

# Counting to Explore and Generalize in Text-based Games

Xingdi Yuan\*

ERIC.YUAN@MICROSOFT.COM

Marc-Alexandre Côté\*

MACOTE@MICROSOFT.COM

Alessandro Sordoni

ALSORDON@MICROSOFT.COM

Romain Laroche

ROMAIN.LAROCHE@MICROSOFT.COM

Rémi Tachet des Combes

REMI.TACHET@MICROSOFT.COM

Matthew Hausknecht

MATTHEW.HAUSKNECHT@MICROSOFT.COM

Adam Trischler

ADAM.TRISCHLER@MICROSOFT.COM

*Microsoft Research*

*2000 McGill College Avenue, Suite #550,  
Montréal, H3A 3H3, Canada*

## Abstract

We propose a recurrent RL agent with an episodic exploration mechanism that helps discovering good policies in text-based game environments. We show promising results on a set of generated text-based games of varying difficulty where the goal is to collect a coin located at the end of a chain of rooms. In contrast to previous text-based RL approaches, we observe that our agent learns policies that generalize to unseen games of greater difficulty.

**Keywords:** Machine Learning, Reinforcement Learning, Exploration, Text-based Games

## 1. Introduction

Text-based games like *Zork* (Infocom, 1980) are complex, interactive simulations. They use natural language to describe the state of the world, to accept actions from the player, and to report subsequent changes in the environment. The player works toward goals which are seldom specified explicitly and must be discovered through exploration. The observation and action spaces in text games are both combinatorial and compositional, and players must contend with partial observability, since descriptive text does not communicate complete, unambiguous information about the underlying game state. Figure 1 depicts *Zork*'s opening scene along with two player commands and the corresponding system responses.

In this paper, we study several methods of exploration in text-based games. Our basic task is a deterministic text-based version of the *chain experiment* (Osband et al., 2016; Plappert et al., 2017) with distractor nodes that are off-chain: the agent must navigate a path composed of discrete locations (rooms) to the goal, ideally without revisiting dead ends. We propose a DQN-based recurrent model for solving text-based games, where the recurrence gives the model the capacity to condition its policy on historical state information. To encourage exploration, we extend count-based exploration approaches (Ostrovski et al., 2017; Tang et al., 2017), which assign an intrinsic reward derived from the count of state visitations during learning, across episodes. Specifically, we propose an episodic count-based exploration scheme, where state counts are reset at the beginning of each episode. This reward plays the role of an *episodic memory* (Gershman and Daw, 2017) that pushes the

```

You are standing in an open field west of a
white house, with a boarded
front door.
There is a small mailbox here.

> open mailbox
Opening the small mailbox reveals a leaflet.

> take leaflet
Taken.

> _

```

Figure 1: Intro to Zork

agent to visit states not previously encountered within an episode. Although the recurrent policy architecture has the capacity to solve the task by remembering and avoiding previously visited locations, we hypothesize that exploration rewards will help the agent learn to utilize its memory.

We generate a set of games of varying difficulty (measured with respect to the path length and the number of off-chain rooms) with a text-based game generator (Côté et al., 2018). We observe that, in contrast to a baseline model and standard count-based exploration methods, the recurrent model with episodic bonus learns policies that not only complete multiple training games at same time successfully but also generalize to *unseen* games of *greater difficulty*.

## 2. Text-based Games as POMDPs

Text-based games are sequential decision-making problems that can be described naturally by the Reinforcement Learning (RL) setting. Fundamentally, text-based games are partially observable Markov decision processes (POMDP) (Kaelbling et al., 1998) where the environment state is never observed directly. To act optimally, an agent must keep track of all observations. Formally, a text-based game is a discrete-time POMDP defined by  $(S, T, A, \Omega, O, R, \gamma)$ , where  $\gamma \in [0, 1]$  is the discount factor.

**Environment States ( $S$ ):** The environment state at turn  $t$  in the game is  $s_t \in S$ . It contains the complete internal information of the game, much of which is hidden from the agent. When an agent issues a command  $\mathbf{c}_t$  (defined next), the environment transitions to state  $s_{t+1}$  with probability  $T(s_{t+1}|s_t, \mathbf{c}_t)$ .

**Actions ( $A$ ):** At each turn  $t$ , the agent issues a text command  $\mathbf{c}_t$ . The interpreter can accept any sequence of characters but will only recognize a tiny subset thereof. Furthermore, only a fraction of recognized commands will actually change the state of the world. The resulting action space is enormous and intractable for existing RL algorithms. In this work, we make the following two simplifying assumptions. (1) **Word-level** Each command is a two-word sequence where the words are taken from a fixed vocabulary  $V$ . (2) **Command syntax** Each command is a (verb, object) pair (direction words are considered objects).




Figure 2: **LSTM-DRQN** processes textual observations word-by-word to generate a fixed-length vector representation. This representation is used by the recurrent policy to estimate Q-values for all verbs  $Q(s, v)$  and objects  $Q(s, o)$ .

**Observations ( $\Omega$ ):** The text information perceived by the agent at a given turn  $t$  in the game is the agent’s observation,  $o_t \in \Omega$ , which depends on the environment state and the previous command with probability  $O(o_t|s_t, c_{t-1})$ . Thus, the function  $O$  selects from the environment state what information to show to the agent given the last command.

**Reward Function ( $R$ ):** Based on its actions, the agent receives reward signals  $r_t = R(s_t, a_t)$ . The goal is to maximize the expected discounted sum of rewards  $E[\sum_t \gamma^t r_t]$ .

### 3. Method

#### 3.1 Model Architecture

In this work, we adopt the LSTM-DQN (Narasimhan et al., 2015) model as baseline. It has two modules: a representation generator  $\Phi_R$ , and an action scorer  $\Phi_A$ .  $\Phi_R$  takes observation strings  $o$  as input, after a stacked embedding layer and LSTM (Hochreiter and Schmidhuber, 1997) encoder, a mean-pooling layer produces a vector representation of the observation. This feeds into  $\Phi_A$ , in which two MLPs, sharing a lower layer, predict the Q-values over all verbs  $w_v$  and object words  $w_o$  independently. The average of the two resulting scores gives the Q-values for the composed actions. The LSTM-DQN does not condition on previous actions or observations, so it cannot memorize what happened previously. We concatenate the previous command  $c_{t-1}$  to the current observation  $o_t$  to lessen this limitation.

To further enhance the agent’s capacity to remember previous states, we replace the shared MLP in  $\Phi_A$  by an LSTM cell. This model is inspired by Hausknecht and Stone (2015); Lample and Chaplot (2016) and we call it LSTM-DRQN. The LSTM cell in  $\Phi_A$  takes the representation generated by  $\Phi_R$  together with history information  $h_{t-1}$  from the previous game step as input. It generates the state information at the current game step, which is then fed into the two MLPs as well as passed forward to next game step. Figure 2 shows the LSTM-DRQN architecture.

#### 3.2 Discovery Bonus

To promote exploration we use an intrinsic reward by counting game state visits (Kolter and Ng, 2009; Tang et al., 2017; Martin et al., 2017; Ostrovski et al., 2017). We investigate two approaches to counting rewards. The first is inspired by Kolter and Ng (2009), where

we define the **cumulative counting bonus** as  $r^+(o_t) = \beta \cdot n(o_t)^{-1/3}$ , where  $n(o_t)$  is the number of times the agent has observed  $o_t$  since the beginning of training (across episodes), and  $\beta$  is the bonus coefficient. During training, as the agent observes new game states more and more, the cumulative counting bonus gradually converges to 0.

The second approach is the **episodic discovery bonus**, which encourages the agent to discover unseen game states by assigning a positive reward whenever it sees a new state. It is defined as:  $r^{++}(o_t) = \begin{cases} \beta & \text{if } n(o_t) = 1 \\ 0.0 & \text{otherwise} \end{cases}$ , where  $n(\cdot)$  is reset to zero at the beginning of each episode. Taking inspiration from Gershman and Daw (2017), we hope this behavior pushes the agent to visit states not previously encountered in the current episode and teaches the agent how to use its memory for this purpose so it may generalize to unseen environments.

## 4. Related Work

**RL Applied to Text-based Games:** Narasimhan et al. (2015) test their LSTM-DQN in two text-based environments: Home World and Fantasy World. They report the quest completion ratio over multiple runs but not how many steps it takes to complete them. He et al. (2015) introduce the Deep Reinforcement Relevance Network (DRRN) for tackling choice-based (as opposed to parser-based) text games, evaluating the DRRN on one deterministic game and one larger-scale stochastic game. The DRRN model converges on both games; however, this model must know in advance the valid commands at each state. Fulda et al. (2017) propose a method to reduce the action space for parser-based games by training word embeddings to be aware of verb-noun affordances. One drawback of this approach is it requires pre-trained embeddings.

**Count-based Exploration:** The exploration-exploitation trade-off has been studied extensively in RL literature, algorithms have been proposed to act near-optimally on MDPs that requires only polynomial amount of resources Kearns and Singh (2002); Brafman and Tenenbholz (2003).

The Model Based Interval Estimation-Exploration Bonus (MBIE-EB) Strehl and Littman (2008) derives an intrinsic reward by counting state-action pairs with a table  $n(s, a)$ . Their exploration bonus has the form  $\beta/\sqrt{n(s, a)}$  to encourage exploring less-visited pairs. In this work, we use  $n(s)$  rather than  $n(s, a)$ , since the majority of actions leave the agent in the same state (*i.e.*, unrecognized commands). Using the latter would reward the agent for trying invalid commands, which is not sensible in our setting.

Tang et al. (2017) propose a hashing function for count-based exploration in order to discretize high-dimensional, continuous state spaces. Their exploration bonus  $r^+ = \beta/\sqrt{n(\phi(s))}$ , where  $\phi(\cdot)$  is a hashing function that can either be static or learned. This is similar to the *cumulative counting bonus* defined above.

**Deep Recurrent Q-Learning:** Hausknecht and Stone (2015) propose the Deep Recurrent Q-Networks (DRQN), adding a recurrent neural network (such as an LSTM (Hochreiter and Schmidhuber, 1997)) on top of the standard DQN model. DRQN estimates  $Q(o_t, h_{t-1}, a_t)$  instead of  $Q(o_t, a_t)$ , so it has the capacity to memorize the state history. Lample and Chaplot (2016) use a model built on the DRQN architecture to learn to play FPS games.

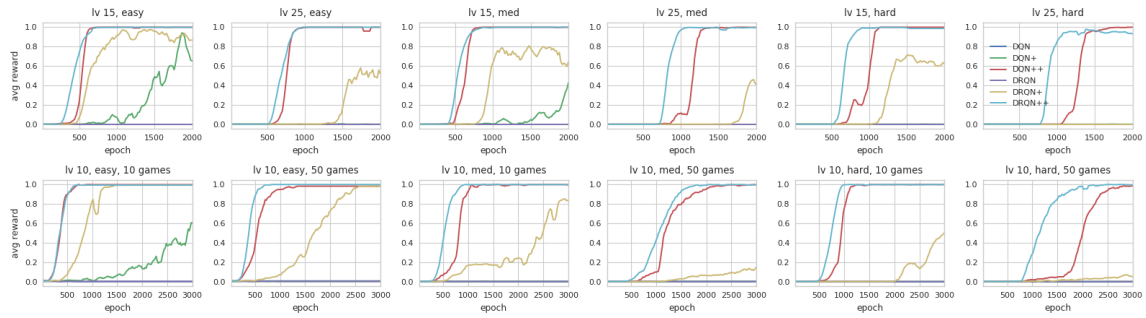


Figure 3: Model performance on single games (top row) and multiple games (bottom row).

A major difference between the work presented in this paper and the related work is that we test on unseen games and train on a set of similar (but not identical) games rather than training and testing on the same game.

## 5. Experiments

### 5.1 Coin Collector Game Setup

To evaluate the two models described above and the proposed discovery bonus, we designed a set of simple text-based games inspired by the chain experiment (Osband et al., 2016; Plappert et al., 2017). Each game contains a given number of rooms that are randomly connected to each other to form a chain (see figures in Appendix C). The goal is to find and collect a “coin” placed in one of the rooms. The player’s initial position is at one end of the chain and the coin is at the other. These games have deterministic state transitions, and in this work there is no stochasticity in observations (i.e., given the same state, the observation text is always the same).

Games stop after a set number of steps or after the player has collected the coin. The game interpreter understands only five commands (`go north`, `go east`, `go south`, `go west` and `take coin`), while the action space is twice as large:  $\{\text{go, take}\} \times \{\text{north, south, east, west, coin}\}$ . See Figure 11, Appendix C for an example of what the agent observes in-game.

Our games have 3 **modes**: *easy* (mode 0), there are no distractor rooms (dead ends) along the path; *medium* (mode 1), each room along the optimal trajectory has one distractor room randomly connected to it; *hard* (mode 2), each room on the path has two distractor rooms, i.e., within a room on the optimal trajectory, all 4 directions lead to a connected room. We use difficulty **levels** to indicate the optimal trajectory’s length of a game.

To solve easy games, the agent must learn to recall its previous directional action and to issue the command that does not reverse it (e.g., if the agent entered the current room by going `east`, do not now go `west`). Conversely, to solve medium and hard games, the agent must reverse its previous action when it enters distractor rooms to return to the chain, and also recall farther into the past to track which exits it has already passed through. Alternatively, since there are no cycles, it can learn a less memory intensive “wall-following” strategy by, e.g., taking exits in a clockwise order from where it enters a room.




Figure 4: Zero-shot evaluation: Average rewards of DQN++ (left) and DRQN++ (right) as a function of the number of games in the training set.

We refer to models with the cumulative counting bonus as MODEL+, and models with episodic discovery bonus as MODEL++, where MODEL  $\in$  {DQN, DRQN}<sup>1</sup> (implementation details in Appendix A). In this section we cover part of the experiment results, the full extent of our experiment results are provided in Appendix B.

## 5.2 Solving Training Games

We first investigate whether the variant models can learn to solve single games with different difficulty modes (easy, medium, hard) and levels  $\{L5, L10, L15, L20, L25, L30\}$ <sup>2</sup>. As shown in Figure 3 (top row), when the games are simple, vanilla DQN and DRQN already fail to learn. Adding the cumulative bonus helps somewhat and models perform similarly with and without recurrence. When the games become harder, the cumulative bonus helps less, while episodic bonus remains very helpful and recurrence in the model becomes very helpful.

Next, we are interested to see whether models can learn to solve a distribution of games. Note that each game has its own counting memory, *i.e.*, the states visited in one game do not affect the counters for other games. Here, we fix the game difficulty level to 10, and randomly generate training sets that contain  $\{2, 5, 10, 30, 50, 100\}$  games in each mode. As shown in Figure 3 (bottom row), when the game mode becomes harder, the episodic bonus has an advantage over the cumulative bonus, and recurrence becomes more crucial for memorizing the game distribution. It is also clear that the episodic bonus and recurrence help significantly when more training games are provided.

## 5.3 Zero-shot Evaluation

Finally, we want to see if a pre-trained model can generalize to unseen games. The generated training set contains  $\{1, 2, 5, 10, 30, 50, 100, 500\}$  L10 games for each mode. Then, for each corresponding mode the test set contains 10 *unseen*  $\{L5, L10, L15, L20, L30\}$  games. There is no overlap between training and test games in either text descriptions or optimal trajectories. At test time, the counting modules are disabled, the agent is not updated, and it generates verb and noun actions based on the *argmax* of their Q-values.

As shown in Figure 4, when the game mode is easy, both models with and without recurrence can generalize well on unseen games by training on a large training set. It is worth noting that by training on 500 L10 easy games, both models can almost perfectly

1. Since all models use the LSTM representation generator, we omit “LSTM” for abbreviation.  
 2. We use  $Lk$  to indicate level  $k$  game.



Figure 5: Average rewards and steps used corresponding to best validation performance in hard games.

solve level 30 unseen easy games. We also observe that models with recurrence are able to generalize better when trained on fewer games.

When testing on hard mode games, we observe that both models suffer from overfitting (after a certain number of episodes, average test reward starts to decrease while training reward increases). Therefore, we further generated a validation set that contains 10 L10 hard games, and report test results corresponding to best validation performance. In addition, we investigated what happens when concatenating the previous 4 steps’ history observation into the input. In Figure 5, we add  $H$  to model names to indicate this variant.

As shown in Figure 5, all models can memorize the 500 training games, while DQN++ and DRQN++H are able to generalize better on unseen games. In particular, the former performs near perfectly on test games. To investigate this, we looked into all the bi-grams of generated commands (*i.e.*, two commands from adjacent game steps) from DQN++ model. Surprisingly, except for moving back from dead end rooms, the agent always explores exits in anti-clockwise order. This means the agent has learned a general strategy that does not require history information beyond the previous command. This strategy generalizes perfectly to all possible hard games because there are no cycles in the maps.

## 6. Final Remarks

We propose an RL model with a recurrent component, together with an episodic count-based exploration scheme that promotes the agent’s discovery of the game environment. We show promising results on a set of generated text-based games of varying difficulty. In contrast to baselines, our approach learns policies that generalize to unseen games of greater difficulty.

In future work, we plan to experiment on games with more complex topology, such as cycles (where the “wall-following” strategy will not work). We would like to explore games that require multi-word commands (*e.g.*, unlock red door with red key), necessitating a model that generates sequences of words. Other interesting directions include agents that learn to map or to deal with stochastic transitions in text-based games.

## References

- Ronen I. Brafman and Moshe Tennenholtz. R-max - a general polynomial time algorithm for near-optimal reinforcement learning. *J. Mach. Learn. Res.*, 3:213–231, March 2003. ISSN 1532-4435. doi: 10.1162/153244303765208377. URL <https://doi.org/10.1162/153244303765208377>.
- Marc-Alexandre Côté, Ákos Kádár, Xingdi Yuan, Ben Kybartas, Tavian Barnes, Emery Fine, James Moore, Matthew Hausknecht, Layla El Asri, Mahmoud Adada, Wendy Tay, and Adam Trischler. Textworld: A learning environment for text-based games. Computer Games Workshop at IJCAI 2018, Stockholm, 2018.
- Nancy Fulda, Daniel Ricks, Ben Murdoch, and David Wingate. What can you do with a rock? affordance extraction via word embeddings. *arXiv preprint arXiv:1703.03429*, 2017.
- Samuel J Gershman and Nathaniel D Daw. Reinforcement learning and episodic memory in humans and animals: an integrative framework. *Annual review of psychology*, 68:101–128, 2017.
- Matthew J. Hausknecht and Peter Stone. Deep recurrent q-learning for partially observable mdps. *CoRR*, abs/1507.06527, 2015. URL <http://arxiv.org/abs/1507.06527>.
- Ji He, Jianshu Chen, Xiaodong He, Jianfeng Gao, Lihong Li, Li Deng, and Mari Ostendorf. Deep reinforcement learning with a natural language action space. *arXiv preprint arXiv:1511.04636*, 2015.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL <http://dx.doi.org/10.1162/neco.1997.9.8.1735>.
- Infocom. Zork I, 1980. URL <http://ifdb.tads.org/viewgame?id=0dbnuxunq7fw5ro>.
- Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- Michael Kearns and Satinder Singh. Near-optimal reinforcement learning in polynomial time. *Mach. Learn.*, 49(2-3):209–232, November 2002. ISSN 0885-6125. doi: 10.1023/A:1017984413808. URL <https://doi.org/10.1023/A:1017984413808>.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- J Zico Kolter and Andrew Y Ng. Near-bayesian exploration in polynomial time. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 513–520. ACM, 2009.
- Guillaume Lample and Devendra Singh Chaplot. Playing FPS games with deep reinforcement learning. *CoRR*, abs/1609.05521, 2016. URL <http://arxiv.org/abs/1609.05521>.



- Jarryd Martin, Suraj Narayanan Sasikumar, Tom Everitt, and Marcus Hutter. Count-based exploration in feature space for reinforcement learning. *arXiv preprint arXiv:1706.08090*, 2017.
- Karthik Narasimhan, Tejas Kulkarni, and Regina Barzilay. Language understanding for text-based games using deep reinforcement learning. *arXiv preprint arXiv:1506.08941*, 2015.
- Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped dqn. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 4026–4034. Curran Associates, Inc., 2016. URL <http://papers.nips.cc/paper/6501-deep-exploration-via-bootstrapped-dqn.pdf>.
- Georg Ostrovski, Marc G Bellemare, Aaron van den Oord, and Rémi Munos. Count-based exploration with neural density models. *arXiv preprint arXiv:1703.01310*, 2017.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- Matthias Plappert, Rein Houthoof, Prafulla Dhariwal, Szymon Sidor, Richard Y Chen, Xi Chen, Tamim Asfour, Pieter Abbeel, and Marcin Andrychowicz. Parameter space noise for exploration. *arXiv preprint arXiv:1706.01905*, 2017.
- Alexander L Strehl and Michael L Littman. An analysis of model-based interval estimation for markov decision processes. *Journal of Computer and System Sciences*, 74(8):1309–1331, 2008.
- Haoran Tang, Rein Houthoof, Davis Foote, Adam Stooke, Xi Chen, Yan Duan, John Schulman, Filip DeTurck, and Pieter Abbeel. # exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2750–2759, 2017.

## Appendix A. Implementation Details

Implementation details of our neural baseline agent are as follows<sup>3</sup>. In all experiments, the word embeddings are initialized with 20-dimensional random matrices; the number of hidden units of the encoder LSTM is 100. In the non-recurrent action scorer we use a 1-layer MLP which has 64 hidden units, with *ReLU* as non-linear activation function, in the recurrent action scorer, we use an LSTM cell which hidden size is 64.

In replay memory, we used a memory with capacity of 500000, a mini-batch gradient update is performed every 4 steps in the gameplay, the mini-batch size is 32. We apply prioritized sampling in all experiments, in which, we used  $\rho = 0.25$ . In LSTM-DQN and LSTM-DRQN model, we used discount factor  $\gamma = 0.9$ , in all models with discovery bonus, we used  $\gamma = 0.5$ .

When updating models with recurrent components, we follow the update strategy in Lample and Chaplot (2016), *i.e.*, we randomly sample sequences of length 8 from the replay memory, zero initialize hidden state and cell state, use the first 4 states to bootstrap a reliable hidden state and cell state, and then update on rest of the sequence.

We anneal the  $\epsilon$  for  $\epsilon$ -greedy from 1 to 0.2 over 1000 epochs, it remains at 0.2 afterwards. In both cumulative and episodic discovery bonus, we use coefficient  $\beta$  of 1.0.

When zero-shot evaluating hard games, we use  $max\_train\_step = 100$ , in all other experiments we use  $max\_train\_step = 50$ ; during test, we always use  $max\_test\_step = 200$ .

We use *adam* (Kingma and Ba, 2014) as the step rule for optimization. The learning rate is  $1e^{-3}$ . The model is implemented using *PyTorch* (Paszke et al., 2017).

All games are generated using TextWorld framework (Côté et al., 2018), we used the house grammar.

---

3. A PyTorch implementation of the proposed method can be found at <https://github.com/xingdi-eric-yuan/TextWorld-Coin-Collector>

Appendix B. More Results

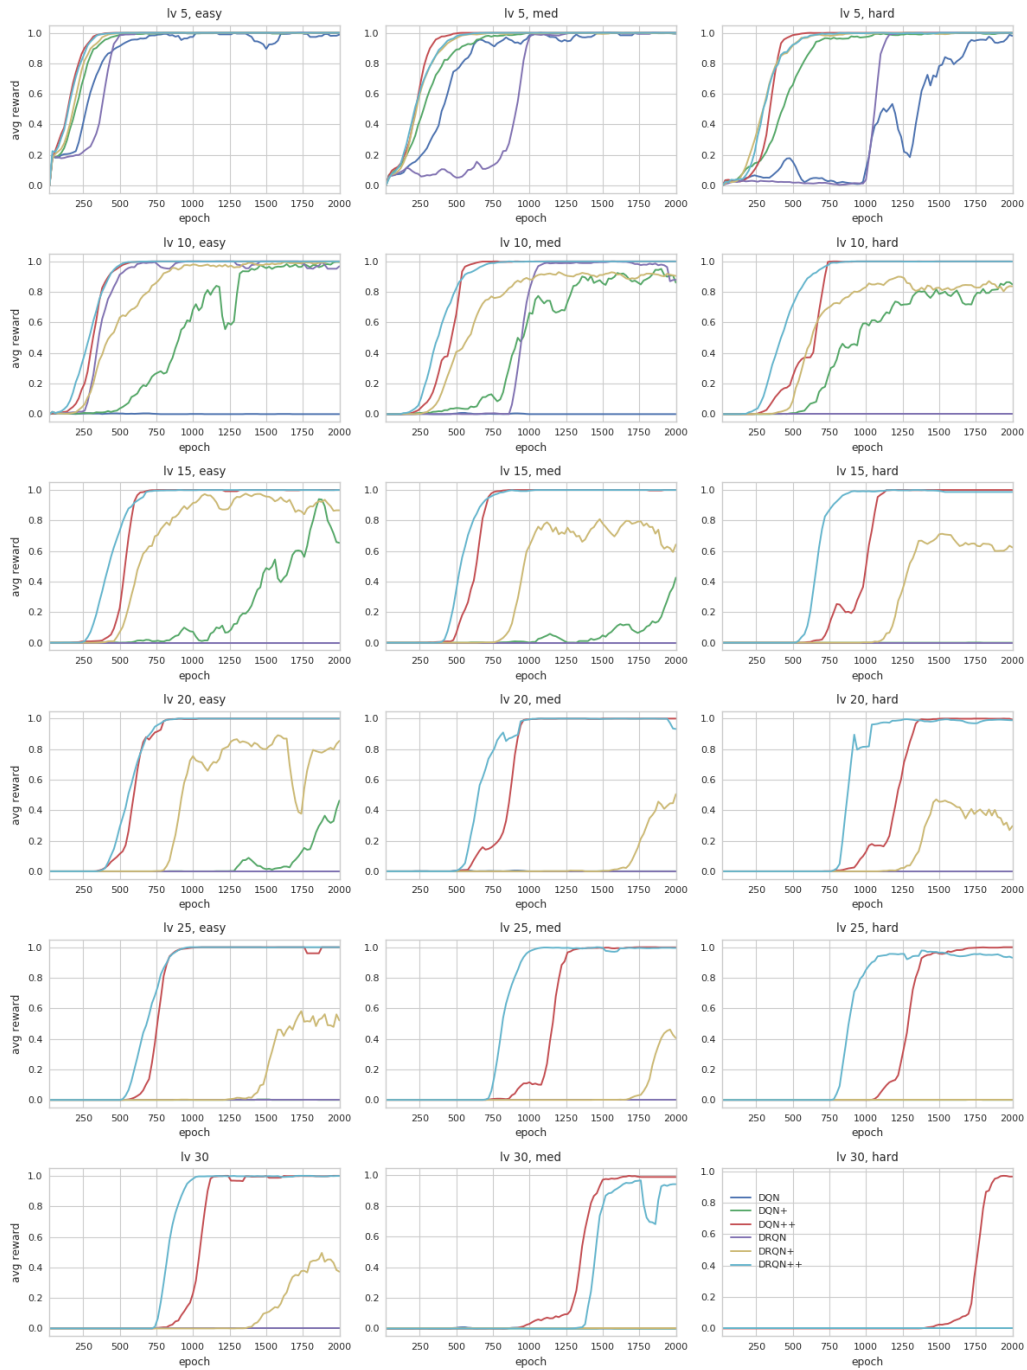


Figure 6: Model performance on single games.

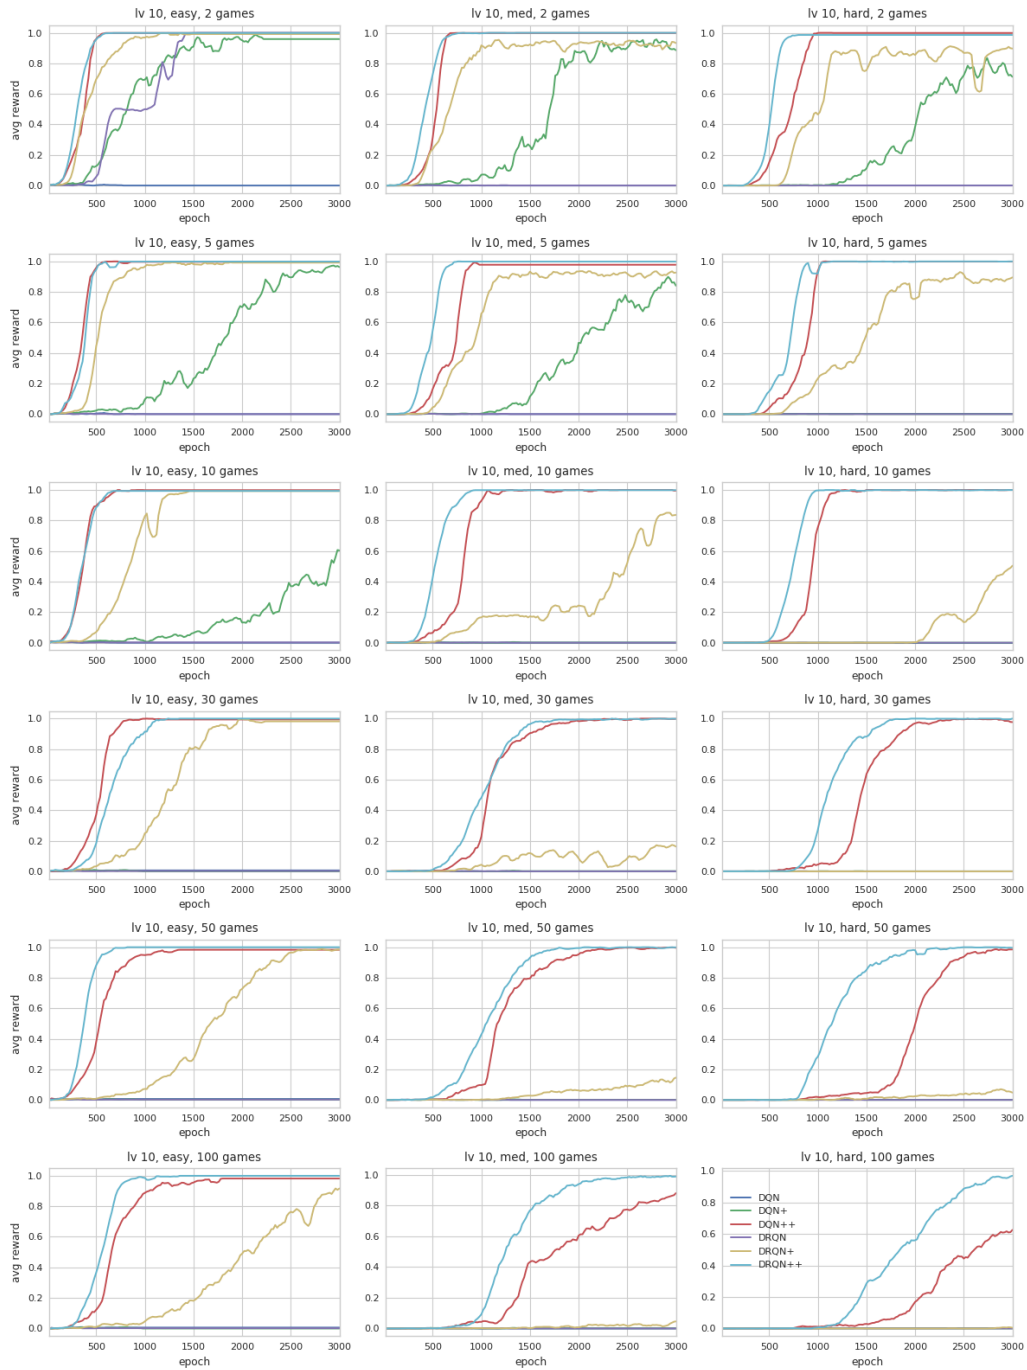


Figure 7: Model performance on multiple games.

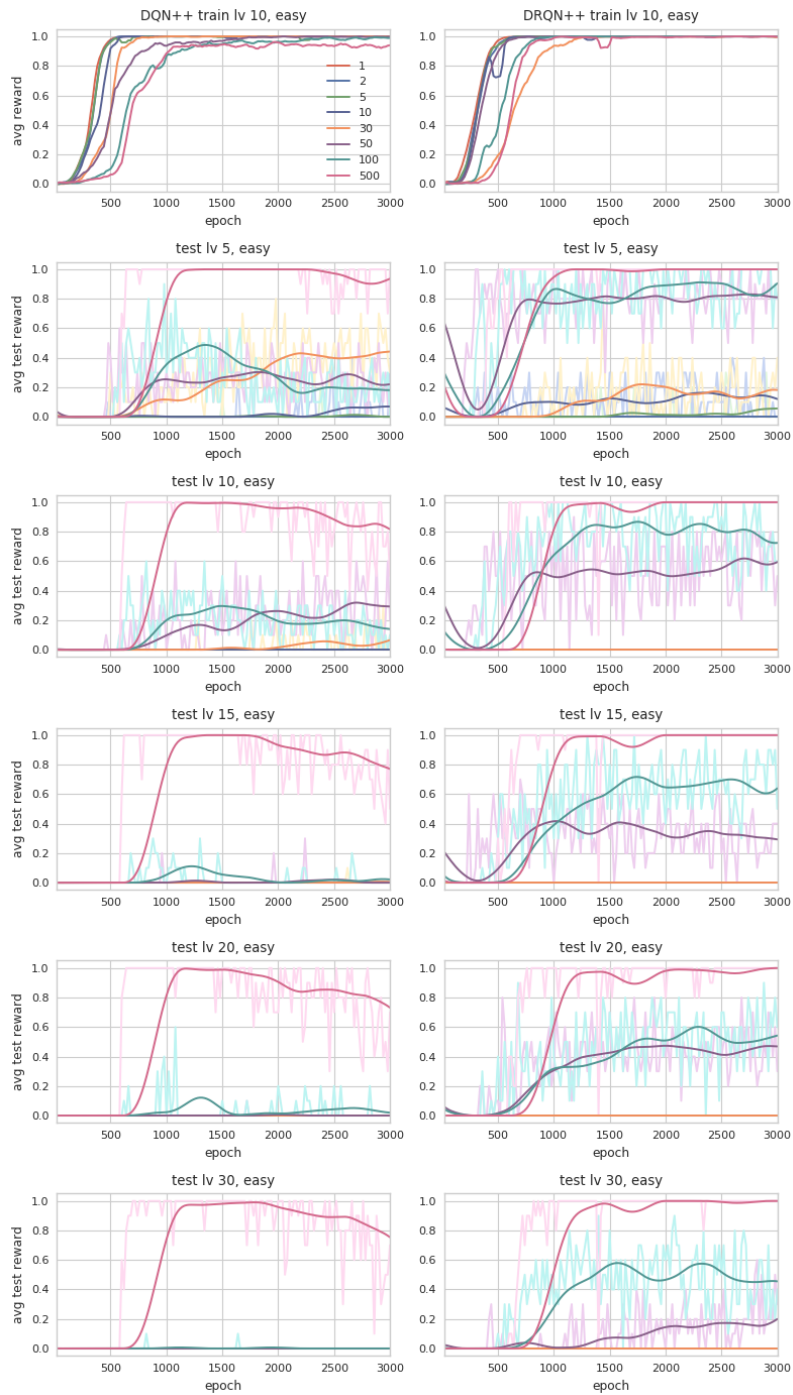


Figure 8: Model performance on unseen easy test games when pre-trained on easy games.

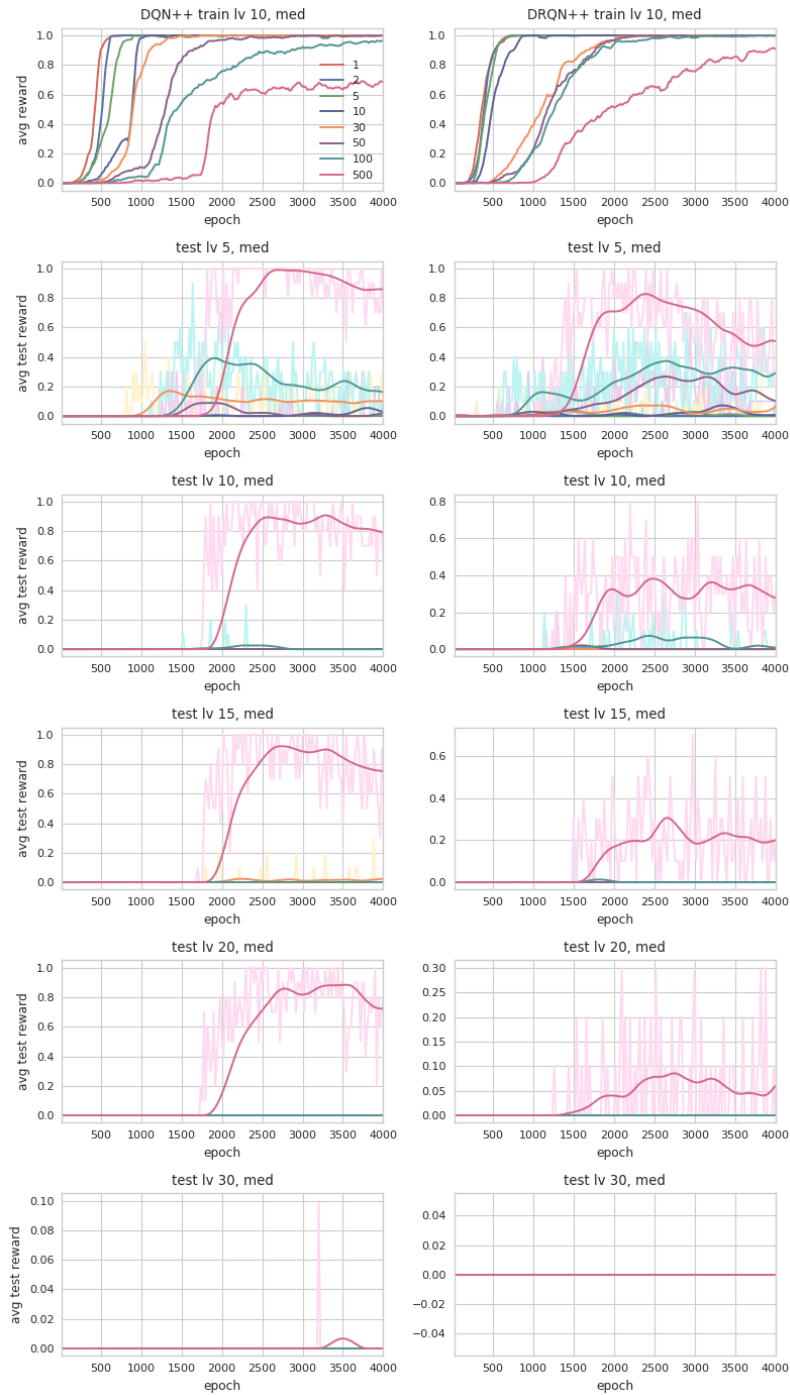


Figure 9: Model performance on unseen medium test games when pre-trained on medium games.

Appendix C. Text-based Chain Experiment

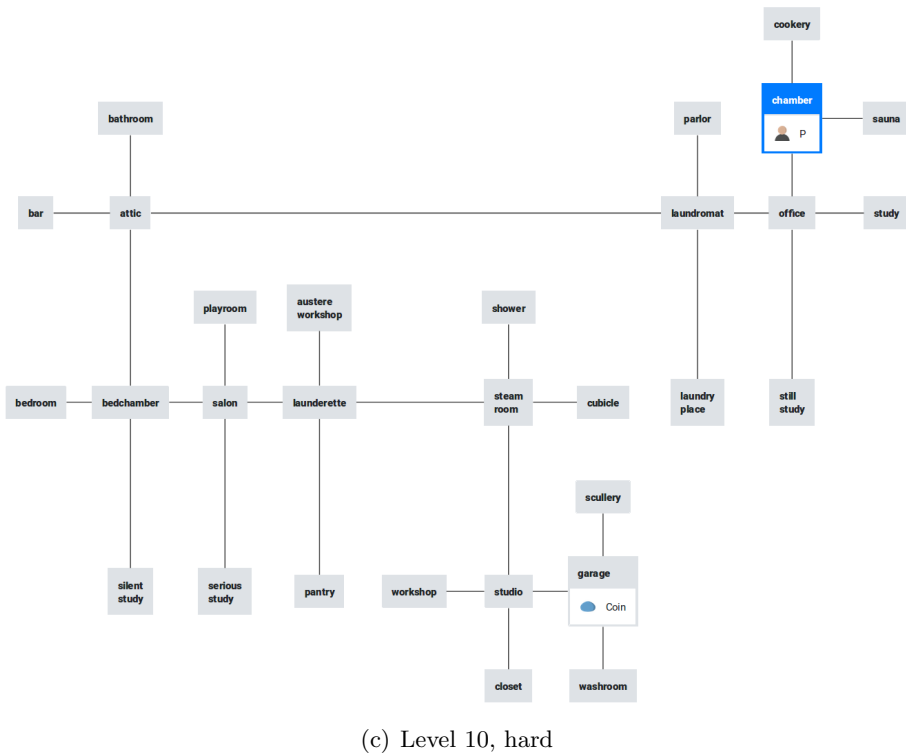
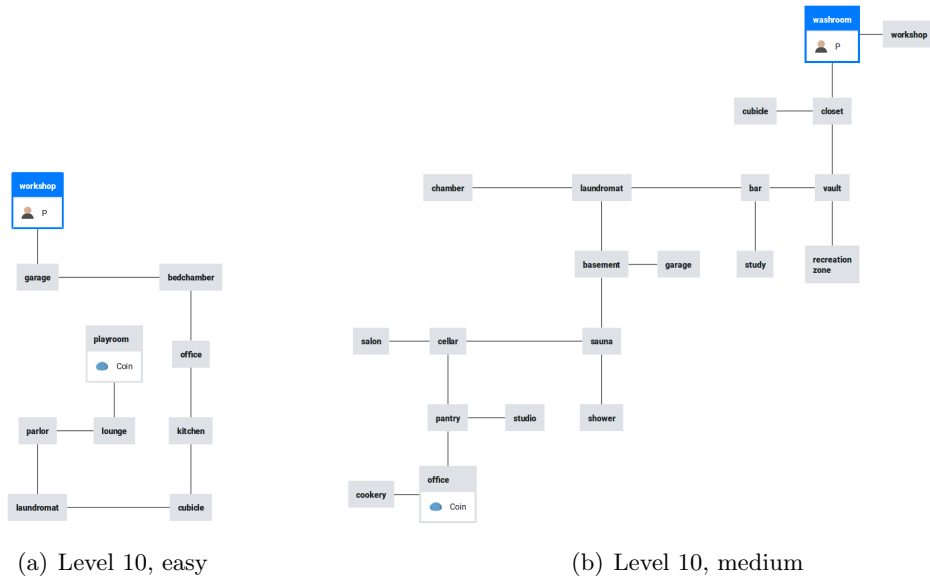


Figure 10: Examples of the games used in the experiments

```
I hope your ready to go into rooms and interact with objects, because you've
just entered TextWorld! take the coin that's in the launderette.

-= Cookhouse -=
Fancy seeing you here. Here, by the way, being the cookhouse.

There is an unguarded exit to the south.
> go south

-= Study -=
You've entered a study. The room seems oddly familiar, as though it were only
superficially different from the other rooms in the building.

You need an unguarded exit? You should try going west. You need an unguarded
exit? You should try going north.

>
```

Figure 11: Text the agent gets to observe for one of the level 10 easy games.