

Graph augmented Deep Reinforcement Learning in the GameRLand3D environment

Edward Beeching,^{1,2} Maxim Peter¹ Philippe Marcotte¹ Jilles Debangoye,² Olivier Simonin²
Joshua Romoff¹ Christian Wolf³

¹ Ubisoft La Forge, Montreal

² INRIA Chroma team, CITI Laboratory. INSA-Lyon, France.

³ Université de Lyon, INSA-Lyon, LIRIS, CNRS, France.

Abstract

We address planning and navigation in challenging 3D video games featuring maps with disconnected regions reachable by agents using special actions. In this setting, classical symbolic planners are not applicable or difficult to adapt. We introduce a hybrid technique combining a low level policy trained with reinforcement learning and a graph based high level classical planner. In addition to providing human-interpretable paths, the approach improves the generalization performance of an end-to-end approach in unseen maps, where it achieves a 20% absolute increase in success rate over a recurrent end-to-end agent on a point to point navigation task in yet unseen large-scale maps of size 1km×1km. In an in-depth experimental study, we quantify the limitations of end-to-end Deep RL approaches in vast environments and we also introduce “GameRLand3D”, a new benchmark and soon to be released environment can generate complex procedural 3D maps for navigation tasks.

Introduction

Long term planning and navigation is an important component when looking to realistically control a non-player character (NPC) in a video game. Whether the NPC is player facing or used for automated testing to find bugs and exploits, the ability to navigate in a complex environment is paramount and is the foundation of every NPC. In the field of video games the Navigation Mesh (NavMesh) (Snook 2000) has become the ubiquitous approach for planning and navigation. In recent years, the sheer scale of video game environments have highlighted the limitations of the NavMesh, particularly in games with complex navigation abilities such as double jumping, teleportation, jump pads, grappling hooks, and wall runs. Given the successes of Deep Reinforcement Learning (RL) in domains of navigation (Mirowski et al. 2016; Jaderberg et al. 2016, 2019), board games (Tesauro 1994; Silver et al. 2016, 2017) and Atari video games (Mnih et al. 2013; Hessel et al. 2018; Hafner et al. 2020), recent works (Alonso et al. 2021) have looked to exploit these approaches for navigation in large 3D video game environments. These works have shown promise




Figure 1: The proposed hybrid planner is capable of performing real-time navigation in complex 3D gaming environments featuring geometrically disconnected regions only reachable by agents through special actions. These environments are difficult for approaches based on the NavMesh.

in smaller game worlds, but have highlighted that as environments scale to hundreds of thousands of square metres, the limitations of end-to-end learning-based methods are encountered.

Currently, with few exceptions, navigation in video games relies on the NavMesh, which is created and used as follows.

1. A graph representation of the world is pre-generated from the game geometry.
2. At runtime a pathfinding algorithm such as A* (Hart, Nilsson, and Raphael 1968b) is run on this graph to find the shortest path between two locations in the game.
3. A player controller, attached to the character, is used for point based navigation to follow the sequence of way-points returned by the planner.

While the navigation mesh provides a compact representation of the world, it is somewhat limited and results in disconnected regions in complex environments, as demonstrated in Figure 1. The disconnected regions are due to some assumptions of the underlying NavMesh generation algorithm. The user can provide the maximum inclination and height the agent can walk and jump, but more complex actions such as double jumping or using a jump pad are

not taken into account. Connections due to special character abilities such as jump pads, grappling hooks and teleporters often involve time consuming manual editing of the graph in order to add additional links. While human intervention is possible, it is not feasible when the environments are procedurally generated as there can be a limitless number of new environment instances. Some automatic approaches do exist but are currently limited to simple actions (Axelrod 2008; Roumimper 2017; Budde 2013).

Recent works have looked to address this problem with an end-to-end Deep Reinforcement Learning (RL) approach (Alonso et al. 2021) and have shown that when trained and evaluated on small (50m×50m) to medium (100m×100m) sized maps, end-to-end approaches achieve good performance. However, when evaluated in larger (300m×300m) environments, policies learned using RL perform well on near and medium distance goals but poorly on farther challenging goals. Euclidean distance is not the best measure of complexity in this case, but as the odds of encountering a hard obstacle increase with the size of the map, it is a decent proxy for the complexity of navigation.

In this paper we explore a hybrid approach that aims to combine the strengths of a classical graph-based NavMesh, with the capacity of end-to-end RL to incorporate complex abilities to improve navigation performance. Designed around the stringent constraints of online gaming environments, our method is able to generate new navigable graphs with low computational complexity. The proposed graph augmented RL agent outperforms the generalization performance of a learning-based agent on a set of maps of 250m×250m. In addition, the augmented approach also greatly improves the performance of a map specific agent that is trained and evaluated on a fixed game world of 1km×1km, 10 times larger than those typically used in the research community (Alonso et al. 2021).

We also introduce the “*GameRLand3D*” environment, a new benchmark, which pushes the size and complexity of gaming environments to new limits. It features realistic production-size maps as well as well as complex navigation tasks suited for modern 3D video games featuring navigation actions like jumps, pads, zip lines etc. We think this benchmark is well suited to quantify the performance of end-to-end approaches and to explore potential novel solutions.

Our contributions are as follows:

- We quantify the limitations of the performance of end-to-end Deep RL approaches in complex and large-scale environments.
- We propose a novel hybrid graph and RL approach that decouples long distances planning and low level navigation and outperforms end-to-end Deep RL methods in both a generalization and map specific setup. Notably, this approach provides interpretable plans which can be tweaked by a game designer in order to modify the agent behavior without the need to fully retrain it.
- We release of an open-source 3D procedural environment which can generate complex 3D worlds of more than one million square meters.




Figure 2: A vast, 1km by 1km procedurally generated map generated with in the GameRLand3D environment.

- We introduce a detailed empirical evaluation and ablation of two hybrid RL and graph based approaches, with a special focus on their run time performance, a key requirement of video games operating in real time at 60 frames per second.

Related work

Deep RL environments

There has been a proliferation of new Deep RL environments in the last five years, the most well known being the gym framework (Brockman et al. 2016) which provides an interface to fully observable games such as the Atari-57 benchmark and several continuous control tasks. There are a variety of partially observable 3D simulators available, starting with more game-like environments such as the ViZ-Doom simulator (Wydmuch, Kempka, and Jaśkowski 2018), DeepMind-Lab (Beattie et al. 2016) and Malmö (Johnson et al. 2016). A number of photo-realistic simulators have also been released such as Gibson (Xia et al. 2018) and AI2Thor (Kolve et al. 2017). The highly efficient Habitat-Lab simulator (Manolis Savva et al. 2019) builds on the datasets of Gibson and AI2Thor and implements a number of benchmark 3D navigation tasks. The aforementioned environments were designed to train and evaluate RL agents in relatively small environments of the order of hundreds of square metres, with relatively small action spaces. Moreover, they do not consider the complex actions that are in modern video games, such as double jumping, teleportation, jump pads, grappling hooks, and wall runs. The objective of the GameRLand3D environment, shown in Figure 2, is to provide a realistic benchmark for real world video game environments, to identify the limitations of current learning-based techniques and to allow reasearcher to explore novel alternatives.

Classical planning and navigation in video games

Existing methods for path planning in the video game industry mostly rely on building a Navigation Mesh (NavMesh)(Snook 2000). A NavMesh is a 2D/3D graph

whose nodes represent convex polygonal regions of the walkable space, typically triangles, and whose edges indicate possibilities of navigation between regions. To find a path between two locations in the world, one can then apply path finding algorithms like A* (Hart, Nilsson, and Raphael 1968a) or one of its variants on this graph. This classical approach is robust for many applications but presents some limitations. In particular, when it is possible for the agent to use mobility actions (like dashing, jumping, double jumping, using a jetpack), the high number of places it can reach from each location dramatically increases the connectivity of the NavMesh (Alonso et al. 2021; Gordillo et al. 2021). Therefore, it increases its memory cost as well as the runtime cost of the pathfinding algorithms, to the point where using a NavMesh ends up being too costly, constraining game design. Another important aspect that a NavMesh does not account for is the kind of constraint that a character can have, such as the turn radius of a car or a plane in 3D. This kind of constraint can be impossible to encode in a NavMesh in an efficient manner.

Contrary to NavMesh-based approaches, RL for point-to-point navigation merges the path finding and the character controller (code the agent actually uses to move and follow the path) components by directly outputting actions the agent needs to take at every step to go from point A to point B. This makes using RL especially attractive since it could compress the semantics/topology of the map in a much more efficient way while having a fixed low cost to run.

Neural approaches to navigation

The field of Deep Reinforcement Learning (RL) has gained attention with successes on board games (Silver et al. 2016) and Atari games (Mnih et al. 2015). Recent works have applied Deep RL for the control of an agent in 3D environments (Mirowski et al. 2016; Jaderberg et al. 2016), exploring the use of auxiliary tasks such as depth prediction, loop detection and reward prediction to accelerate learning. Other recent work uses street-view scenes to train an agent to navigate in city environments (Mirowski et al. 2018).

To infer long term dependencies and store pertinent information about the partially observable environment, network architectures typically incorporate recurrent memory such as Gated Recurrent Units (Chung et al. 2015) or Long Short-Term Memory (Hochreiter and Schmidhuber 1997). Extensions to memory based neural approaches began with Neural Turing Machines (Graves, Wayne, and Danihelka 2014) and Differentiable Neural Computers (Graves et al. 2016) and have been adapted to Deep RL agents (Wayne et al. 2018).

Spatially structured memory architectures have been shown to augment an agent’s performance in 3D environments and are broadly split into two categories: metric maps which discretize the environment into a grid based structure and topological maps which produce node embeddings at key points in the environment. Research in learning to use a metric map is extensive and includes spatially structured memory (Parisotto and Salakhutdinov 2017), Neural SLAM based approaches (Zhang et al. 2017) and approaches incorporating projective geometry and neural memory (Gupta et al. 2017; Bhatti et al. 2016). These techniques are com-




Figure 3: Left: An agent’s depth observation in the GameRLand3D environment. Right: the resulting 2D observation in tensor form, that will be processed by the agent’s CNN.

bined, extended and evaluated in (Beeching et al. 2020b; Wani et al. 2020). Research combining learning, navigation in 3D environments and topological representations has been limited in recent years with notable works being (Savinov, Dosovitskiy, and Koltun 2018) who create a graph through random exploration in the ViZDoom RL environment (Wydmuch, Kempka, and Jaśkowski 2018). (Eysenbach, Salakhutdinov, and Levine 2019) also performs planning in 3D environments on a graph-based structure created from randomly sampled observations, with node distances estimated with value estimates. (Beeching et al. 2020a) perform topological planning with a neural approximation of a classical planning algorithm that can be applied in uncertain environments. Other recent approaches such as (Chaplot et al. 2020) build a graph structure, but rely on 360 degree camera measurements. Practically all of the aforementioned approaches are applied in small planar environments of hundreds of square metres, and scaling these methods to the environment where we apply our hybrid approach poses several challenges. Availability of ground truth: most graph based methods require ground truth actions and graph connectivity, as the calculation of optimal actions is often intractable in 3D environments of hundreds of thousands of square metres this precludes this possibility. We are also constrained due to inference time: many of these methods do not operate at real time speed, particularly when controlling hundreds of agents in parallel.

Environment and task definition

In order to address the challenging problem of navigation in vast complex 3D environments we have developed the GameRLand3D environment. GameRLand3D allows the creation of enormous worlds and enables complex interactions such as double jumping and jump-pads, includes buildings, water and lava hazards. The environment is built using the Unity-ML agents (Juliani et al. 2018) framework, which provides a versatile tool to build RL environments. We interact with the environment by controlling an agent, represented as a 1.8 meter humanoid character, shown in Figure 3.

Our environment implementation uses 3D Perlin noise (Perlin 1985) for the procedural generation of the game world, followed by random placements of buildings,

plateaus and jump pads to create a challenging state space for agents to interact with, explore and navigate.

Task The task we consider is point-to-point navigation where the agent is spawned at a location p_{spawn} and must navigate to a goal location p_{goal} . Start and goal locations are sample from navigable terrain, using a NavMesh for simplicity, through a function $f_{navmesh}$ from which we can sample valid navigation locations. The NavMesh contains many disconnected regions and cannot be directly used for long distances navigation, but can be used to identify valid spawn locations. To reduce overall inference time and to allow for planning over longer time horizons, we request a decision from the agent’s policy network every 10 time-steps, repeating the actions for the intermediary 10 steps.

Constraints Video game environments provide certain advantages for training RL agents when compared to learning a model that may be transferred to the real world. The pose of the agent is known and we are not expected to rely on noisy pose estimates and fulfill safety guarantees, which add an additional challenge to robotic navigation. Video games do have other requirements and constraints, model implementations must be highly optimized, parallelized and run at real time speeds, making decisions at a minimum of 60 times per second. If an agent is player facing, its behavior must be credibly “human-like”, a subjective metric that is discussed in (Devlin et al. 2021). Ideally, methods should provide interpretable information, so that game designers can be confident about adopting it in production.

Observations Rendering agent observations from the environment with a virtual monocular camera is prohibitively expensive for NPC control in video games, so the agent’s observations are a depth measurement of a 8×8 cone of raycasts in front of the agent (see Figure 3). We also provide useful information such as the location of the goal in the agent’s frame of reference, the agent’s velocity and acceleration.

Actions The environment actions are continuous and correspond to actions that are typically available to a human player: forward, backward, strafe, turn, jump and double-jump (where a player can execute a second jump while in mid-air). The jump action is treated like a continuous action for the RL policy and is discretized in the environment.

Rewards Building on the work of (Alonso et al. 2021), we combine three rewards: A **sparse reward** of +1 when the agent succeeds in the task and -1 when the agent fails. A **dense reward** every step in which the agent reduces its best euclidean distance to the goal location, divided by a normalization factor λ . A **time-step penalty** of -0.0005 per step, to encourage the agent to perform the task quickly. An episode is terminated when the agent is within 1m of its goal, the agent falls off the map or the agent has not decreased its best euclidean distance to the goal for 300 time-steps.

Empirical evaluation of end-to-end methods

We initially trained a baseline policy, in order to quantify the empirical performance of end-to-end Deep RL approaches

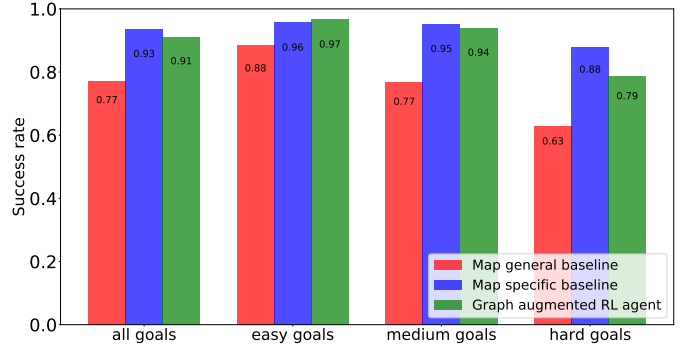


Figure 4: The generalization gap in performance between a map specific recurrent agent trained on a single map and a general recurrent agent trained on 128 maps and then evaluated on the same unseen map. Compared against a Graph augmented RL agents, introduced in this paper.

in such large-scale worlds. In particular, we were interested in generalization, how well a policy trained in one set of environments transfers to a new environment. The baseline policy was trained using a recurrent version of a sample efficient off-policy actor critic RL algorithm, Soft Actor Critic (SAC) (Haarnoja et al. 2018a), with a learned entropy coefficient and no state value network (Haarnoja et al. 2018b). The SAC algorithm aims to find a policy that maximizes cumulative reward and the entropy of each visited state.

$$\pi^* = \arg \max_{\pi} \sum_t \mathbb{E}_{(s_t, a_t) \sim \rho_{\pi}} [r(s_t, a_t) + \alpha \mathcal{H}(\pi(\cdot | s_t))] \quad (1)$$

We include a learnable entropy temperature parameter τ , to enable automatic optimization of the entropy of the objective function. We found that having independent policy and Q-network embeddings improves the reliability of the Q-value estimates and led to more robust, stable training. We trained several versions of the baseline agent architecture. A **reactive** agent that receives the 8×8 raycast observation and a **memory-based** recurrent agent that extends the reactive agent with an LSTM in order to be able to learn longer term dependencies. We use a burn-in of 20 time-steps to initialize the hidden state as recommended in (Kapturowski et al. 2018; Paine et al. 2019).

In order to quantify the reduction in performance when a policy is learned on one specific map or when a policy is learned on a set of maps and evaluated on another, we trained two versions of the memory based agent: a general agent and a map specific agent. The general agent is trained on a large dataset of 128 procedurally generated map instances. During training we evaluate the performance of the policy on a set of 4 validation maps and keep the model with the highest success rate. In addition, independently we train 4 map specific agents, one per test map, the objective being to quantify the decrease in performance, or *generalization gap*, between the map specific and general agent, see Figure 4. The graph augmented methods that we discuss in the following section use the general baseline for point-to-point navigation and

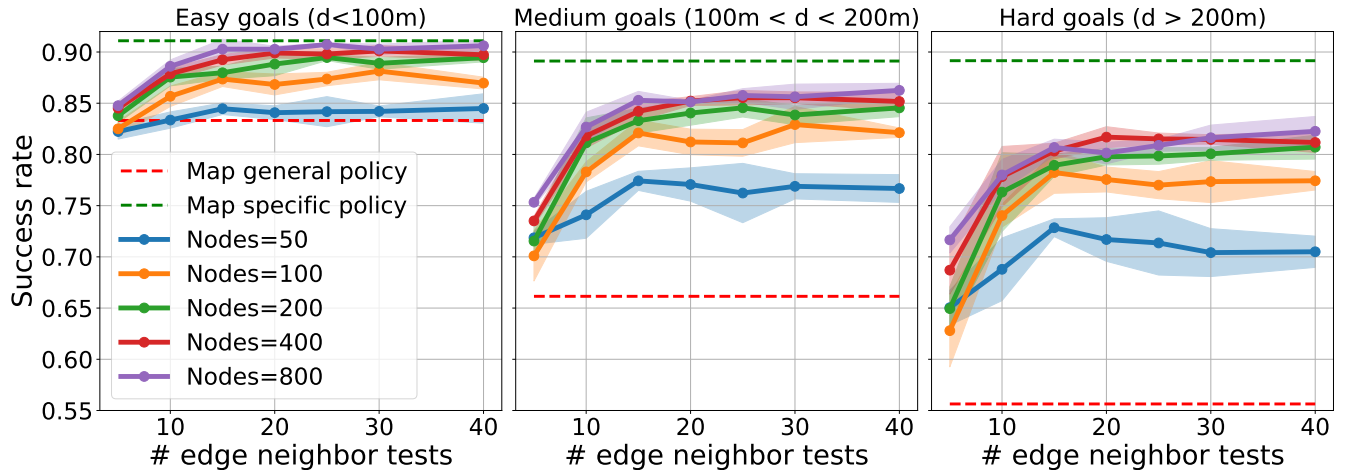


Figure 5: Results for different hyper-parameters for the method *Distance constrained random vertex placement*: graph sizes of 50 to 800 nodes; edge candidates of 5 to 40 neighbors

achieve comparable performance to the map specific agent. We note that a general agent is a crucial requirement in the video game production setting, as it removes the need of re-training on every new map instance.

Graph augmented RL agents

Pure end-to-end training of vanilla agents with recurrent memory leads to sub-optimal solutions for large and complex environments, as planning is performed on memory without spatial structure. While spatial organization of neural memory can be learned implicitly in principle, as shown by the emergence of grid-cells in unstructured agents trained for localization (Cueva and Wei 2018), in practice this remains a hard task, in particular from reward alone. On the other hand, and as mentioned, the NavMesh alone is hardly applicable in the presence of special agent actions leading to disconnected regions.

We address this problem and propose a hybrid method which combines a directed graph (digraph) for planning and way-point to way-point navigation with a low-level policy trained with RL. This approach decomposes the problem of navigation over long distances into sequences of smaller sub-problems. The motivation is that (i) small-scale way-point navigation can be learned easily through reward, even in the presence of disconnected regions accessible with jump actions; (ii) large-scale navigation in a connected graph can be solved efficiently with optimal symbolic algorithms, like Dijkstra’s algorithm or A*. We will show that the proposed method features a high success rate, provides interpretable paths and allows for identification of disconnected regions in the environment topology.

We explored several variants of the proposed method, which differ in the way how they dynamically create graphs when a new environment is spawned. For all variants, we build a graph G comprised of a set of vertices V and edges E , where edge $e(i, j)$ connects vertex v_i to vertex v_j . Each edge has a cost $c(i, j)$. Assigned to each vertex is a 3D po-

sition in the game. Planning from a position p_{start} to a goal position p_{goal} is performed by identifying the nearest vertex in the graph to the starting position v_{start} and the nearest vertex to the goal position v_{goal} . The optimal sequence of vertex locations to traverse in order to reach v_{goal} is found with the classical path planning algorithm Dijkstra (Dijkstra 1959).

The process of graph creation follows three key steps: identify vertex locations, testing edge connectivity and graph building. The following sections describe the different methods we explored for identifying the vertex locations, edges, edge costs and building the graph.

Unconstrained random vertex placement

In the unconstrained random approach we sample n vertices from the function $f_{navmesh}$. This provides vertices that are located on navigable terrain, but the edge connectivity is unknown. Edge connectivity is computed by first identifying the k nearest neighbors of each vertex and then evaluating a short point-goal episode between each vertex and its k -nearest neighbors. If the episode is successful then the directed edge is considered to be connected, with its cost equal to the number of steps in the episode. Using the agent steps as a cost over the euclidean distance leads to graph planning that takes into account local obstacles in the planning process.

Distance constrained random vertex placement

We augment the unconstrained random approach with coverage criteria, in order to spread nodes evenly across the environment and ensure that we can estimate connectivity all over the map. We aim to identify the set of vertex positions V that satisfy the criteria $d_V := \max_{x, y \in V, x \neq y} d(x, y)$, where $d(x, y)$ is the euclidean distance between two vertex positions. We find the set of positions V using a greedy strategy. We first calculating the optimal average distance between vertices, assuming the environment was a planar

surface, and then sampling vertex locations from $f_{navmesh}$ until no nodes are present in the proximity of the sampled vertex location. In order to avoid infinite loops this test is repeated 100 times and after each test the distance criterion is decayed by a factor of 0.98.

Optimizations

Robust estimation of edge connectivity One weakness of performing the edge connectivity test once is that there is a small probability that agents can introduce an edge in the graph that often fails, leading to overconfident planning. This problem is compounded when there are several edges in the graph that are difficult to navigate. In order to identify and prune these edges, we evaluate an option to repeat the edge connectivity tests several times and keep the edges where the percentage of the tests exceed a given threshold. This approach leads to robust plans and higher performance, with the downside of a longer graph building step.

Re-planning Qualitative analysis demonstrated that occasionally a graph augmented agent would fail to reach its next way-point, resulting in an agent that was often stuck as it had fallen off the side of an obstacle. This was resolved with the implementation of a re-planning step, where if the agent does not decrease its distance to the next way-point for 50 time-steps we automatically re-plan from its current location to the goal location. We observed that re-planning increased the success rate of all the examined methods.

Experiments

Setup We generate maps of size $250m \times 250m \times 80m$, 128 training maps, 4 validation maps and 4 tests maps to test the baseline and graph augmented methods. We then scale this approach to maps of $1km \times 1km \times 80m$ to evaluate the performance in vast game worlds. Such a map is notably larger than the largest maps tested in the literature (Alonso et al. 2021).

While other works also include the Success weighted by Path Length (SPL) metric (Anderson et al. 2018), we are unable to compare against this metric as the SPL metric. As SPL requires ground truth shortest paths, which cannot be computed in such large environments with complex action spaces. We separate the goals into three groups, easy, medium and hard, based on distance between the start and end points to break down the performance of each method. The results on hard, long distance goals, are of particular interest as this is where the baseline policy had the worst performance.

Hyper-parameter optimizations We extensively optimized the hyper-parameters of each approach and report the results of the best performing configuration for each technique. Each test was evaluated with three seeds in order to quantify the average and variation of performance that can be achieved with these methods. We evaluate the performance of each method using the success rate metric on 4 test maps, with 2,500 goals per map. As an example, we report in Figure 5 the evaluation of a single method, *Random vertex with coverage constraint*, in numerous hyper-parameter configurations. In this case we have varied the number of nodes




Figure 6: Run-time trade-offs of the coverage constraint approach. Left: Average graph search time in ms in several configurations. Right: Average graph construction and evaluation time in minutes on 2,500 goals.

and the number of nearest neighbors with which to perform edge connectivity tests.

Comparisons We evaluate the performance of the proposed graph augmented methods and compare them to the performance of the recurrent baselines in table 1. We observe that the *Distance Constrained Random Vertex Placement* outperforms the baseline method with a success rate of 85.9%, an improvement of 16.3%. We also implemented flooding and trajectory based methods which are excluded for brevity.

Run-time analysis Run-time performance is critical in real time video games, where the game logic and rendering must be executed 60 times per second. We perform an analysis of the run-time performance of the various approaches. In particular, we are interested in the time it takes to build the graph, to evaluate 2,500 goals on a test map and average duration of a graph search when exploiting the hybrid approach. Figure 6 summarizes the results for key sets of hyper-parameters for the *Distance Constrained Random Vertex Placement* approach. While larger graphs improve the performance of the augmented agent, they also increase the evaluation time and the time taken to perform a graph search. There is clearly a trade-off to be made, with the optimal choice depending on the exact use case, for an offline exhaustive test of a map instance a higher node count is suitable. Whereas if the graph must be created in near real-time a smaller node count would be required.

Scaling to production size maps Previous experiments have focused on generalization of a learned policy to unseen maps. We have extended our method to a vast $1km \times 1km \times 200m$ map. We find that even a map specific baseline RL agent, whose policy is learned directly on the map it is evaluated on (environment overfit), performs poorly compared to our hybrid agent, with an overall success rate of 72%. We augment the map specific agent with the *Distance Constrained Random Vertex Placement approach* and evaluated the performance on 10,000 goals with graphs of size 200 to 1,600 nodes. We achieve an increase in success rate of 20% on all goals, shown in Figure 7.

Method	Success rate			
	All goals	Easy goals	Medium goals	Hard goals
Map specific LSTM baseline (not comparable)	0.899 ± 0.000	0.911 ± 0.000	0.891 ± 0.001	0.892 ± 0.001
General reactive baseline	0.680 ± 0.000	0.821 ± 0.000	0.619 ± 0.000	0.574 ± 0.000
General LSTM baseline	0.696 ± 0.000	0.833 ± 0.000	0.661 ± 0.001	0.556 ± 0.002
Unconstrained random vertex placement	0.824 ± 0.023	0.878 ± 0.013	0.812 ± 0.029	0.767 ± 0.029
Distance constrained random vertex placement	0.859 ± 0.002	0.904 ± 0.002	0.850 ± 0.004	0.809 ± 0.004

Table 1: Evaluation on a test set of 4 maps of size 250m×250m. In addition to general reach and memory-based policies, we include the performance of 4 map specific policies, learned independently on each test map. We compare against the best performance of the unconstrained and constrained coverage-based hybrid methods.




Figure 7: Evaluation of the Random Graph with coverage approach on a single $1km^2$ map with 10,000 goals. Compared against a map specific LSTM baseline policy. Shown is goal distance, in 100m bins, compared with the success rate of each approach. We observe that the method begins to reach diminishing returns at 1,600 vertex locations. Overall average success rate is increased from 72% to 92%.

Conclusions

We have introduced a set of graph augmented RL approaches for planning and navigation in vast 3D video game worlds. To demonstrate the capacity of these approaches we have developed the GameRLand3D environment, an open-source reinforcement learning environment. GameRLand3D enables procedural generation of vast open world environments with hazards such as water and lava, and features such as jump pads, buildings and overhangs. We have identified that end-to-end Deep RL approaches under-perform in large-scale environments. We evaluated two graph augmented RL techniques and performed extensive hyper-parameter tests to find a reasonable compromise between graph building time, run-time performance and success rate. Our hybrid approaches boost generalization performance from 69.6% to 86.2%. We also achieve a 20% increase in success rate in enormous environments.

In addition to support and maintenance of the GameRLand3D environment, our roadmap includes the addition of more complex building types, player abilities such as teleportation and jetpacks, and other features such as ravines and caves. Future hybrid approaches in this environment will in-

vestigate the possibility of online graph building, using function approximation to predict edge connectivity in dynamic environments where the graph must be reconstructed when there is a change in map topology.

References

- Alonso, E.; Peter, M.; Goumard, D.; and Romoff, J. 2021. Deep Reinforcement Learning for Navigation in AAA Video Games. In Zhou, Z.-H., ed., *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, 2133–2139. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Anderson, P.; Chang, A.; Chaplot, D. S.; Dosovitskiy, A.; Gupta, S.; Koltun, V.; Kosecka, J.; Malik, J.; Mottaghi, R.; Savva, M.; et al. 2018. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*.
- Axelrod, R. 2008. *Navigating graph generation in highly dynamic worlds*, volume 4 of *AI Game Programming Wisdom*, chapter 2.6, 125–141. Charles River Media.
- Beattie, C.; Leibo, J. Z.; Teplyashin, 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. 2016. DeepMind Lab. *arxiv pre-print 1612.03801*.
- Beeching, E.; Dibangoye, J.; Simonin, O.; and Wolf, C. 2020a. Learning to plan with uncertain topological maps. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, 473–490. Springer.
- Beeching, E.; Wolf, C.; Dibangoye, J.; and Simonin, O. 2020b. EgoMap: Projective mapping and structured egocentric memory for Deep RL. *arXiv preprint arXiv:2002.02286*.
- Bhatti, S.; Desmaison, A.; Miksik, O.; Nardelli, N.; Sidharth, N.; and Torr, P. H. 2016. Playing Doom with slam-augmented deep reinforcement learning. *arXiv preprint arXiv:1612.00380*.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym. *arxiv pre-print 1606.01540*.
- Budde, S. 2013. *Automatic generation of jump links in arbitrary 3d environments*. Master’s thesis, Humboldt University of Berlin.

- Chaplot, D. S.; Salakhutdinov, R.; Gupta, A.; and Gupta, S. 2020. Neural Topological SLAM for Visual Navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Chung, J.; Gulcehre, C.; Cho, K.; and Bengio, Y. 2015. Gated feedback recurrent neural networks. In *International conference on machine learning*, 2067–2075. PMLR.
- Cueva, C.; and Wei, X. 2018. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. In *International Conference on Learning Representations*.
- Devlin, S.; Georgescu, R.; Momennejad, I.; Rzepecki, J.; Zuniga, E.; Costello, G.; Leroy, G.; Shaw, A.; and Hofmann, K. 2021. Navigation Turing Test (NTT): Learning to Evaluate Human-Like Navigation. In *2021 International Conference on Machine Learning*. Source code available at: <https://github.com/microsoft/NTT>.
- Dijkstra, E. W. 1959. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1): 269–271.
- Eysenbach, B.; Salakhutdinov, R. R.; and Levine, S. 2019. Search on the replay buffer: Bridging planning and reinforcement learning. In *Advances in Neural Information Processing Systems*.
- Gordillo, C.; Bergdahl, J.; Tollmar, K.; and Gisslén, L. 2021. Improving Playtesting Coverage via Curiosity Driven Reinforcement Learning Agents. *arXiv:2103.13798*.
- Graves, A.; Wayne, G.; and Danihelka, I. 2014. Neural Turing machines. *arXiv preprint arXiv:1410.5401*.
- Graves, A.; Wayne, G.; Reynolds, M.; Harley, T.; Danihelka, I.; Grabska-Barwińska, A.; Colmenarejo, S. G.; Grefenstette, E.; Ramalho, T.; Agapiou, J.; et al. 2016. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538(7626): 471–476.
- Gupta, S.; Davidson, J.; Levine, S.; Sukthankar, R.; and Malik, J. 2017. Cognitive mapping and planning for visual navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018a. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*.
- Haarnoja, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; et al. 2018b. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*.
- Hafner, D.; Lillicrap, T.; Norouzi, M.; and Ba, J. 2020. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*.
- Hart, P.; Nilsson, N.; and Raphael, B. 1968a. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2): 100–107.
- Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968b. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, SSC-4(2): 100–107.
- Hessel, M.; Modayil, J.; van Hasselt, H. P.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2018. Rainbow: Combining improvements in deep reinforcement learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Jaderberg, M.; Czarnecki, W. M.; Dunning, I.; Marris, L.; Lever, G.; Castaneda, A. G.; Beattie, C.; Rabinowitz, N. C.; Morcos, A. S.; Ruderman, A.; et al. 2019. Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science*, 364(6443): 859–865.
- Jaderberg, M.; Mnih, V.; Czarnecki, W. M.; Schaul, T.; Leibo, J. Z.; Silver, D.; and Kavukcuoglu, K. 2016. Reinforcement Learning with Unsupervised Auxiliary Tasks. *arXiv preprint arXiv:1611.05397*.
- Johnson, M.; Hofmann, K.; Hutton, T.; and Microsoft, D. B. 2016. The Malmo Platform for Artificial Intelligence Experimentation. In *IJCAI*.
- Juliani, A.; Berges, V.-P.; Vckay, E.; Gao, Y.; Henry, H.; Mattar, M.; and Lange, D. 2018. Unity: A general platform for intelligent agents. *arXiv preprint arXiv:1809.02627*.
- Kapturowski, S.; Ostrovski, G.; Quan, J.; Munos, R.; and Dabney, W. 2018. Recurrent experience replay in distributed reinforcement learning. In *International conference on learning representations*.
- Kolve, E.; Mottaghi, R.; Gordon, D.; Zhu, Y.; Gupta, A.; and Farhadi, A. 2017. AI2-THOR: An Interactive 3D Environment for Visual AI. *arxiv pre-print 1712.05474v1*.
- Manolis Savva; Abhishek Kadian; Oleksandr Maksymets; Zhao, Y.; Wijmans, E.; Jain, B.; Straub, J.; Liu, J.; Koltun, V.; Malik, J.; Parikh, D.; and Batra, D. 2019. Habitat: A Platform for Embodied AI Research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Mirowski, P.; Grimes, M.; Malinowski, M.; Hermann, K. M.; Anderson, K.; Teplyashin, D.; Simonyan, K.; Zisserman, A.; Hadsell, R.; et al. 2018. Learning to navigate in cities without a map. *Advances in Neural Information Processing Systems*, 31: 2419–2430.
- Mirowski, P.; Pascanu, R.; Viola, F.; Soyer, H.; Ballard, A. J.; Banino, A.; Denil, M.; Goroshin, R.; Sifre, L.; Kavukcuoglu, K.; Kumaran, D.; and Hadsell, R. 2016. Learning to navigate in complex environments. *arXiv preprint arXiv:1611.03673*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- 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. 2015. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533.

Paine, T. L.; Gulcehre, C.; Shahriari, B.; Denil, M.; Hoffman, M.; Soyer, H.; Tanburn, R.; Kapturowski, S.; Rabinowitz, N.; Williams, D.; et al. 2019. Making efficient use of demonstrations to solve hard exploration problems. *arXiv preprint arXiv:1909.01387*.

Parisotto, E.; and Salakhutdinov, R. 2017. Neural map: Structured memory for deep reinforcement learning. *arXiv preprint arXiv:1702.08360*.

Perlin, K. 1985. An image synthesizer. *ACM Siggraph Computer Graphics*, 19(3): 287–296.

Roumimper, N. 2017. *Mesh Navigation Through Jumping*. Master’s thesis, Utrecht University.

Savinov, N.; Dosovitskiy, A.; and Koltun, V. 2018. Semi-parametric topological memory for navigation. *arXiv preprint arXiv:1803.00653*.

Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; Dieleman, S.; Grewe, D.; Nham, J.; Kalchbrenner, N.; Sutskever, I.; Lillicrap, T.; Leach, M.; Kavukcuoglu, K.; Graepel, T.; and Hassabis, D. 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587): 484–489.

Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017. Mastering the game of go without human knowledge. *nature*, 550(7676): 354–359.

Snook, G. 2000. Simplified 3D movement and pathfinding using navigation meshes. In *Game Programming Gems*, 288–304. Charles River Media.

Tesauro, G. 1994. TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural computation*, 6(2): 215–219.

Wani, S.; Patel, S.; Jain, U.; Chang, A.; and Savva, M. 2020. MultiON: Benchmarking Semantic Map Memory using Multi-Object Navigation. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 9700–9712. Curran Associates, Inc.

Wayne, G.; Hung, C.-C.; Amos, D.; Mirza, M.; Ahuja, A.; Grabska-Barwinska, A.; Rae, J.; Mirowski, P.; Leibo, J. Z.; Santoro, A.; et al. 2018. Unsupervised predictive memory in a goal-directed agent. *arXiv preprint arXiv:1803.10760*.

Wydmuch, M.; Kempka, M.; and Jaśkowski, W. 2018. ViZ-Doom Competitions: Playing Doom from Pixels. *arXiv preprint arXiv:1809.03470*.

Xia, F.; R. Zamir, A.; He, Z.-Y.; Sax, A.; Malik, J.; and Savarese, S. 2018. Gibson env: real-world perception for embodied agents. In *Computer Vision and Pattern Recognition (CVPR), 2018 IEEE Conference on*. IEEE.

Zhang, J.; Tai, L.; Boedecker, J.; Burgard, W.; and Liu, M. 2017. Neural Slam: Learning to explore with external memory. *arXiv preprint arXiv:1706.09520*.