remesh to accommodate the change in the mesh topology, as shown in 2 (d).

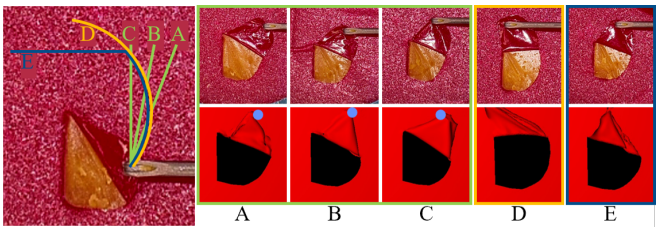

Fig. 3: Validation of the tearing simulation. Starting with the initial flap grasped with a cystotome, five robot-guided tool tip trajectories (A-E) are performed on the SimuloRhexis® phantom eye (top row) and in simulation (bottom row), starting with the initial flap grasped by a cystotome.

C. Validation of the Tearing Simulation

We verify the accuracy of the tearing simulation by comparing the simulated tearing behavior with the tearing behavior observed in the SimuloRhexis® phantom eye.

To conduct defined and repeatable instrument movements on the phantom eye, a surgical needle (cystotome) is mounted on and controlled by the Meca500® robotic manipulator (MECADEMIC®, Canada). Simplifying the experiment, the cornea of the phantom eye is removed to ease accessibility to the foil.

For evaluation, a small initial cut is made into the foil and propagated to create the initial flap, the starting point for tearing. Five tearing trajectories are performed, with the capsule flap grasped by the instrument. The same orientation of the cystotome is maintained over the complete trajectory. For three of the trajectories, the instrument is moved on a straight line of $4\,\mathrm{mm}$, with different angles relative to the initial cut. The fourth trajectory consists of a circular instrument tip motion until it reaches an angle of $180°$ relative to the initial cut. The last trajectory is a combination

of both, a circular and a straight-line movement. The tearing experiments are recorded in a top-down view with the built-in camera of a surgical microscope (ARTEVO 700, Carl Zeiss AG). For a side-by-side analysis, identical trajectories are executed in the SOFA simulation, and the shapes of the capsule bag tears are compared. Fig. 3 shows results.

To ensure comparability between initial conditions of the tearing experiments in reality and simulation, we average all initial flap openings from the SimuloRhexis® phantom eye experiments. The averaged flap opening geometry is then transferred to the simulation for all experiments. A predetermined mesh node serves as the first grasping point.

We use the Root Mean Squared Error (RMSE) to compare the radius of the tear line in the simulation with that in the real world (relative to the center of the lens capsule, sampled at intervals of $0.1°$). The average deviation between the tear lines in simulation and the SimuloRhexis® phantom eye is $0.13\,\mathrm{mm}$ (RMSE), which is in the range of a single mesh element. Minor differences observed between two trials of the same trajectory are attributed to potential variations in the initial flap opening, a different attachment point of the flap to the instrument, measurement errors, or varying material properties in the phantom eye foil. For all test trajectories, the *final angle* deviates in the range of $3°$ to $8°$ and is shown to be highly comparable. Summarizing, the tearing simulation has proven sufficiently similar to real-world experiments and can serve as a realistic baseline for later RL.

III. Learning Soft-tissue Tearing

We train a RL policy to predict the trajectory that the instrument tip needs to follow in order to generate a circular and centered tear with a radius of $2.75\,\mathrm{mm}$. Every training episode starts after a pre-determined initial flap is created (Fig. 1, phase (i)). The task is considered complete when a circular tear of $270°$ is achieved (Fig. 1, phase (iv)), excluding the finalization of the rhexis. The policy is learned in simulation, which allows for safe and efficient exploration during the training process.

A. Reinforcement Learning

Our scene uses SOFA (version 21.12) accessed via a modified training environment of LapGym, a framework for RL in robot-assisted surgery that supports deformable objects and topological changes [23]. We train using Proximal Policy Optimization (PPO) via stable-baselines3 [49], which has shown promising results for related soft-tissue manipulation tasks like tissue dissection or cutting in [23]. Agents are trained with a discount of $\gamma = 0.995$ and $\gamma_{\text{GAE}} = 0.95$. The batch size is set to 128, with 256 steps in 16 parallel environments and 4 epochs per iteration. The learning rate and clip ratio are set to 0.0002 and 0.2, respectively, with a linear scheduler applied to both values. Advantages are normalized after each iteration. The "Nature CNN" from [50] is used as the visual feature extractor.

1) Action Space: We define our actions $\vec{a} = (a_x, a_y, a_z)$ as a three-dimensional vector within the action space $\mathcal{A} = [-25, 25]^3 \mu\mathrm{m}$. Each action corresponds to a linear movement of the instrument tip in Cartesian space. Limiting the action space to $\pm 25\mu\mathrm{m}$ prevents the RL agent from taking excessively large or erratic actions and provides smooth instrument movements.

2) Observation Space: We use stacks of the four latest RGB images from a top-down perspective, which represents the view from an ophthalmic surgical microscope, as observations for our policy. This allows the agent to capture motion information over the previous $4 \times 0.01\mathrm{s}$. Each observation is of size $128 \times 128 \times 12$.

3) Reward Function: The reward function $\mathcal{R}$ consists of four components that encourage a tear propagation following the ideal tear line. Towards this end, we define our reward as

$$\mathcal{R} = \mathcal{R}_{\delta(\theta)} + \mathcal{R}_{\delta(\mathtt{r})} + \mathcal{R}_{\mathtt{progress}} + \mathcal{R}_{\mathtt{success}}.$$

We define r as the radius of the last tear point, and θ as the angle between the last tear point and the direction of the initial cut, both with respect to the center of the capsule bag. The change in the radius and angle between consecutive frames is denoted by $\delta(r)$ and $\delta(\theta)$, respectively.

$\mathcal{R}_{\delta(\theta)}$ gives a reward if a change in the tear angle is observed ($\delta(\theta) \neq 0$). If the tear angle increases towards our target angle, a positive reward is given and vice versa. The total achievable reward is normalized, such that, once the CCC is complete ($\theta = 270°$), a total reward of 1 is achieved.

In similar fashion, $\mathcal{R}_{\delta(\mathtt{r})}$ rewards a change in radius on tear progression. If the new radius is closer to the optimal tear radius $r_{opt} = 2.75\,\mathrm{mm}$, a positive reward is given. If the deviation increases, the agent is penalized with a negative reward. Again, this reward is normalized in such a way that for reaching the optimal tear radius, a cumulative reward of 1 is given.

$\mathcal{R}_{\mathtt{progress}}$ is defined as 1.01^{θ} and entices the agent to progress. Because the reward is dependent on the total angle, larger rewards are given the closer the agent is to completing the CCC. To avoid reward hacking, it is only given if the tear angle has changed ($\delta(\theta) \neq 0$) and the tear is close to the optimal tear ($\|\mathtt{r} - \mathtt{r}_{opt}\| < 0.5\,\mathrm{mm}$).

$\mathcal{R}_{\mathtt{success}} = 100$ is only given once the agent has reached the target angle $\theta_{target} = 270°$.

4) Termination Conditions: We terminate a training episode if $\theta \geq 270°$ or after a time limit of 1000 steps $\hat{=}$ $100\mathrm{s}$.

B. Regrasping

CCC requires the surgeon to regrasp the capsular bag several times in order to achieve a $270°$ tear. In our work, the task to be learned focuses solely on the decision on *when* and *where* to regrasp, without considering the required instrument movement from the old to the new grasping location or the grasping motion itself. We define heuristic rules for the regrasp decision based on domain knowledge acquired during discussions with ophthalmic experts and the analysis of surgical videos. Regrasping is performed for two main reasons: First, when the instrument tip is closer to the tear point, controlling the tear's trajectory becomes easier. This

implies that a change in the instrument's movement direction has a more pronounced effect on the tear's direction. Secondly, if the distance between the instrument tip and the tear point becomes too large, the workspace restrictions cannot be maintained. Thus, we define the following regrasping rule: If the distance between the instrument tip and the tear position exceeds a threshold of $3\,\mathrm{mm}$, a regrasp is performed.

To determine the new grasp location, we use the path finding `A*` algorithm [51] to find the shortest path along the mesh edges between the tear point and the old grasping position. The new instrument position is chosen at the node at $\approx 25\,\%$ along this path. This ensures that the new grasp point is on the capsule flap and in close proximity to the current tear point. On implementation level, the connection between the instrument and the mesh is separated and reinstated on the location of the new attachment point during the same simulation step.

C. Domain Randomization

We use DR to increase the robustness of our learned model to changes in the environment. Besides reducing the sim-to-real gap, DR also allows us to model inter-patient differences in tissue characteristics and geometry, as well as variability in the geometry of the initial flap, i.e., the agent's starting state. For each new training, we select a set of random tissue and state parameters. The most significant parameters that affect the tearing simulation are the Young's modulus and the tear stress threshold of the capsule bag, both of which are selected based on a uniform distribution within $5\,\%$ of the corresponding mean values defined in Section II. The range for DR is oriented to the expected variance of geometrical and biomechanical properties for human eye CCC.

IV. Evaluation

A. Metrics

We evaluate the trained policies across 100 environments in simulation and examine the tearing trajectories once the agent terminates. The evaluation of a policy's performance is based on metrics reflecting progress and accuracy. Progress reflects how far the tear has progressed along the circle. Accuracy-based metrics assess how well the tearing trajectory aligns with the optimal tear, i.e., a centered circle with a radius of $2.75\,\mathrm{mm}$.

1) Completion Rate: An environment is defined as `complete` if the tear reaches the target angle of $270°$.

2) Success Rate: An environment is defined as successful if the generated tear reaches the target angle of $270°$ (i.e., the environment is `complete`) and the maximum distance between the generated optimal tear is less than $1\,\mathrm{mm}$.

3) Robustness: To identify whether a given policy could be trained robustly (i.e., produces reliable and repeatable results), we investigate the mean and maximum distance to the optimal tear for `completed` results.

B. Experimental Setup

We define four experiments to investigate the influence of DR and regrasping on the training in our simulation environment.

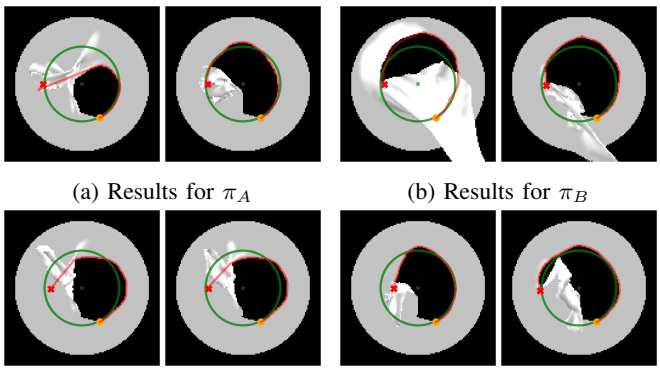

(a) Results for π_A (b) Results for π_B

(c) Results for π_C (d) Results for π_D

Fig. 4: Overview of resulting CCC lines (red) as generated by our policies. On the left side is the best completed but not successful CCC and on the right, the best successful CCC, where best is defined as the smallest maximum distance to the optimal tear line (green circle). The yellow dot marks the start of the CCC line and the red x identifies its end.

In the first experiment, we evaluate the policy π_A with enabled DR during training and evaluation. The agent is allowed to perform regrasps. This experiment aims to investigate whether the agent can learn the complex relationship between instrument tip movement and tear propagation in simulation and gains control over the tearing task.

The second experiment (policy π_B) focuses on learning in a randomized environment without regrasping. This experiment investigates whether regrasping is necessary.

Experiments three (policy π_C) and four (policy π_D) investigate the effect of DR. Regrasping is enabled. For policy π_C, the agent is trained and evaluated without DR. This is deemed the simplest task, as the environment does not change but the agent still has the best control over the tear due to regrasping. In contrast, with policy π_D we evaluate the agent trained without DR in a randomized environment. This yields insights into how much alterations in tissue parameter and capsule bag geometry influence the performance.

V. Results and Discussion

It is essential to balance the conflicting reward components to avoid the agent prioritizing either progress (completion without an accurate tear circle) or accuracy (tear line close to the circle but slow or no movement). We discovered that a promising approach is to use a combination of sparse rewards, such as a $\mathcal{R}_{\texttt{success}}$ at the end of an episode, and dense rewards, such as $\mathcal{R}_{\delta(\theta)}$ and $\mathcal{R}_{\delta(\mathbf{r})}$.

Table I presents the results of our experiments in the tearing simulation. Fig. 4 showcases both the best completed but not successful (left pictures) and best successful trajectory (right pictures) for each policy. With our baseline policy π_A, we achieve a completion and success rate of $62\,\%$ and $25\,\%$, respectively. Noticeable is the high standard deviation of the maximum distance of the target CCC radius of $21.74\,\mathrm{mm}$. The qualitative investigation of the results confirm that the policy is able to generate accurate and completed tearing

TABLE I: Results of CCC in simulation using Reinforcement Learning ($n = 100$). Tear distances to the target CCC radius is evaluated for `completed` environments only.

Scenario				Tear distance to target CCC radius [mm]		Completion	Success
Policy	**DR_{train}**	**DR_{eval}**	**Regrasp**	**Mean distance**	**Mean max distance**	**rate** [%]	**rate** [%]
π_A	✓	✓	✓	$+0.223$ $(SD = 2.104)$	4.94 $(SD = 21.74)$	62	25
π_B	✓	✓	✗	$+0.585$ $(SD = 1.288)$	2.48 $(SD = 2.67)$	40	2
π_C	✗	✗	✓	$+0.046$ $(SD = 0.068)$	1.18 $(SD = 0.39)$	85	5
π_D	✗	✓	✓	-0.129 $(SD = 0.379)$	1.82 $(SD = 1.67)$	57	19

lines for some environments. However, in the cases that the policy fails, the deviation towards the optimal tear is high. Further investigation is needed to find the failure reasons. Possible reasons might be instabilities of the simulation resulting in unexpected behavior or a high dependency on the varying tissue and geometry parameter.

Our results for π_B suggest that automation without regrasping is not a feasible option. When comparing the performance of policy π_A and π_B with and without the regrasping feature, respectively, π_A outperformed π_B with respect to the success and completion rate. π_B achieves an insufficient success rate of $2\,\%$. Regrasping is a particularly challenging task to automate due to the high perception requirements to localize the 6D regrasp pose and high control requirements to grasp the deformable lens capsule. Thus, an instrument trajectory without the need of regrasping would be desirable for automation of CCC. However, if we were to consider cornea boundaries to restrict instrument movements in the anterior eye, we anticipate a further decrease in performance without regrasping.

The training task that involved static boundary conditions during both training and evaluation (π_C) yielded the highest completion rate. This can be referred to the less complex training task (no DR during training and evaluation) and potentially the overfitting effect. For π_C, the success rate is unexpectedly low ($5\,\%$). We observe several cases where the maximum distance is only slightly over our threshold of $1\,\mathrm{mm}$ for successful completion. These cases could be considered "marginally unsuccessful". The mean maximum distance for these cases was 1.18 $(SD = 0.39)$. Therefore, and based on qualitative results, we conclude that the success metric with this threshold does not fully capture the performance of π_C.

π_D performs worse than π_A with respect to completion and success rate. This shows that domain randomization during training of π_A helped to make the policy more robust for a variation of lens capsule properties.

Future work is the transfer of the trained agent to physical eye models. Our observations during robot-guided tearing on the SimuloRhexis® phantom eye exhibited similar tear lines as the simulation. Therefore, we anticipate comparable tearing behavior while controlling the robot based on the RL policies, too. However, the sim-to-real transfer presents a challenge due to the domain gap between the visual appearance of the simulation and the phantom model. Investigating the use of vector observations including robot states and visual key points from the tear geometry through computer vision shows promise in addressing this challenge. Furthermore, our preliminary results indicate that the current policies generate trajectories where the instrument is significantly lifted from the capsule bag, making it harder to efficiently control and hold the flap. In contrast, when observing human-guided trajectories, the instrument tip touches the lens capsule while tearing which makes it easier to hold the grasp. To overcome this, we propose including a penalty term in the reward function to discourage the instrument from moving away from the foil surface. This approach would also be beneficial when restricting the allowed workspace to the anterior chamber, which is limited by the cornea.

VI. Conclusion

We demonstrate the implementation of a realistic tearing simulation using optimized parameters from a real phantom model, which enables the investigation of RL in simulation for CCC. We trained and compared several policies using RL for automating planning of the instrument trajectory for CCC. Our tearing simulation, despite its high level of abstraction, accurately models the tearing behavior of a phantom eye. Successful validation of the similarity between simulation and reality provides a crucial step towards future sim-to-real transfer of automated instrument movement.

Our experiments demonstrate that successful training on the CCC task of an RL agent is possible in simulation. Additional adaptations of the reward functions and fine-tuning of training parameters, such as the resolution and length of the image observation stack, and learning rate, or the level of DR, are required to address the remaining challenges and achieve higher success rates. Future work will focus on increasing the realism and stability of the simulation, on investigating sim-to-real strategies, and on validating the predicted trajectory on the SimuloRhexis® phantom eye. In particular, challenges to be solved include the inter-patient variance in boundary conditions, and challenging automation of the grasping process.

In conclusion, the combination of a validated tearing simulation and RL to learn optimal instrument trajectories holds great potential for improving and automating CCC procedures in the anterior eye. Further research is necessary to fully realize this potential.

References

[1] J. D. Steinmetz, R. R. Bourne, P. S. Briant, S. R. Flaxman, H. R. Taylor, J. B. Jonas, A. A. Abdoli, W. A. Abrha, A. Abualhasan, E. G. Abu-Gharbieh *et al.*, "Causes of blindness and vision impairment in 2020 and trends over 30 years, and prevalence of avoidable blindness

in relation to vision 2020: the right to sight: an analysis for the global burden of disease study," *The Lancet Global Health*, vol. 9, no. 2, pp. e144–e160, Dec 2021.

[2] National Eye Institute (NIH). (2020, Feb.) Cataract tables. [Online]. Available: https://www.nei.nih.gov/learn-about-eye-health/eye-health-data-and-statistics/cataract-data-and-statistics/cataract-tables

[3] World Health Organization (WHO), "Ageing and health," Oct. 2022. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/ageing-and-health

[4] J. Dong, X. Wang, X. Wang, and J. Li, "A practical continuous curvilinear capsulorhexis self-training system," *Ind. J. of Ophthalmology*, vol. 69, no. 10, pp. 2678–2686, 2021.

[5] T. Edwards, K. Xue, H. Meenink, M. Beelen, G. Naus, M. Simunovic, M. Latasiewicz, A. Farmery, M. De Smet, and R. MacLaren, "First-in-human study of the safety and viability of intraocular robotic surgery," *Nature Biomed. Eng.*, vol. 2, no. 9, pp. 649–656, 2018.

[6] M. Zhou, Q. Yu, K. Huang, S. Mahov, A. Eslami, M. Maier, C. P. Lohmann, N. Navab, D. Zapp, A. Knoll, and M. A. Nasseri, "Towards robotic-assisted subretinal injection: A hybrid parallel–serial robot system design and preliminary evaluation," *IEEE Trans. on Ind. Electron.*, vol. 67, no. 8, pp. 6617–6628, 2020.

[7] D. H. Bourla, J. P. Hubschman, M. Culjat, A. Tsirbas, A. Gupta, and S. D. Schwartz, "Feasibility study of intraocular robotic surgery with the da vinci surgical system," *Retina*, vol. 28, no. 1, pp. 154–158, 2008.

[8] T. Nakano, N. Sugita, T. Ueta, Y. Tamaki, and M. Mitsuishi, "A parallel robot to assist vitreoretinal surgery," *Int. J. of Comput. Assisted Radiol. and Surgery*, vol. 4, pp. 517–526, 2009.

[9] H. Meenink, R. Hendrix, P. Rosielle, M. Steinbuch, H. Nijmeijer, and M. de Smet, "A master-slave robot for vitreo-retinal eye surgery," in *Int. Conf. Eur. Soc. for Precis. Eng. and Nanotechnol.*, 2010, pp. 408–411.

[10] R. A. MacLachlan, B. C. Becker, J. C. Tabares, G. W. Podnar, L. A. Lobes, and C. N. Riviere, "Micron: An actively stabilized handheld tool for microsurgery," *IEEE Trans. on Robot.*, vol. 28, no. 1, pp. 195–212, 2012.

[11] R. Taylor, P. Jensen, L. Whitcomb, A. Barnes, R. Kumar, D. Stoianovici, P. Gupta, Z. Wang, E. Dejuan, and L. Kavoussi, "A steady-hand robotic system for microsurgical augmentation," *Int. J. of Robot. Res.*, vol. 18, no. 12, pp. 1201–1210, 1999.

[12] A. Üneri, M. A. Balicki, J. Handa, P. Gehlbach, R. H. Taylor, and I. Iordachita, "New steady-hand eye robot with micro-force sensing for vitreoretinal surgery," in *Int. Conf. on Biomed. Robot. and Biomechatronics*, 2010, pp. 814–819.

[13] T. Meenink, G. Naus, M. de Smet, M. Beelen, and M. Steinbuch, "Robot assistance for micrometer precision in vitreoretinal surgery," *Investigative Ophthalmology & Visual Sci.*, vol. 54, no. 15, pp. 5808–5808, 2013.

[14] M. A. Nasseri, M. Eder, S. Nair, E. Dean, M. Maier, D. Zapp, C. P. Lohmann, and A. Knoll, "The introduction of a new robot for assistance in ophthalmic surgery," in *Int. Conf. of the IEEE Eng. in Medicine and Biol. Soc.*, 2013, pp. 5682–5685.

[15] A. Gijbels, J. Smits, L. Schoevaerdts, K. Willekens, E. B. Vander Poorten, P. Stalmans, and D. Reynaerts, "In-human robot-assisted retinal vein cannulation, a world first," *Ann. of Biomed. Eng.*, vol. 46, pp. 1676–1685, 2018.

[16] J. Cehajic-Kapetanovic, K. Xue, T. L. Edwards, T. C. Meenink, M. J. Beelen, G. J. Naus, M. D. de Smet, and R. E. MacLaren, "First-in-human robot-assisted subretinal drug delivery under local anesthesia," *Ameri. J. of Ophthalmology*, vol. 237, pp. 104–113, 2022.

[17] E. Tagliabue, A. Pore, D. Dall'Alba, E. Magnabosco, M. Piccinelli, and P. Fiorini, "Soft Tissue Simulation Environment to Learn Manipulation Tasks in Autonomous Robotic Surgery," in *Int. Conf. on Intell. Robots and Syst.*, Oct. 2020, pp. 3261–3266, iSSN: 2153-0866.

[18] P. M. Scheikl, E. Tagliabue, B. Gyenes, M. Wagner, D. Dall'Alba, P. Fiorini, and F. Mathis-Ullrich, "Sim-to-Real Transfer for Visual Reinforcement Learning of Deformable Object Manipulation for Robot-Assisted Surgery," *IEEE Robot. and Automat. Lett.*, vol. 8, no. 2, pp. 560–567, Feb. 2023.

[19] T. Huang, K. Chen, B. Li, Y. Liu, and Q. Dou, "Demonstration-guided reinforcement learning with efficient exploration for task automation of surgical robot," *Int. Conf. on Robotics and Automat.*, pp. 4640–4647, 2023.

[20] B. Thananjeyan, A. Garg, S. Krishnan, C. Chen, L. Miller, and K. Goldberg, "Multilateral surgical pattern cutting in 2D orthotropic gauze with deep reinforcement learning policies for tensioning," *Int. Conf. on Robotics and Automat.*, pp. 2371–2378, May 2017.

[21] N. D. Nguyen, T. Nguyen, S. Nahavandi, A. Bhatti, and G. Guest, "Manipulating Soft Tissues by Deep Reinforcement Learning for Autonomous Robotic Surgery," *Int. Syst. Conf.*, pp. 1–7, Apr. 2019.

[22] T. Nguyen, N. D. Nguyen, F. Bello, and S. Nahavandi, "A New Tensioning Method using Deep Reinforcement Learning for Surgical Pattern Cutting," *Int. Conf. on Ind. Technol.*, pp. 1339–1344, 2019.

[23] P. M. Scheikl, B. Gyenes, R. Younis, C. Haas, G. Neumann, M. Wagner, and F. Mathis-Ullrich, "LapGym - an open source framework for reinforcement learning in robot-assisted laparoscopic surgery," *J. of Mach. Lear. Res.*, vol. 24, no. 368, pp. 1–42, 2023.

[24] F. Richter, R. K. Orosco, and M. C. Yip, "Open-sourced reinforcement learning environments for surgical robotics," 2020.

[25] J. Nayak, "Automated classification of normal, cataract and post cataract optical eye images using svm classifier," in *Proc. of the World Congr. on Eng. and Comput. Sci.*, vol. 1, 2013, pp. 23–25.

[26] J.-J. Yang, J. Li, R. Shen, Y. Zeng, J. He, J. Bi, Y. Li, Q. Zhang, L. Peng, and Q. Wang, "Exploiting ensemble learning for automatic cataract detection and grading," *Comput. Methods and Programs in Biomedicine*, vol. 124, pp. 45–57, 2016.

[27] E. Rahimy, J. Wilson, T.-C. Tsao, S. Schwartz, and J.-P. Hubschman, "Robot-assisted intraocular surgery: development of the IRISS and feasibility studies in an animal model," *Eye*, vol. 27, no. 8, pp. 972–978, Aug. 2013.

[28] C.-W. Chen, Y.-H. Lee, M. Gerber, H. Cheng, Y.-C. Yang, A. Govetto, A. Francone, S. Soatto, W. Grundfest, J.-P. Hubschman, and T.-C. Tsao, "Intraocular robotic interventional surgical system (IRISS): Semi-automated OCT-guided cataract removal," *Int. J. of Med. Robot. and Comput. Assisted surgery*, vol. 14, p. e1949, Aug. 2018.

[29] S. J. Bali, C. Hodge, M. Lawless, T. V. Roberts, and G. Sutton, "Early experience with the femtosecond laser for cataract surgery," *Ophthalmology*, vol. 119, no. 5, pp. 891–899, 2012.

[30] K. E. Donaldson, R. Braga-Mele, F. Cabot, R. Davidson, D. K. Dhaliwal, R. Hamilton, M. Jackson, L. Patterson, K. Stonecipher, S. H. Yoo *et al.*, "Femtosecond laser–assisted cataract surgery," *J. of Cataract & Refractive Surgery*, vol. 39, no. 11, pp. 1753–1763, 2013.

[31] T. V. Roberts, M. Lawless, S. J. Bali, C. Hodge, and G. Sutton, "Surgical outcomes and safety of femtosecond laser cataract surgery: a prospective study of 1500 consecutive cases," *Ophthalmology*, vol. 120, no. 2, pp. 227–233, 2013.

[32] H. W. Roberts, A. C. Day, and D. P. O'Brart, "Femtosecond laser–assisted cataract surgery: a review," *Eur. J. of Ophthalmology*, vol. 30, no. 3, pp. 417–429, 2020.

[33] S. Daya, S.-P. Chee, S.-E. Ti, R. Packard, and D. H. Mordaunt, "Comparison of anterior capsulotomy techniques: continuous curvilinear capsulorhexis, femtosecond laser-assisted capsulotomy and selective laser capsulotomy," *Brit. J. of Ophthalmology*, vol. 104, no. 3, pp. 437–442, 2020.

[34] M. Agus, E. Gobbetti, G. Pintore, G. Zanetti, and A. Zorcolo, "Real-time cataract surgery simulation for training," in *Eurographics Italian Chapter Conf.*, vol. 6, 2006, pp. 183–187.

[35] J. Grimm, "Tearing of membranes for interactive real-time surgical training," *Stud. Health Technol. Inform.*, vol. 111, pp. 153–159, 2005.

[36] R. Webster, J. Sassani, R. Shenk, M. Harris, J. Gerber, A. Benson, J. Blumenstock, C. Billman, and R. Haluck, "Simulating the continuous curvilinear capsulorhexis procedure during cataract surgery on the eyesi system," *Stud. Health Technol. Inform.*, vol. 111, no. 592, p. 5, 2005.

[37] K. Weber, C. Wagner, and R. Männer, "Simulation of the continuous curvilinear capsulorhexis procedure," in *Int. Symp. on Biomed. Simul.*, 2006, pp. 113–121.

[38] K. Weber, "Interaktive Echtzeitsimulation deformierbarer Oberflächen für Trainingssysteme in der Augenchirurgie," Ph.D. dissertation, 2009.

[39] S. Sikder, K. Tuwairqi, E. Al-Kahtani, W. G. Myers, and P. Banerjee, "Surgical simulators in cataract surgery training," *Brit. J. of Ophthalmology*, vol. 98, no. 2, pp. 154–158, 2014.

[40] B. Le Gouis, F. Lehericey, M. Marchai, B. Arnaldi, V. Gouranton, and A. Lécuyer, "Haptic rendering of fem-based tearing simulation using clusterized collision detection," in *IEEE World Haptics Conf.*, 2017, pp. 406–411.

[41] J. Allard, M. Marchal, and S. Cotin, "Fiber-based fracture model for simulating soft tissue tearing," in *Medicine Meets Virtual Reality*, 2009, vol. 17, pp. 13–18.

[42] J. Dequidt, H. Courtecuisse, O. Comas, J. Allard, C. Duriez, S. Cotin, E. Dumortier, O. Wavreille, and J.-F. Rouland, "Computer-based training system for cataract surgery," *Simulation*, vol. 89, no. 12, pp. 1421–1435, 2013.

[43] O. Comas, Z. A. Taylor, J. Allard, S. Ourselin, S. Cotin, and J. Passenger, "Efficient nonlinear FEM for soft tissue modelling and its GPU implementation within the open source framework SOFA," in *Biomed. Simul.: Int. Symp.*, 2008, pp. 28–39.

[44] F. Faure, C. Duriez, H. Delingette, J. Allard, B. Gilles, S. Marchesseau, H. Talbot, H. Courtecuisse, G. Bousquet, I. Peterlik, and S. Cotin, "SOFA: A Multi-Model Framework for Interactive Physical Simulation," in *Soft Tissue Biomechanical Modeling for Computer Assisted Surgery*, Jun. 2012, vol. 11, pp. 283–321.

[45] M. Nesme, Y. Payan, and F. Faure, "Efficient, physically plausible finite elements," in *Eurographics*, 2005.

[46] S. Lees and J. D. Heeley, "Density of a sample bovine cortical bone matrix and its solid constituent in various media," *Calcified Tissue International*, vol. 33, pp. 499–504, 1981.

[47] S. Krag and T. T. Andreassen, "Mechanical properties of the human lens capsule," *Progress in Retinal and Eye Research*, vol. 22, no. 6, pp. 749–767, 2003.

[48] R. F. Fisher, "Elastic constants of the human lens capsule," *J. Physiol.*, vol. 201, no. 1, pp. 1–19, Mar. 1969.

[49] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, "Stable-baselines3: Reliable reinforcement learning implementations," *J. of Mach. Learn. Res.*, vol. 22, no. 268, pp. 1–8, 2021.

[50] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015.

[51] P. E. Hart, N. J. Nilsson, and B. Raphael, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," *IEEE Trans. on Syst. Sci. and Cybernetics*, vol. 4, no. 2, pp. 100–107, Jul. 1968.