

Challenges of Context and Time in Reinforcement Learning: Introducing Space Fortress as a Benchmark

Akshat Agarwal
Robotics Institute
Carnegie Mellon University
aa7@cmu.edu

Ryan Hope
Department of Psychology
Carnegie Mellon University
rmh3093@gmail.com

Katia Sycara
Robotics Institute
Carnegie Mellon University
katia@cs.cmu.edu

Abstract

Research in deep reinforcement learning (RL) has coalesced around improving performance on benchmarks like the Arcade Learning Environment. However, these benchmarks do not emphasize two important characteristics that are often present in real-world domains: requirement of changing strategy conditioned on latent contexts, and temporal sensitivity. As a result, research in RL has not given these challenges their due, resulting in algorithms which do not understand critical changes in context, and have little notion of real world time. This paper introduces the game of Space Fortress as a RL benchmark which specifically targets these characteristics. We show that existing state-of-the-art RL algorithms are unable to learn to play the Space Fortress game, and then confirm that this poor performance is due to the RL algorithms' context insensitivity. We also identify independent axes along which to vary context and temporal sensitivity, allowing Space Fortress to be used as a testbed for understanding both characteristics in combination and also in isolation. We release Space Fortress as an open-source Gym environment.

1 Introduction

Recent advances in computer vision (Krizhevsky, Sutskever, and Hinton 2012) and natural language processing (Sutskever, Vinyals, and Le 2014) can be attributed to the advent of deep learning and the presence of robust benchmarks to quantitatively measure progress, such as the ImageNet challenge (Russakovsky et al. 2015). In the last few years, neural network-based function approximation has also proven successful in reinforcement learning, with AI agents now able to perform at superhuman levels in games like Go (Silver et al. 2016) and the Atari (Mnih et al. 2015) suite. Once again, research in Deep RL has been steered by the establishment of benchmarks like the Arcade Learning Environment (Bellemare et al. 2013), along with the OpenAI Gym interface (Brockman et al. 2016), which has been widely adopted by the research community.

These benchmarks have not emphasized 2 challenging characteristics: (a) abrupt context-dependent switching of strategy and (b) temporal sensitivity. For agents to operate in the real world, they need to be able to switch behaviors

very abruptly, often dependent on broader contexts which might never be revealed to the agent. This necessitates (i) learning to identify critical points where behavior needs to change, and (ii) learning the different behaviors required in each context. For example, driving at 60 kmph would be acceptable and safe on a foggy day with regular roads (context), but would be dangerous on a foggy day with icy roads (different context; difficult to observe the change if the car is not equipped with a sensor for measuring road condition). A shift in context can happen abruptly with little to no change in the observed state, but necessitate completely different strategies.

Temporal sensitivity can be defined as the agent having an understanding of time as a quantity independent of its rate of decision making, i.e., as something that is always ticking. Currently, most RL agents make decisions at a fixed frame-rate (some papers skip 1 to 4 frames randomly). With respect to the real world, this frame rate could be 4fps or 40fps, the agent would neither know nor care. However, we argue that this is not acceptable for real-world dynamic environments where decisions might have to be adaptively taken either very quickly or very slowly. To enable the agent to vary its speed of decision making, it must understand that the environment moves at a pace independent of its own.

In this paper, we introduce a challenging RL environment based on Space Fortress (SF) (Mané and Donchin 1989), an arcade-style game which was developed by psychologists in the 80s to study human skill acquisition, and is still used quite frequently (Towne, Boot, and Ericsson 2016; Destefano and Gray 2016). The objective of the game is to fly a ship and destroy a fortress by firing missiles at it. The ship has to respect a minimum time difference between successive shots, while building up the fortress' vulnerability, and once the fortress becomes vulnerable, destroy it with a rapid double shot. As a RL testbed, Space Fortress possesses both the characteristics discussed above: context-dependent strategy change (change in required firing rate after the fortress becomes vulnerable) and time sensitivity (firing rate requirements independent of the agent's decision speed i.e., the frame rate). It also has a sparse reward structure, and, as we show, is not solved by any state-of-the-art RL algorithms such as Rainbow (Hessel et al. 2018), Proximal Policy Optimization (PPO) (Schulman et al. 2017) and Advantage Actor-Critic (A2C) (Mnih et al. 2016).

While being an interesting and relevant challenge for reinforcement learning, the rich background on human skill acquisition research based on Space Fortress also makes it an attractive tool to study human-AI collaboration in a dynamic environment, compare skill acquisition techniques of humans vs artificial agents, and work on few-shot learning by leveraging lessons from cognitive architectures like ACT-R (Anderson 2009) which have previously learned the game with extremely high sample efficiency, albeit using hand-crafted features and extensive domain knowledge.

We make the following contributions. First, we present a new RL testbed that requires the agent to switch strategies abruptly based on context, and develop a conceptualization of time independent of its speed of decision making, and demonstrate empirically that performance on par with humans is beyond the capability of current state-of-the-art RL algorithms, even after relaxing the reward sparsity through shaping (Ng, Harada, and Russell 1999). We identify the aspects of the game which can be varied to control both temporal and context sensitivity, allowing research on either *in isolation*. Finally, we demonstrate that *after introducing modifications* to ease identification of critical contexts, the PPO algorithm learns to play the game well enough to outperform humans, verifying that context insensitivity is the primary driver behind the poor performance of RL algorithms. We also present robust human benchmark results for Space Fortress, allowing future researchers to place new experimental results in context. We open-source¹ the OpenAI Gym environment for Space Fortress as well as all the code used to run our experiments, to promote research in temporal and context-sensitive reinforcement learning algorithms.

2 Related Work

The Arcade Learning Environment (ALE) (Bellemare et al. 2013) poses the challenge of building AI agents with competency across dozens of Atari 2600 games, like Space Invaders, Asteroids, Bowling and Enduro. Following the development of Deep Q Networks (Mnih et al. 2015), a lot of research in the RL community has focused on improving performance in one or more of the games with improvements like massive parallelization, sample efficiency (Wang et al. 2015; Schaul et al. 2015), better exploration (Fortunato et al. 2017; Plappert et al. 2017), solving tasks with sparse rewards (Pathak et al. 2017; Andrychowicz et al. 2017) and long-term strategies (Bacon, Harb, and Precup 2017; Kulkarni et al. 2016). In continuous control tasks on the MuJoCo testbed (Todorov, Erez, and Tassa 2012), on-policy actor critic methods (Schulman et al. 2017; Mnih et al. 2016) have shown promise. Bellemare, Dabney, and Munos (2017) estimated a probability distribution over the Q-value of a state (instead of just the mean of the Q-value), with greatly improved results. Rainbow (Hessel et al. 2018) combined a lot of orthogonal improvements in DQNs to achieve state of the art results. However, we show below that these algorithms fail to learn anything on Space Fortress.

Games like Ms. Pacman and Seaquest in the ALE have previously required some context or temporal sensitivity, but

these characteristics can't be directly controlled or varied. As a RL testbed, Space Fortress relies heavily on both context and temporal sensitivity, as we show in Section 4, and both characteristics can be controlled directly to enable their study in isolation.

There has also been a fair amount of prior work on reinforcement learning with sparse rewards. Pathak et al. (2017) use curiosity as an intrinsic reward signal to efficiently direct exploration. State visitation counts have also been investigated for exploration (Bellemare et al. 2016), and Osband et al. (2016) train multiple value functions and make use of bootstrapping and Thompson sampling for exploration. These works focus on learning with sparse rewards through better exploration of the state space, which does not help with Space Fortress where exploration is required in time and in latent contexts.

Zambrano, Roelfsema, and Bohte (2015) trained agents to deal with actions that take a finite amount of time through neural reinforcement learning in grid worlds, which still did not require a conceptualization of time independent of the internal speed of decision making, hence differing from the proposed work. Finally, van der Linden et al.; van Oijen et al. (2017; 2017) previously used A3C on a simple control task abstracted from Space Fortress, with no fortress destruction required. Crucially, this task removed the interesting characteristics of Space Fortress, namely contextual and temporal sensitivity, as well as reward sparsity. We release an implementation of the game as an OpenAI Gym environment to promote research, conduct an ablation study to ascertain the roles of context and temporal sensitivity and reward sparsity in poor performance, and then present results showing existing RL algorithms outperforming humans after we control for the above factors.

3 The Space Fortress RL Environment

We now describe the Space Fortress game, discuss its utility as a testbed for reinforcement learning, and present results from humans learning to play the game, intended as a baseline. The game environment can be seen in Fig. 1.

3.1 Game Description

The player/AI agent controls a ship, which has to fly around in a frictionless arena, firing missiles to destroy a fortress located centrally within the arena. Hitting the walls on either sides or being hit by shells fired by the fortress results in immediate ship death, which incurs a penalty on the agent. Destroying the fortress, however, requires a *context-aware strategy*. Each missile that successfully hits the fortress increases its vulnerability v by one. When $v < 10$, the fortress is 'not vulnerable', and the ship must fire its missiles spaced *more than 250ms* apart. Firing faster than this while $v < 10$ leads to the fortress vulnerability getting reset back to zero. This is obviously undesirable and the agent must learn to shoot slowly. However, once $v = 10$, the fortress becomes vulnerable, and a rapid double fire (2 shots spaced *less than 250ms* apart) is required to destroy the fortress. We refer to this 250ms time specification as the "critical time interval". It is important to note that once $v = 10$, shooting

¹<https://github.com/agakshat/spacefortress>

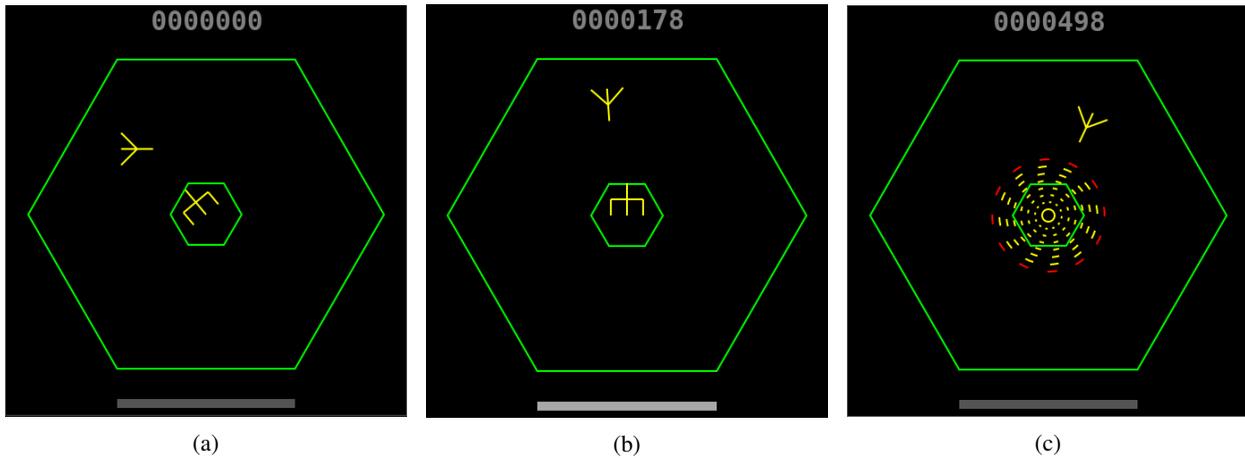


Figure 1: Game Screens in **Space Fortress**. The ship has to fly between the two hexagons, while the fortress can only change its orientation at a fixed position. The game score is displayed at the top, and the fortress’ vulnerability is displayed as a bar which fills up on each shot. (a) The bar is empty, indicating that the fortress’ vulnerability is 0, (b) The bar is now full, indicating that vulnerability is equal to 10 and a rapid double shot will now destroy the fortress and (c) The fortress has been destroyed. This is followed by a reset of the fortress and continuation of the game until the end of the episode (3 minutes game time)

further at the fortress at a rate less than 4Hz will lead to no change in vulnerability. Hence the firing *strategy completely reverses* at the point when vulnerability reaches 10, and the agent must learn to identify this critical point to perform well. Since the game is simply reset (without ending the episode) when the fortress is destroyed, it is crucial that the agent also recognize this second critical point of fortress destruction, and switch back its firing rate to continue playing well.

A single game lasts for 3 minutes. The game does not end in the event of either a fortress or ship destruction, and points are scored by destroying the fortress as many times as possible in those 3 minutes while avoiding getting shot down by the fortress or colliding with the arena. When a fortress is destroyed, its vulnerability resets to zero, and the game continues. When the ship is destroyed, it respawns at a random position and orientation, but the fortress’ vulnerability is preserved.

3.2 Game Versions

Space Fortress requires the agent to master advanced controls in a frictionless environment, orienting and firing missiles at the fortress while avoiding shells and not colliding with the walls. Since current RL algorithms proved unable to solve the game in its entirety (see experiments in Section 4.2), we introduced another version of the game to reduce navigation complexity by having the ship automatically pointed at the fortress. Throughout the rest of the paper, the simpler version is called ‘Autoturn’, while the original game is referred to as ‘Youturn’.

3.3 Human Evaluations

The human player results were collected in the context of a study on human skill acquisition (Anderson et al. 2018). 117 people were asked to play 20 games of Space Fortress, with

52 playing Autoturn and 65 playing Youturn. They were all *given instructions about the rules of the game beforehand*, and told about the change in firing rate required when the fortress vulnerability reaches 10. Considering that humans would require some turns to learn to play the game, we report the following results in Table 1: (1) Best performance of any subject in any game, (2) Average performance of all subjects in the last 5 games, considering the first 15 as a learning phase, (3) Average performance of all subjects in the last 10 games, considering the first 10 as a learning phase, (4) Average performance of all subjects in the last 15 games, considering the first 5 as a learning phase and (5) Average performance of all subjects in all 20 games. The scores shown to the humans (and reported in Table 1) were as follows: +100 for fortress destruction, -100 for ship death and -2 for each missile shot to penalize excessive firing.

3.4 RL Setup

We now describe the exact game setup used for reinforcement learning on Space Fortress.

- **Observations:** The observations are in the form of pixel-level grayscale 84x84 size renderings of the game screen (similar to Fig. 1. Important information such as the time lapsed since the last shot is not a part of this observation, making the task partially observed. We provide the agent with a stack of the last 4 observations as input at each time step, allowing it to infer direction of movement of the ship and fortress using the difference between successive frames.
- **Actions:** The agent chooses from 5 actions: (i) No Operation, (ii) Fire (a missile), (iii) Thrust Forward (in the direction of current orientation), (iv) Thrust Right (rotate right without changing position) and (v) Thrust Left (rotate left without changing position). The game operates at a default frame rate of 30 FPS and there is no action repeat,

Game	N	Metric	Best	Last 5	Last 10	Last 15	All
Autoturn	52	Score	3000	1989	1978	1940	1810
		FortressDeath	40	30.311	30.044	29.591	28.181
Youturn	65	Score	2314	216	153	43	-169
		FortressDeath	32	14.36	13.704	12.882	11.4

Table 1: Aggregated results for 102 humans playing Space Fortress. After being provided with instructions about the rules beforehand, each player played the game for 1 hour, or 20 games. Allowing for a few practice games, we report the average scores on the Last ‘ K ’ games ($K \in \{5, 10, 15\}$), as well as the best individual score.

S.No.	Algorithm	Game	Avg. Score	Best Score	Fortress Death
1	A2C	Autoturn	-2685	-2242	0
2	A2C	Youturn	-5859	-5604	0
3	PPO	Autoturn	-2502	-2178	0
4	PPO	Youturn	-5269	-4698	0
3	Rainbow	Autoturn	-8327	-8264	0
4	Rainbow	Youturn	-9378	-9245	0

Table 2: Average game scores for RL agents, trained **with default (sparse) rewards**, for 45M steps

which means an action is chosen every 33ms. Note that the Autoturn version only has 3 actions (since no turning is required).

- **Rewards:** In line with Mnih et al. (2015), we found that learning was more stable when using clipped rewards. The fortress and ship destruction rewards were clipped to +1 and -1, respectively, and the missile penalty reduced to -0.05. Note that the results used for evaluation and reporting were not clipped, in order to follow the same scheme as described in Section 3.3.

4 Experiments and Results

In this section, we experimentally show that (a) no state-of-the-art reinforcement learning algorithm (Rainbow (Hessel et al. 2018), A2C (Mnih et al. 2016) and PPO (Schulman et al. 2017)) can learn to play Space Fortress, (b) removing reward sparsity does not improve the performance and (c) making context identification easier through specific alterations in the reward structure allow PPO to achieve superhuman performance. We also discuss temporal sensitivity by examining effectiveness of transfer of learning across different settings of the game’s critical time interval.

4.1 Network Architecture

For Rainbow, the Q-network architecture was identical to that in Hessel et al. (2018). For PPO and A2C, we present two policy network architectures:

- **SF-GRU:** The agent’s policy network takes the 1x84x84 environment observations as input, and outputs (a) a probability distribution over the actions, and (b) a value function estimate of the expected return. The input goes through two convolutional layers with 16, 32 filters of size 8,4 and stride 4,2 respectively, and ReLU activation. The output is flattened and passed through a linear layer with a ReLU non-linearity to get an output vector of size

256. This is then passed through a unidirectional Gated Recurrent Unit (GRU) cell (Cho et al. 2014) with a tanh non-linearity giving an output of size 256. Finally, this vector is passed as input to two linear layers that output the probability distribution over actions (using a softmax activation) and the value estimate of the expected return.

- **SF-FF:** Same as above, but with a fully connected layer of size 256 with ReLU non-linearity instead of the recurrent GRU cell.

For all experiments, we ran 16 processes collecting game experience in parallel, with discount factor $\gamma = 0.99$ and Generalized Advantage Estimation (GAE) (Schulman et al. 2015) parameter $\lambda = 0.95$. PPO used value loss coefficient $c_1 = 0.5$, entropy regularization coefficient $c_2 = 0.05$ and learning rate 1×10^{-3} , while A2C used $c_1 = 0.5$, $c_2 = 0.01$ and learning rate 5×10^{-4} . Both A2C and PPO used $n = 1024$ -step returns. These hyperparameters were found after extensive tuning. We also clipped the gradients of all the network parameters to 0.5, to prevent catastrophic updates from outlying samples of the expected gradient value. Since the PPO algorithm is more stable, we updated the policy 4 times every epoch - while A2C made only 1 update every epoch. We would also like to note that we did extensive experimentation with different policy network architectures and found that increasing the size or number of layers did not really improve performance.

4.2 With Default (Sparse) Rewards

With the default sparse reward structure which rewards fortress destruction and penalizes ship destruction and missile firing, no algorithm is able to learn to destroy the fortress. A visual inspection of the game play revealed that the PPO and A2C agents (with both architectures) just learned to stop firing, since that leads to an immediate penalty. The Rainbow agent did not learn anything. Table 2 presents the aggregated results for PPO, A2C and Rainbow

S.No.	Algorithm	Game	Avg. Score	Best Score	Fortress Death
1	A2C	Autoturn	-4116	-2100	0
2	A2C	Youturn	-4781	-3890	1.3
3	PPO	Autoturn	-1294	-1108	1
4	PPO	Youturn	-1435	-1206	0.94
3	Rainbow	Autoturn	-6161	-5960	0
4	Rainbow	Youturn	-4894	-4577	0

Table 3: Average game scores for RL agents, trained **with dense rewards**, for 45M steps.

S.No.	Algorithm	Architecture	Game	Avg. Score	Best Score	Fortress Death
1	A2C	SF-GRU	Autoturn	-1641	-718	3
2	A2C	SF-GRU	Youturn	-2444	-1700	11
3	PPO	SF-FF	Autoturn	2337	2818	41
4	PPO	SF-FF	Youturn	2235	2880	40
5	PPO	SF-GRU	Autoturn	2510	2870	43
6	PPO	SF-GRU	Youturn	2356	2932	41
7	Rainbow	-	Autoturn	-2973	-2330	1.2
8	Rainbow	-	Youturn	-4112	-3934	0.0

Table 4: Average game scores for RL agents, trained **after making context identification easier**, for 45M steps

on both versions of the game. The ‘Fortress Death’ column in Table 2 indicates the number of times the agent was able to destroy the fortress per game, on average.

4.3 With Dense Rewards

Considering the results in Section 4.2 and to understand how reward sparsity is impacting performance, we introduce an additional reward of +1 each time the fortress is hit by a missile, and a penalty of -1 if the fortress’ vulnerability gets reset due to a faster firing speed than the context demands. This makes the reward density comparable to the Atari games on which PPO, A2C and Rainbow have all been shown to perform well (like Pong, Enduro etc.). Their performance on Space Fortress with dense rewards can be seen in Table 3, where the scores for PPO and Rainbow have improved. From watching a video of the trained agent playing the game, we observed that the improvement stemmed from having learned to avoid ship death and to fire at the fortress, albeit without knowledge of the critical time interval and context-dependent strategy shifts, resulting in an inability to destroy the fortress with any consistency.

Hence, Space Fortress clearly presents a challenge to the state of the art in reinforcement learning, and is a useful and relevant benchmark for further research.

We now move on to studying the impact of context insensitivity of RL algorithms on the task performance (Section 4.4), and show that by making the identification of critical contexts easier with 2 simple modifications to the reward, PPO learns to play the game very well - outperforming humans comfortably. This clearly indicates that it is context insensitivity and the inability to identify critical points which is hampering performance, further making the case that the Space Fortress game is a useful benchmark for studying con-

text sensitivity.

4.4 After Making Context Identification Easier

As discussed in Section 3.1, there are 2 critical points which the agent has to learn to identify and switch strategies at. The first is when the fortress becomes vulnerable, i.e. $v = 10$ and the agent has to switch from firing more than 250ms apart to a rapid double shot fired less than 250ms apart. The second is when the fortress is destroyed, and the agent has to switch back to its slow firing speed. To confirm our intuition that it is indeed the algorithms’ inability to identify these critical points and accordingly adapt its firing strategy which lead to poor performance, we introduce 2 changes to the reward structure (with respect to the dense reward from Section 4.3) which make it trivial for the agent to identify the critical points where context changes:

- Instead of rewarding fortress hits (as in Section 4.3), we switch to rewarding fortress vulnerability change, by giving a reward of +1 for unit increase in vulnerability, and giving a penalty of -1 for decrease in vulnerability. This has the effect of rewarding fortress hits only until the fortress’ vulnerability is building up to 10, at which point further hits are not rewarded. This clearly helps it identify the critical context at which the fortress becomes vulnerable.
- We give the agent a bonus reward of +2 for fortress destruction, to help it identify when the fortress is destroyed. This bonus is applied after reward clipping has been done.

Table 4 presents the results for agents trained after these 2 changes have been introduced to the reward structure to ease context identification. PPO with the recurrent architecture SF-GRU achieves the best performance in both score and

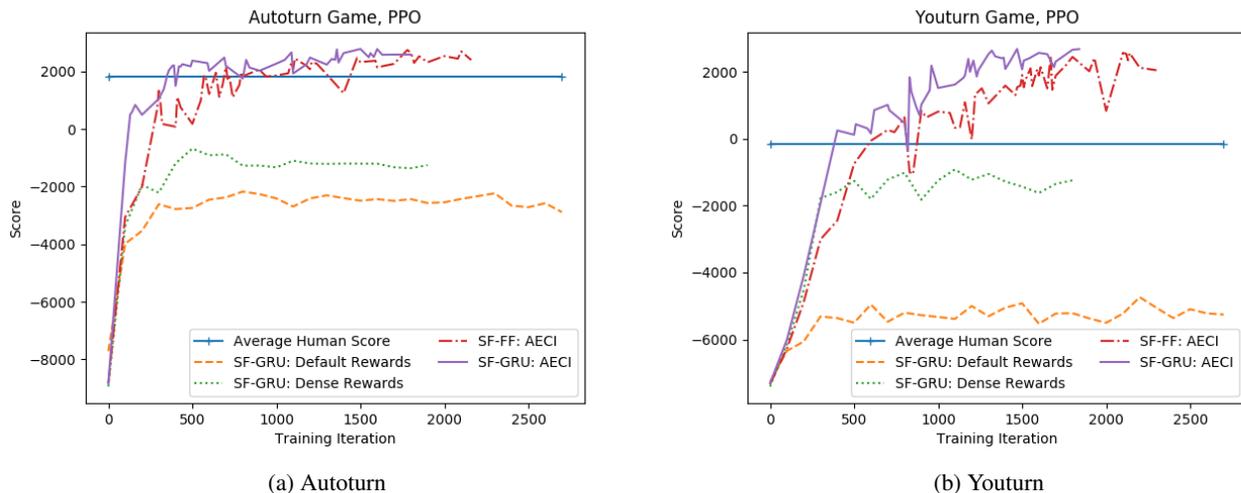


Figure 2: Learning Curves for PPO on Space Fortress, for different reward structures and architectures. ‘Average Human Score’ refers to the average score over all 20 games, provided as a point of comparison. ‘Default Rewards’ discussed in Section 4.2, ‘Dense Rewards’ in Section 4.3 and ‘AECI’ (After Easing Context Identification) in Section 4.4. Both SF-FF and SF-GRU architectures are able to achieve superhuman performance *after making context identification easier*. The agent’s performance is very poor with both default (sparse) rewards and dense rewards.

number of fortress deaths, learning faster than SF-FF, and achieving a higher final score. The performance of both A2C and Rainbow also improves, although they are still unable to outperform humans. Fig. 2 tracks the learning curves for PPO learning with all 3 reward settings (default, dense and after making context identification easier) in both game versions - Autoturn and Youturn.

We also experimented with introducing various other rewards which were not specifically directed at making context identification easier, and found none of them to help learning. Hence we inferred that the reward shaping we introduced in this section helped specifically because it made context identification easier.

4.5 Temporal Sensitivity

Having established that context insensitivity is the primary driver of poor performance of state of the art RL algorithms on Space Fortress, we now present some analysis of the temporal sensitivity of the PPO algorithm, while keeping the reward structure as described in Section 4.4. As described in Section 3.1, Space Fortress has a dominant temporal aspect - missiles must hit the fortress at least 250ms apart when it is not vulnerable, and then the strategy must reverse to hit the fortress *twice within 250ms* when it is vulnerable, in order to destroy it. In order to understand whether the RL algorithms had developed any understanding of time as an independent dimension, we modified the critical time interval from 250ms to other values, and checked for positive transfer of learning from the policy trained with 250ms as the critical time interval. We achieve transfer of learning by simply initializing the weights of the transferee with the learned weights of the transferer.

Figure 3 compares the learning curves for an agent learn-

ing with PPO (using the SF-FF architecture) on Youturn, when the critical time interval is changed from 250ms to {125, 400, 600}ms. The blue line is for an agent learning from scratch, while the orange line is for an agent transferring learning from the PPO SF-FF agent trained on the 250ms interval. From Figure 3b and 3c, it can be seen that while the transfer of learning helps by initializing the weights in a favorable corner of the parameter space, the learning saturates very quickly and ends up with a final score much lower than achieved when the critical interval was 250ms. Modifying the critical time interval in Space Fortress is a useful technique to study temporal sensitivity of reinforcement learning algorithms.

5 Conclusion

This paper introduced Space Fortress as a new testbed for deep reinforcement learning research, complementing the existing suite of benchmarks while focusing specifically on time-sensitive game play, abrupt context-dependent shift of strategy and sparse rewards. We showed that state of the art RL algorithms (PPO, A2C and Rainbow) were unable to learn to play the game with neither the default sparse rewards nor the dense reward structure we defined. After making context identification easier through two minor tweaks in the reward structure, however, PPO was able to learn to play the game, outperforming humans comfortably. This ablation study allowed us to conclude that context insensitivity was the primary reason behind the poor performance of RL algorithms on Space Fortress, along with the inability to learn with sparse rewards. We then looked at whether PPO develops a concept of time as an independent variable - by checking for positive transfer of learning while changing the critical time interval of 250ms in Space Fortress. We found

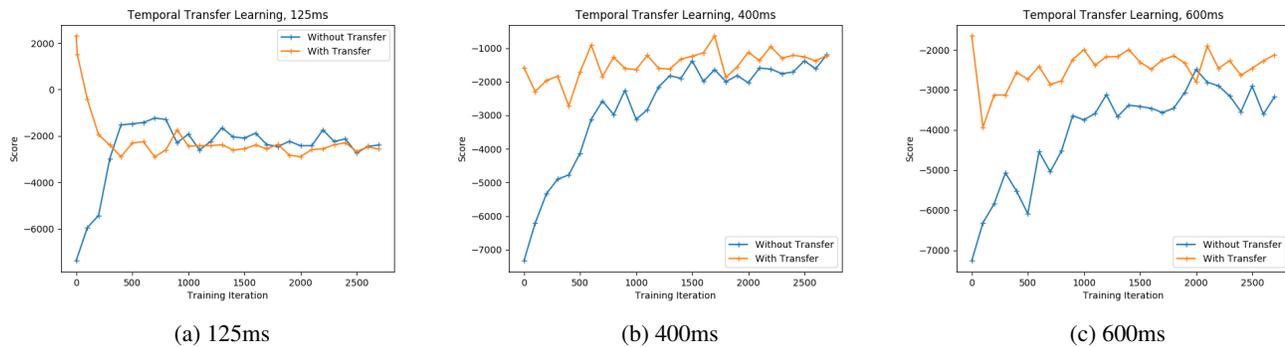


Figure 3: Checking for positive transfer of learning while changing the critical time interval in Space Fortress. In each figure, we changed the time interval to a different value and verified whether transferring learned weights from the agent trained on 250ms as the critical interval led to any speedup in the learning process, and an improvement in the final performance.

that while there was some positive transfer of learning, the agents saturated very quickly and did not achieve a good final score. By studying generalization and transfer across different settings of the critical time interval, Space Fortress can hence also be used as a benchmark to study temporal sensitivity of reinforcement learning algorithms.

Learning to play Space Fortress without making any modifications to the reward structure will require reinforcement learning algorithms to be able to identify various latent contexts and adapt their strategies suitably. It will also require being able to learn with very sparse rewards. This is beyond the capability of current state of the art reinforcement learning algorithms, making Space Fortress a useful benchmark for research.

Acknowledgments This research was sponsored by AFOSR Grants FA9550-15-1-0442 and FA9550-18-1-0251. The collection of human data and development of the OpenAI Gym interface for Space Fortress was supported by ONR grant N00014-15-1-2151. We would like to thank Shawn Betts and John Anderson for insightful discussions on the game of Space Fortress, and for the OpenAI Gym interface for Space Fortress which was used to run the experiments in this paper.

References

- Anderson, J.; Betts, S.; Bothell, D.; Hope, R. M.; and Lebiere, C. 2018. Three aspects of skill acquisition.
- Anderson, J. R. 2009. *How can the human mind occur in the physical universe?* Oxford University Press.
- Andrychowicz, M.; Wolski, F.; Ray, A.; Schneider, J.; Fong, R.; Welinder, P.; McGrew, B.; Tobin, J.; Abbeel, O. P.; and Zaremba, W. 2017. Hindsight experience replay. In *Advances in Neural Information Processing Systems*, 5048–5058.
- Bacon, P.-L.; Harb, J.; and Precup, D. 2017. The option-critic architecture. In *AAAI*, 1726–1734.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research* 47:253–279.
- Bellemare, M.; Srinivasan, S.; Ostrovski, G.; Schaul, T.; Saxton, D.; and Munos, R. 2016. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, 1471–1479.
- Bellemare, M. G.; Dabney, W.; and Munos, R. 2017. A distributional perspective on reinforcement learning. *arXiv preprint arXiv:1707.06887*.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. Openai gym. *arXiv preprint arXiv:1606.01540*.
- Cho, K.; Van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Destefano, M., and Gray, W. D. 2016. Where should researchers look for strategy discoveries during the acquisition of complex task performance? the case of space fortress. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, 668–673.
- Fortunato, M.; Azar, M. G.; Piot, B.; Menick, J.; Osband, I.; Graves, A.; Mnih, V.; Munos, R.; Hassabis, D.; Pietquin, O.; et al. 2017. Noisy networks for exploration. *arXiv preprint arXiv:1706.10295*.
- Hessel, M.; Modayil, J.; van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2018. Rainbow: Combining improvements in deep reinforcement learning. *AAAI Conference on Artificial Intelligence*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.
- Kulkarni, T. D.; Narasimhan, K.; Saeedi, A.; and Tenenbaum, J. 2016. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In *Advances in neural information processing systems*, 3675–3683.

- Mané, A., and Donchin, E. 1989. The space fortress game. *Acta psychologica* 71(1-3):17–22.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540):529.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, 1928–1937.
- Ng, A. Y.; Harada, D.; and Russell, S. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, 278–287.
- Osband, I.; Blundell, C.; Pritzel, A.; and Van Roy, B. 2016. Deep exploration via bootstrapped dqn. In *Advances in neural information processing systems*, 4026–4034.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning (ICML)*, volume 2017.
- Plappert, M.; Houthoofd, R.; Dhariwal, P.; Sidor, S.; Chen, R. Y.; Chen, X.; Asfour, T.; Abbeel, P.; and Andrychowicz, M. 2017. Parameter space noise for exploration. *arXiv preprint arXiv:1706.01905*.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115(3):211–252.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- Schulman, J.; Moritz, P.; Levine, S.; Jordan, M.; and Abbeel, P. 2015. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of go with deep neural networks and tree search. *nature* 529(7587):484–489.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, 3104–3112.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, 5026–5033. IEEE.
- Towne, T. J.; Boot, W. R.; and Ericsson, K. A. 2016. Understanding the structure of skill through a detailed analysis of individuals’ performance on the space fortress game. *Acta psychologica* 169:27–37.
- van der Linden, P. A.; Intelligentie, B. O. K.; Poppinga, G.; Roessingh, J.; and van Splunter, S. 2017. On higher-order control tasks: The application of a3c on space fortress.
- van Oijen, J.; Poppinga, G.; Brouwer, O.; Aliko, A.; and Roessingh, J. J. 2017. Towards modeling the learning process of aviators using deep reinforcement learning. In *Systems, Man, and Cybernetics (SMC), 2017 IEEE International Conference on*, 3439–3444. IEEE.
- Wang, Z.; Schaul, T.; Hessel, M.; Van Hasselt, H.; Lanctot, M.; and De Freitas, N. 2015. Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv:1511.06581*.
- Zambrano, D.; Roelfsema, P. R.; and Bohte, S. M. 2015. Continuous-time on-policy neural reinforcement learning of working memory tasks. In *Neural Networks (IJCNN), 2015 International Joint Conference on*, 1–8. IEEE.