

Analysis of Randomization Effects on Sim2Real Transfer in Reinforcement Learning for Robotic Manipulation Tasks

Josip Josifovski^{1*}, Mohammadhossein Malmir¹, Noah Klarmann², Bare Luka Žagar¹,
Nicolás Navarro-Guerrero³ and Alois Knoll¹

Abstract—Randomization is currently a widely used approach in Sim2Real transfer for data-driven learning algorithms in robotics. Still, most Sim2Real studies report results for a specific randomization technique and often on a highly customized robotic system, making it difficult to evaluate different randomization approaches systematically. To address this problem, we define an easy-to-reproduce experimental setup for a robotic reach-and-balance manipulator task, which can serve as a benchmark for comparison. We compare four randomization strategies with three randomized parameters both in simulation and on a real robot. Our results show that more randomization helps in Sim2Real transfer, yet it can also harm the ability of the algorithm to find a good policy in simulation. Fully randomized simulations and fine-tuning show differentiated results and translate better to the real robot than the other approaches tested.

I. INTRODUCTION

Current research in reinforcement learning (RL) shows impressive results in simulations, as RL agents win in video games [1], outperform the world champion in the game of Go [2], or develop creative strategies for multi-agent hide-and-seek environments [3]. However, applying such models to control real robots is even more challenging. For example, real-world training of a data-driven model [4] can be impractical or unsafe in many cases where environments are dynamic or computational, and hardware resources are limited, which is often the case in robotics. Simulations seem like a viable alternative for training robots – with the ever-increasing simulation accuracy and computational speedup, virtually unlimited training data is available to train RL agents for robotic tasks. There are many simulation platforms [5], [6] [7], [8], [9] that target robotic simulations both in research and in industrial context. However, simulations are always an approximation of the real environment. Thus, models trained in simulation often fail in the real world, or their performance decreases considerably, which is known as the reality gap problem.

Currently, there are two predominant approaches for minimizing the reality gap: developing highly realistic simula-

tions or using imperfect simulations in combination with randomization strategies [10]. Developing highly realistic simulations depends on processes like system identification, which are often hardware-specific, and the parameters identified with such processes might vary significantly based on conditions in the environment [11]. Furthermore, it is not always easy to determine which real-world phenomena and to what level of precision they should be modeled in highly realistic simulations.

On the contrary, randomization strategies take advantage of the generalization capabilities of neural networks and randomize aspects of the task that are challenging to measure or model precisely [12]. The general assumption, in this case, is that the true real-world parameters are covered in the randomization ranges or that models trained on randomized data learn more robust policies. Arguably, randomization strategies are more applicable in general than learning on highly realistic simulations because the policies learned under randomization might adapt and work in dynamic environments or settings not encountered during training.

However, randomization strategies come at the cost of a more significant number of training samples to learn a task. Moreover, randomization can lead to sub-optimal solutions or even prevent learning, which can be attributed to the inherent unpredictability of the randomized simulations. Therefore, it is essential to quantify the effects of randomization on the learning of policies that would transfer best to the real world. Such systematic studies are still lacking in the field – while there are many successful examples of Sim2Real transfer, they are usually a one-time implementation on a highly-customized system, which hinders experiment reproducibility, benchmarking, or comparative analysis of Sim2Real transfer approaches.

Thus, here we present a comparative analysis of training strategies with different levels of randomization for a robotic reaching task. The chosen task is sufficiently simple to train RL models quickly, while it is still complex enough to capture the effects of the reality gap. Moreover, deterministic solution for the task exists and can be used as a baseline. The simulation environments and code, together with a demonstration video, have been made available to encourage reproducibility at <https://josifovski.github.io/sim2real-randomization-effects/>.

II. BACKGROUND AND RELATED WORK

Training in simulation for robotic control and the reality gap problem have been addressed in many earlier research

* Corresponding Author: Josip Josifovski
josip.josifovski@tum.de

¹ Department of Informatics, Technical University of Munich, Germany.

² Rosenheim University of Applied Sciences, Rosenheim, Germany.

³ Deutsches Forschungszentrum für Künstliche Intelligenz GmbH, Germany.

This work has been financially supported by AI4DI project, which has received funding from the ECSEL Joint Undertaking (JU) under grant agreement No. 826060. The JU receives support from the European Union's Horizon 2020 research and innovation programme, Germany, Austria, Czech Republic, Italy, Latvia, Belgium, Lithuania, France, Greece, Finland, and Norway.

studies [13], [14], [15], [16]. More recently, with the rise of deep learning and the advances in computation, Sim2Real transfer became a common practice in fields like computer vision [12], [17], where synthetic data can be used to train models in simulation completely or to improve their performance when combined with real data. In the last few years, bringing reinforcement learning in combination with deep learning and improved simulations opened the possibility of using Sim2Real transfer effectively in the robotics domain and allowed simulated experiences to serve as a replacement for the hard-to-acquire robotic experience on a real robot. An overview of the Sim2Real approaches in robotics was presented by Zhao et al. [11].

Some of the Sim2Real approaches aim at precise reproduction of the relevant real-world phenomena in simulation [18], [19]. For instance, James and Johns [18] trained a 7-DOF robotic arm in simulation for locating and grasping a cube by adjusting the simulation to resemble the real world as closely as possible. Similarly, Kaspar et al. [19] show successful Sim2Real transfer for a peg in-hole task with a 7-DOF robotic arm by training only in simulation, where they perform system identification to resemble the real robot’s dynamics in simulation accurately.

As matching and correctly simulating all relevant real-world phenomena for a robotic task is often complicated to achieve, other approaches [20], [12], [21], [22] rely on randomization of the critical simulation parameters during model training to increase model robustness and minimize the effects of the mismatch between the simulation and the real world. For example, Peng et al. [20] trained a robotic arm to perform a pushing task by randomizing the physical properties of the robot and the dynamics of the simulation. Similarly, Tobin et al. [12] employed domain randomization to create millions of unrealistically-appearing training images for robust object localization in a robotic grasping task. In the same direction, James et al. [22] combined a domain randomization technique with a generative model to facilitate the learning to grasp unknown objects with a robotic arm. OpenAI [21] successfully transferred policies learned only in simulation to a real robotic hand to solve a Rubic’s cube. They achieved this by automated domain randomization, a form of curriculum learning where the randomization ranges become increasingly wider to make the model learn more general solutions instead of overfitting to a specific randomization seen during training. An overview of randomization techniques in robot learning for Sim2Real transfer was presented by Muratore et al. [23].

Many Sim2Real approaches employ domain adaptation, where the model trained in simulation is fine-tuned on the real robot as a final step to improve its performance. The domain adaptation can also be framed as a continual learning problem, by considering the same robotic task in simulation and the real world as two separate tasks to solve. For example, Rusu et al. [24] trained a model for a reaching task in simulation and then reused the features of the model to train a second model in the real world. They reported that due to the reusability of the skills learned in simulation,

the additional training steps in the real world are drastically decreased.

While domain adaptation is a viable last step, it is always important to minimize the time needed to train on a real robot or circumvent it completely. Thus, it is necessary to measure how different algorithms’ properties or training strategies influence the difficulty to learn a task in simulation and how the performance in simulation relates to the performance on a real robot. In this direction, Kadian et al. [25] developed a complex setup of a real environment and 3D scan it to create a virtual replica in order to measure Sim2Real predictivity for a visual navigation robotic task.

Similarly, here we use a real environment and a virtual replica and quantify the Sim2Real transfer. We also describe how well the simulation performance predicts the performance in a real environment. Unlike Kadian et al. [25], we focus on randomization strategies and their effects on learning a robotic manipulator task, with an accent on environment simplicity in order to mitigate confounding effects and ease experiment reproducibility.

III. METHODOLOGY

We first formally specify a robotic manipulation task suitable for comparing of Sim2Real transfer. Then we describe the compared randomization strategies used. Finally, we explain the evaluation procedure in both simulation and real-world setups.

A. Task Formulation

The task is to learn the differential inverse kinematics of a 2-joint robot, always starting from the same ‘home’ position – the robot is upright with both joints at 0 degrees rotation. The end-effector should be moved to a specific reachable target position by continuously controlling the joint velocities. Typical RL reaching benchmarks and implementations (e.g., [26]–[28]), only consider the goal of reaching the target location. However, there is no consideration for what happens after achieving the goal, which poses potential safety issues when transferring the behavior to the real world. Thus, we require the agent to learn to reach and stay at the specified target until the end of the episode. We call this the *reach-and-balance* task:

$$\mathbf{s}_t = [\Delta x_t, \Delta y_t, \Delta z_t, q_1, q_2] \in \mathbb{R}^5 \quad (1)$$

s_t is the state of the agent at a specific timestep, Δx , Δy and Δz are the distances between the end-effector and the target along Cartesian the corresponding coordinate axis. At the same time, q_1 and q_2 are the joint positions. The actions the agent can perform are specified by:

$$\mathbf{a}_t = [\dot{q}'_1, \dot{q}'_2] \in \mathbb{R}^2 \quad (2)$$

where \dot{q}'_i are the angular velocities for the first and the second joint starting from the robot’s home position. The reward function is defined as:

$$\mathbf{r}_t = \begin{cases} -\mathbf{d}_t & \text{if } C \text{ is false} \\ -\mathbf{d}_t * (\mathbf{T} - \mathbf{t}) & \text{otherwise} \end{cases} \in \mathbb{R} \quad (3)$$

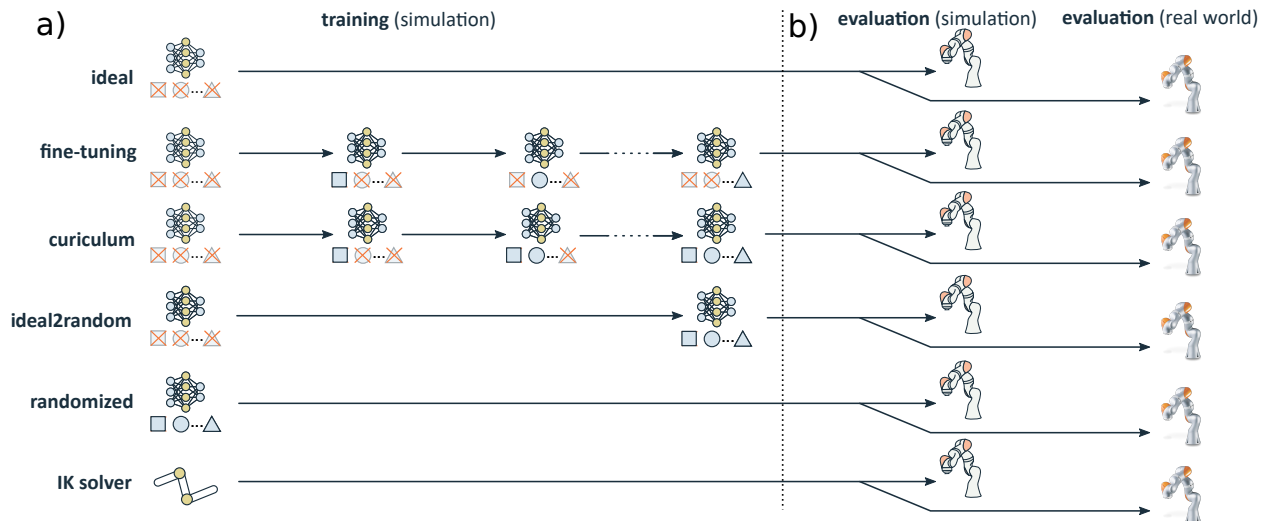


Fig. 1. Methodology a) Different training strategies. Each shape represents a different randomization parameter and whether it is enabled or not during training. b) The simulation-trained model is evaluated under equivalent conditions in simulation and on the real robot.

where \mathbf{d}_t is the Euclidean distance between the target and the end-effector, \mathbf{T} is the episode duration, and C is the terminal condition. During training, C becomes true when the robot collides with the floor. However, during evaluation, C becomes true if any of the joints reach a predefined safety rotation limit. The goal of the agent is to maximize the episode reward:

$$\mathbf{R} = \sum_{t=1}^T \mathbf{r}_t \quad (4)$$

B. Training Strategies

To investigate randomization effects, we selected different randomization strategies based on their conceptual differences, see Figure 1:

- **Ideal simulation** – The simulation environment does not include any randomization parameter relevant for Sim2Real transfer.
- **Fine-tuning** – It uses a model pre-trained on an ideal simulation as the starting point. The model is further trained sequentially for each of the randomization parameters of interest, one at a time.
- **Curriculum learning** – It uses a model pre-trained on an ideal simulation as the starting point. The model is further trained sequentially with an increasing number of randomized parameters.
- **ideal2randomized** – It uses a model pre-trained on an ideal simulation as the starting point. The model is further trained with all the randomized parameters of interest at once.
- **Randomized simulation** – A model is trained from scratch, with all the randomization parameters of interest at once.
- **Inverse kinematics** – It is a deterministic strategy. We use an inverse kinematics solver as a baseline.

IV. EXPERIMENT SETUP

The simulation environment [29] is based on Unity [30] and has a comparable API to OpenAI Gym [5]. It was chosen

because it allows parallel simulation of many robots at once. We used the KUKA LBR iiwa 14 robot manipulator [31] due to its ubiquity in both the research and industrial domain. The simulated robot parameters and meshes are obtained from the URDF data provided by ROS-Industrial [32]. The real robot is controlled via ROS using the IIWA stack [33].

A. Randomization Parameters

We select three parameters that are commonly used in Sim2Real transfer in robotics [20] [34], [11], i.e., Latency (L), Torque (T), and Noise (N).

Latency (L): In an ideal simulation condition, there is no latency. For the randomized case, inspired by Tan et al. [34], we collect a list of ideal observations for every timestep. Then, a random latency value is drawn from a uniform distribution. A linear interpolation of the two consecutive ideal observations closest to the timestep (t') resulting from the current timestep (t) minus the latency is used to determine the latency-affected state of the agent. If the resulting state ($s_{t'}$) precedes the latency-affected state from the previous timestep ($s_{t'-1}$), then the same latency-affected state from $s_{t'-1}$ is used.

Torque (T): The joint velocity control is implemented using the spring-damper equation as explained in Unity Robotics Hub [35]. Here, we randomize the stiffness and damping coefficients for the joints, which results in a randomized torque profile for reaching the commanded velocity, similarly to Peng et al. [20]. On the contrary, the commanded joint velocity is reached instantaneously in an ideal simulation condition.

Noise (N): Random noise drawn from a uniform distribution is added to simulate sensor reading and action noise [20], [34].

The ranges for the randomization parameters can be defined based on system identification, or unrealistic ranges can be used, expecting that the true system values are included within the range. The latter is used because we believe this technique can generalize better and thus more effectively

cope with the Sim2Real transfer. An agent trained in an ideal simulation was used to determine the range of the randomization parameters. The range of the corresponding parameter was determined by increasing the parameter range until the agent’s performance degraded by more than 10%. Table I summarizes the values for all three randomization parameters. Figure 2 shows a visual representation of the randomization parameters.

TABLE I
RANDOMIZATION PARAMETERS.

Param.	Range	Target Domain	Correlation to Ideal Obs.	Effect Freq.
L	[0, 1] sec	Obs.	linear	timestep
T	stiffness: [1, 100] damping: [1, 100]	action	non-linear	episode
N	[0, 10] %	Obs.	uncorrelated	timestep

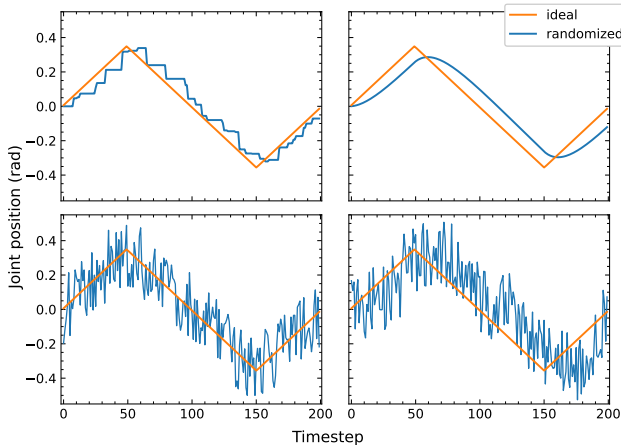


Fig. 2. Effects of different randomization parameters on the observation. Plots show the agent’s observation for the second joint angle after requesting a fixed 20 deg/sec rotational speed for 1 second clockwise, 2 seconds counter-clockwise and 1 second clockwise. Top left: randomized latency. Top right: randomized torque. Bottom left: randomized noise. Bottom right: combined randomization of all parameters.

B. Training Procedure

We used Proximal Policy Optimization (PPO) [36] for all the experiments. PPO is commonly used across many different domains, thus making it suitable for comparative analysis. Another benefit of PPO is that it uses a clipped objective function, which is conservative and prevents large policy changes between two consecutive updates. Consecutive updates are especially useful since the inherent uncertainty in the observations due to randomization might lead to large policy updates based on incorrect observations. We use the implementation from StableBaselines [37], which allows parallelization of the training on many environments simultaneously. The hyperparameters for a similar robotic task, the ReacherBulletEnv-v by Raffin et al. [38] were used. However, after preliminary tests, the clipping range was changed to 0.1 to reduce the variance during randomization training. The number of parallel timesteps per update was also updated from 2048 to 256, as we train on 64 environments in parallel.

Episode length is 250 timesteps, where one timestep is 20ms of simulation time, which amounts to 5 seconds of simulation time per episode during training with a control frequency of 50 Hz.

The max joint speed was set to 1 rad/sec . However, this value can be exceeded when using the Torque randomization parameter.

All models are trained for a total of 40 million timesteps for five different initializations. For each strategy and initialization, the targets were randomly sampled along the half-sphere around the robot, starting from a height of 10 cm above the floor. While the number of total training steps for each run is the same, a different number of training targets might be used since some episodes might terminate earlier due to collisions. Under the ideal strategy the model is trained completely on non-randomized simulation, while for the randomized strategy all parameters are randomized for the whole duration of the training. For the fine-tuning, curriculum, and ideal2randomized strategies, a model pre-trained for 31 million timesteps under the ideal strategy is used as the starting point and trained for additional 9 million timesteps. Under the fine-tuning strategy, the pre-trained model is sequentially trained on three randomized simulations for 3 million timesteps per simulation, where only one randomization parameter is enabled at a time. Similarly, for the curriculum strategy the pre-trained model is further trained on a sequence of three randomized simulations for 3 million timesteps per simulation, where an additional randomization parameter is enabled each time, leading from less to more randomized training. Under the fine-tuning and curriculum strategies the parameter randomization sequence might affect the final performance; hence we test all the permutations of the randomization parameters. Lastly, under the ideal2randomized strategy the pre-trained model on ideal simulation is further trained directly on fully-randomized simulation for 9 million timesteps.

C. Evaluation Procedure

The evaluation procedure is illustrated in Figure 1b. While the training takes place only in simulation, the evaluation takes place both in simulation (ideal simulation setting) and on the real robot, see Figure 3.

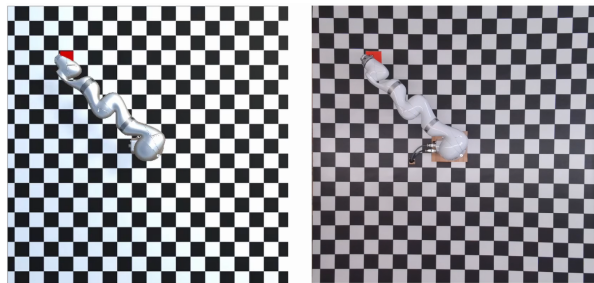


Fig. 3. Left: The simulated environment. Right: The real environment.

The conditions for evaluating the trained model both on the simulated and the real robot are identical. Namely, the robot always starts at the ‘home’ position and the same 50

randomly generated target positions are used in simulation and on the real robot. The same maximal joint speed of $\pi/9 \text{ rad/sec}$ (20 deg/sec) is used. The maximum allowed position for joint 1 is ± 150 degrees and ± 100 degrees for joint 2. Episode duration of 500 timesteps (10 seconds) is used to give enough time to the robot to reach and balance on the farthest targets. However, the episode can also terminate if any joint reaches the predefined safety range, 20 degrees off the corresponding mechanical limit of the joint.

V. EXPERIMENTAL RESULTS

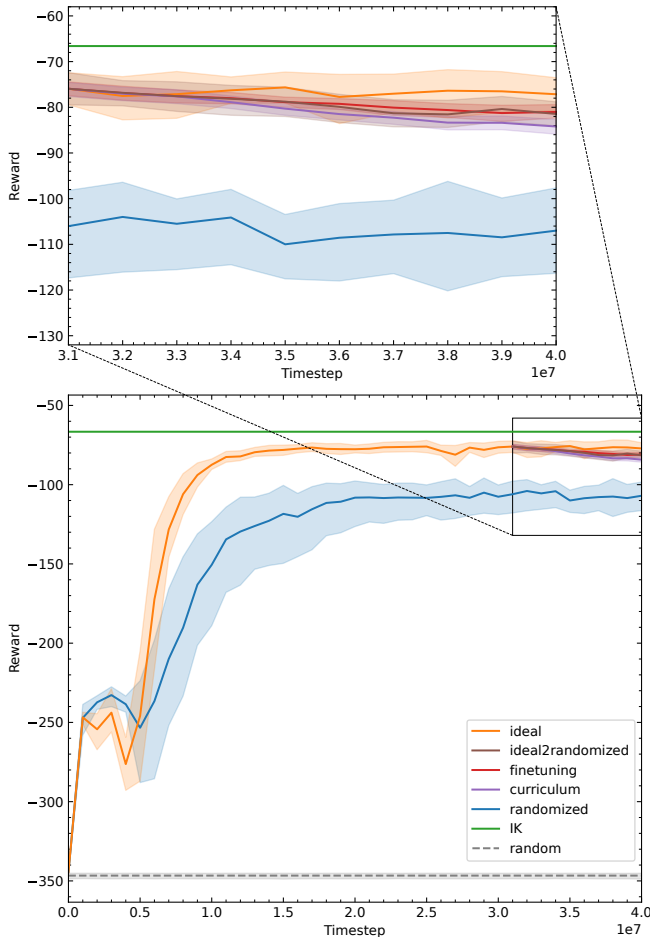


Fig. 4. Training progress for 1 rad/sec measured every 10^6 timesteps with the same random targets. The IK solver in green is the baseline. The shaded area represents 95% confidence intervals.

Figure 4 shows the learning evolution of the different strategies measured as a cumulative reward. The results consist of an average of 5 different random initializations. Except for the fully randomized strategy, all strategies present a small standard deviation and converge to a comparable value, approximately 4% off the IK solver performance w.r.t. random agent. The fully randomized strategy reaches a performance within 14% of the baseline and shows a more substantial standard deviation between different initializations. The rate of improvement of the fully randomized strategy is also lower than the other strategies.

All strategies show a performance decline early in training, starting when the agent collides with the floor. The duration of this performance decline indicates how long each of the randomization strategies takes to avoid collisions.

Table II shows a summary of the results both in simulation and on the real robot. While the fully randomized strategy had the worst performance during training, it experienced the smallest differential between simulation and reality. Moreover, its results on the real robot are within the confidence intervals of the training results. Furthermore, it leads to the best performance on the real robot.

Meanwhile, the real-world performance for all other randomization strategies is substantially different from their simulation results, which the standard deviation values cannot account for.

From the randomization strategies that start with a pre-trained model as a basis, the fine-tuning strategy consistently achieves the best or comparable performance to the other strategies on the real robot regardless of the order of randomization. Only the TLN randomization order leads to substantially worse performance for the fine-tuning strategy, but this result is still within the range of results obtained with the other randomization strategies.

Curriculum learning and ideal2randomized show a substantial difference with their simulation results and lead to worse real-world performance than the ideal simulation condition.

TABLE II
MEAN REWARD FOR DIFFERENT RANDOMIZATION STRATEGIES.

Strategy	Seq.	Avg. in Simulation	Best Model (Simulation)	Best Model (Real)
ideal	N/A	-77.14 ± 4.88	-72.03	-142.50
fine-tuning	N/A	-81.02 ± 4.31	N/A	N/A
	<i>TNL</i>	-79.45 ± 1.58	-76.66	-113.82
	<i>TLN</i>	-83.92 ± 3.81	-80.38	-132.71
	<i>NTL</i>	-84.30 ± 5.91	-77.69	-114.94
	<i>NLT</i>	-78.58 ± 2.72	-74.90	-115.21
curriculum	<i>LTN</i>	-80.98 ± 2.27	-78.74	-119.74
	<i>LNT</i>	-78.89 ± 5.49	-74.55	-113.01
	N/A	-84.21 ± 4.35	N/A	N/A
	<i>TNL</i>	-84.05 ± 3.33	-79.85	-118.85
	<i>TLN</i>	-87.61 ± 6.09	-81.83	-139.47
ideal2random	<i>NTL</i>	-85.55 ± 3.75	-79.40	-154.58
	<i>NLT</i>	-83.74 ± 5.68	-76.99	-129.13
	<i>LTN</i>	-82.59 ± 2.99	-77.67	-171.78
	<i>LNT</i>	-81.71 ± 2.42	-78.99	-159.85
	N/A	-81.46 ± 3.48	-76.54	-167.53
randomized	N/A	-107.00 ± 12.19	-96.01	-109.10
IK solver	N/A	-66.60	-66.60	-77.52

The two best randomization strategies, i.e., full randomization and fine-tuning, were further investigated. The experimental setup was kept identical, except for the maximal joint velocities, which were decreased to $\pi/9 \text{ rad/sec}$ during training to match the joint velocity during testing. The episode length was increased to 500 timesteps to allow enough time for solving the task at a lower speed.

Figure 5 shows that a lower joint velocity makes the problem more challenging because it takes longer to explore the whole search space. However, the same trends for the

1 rad/sec training condition were observed. The fine-tuning strategy converges to a comparable value to the ideal simulation case and is approximately 5% off the baseline w.r.t. the random agent performance. The fully randomized strategy converges to a substantially lower cumulative reward value of approximately 51% off the baseline and at a slower rate than when trained at a speed of 1 rad/sec. All the strategies also show a decline in performance early on in training due to the collision to the floor and take a longer time to recover than the initial setup due to the decreased joint speeds.

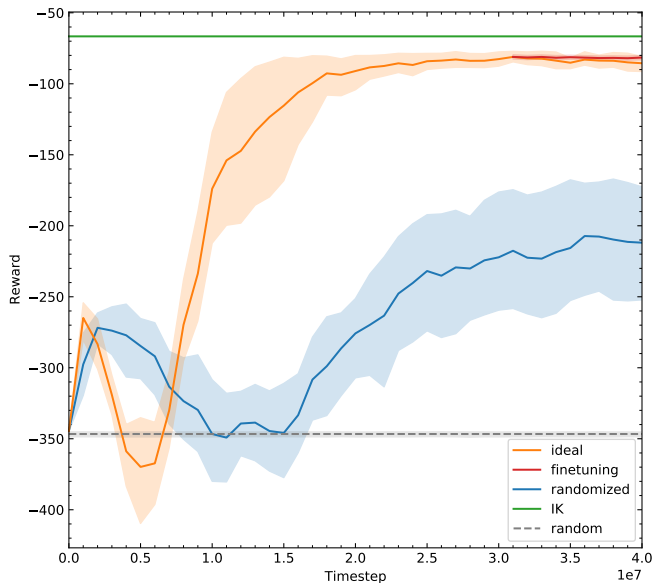


Fig. 5. Training progress for $\pi/9\text{rad/sec}$ joint speed measured every 10^6 timesteps with the same random targets. The IK solver in green is the baseline. The shaded area represents 95% confidence intervals.

Table III summarizes the results for $\pi/9\text{rad/sec}$ joint speed in training for both in simulation and on the real robot.

The performance for the fine-tuning strategy remains practically unchanged both in simulation and on the real robot. In this case, the performance on the real robot also falls outside the confidence intervals determined in the simulations. However, the fully-randomized agents cannot achieve the same performance in either the simulation or the real robot. Moreover, the simulation and the real robot performance substantially deteriorate when training at a lower joint speed. However, the performance on the real robot still falls within the confidence intervals determined in the simulations.

VI. DISCUSSION

In qualitative terms, we observed that the fully-randomized model reached and stabilized near the target location for most of the targets during the evaluation on the real robot. In contrast, the models of the other strategies showed a visible jitter around the target location, which was subjectively worse with the curriculum and the ideal2randomized models. In addition, there was a significant difference in the number of times the episode terminated due to joint limits being reached in the first experiment: 52% for ideal2randomized,

TABLE III
MEAN REWARD FOR THE BEST RANDOMIZATION STRATEGIES UNDER LOWER MAXIMAL JOINT SPEEDS.

Strategy	Seq.	Avg. in Simulation	Best Model (Simulation)	Best Model (Real)
ideal	N/A	-84.16 ± 5.01	-79.05	-129.10
fine-tuning	N/A	-81.71 ± 6.41	N/A	N/A
	<i>TNL</i>	-79.29 ± 3.96	-74.20	-115.17
	<i>TLN</i>	-84.19 ± 12.14	-75.60	-125.63
	<i>NTL</i>	-79.05 ± 5.36	-73.74	-117.33
	<i>NLT</i>	-80.42 ± 4.08	-75.83	-121.82
	<i>LTN</i>	-81.76 ± 3.91	-78.25	-117.41
randomized	<i>LNT</i>	-85.53 ± 5.50	-77.79	-135.18
	N/A	-210.26 ± 47.85	-158.69	-197.22

46% for ideal simulation, and 28.7% for curriculum learning, while only 8% for the fully randomized simulation and 4.33% for the fine-tuning models. For the second experiment on the other hand, when the training scenario was more similar to the evaluation scenario due to the same joint velocity, the model trained under ideal simulation conditions terminated the least due to reaching joint limits (10%), followed by 16% for randomized and 20.3% for the fine-tuning strategy. The qualitative differences are visible in the supplementary video material.

While the randomized strategy was able to find a suitable policy under the first experiment, reducing the joint speeds in combination with the unchanged randomization ranges made the exploration under fully-randomized settings significantly more challenging in the second experiment. The lower joint velocity leads to notably worse model performance both in simulation and on the real robot. This worse performance indicates that strong randomization can reduce the model convergence rate or the ability to find a suitable policy under some conditions.

Interestingly, the fine-tuning strategy performed similarly under both joint velocity conditions. Moreover, it led to overall lower episode terminations due to safety limits when training and testing velocity conditions were different. Therefore, starting from a naive solution in ideal simulation followed by sequential randomization can be an alternative to fully-randomized training for Sim2Real transfer. Especially when the ranges are not based on system identification and might be unrealistic, or many randomization parameters are combined simultaneously.

VII. CONCLUSION

We investigated the randomization effects on Sim2Real Transfer in reinforcement learning for a 2 DoFs robot reaching task. Four different randomization strategies were tested: fine-tuning, curriculum learning, ideal2randomized, and fully randomized simulations. Three commonly used parameters in Sim2Real transfer in robotics were randomized: Latency, Torque, and Sensor and Motor noise.

Curriculum learning and ideal2randomized despite obtaining comparable results in simulation as fine-tuning and the naive ideal simulation, did not transfer well onto the real robot and consistently lead to worse performance than the ideal simulation.

Our results show that fully-randomized simulations can lead to the best real robot performance, although not always. Despite that, we believe it is a promising strategy to bridge the reality gap because the real robot performance consistently falls within the confidence interval obtained in simulations. However, inappropriately intensive randomization might negatively affect the overall performance, even below the ideal simulation condition.

On the contrary, the fine-tuning strategy leads to a consistent real robot's performance regardless of randomization order or the intensity of the randomization. However, the real robot's performance cannot be anticipated based on the simulated results of the fine-tuning strategy.

Finally, additional experiments on the real robot are necessary to determine whether these results are statistically significant.

a) Future Work: We plan to determine the effect size and generalization of our conclusions. We would also like to explore the effect of design decisions on Sim2Real performance, akin to the joint velocity used for training, which can severely impact the performance of some strategies. We want to explore sequential randomization strategies further because they seemingly lead to a more consistent real-world performance but have a larger reality gap than fully-randomized strategies. Finally, we want to work with end-to-end robotic models, including visual input, and analyze visual domain randomization under the same methodology.

REFERENCES

- [1] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing Atari with Deep Reinforcement Learning," in *NIPS: Deep Learning Workshop*, 2013.
- [2] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [3] B. Baker, I. Kanitscheider, T. Markov, Y. Wu, G. Powell, B. McGrew, and I. Mordatch, "Emergent Tool Use from Multi-Agent Autocurricula," in *International Conference on Learning Representations (ICLR)*, ser. Eight, Virtual from Addis Ababa, Ethiopia, 2020.
- [4] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection," *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [5] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "OpenAI Gym," *arXiv:1606.01540 [cs]*, 2016.
- [6] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine, "Meta-World: A Benchmark and Evaluation for Multi-Task and Meta Reinforcement Learning," in *Conference on Robot Learning (CoRL)*, vol. 100. Virtual Event: PMLR, 2020, pp. 1094–1100.
- [7] 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," in *Conference on Neural Information Processing Systems (NeurIPS)*, ser. Datasets and Benchmarks Track, Virtual Event, 2021.
- [8] E. Falotico, L. Vannucci, A. Ambrosano, U. Albanese, S. Ulbrich, J. C. Vasquez Tieck, G. Hinkel, J. Kaiser, I. Peric, O. Denninger, N. Cauli, M. Kirtay, A. Roennau, G. Klinker, A. Von Arnim, L. Guyot, D. Peppicelli, P. Martínez-Cañada, E. Ros, P. Maier, S. Weber, M. Huber, D. Plecher, F. Röhrbein, S. Deser, A. Roitberg, P. van der Smagt, R. Dillman, P. Levi, C. Laschi, A. C. Knoll, and M.-O. Gewaltig, "Connecting Artificial Brains to Robots in a Comprehensive Simulation Framework: The Neurobotics Platform," *Frontiers in Neurobotics*, vol. 11, no. 2, 2017.
- [9] 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 (IROS)*, Vilamoura-Algarve, Portugal, 2012, pp. 5026–5033.
- [10] M. Malmir, J. Josifovski, N. Klarmann, and A. Knoll, "Robust Sim2Real Transfer by Learning Inverse Dynamics of Simulated Systems," in *2nd R:SS Workshop on Closing the Reality Gap in Sim2Real Transfer for Robotics*, Corvallis, OR, USA, 2020, p. 3.
- [11] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-Real Transfer in Deep Reinforcement Learning for Robotics: A Survey," in *IEEE Symposium Series on Computational Intelligence (SSCI)*, Canberra, ACT, Australia, 2020, pp. 737–744.
- [12] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vancouver, BC, Canada, 2017, pp. 23–30.
- [13] N. Jakobi, P. Husbands, and I. Harvey, "Noise and the Reality Gap: The Use of Simulation in Evolutionary Robotics," in *European Conference on Artificial Life (ECAL)*, ser. LNCS, vol. 929. Granada, Spain: Springer, 1995, pp. 704–720.
- [14] N. Jakobi, "Running Across the Reality Gap: Octopod Locomotion Evolved in a Minimal Simulation," in *European Workshop on Evolutionary Robotics (EvoRobots)*, ser. LNCS, vol. 1468. Paris, France: Springer, 1998, pp. 39–58.
- [15] J. C. Zagal, J. Ruiz-del-Solar, and P. Vallejos, "Back to Reality: Crossing the Reality Gap in Evolutionary Robotics," *IFAC Proceedings Volumes*, vol. 37, no. 8, pp. 834–839, 2004.
- [16] S. Koos, J.-B. Mouret, and S. Doncieux, "The Transferability Approach: Crossing the Reality Gap in Evolutionary Robotics," *IEEE Transactions on Evolutionary Computation*, vol. 17, no. 1, pp. 122–145, 2013.
- [17] J. Josifovski, M. Kerzel, C. Pregizer, L. Posniak, and S. Wermter, "Object Detection and Pose Estimation Based on Convolutional Neural Networks Trained with Synthetic Data," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Madrid, Spain, 2018, pp. 6269–6276.
- [18] S. James and E. Johns, "3D Simulation for Robot Arm Control with Deep Q-Learning," in *NIPS Workshop: Deep Learning for Action and Interaction*, Barcelona, Spain, 2016.
- [19] M. Kaspar, J. D. Muñoz Osorio, and J. Bock, "Sim2Real Transfer for Reinforcement Learning Without Dynamics Randomization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Las Vegas, NV, USA, 2020, pp. 4383–4388.
- [20] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-Real Transfer of Robotic Control with Dynamics Randomization," in *IEEE International Conference on Robotics and Automation (ICRA)*, Brisbane, QLD, Australia, 2018, pp. 3803–3810.
- [21] OpenAI, I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, J. Schneider, N. Tezak, J. Tworek, P. Welinder, L. Weng, Q. Yuan, W. Zaremba, and L. Zhang, "Solving Rubik's Cube with a Robot Hand," *arXiv:1910.07113 [cs, stat]*, 2019.
- [22] S. James, P. Wohlhart, M. Kalakrishnan, D. Kalashnikov, A. Irpan, J. Ibarz, S. Levine, R. Hadsell, and K. Bousmalis, "Sim-to-Real Via Sim-to-Sim: Data-Efficient Robotic Grasping Via Randomized-to-Canonical Adaptation Networks," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, USA, 2019, pp. 12 619–12 629.
- [23] F. Muratore, F. Ramos, G. Turk, W. Yu, M. Gienger, and J. Peters, "Robot Learning from Randomized Simulations: A Review," *arXiv:2111.00956 [cs]*, 2022.
- [24] A. A. Rusu, M. Večerík, T. Rothörl, N. Heess, R. Pascanu, and R. Hadsell, "Sim-to-Real Robot Learning from Pixels with Progressive Nets," in *Annual Conference on Robot Learning (CoRL)*, vol. 78. Mountain View, CA, USA: PMLR, 2017, pp. 262–270.
- [25] A. Kadian, J. Truong, A. Gokaslan, A. Clegg, E. Wijmans, S. Lee, M. Savva, S. Chernova, and D. Batra, "Sim2Real Predictivity: Does Evaluation in Simulation Predict Real-World Performance?" in *2nd R:SS Workshop on Closing the Reality Gap in Sim2Real Transfer for Robotics*, Corvallis, OR, USA, 2020, p. 3.
- [26] "Open ai reacher-v2 environment," <https://gym.openai.com/envs/Reacher-v2/>, accessed: 2022-2-23.

- [27] N. Navarro-Guerrero, R. Lowe, and S. Wermter, "Improving Robot Motor Learning with Negatively Valenced Reinforcement Signals," *Frontiers in Neurobotics*, vol. 11, no. 10, 2017.
- [28] C. Stahlhut, N. Navarro-Guerrero, C. Weber, and S. Wermter, "Interaction in Reinforcement Learning Reduces the Need for Finely Tuned Hyperparameters in Complex Tasks," *Kognitive Systeme*, vol. 3, no. 2, 2015.
- [29] J. Josifovski, M. Malmir, N. Klarmann, and A. Knoll, "Continual Learning on Incremental Simulations for Real-World Robotic Manipulation Tasks," in *2nd R:SS Workshop on Closing the Reality Gap in Sim2Real Transfer for Robotics*, Corvallis, OR, USA, 2020, p. 3.
- [30] "Unity 3d," <https://unity.com/>, accessed: 2020-6-26.
- [31] "Kuka lbr-iiwa," <https://www.kuka.com/products/robot-systems/industrial-robots/lbr-iiwa>, accessed: 2020-6-26.
- [32] "Ros industrial," https://github.com/ros-industrial/kuka_experimental, accessed: 2022-2-23.
- [33] C. Hennersperger, B. Fuerst, S. Virga, O. Zettinig, B. Frisch, T. Neff, and N. Navab, "Towards MRI-Based Autonomous Robotic Us Acquisitions: A First Feasibility Study," *IEEE Transactions on Medical Imaging*, vol. 36, no. 2, pp. 538–548, 2017.
- [34] J. Tan, T. Zhang, E. Coumans, A. Iscen, Y. Bai, D. Hafner, S. Bohez, and V. Vanhoucke, "Sim-to-Real: Learning Agile Locomotion for Quadruped Robots," in *Robotics: Science and Systems (R:SS)*, vol. 14, Pittsburgh, PA, USA, 2018.
- [35] "Unity robotics hub," https://github.com/Unity-Technologies/Unity-Robotics-Hub/blob/main/tutorials/urdf_importer/urdf_appendix.md, accessed: 2022-2-23.
- [36] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal Policy Optimization Algorithms," Tech. Rep. arXiv: 1707.06347, 2017.
- [37] A. Hill, A. Raffin, M. Ernestus, A. Gleave, A. Kanervisto, R. Traore, P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, and Y. Wu, "Stable Baselines," 2018.
- [38] A. Raffin, "Rl baselines zoo," <https://github.com/araffin/rl-baselines-zoo>, 2018.