

IGLU Gridworld: Simple and Fast Environment for Embodied Dialog Agents

Artem Zholus^{*†} Alexey Skrynnik[‡] Shrestha Mohanty[¶] Zoya Volovikova[§] Julia Kiseleva[¶]
 Artur Szlam[¶] Marc-Alexandre Côté[¶] Aleksandr I. Panov[‡]

Abstract

We present the IGLU Gridworld: a reinforcement learning environment for building and evaluating language conditioned embodied agents in a scalable way. The environment features visual agent embodiment, interactive learning through collaboration, language conditioned RL, and combinatorically hard task (3d blocks building) space.

1. Introduction

In this work, we propose a new reinforcement learning environment called Interactive Grounded Language Understanding (IGLU) Gridworld¹. This environment is at the core of the IGLU 2022 competition hosted at NeurIPS² [6].

The environment consists in an asymmetric collaboration between the *Architect* who has oracle access to the target structure and has to provide an instruction to the *Builder* that has to follow the instruction in a 3D blocks gridworld (Figure 1). While the overall task is *trivially-sided* (i.e., the collaboration starts and ends with simply a small integer 3D array with just 7 possible block colors), it encompasses a nontrivial process which involves contextualized collaborative instructions and an embodied behavior. The IGLU Gridworld enables research in 1) collaborative learning (due to the nature of the task); 2) different kinds of generalization in imitation/reinforcement learning (since the space of tasks is large and grids/dialogues can be easily augmented); 3) hierarchical RL (given the color-shape patterns in structures and the contextualized nature of dialogs); 4. skill learning (since building different sorts of structures requires different action patterns); 5) open-ended learning (since the overall input and target, a 3D grid, is easily formalizable for a co-evolution of the Architect and Builder); and 6) embodied AI (since the environment represents an ego-centric agent




Figure 1. Example of an Architect-Builder task.

that lives in a 3D visual world with combinatorically rich space of states). IGLU Gridworld is implemented in Python to facilitate transparency and simplicity of the game, flexibility of the rules, and uniformity with our data collection tool [5]. The last one is especially important for collecting the dataset with behaviours that are valid for our environment. The contributions of this work are two fold: 1) IGLU Gridworld, a new lightweight and fast RL environment written in pure Python for learning embodied agents that follow dialog instructions with collaborative context; and 2) we propose a multitask hierarchical RL baseline and evaluate its performance on various tasks from [10] grouped by embodiment skills needed to solve that tasks.

	gridworld	embodiment	rich task space	language (state)	API	speed (SPS) ³
Mujoco [16]	✗	2D/3D	✗	✗	C++	2.4k
Baby AI [1]	2D	2D	✓	CFG	python	3k
Nethack [8,9,13]	2D	2D	✓	natural	DSL	14.4k
TextWorld [2]	✗	✗	✓	CFG	python	300
AI2Thor [7]	✗	3D	✓	natural	python	30
Habitat 2.0 [14]	✗	3D	✗	✗	C++	1.4k
Megaverse [12]	3D	3D	✗	✗	C++	327k
Xland [15]	3D	3D	✓	CFG	✗	30
MineRL [4]	3D	3D	✓	✗	python	180
IGLU Gridworld	3D	3D	✓	natural	python	4.4k

Table 1. IGLU Gridworld vs. related RL environments. CFG: context free grammar. DSL: domain-specific language. Rich task space: can compose tasks to yield a combinatorial task space. API: language that can be used to extend the environment.

2. Related Work

IGLU Gridworld and IGLU contest [5] are inspired by the collaborative task proposed in [10]. Several environ-

*Corresponding author: zholus.aa@phystech.edu

†Moscow Institute of Physics and Technology (MIPT)

‡Artificial Intelligence Research Institute (AIRI)

§ITMO University

¶Microsoft Research

¶Meta AI

¹<https://github.com/iglu-contest/gridworld>

²<https://www.iglu-contest.net/>

ments offer similar features to the ones proposed in IGLU Gridworld. We summarize them in Table 1. In particular, IGLU is created for testing 3D embodied agents that are conditioned on natural language instructions. Combined with a rich task space (combination of 3D colored blocks), the environment offers a scalable (i.e., fast and lightweight environment) and extendable (i.e., simple Python API) platform for language and visual embodiment understanding.

3. IGLU Gridworld Description

We implement the environment in Python using a simplified version of an open-source Minecraft clone⁴. The renderer is decoupled from the core environment and can be disabled with zero overhead enabling the environment to run headless with GPU-acceleration. **The environment runs at 4400 steps-per-second (SPS) but can reach 17000 when rendering is disabled.** Such speed, for a Python RL environment, is possible due to its simplicity. The observation space consists of a point-of-view (POV) image (64, 64, 3), an inventory list (6,), a snapshot of the building zone (11, 9, 11), and the agent’s position with pitch and yaw angles (5,). The task information is provided in string instruction. The building zone is represented via a 3D tensor with block ids (0 for air, 1 for blue block, etc.). The agent can navigate the building zone, place and break blocks, and switch between block types. The action space is a discrete space of 14 actions: *noop*, *step [forward—backward—right—left]*, *jump*, *break block*, *place block*, *choose block type [1-6]* and two-dimensional continuous camera movement action. The reward reflects how complete the built structure is irrespective of its location. We run a convolution-like procedure between the state grid and the target grid. We find the best alignment (translation and rotation of state grid onto the target grid) that yields the maximum per-block intersection⁵. We reward the agent for change in the maximal intersection over time.

4. Tasks in IGLU Gridworld

We split tasks into subtasks in IGLU Gridworld using a notion of collaboration *segments*. Each segment is represented by a (collaborative context, target) pair. Context is a (dialog of instructions and clarifying questions, blocks placed in response to that dialog) pair. Target is a (most recent instruction, target blocks placed in response to that instruction) pair. For each task, the environment is initialized with context blocks and the agent is tasked to extend or shrink it to the target blocks. The only source of target information provided to the agent is the target instruction and

the context dialog. The segments serve as building blocks for each research direction mentioned in Section 1. As one of them, we describe the interactive collaboration used in the IGLU competition 2022. Each task is a sequence of segments and the environment always starts with a segment with empty starting world. For each segment, the agent acts and once the segment is “solved”, the environment internally switches to the next segment until the last one is solved. We view this process as resembling to the teacher forcing technique used in NLP. In contrast to the evaluation, during training we can reset the env at any segment.

5. Baseline

Training an agent to build any language-defined structure is a challenging task. To overcome this, we have developed a multitask hierarchical builder (MHB) with three modules: task generator (NLP part), subtask generator (heuristic part), and subtask solving module (RL part).

First, we finetuned on the dialog dataset [10] using BERT-based [3] 3D transposed convolutional head for the target prediction. Second, we developed the subtasks heuristic (with predefined blocks order) module, which uses the target (3D voxel) as input to predict next cube to add or remove. Third, we trained the agent, using the high-performance APPO algorithm [11] extended for goal-based policy. The agent curriculum generated compact goal structures (not related to dataset [10]). And finally, the complete MBH algorithm was evaluated on all tasks from the dataset [10] (see Table 2). We also report the results for the random agent. We labeled each task with the embodied skills required to solve that task. For the skill description and labeling, see gridworld repository⁶.

F_1 score	flying	tall	diagonal	flat	tricky	all
MHB agent (NLP)	0.292	0.322	0.242	0.334	0.295	0.313
MHB agent (full)	0.233	0.243	0.161	0.290	0.251	0.258
Random agent (full)	0.039	0.036	0.044	0.038	0.043	0.039

Table 2. Per-skill aggregation of the baselines performance metrics. For each task, we calculate F_1 score between built and target structures. For each skill, we average the performance on all targets requiring that skill.

6. Conclusion

In this work, we presented IGLU Gridworld, a simulator for testing embodied agent in interactive blocks building task with dialog context. The environment offers its speed and scalability for many research directions related to RL, Embodied AI, NLP, Lifelong learning. We presented an end-to-end multitask hierarchical RL baseline, with notable performance which is higher than the performance of any solution of the IGLU competition at NeurIPS 2021 [5].

³For Megaverse, we report the one GPU simulation throughput [12], note that this environment can be only run in parallelization mode whereas, for the rest (including ours), we report single env speed. For XLand, we report the speed from [17].

⁴<https://github.com/fogleman/Minecraft>

⁵github.com/iglu-contest/gridworld/blob/master/gridworld/task.py

⁶<https://github.com/iglu-contest/gridworld/tree/master/skills>

References

- [1] Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. BabyAI: First steps towards grounded language learning with a human in the loop. In *International Conference on Learning Representations*, 2019. [1](#)
- [2] Marc-Alexandre Côté, Ákos Kádár, Xingdi Yuan, Ben Kybartas, Tavian Barnes, Emery Fine, James Moore, Ruo Yu Tao, Matthew Hausknecht, Layla El Asri, Mahmoud Adada, Wendy Tay, and Adam Trischler. Textworld: A learning environment for text-based games. *CoRR*, abs/1806.11532, 2018. [1](#)
- [3] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018. [2](#)
- [4] William H. Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden Codell, Manuela Veloso, and Ruslan Salakhutdinov. Miner1: A large-scale dataset of minecraft demonstrations. In *IJCAI*, pages 2442–2448, 2019. [1](#)
- [5] Julia Kiseleva, Ziming Li, Mohammad Aliannejadi, Shrestha Mohanty, Maartje ter Hoeve, Mikhail S. Burtsev, Alexey Skrynnik, Artem Zholus, Aleksandr I. Panov, Kavya Srinet, Arthur D. Szlam, Yuxuan Sun, Marc Hofmann, Ahmed Hassan Awadallah, Linar Abdrazakov, Igor Churin, Putra Mangala, Kata Naszádi, Michiel van der Meer, and Taewoon Kim. Interactive grounded language understanding in a collaborative environment: Iglu 2021. 2021. [1](#), [2](#)
- [6] Julia Kiseleva, Alexey Skrynnik, Artem Zholus, Shrestha Mohanty, Negar Arabzadeh, Marc-Alexandre Côté, Mohammad Aliannejadi, Milagro Teruel, Ziming Li, Mikhail Burtsev, Maartje ter Hoeve, Zoya Volovikova, Aleksandr Panov, Yuxuan Sun, Kavya Srinet, Arthur Szlam, and Ahmed Awadallah. Iglu 2022: Interactive grounded language understanding in a collaborative environment at neurips 2022, 2022. [1](#)
- [7] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. AI2-THOR: An Interactive 3D Environment for Visual AI. *arXiv*, 2017. [1](#)
- [8] Heinrich Küttler, Nantas Nardelli, Alexander Miller, Roberta Raileanu, Marco Selvatici, Edward Grefenstette, and Tim Rocktäschel. The nethack learning environment. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 7671–7684. Curran Associates, Inc., 2020. [1](#)
- [9] Michael Matthews, Mikayel Samvelyan, Jack Parker-Holder, Edward Grefenstette, and Tim Rocktäschel. Skillhack: A benchmark for skill transfer in open-ended reinforcement learning. In *ICLR Workshop on Agent Learning in Open-Endedness*, 2022. [1](#)
- [10] Anjali Narayan-Chen, Prashant Jayannavar, and Julia Hockenmaier. Collaborative dialogue in Minecraft. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5405–5415, Florence, Italy, July 2019. Association for Computational Linguistics. [1](#), [2](#)
- [11] Aleksei Petrenko, Zhehui Huang, Tushar Kumar, Gaurav Sukhatme, and Vladlen Koltun. Sample factory: Egocentric 3d control from pixels at 100000 fps with asynchronous reinforcement learning. In *ICML*, 2020. [2](#)
- [12] Aleksei Petrenko, Erik Wijmans, Brennan Shacklett, and Vladlen Koltun. Megaverse: Simulating embodied agents at one million experiences per second. In *ICML*, 2021. [1](#), [2](#)
- [13] Mikayel Samvelyan, Robert Kirk, Vitaly Kurin, Jack Parker-Holder, Minqi Jiang, Eric Hambro, Fabio Petroni, Heinrich Küttler, Edward Grefenstette, and Tim Rocktäschel. Minihack the planet: A sandbox for open-ended reinforcement learning research. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. [1](#)
- [14] Andrew Szot, Alex Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimir Vondrus, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel Chang, Zsolt Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra. Habitat 2.0: Training home assistants to rearrange their habitat. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. [1](#)
- [15] Open Ended Learning Team, Adam Stooke, Anuj Mahajan, Catarina Barros, Charlie Deck, Jakob Bauer, Jakub Sygnowski, Maja Trebacz, Max Jaderberg, Michaël Mathieu, Nat McAleese, Nathalie Bradley-Schmieg, Nathaniel Wong, Nicolas Porcel, Roberta Raileanu, Steph Hughes-Fitt, Valentin Dalibard, and Wojciech Marian Czarnecki. Open-ended learning leads to generally capable agents. *CoRR*, abs/2107.12808, 2021. [1](#)
- [16] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012. [1](#)
- [17] Tom Ward, Andrew Bolt, Nik Hemmings, Simon Carter, Manuel Sanchez, Ricardo Barreira, Seb Noury, Keith Anderson, Jay Lemmon, Jonathan Coe, Piotr Trochim, Tom Handley, and Adrian Bolton. Using unity to help solve intelligence. *CoRR*, abs/2011.09294, 2020. [2](#)