

Sample-Efficient Learning to Solve a Real-World Labyrinth Game Using Data-Augmented Model-Based Reinforcement Learning

Thomas Bi¹ and Raffaello D’Andrea¹

Abstract—Motivated by the challenge of achieving rapid learning in physical environments, this paper presents the development and training of a robotic system designed to navigate and solve a labyrinth game using model-based reinforcement learning techniques. The method involves extracting low-dimensional observations from camera images, along with a cropped and rectified image patch centered on the current position within the labyrinth, providing valuable information about the labyrinth layout. The learning of a control policy is performed purely on the physical system using model-based reinforcement learning, where the progress along the labyrinth’s path serves as a reward signal. Additionally, we exploit the system’s inherent symmetries to augment the training data. Consequently, our approach learns to successfully solve a popular real-world labyrinth game in record time, with only 5 hours of real-world training data.

I. INTRODUCTION

The labyrinth game, pictured in Figure 1, is a marble game consisting of a maze with walls and holes. The objective is to steer a steel ball from a given start point to the end point by tilting the playfield using two knobs. In doing so, the player must prevent the ball from falling down any of the holes. While it is a relatively straightforward game, it requires fine motor skills and spatial reasoning abilities, and, from experience, humans require a great amount of practice to become proficient at the game. Moreover, the labyrinth presents additional challenges, including stiction effects between the ball and the walls/floor, irregularities in the labyrinth’s surface, and the presence of nonlinear coupling between the control knobs and the labyrinth board. Due to these reasons, we believe that the labyrinth provides an excellent platform to apply and evaluate state-of-the-art robotic learning approaches, and advance our understanding of fine motor control, spatial reasoning, and learning for dynamic systems.

In this paper, we present a method that efficiently learns to solve a popular labyrinth game, the BRIO labyrinth, faster than any previously recorded time (to the best of our knowledge), using only 5 hours of experience collected on the physical system. The following aspects of our method are key to its success. First, observations that efficiently encode the relevant information necessary to infer the optimal strategy are extracted from images captured by a camera positioned above the labyrinth. These observations include the ball position, plate inclination angles, directional information about the path through the labyrinth, and a rectified image patch centered around the current ball position. This image patch



Fig. 1: The labyrinth game is a marble game with the goal of steering a ball from a start to an end position while avoiding letting the ball fall down the holes. Pictured above is the BRIO Labyrinth, introduced almost 80 years ago, and with millions sold.

provides information about the presence of walls and holes in the vicinity of the ball which can be used by the learned policy to efficiently avoid holes and exploit the collision of the ball with walls to redirect the ball. Second, a dense reward function is defined based on the progress through the labyrinth. Lastly, we combine the state-of-the-art model-based reinforcement learning (RL) approach DreamerV3[1] with a data augmentation technique, leveraging the system’s presumed symmetry. Specifically, during training, we randomly mirror the observations about two planes of symmetry, thereby generating more diverse training data and improving the generalization to unseen parts of the labyrinth during training.

A. Related work

In recent years, deep RL has excelled at attaining super-human performance amongst a variety of different games such as chess, Go, and shogi [2], Atari video games [3], 2D and 3D navigation in mazes ([4], [5], [6]), Minecraft [7], and Gran Turismo [8]. However, similar success in systems that require physical interactions with the real world is limited; examples include autonomous drones that have been able to beat human world champions at drone racing [9], and a curling robot that reached the level of top-ranked women’s teams [10]. The few examples that do exist require expensive and bespoke hardware and are thus very difficult

¹ Institute for Dynamic Systems and Control, ETH Zurich, Switzerland, bit@ethz.ch

to reproduce, let alone exploit.

To further study the potential of artificial intelligence in the real world, we believe the labyrinth game to be an ideal testbed due to its cost-efficiency, wide availability, and challenging gameplay. Previous research includes [11], where a model-based path-following controller is employed to solve a real-world labyrinth, whereby the inclination angles of the board are controlled by two servos, and a camera tracks the ball position. In [12], the inverse dynamics of the real-world board are learned using the mixture of several locally weighted linear models, where the parameters of each model are learned online. In order to succeed, the two mentioned works both replace the original labyrinth with custom boards with simpler layouts and attempt to remove any unevenness from the surface. In contrast, the authors of [13] do not make any direct modifications to the game and aim to solve it in its original state. A camera captures the position of the ball and the inclination angles of the board. A Kalman Filter [14] and state-feedback controller are then designed based on a linear ball-and-plate model. The resulting closed-loop system is able to successfully complete the labyrinth with a 75% success rate. Approaches employing RL to solve the labyrinth include [15] and [16]. However, the validity of these approaches was only verified in simulation.

In this work, we exploit recent advances in model-based RL and its ability to make informed decisions about potentially successful behaviours by planning into the future. The effectiveness of learned models has previously been demonstrated by the authors of [17], where world models enable a physical quadrupedal robot to learn to stand up and walk within 1 hour. By leveraging such an approach we obtain a highly performant policy that solves a popular real-world labyrinth game, the BRIO Labyrinth, in record time. In a similar spirit to the work in [13], we do not modify the game itself and use the original labyrinth without modifying any of its dynamics.

B. Outline

The physical setup is presented in Section II. In Section III, the proposed method and its components are described. The performance of the policy learned on the real-world system, as well as an ablation study, are presented in Section IV. Finally, Section V draws conclusions and gives an outlook on future work.

II. HARDWARE

The robotic system considered in this paper consists of four main parts; (i) the labyrinth itself, (ii) two motors that control the inclination angles of the labyrinth plate, (iii) a wide-angle camera capturing a full view of the system, and finally, (iv) a workstation that concurrently runs all the presented algorithms. The system is pictured in Figure 2.

A. Motors

In order to control the orientation of the labyrinth board, the original BRIO labyrinth has two plastic knobs press fitted onto two shafts which control the orientation around




Fig. 2: The physical setup of our system. Two Dynamixel MX-12W motors are directly coupled to the shafts controlling the inclination angles of the labyrinth plate, while a wide-angle USB3 camera captures top-down images of the labyrinth.

the two axes. We remove the knobs, couple two Dynamixel MX-12W motors directly to the two shafts, and mount the motors to the sides of the labyrinth. Furthermore, the motors are daisy-chained and connected to a Dynamixel U2D2 USB converter which allows for commanding the motors using serial communication over a single USB interface. The motors contain internal controllers which are set to velocity control mode.

B. Camera

Next, we employ the wide-angle See3CAM_24CUG camera which provides a USB3 interface. It is placed 20cm above the center of the labyrinth plate and is mounted to the base of the labyrinth through the use of two aluminum slot profiles such that each frame captures a view of the entire labyrinth plate. The camera outputs a 1920x1200 RGB image at 55 frames per second.

C. Computing

The motors and camera are connected to a desktop workstation with an AMD Threadripper Pro 5955WX CPU and 2x RTX4090 GPUs. All computational workload, including sensing, actuating, and learning, is performed by said workstation. All software is written using ROS2 [18] to allow for easy intra-process communication and separation of workloads.

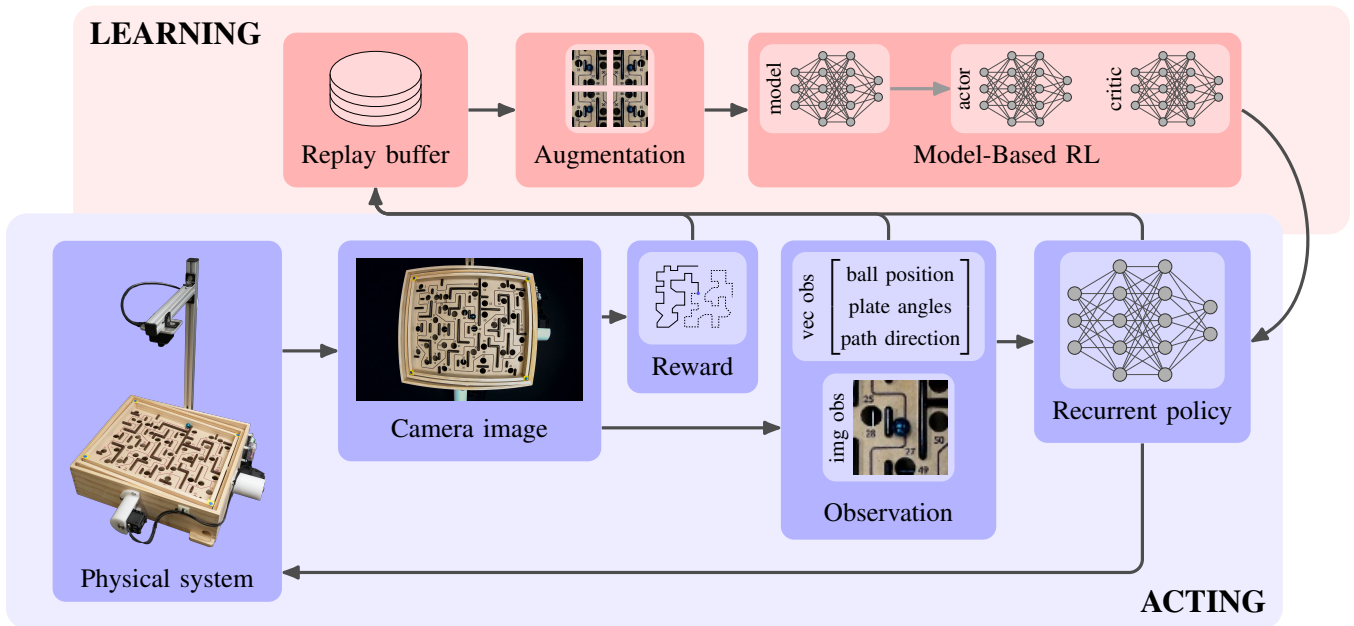


Fig. 3: Overview of the proposed method. On the real-world system, rollouts are collected using the current policy, where observations and reward signals are extracted from the camera images. The collected rollouts are added to the replay buffer using which a model-based RL algorithm aims to optimize a recurrent policy. The learning and acting take place in parallel, and the replay buffer and policy are synchronized between the two processes at the end of each episode.

III. METHOD

We describe our method in three distinct steps. First, the problem of solving the labyrinth is posed as an infinite-horizon optimal control problem under partial observability using a recurrent policy. Second, we detail how relevant observations and reward information are extracted from the camera images. Lastly, we use model-based RL and data augmentation to optimize the learning objective in a sample-efficient manner. Figure 3 shows an overview of our method.

A. Problem Statement

Our aim is to use the presented system to complete the labyrinth as quickly as possible. We pose this problem as an infinite-horizon optimal control problem. Formally, the goal is to find a recurrent policy

$$\pi: (\mathbf{o}(0:k), \mathbf{a}(0:k-1)) \rightarrow \mathbf{a}(k) \quad (1)$$

that is optimal in the sense of maximizing the expected sum of discounted rewards, i.e.,

$$\pi = \max_{\pi} \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r(\mathbf{o}(k), \mathbf{a}(k), \mathbf{o}(k+1)) \right], \quad (2)$$

where γ is the discount factor, k is the discrete time step, $\mathbf{o}(k)$ is the observation at time k , $\mathbf{a}(k)$ is the control action to take at time k , and $r(k)$ is the scalar reward at time k .

As described in Section II, the two employed motors that control the inclination angles of the plate use an internal controller to track angular velocity commands. The control action $\mathbf{a}(k)$ thus consists of the two angular velocity commands for the two motors.

So as to define the reward, we first define a path through the labyrinth by means of defining a sequence of waypoints that are connected linearly. The waypoints are chosen such that the resulting path closely matches the path drawn on the labyrinth board (see Figure 1). Next, we define the reward to be the progress along this path at each time step, i.e.,

$$r(\mathbf{o}(k), \mathbf{a}(k), \mathbf{o}(k+1)) = l(k+1) - l(k), \quad (3)$$

where $l(k)$ is the distance along the path of the point on the path nearest to the ball position at time k . The chosen reward encourages the policy to progress through the labyrinth as quickly as possible. Moreover, it is important to note that the chosen reward does not penalize any deviations of the ball from the path as the optimal trajectory of the ball does not necessarily coincide with the given path.

Additionally, we identify two scenarios that the policy should avoid. First, the ball may not fall down any of the holes. Second, the policy may not cheat by skipping over certain parts of the labyrinth. When either of those two scenarios occurs, the corresponding state is set to be a terminal state for which the sum of future rewards is set to 0.

From the chosen observations we need to be able to recover a Markovian representation of the state. In the absence of any walls or holes, the labyrinth can be interpreted as a ball-and-plate system [19], where the system's state can be defined as

$$\mathbf{x}_{bp} = [x_b \ \dot{x}_b \ y_b \ \dot{y}_b \ \alpha \ \dot{\alpha} \ \beta \ \dot{\beta}]^T, \quad (4)$$

where (x_b, y_b) is the ball's position within the plate frame, and α and β are the inclination angles of the plate along the




Fig. 4: A diagrammatic overview of the labyrinth and the extracted observations.

two fixed axes of rotation (see Figure 4). Since the recurrent policy is conditioned on all past observations and control actions, linear and angular velocities can be recovered from the positional information alone. We thus include only the ball’s position and the inclination angles in the observation. Furthermore, we encode the direction of the path through the labyrinth with $n_p = 5$ equally spaced points along the path in the direction towards the goal starting at the point on the path closest to the current ball position. While two points are theoretically sufficient to encode the path direction, we find that setting $n_p = 5$ yields slight improvements in the sample efficiency of the policy learning. In summary, the low-dimensional vector observations are given by

$$\mathbf{o}_{\text{vec}} = [x_b \quad y_b \quad \alpha \quad \beta \quad \mathbf{p}_{1:n_p}]^T, \quad (5)$$

where $\mathbf{p}_{1:n_p}$ are the n_p points sampled from the path.

To include information about the position of walls and holes at the current position within the labyrinth, we use a top-down view of a $6\text{cm} \times 6\text{cm}$ window of the labyrinth centered at the ball position. The choice of the window size is an exercise in compromise. Too small a window, and insufficient information is retained, too large a window, and irrelevant information gets encoded into the observation, rendering the task of learning an effective policy more difficult. The final observation is then given by the tuple

$$\mathbf{o}(k) = (\mathbf{o}_{\text{vec}}(k), \mathbf{o}_{\text{img}}(k)), \quad (6)$$

where $\mathbf{o}_{\text{img}}(k)$ denotes the view of the labyrinth surrounding the ball.

B. Image Processing

The reward and the observations defined above are extracted from the current camera image at each time step. In a first step, a camera model is obtained using a state-of-the-art calibration technique [20]. The calibration yields the projection function of 3D points to 2D pixel coordinates

(and the inverse), as well as the position and orientation of the camera center w.r.t. the labyrinth base.

To facilitate the estimation process, we (i) use a blue-colored steel ball (same dimension, weight, and material as the original ball), and (ii) place blue markers on the corners of the labyrinth plate. Blue is chosen as it is in great contrast to the mostly green and red dominant color of the labyrinth.

The extraction of the observation $\mathbf{o}(k)$ from a given camera image can then be summarized in the following steps:

- 1) The ball and the markers placed on the corners of the plate are tracked using predictive windows centered on the previous estimate of their positions. Within those windows, the ball and markers are detected by masking the image using hand-tuned hue, saturation, value (HSV) ranges, extracting contours from the masked image, and finally computing the center of mass of the extracted contours.
- 2) Using the pixel coordinates of the tracked corner markers, together with the calibrated camera model and the knowledge of the position of the corners w.r.t. the labyrinth plate, we can estimate the inclination angles α and β of the labyrinth plate using infinitesimal plane-based pose estimation [21].
- 3) Using the pixel coordinates of the tracked ball, together with the pose of the labyrinth plate and the calibrated camera model, an estimate of the ball position (x_b, y_b) in the labyrinth plate frame is obtained.
- 4) Using the ball position, the pose of the labyrinth plate, and the calibrated camera model, the rectified view of the labyrinth layout surrounding the ball (as defined in Section III-A) is extracted. The view is encoded as a 64×64 RGB image.
- 5) From the estimated ball position, the point on the pre-defined path closest to the current ball position is obtained, yielding the current progress $l(k)$ along the labyrinth path.

C. Learning and Data Augmentation

In order to learn the policy, we employ the model-based RL algorithm DreamerV3 [1] with default hyperparameters. There, a world model learns a Markovian representation of the system as well as the system dynamics. The learned world model is used to generate imagined trajectories which are used to train actor and critic networks. Akin to [17], we deploy DreamerV3 in an asynchronous manner; the learning of the world model, the critic, and the actor, is performed in parallel to acting and collecting experience on the real system. At the end of each episode on the physical setup, policy networks are synced, and the collected experience from the episode is added to the replay buffer.

To improve the sample efficiency of the learning process, sequences sampled from the replay buffer during training are augmented by randomly mirroring the trajectories about the two vertical planes that coincide with the two axes of rotation. This includes mirroring both the low-dimensional and the image observations, as well as the control actions. In




Fig. 5: Observations are augmented during training by mirroring about two planes of symmetry.

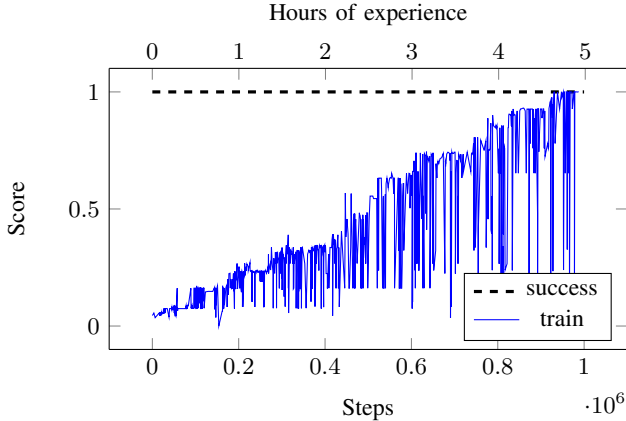


Fig. 6: Single training run on the physical system. A policy that successfully completes the labyrinth is learned within 5 hours. The x -axis represents the number of environment interactions, while the y -axis represents the normalized accumulated reward for each episode.

essence, we obtain trajectories on vertically and/or horizontally flipped views of the labyrinth (see Figure 5), yielding four distinct versions of the labyrinth. We argue that learning from mirrored versions of the labyrinth translates to a more robust policy that is able to generalize to unseen parts of the labyrinth.

IV. RESULTS

The validity of the methods presented is verified by (i) running a single training run on the physical system and evaluating the performance of the policy, and (ii) performing an ablation study using a simulated version of the labyrinth.

The learning of the policy on the real-world labyrinth is conducted with a budget of 1 million time steps (equivalent to 5.05h at a control rate of 55Hz). During learning on the physical system, we manually place the ball back to its starting position by hand whenever an episode terminates. The normalized accumulated reward for each episode during training is plotted in Figure 6. As can be seen, the policy successfully navigates the ball to the end of the labyrinth with less than 5 hours of collected data.

Additionally, we run the resulting policy on the physical setup $N = 50$ times and report the success rate and the average completion time for successful runs in Table I. On average, the policy outperforms the previously fastest recorded time by 0.22s and does so with a 76% success




Fig. 7: A rectified view of the ball trajectory through the labyrinth during a successful run of the final policy.

TABLE I: Comparison between our method and the previous fastest recorded completion time.

Method	Success rate	Completion time
This paper ($N = 50$)	76%	15.73s \pm 0.36s
Human (best) ¹	-	15.95s

rate. Figure 7 illustrates a successful run of the final policy. As can be seen, the policy effectively exploits the walls to quickly redirect the ball. It can also be observed that the final policy navigates the ball in the close vicinity of the holes in order to maximize its performance (such as near holes 10 and 12 in Figure 7). We found that this greediness of the policy leads to failed runs and that training the policy for more than the allocated 1 million time steps increases the robustness and success rate of the policy.

To prove the merits of different aspects of the proposed method, we run ablation studies using a simulated version of the labyrinth; we recreate the system, including the labyrinth, the motors, and the camera, in simulation using the MuJoCo [22] physics engine. Then, three different training settings are assessed in simulation; (i) using only the vector observations as described in Equation (5) without data augmentation, (ii) using both vector and image observations (Equation (6)) without data augmentation, and (iii) the full proposed method including vector and image observations as well as random mirroring of the observations during training. For each scenario, the learning is performed across 16 different seeds, with the same hyperparameters as used on the physical system. The results are shown in Figure 8. As can be seen, the addition of each component leads to a higher final performance given the same number of environment interactions. As mentioned in Section III, we interpret this result as follows. The image patch observations encode the position of the walls and the holes near the ball, helping

¹According to <https://recordsetter.com/world-record/complete-game-labyrinth/12617>.

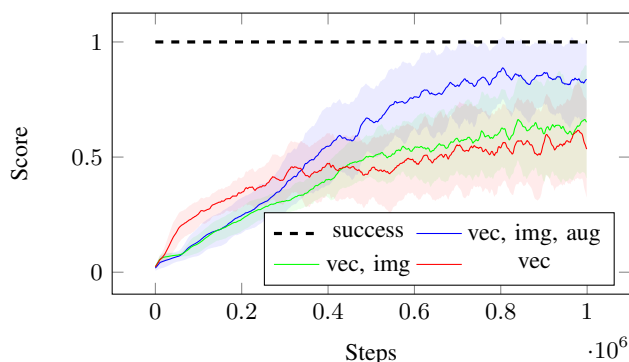


Fig. 8: The mean and standard deviation of the scores averaged over 16 seeds for variations of our method incorporating different components. "vec" denotes the low-dimensional observation \mathbf{o}_{vec} , "img" denotes the use of the image patch \mathbf{o}_{img} , and "aug" denotes the usage of the proposed data augmentation. The addition of the image patch and data augmentation leads to a significant increase in the performance of the policy. The x -axis represents the number of environment interactions, while the y -axis represents the normalized accumulated reward for each episode.

the policy make more informed decisions, while the data augmentation significantly improves sample efficiency by diversifying the training data and improving generalization to previously unseen parts of the labyrinth. While the above-mentioned ablation study could in theory be performed on the physical system, the time investment required to run the three scenarios multiple times in the real world is deemed too costly.

V. CONCLUSION

In this work, we have presented a novel approach to solving a popular and widely-available real-world labyrinth game, leveraging state-of-the-art deep reinforcement learning techniques. Our method demonstrates remarkable proficiency in navigating the labyrinth, achieving completion times faster than any previously recorded. Furthermore, the combination of appropriately chosen observations that encode all relevant information, model-based RL, and data augmentation exploiting inherent symmetries of the system, result in a sample-efficient learning method that is directly employed on the real system.

Looking further ahead, we aim to open-source the hardware and software of the robotic system so as to provide other researchers the opportunity to utilize the labyrinth as a testbed, as we believe it provides a great real-world benchmark with low space requirements, low hardware complexity, and modest cost, for both model-based control as well as reinforcement learning.

ACKNOWLEDGMENT

The authors would like to thank Daniel Wagner, Matthias Mueller, and Tim Flueckiger for their contributions to the development of the system, and Pete Wurman for fruitful discussions.

REFERENCES

- [1] D. Hafner, J. Pasukonis, J. Ba, and T. Lillicrap, "Mastering diverse domains through world models," *arXiv preprint arXiv:2301.04104*, 2023.
- [2] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, *et al.*, "Mastering atari, go, chess and shogi by planning with a learned model," *Nature*, vol. 588, no. 7839, pp. 604–609, 2020.
- [3] Y. Burda, H. Edwards, D. Pathak, A. Storkey, T. Darrell, and A. A. Efros, "Large-scale study of curiosity-driven learning," *International Conference on Learning Representations*, 2019.
- [4] A. Juliani, V.-P. Berges, E. Teng, A. Cohen, J. Harper, C. Elion, C. Goy, Y. Gao, H. Henry, M. Mattar, and D. Lange, "Unity: A general platform for intelligent agents," *arXiv preprint arXiv:1809.02627*, 2020.
- [5] A. Devo, G. Costante, and P. Valigi, "Deep reinforcement learning for instruction following visual navigation in 3d maze-like environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1175–1182, 2020.
- [6] E. Parisotto and R. Salakhutdinov, "Neural map: Structured memory for deep reinforcement learning," *arXiv preprint arXiv:1702.08360*, 2017.
- [7] B. Baker, I. Akkaya, P. Zhokhov, J. Huizinga, J. Tang, A. Ecoffet, B. Houghton, R. Sampedro, and J. Clune, "Video pretraining (vpt): Learning to act by watching unlabeled online videos," *arXiv preprint arXiv:2206.11795*, 2022.
- [8] P. R. Wurman, S. Barrett, K. Kawamoto, J. MacGlashan, K. Subramanian, T. J. Walsh, R. Capobianco, A. Devlic, F. Eckert, F. Fuchs, *et al.*, "Outracing champion gran turismo drivers with deep reinforcement learning," *Nature*, vol. 602, no. 7896, pp. 223–228, 2022.
- [9] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Müller, V. Koltun, and D. Scaramuzza, "Champion-level drone racing using deep reinforcement learning," *Nature*, vol. 620, no. 7976, pp. 982–987, 2023.
- [10] D.-O. Won, K.-R. Müller, and S.-W. Lee, "An adaptive deep reinforcement learning framework enables curling robots with human-like performance in real-world conditions," *Science Robotics*, vol. 5, no. 46, p. eabb9764, 2020.
- [11] E. Frid and F. Nilsson, "Path following using gain scheduled lqr control: with applications to a labyrinth game," 2020.
- [12] K. Öfjäll and M. Felsberg, "Combining vision, machine learning and automatic control to play the labyrinth game," in *Swedish Symposium on Image Analysis*, 2012.
- [13] N. A. Andersen, O. Ravn, and A. T. Sørensen, "Real-time vision based control of servomechanical systems," in *Experimental Robotics II: The 2nd International Symposium, Toulouse, France, June 25–27 1991*. Springer, 1993, pp. 388–402.
- [14] R. E. Kalman, "A new approach to linear filtering and prediction problems," 1960.
- [15] L. Abdenebaoui, E. A. Kirchner, Y. Kassahun, and F. Kirchner, "A connectionist architecture for learning to play a simulated brio labyrinth game," in *KI 2007: Advances in Artificial Intelligence*. Springer, 2007, pp. 427–430.
- [16] J. H. Metzen, E. Kirchner, L. Abdenebaoui, and F. Kirchner, "The brio labyrinth game—a testbed for reinforcement learning and for studies on sensorimotor learning," in *Multidisciplinary Symposium on Reinforcement Learning*, 2009.
- [17] P. Wu, A. Escontrela, D. Hafner, P. Abbeel, and K. Goldberg, "Daydreamer: World models for physical robot learning," in *Conference on Robot Learning*. PMLR, 2023, pp. 2226–2240.
- [18] S. Macenski, T. Foote, B. Gerkey, C. Lalancette, and W. Woodall, "Robot operating system 2: Design, architecture, and uses in the wild," *Science Robotics*, vol. 7, no. 66, p. eabm6074, 2022.
- [19] M. Nokhbeh and D. Khashabi, "Modelling and control of ball-plate system."
- [20] D. Scaramuzza, A. Martinelli, and R. Siegwart, "A toolbox for easily calibrating omnidirectional cameras," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2006, pp. 5695–5701.
- [21] T. Collins and A. Bartoli, "Infinitesimal plane-based pose estimation," *International journal of computer vision*, vol. 109, no. 3, pp. 252–286, 2014.
- [22] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 5026–5033.