

Dynamic Throwing with Robotic Material Handling Machines

Lennart Werner*, Fang Nan*, Pol Eyschen, Filippo A. Spinelli, Hongyi Yang, and Marco Hutter

Abstract—Automation of hydraulic material handling machinery is currently limited to semi-static pick-and-place cycles. Dynamic throwing motions which utilize the passive joints, can greatly improve time efficiency as well as increase the dumping workspace. In this work, we use Reinforcement Learning (RL) to design dynamic controllers for material handlers with underactuated arms as commonly used in logistics. The controllers are tested both in simulation and in real-world experiments on a 12-ton test platform. The method is able to exploit the passive joints of the gripper to perform dynamic throwing motions. With the proposed controllers, the machine is able to throw individual objects to targets outside the static reachability zone with good accuracy for its practical applications. The work demonstrates the possibility of using RL to perform highly dynamic tasks with heavy machinery, suggesting a potential for improving the efficiency and precision of autonomous material handling tasks.

I. INTRODUCTION

Automation of hydraulic machinery has recently gained increased attention in the research domain, driven by its significant relevance to construction, agriculture, forestry, mining and logistics [1], [2], [3]. The primary advantages of automation in these areas include reducing risks for human operators, enhancing accuracy, and addressing labor shortages. However, despite these benefits, the performance of automated systems often falls short when compared to that of skilled human operators, particularly across a variety of tasks.

Bulk material handling is typically carried out by large machines equipped with hydraulic actuators and underactuated joints that connect the gripper to the arm. These machines are commonly used for tasks such as loading, unloading, or sorting. The underactuated design plays a crucial role in enabling the manipulation of materials of various shapes and sizes, allowing human operators to perform agile maneuvers that significantly enhance operational efficiency. As depicted in Fig. 2, human operators can exploit the dynamics of the freely swinging gripper to throw payloads toward specific targets. This skill is achieved by precisely coordinating minimal arm movement with the gripper’s oscillations, allowing operators to target locations that lie beyond the conventional operational range of the arm.

While these maneuvers offer benefits in terms of efficiency and expanded dumping workspace, they are challenging to execute with precision and require considerable practice,

*The authors contributed equally to this work.

The authors are with Robotic Systems Lab, ETH Zurich, 8092 Zurich, Switzerland. Correspondence should be addressed to Lennart Werner lwerner@ethz.ch

This project has received funding from the European Union’s Horizon Europe Framework Programme under grant agreement No 101070405.



Fig. 1. Two-dimensional throwing controller deployed on the M545 excavator with overlaid payload trajectory.

making them extremely difficult for inexperienced operators. The dynamic nature of the throwing motion also introduces a high risk of collisions with obstacles or the machinery itself. Moreover, the intricate dynamics of the hydraulic actuation as well as the passive gripper pose significant challenges to the modeling process, thereby complicating task planning and control. Even small control errors could be magnified by the passive joints and lead to very different end-effector trajectories. Because of these challenges, automation of underactuated material handling machines has rarely been addressed in the past. Previous work [3] presented an RL-based approach to design an anti-swing controller for such machines. Drawing inspiration from it, our research leverages the dynamics of the free-swinging gripper to learn dynamic handling skills, rather than suppressing the oscillation. Such an ability could reduce the cycle time, increase the operational space, reduce energy consumption, and therefore increase the efficiency of the overall process.

A. Contributions

This work presents the first dynamic throwing controller for hydraulic material handling machinery. The controller leverages the underactuated joints of material-handling machines to perform dynamic throwing motions, as illustrated in Fig. 1. With appropriate sim-to-real techniques, the policy, purely trained in simulation, was successfully deployed on a 12-ton material handling test platform. An in-depth accuracy analysis benchmarks the synthesized controllers. Real world experiments confirm the identified precision. The experiments show that the controller is able to perform accurate, repeatable dumping of the payload to targets even outside of the regular machine workspace.




Fig. 2. Human operators using material handling machines to perform throwing in different applications.¹

II. RELATED WORK

A. Robot Throwing

Manipulation has always been a key direction in robotic research, with various applications in factories, warehouses, and households [4]. Instead of static pick-and-place of objects, a number of research works focused on using robotic throwing to increase time and energy efficiency. For example, [5] presents a throwing trajectory planning algorithm for solid waste handling, taking into account robot kinematic constraints. A later method plans throwing trajectories also taking into account the objects' flying dynamics [6]. Through fast online replanning, the system is robust under disturbance. These works however have not been implemented in a real-world experiment.

Besides classical model-based methods, learning-based approaches are also used for robotic throwing. An early work is done in [7], which uses a hierarchical learning framework to train a planar robotic arm to throw balls at one of three fixed targets. Inverse RL has been used in [8] to train a humanoid robot to throw balls at a target based on human feedback. A learning-based method to predict a correction to the throwing velocity was presented in [9], enabling operations with arbitrary objects. Recently, imitation learning has also been used to learn this skill from human demonstrations [10]. Both the last two works showed a remarkable success rate in real-world experiments, but they relied on an existing arm controller to throw the object with the planned release velocity. Reaching a certain throwing velocity at a specific release position, however, is a non-trivial problem for underactuated robot manipulators such as material handling machines.

There have been some studies of throwing with an underactuated robotic arm, inspired mainly by ball-pitching motions of humans [11], [12]. A similar problem has been studied in [13], which focused on a robot arm with a flexible link. The authors showed that flexibility helps the arm to perform fast motion within a short operation time. All of these studies, however, use a model-based approach to design controllers that output torque commands. In real-world applications, such models can hardly capture the full dynamics of the hardware. For this reason, some works address modeling with learning methods [14]. Using a Neural Network (NN) to learn the direct model of the throwing

task, accounting for non-linearities and delays typical of pneumatic soft robots, Bianchi et al. show how an RL policy is able to effectively learn dynamic tasks even when the system is subject to large uncertainties.

B. Automation of Heavy Machinery

Many of the techniques developed in robotic manipulation have been transferred to large scale heavy machines. However, material handling automation has rarely been the subject of research, despite being a crucial component in any construction task since these machines are the most efficient ones to manage equipment and material transport.

Existing works primarily focus on construction, including trenching [15], excavation [16], [17], [18], or drywall stacking [1]. Mining and earthwork, conceptually similar tasks, have also been studied [19], [20]. Farming and forestry applications are presented in literature as well [2], [21].

Among the previously investigated approaches, learning-based methods [16], [22], [3] have been proven suitable to learn and leverage complex machine dynamics while being robust to disturbances under different deployment scenarios. We decided to pursue the same direction to address the control of hydraulic machines with redundant and underactuated kinematics.

Recent advancements in material handling automation show preliminary success. [3] developed an RL controller able to achieve position control of large material handlers with speed and accuracy comparable to human drivers. Their methodology surpasses human capabilities in dampening gripper oscillations, thereby enhancing safety during repetitive working routines. However, this approach doesn't leverage the passive end-effector dynamics. While being necessary for accurate gripper positioning, it prevents the full utilization of the machinery's capabilities and limits operational efficiency. RL has also been used in [21] on a forest forwarder with an underactuated wood log gripper. Their work, despite being tested only in simulation, shows that a learning-based approach can successfully exploit the gripper's underactuated dynamics to complete the log grasping task and even be tuned to reduce energy consumption.

Building upon existing attempts for material handling automation, our work presents a learned controller that uses tool oscillation as an exploitable component to increase machine reach and operational efficiency.

III. METHODOLOGY

We use RL to train the throwing controllers purely in simulation and later deploy them on the real machine. Due to the lack of fast and accurate simulation environments for hybrid hydraulic and rigid-body dynamics, a torque-based robot simulator is used for training. As shown in Fig. 3, a low-level joint velocity controller is introduced in the simulation environment, while in the real world a PID joint velocity controller with feed-forward compensation commands the proportional valves [23]. As the rigid-body dynamics in the training environment can be obtained accurately, we could design a more powerful Inverse Dynamics

¹Picture source: <https://www.youtube.com/watch?v=i46xzaQB50>, <https://www.youtube.com/watch?v=KFxRoSshLGI>

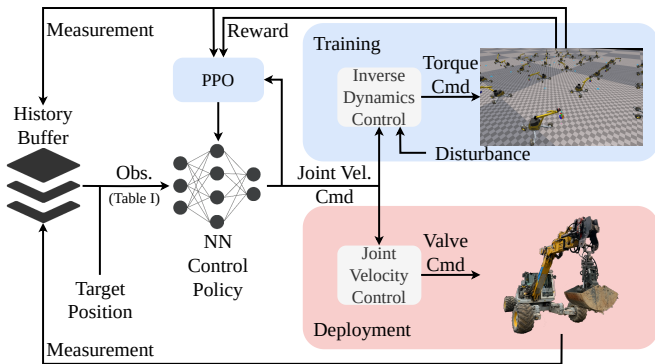


Fig. 3. The pipeline for training and deployment of the controller. The learned throwing controller outputs joint velocity commands. During training, a torque-based simulator is used with a joint velocity controller. During deployment, a low-level controller replaces the joint velocity controller in simulation to command the hydraulic valves.

(ID) low-level controller for the simulated machine. To emulate the delay and inaccuracies of the velocity control loop on the real system, we artificially degrade the control performance in simulation by adding random disturbances to the joint velocity commands. This approach eliminates the need to simulate complex hydraulic system dynamics during training. It significantly decreases training time and has been used in previous works that achieved successful sim-to-real transfers [24], [3]. In this section are described the proposed training setup and necessary real-to-sim and sim-to-real steps for a successful real-world deployment.

We introduce two different control policies for the task of payload throwing. The 3D policy utilizes all joints on the upper carriage of the machine. Combining the motion of cabin turn, boom, dipper, and telescope (Fig. 4 for reference), it can achieve high end-effector velocities. The usability of this controller is constrained by the available free space around the machine. Therefore, we additionally introduce a 2D policy with the limitation of a fixed cabin turn joint. In this case, only in-plane motions of boom, dipper, and telescope are performed for throwing. Both controllers are trained to reach targets at different distances beyond the conventional reach of the arm.

A. System Description

The test platform used in this work is a modified version of the autonomous excavator HEAP [23]. A material-handling gripper is installed, as shown in Fig. 4. The gripper has two passive DoFs around the pitch and roll axes. As indicated in the figure, we only use four actuators on HEAP’s upper body for the throwing task.

A small, battery-powered, wireless IMU module was developed for the purpose of identifying gripper dynamics and passive joint state estimation. The PCB carries an ST ASM330 automotive-grade IMU and an ESP32 for signal filtering, pose estimation, and communication. Rosserial¹ is used to communicate the time-synchronized pose readings to the main computer on the excavator.

¹<http://wiki.ros.org/roscserial>




Fig. 4. The material handling test platform is a 12-ton multipurpose excavator with a material handling gripper. A wireless IMU module is attached to the gripper for state estimation of the passive joints. The DoFs considered in this work are marked in the picture.

B. Simulation for Training

We built the simulation environment in Isaac Gym [25], which performs fast GPU-based physics simulation. Assuming the throwing target is reachable without relocating the machine’s chassis, we only simulate the six movable joints of HEAP. In the training simulation environment, we implement the low-level joint velocity controller with an ID approach.

During training, the excavator’s active DoFs are initialized at random positions without initial velocity, while the tool is always aligned with gravity. The throwing target positions are also randomly sampled at ground level up to 45% farther than the maximum reach of the arm. The horizontal distance from the target to the cabin turn center is constrained to have a minimum value so that targets leading to self-collisions are avoided. During simulation, the RL agent provides joint velocity commands to the ID controller at every step. Payload release is handled by spawning a negligible mass ball at the gripper’s center with the same initial linear velocity. Releasing the payload is delayed from the agent command by a constant time that has been previously identified on the machine, as explained in Section III-F. Subsequently, the ball follows a ballistic trajectory. The release of the payload is only triggered once per rollout. Episodes terminate whenever a self- or ground collision is detected, or when the ball hits the ground.

C. Observations and Actions

The control policy is an Multi-Layer Perceptron (MLP) network with hidden dimension size [256, 128]. The input to the policy network mostly consists of the machine’s proprioceptive information, including its joint positions, joint

velocities, previous commands, and target position. A full list of the 65-dimensional observation space for the 3D policy is given in Table I. Subscripts of symbols in the table indicate joint indices in a serial order of the kinematic chain. Superscripts indicate the time step with k being the current one. The joint states are fed to the policy with a history of three time steps, allowing the agent to infer the gripper dynamics and the behavior of the joint velocity controller. We also provide the previous actions, encouraging the controller to learn consistent and smooth motions. Similar observation vectors have been demonstrated helpful for successful deployment in the real world [16] in order to adapt to low bandwidths and large delays of hydraulic actuators. Additionally, the measurements from a hydraulic machine are usually noisy and the history helps the policy to infer the real joint states.

To train the 2D control policy, we removed the cabin turn joint position, joint velocity, and previous command from the observations, resulting in a 56-dimensional observation space. The same MLP architecture is used for both policies.

TABLE I
OBSERVATIONS OF THE CONTROL POLICY, 3D

Observation	Notation	Dimension
Joint positions	$\mathbf{q}_{1:6}^{k,k-1,k-2,k-3}$	24
Joint velocities	$\dot{\mathbf{q}}_{1:6}^{k,k-1,k-2,k-3}$	24
Previous commands	$\mathbf{u}_{1:5}^{k-1}$	5
Gripper position	\mathbf{r}_{EE}	3
Gripper velocity	$\dot{\mathbf{r}}_{EE}$	3
Gripper position error X-Y plane	err_{2D}	1
Gripper position error 3D	err_{3D}	1
Target position	\mathbf{r}^*	3
Gripper opened (binary)	i_{opened}	1

Our 3D throwing controller outputs a command vector \mathbf{a} in the following form:

$$\mathbf{a} = \left[\hat{q}_1, \hat{q}_2, \hat{q}_3, \hat{q}_4, \hat{q}_{\text{release}} \right]^T, \quad (1)$$

while the 2D policy does not include the cabin turn command \hat{q}_1 . In addition to the joint velocity references $\hat{q}_{1:4}$, the control policy outputs a gripper opening command \hat{q}_{release} . When the command value exceeds a predefined threshold 0.9, it triggers the release of the payload. Furthermore, this action is overwritten by a constant releasing command, with a 5% probability, when the gripper reaches a neighborhood of the target location with low speed. This modification, despite violating the rigorous RL update rule and possibly leading to slower policy updates, allows for easier and better exploration over the discrete action space, helping to escape from local minima.

D. Reward and Terminations

The per-step reward function is in the following form:

$$r = c_1 r_{\Delta_{err}} + c_2 r_{err_{3D}} - c_3 r_{act_{diff}} - c_4 r_{act} \quad (2)$$

with

$$\begin{aligned} r_{\Delta_{err}} &= \max(0, err_{2D}^* - err_{2D}), \\ &\quad + \max(0, err_{3D}^* - err_{3D}), \\ r_{err_{3D}} &= \exp\left(-b_1 \|\mathbf{r}^* - \mathbf{r}_{ball}\|_2^2\right), \\ r_{act_{diff}} &= \|\mathbf{u}_{1:5} - \mathbf{u}_{1:5}^{k-1}\|_2^2, \\ r_{act} &= \|\mathbf{u}_{1:5}\|_2^2, \text{ if } i_{\text{opened}} = 1, \end{aligned}$$

where err_{2D}^* and err_{3D}^* are the minimum X-Y plane and 3D distances from the ball to the target position ever reached during each rollout. All weights b, c are positive. The weights c_1 to c_3 were tuned independently for the two control policies due to the differences in action dimension and scale of position errors. $r_{err_{3D}}$ and $r_{\Delta_{err}}$ are provided at each step to encourage transporting the ball towards the target position. $r_{\Delta_{err}}$ is necessary because a high value of $r_{err_{3D}}$ can be attained by keeping the gripper next to the target position without releasing the ball. With a higher weight on $r_{\Delta_{err}}$, such behavior can be avoided. Over the entire process, the penalization term $r_{act_{diff}}$ incentivizes the policy to perform smooth motions. A very small regularization term r_{act} is used to prevent meaningless actions after the ball is released, preparing the machine for the next sequence. A negative termination reward is given in case of self-collision, ground collision, or reaching joint limits. When the rollout terminates due to the ball dropping to the ground, a positive termination reward in the form of $\exp(b_2 err_{3D}^2)$ is provided to the agent.

E. Training

The control policy is trained using the Proximal Policy Optimization (PPO) algorithm [26] with generalized advantage estimation¹. The simulator runs at 100 Hz, and the control frequency during training is 12.5 Hz. We train with 8192 parallel environments and 4000 iterations with 32 control steps per iteration. The policy training takes about 1.7 hours on a desktop computer with an Intel i9-13900K CPU and an NVIDIA RTX4090 GPU, equivalent to 2.66 years of experience.

F. Sim-to-Real Transfer

Successful execution of the learned motions on hardware requires additional effort to bridge the sim-to-real gap. Besides random noises added to the observation vector and disturbances in the low-level controller, we performed parameter identification on the hardware to remove the repeatable mismatches between the simulation and the real world. Domain randomization was applied to the identified parameters to further robustify the policy against noise and unmodeled dynamics.

¹The implementation can be found at https://github.com/leggedrobotics/rs1_rl




Fig. 5. Identification of the passive joint friction in pitch / Y direction. Comparison of simulated oscillation and measurements from the IMU.

1) *Actuation Delay*: On a hydraulic machine with proportional valves, the performance of actuator control is typically limited by large delays and low bandwidth. The actuation delay of the gripper has been experimentally identified by dropping the wireless IMU from the gripper and recording the time from command to free-fall. It was identified as 258 ms with a standard deviation of 15 ms after 39 releases.

2) *Passive Joint Friction*: Having a good model of friction and damping for the passive joints is critical, as the controller has to move the arm accordingly to excite these joints. The free motion of the gripper is recorded with the wireless IMU, and the data is fitted to the friction model outlined in Eq. (3). $\ddot{\Theta}_{\text{fric}}$ represents the joint acceleration caused by friction and damping effects, v the damping coefficient, $\dot{\Theta}$ the joint velocity and η the velocity independent friction component. A damping coefficient v of 0.03 and a friction coefficient η of 0.1 are identified for our gripper, indicating a dominant velocity-independent friction part. Both passive axes behave similarly. Figure 5 compares the measured motion of a passive joint when releasing from an initial angle and the predicted motion from the identified parameters. The coinciding curves indicate a small sim-to-real gap with the identified model.

$$\ddot{\Theta}_{\text{fric}} = v \cdot \dot{\Theta} + \eta \cdot \text{sign}(\dot{\Theta}) \quad (3)$$

IV. EXPERIMENTS

A. Simulation Experiments

The quantitative throwing performance was primarily evaluated in simulation to gather relevant statistics. Later, a smaller number of real-world experiments was carried out to validate the result. A ball shaped object was used to make precision evaluation in simulation and reality comparable. To avoid benchmarking the trained policies on the seen environment and dynamics, the evaluation was performed in a different simulator, Gazebo.

The Gazebo simulator uses a different, PID-based joint controller which represents the real world joint dynamics with higher fidelity compared to Isaac Gym. Even though this sim-to-sim transfer can only partially emulate the parameter change experienced during real-world deployment, it still challenges the robustness of the controller and allows for a




Fig. 6. Visualization of the evaluation in sim. 3D policy throwing at a target at 9.5 m distance with a static reach of 7.5 m. Target points are indicated in blue, payload trajectory in red with the initial velocity vector on release.

training-independent performance evaluation. The two presented controllers are independently evaluated for precision and later compared for practicality.

1) *Gazebo Simulation Setup*: To benchmark the controllers, the setup from [23] is used. The passive gripper was added to the simulation environment using the same identified friction and damping as described in III-F.

For each evaluation run, the machine’s arm joints get randomly re-initialized at collision-free positions. This ensures, that the measured performance is independent of the initial arm configuration. Target points get sampled along the downrange axis of the machine at a distance from the cabin. We focus on distances at the border and outside of the arm’s static reach as the throwing motion is particularly useful for these targets. Due to the symmetry of the task and training with arbitrary targets, sampling along the downrange axis only is sufficient even for the 3D policy evaluation.

Figure 6 shows the evaluation environment with a randomly selected target point and throwing trajectory from the 3D policy. Here, the cabin turn is initialized nearly 180° away from the target. The 3D controller turns the cabin joint rapidly to swing up the passive gripper and generate a high radial velocity for throwing at the target.

2) *2D Policy*: To evaluate the 2D Policy, the arm of the excavator is aligned with the downrange axis. Target distances from 7.5 m to 10 m in steps of 0.5 m are evaluated. Each target distance is repeated 200 times from different starting configurations. The left plot in Fig. 7 shows the achieved range accuracy with mean impact point and standard deviation.

It is visible, that the controller achieves good accuracy and repeatability up to 9.5 m. Targets close to the conventional reach of the arm are overshoot but the spread of impact points remains low with a 1σ spread of 11 cm and 15 cm for the three close targets.

The learned motion raises the arm to a certain height and uses the stored potential energy to swing the passive pitch joint of the gripper. A sudden stop of the boom joint accelerates the passive joint and the gripper is opened at the correct time. Fig. 1 depicts the motion of the 2D policy




Fig. 7. Visualization of the achieved extended range through agile material handling. The plot on the left side shows the evaluation target accuracy results from simulation for the 2D policy. On the right, the equivalent plot for the 3D policy is shown.

throwing a ball into a bucket.

3) *3D Policy*: Adding cabin turn motion enables the controller to build up kinetic energy without the need to lift the gripper. As shown in Fig. 6, the learned motion accelerates the cabin joint and releases the payload with significant tangential velocity. This way, targets not limited to the downrange axis can be reached.

The simultaneous cabin motion spreads the impacts along the crossrange axis. Fig. 7 on the right plot shows the evaluation from Gazebo for targets between 7.0 m and 9.5 m. In this evaluation, the cabin motion is clockwise for all samples. Impacts for targets up to 8.0 m show a distinct offset in the direction of motion. The release point for these throws is close to the target itself, which increases the sensitivity to inaccurate release timing. Further targets show better precision with an overall similar accuracy of approximately 30 cm standard deviation in both directions for all target distances.

4) *Evaluation Conclusion*: The two learned motions are meant for different initial configurations. While the 2D policy is more precise in all the evaluated distances, it requires previous alignment with the target. Arbitrary starting and target positions can directly be handled by the learned 3D policy. The 3D motion is much more dynamic and might pose a significant risk for objects close to the machine. By using either presented controller, the effective dumping range of the machine can be significantly increased from 7.5 m to 9.5 m as visualized in Fig. 7

B. Real World Experiments

Real-world experiments are conducted to verify the accuracy of the simulation results. Both policies are evaluated with a target at 9 m downrange and 0 m crossrange. Each policy is tested 10 times with randomized initial positions. As a payload, a standard size 7 basketball is used for good visibility. Target impact positions are recorded through photogrammetry. The 2D policy validation setup can be seen in Fig. 1. Figure 8 shows the results of the experiments. Both policies can hit a downrange distance of 9 m with a mean downrange impact location of 8.92 m for the 2D policy, and 8.9 m for the 3D policy. As expected, the 2D policy shows superior crossrange accuracy, showing a mean crossrange




Fig. 8. Results of the real-world evaluation. The crosses represent the impact locations for the 2D and 3D policies. The ellipsoids mark one standard deviation. The target for both policies is denoted by the red dot.

error of 0.05 m compared to the -0.14 m achieved by the 3D policy.

The 2D policy shows standard deviations of 0.37 m and 0.3 m in the downrange and crossrange directions respectively. The observed crossrange error is caused by the cabin slowly drifting during the evaluation run. Since the policy does not command the cabin turn joint, vibrations during the individual throws slowly change the cabin orientation. Furthermore, small differences in the payload grasp, such as a slight off-center grasp of the basketball, also contribute to this error.

While the 3D policy achieves a similar downrange precision with a standard deviation of 0.42 m, the crossrange precision suffers, showing a standard deviation of 0.72 m. As the 3D policy is making full use of the cabin turn joint, this error is expected. Even small differences in the release timing will have a significant effect on the crossrange impact location due to the additional velocity gained by the rotation of the cabin. Differences in the grasp strength of the payload may also introduce an additional release delay, which will have a larger impact on the crossrange error using the 3D policy. Since the achieved standard deviation is similar to the size size of the bucket itself, the expected accuracy

when handling bulk material will not be degraded by the performance of our controller.

V. CONCLUSION AND FUTURE WORK

In this work we show a controller trained with RL that performs dynamic throwing with automated heavy machinery. Through training in simulation, the obtained control policies learned to exploit the underactuated kinematics of the machine to complete throwing tasks. The controllers were successfully deployed in a different simulation environment and in the real world. The learned controllers perform throwing motions in unconstrained environments, extending the reach of the machine with a precision similar to the size of the gripper itself. This accuracy renders our controller not only suitable for dumping material in free space, but also to accomplish common logistic tasks on a construction site. Such throwing controllers could remove the need to dampen the passive joints before material release, leading to a potential boost in efficiency by reducing cycle times. We believe these benefits will significantly improve the productivity of relevant industries and lay the path to fully autonomous machines.

In upcoming research we plan to extend the presented method to tackle more realistic scenarios. We focused on showing the capability to throw with underactuated cranes, but for practical applications, the controller should allow the user to trade-off between more dynamic throws and more accurate ones. Another interesting challenge will be the addition of throwing to 3D targets, e.g., throwing bulk material onto the top of a pile. Moreover, precision throwing can be extended to different payload weights and sizes with necessary adaptability. We aim to include a simulation that mimics real-world dynamics of bulk materials, focusing on the efficient deposition of granular material. State of the art particle simulators can only partially match the needed fidelity and speed to be directly utilized. To enhance this process, we will integrate the presented method with a sophisticated planning algorithm that selects the most effective points for material placement, optimizing the dumping strategy.

ACKNOWLEDGMENT

The authors would like to thank Edo Jelavic and Andreas Dietsche for their valuable support in setting up the simulation and hardware.

REFERENCES

- [1] R. L. Johns, M. Wermelinger, R. Mascaro, D. Jud, I. Hurkxkens, L. Vasey, M. Chli, F. Gramazio, M. Kohler, and M. Hutter, "A framework for robotic excavation and dry stone construction using on-site materials," *Science Robotics*, vol. 8, no. 84, p. eabp9758, 2023.
- [2] E. Jelavić, D. Jud, P. Egli, and M. Hutter, "Robotic Precision Harvesting: Mapping, Localization, Planning and Control for a Legged Tree Harvester," *Field Robotics*, vol. 2, no. 1, pp. 1386–1431, 2022.
- [3] F. A. Spinelli, P. Egli, J. Nubert, F. Nan, T. Bluemer, P. Goegler, S. Brockes, F. Hofmann, and M. Hutter, "Reinforcement learning control for autonomous hydraulic material handling machines with underactuated tools," *Accepted for IROS24*, 2024.
- [4] A. Billard and D. Kragic, "Trends and challenges in robot manipulation," *Science*, vol. 364, no. 6446, p. eaat8414, 2019.
- [5] H. Chen, B. Zhang, and T. Fuhlbrigge, "Robot Throwing Trajectory Planning for Solid Waste Handling," in *2019 IEEE 9th Annual International Conference on CYBER Technology in Automation, Control, and Intelligent Systems (CYBER)*, 2019, pp. 1372–1375.
- [6] Y. Liu, A. Nayak, and A. Billard, "A Solution to Adaptive Mobile Manipulator Throwing," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 1625–1632.
- [7] J. Kober and J. Peters, "Learning elementary movements jointly with a higher level task," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. San Francisco, CA: IEEE, 2011, pp. 338–343.
- [8] Z. Loncarevic, R. Pahic, M. Simonic, A. Ude, and A. Gams, "Learning of Robotic Throwing at a Target using a Qualitative Learning Reward," in *IEEE EUROCON 2019 -18th International Conference on Smart Technologies*. Novi Sad, Serbia: IEEE, 2019, pp. 1–6.
- [9] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. Funkhouser, "TossingBot: Learning to Throw Arbitrary Objects With Residual Physics," *IEEE Transactions on Robotics*, vol. 36, no. 4, pp. 1307–1319, 2020.
- [10] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song, "Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots," *arXiv:2402.10329*, 2024.
- [11] U. Mettin, A. S. Shiriaev, L. B. Freidovich, and M. Sampei, "Optimal ball pitching with an underactuated model of a human arm," in *IEEE International Conference on Robotics and Automation*, 2010.
- [12] S. Katsumata, S. Ichinose, T. Shoji, S. Nakaura, and M. Sampei, "Throwing motion control based on output zeroing utilizing 2-link underactuated arm," in *2009 American Control Conference*. St. Louis, MO, USA: IEEE, 2009, pp. 3057–3064.
- [13] Y. Gai, Y. Kobayashi, Y. Hoshino, and T. Emaru, "Motion Control of a Ball Throwing Robot with a Flexible Robotic Arm," *International Journal of Computer and Information Engineering*, vol. 7, no. 7, pp. 937–945, 2013.
- [14] D. Bianchi, M. G. Antonelli, C. Laschi, A. M. Sabatini, and E. Falotico, "SofToss: Learning to Throw Objects With a Soft Robot," *IEEE Robotics & Automation Magazine*, pp. 2–12, 2023.
- [15] D. Jud, P. Leemann, S. Kerscher, and M. Hutter, "Autonomous Free-Form Trenching Using a Walking Excavator," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3208–3215, 2019.
- [16] P. Egli and M. Hutter, "A General Approach for the Automation of Hydraulic Excavator Arms Using Reinforcement Learning," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5679–5686, 2022.
- [17] L. Terenzi and M. Hutter, "Towards Autonomous Excavation Planning," *arXiv:2308.11478*, 2023.
- [18] S. Jin, Z. Ye, and L. Zhang, "Learning Excavation of Rigid Objects with Offline Reinforcement Learning," *arXiv:2303.16427*, 2023.
- [19] L. Zhang, J. Zhao, P. Long, L. Wang, L. Qian, F. Lu, X. Song, and D. Manocha, "An autonomous excavator system for material loading tasks," *Science Robotics*, vol. 6, no. 55, p. eabc3164, 2021.
- [20] K. Aoshima, D. Lindmark, and M. Servin, "Examining the simulation-to-reality-gap of a wheel loader interacting with deformable terrain," in *11th Asia-Pacific Regional Conference of the ISTVS September 26-28, 2022, Harbin, China*. International Society for Terrain-Vehicle Systems, 2022.
- [21] J. Andersson, K. Bodin, D. Lindmark, M. Servin, and E. Wallin, "Reinforcement Learning Control of a Forestry Crane Manipulator," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 2121–2126.
- [22] M. Lee, H. Choi, C. Kim, J. Moon, D. Kim, and D. Lee, "Precision Motion Control of Robotized Industrial Hydraulic Excavators via Data-Driven Model Inversion," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1912–1919, 2022.
- [23] D. Jud, S. Kerscher, M. Wermelinger, E. Jelavic, P. Egli, P. Leemann, G. Hottiger, and M. Hutter, "HEAP - The autonomous walking excavator," *Automation in Construction*, vol. 129, p. 103783, 2021.
- [24] P. Egli, D. Gaschen, S. Kerscher, D. Jud, and M. Hutter, "Soil-Adaptive Excavation Using Reinforcement Learning," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9778–9785, 2022.
- [25] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, and G. State, "Isaac Gym: High Performance GPU-Based Physics Simulation For Robot Learning," *arXiv:2108.10470*, 2021.
- [26] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal Policy Optimization Algorithms," *arXiv:1707.06347*, 2017.