

Obstacle Tower Without Human Demonstrations: How Far a Deep Feed-Forward Network Goes with Reinforcement Learning

Marco Pleines^{*}, Jenia Jitsev[†], Mike Preuss[‡], Frank Zimmer[§]

^{*}*Faculty of Computer Science, Technische Universität Dortmund, Dortmund, Germany*
marco.pleines@tu-dortmund.de

[†]*Juelich Supercomputing Center (JSC), Helmholtz Research Center Juelich, Juelich, Germany*
j.jitsev@fz-juelich.de

[‡]*LIACS, Universiteit Leiden, Leiden/NL*
m.preuss@liacs.leidenuniv.nl

[§]*Faculty of Communication and Environment, Hochschule Rhein-Waal, Kamp-Lintfort, Germany*
frank.zimmer@hochschule-rhein-waal.de

Abstract—The Obstacle Tower Challenge is the task to master a procedurally generated chain of levels that subsequently get harder to complete. Whereas the most top performing entries of last year’s competition used human demonstrations or reward shaping to learn how to cope with the challenge, we present an approach that performed competitively (placed 7th) but starts completely from scratch by means of Deep Reinforcement Learning with a relatively simple feed-forward deep network structure. We especially look at the generalization performance of the taken approach concerning different seeds and various visual themes that have become available after the competition, and investigate where the agent fails and why. Note that our approach does not possess a short-term memory like employing recurrent hidden states. With this work, we hope to contribute to a better understanding of what is possible with a relatively simple, flexible solution that can be applied to learning in environments featuring complex 3D visual input where the abstract task structure itself is still fairly simple.

Index Terms—deep reinforcement learning, artificial general intelligence, visual debugging

I. INTRODUCTION

Deep Reinforcement Learning (DRL) had tremendous successes during the last years. Very often it has been employed as direct end-to-end learning from high-dimensional raw pixel images for difficult tasks such as playing Atari games [1], Doom [2], or the more cooperative games capture-the-flag [3] and Dota 2 [4]. Such game environments are also more and more combined with additional information beyond pure pixels as for hide-and-peek [5] and AlphaStar [6]. The latter one plays the complex real-time strategy game StarCraft II on human grandmaster level, a milestone that has been presumed not reachable for years not long ago.

Concerning Atari games, feed-forward convolutional neural networks (FFCNN) can be successfully trained to solve those using basic policy gradient methods as REINFORCE or value based ones as Q-Learning. For more complex environments featuring 3D worlds or sparse long horizon rewards, more advanced network architectures are considered, like various

convolutional recurrent neural networks. This paper shows that up to a certain degree it is also possible to solve complex 3D environments using a rather simple FFCNN when training those with state-of-the art DRL algorithms like Proximal Policy Optimization (PPO) [7].

Furthermore, evaluating a model’s generalization capability, obtained through DRL, usually lacks in a clear split between training and testing phases due to the widely used standard benchmark game environments that suffer from fixed structures. To overcome this issue, environments that utilize procedural content generation (PCG) approaches shall be employed here. Hence, the model can be trained and evaluated on distinct seeds, each defining a unique instance of an environment, guaranteeing a clear split.

One example for a procedurally generated environment is Obstacle Tower (OT) [8]. In OT, the agent is challenged in terms of vision, control, planning, and generalization, while its goal is to ascend a tower of floors that get more difficult as the agent progresses [8]. The first 5 floors do not involve any special puzzles for the agent. After that, the agent has to find keys to get past locked doors. Once floor 10 is reached, a difficult sokoban puzzle is introduced. In 2019, the developers of OT held a challenge where the top entries moved beyond floor 10 only with the help of domain knowledge such as adding human demonstrations to the training data [9].

In this work, we demonstrate that OT can be solved up to floor 10 using a rather simple FFCNN when trained with advanced DRL techniques (PPO) without the use of human demonstrations. In the original paper, which introduced OT, the highest floor reached using a FFCNN and PPO was 5. Reaching Level 10 is quite challenging given the rather complex OT 3D world environment and tasks like key-door puzzles or double jumps introduced from level 5 on. Overall, our FFCNN-PPO algorithm performed competitive in the official

OT challenge, ranking 7th¹.

To further study generalization of our learning algorithm, we train our model on 3 and evaluate on 2 different visual themes (or skins) that are offered by the OT environment (Figure 1). Training on selected skin sets and testing on ones the algorithm has not seen before allows us to draw conclusions about its generalization capability with regard to the environment’s vision challenge. While we can state that the FFCNN is able to cope well with different visual themes during training without collapsing, we observe a clear drop in testing performance on the novel themes, which shows obvious generalization limits of the FFCNN.

This paper proceeds as follows: the next section highlights the measures taken by the top competitors of the OT challenge, while showcasing other work in the broader context of generalization. Section 2 describes the taken approach concerning its details about the environment configuration, the architecture of the FFCNN, and the PPO training. After that, the conducted experiment is described. Results are shown for the generalization performance of the trained model as well as a detailed examination of the learned policy. To further elaborate our achievements, section 5 discusses the observed peculiarities of our results and approach. The last section concludes our findings and describes consecutive work.

II. RELATED WORK

We start by relating our work to the top competitors of the OT challenge before putting it into the wider context of generalization. The challenge’s organizers state that average human performance is around floor 15 [8], and this has only been surpassed by the winner (average floor 19.4) and the runner-up (average floor 16) of the challenge. The top 4 entries were able to get past floor 10 by means of a PPO approach that was augmented with human demonstrations [9]. Common measures, shared by several approaches, to reduce the problem complexity are:

- **Reduced action space:** OT features a multi-discrete action space containing 3 subspaces comprising 11 actions in total. If a regular single discrete space is used, 54 possible action combinations become available. Many of those may not be necessary for achieving good results on the one hand. On the other hand, large action numbers also make the learning problem more difficult. It becomes conceivably easier to learn with only around 10 usable actions, by also preselecting reasonable action combinations of those that are potentially possible.
- **Use of memory cells:** simple neural networks, like the one employed in this work, have no means of treating developments over time in a meaningful way. However, this is possible with GRU or LSTM cells which are used by several related approaches.
- **Frame stacks:** adding past frames to the agent’s observation is another popular measure, especially if no memory cell is used.

Environment	Steps per Second
Procgen CoinRun	5375
MiniGrid FourRoom [23]	634
Atari Breakout [24]	4041
Obstacle Tower (100 floors)	43
Obstacle Tower (10 floors)	51

TABLE I
THE NUMBER OF STEPS PER SECOND IS AVERAGED OVER 100 EPISODES PER ENVIRONMENT.

- **Data augmentation:** such techniques improve data quality and simplify the learning process. Mirroring left to right and vice versa is one example.
- **Reward shaping:** the learning process can be sped up by providing more and better suited rewards to the agent.

As of now, the 1st and 4th place shared technical details via blog posts and a preprint [10], [11]. The competition winner, Alex Nichol, trained a classifier to help the agent detect various environment entities such as doors and keys. Throughout 50 consecutive frames, Nichol adds the received reward, the executed action, the possession of the key, and the classifier’s output to the agent’s observation space. Instead of using an entropy bonus for exploration, he applies KL-Divergence to push the agent’s policy towards a prior, which was trained with behavioral cloning beforehand. This way the agent was able to solve the complex sokoban puzzle.

One of the main motivations for setting up the OT challenge was to see how learning methods can deal with generalization. By exposing learning agents to highly variable environments, overfitting [12], [13] shall be reduced and the agents should focus on learning the underlying major factors and not specific details of a single problem instance that are often misleading in general cases. Encouraging generalization of the learning process by injecting more diversity into training environments has also been the main motivation to set up the GVGAI [14] environment. There, the diverse set of games and levels has been created manually at first. However, the setup also blends well with procedural content generation (PCG) techniques [15] which are designed to provide controllable content variations that can be introduced systematically and automatized.

Several works have investigated how PCG can be used in order to strengthen generalization [16], [17]. Learning environments, like Procgen [18], explicitly focus on enabling this and offer benchmarks for testing the generalization ability of RL algorithms. Other approaches for achieving stronger generalization consist of, for example, adding different types of memory to the neural networks [19], inject noise [20], randomize the network’s feature space [21], or randomize and distort the raw visual input from the training domain [22].

III. APPROACH

Before elaborating the taken approach in greater detail, it is important to show that the simulation speed of the OT environment limits the number of training sessions and experiments. As seen in Table I, OT runs much slower than other environments. The benchmarks were run on an Ubuntu

¹<https://youtu.be/P2rBDHBHxcM>



Fig. 1. The 5 visual themes featured by Obstacle Tower from left to right: Ancient, Industrial, Modern, Moorish, and Future.

machine (nVidia Quadro K1200, 2x Intel Xeon E5-2640v4 CPUs, 64GB RAM). In order to speed up the training process, the floor generation is limited to 10 floors. This means that once the agent completes floor 9, the episode terminates with the result, that the agent reached floor 10. Therefore, the difficult sokoban puzzle is not part of the training.

A. Environment Properties

Besides limiting the number of generated floors, we apply further changes to the environment. To simplify the challenge, the agent executes one action for two consecutive frames (i.e. frame skipping). Its action space is reduced from 11 to 7 actions, which are represented by 3 subspaces:

- Subspace A:
 - No action
 - Move forward
- Subspace B:
 - No action
 - Jump
- Subspace C:
 - No action
 - Rotate left
 - Rotate right

Moving left, right and backward are removed from the original action space, because these are not mandatory to solve OT. At last, the reward function remains unchanged:

- +1.0 for reaching the next floor,
- +0.1 for opening a door,
- and +0.1 for collecting a key.

Concerning the observation space, the agent receives the current and the past two visual observations of the environment. The stacked image frames shall enable the agent to derive its velocity and acceleration. Turning the frames into gray-scale is not considered, because it might raise the difficulty for the agent to identify a key. However, the RGB image frames were unintentionally divided by 255 twice instead of just once. Finally, the agent receives a vector of game state variables, featuring the remaining time and whether the agent has a key or not.

B. Model Architecture

The environment properties affect the architecture of the trained FFCNN, which is illustrated by Figure 2. The model receives few ($n=3$) temporally stacked image frames (each with 3 RGB color channels) and a vector of game state variables (has key, remaining time) as input. Once the visual observation is processed by the convolutional layers (i.e. visual encoder), the flattened results and the game state vector input

Training Parameter	Value
Discount Factor	0.99
Lamda (GAE)	0.95
Value Function Coefficient	0.5
Entropy Bonus Coefficient	0.01
PPO Updates	50,000
Epochs	4
Number of Environments	16
Trajectory Length	8,192
Minibatches	4
Learning Rate	3.25e-4
Clip Range	0.2
Activations	ReLU
Optimizer	Adam

TABLE II
TRAINING PARAMETERS

are concatenated. The concatenation is then fed into a fully connected hidden layer, after which the neural net is split into two branches. Those follow then the actor-critic architecture design [25], [26]. One branch is used for the value function that predicts the expected long-term reward with a single scalar output and the another one represents the policy signaling action probabilities. Each branch contains a fully connected hidden layer. Due to this setup, the value function and the policy share parameters of a common hidden layer while also maintaining a separate one in their respective branch. As the action space is decomposed into 3 subspaces, the policy is composed of 3 branches as inspired by action branching [27]. Each action branch in turn contains different numbers of available actions that predefine generic reasonable combinations (see Figure 2). Hence, the model supports multi-discrete action spaces and avoids the necessity of implementing all action combinations causing a higher dimensional output.

C. PPO Training

The implementation² of the training algorithm PPO is closely related to its publication from Schulman et al. (2017). Generalized advantage estimation (GAE) is used by the value loss function. The objectives for the value function and the policy are clipped. Further, the final loss function comprises an entropy bonus term to encourage exploration. However, action branching requires one small adjustment to the policy and the entropy bonus. For each policy branch, all outputs are concatenated leading to a flattened view of these. These are then processed by the loss function without further adjustments

²<https://github.com/MarcoMeter/neroRL>

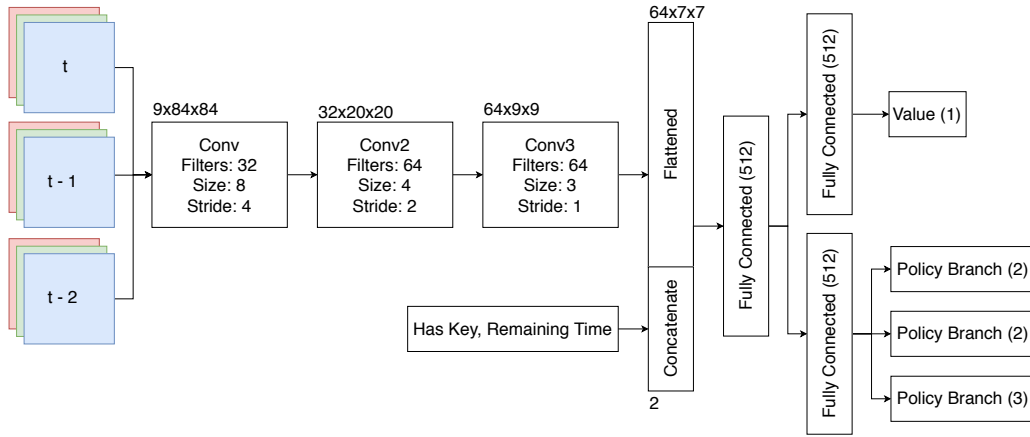


Fig. 2. The architecture of the utilized feed-forward convolutional neural network.

needed. Concerning the entropy bonus, the mean of the policy branches' entropies is made use of.

Because of the high computational cost of running the OT environment, training parameters cannot undergo further optimization. The ones provided by Table II are derived from the experiences made during the OT challenge. It has to be noted that the learning rate, the entropy bonus coefficient, and the clip range decay linearly dependent on the remaining PPO updates. One PPO update optimizes the model using 4 minibatches per epoch across the training data, which was collected by the agents sampling actions from the current policy. The intention behind annealing training parameters is to boost the agent's training performance in the beginning and then later to take smaller steps to fine-tune its policy.

IV. EXPERIMENTAL ANALYSIS

Under the standard conditions of the OT environment, every 10th floor enables another visual theme to confront the agent. So starting from floor 0 the agent only faces the ancient theme. Once floor 10 is reached, the ancient and the moorish themes are alternated randomly. Throughout the OT competition, 100 training seeds were available, while 5 distinct seeds were kept hidden to evaluate the agent's ability to generalize. Therefore, the agents faced only the ancient and moorish theme during the competition's evaluation. In order to fully assess the agent's generalization capability, we utilize three skins (ancient, industrial, and modern) for training, while leaving out the other ones (moorish, and future) for evaluation. Due to the simulation speed constraints, we only train on this set as denoted by Figure 1.

A. Generalization Performance

We ran three training sessions using the same training parameters where the agent faced 100 seeds, while all three training skins were randomly alternated. The only other modification to the environment is the limited floor number of 9. While training for 50,000 PPO updates, every 200th update, the model was evaluated on 5 distinct seeds, which were not used during training. These seeds were evaluated 3 times

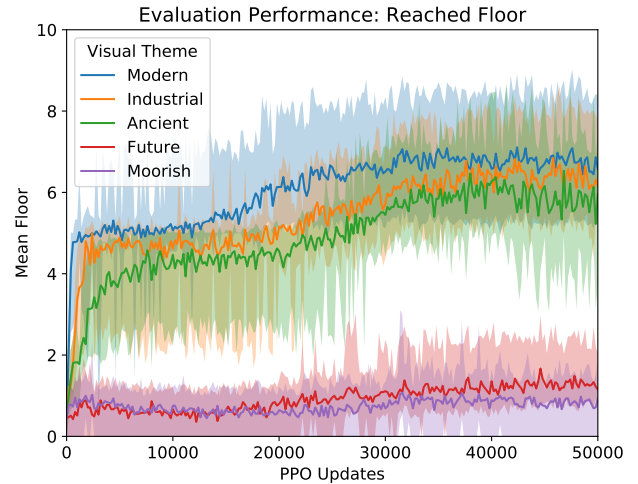


Fig. 3. The achieved mean floor on all 5 visual themes across three training runs. Less opaque colored regions visualize their respective variance based on the asymmetric deviation. The moorish and the future theme were not part of the training.




Fig. 4. This bar chart visualizes the number of times an episode terminated at a certain floor on each theme set given the final model after training. A much lower performance is observed for the evaluation theme set.

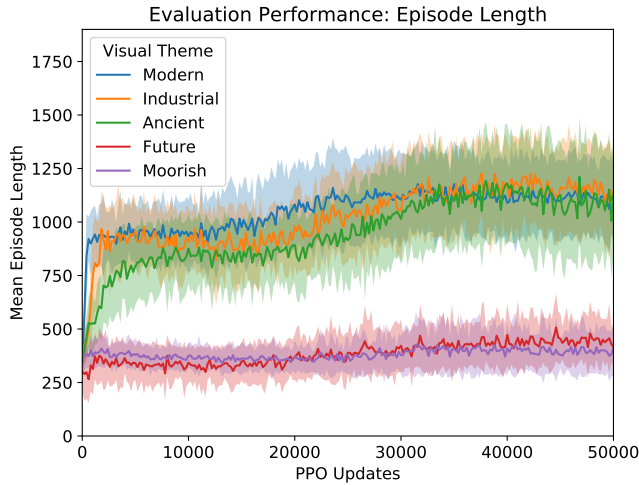


Fig. 5. The achieved mean episode length on all 5 visual themes across three training runs. Less opaque colored regions visualize their respective variance based on the standard deviation. The moorish and the future theme were not part of the training.

for all 5 visual themes. Thus, the agent was evaluated by 15 episodes per theme. Multiplied by three training runs, the results of 75 episodes were collected for each interval. Figure 3 shows the achieved mean floor for each particular skin, whereas Figure 5 illustrates the mean episode length.

It can be observed that the agent poorly performs on the moorish and the industrial theme, which were not seen by the agent during training. By the end of the training, the agent reaches about mean floor 1.15 and 0.93 on these themes. A different picture emerges on the training themes that the agent has seen during training. On these, the agent’s performance converges at about mean floor 6, although it did not encounter these seeds during training. On the training seeds, the agent achieved a mean floor of about 8.66. Overall, a high variance can be observed. For instance on the ancient theme given the final model after training, the agent got stuck at floor 0 on one of the 15 evaluation episodes. This variance becomes clearer by examining the number of times an episode terminated on a certain floor. In the provided bar chart (Figure 4), most of the time the episode ends on floor 5 on the training themes. Terminations on floor 0, 2, 3, 9, and 10 can be understood as outliers. Thus, there is always a chance that the agent accomplishes all 9 floors or gets stuck already at the first one, even if the same seed is tried over and over again.

During the first 10,000 PPO updates, the agent rapidly learns to reach floor 5 on the training skins. Due to the introduction of the key puzzle tasks, the agent’s policy is stuck for approximately 10,000 PPO updates on a plateau. After that, the policy slowly improves over time.

Concerning the episode length, it is correlated with the successful tower ascend of the agent. Reaching a new floor is rewarded with a time extension. As the episode ends once floor 9 is finished, the episode length should decrease as the policy




Fig. 6. The agent immediately solves the key puzzle and exits the floor.

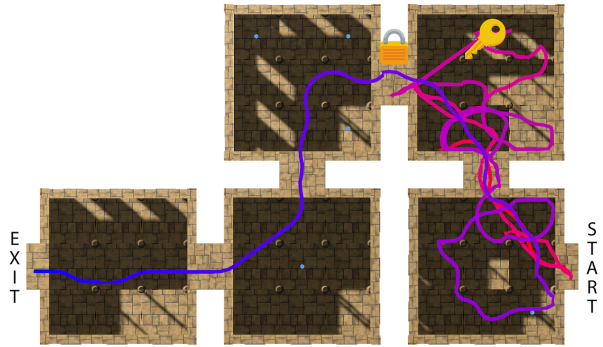


Fig. 7. The agent wanders around the entire accessible rooms, but eventually collects the key and exits the floor.

improves. A slight trend for such decrease can be observed on the training skins, indicating that the agent becomes more proficient during its tower ascend.

B. Agent Behavior

Multiple observations can be done by watching how the agent utilizes its learned policy to solve OT on the training themes. First of all, it can be noticed that the agent’s locomotion is rather shaky. While moving forward, the agent tends to continuously execute the actions “rotate left” and “rotate right”. Another observation is that the agent likes to go through doors in general, no matter whether required or not for the current task. Therefore, it may also happen that the agent moves all the way back to the beginning instead of heading to the floor exit, but still the agent might be able to finish the current floor. Regarding the jump action, the agent usually jumps when required, like in situations where the agent approaches an obstacle. On very rare occasions, the agent gets stuck on corners of a door or on similar environmental structures while experiencing a rewarding stimulus in its visual field. In this situation, the agent keeps uselessly moving forward and is therefore not able to finish the current floor.

An important subtask, tackled by the agent, is the key puzzle, which can be partially visualized by the agent’s taken path as seen in Figures 6, 7, and 8. The drawn path of the agent

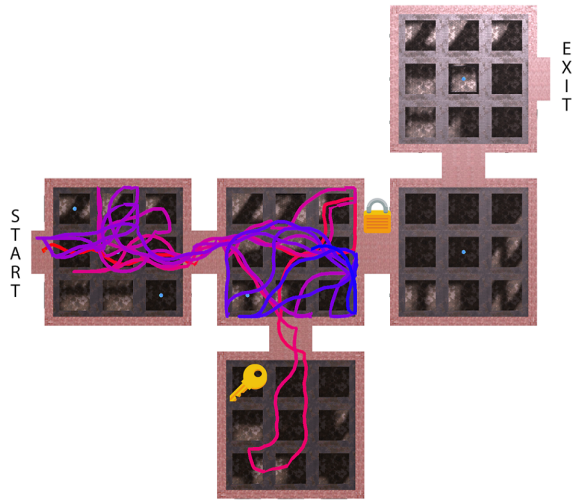


Fig. 8. The agent wanders around the entire first two rooms. The key room is visited once, but the key was not collected. Thus, the floor was not solved.

starts out red and turns blue over time. Further, the positions of the key and the door are marked by their respective icons. Watching the agent’s behavior, the paths shown by Figure 7, and 8 are more likely to occur. Most of the time the agent walks across the entire accessible space to grab the key by chance. Also, it can be observed that the agent tries to get past locked doors without being in possession of the key. It may happen that the agent attempts to pass the same locked door numerous times. Such an attempt is also present in the more goal directed path as seen in Figure 6. One important observation is that sometimes the agent seems to ignore the key, even if the key is very close and in the vision field of the agent. However, we observed that the agent is able to solve the difficult key puzzle (Figure 9) that requires the precise execution of two consecutive jumps. While running 150 episodes per theme and per 3 unseen seeds, these are the measured probabilities for the agent to successfully master this double jump key puzzle using the final model: Ancient: 0%, Industrial: 30%, Modern: 33%. Based on a shallow search, the agent starts to show the ability of solving this puzzle after about 20,000 PPO updates, with a very low chance of success at the beginning. With learning progress, the success chance rises substantially on the modern and industrial themes, while staying zero for the ancient one. Besides the agent’s imminent goals, blue time orbs, which extend the agent’s time upon collection, are not really of interest to the agent. In general, a different behavior is not observed in the case of the agent running out of time.

V. DISCUSSION

The performance shown in the evaluation hints that while the agent is able to generalize well on novel seeds on the training themes, it fails to generalize on previously unseen themes. Therefore, the generalization capability of the employed FFCNN is limited. On the one hand, it cannot cope well with such strong visual variations as challenged by different



Fig. 9. These are double jump modules used by the ancient, industrial, and modern theme, which were explored by the agent during training.

OT themes. On the other hand, the utilized approach manages to train a policy that is able to cope with environment shifts during training. To some degree, concerning unseen seeds given the 3 training themes, the rather challenging key puzzles that require action sequences like the double jump are solved. However, high variances and extreme outliers were observed. These observations and peculiarities of the agent’s behavior are discussed and put into context using the subsequent sections to further elaborate the outcome of our work.

A. Limits of Learned Representations

Multiple results indicate that the agent’s representations are not rich enough to perceive relevant properties and structures of its environment. As the agent’s visual observation is limited to the current and the past two image frames, the policy tells the agent to continuously rotate left and right, while moving forward. As recurrent hidden states, summarizing previous perceptual history, are missing in the FFCNN, the agent only lives in the moment and thus strives to capture as much relevant information as possible. Lacking short-term memory, it has to exploit the capacity of the available RGB frames, which is accomplished by its shaky locomotion behavior. Besides the game state variables, the image frame stack is the only data available to the agent to make policy decisions. Therefore, the agent can only react to immediate issues that are contained in its momentary input. For example, whenever a door is spotted, the agent’s imminent behavior is to approach this door no matter where that door leads to. This may cause the agent to run into doors that are disadvantageous to the agent’s goals. With more information from the short-term past, the agent would gain more potential to improve its navigation.

Another related problem is concerned with the agent being unable to pursue the key or to establish a clear key-door connection. As previously shown, the agent wanders around entire spaces until it picks up the required key by chance. By retaining information from the past, such inefficient behavior could be mitigated in a way that the agent does not visit places that he has already visited before. It could also execute goal-oriented behaviors, for instance recalling positions of an already encountered door after gathering a respective key.

Furthermore, the agent sometimes seems to ignore keys that are close to him. Due to the utilized frame skip, it might be possible that the key is not present on any of the visual observations made by the agent. This becomes more problematic as the key rotates continuously. Especially due to

the low resolution of the visual observations, the key could be missed, because its current visible surface is too small.

Mostly, these issues can be traced back to a lack of short-term memory. One potential measure towards resolving this problem, without introducing recurrent cells, is to increase the number of image frames and adding the skipped frames to the agent’s observation. Though, collecting more frames raises the dimensionality of the visual observations and makes it more difficult for the training process to efficiently make use of it.

B. Impact of Game State Input Variables

Frequently, the agent tries to get past locked doors without being in possession of a key. As the agent directly receives the information whether he has a key or not, it should be able to learn the relationship between these components. A similar problem can be observed for the time management of the agent. It neither collects time orbs on purpose nor changes its behavior while running out of time. It could be possible that the two game state variables alone cannot affect the policy noticeably, because these are just two input features that share the same layer with 3136 dimensional outputs delivered by the visual encoder. In the future, it has to be examined if this hypothesis applies and how to cope with it. One solution might be to project the game state input to as many units as used by the output of the visual encoder, while sharing the weights among all of those replicated units to avoid parameter explosion. Solving this issue may lead to a more robust and better performing policy that is able to consistently exploit the key-door and time bonus relationships.

C. Visual Encoder Complexity

Another question emerges concerning the complexity of the visual encoder. As seen previously, the double jump key puzzle is only solvable for the industrial and modern theme. The chance of success for the ancient theme is zero. Further, the agent’s performance on this theme is in general inferior to the other ones as illustrated by Figure 3. One reason might be that it is more difficult to distinguish the key from ancient themed floor components visually, as in this theme, everything is dyed in brownish hues which impairs object detection.

However, it seems that the difference in performance between the themes becomes expressed very early on in the training, where key puzzles are not introduced yet. As we performed 3 distinct training runs, we can observe the same scheme each time. Performance improvements on the modern theme kick in early and rapidly, followed by the industrial and ancient one. This suggests that this distinct performance is due to differences in some generic visual properties of the training themes. This visual discrepancy results in a dominating theme that is generally solved better than the others.

Now, it can be argued that the model is more likely to extract features from the dominant themes, because those might be more distinguishable, e.g in terms of a stronger contrast. A more general take on this would be to assume a lack of capacity that impairs the visual encoder to deal with so many distinct themes simultaneously during training.

D. Annealing Schedule for Training Parameters

A problem which was not discussed yet is the agent getting stuck in certain situations by repeating the same action endlessly. As the entropy bonus coefficient linearly declines over time without a lower bound threshold, the policy may become too deterministic. If the agent gets stuck at a corner of an exit door, the probability of the action “moving forward” is extremely high compared to the other ones. Thus, the agent can barely escape this situation, which is a possible explanation for the observed negative outliers.

The approach of using annealing training parameters can be questioned in general. Its intention is to optimize the model stronger in the beginning and fine-tune it towards the end using a lower bound threshold. This bound was not implemented in this approach and therefore the learning rate, the clip range, and the entropy bonus coefficient equal 0 at the end of the training. As new subtasks are continually introduced in the environment dependent on reaching higher floors, the agent is required to explore new ways to act, based on how its performance is changing due to environmental shifts. Making training parameters subject to adaptation from data would further increase the agent’s ability to deal with changing tasks.

VI. CONCLUSION AND FUTURE WORK

While our original approach performed competitively in the OT challenge, this paper shows that the underlying - rather simple - FFCNN can solve novel seeds on three visual themes, which were faced by the agent during training. However, the trained model has clear limitations concerning its weak performance across the two left-out visual themes. Therefore, the visual encoder of the model does not generalize well to unseen visual themes. By analyzing the agent’s behavior on the training themes, it becomes apparent that the agent solves its task in stimulus-response schemes. This is due to the limited observation space of the agent, where the agent operates on the current and the past two image frames. By stacking more frames or adding a memory cell to the FFCNN a more proficient agent behavior can be expected.

To improve the generalization capability of the visual encoder, its capacity could be increased by more sophisticated network architectures containing components such as residual blocks [28] and attention mechanisms [29]. It is also advisable to put further effort into understanding the agent’s policy. Using visualization techniques like layer-wise relevance propagation [30], it could be possible to derive more insights from saliency maps, which show what inputs are of relevance to the agent [31]. Another concern for further research are adaptive training parameters, like the learning rate. This is especially challenging for environments like OT, that encompass multiple subtasks introduced while progressing in the environment.

One drawback to cope with is the slow simulation speed of OT, because it constraints rapid experimenting. Besides optimizing the environment itself, the developed training process could be made more efficient by augmenting the collected training data. A further option to accelerate learning experiments is to use distributed training to run DRL algorithm

on multiple machines [32], [33]. Another approach would be to introduce an adaptive environment that allows the agent to collect the data that is most useful at the current point of training. For example, if the agent struggles to learn the double jump key puzzle, the puzzle could be made more likely to be visited by the agent, which in general may lead towards exploring different techniques for adaptive sampling.

Overall, following the envisioned research directions may result in potent learning algorithms that are also able to cope with more challenging subtasks from scratch, like the difficult sokoban puzzle, that has only been solved using human demonstrations so far.

REFERENCES

- [1] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, Feb 2015.
- [2] M. Wydmuch, M. Kempka, and W. Jaskowski, "Vizdoom competitions: Playing doom from pixels," *IEEE Trans. Games*, vol. 11, no. 3, pp. 248–259, 2019.
- [3] M. Jaderberg, W. M. Czarnecki, I. Dunning, L. Marris, G. Lever, A. G. Castañeda, C. Beattie, N. C. Rabinowitz, A. S. Morcos, A. Ruderman, N. Sonnerat, T. Green, L. Deason, J. Z. Leibo, D. Silver, D. Hassabis, K. Kavukcuoglu, and T. Graepel, "Human-level performance in 3d multiplayer games with population-based reinforcement learning," *Science (New York, N.Y.)*, vol. 364, pp. 859–865, May 2019.
- [4] OpenAI, C. Berner, G. Brockman, B. Chan, V. Cheung, P. Debiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse, R. Józefowicz, S. Gray, C. Olsson, J. Pachocki, M. Petrov, H. P. de Oliveira Pinto, J. Raiman, T. Salimans, J. Schlatter, J. Schneider, S. Sidor, I. Sutskever, J. Tang, F. Wolski, and S. Zhang, "Dota 2 with large scale deep reinforcement learning," *arXiv:1912.06680*, 2019.
- [5] B. Baker, I. Kanitscheider, T. Markov, Y. Wu, G. Powell, B. McGrew, and I. Mordatch, "Emergent tool use from multi-agent autotutorials," *arXiv:1909.07528*, 2019.
- [6] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, J. Oh, D. Horgan, M. Kroiss, I. Danihelka, A. Huang, L. Sifre, T. Cai, J. P. Agapiou, M. Jaderberg, A. S. Vechnyevets, R. Leblond, T. Pohlen, V. Dalibard, D. Budden, Y. Sulsky, J. Molloy, T. L. Paine, C. Gulcehre, Z. Wang, T. Pfaff, Y. Wu, R. Ring, D. Yogatama, D. Wünsch, K. McKinney, O. Smith, T. Schaul, T. Lillicrap, K. Kavukcuoglu, D. Hassabis, C. Apps, and D. Silver, "Grandmaster level in starcraft ii using multi-agent reinforcement learning," *Nature*, Oct. 2019.
- [7] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv:1707.06347*, 2017.
- [8] A. Juliani, A. Khalifa, V. Berges, J. Harper, E. Teng, H. Henry, A. Crespi, J. Togelius, and D. Lange, "Obstacle tower: A generalization challenge in vision, control, and planning," in *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI 2019*, 2019, pp. 2684–2691.
- [9] A. Juliani and J. Shih, "Announcing the obstacle tower challenge winners and open source release," 2019, available at <https://blogs.unity3d.com/2019/08/07/announcing-the-obstacle-tower-challenge-winners-and-open-source-release/> retrieved March 20, 2020.
- [10] A. Nichol, "Competing in the obstacle tower challenge," 2019, available at <https://blog.aqnichol.com/2019/07/24/competing-in-the-obstacle-tower-challenge/> retrieved March 24, 2020.
- [11] J. Booth, "PPO dash: Improving generalization in deep reinforcement learning," *arXiv:1907.06704*, 2019.
- [12] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.
- [13] C. Zhang, O. Vinyals, R. Munos, and S. Bengio, "A study on overfitting in deep reinforcement learning," *arXiv:1907.06704*, 2018.
- [14] D. Perez-Liebana, J. Liu, A. Khalifa, R. D. Gaina, J. Togelius, and S. M. Lucas, "General video game ai: A multitrack framework for evaluating agents, games, and content generation algorithms," *IEEE Transactions on Games*, vol. 11, no. 3, pp. 195–214, 2019.
- [15] N. Shaker, J. Togelius, and M. J. Nelson, *Procedural Content Generation in Games*. Springer, 2016.
- [16] N. Justesen, R. R. Torrado, P. Bontrager, A. Khalifa, J. Togelius, and S. Risi, "Illuminating generalization in deep reinforcement learning through procedural level generation," *arXiv:1806.10729*, 2018.
- [17] K. Cobbe, O. Klimov, C. Hesse, T. Kim, and J. Schulman, "Quantifying generalization in reinforcement learning," in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, ser. Proceedings of Machine Learning Research, vol. 97. PMLR, 2019, pp. 1282–1289.
- [18] K. Cobbe, C. Hesse, J. Hilton, and J. Schulman, "Leveraging procedural generation to benchmark reinforcement learning," *arXiv:1912.01588*.
- [19] M. Fortunato, M. Tan, R. Faulkner, S. Hansen, A. Puigdomènech Badia, G. Buttimore, C. Deck, J. Z. Leibo, and C. Blundell, "Generalization of reinforcement learners with working and episodic memory," in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 12469–12478.
- [20] M. Igl, K. Ciosek, Y. Li, S. Tschjatschek, C. Zhang, S. Devlin, and K. Hofmann, "Generalization in reinforcement learning with selective noise injection and information bottleneck," in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 13978–13990.
- [21] K. Lee, K. Lee, J. Shin, and H. Lee, "Network randomization: A simple technique for generalization in deep reinforcement learning," *arXiv:1910.05396*, 2019.
- [22] S. James, P. Wohlhart, M. Kalakrishnan, D. Kalashnikov, A. Irpan, J. Ibarz, S. Levine, R. Hadsell, and K. Bousmalis, "Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12627–12637.
- [23] M. Chevalier-Boisvert, L. Willems, and S. Pal, "Minimalistic gridworld environment for openai gym," <https://github.com/maximecb/gym-mini-grid>, 2018.
- [24] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling, "The arcade learning environment: An evaluation platform for general agents," *J. Artif. Intell. Res.*, vol. 47, pp. 253–279, 2013.
- [25] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [26] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *International conference on machine learning*, 2016, pp. 1928–1937.
- [27] A. Tavakoli, F. Pardo, and P. Kormushev, "Action branching architectures for deep reinforcement learning," in *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*. AAAI Press, 2018, pp. 4131–4138.
- [28] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, Jun. 2016, pp. 770–778.
- [29] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, 2017, pp. 5998–6008.
- [30] G. Montavon, A. Binder, S. Lapuschkin, W. Samek, and K. Müller, "Layer-wise relevance propagation: An overview," in *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, ser. Lecture Notes in Computer Science. Springer, 2019, vol. 11700, pp. 193–209.
- [31] S. Greydanus, A. Koul, J. Dodge, and A. Fern, "Visualizing and understanding atari agents," in *International Conference on Machine Learning*, 2018, pp. 1792–1801.
- [32] L. Espeholt, H. Soyer, R. Munos, K. Simonyan, V. Mnih, T. Ward, Y. Doron, V. Firoiu, T. Harley, I. Dunning, S. Legg, and K. Kavukcuoglu, "IMPALA: scalable distributed deep-rl with importance weighted actor-learner architectures," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 2018, pp. 1406–1415.
- [33] L. Espeholt, R. Marinier, P. Stanczyk, K. Wang, and M. Michalski, "Seed rl: Scalable and efficient deep-rl with accelerated central inference," in *International Conference on Learning Representations*, 2020.