
Megaverse: Simulating Embodied Agents at One Million Experiences per Second

Aleksei Petrenko^{1,2} Erik Wijmans^{1,3} Brennan Shacklett⁴ Vladlen Koltun¹




Figure 1. Megaverse is a new research platform that supports simulating embodied agents in immersive interactive 3D environments at over 1,000,000 experiences per second on a single 8-GPU server. Video demonstrations of Megaverse scenarios and behaviors of trained agents are available at www.megaverse.info

Abstract

We present Megaverse, a new 3D simulation platform for reinforcement learning and embodied AI research. The efficient design of our engine enables physics-based simulation with high-dimensional egocentric observations at more than 1,000,000 actions per second on a single 8-GPU node. Megaverse is up to 70x faster than DeepMind Lab in fully-shaded 3D scenes with interactive objects. We achieve this high simulation performance by leveraging batched simulation, thereby taking full advantage of the massive parallelism of modern GPUs. We use Megaverse to build a new benchmark that consists of several single-agent and multi-agent tasks covering

a variety of cognitive challenges. We evaluate model-free RL on this benchmark to provide baselines and facilitate future research. The source code is available at www.megaverse.info

1. Introduction

Since the advent of deep learning, progress in the field has been driven by two factors: the availability of large amounts of data and the ability to train big and expressive parametric models on this data. Deep reinforcement learning is no exception: the existence of fast, efficient, and challenging simulation environments that can generate millions of trajectories allowed researchers to quickly iterate, experiment with learning at scale, and ultimately develop improved models and learning algorithms.

The majority of RL environments used in contemporary Deep RL research with high-dimensional observations are based on repurposed game engines. For example, the Arcade Learning Environment (ALE) (Bellemare et al., 2013),

¹Intel Labs ²University of Southern California ³Georgia Institute of Technology ⁴Stanford University. Correspondence to: Aleksei Petrenko <petrenko@usc.edu>.

VizDoom (Kempka et al., 2016), and DeepMind Lab (Beattie et al., 2016) use the engines of Atari 2600, Doom, and Quake III, respectively. While reusing existing graphics pipelines reduces development costs, it can be a suboptimal basis for RL environments.

Game engines have been designed and optimized to render a single complex scene at resolutions and framerates tuned for human perception. The simulation of game physics proceeds in very small steps such that the resulting motion and animations appear smooth when the end user interacts with the game world in real time. Environments for RL research, however, have drastically different requirements. State-of-the-art learning algorithms typically consume large amounts of low-resolution observations rendered much faster than real time. Turning a game engine into a high-throughput simulator for RL thus requires using OS-level parallelism, where individual game instances are run at the same time in parallel processes. This prevents the engines from efficiently sharing resources, increases memory consumption, and ultimately fails to fully utilize the throughput of modern hardware accelerators. In addition, due to the high granularity of physics simulation, researchers often use techniques like frame-skipping (Bellemare et al., 2013) to simplify credit assignment, thereby wasting computation on synthesizing views that are never observed by agents.

We argue that a purpose-built rendering and simulation engine for RL research can eliminate inefficiencies of the existing environments and therefore accelerate research in reinforcement learning and artificial intelligence. We present Megaverse, a new platform for embodied AI. Built from scratch as a lightweight simulator, it leverages batched rendering (Shacklett et al., 2021) to fully utilize the throughput of modern GPUs and render up to 1.1×10^6 observations per second on a single 8-GPU node at $128 \times 72 \times 3$ resolution. By default we simulate physics at lower frequency compared to traditional 3D engines, which eliminates the need for frame-skipping and thus enables an experience collection rate up to $\sim 20x$ faster than Atari and $\sim 70x$ faster than DeepMind Lab on hardware commonly found in deep learning research labs.

Since Megaverse was designed to render an arbitrary number of viewpoints and scenes in parallel, it naturally supports multi-agent simulation. While traditional RL engines either do not support multi-agent training (Bellemare et al., 2013) or require a slow network setup to enable it (Kempka et al., 2016), Megaverse can simulate the experiences of dozens of agents in the same environment without loss of performance, in both collaborative and competitive (self-play) settings.

We introduce Megaverse as a simulation platform that can be used to create virtual worlds for deep RL and embodied AI research. We release eight environments built on top of Megaverse that cover an array of embodied cogni-

tive tasks and prove to be hard for modern RL algorithms. Our benchmark, *Megaverse-8*, addresses challenges such as navigation, exploration, and memory. All of the environments are procedurally generated, and can therefore be used to investigate generalization of trained agents (Cobbe et al., 2019). Many of the challenges require our agents to learn nontrivial physics-based environment manipulation, which has previously been a feature of resource-intensive high-fidelity simulators (Kolve et al., 2017).

A key goal of Megaverse is to democratize deep RL research. State-of-the-art results in RL have primarily been a prerogative of large research labs with access to vast computational resources. A fast and efficient simulation engine that supports interactive immersive environments that call for advanced embodied cognition can enable rapid community-wide experimentation and iteration, thus accelerating progress in the field.

2. Prior work

The first artificial agents were typically confined to miniature grid worlds or board games (Tesauro, 1995). These relatively simple environments were nevertheless a challenge for early intelligent systems. They allowed researchers to hone the foundations of reinforcement learning theory and general-purpose learning algorithms (Sutton & Barto, 1992).

With the advent of powerful function approximators, researchers turned their attention to 2D computer games as a new challenge for artificial agents. The DQN algorithm (Mnih et al., 2015) demonstrated the ability to learn directly from high-dimensional pixel observations, matching or exceeding human-level performance on multiple Atari games.

Rapid progress on Atari-like benchmarks led to a phase transition: a new generation of AI research platforms brought *immersive* simulators. In contrast to the flatland of arcade games, the real world is immersive and 3-dimensional. In order to successfully operate in the real world, artificial agents have to master skills such as spatiotemporal reasoning and object manipulation. Simulators derived from first-person 3-dimensional video games, such as VizDoom (Kempka et al., 2016), DeepMind Lab (Beattie et al., 2016), and MineRL (Guss et al., 2019) were among the first to offer virtual embodiment and egocentric perception.

Immersive simulators vary in their fidelity and throughput. Advanced simulation platforms built on top of modern 3D engines, such as Unity (Kolve et al., 2017) or Unreal (Doso-vitskiy et al., 2017), trade simulation speed for high-fidelity graphics. These are useful in studying sim-to-real transfer and perceptual aspects of embodied AI, but end-to-end learning of non-trivial skills and behaviors in these environments requires massive computational resources (Wijmans et al., 2020).

Other efforts focus on increasing the behavioral complexity of simulated worlds. Interactions arising in such environments allow researchers to study non-trivial behaviors and strategies learned by artificial agents. Environments such as Dota 2 (Berner et al., 2019) and StarCraft II (Vinyals et al., 2019) are complex modern strategy games, and playing these games at a human level requires advanced long-term planning. However, these games are not fully immersive, they provide only a structured top-down view of the environment. The OpenAI Hide-and-Seek project (Baker et al., 2020) investigates sophisticated behaviors emerging in an environment with full physics simulation. Unlike the strategy games mentioned above, this environment simulates egocentric perception, although only 1D Lidar-like sensing is supported.

Capture the Flag (Jaderberg et al., 2019), which is based on the Quake III engine, demonstrated great potential of reinforcement learning in immersive environments, but the project relied on complex closed-source compute infrastructure. With Megaverse we aim to build a platform that enables the exploration of advanced embodied cognition at or beyond the level of Capture the Flag and Hide-and-Seek, with full physics simulation and high-dimensional pixel observations, without requiring tens or hundreds of servers for experience collection.

Megaverse is not the first initiative that aims at building an RL simulator from scratch. For example, MINOS (Savva et al., 2017) and its successor Habitat (Savva et al., 2019) use a purpose-built rendering engine and support high-resolution textured scenes based on 3D scans of real environments (Chang et al., 2017). These and other environments can be used as testbeds for embodied challenges such as navigation (Anderson et al., 2018) and rearrangement (Batra et al., 2020). Most related to our work is Shacklett et al. (2021) who demonstrate that considerable speedups can be gained via batched simulation, synthesizing the observations in many environments simultaneously. Whereas Shacklett et al. (2021) only examined very simplistic physics (simple collisions with static geometry), with a single agent per environment, we apply batched simulation to multi-agent environments with complex interactivity.

The idea of batched GPU-accelerated simulation has also been applied in the context of continuous control such as robot locomotion and dexterous manipulation (Liang et al., 2018). Perhaps a combination of batched rendering and batched physics can pave the way for a new generation of fast high-fidelity simulators that will become a core part of future AI research infrastructure.

3. Megaverse

Megaverse is a purpose-built simulation platform for embodied AI research. Our engine can simulate fully immersive 3D worlds with multiple agents interacting with each other and manipulating physical objects, at more than 10^6 experiences per second on a single node. The agents perceive the world through high-dimensional observations, rendered with dynamic shading and simulated lighting.

We introduce a number of performance optimizations that are instrumental in unlocking this performance regime. Our discretized physics approach (Section 3.1) allows us to streamline the computation of non-trivial physical interactions between simulated objects. The parallel architecture built around the batched Vulkan renderer (Shacklett et al., 2021) (Section 3.2) enables the production of hundreds of observations in a single pass, drastically reducing the required amount of communication between hardware components. Another algorithmic optimization is 3D geometry simplification (Section 3.3). The following sections describe these architectural choices and optimizations in more detail.

3.1. Discretized continuous physics

Full simulation of physical contact and collisions between dozens of objects can be computationally expensive. Most physics engines require a small simulation step to avoid unrealistic interpenetration of objects caused by movement interpolation. Megaverse works around this by discretizing some of the physical interactions. Even though movement and collision checking for agents are fully continuous, we simplify the simulation of more complex interactions, such as object stacking. This enables non-trivial object manipulation, *e.g.* building staircases and bridges using physical objects, or moving objects around while solving rearrangement tasks without prohibitively expensive high-frequency simulation of full contact forces.

Specifically, we use a voxel grid data structure: the agents are free to pick up objects anywhere and move them continuously, but placement of objects is only allowed at discretized locations in space, similar to MineRL (Guss et al., 2019). This has two major benefits. First, proximity checks and other spatial queries become trivial $\mathcal{O}(1)$ operations. More importantly, collision checking even with hundreds of objects also becomes extremely fast. We take advantage of caching mechanisms based on axis-aligned bounding boxes implemented in the Bullet physics engine (Coumans & Bai, 2016–2020). Since most of the interactive objects reside in axis-aligned voxels, a simple check based on bounding box intersections eliminates the vast majority of potential collision candidates.




Figure 2. **Megaverse parallel architecture** in the context of a reinforcement learning system. Synchronous or asynchronous training sequentially interacts with $K \geq 1$ instances of Megaverse, each hosting $N \geq 1$ environments with $M \geq 1$ agents. Megaverse parallelizes physics computations on CPU, after which the entire vector of observations is rendered in a single pass.

3.2. Large batch simulation and rasterization

One of the key design features of our implementation is an efficient parallel architecture (Figure 2). A single instance of the traditional game-based RL environment can simulate only a single virtual world and render a single observation per simulation step. One instance of Megaverse can advance hundreds of environments, each containing multiple agents, in parallel. This allows our engine to take full advantage of the massive parallelism of modern computing platforms.

Formally, for environment i our system maintains an internal state $\mathbf{s}_i^{(t)} = [s_{(i,1)}^{(t)}, \dots, s_{(i,N_i^{(t)})}^{(t)}]$, where $s_{(i,n)}^{(t)} \in \text{Sim}(3)^1$ and $N_i^{(t)}$ is the number of entities in the environment i at time step t (*i.e.* interactive objects, walls, agent bodies, cameras, etc.). Then, given a tensor of actions $\mathbf{A}^{(t)} \in \mathcal{A}^{\text{NumEnvs} \times \text{NumAgents}}$, where \mathcal{A} is the set of all available actions, we update the internal state to produce $\mathbf{S}^{(t+1)} = [s_i^{(t+1)}, \dots, s_{\text{NumEnvs}}^{(t+1)}]$. Taken together, the dynamics can be summarized as follows:

$$\mathbf{S}^{(t+1)} = \text{Physics}(\mathbf{S}^{(t)}, \mathbf{A}^{(t)}) \quad (1)$$

$$\mathbf{O}^{(t+1)} = \text{Render}(\mathbf{S}^{(t+1)}) \quad (2)$$

where $\mathbf{O}^{(t+1)}$ is the rendered observations while *Physics* and *Render* are *batched* modules that are responsible for their own parallelization. See Figure 2 for a visual depiction.

Physics simulation is parallelized on the CPU by scheduling state updates for individual environments on a thread pool with a configurable number of threads. In order to parallelize the rasterization step we adopt the optimized batched rendering pipeline proposed by Shacklett et al. (2021). This technique takes advantage of the fact that modern GPUs excel at rendering relatively small numbers of ultra-high-resolution images. The renderer bundles together all the rendering commands corresponding to individual agents and makes a single request to the GPU to render all the observations. This massively cuts down the required amount of communication between CPU and GPU, and helps the renderer achieve high GPU utilization. The impact of this

technique on sampling performance is investigated in Section 5.1. For compatibility with existing training systems, by default we transfer the resulting rendered images to CPU memory, although it is also possible to expose them directly as PyTorch GPU-side tensors (Paszke et al., 2019).

3.3. 3D geometry optimization

Voxelized geometry allows us to speed up physics calculations, but it is not the most efficient way to represent procedurally generated environment layouts, especially in scenarios with non-trivial 3D landscapes. A naive way to visualize such layouts would require rendering thousands of individual voxels that make up the environment. Instead, at the beginning of every episode, after the random landscape is generated, we merge adjacent voxels into a small number of enclosing parallelepipeds (Figure 3). While finding the optimal solution for this problem is NP-hard, a greedy $\mathcal{O}(n)$ algorithm (where n is a total number of non-empty voxels) is sufficient to significantly reduce the number of geometric primitives in the environment. We study the impact of this technique on sampling throughput in Section 5.1.




Figure 3. After procedurally generating the environment layout, we minimize the number of primitives by coalescing adjacent voxels where possible.

4. Megaverse-8 benchmark

Using the Megaverse simulation platform we created a benchmark called *Megaverse-8* (Figure 4), designed for training and evaluation of embodied agents. The bench-

¹The group of all rotations, translations, and scalings in 3D.

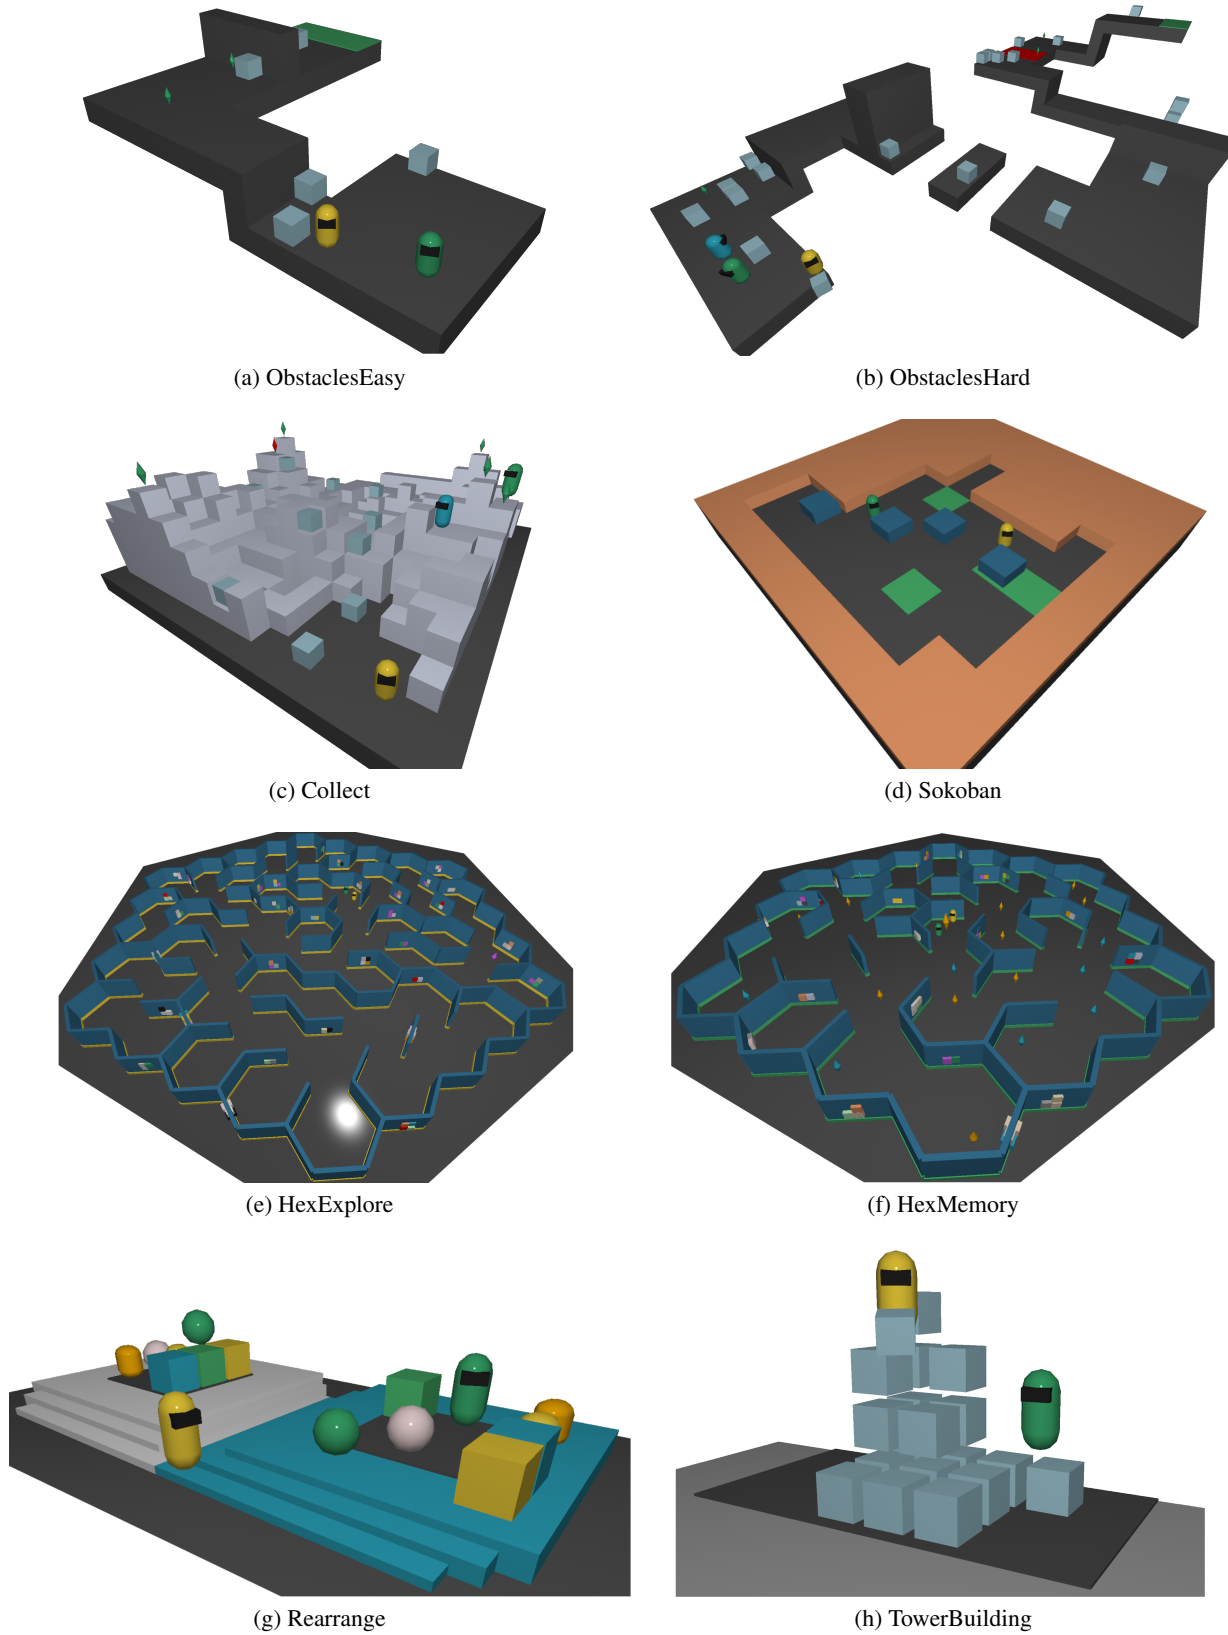


Figure 4. Overview of procedurally generated environments in Megaverse-8. At the beginning of each episode, the visual appearance and layout of environments are sampled randomly.

mark comprises eight tasks, each aiming to test different cognitive abilities of intelligent agents including exploration and navigation in 3-dimensional spaces, tool use and object manipulation, and long-term planning and memory. The benchmark is built with multi-agent support in mind, and all scenarios are suitable for teams of agents.

The environments in the benchmark are relatively simple for human-level intelligence, yet present a serious challenge for artificial agents (Section 5). To score well, agents must demonstrate common-sense comprehension of the physical world, such as understanding object permanence and occlusion, and the ability to interactively manipulate objects. Creating a research platform that provides scenarios that elicit these skills is one of our motivations.

Megaverse-8 environments are procedurally generated and each task has a practically infinite number of instantiations. We randomize the parameters of the task, 3D geometry of the environment, starting positions of interactable objects and agents, *etc.* We randomize the visual appearance by sampling random monochrome materials. Procedural synthesis mitigates overfitting, allows us to evaluate the performance of the agents on unseen environments, and may facilitate the emergence of generalizable skills.

The following paragraphs provide high-level descriptions of the tasks in Megaverse-8. Please refer to the supplementary materials and www.megaverse.info for detailed specifications and video demonstrations.

Megaverse “HexExplore” (Fig. 4e). Agents are placed in a randomized hexagonal maze and tasked to find a target object. The episode is considered solved when the target object is touched by the agent. This environment tests agents’ episodic exploration abilities (Savinov et al., 2019).

Megaverse “Collect” (Fig. 4c). Agents navigate in a procedurally generated 3D landscape and collect objects. Green objects provide positive reward, while red objects generate an equal-magnitude penalty. The episode is considered solved when all positive-reward objects are collected. To generate the level geometry a 2D fractal noise texture (Perlin, 1985) is synthesized and then interpreted as an elevation map at each location in a discretized space. This environment tests agents’ skills in traversing 3D landscapes, including the ability to control the gaze direction both vertically and horizontally.

Megaverse “TowerBuilding” (Fig. 4h). Agents are challenged to construct structures made of interactable boxes. Agents can pick up boxes scattered in the environment and place them in the building zone, which is marked by a distinct color. Agents receive positive reward for placing a new block in the building zone. The reward function grows as $r^{(t)} = 2^h$ with the height h at which the block is placed, thus in order to maximize the score the agents are incen-

tivized to construct structures with as many levels as possible. The building process is subject to realistic constraints: blocks can only be placed on top of other blocks and can only be removed if they have no blocks above them. To build non-trivial tall structures the agents need to maintain scaffolding pathways that allow them to carry the blocks to higher levels.

Megaverse “Sokoban” (Fig. 4d). Fast immersive version of the classic Sokoban puzzle, inspired by the Mujoban environment (Mirza et al., 2020). At the beginning of every episode, a random puzzle with 4 boxes is sampled from the Boxoban dataset (Guez et al., 2018). The agents are required to push the boxes into target positions, marked green. As in classic Sokoban, some of the moves are irreversible, therefore the agents must strategically plan ahead in order to succeed.

Megaverse “HexMemory” (Fig. 4f). Agents are placed in a randomized hexagonal maze in front of a reference object with randomly sampled shape and material. Smaller copies of the reference object are scattered throughout the environment, alongside other objects that do not match the reference. The agent’s task is to collect objects matching the reference, while avoiding other types of objects. When all matching objects are collected, the episode is terminated and the puzzle is considered solved. The scenario requires agents to memorize the visual appearance of an object and keep it in memory for long periods of time, as the reference object inevitably disappears from view as the agent navigates the maze. HexMemory challenges agents’ ability to form and retain memories. This environment is inspired by Beeching et al. (2019).

Megaverse “Obstacles” (Figures 4a and 4b). Procedurally generated 3D obstacle course presented in two versions: ObstaclesEasy and ObstaclesHard. Agents are spawned on one side of the course and are required to reach a target location on the other side. In order to get there they need to overcome different types of obstacles, such as pits, lava lakes, and high walls. Good coordination and movement is not sufficient to overcome most obstacles. For example, a wall can be too tall for the agent to jump over. Agents must use interactive objects placed in the environment to build bridges, staircases, and other artificial tools that help them accomplish the task. ObstaclesEasy and ObstaclesHard differ in both the length of the obstacle course and the difficulty of individual puzzles. ObstaclesHard is particularly difficult due to reward sparsity. As the agents are unlikely to discover the sophisticated construction behaviors by mere random exploration, the obstacle course environments can be a good test for advanced exploration strategies such as intrinsic curiosity (Pathak et al., 2017).

Megaverse “Rearrange” (Fig. 4g). Inspired by the classic

Environment	Simulation throughput	Training throughput
System #1 (12xCPU, 1xRTX3090)		
Atari (84 × 84 grayscale)	19.4k (16.8x)	15.0k (2.8x)
VizDoom (128 × 72 RGB)	38.1k (8.6x)	18.9 (2.3x)
DMLab (96 × 72 RGB)	6.1k (53.5x)	4.6k (9.3x)
Megaverse (128 × 72 RGB)	327k	42.7k
System #2 (36xCPU, 4xRTX2080Ti)		
Atari (84 × 84 grayscale)	47.1K (18.2x)	31.4K (2.9x)
VizDoom (128 × 72 RGB)	79.5K (10.8x)	38.5K (2.3x)
DMLab (96 × 72 RGB)	12.4K (68.7x)	7.7K (11.6x)
Megaverse (128 × 72 RGB)	856K	90.1K
System #3 (48xCPU, 8xRTX2080Ti)		
Atari (84 × 84 grayscale)	53.7K (21.4x)	34.6K (3.9x)
VizDoom (128 × 72 RGB)	100.1K (11.5x)	44.7K (3x)
DMLab (96 × 72 RGB)	63.3K (72.6x)	9.8K (13.7x)
Megaverse (128 × 72 RGB)	1148K	134K

Table 1. Pure sampling and training throughput with mainstream RL simulators vs. Megaverse. The performance is reported in observations per second observed by the agent, i.e. after frameskip.

MIT Copy Demo, this environment challenges the agents to replicate a reference structure made out of colored objects. In order to successfully complete the task the agents have to recognise and remember the object arrangement and replicate it in a designated area by rearranging interactive objects in a specific way. The task is considered solved when the reference arrangement is replicated precisely.

We expect that some of the environments, such as “Rearrange” and “ObstaclesHard”, will be too challenging for present-day end-to-end learning systems. Agents not only need to discover the low-level object manipulation skills, but must also form appropriate internal representations and explore *compositions* of skills in order to succeed.

5. Experiments

5.1. System performance

We start by benchmarking the performance of the Megaverse platform. We examine pure simulation speed, when no inference or learning are done, as well as performance of Megaverse environments as a part of a full RL training system. In order to measure training performance, we use Sample Factory (Petrenko et al., 2020), a fast off-the-shelf RL implementation. We use three different hardware setups that are representative of systems commonly found in deep learning research labs. We compare performance of Megaverse to other fast environments used in reinforcement learning, namely Atari (Bellemare et al., 2013), VizDoom (Kempka et al., 2016), and DMLab (Beattie et al., 2016).

We find that in simulation throughput, Megaverse is an

Optimized geometry	Batched rendering (Shacklett et al., 2021)	Simulation throughput
X	X	20.7K (10.1x)
✓	X	29.6K (7.1x)
X	✓	45.7K (4.6x)
✓	✓	210K

Table 2. Influence of optimized geometry and batched rendering optimizations on the overall sampling throughput. Performance measured on a 10-core 1xGTX1080Ti system in Megaverse-8 “Collect” scenario.

order of magnitude faster than the next fastest environment, VizDoom (Table 1), while supporting considerably more complex interactions. Our platform is between 50x and 70x faster than the most comparable environment, DMLab. In end-to-end training, Megaverse is entirely bottlenecked by learning and inference throughput. However, it still enables training speeds 2-3 times faster than VizDoom and up to 14x faster than DMLab.

Ablation study. We examine the impact of two key performance optimizations in Megaverse: batched rendering (Section 3.2) and geometry optimization (Section 3.3). The results show that both of these techniques are required to achieve high throughput (Table 2). Without geometry optimization the system would be heavily bottlenecked by physics calculations on the CPU, and without batched rendering the communication between CPU and GPU is a major bottleneck.

5.2. Single-agent baseline

Setup. In this section we present RL training results on the Megaverse-8 benchmark. We train agents using asynchronous proximal policy optimization (PPO) (Schulman et al., 2017) with V-trace off-policy correction (Espeholt et al., 2018) using the Sample Factory implementation (Petrenko et al., 2020). Given the challenges of learning good representation with model-free RL from scratch, we also experiment with using Action Conditional Contrastive Predictive Coding (CPCIA) (Guo et al., 2018) as an auxiliary loss. We train both standard PPO and a CPCIA-augmented version on 2×10^9 environment steps. We find that CPCIA augmentation leads to considerable performance improvements on TowerBuilding and Exploration tasks without significantly affecting other scenarios, therefore we decided to use it in all other experiments.

Results. We establish that the proposed benchmark has considerable diversity in task difficulty. While all tasks are far from being solved, reasonable progress can be made on four of the eight (Figure 5). ObstaclesEasy, Collect, and HexExplore require robust 3D navigation skills and basic object manipulation abilities. Model-free RL was able to achieve non-trivial performance in these scenarios, although

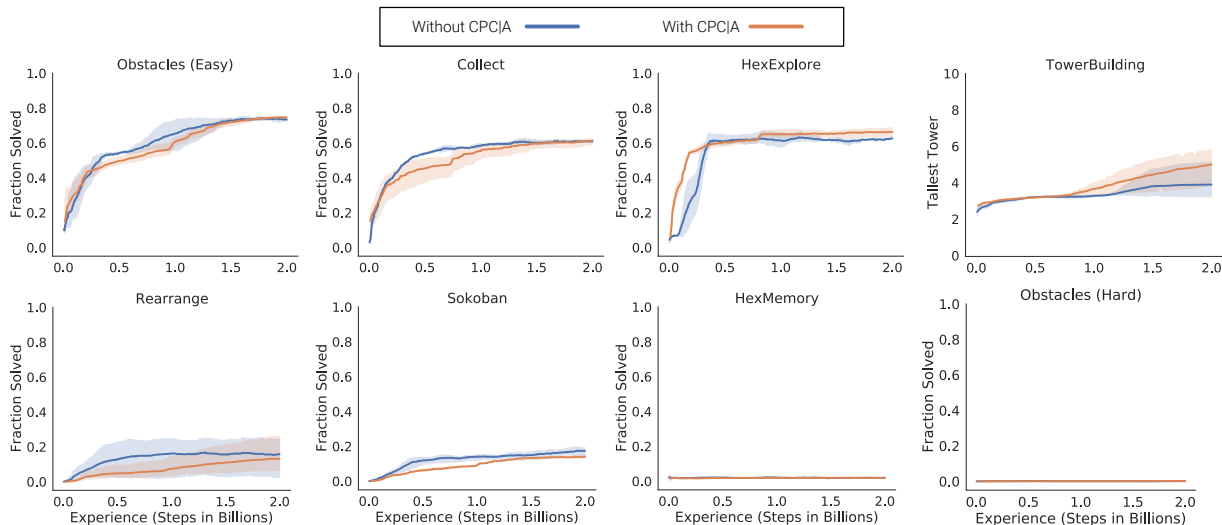


Figure 5. **Single-agent performance.** Megaverse-8 has a diversity of challenges, with environments where model-free RL is able to make progress (top four) and environments where it fails to achieve non-trivial performance (bottom four). Even for the simplest environments, there is still considerable room for improvement. The results are reported for 3 random seeds.

none of the agents approached 100% success.

Our agents demonstrated surprisingly high level of performance in the TowerBuilding scenario. Comparatively dense reward allowed the agents to master object stacking and consistently construct structures up to ten levels high. Video demonstrations of agent performance can be found at www.megaverse.info

Other scenarios in the benchmark have proven to be a much harder challenge. Both Rearrange and Sokoban agents improved at the beginning of training, but ultimately failed to reach satisfactory performance levels. In the Rearrange scenario, the agents learned to randomly shuffle the objects in the hope of matching the target arrangement by accident, and never learned to pay attention to the reference arrangement. In Sokoban, the agents tend to push blocks to nearby targets which happens to be sufficient to solve some of the puzzles. Full completion of the task requires long-term planning, and the agents ultimately failed to demonstrate the ability to do that.

HexMemory and ObstaclesHard turned out to be the most challenging scenarios. HexMemory is a hard credit assignment and memory challenge. Agents failed to capture the relationship between temporally distal events, such as observation of the reference object and collection of similar/dissimilar objects in the environment. This experiment shows that GRU policy networks that we used in our experiments are not sufficient for this type of task, although there is potential for other policy architectures, such as transformers (Parisotto et al., 2020).

ObstaclesHard is perhaps the most challenging scenario in the Megaverse-8 benchmark. To traverse the obstacle

course completely the agents need to master multiple skills and combine them in intelligent ways to overcome obstacles. Learning individual skills via a curriculum of simpler environments may be a promising research direction.

5.3. Multi-agent baseline

Setup. We continue by examining multi-agent performance with two and four agents on our proposed benchmark. In all Megaverse-8 tasks, the agents must work together to perform well. To encourage cooperation, we experiment with Team Spirit reward shaping, inspired by OpenAI Dota 2 experiments (Berner et al., 2019). Team Spirit modifies the credit assignment such that agents are rewarded both for their own actions and the actions of other agents in the team. Formally, the reward for agent i as time t is $r_i^{(t)} = (1 - \text{TeamSpirit})\tilde{r}_i^{(t)} + \frac{\text{TeamSpirit}}{\text{NumAgents}} \sum_j \tilde{r}_j^{(t)}$, where $\tilde{r}_i^{(t)}$ is the individual agent rewards before incorporating Team Spirit. This reward makes credit assignment harder, thus we gradually increase Team Spirit from 0.0 to 1.0 over the first one billion steps of training as a form of curriculum.

Results. For two tasks, HexExplore and Collect, we find that having more agents is beneficial (Figure 6). In these tasks, relatively high score can be achieved even if agents ignore each other and focus on maximizing their own reward. This is confirmed by the fact that Team Spirit hurts performance. For HexMemory, a team of agents has a higher chance of completing the task randomly. Even though the results are better than for a single agent, teams of agents fail to make progress. In TowerBuilding, we discover that two agents perform approximately as well as one agent, and four agents consistently perform worse. The agents end up

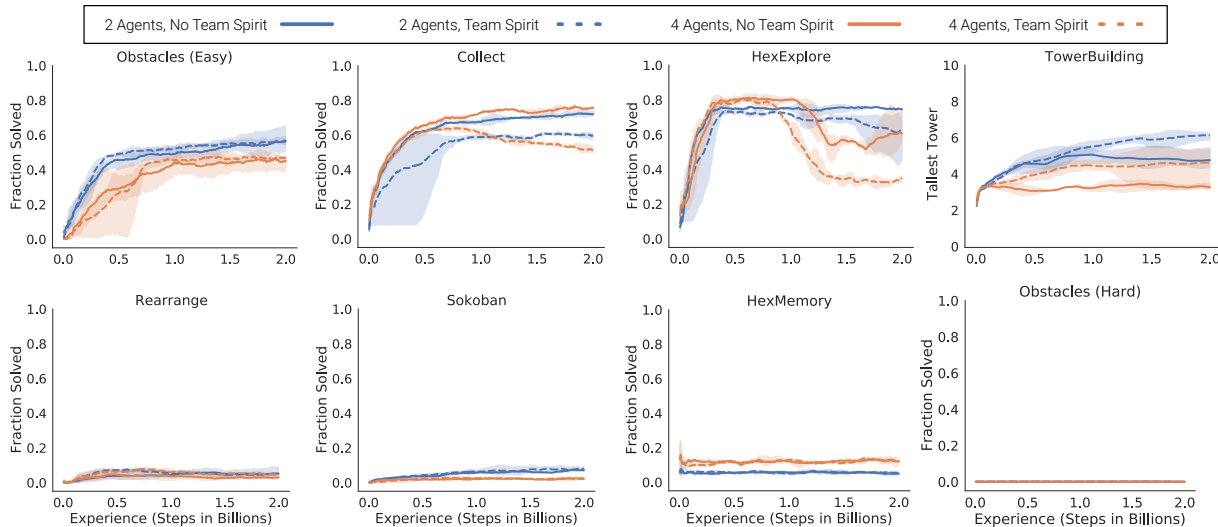


Figure 6. **Multi-agent performance.** For the majority of tasks in Megaverse-8, increasing the number of agents does not yield better results with our model-free RL framework. For all tasks except tower building, a simple Team Spirit strategy of sharing rewards has either no impact or negative impact on performance due to the increased difficulty of credit assignment. The results are reported for 3 random seeds.

competing for rewards with each other instead of working together. In this case the addition of Team Spirit encourages cooperation and improves performance in both two- and four-agent teams.

6. Discussion

We presented a new research platform *Megaverse*, capable of achieving simulation throughput over 1,000,000 observations per second on a single node – unprecedented for immersive 3D simulators. Aside from the ability to simulate embodied agents tens of thousands of times faster than real time, our engine can also match the throughput of existing simulators while using only a fraction of computational resources. Our dedicated simulation platform can make large-scale RL experiments more accessible, thus accelerating progress in AI research.

Hard problems and good metrics for evaluating progress on these problems are instrumental for science. While the traditional benchmarks in Deep RL are definitely not trivial, they are getting tantalizingly close to being solved (Badia et al., 2020; Parisotto et al., 2020; Petrenko et al., 2020). We use the Megaverse platform to build *Megaverse-8*, a new suite of hard challenges for embodied AI. Complete solution of all tasks in Megaverse-8 requires the agents to master object manipulation, rearrangement, and composition of different low-level skills. We hope that solving these challenges in a robust and principled way will advance our understanding of embodied intelligence.

Extremely fast simulation provided by Megaverse can have impact beyond deep reinforcement learning. For exam-

ple, contemporary derivative-free optimization methods are known for their supremacy in Mujoco-like environments (Sener & Koltun, 2020), but evaluating them in scenarios with high-dimensional observations has previously been very costly. With more than an order-of-magnitude improvement in simulation throughput, evaluation of derivative-free methods in immersive 3D environments may be feasible.

Another research direction that can leverage fast simulation is meta-learning. With highly optimized learning systems (Stooke & Abbeel, 2019; Espeholt et al., 2019; Petrenko et al., 2020), entire training sessions in simple Megaverse environments can be completed in mere seconds, and thus can be used as a part of a larger meta-learning process. While partial learning of optimizer features and loss functions has been demonstrated (Bechtle et al., 2021), access to sufficiently fast training may enable the optimization of whole learners parameterized by neural networks.

Megaverse opens new possibilities in the field of multi-agent learning. Megaverse is one of the first open-source platforms that allows fast simulation of multiple agents interacting in immersive environments. The accessibility of such a platform can have important implications for studying multi-agent cooperation, autotricula emerging from self-play (Baker et al., 2020), and the emergence of communication and language (Mordatch & Abbeel, 2018).

References

Anderson, P., Chang, A. X., Chaplot, D. S., Dosovitskiy, A., Gupta, S., Koltun, V., Kosecka, J., Malik, J., Mottaghi, R., Savva, M., and Zamir, A. R. On evaluation of embodied

- navigation agents. *arXiv:1807.06757*, 2018.
- Badia, A. P., Piot, B., Kapturowski, S., Sprechmann, P., Vitvitskiy, A., Guo, Z. D., and Blundell, C. Agent57: Outperforming the atari human benchmark. In *ICML*, 2020.
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., and Mordatch, I. Emergent tool use from multi-agent autotutorials. In *ICLR*, 2020.
- Batra, D., Chang, A. X., Chernova, S., Davison, A. J., Deng, J., Koltun, V., Levine, S., Malik, J., Mordatch, I., Mottaghi, R., Savva, M., and Su, H. Rearrangement: A challenge for embodied AI. *arXiv:2011.01975*, 2020.
- Beattie, C., Leibo, J. Z., Teplyaev, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., Schrittwieser, J., Anderson, K., York, S., Cant, M., Cain, A., Bolton, A., Gaffney, S., King, H., Hassabis, D., Legg, S., and Petersen, S. DeepMind Lab. *arXiv:1612.03801*, 2016.
- Behtle, S., Molchanov, A., Chebotar, Y., Grefenstette, E., Righetti, L., Sukhatme, G., and Meier, F. Meta learning via learned loss. In *International Conference on Pattern Recognition*, 2021.
- Beeching, E., Wolf, C., Dibangoye, J., and Simonin, O. Deep reinforcement learning on a budget: 3D control and reasoning without a supercomputer. *arXiv:1904.01806*, 2019.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. In *IJCAI*, 2013.
- Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., de Oliveira Pinto, H. P., Raiman, J., Salimans, T., Schlatter, J., Schneider, J., Sidor, S., Sutskever, I., Tang, J., Wolski, F., and Zhang, S. Dota 2 with large scale deep reinforcement learning. *arXiv:1912.06680*, 2019.
- Chang, A. X., Dai, A., Funkhouser, T. A., Halber, M., Nießner, M., Savva, M., Song, S., Zeng, A., and Zhang, Y. Matterport3D: Learning from RGB-D data in indoor environments. In *3DV*, 2017.
- Cobbe, K., Hesse, C., Hilton, J., and Schulman, J. Leveraging procedural generation to benchmark reinforcement learning. *arXiv:1912.01588*, 2019.
- Coumans, E. and Bai, Y. PyBullet, a Python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2020.
- Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., and Koltun, V. CARLA: An open urban driving simulator. In *Conference on Robot Learning (CoRL)*, 2017.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., Legg, S., and Kavukcuoglu, K. IMPALA: Scalable distributed Deep-RL with importance weighted actor-learner architectures. In *ICML*, 2018.
- Espeholt, L., Marinier, R., Stanczyk, P., Wang, K., and Michalski, M. SEED RL: Scalable and efficient Deep-RL with accelerated central inference. *arXiv:1910.06591*, 2019.
- Guez, A., Mirza, M., Gregor, K., Kabra, R., Racaniere, S., Weber, T., Raposo, D., Santoro, A., Orseau, L., Eccles, T., Wayne, G., Silver, D., Lillicrap, T., and Valdes, V. An investigation of model-free planning: Boxoban levels. <https://github.com/deepmind/boxoban-levels/>, 2018.
- Guo, Z. D., Azar, M. G., Piot, B., Pires, B. A., and Munos, R. Neural predictive belief representations. *arXiv:1811.06407*, 2018.
- Guss, W. H., Houghton, B., Topin, N., Wang, P., Codel, C., Veloso, M., and Salakhutdinov, R. MineRL: A large-scale dataset of Minecraft demonstrations. In *IJCAI*, 2019.
- Jaderberg, M., Czarnecki, W. M., Dunning, I., Marris, L., Lever, G., Castañeda, A. G., Beattie, C., Rabinowitz, N. C., Morcos, A. S., Ruderman, A., Sonnerat, N., Green, T., Deason, L., Leibo, J. Z., Silver, D., Hassabis, D., Kavukcuoglu, K., and Graepel, T. Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science*, 364(6443), 2019.
- Kempka, M., Wydmuch, M., Runc, G., Toczek, J., and Jaskowski, W. Vizdoom: A Doom-based AI research platform for visual reinforcement learning. In *IEEE Conference on Computational Intelligence and Games*, 2016.
- Kolve, E., Mottaghi, R., Han, W., VanderBilt, E., Weihs, L., Herrasti, A., Gordon, D., Zhu, Y., Gupta, A., and Farhadi, A. AI2-THOR: An interactive 3D environment for visual AI. *arXiv:1712.05474*, 2017.
- Liang, J., Makovychuk, V., Handa, A., Chentanez, N., Macklin, M., and Fox, D. GPU-accelerated robotic simulation for distributed reinforcement learning. In *Conference on Robot Learning (CoRL)*, 2018.
- Mirza, M., Jaegle, A., Hunt, J. J., Guez, A., Tunyasuvunakool, S., Muldal, A., Weber, T., Karkus, P., Racanière, S., Buesing, L., Lillicrap, T. P., and Heess, N. Physically embedded planning problems: New challenges for reinforcement learning. *arXiv:2009.05524*, 2020.

- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. Human-level control through deep reinforcement learning. *Nature*, 518(7540), 2015.
- Mordatch, I. and Abbeel, P. Emergence of grounded compositional language in multi-agent populations. In *AAAI*, 2018.
- Parisotto, E., Song, H. F., Rae, J. W., Pascanu, R., Gülçehre, Ç., Jayakumar, S. M., Jaderberg, M., Kaufman, R. L., Clark, A., Noury, S., Botvinick, M., Heess, N., and Hassel, R. Stabilizing transformers for reinforcement learning. In *ICML*, 2020.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. PyTorch: An imperative style, high-performance deep learning library. In *Neural Information Processing Systems*, 2019.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *ICML*, 2017.
- Perlin, K. An image synthesizer. In *SIGGRAPH*, 1985.
- Petrenko, A., Huang, Z., Kumar, T., Sukhatme, G., and Koltun, V. Sample factory: Egocentric 3D control from pixels at 100000 FPS with asynchronous reinforcement learning. In *ICML*, 2020.
- Savinov, N., Raichuk, A., Marinier, R., Vincent, D., Pollefeys, M., Lillicrap, T., and Gelly, S. Episodic curiosity through reachability. In *ICLR*, 2019.
- Savva, M., Chang, A. X., Dosovitskiy, A., Funkhouser, T., and Koltun, V. MINOS: Multimodal indoor simulator for navigation in complex environments. *arXiv:1712.03931*, 2017.
- Savva, M., Kadian, A., Maksymets, O., Zhao, Y., Wijmans, E., Jain, B., Straub, J., Liu, J., Koltun, V., Malik, J., Parikh, D., and Batra, D. Habitat: A Platform for Embodied AI Research. In *ICCV*, 2019.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv:1707.06347*, 2017.
- Sener, O. and Koltun, V. Learning to guide random search. In *ICLR*, 2020.
- Shacklett, B., Wijmans, E., Petrenko, A., Savva, M., Batra, D., Koltun, V., and Fatahalian, K. Large batch simulation for deep reinforcement learning. In *ICLR*, 2021.
- Stooke, A. and Abbeel, P. rlpyt: A research code base for deep reinforcement learning in PyTorch. *arXiv:1909.01500*, 2019.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 1992.
- Tesauro, G. Temporal difference learning and TD-gammon. *Commun. ACM*, 38(3), 1995.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., et al. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782), 2019.
- Wijmans, E., Kadian, A., Morcos, A., Lee, S., Essa, I., Parikh, D., Savva, M., and Batra, D. DD-PPO: Learning near-perfect PointGoal navigators from 2.5 billion frames. In *ICLR*, 2020.