

Algorithmic Improvements for Deep Reinforcement Learning Applied to Interactive Fiction

Vishal Jain,^{1,3} William Fedus,^{1,2} Hugo Larochelle,^{1,2,5} Doina Precup,^{1,3,4,5} Marc G. Bellemare^{1,2,3,5}

¹Mila, ²Google Brain, ³McGill University, ⁴DeepMind, ⁵CIFAR Fellow
vishal.jain@mail.mcgill.ca, {liamfedus, hugolarochelle, bellemare}@google.com, dprecup@cs.mcgill.ca

Abstract

Text-based games are a natural challenge domain for deep reinforcement learning algorithms. Their state and action spaces are combinatorially large, their reward function is sparse, and they are partially observable: the agent is informed of the consequences of its actions through textual feedback. In this paper we emphasize this latter point and consider the design of a deep reinforcement learning agent that can play from feedback alone. Our design recognizes and takes advantage of the structural characteristics of text-based games. We first propose a contextualisation mechanism, based on accumulated reward, which simplifies the learning problem and mitigates partial observability. We then study different methods that rely on the notion that most actions are ineffectual in any given situation, following Zahavy et al.'s idea of an admissible action. We evaluate these techniques in a series of text-based games of increasing difficulty based on the TextWorld framework, as well as the iconic game ZORK. Empirically, we find that these techniques improve the performance of a baseline deep reinforcement learning agent applied to text-based games.

1 Introduction

In a text-based game, also called *interactive fiction* (IF), an agent interacts with its environment through a natural language interface. Actions consist of short textual commands, while observations are paragraphs describing the outcome of these actions (Figure 1). Recently, interactive fiction has emerged as an important challenge for AI techniques (Atkinson et al. 2018), in great part because the genre combines natural language with sequential decision-making.

From a reinforcement learning perspective, IF domains pose a number of challenges. First, the state space is typically combinatorial in nature, due to the presence of objects and characters with which the player can interact. Since any natural language sentence may be given as a valid command, the action space is similarly combinatorial. The player observes its environment through feedback in natural language, making this a partially observable problem. The reward structure is usually sparse, with non-zero rewards only received when the agent accomplishes something meaningful,

>look

West of House

This is an open field west of a white house, with a boarded front door.
There is a small mailbox here.
A rubber mat saying 'Welcome to Zork!' lies by the door.

>open mailbox

You open the mailbox, revealing a small leaflet.

>take leaflet

Taken.

>inventory

You are carrying:
A small leaflet

Figure 1: The introductory gameplay from ZORK.

such as retrieving an important object or unlocking a new part of the domain.

We are particularly interested in bringing deep reinforcement learning techniques to bear on this problem. In this paper, we consider how to design an agent architecture that can learn to play text adventure games from feedback alone. Despite the inherent challenges of the domain, we identify three structural aspects that make progress possible:

- **Rewards from subtasks.** The optimal behaviour completes a series of subtasks towards the eventual game end;
- **Transition structure.** Most actions have no effect in a given state;
- **Memory as state.** Remembering key past events is often sufficient to deal with partial observability.

While these properties have been remarked on in previous work (Narasimhan, Kulkarni, and Barzilay 2015; Zahavy et al. 2018), here we relax some of the assumptions previously made and provide fresh tools to more tractably solve IF domains. More generally, we believe these tools to be useful in partially observable domains with similar structure.

Our first contribution takes advantage of the special reward structure of IF domains. In IF, the accumulated reward within an episode correlates with the number of completed

subtasks and provides a good proxy for an agent’s progress. Our *score contextualisation* architecture makes use of this fact by defining a piecewise value function composed of different deep network heads, where each piece corresponds to a particular level of cumulative reward. This separation allows the network to learn separate value functions for different portions of the complete task; in particular, when the problem is linear (i.e., there is a fixed ordering in which subtasks must be completed), our method can be used to learn a separate value function for each subtask.

Our second contribution extends the work of Zahavy et al. (2018) on action elimination. We make exploration and action selection more tractable by determining which actions are *admissible* in the current state. Formally, we say that an action is admissible if it leads to a change in the underlying game state. While the set of available actions is typically large in IF domains, there are usually few commands that are actually admissible in any particular context. Since the state is not directly observable, we first learn an LSTM-based *auxiliary classifier* that predicts which actions are admissible given the agent’s history of recent feedback. We use the predicted probability of an action being admissible to modulate or *gate* which actions are available to the agent at each time step. We propose and compare three simple modulation methods: masking, drop out, and finally consistent Q-learning (Bellemare et al. 2016b). Compared to Zahavy et al.’s algorithm, our techniques are simpler in spirit and can be learned from feedback alone.

We show the effectiveness of our methods on a suite of seven IF problems of increasing difficulty generated using the TextWorld platform (Côté et al. 2018). We find that combining the score contextualisation approach to an otherwise standard recurrent deep RL architecture leads to faster learning than when using a single value function. Furthermore, our action gating mechanism enables the learning agent to progress on the harder levels of our suite of problems.

2 Problem Setting

We represent an interactive fiction environment as a partially observable Markov decision process (POMDP) with deterministic observations. This POMDP is summarized by the tuple $(\mathcal{S}, \mathcal{A}, P, r, \mathcal{O}, \psi, \gamma)$, where \mathcal{S} is the state space, \mathcal{A} the action space, P is the transition function, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function, and $\gamma \in [0, 1)$ is the discount factor. The function $\psi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathcal{O}$ describes the observation $o = \psi(s, a, s')$ provided to the agent when action a is taken in state s and leads to state s' .

Throughout we will make use of standard notions from reinforcement learning (Sutton and Barto 1998) as adapted to the POMDP literature (McCallum 1995; Silver and Veness 2010). At time step t , the agent selects an action according to a policy π which maps a *history* $h_t := o_1, a_1, \dots, o_t$ to a distribution over actions, denoted $\pi(\cdot | h_t)$. This history is a sequence of observations and actions which, from the agent’s perspective, replaces the unobserved environment state s_t . We denote by $B(s | h_t)$ the probability or *belief* of being in state s after observing h_t . Finally, we will find it convenient to rely on time indices to indicate the relationship between a history h_t and its successor, and denote by h_{t+1} the history

resulting from taking action a_t in h_t and observing o_{t+1} as emitted by the hidden state s_{t+1} .

The action-value function Q^π describes the expected discounted sum of rewards when choosing action a after observing history h_t , and subsequently following policy π :

$$Q^\pi(h_t, a) = \mathbb{E} \left[\sum_{i \geq 0} \gamma^i r(s_{t+i}, a_{t+i}) \mid h_t, a \right],$$

where we assume that the action at time $t + j$ is drawn from $\pi(\cdot | h_{t+j})$; note that the reward depends on the sequence of hidden states s_{t+1}, s_{t+2}, \dots implied by the belief state $B(\cdot | h_t)$. The action-value function satisfies the Bellman equation over histories

$$Q^\pi(h_t, a) = \mathbb{E}_{s_t, s_{t+1}} \left[r(s_t, a) + \gamma \max_{a' \in \mathcal{A}} Q^\pi(h_{t+1}, a') \right].$$

When the state is observed at each step ($\mathcal{O} = \mathcal{S}$), this simplifies to the usual Bellman equation for Markov decision processes:

$$Q^\pi(s_t, a) = r(s_t, a) + \gamma \mathbb{E}_{s_{t+1} \sim P} \max_{a' \in \mathcal{A}} Q^\pi(s_{t+1}, a'). \quad (1)$$

In the fully observable case we will conflate s_t and h_t .

The Q-learning algorithm (Watkins 1989) over histories maintains an approximate action-value function Q which is updated from samples h_t, a_t, r_t, o_{t+1} using a step-size parameter $\alpha \in [0, 1)$:

$$\begin{aligned} Q(h_t, a_t) &\leftarrow Q(h_t, a_t) + \alpha \delta_t \\ \delta_t &= r_t + \gamma \max_{a \in \mathcal{A}} Q(h_{t+1}, a) - Q(h_t, a_t). \end{aligned} \quad (2)$$

Q-learning is used to estimate the *optimal action-value function* attained by a policy which maximizes Q^π for all histories. In the context of our work, we will assume that this policy exists. Storing this action-value function in a lookup table is impractical, as there are in general an exponential number of histories to consider. Instead, we use recurrent neural networks approximate the Q-learning process.

2.1 Consistent Q-Learning

Consistent Q-learning (Bellemare et al. 2016b) learns a value function which is consistent with respect to a local form of policy stationarity. Defined for a Markov decision process, it replaces the term δ_t in (2) by

$$\delta_t^{\text{cql}} = r_t + \begin{cases} \gamma \max_{a \in \mathcal{A}} Q(s_{t+1}, a) - Q(s_t, a_t) & s_{t+1} \neq s_t \\ (\gamma - 1)Q(s_t, a_t) & s_{t+1} = s_t. \end{cases} \quad (3)$$

Consistent Q-learning can be shown to decrease the action-value of suboptimal actions while maintaining the action-value of the optimal action, leading to larger *action gaps* and a potentially easier value estimation problem.

Observe that consistent Q-learning is not immediately adaptable to the history-based formulation, since h_{t+1} and h_t are sequences of different lengths (and therefore not comparable). One of our contributions in this paper is to derive a related algorithm suited to the history-based setting.

2.2 Admissible Actions

We will make use of the notion of an *admissible action*, following terminology by Zahavy et al. (2018).¹

Definition 1. An action a is admissible in state s if

$$P(s | s, a) < 1.$$

That is, a is admissible in s if its application may result in a change in the environment state. When $P(s | s, a) = 1$, we say that an action is inadmissible.

We extend the notion of admissibility to histories as follows. We say that an action a is admissible given a history h if it is admissible in *some* state that is possible given h , or equivalently:

$$\sum_{s \in S} B(s | h) P(s | s, a) < 1.$$

We denote by $\xi(s) \subseteq \mathcal{A}$ the set of admissible actions in state s . Abusing notation, we define the *admissibility function*

$$\begin{aligned} \xi(s, a) &:= \mathbb{I}_{[a \in \xi(s)]} \\ \xi(h, a) &:= \Pr\{a \in \xi(S)\}, S \sim B(\cdot | h). \end{aligned}$$

We write \mathcal{A}_t for the set of admissible actions given history h_t , i.e. the actions whose admissibility in h_t is strictly greater than zero. In IF domains, inadmissible actions are usually dominated, and we will deprioritize or altogether rule them out based on our estimate of $\xi(h, a)$.

3 More Efficient Learning for IF Domains

We are interested in learning an action-value function which is close to optimal and from which can be derived a near-optimal policy. We would also like learning to proceed in a sample-efficient manner. In the context of IF domains, this is hindered by both the partially observable nature of the environment and the size of the action space. In this paper we propose two complementary ideas that alleviate some of the issues caused by partial observability and large action sets. The first idea contextualizes the action-value function on a surrogate notion of progress based on total reward so far, while the second seeks to eliminate inadmissible actions from the exploration and learning process.

Although our ideas are broadly applicable, for concreteness we describe their implementation in a deep reinforcement learning framework. Our agent architecture (Figure 2) is derived from the LSTM-DRQN agent (Yuan et al. 2018) and the work of Narasimhan, Kulkarni, and Barzilay (2015).

3.1 Score Contextualisation

In applying reinforcement learning to games, it is by now customary to translate the player’s score differential into rewards (Bellemare et al. 2013; OpenAI 2018). Our setting is similar to Arcade Learning Environment in the sense that the environment provides the score. In IF, the player is

¹Note that our definition technically differs from Zahavy et al. (2018)’s, who define an admissible action as one that is not ruled out by the learning algorithm.

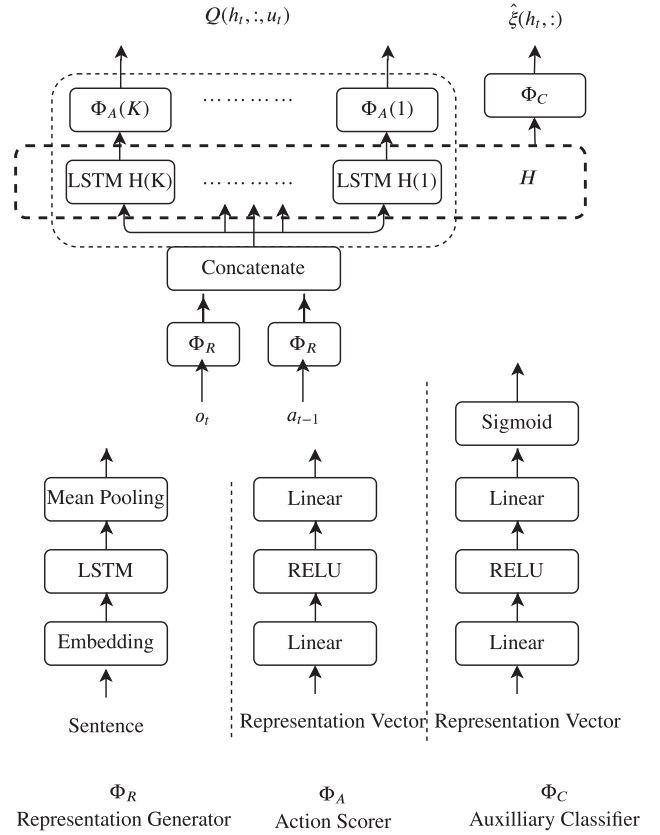


Figure 2: Our IF architecture consists of three modules: a representation generator Φ_R that learns an embedding for a sentence, an action scorer Φ_A that chooses a network head i (a feed-forward network) conditional on score u_t , learns its Q-values and outputs $Q(h_t, :, u_t)$ and finally, an auxiliary classifier Φ_C that learns an approximate admissibility function $\hat{\xi}(h_t, :)$. The architecture is trained end-to-end.

awarded points for acquiring an important object, or completing some task relevant to progressing through the game. These awards occur in a linear, or almost linear structure, reflecting the agent’s progression through the story, and are relatively sparse. We emphasize that this is in contrast to the more general reinforcement learning setting, which may provide reward for surviving, or achieving something at a certain rate. In the video game SPACE INVADERS, for example, the notion of “finishing the game” is ill-defined: the player’s objective is to keep increasing their score until they run out of lives.

We make use of the IF reward structure as follows. We call *score* the agent’s total (undiscounted) reward since the beginning of an episode, remarking that the term extends beyond game-like domains. At time step t , the score u_t is

$$u_t := \sum_{i=0}^{t-1} r_i.$$

In IF domains, where the score reflects the agent’s progress,

it is reasonable to treat it as a state variable. We propose maintaining a separate action-value function for each possible score. This action-value function is denoted $Q(h_t, a_t, u_t)$. We call this approach *score contextualisation*. The use of additional context variables has by now been demonstrated in a number of settings (Rakelly et al. (2019); Icarte et al. (2018); Ghosh et al. (2018)). First, credit assignment becomes easier since the score provides clues as to the hidden state. Second, in settings with function approximation we expect optimization to be simpler since for each u , the function $Q(\cdot, \cdot, u)$ needs only be trained on a subset of the data, and hence can focus on features relevant to this part of the environment.

In a deep network, we implement score contextualisation using K network heads and a map $\mathcal{J} : \mathbb{N} \rightarrow \{1, \dots, K\}$ such that the $\mathcal{J}(u_t)^{th}$ head is used when the agent has received a score of u_t at time t . This provides the flexibility to either map each score to a separate network head, or multiple scores to one head. Taking $K = 1$ uses one monolithic network for all subtasks, and fully relies on this network to identify state from feedback. In our experiments, we assign scores to networks heads using a round-robin scheme with a fixed K . Using Narasimhan, Kulkarni, and Barzilay (2015)’s terminology, our architecture consists of a shared *representation generator* Φ_R with K independent LSTM heads, followed by a feed-forward *action scorer* $\Phi_A(i)$ which outputs the action-values (Figure 2).

3.2 Action Gating Based on Admissibility

In this section we revisit the idea of using the admissibility function to eliminate or more generally *gate* actions. Consider an action a which is inadmissible in state s . By definition, taking this action does not affect the state. We further assume that inadmissible actions produce a constant level of reward, which we take to be 0 without loss of generality:

$$a \text{ inadmissible in } s \implies r(s, a) = 0.$$

This assumption is reasonable in IF domains, and more generally holds true in domains that exhibit subtask structure, such as the video game MONTEZUMA’S REVENGE (Belle-mare et al. 2016a). We can combine knowledge of P and r for inadmissible actions with Bellman’s equation (1) to deduce that for any policy π ,

$$a \text{ inadmissible in } s \implies Q^\pi(s, a) \leq \max_{a' \in \mathcal{A}} Q^\pi(s, a') \quad (4)$$

If we know that a is inadmissible, then we do not need to learn its action-value.

We propose learning a classifier whose purpose is to predict the admissibility function. Given a history h , this classifier outputs, for each action a , the probability $\hat{\xi}(h, a)$ that this action is admissible. Because of state aliasing, this probability is in general strictly between 0 and 1; furthermore, it may be inaccurate due to approximation error. We therefore consider action gating schemes that are sensitive to intermediate values of $\hat{\xi}(h, a)$. The first two schemes produce an approximately admissible set $\hat{\mathcal{A}}_t$ which varies from time step to time step; the third directly uses the definition of admissibility in a history-based implementation of the consistent Bellman operator.

Dropout. The *dropout* method randomly adds each action a to $\hat{\mathcal{A}}_t$ with probability $\hat{\xi}(h_t, a)$.

Masking. The *masking* method uses an elimination threshold $c \in [0, 1]$. The set $\hat{\mathcal{A}}_t$ contains all actions a whose estimated admissibility is at least c :

$$\hat{\mathcal{A}}_t := \{a : \hat{\xi}(h_t, a) \geq c\}.$$

The masking method is a simplified version of Zahavy et al. (2018)’s action elimination algorithm, whose threshold is adaptively determined from a confidence interval, itself derived from assuming a value function and admissibility functions that can be expressed linearly in terms of some feature vector.

In both the dropout and masking methods, we use the action set $\hat{\mathcal{A}}_t$ in lieu of the full action set \mathcal{A} when selecting exploratory actions.

Consistent Q-learning for histories (CQLH). The third method leaves the action set unchanged, but instead drives the action-values of purportedly inadmissible actions to 0. This is done by adapting the consistent Bellman operator (3) to the history-based setting. First, we replace the indicator $\mathbb{I}_{[s_{t+1} \neq s_t]}$ by the probability $\hat{\xi}_t := \hat{\xi}(h_t, a_t)$. Second, we drive $Q(s_t, a_t)$ to 0 in the case when we believe the state is unchanged, following the argumentation of (4). This yields a version of consistent Q-learning which is adapted to histories, and makes use of the predicted admissibility:

$$\begin{aligned} \delta_t^{\text{CQLH}} = & r_t + \gamma \max_{a \in \mathcal{A}} Q(h_{t+1}, a) \hat{\xi}_t \\ & + \gamma Q(h_t, a_t) (1 - \hat{\xi}_t) - Q(h_t, a_t). \end{aligned} \quad (5)$$

One may ask whether this method is equivalent to a belief-state average of consistent Q-learning when $\hat{\xi}(h_t, a_t)$ is accurate, i.e. equals $\xi(h_t, a_t)$. In general, this is not the case: the admissibility of an action depends on the hidden state, which in turn influences the action-value at the next step. As a result, the above method may underestimate action-values when there is state aliasing (e.g., $\hat{\xi}(h_t, a_t) \approx 0.5$), and yields smaller action gaps than the state-based version when $\hat{\xi}(h_t, a_t) = 1$. However, when a_t is known to be inadmissible ($\hat{\xi}(h_t, a_t) = 0$), the methods do coincide, justifying its use as an action gating scheme.

We implement these ideas using an *auxiliary classifier* Φ_C . For each action a , this classifier outputs the estimated probability $\hat{\xi}(h_t, a)$, parametrized as a sigmoid function. These probabilities are learned from bandit feedback: after choosing a from history h_t , the agent receives a binary signal e_t as to whether a was admissible or not. In our setting, learning this classifier is particularly challenging because the agent must predict admissibility solely based on the history h_t . As a point of comparison, using the information-gathering commands LOOK and INVENTORY to establish the state, as proposed by Zahavy et al. (2018), leads to a simpler learning problem, but one which does not consider the full history. The need to learn $\hat{\xi}(h_t, a)$ from bandit feedback also encourages methods that generalize across histories and textual descriptions.

4 A Synthetic IF Benchmark

Both score contextualisation and action gating are tailored to domains that exhibit the structure typical of interactive fiction. To assess how useful these methods are, we will make use of a synthetic benchmark based on the TextWorld framework (Côté et al. 2018). TextWorld provides a reinforcement learning interface to text-based games along with an environment specification language for designing new environments. Environments provide a set of locations, or *rooms*, objects that can be picked up and carried between locations, and a reward function based on interacting with these objects. Following the genre, special *key* objects are used to access parts of the environment.

Our benchmark provides seven environments of increasing complexity, which we call *levels*. We control complexity by adding new rooms and/or objects to each successive level. Each level also requires the agent to complete a number of subtasks (Table 1), most of which involve carrying one or more items to a particular location. Reward is provided only when the agent completes one of these subtasks. Thematically, each level involves collecting food items to make a salad, inspired by the first TextWorld competition. Example objects include an apple and a head of lettuce, while example actions include `get apple` and `slice lettuce with knife`. Accordingly we call our benchmark SaladWorld.

SaladWorld provides a graded measure of an agent architecture’s ability to deal with both partial observability and large action spaces. Indeed, completing each subtask requires memory of what has previously been accomplished, along with where different objects are. Together with this, each level in the SaladWorld involves some amount of history-dependent admissibility i.e the admissibility of the action depends on the history rather than the state. For example, `put lettuce on counter` can only be accomplished once `take lettuce` (in a different room) has happened. Keys pose an additional difficulty as they do not themselves provide reward. As shown in Table 1, the number of possible actions rapidly increases with the number of objects in a given level. Even the small number of rooms and objects considered here preclude the use of tabular representations, as the state space for a given level is the exponentially-sized cross-product of possible object and agent locations. In fact, we have purposefully designed SaladWorld as a small challenge for IF agents, and even our best method falls short of solving the harder levels within the allotted training time. Full details are given in Table 2 in the appendix.

5 Empirical Analysis

In the first set of experiments, we use SaladWorld to establish that both score contextualisation and action gating provide positive benefits in the context of IF domain. We then validate these findings on the celebrated text-based game ZORK used in prior work (Fulda et al. 2017; Zahavy et al. 2018).

Our baseline agent is the LSTM-DRQN agent (Yuan et al. 2018) but with a different action representation. We augment

Table 1: Main characteristics of each level in our synthetic benchmark.

LEVEL	# ROOMS	# OBJECTS	# SUB-TASKS	$ \mathcal{A} $
1	4	2	2	8
2	7	4	3	15
3	7	4	3	15
4	9	8	4	50
5	11	15	5	141
6	12	20	6	283
7	12	20	7	295

this baseline with either or both score contextualisation and action gating, and observe the resulting effect on agent performance in SaladWorld. We measure this performance as the fraction of subtasks completed during an episode, averaged over time. In all cases, our results are generated from 5 independent trials of each condition. To smooth the results, we use moving average with a window of 20,000 training steps. The graphs and the histograms report average \pm std. deviation across the trials.

Score contextualisation uses $K = 5$ network heads; the baseline corresponds to $K = 1$. Each head is trained using the Adam optimizer (Kingma and Ba 2015) with a learning rate $\alpha = 0.001$ to minimize a Q-learning loss (Mnih et al. 2015) with a discount factor of $\gamma = 0.9$. The auxiliary classifier Φ_C is trained with the binary cross-entropy loss over the selected action’s admissibility (recall that our agent only observes the admissibility function for the selected action). Training is done using a balanced form of prioritized replay which we found improves baseline performance appreciably. Specifically, we use the sampling mechanism described in Hausknecht and Stone (2015) with prioritization i.e we sample τ_p fraction of episodes that had atleast one positive reward, τ_n fraction with atleast one negative reward and $1 - \tau_p - \tau_n$ from whole episodic memory \mathcal{D} . For prioritization, $\tau_p = \tau_n = 0.25$.

Actions are chosen from the estimated admissible set $\hat{\mathcal{A}}_t$ according to an ϵ -greedy rule, with ϵ annealed linearly from 1.0 to 0.1 over the first million training steps. To simplify exploration, our agent further takes a forced LOOK action every 20 steps. Each episode lasts for a maximum T steps. For Level 1 game, $T = 100$, whereas for rest of the levels $T = 200$.

5.1 Score Contextualisation

We first consider the effect of score contextualisation on our agents’ ability to complete tasks in SaladWorld. We ask,

Does score contextualisation mitigate the negative effects of partial observability?

We begin in a simplified setting where the agent knows the admissible set \mathcal{A}_t . We call this setting *oracle gating*. This setting lets us focus on the impact of contextualisation alone. We compare our score contextualisation (SC) to the baseline and also to two “tabular” agents. The first tabular agent treats the most recent feedback as state, and hashes each unique description-action pair to a Q-value. This results in

Table 2: Subtasks information and scores possible for each level of the suite.

Level	Subtasks	Possible Scores
1	Following subtasks with reward and fulfilling condition: <ul style="list-style-type: none"> • 10 points when the agent first enters the vegetable market. • 5 points when the agent gets lettuce from the vegetable market and puts lettuce on the counter. 	10, 15
2	All subtasks from previous level plus this subtask: <ul style="list-style-type: none"> • 5 points when the agent takes the blue key from open space, opens the blue door, gets tomato from the supermarket and puts it on the counter in the kitchen. 	5, 10, 15, 20
3	All subtasks from level 1 plus this subtask: <ul style="list-style-type: none"> • 5 points when the agent takes the blue key from open space, goes to the garden, opens the blue door with the blue key, gets tomato from the supermarket and puts it on the counter in the kitchen. <p>Remark: Level 3 game differs from Level 2 game in terms of number of steps required to complete the additional sub-task (which is greater in case of Level 3)</p>	5, 10, 15, 20
4	All subtasks from previous level plus this subtask: <ul style="list-style-type: none"> • 5 points when the agent takes parsley from the backyard and knife from the cutlery shop to the kitchen, puts parsley into fridge and knife on the counter. 	5, 10, 15, 20, 25
5	All subtasks from previous level plus this subtask: <ul style="list-style-type: none"> • 5 points when the agent goes to fruit shop, takes chest key, opens container with chest key, takes the banana from the chest and puts it into the fridge in the kitchen. 	5, 10, 15, 20, 25, 30
6	All subtasks from previous level plus this subtask: <ul style="list-style-type: none"> • 5 points when the agent takes the red key from the supermarket, goes to the playroom, opens the red door with the red key, gets the apple from cookhouse and puts it into the fridge in the kitchen. 	5, 10, 15, 20, 25, 30, 35
7	All subtasks from previous level plus this subtask: <ul style="list-style-type: none"> • 5 points when the agent prepares the meal. 	5, 10, 15, 20, 25, 30, 35, 40

a memoryless scheme that ignores partial observability. The second tabular agent performs the information-gathering actions LOOK and INVENTORY to construct its state description, and also hashes these to unique Q-values. Accordingly, we call this the “LI-tabular” agent. This latter scheme has proved to be a successful heuristic in the design of IF agents (Fulda et al. 2017), but can be problematic in domains where taking information-gathering actions can have negative consequences (as is the case in ZORK).

Figure 3 shows the performance of the four methods across SaladWorld levels, after 1.3 million training steps. We observe that the tabular agents’ performance suffers as soon as there are multiple subtasks, as expected. The baseline agent performs well up to the third level, but then shows significantly reduced performance. We hypothesize that this occurs because the baseline agent must estimate the hidden state from longer history sequences and effectively learn an implicit contextualisation. Beyond the fourth level, the performance of all agents suffers, suggesting the need for

a better exploration strategy, for example using expert data (Tessler et al. 2019).

We find that score contextualisation performs better than the baseline when the admissible set is unknown. Figure 4 compares learning curves of the SC and baseline agents with oracle gating and using the full action set, respectively, in the simplest of levels (Level 1 and 2). We find that score contextualisation can learn to solve these levels even without access to \mathcal{A}_t , whereas the baseline cannot. Our results also show that oracle gating simplifies the problem, and illustrate the value in handling inadmissible actions differently.

We hypothesize that score contextualisation results in a simpler learning problem in which the agent can more easily learn to distinguish which actions are relevant to the task, and hence facilitate credit assignment. Our result indicates that it might be unreasonable to expect contextualisation to arise naturally (or easily) in partially observable domains with large actions sets. We conclude that score contextualisation mitigates the negative effects of partial observability.




Figure 3: Fraction of tasks solved by each method at the end of training for 1.3 million steps. The tabular agents, which do not take history into account, perform quite poorly. LI stands for “look, inventory” (see text for details).

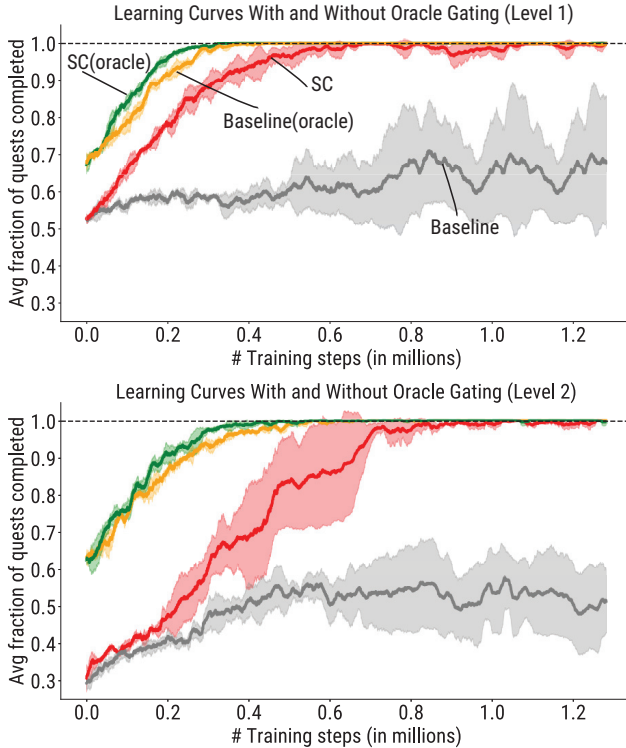


Figure 4: Comparing whether score contextualisation as an architecture provides a useful representation for learning to act optimally. Row 1 and 2 correspond to Level 1 and 2 respectively.

5.2 Score Contextualisation with Learned Action Gating

The previous experiment (in particular, Figure 4) shows the value of restricting action selection to admissible actions. With the goal in mind of designing an agent that can operate from feedback alone, we now ask:

Can an agent learn more efficiently when given bandit feedback about the admissibility of its chosen actions?

We address this question by comparing our three action gating mechanisms. As discussed in Section 3.2, the output of




Figure 5: Fraction of tasks solved by each method at the end of training for 1.3 million steps. Except in Level 1, action gating by itself does not improve end performance.




Figure 6: Effectiveness of action gating with score contextualisation in Level 3. Of the three methods, masking performs best.

the auxiliary classifier describes our estimate of an action’s admissibility for a given history.

As an initial point of comparison, we tested the performance of the baseline agent when using the auxiliary classifier’s output to gate actions. For the masking method, we selected $c = 0.001$ from a larger initial parameter sweep. The results are summarized in Figure 5. While action gating alone provides some benefits in the first level, performance is equivalent for the rest of the levels.

However, when combined with score contextualisation (see Fig 6, 7), we observe some performance gains. In Level 3 in particular, we almost recover the performance of the SC agent with oracle gating. From our results we conclude that masking with the right threshold works best, but leave as an open question whether the other action gating schemes can be improved.

Figure 8 shows the final comparison between the baseline LSTM-DRQN and our new agent architecture which incorporates action gating and score contextualisation. Our results show that the augmented method significantly outperforms the baseline, and is able to handle more complex IF domains. From level 4 onwards, the learning curves in the appendix show that combining score contextualisation with masking results in faster learning, even though final performance is unchanged. We posit that better exploration schemes are required for further progress in SaladWorld.




Figure 7: Fraction of tasks solved by each method at the end of training for 1.3 million steps. For first 3 levels, SC + Masking is better or equivalent to SC. For levels 4 and beyond, better exploration strategies are required.




Figure 8: Score contextualisation and masking compared to the baseline agent. We show the fraction of tasks solved by each method at the end of training for 1.3 million steps.

5.3 Zork

As a final experiment, we evaluate our agent architecture on the interactive fiction ZORK I, the first installment of the popular trilogy. ZORK provides an interesting point of comparison for our methods, as it is designed by and for humans – following the ontology of Bellemare et al. (2013), it is a domain which is both *interesting* and *independent*. Our main objective is to compare the different methods studied with Zahavy et al. (2018)’s AE-DQN agent. Following their experimental setup, we take $\gamma = 0.8$ and train for 2 million steps. All agents use the smaller action set (131 actions). Unlike AE-DQN, however, our agent does not use information-gathering actions (LOOK and INVENTORY) to establish the state.

Figure 9 shows the corresponding learning curves. Despite operating in a harder regime than AE-DQN, the score contextualizing agent reaches a score comparable to AE-DQN, in about half of the training steps. All agents eventually fail to pass the 35-point benchmark, which corresponds to a particularly difficult in-game task (the “troll quest”) which involves a timing element, and we hypothesize requires a more intelligent exploration strategy.

6 Related Work

RL applied to Text Adventure games: LSTM-DQN by Narasimhan, Kulkarni, and Barzilay (2015) deals with parser-based text adventure games and uses an LSTM to generate feedback representation. The representation is then




Figure 9: Learning curves for different agents in Zork.

used by an action scorer to generate scores for the action verb and objects. The two scores are then averaged to determine Q-value for the state-action pair. In the realm of choice-based games, He et al. (2016) uses two separate deep neural nets to generate representation for feedback and action respectively. Q-values are calculated by dot-product of these representations. None of the above approaches deals with partial observability in text adventure games.

Admissible action set learning: Tao et al. (2018) approach the issue of learning admissible set given context as a supervised learning one. They train their model on (input, label) pairs where input is context (concatenation of feedbacks by LOOK and INVENTORY) and label is the list of admissible commands given this input. AE-DQN (Zahavy et al. 2018) employs an additional neural network to prune in-admissible actions from action set given a state. Although the paper doesn’t deal with partial observability in text adventure games, authors show that having a tractable admissible action set led to faster convergence. Fulda et al. (2017) work on bounding the action set through affordances. Their agent is trained through tabular Q-Learning.

Partial Observability: Yuan et al. (2018) replace the shared MLP in Narasimhan, Kulkarni, and Barzilay (2015) with an LSTM cell to calculate context representation. However, they use concatenation of feedbacks by LOOK and INVENTORY as the given state to make the game more observable. Their work also doesn’t focus on pruning in-admissible actions given a context. Finally, Ammanabrolu and Riedl (2019) deal with partial observability by representing state as a knowledge graph and continuously updating it after every game step. However, the graph update rules are hand-coded; it would be interesting to see they can be learned during gameplay.

7 Conclusions and Future work

We introduced two algorithmic improvements for deep reinforcement learning applied to interactive fiction (IF). While naturally rooted in IF, we believe our ideas extend more generally to partially observable domains and large discrete action spaces. Our results on SaladWorld and ZORK show the usefulness of these improvements. Going forward, we believe better contextualisation mechanisms should yield fur-

ther gains. In ZORK, in particular, we hypothesize that going beyond the 35-point limit will require more tightly coupling exploration with representation learning.

8 Acknowledgments

This work was funded by the CIFAR Learning in Machines and Brains program. Authors thank Compute Canada for providing the computational resources.

References

- Ammanabrolu, P., and Riedl, M. 2019. Playing text-adventure games with graph-based deep reinforcement learning. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 3557–3565. Minneapolis, Minnesota: Association for Computational Linguistics.
- Atkinson, T.; Baier, H.; Copplestone, T.; Devlin, S.; and Swan, J. 2018. The text-based adventure AI competition. *IEEE Transactions on Games*.
- 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. 2016a. Unifying count-based exploration and intrinsic motivation. In Lee, D. D.; Sugiyama, M.; Luxburg, U. V.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems* 29. Curran Associates, Inc. 1471–1479.
- Bellemare, M. G.; Ostrovski, G.; Guez, A.; Thomas, P. S.; and Munos, R. 2016b. Increasing the action gap: New operators for reinforcement learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, 1476–1483. AAAI Press.
- Côté, M.-A.; Kádár, A.; Yuan, X.; Kybartas, B.; Barnes, T.; Fine, E.; Moore, J.; Hausknecht, M.; Asri, L. E.; Adada, M.; Tay, W.; and Trischler, A. 2018. Textworld: A learning environment for text-based games. *CoRR* abs/1806.11532.
- Fulda, N.; Ricks, D.; Murdoch, B.; and Wingate, D. 2017. What can you do with a rock? affordance extraction via word embeddings. *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*.
- Ghosh, D.; Singh, A.; Rajeswaran, A.; Kumar, V.; and Levine, S. 2018. Divide-and-conquer reinforcement learning. *ICLR 2018*.
- Hausknecht, M., and Stone, P. 2015. Deep recurrent q-learning for partially observable mdps. In *AAAI Fall Symposium on Sequential Decision Making for Intelligent Agents (AAAI-SDMIA15)*.
- He, J.; Chen, J.; He, X.; Gao, J.; Li, L.; Deng, L.; and Ostendorf, M. 2016. Deep reinforcement learning with a natural language action space. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1621–1630. Berlin, Germany: Association for Computational Linguistics.
- Icarte, R. T.; Klassen, T. Q.; Valenzano, R.; and McIlraith, S. A. 2018. Using reward machines for high-level task specification and decomposition in reinforcement learning. In *Proceedings of the International Conference on Machine Learning*.
- Kingma, D. P., and Ba, J. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*.
- McCallum, A. K. 1995. *Reinforcement learning with selective perception and hidden state*. Ph.D. Dissertation, University of Rochester.
- 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.
- Narasimhan, K.; Kulkarni, T.; and Barzilay, R. 2015. Language understanding for text-based games using deep reinforcement learning. *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- OpenAI. 2018. Gym retro. <https://github.com/openai/retro>.
- Rakelly, K.; Zhou, A.; Quillen, D.; Finn, C.; and Levine, S. 2019. Efficient off-policy meta-reinforcement learning via probabilistic context variables. *ICML 2019*.
- Silver, D., and Veness, J. 2010. Monte-carlo planning in large pomdps. In *Advances in Neural Information Processing Systems*.
- Sutton, R. S., and Barto, A. G. 1998. *Reinforcement learning: An introduction*. MIT Press.
- Tao, R. Y.; Côté, M.-A.; Yuan, X.; and Asri, L. E. 2018. Towards solving text-based games by producing adaptive action spaces.
- Tessler, C.; Zahavy, T.; Cohen, D.; Mankowitz, D. J.; and Mannor, S. 2019. Sparse imitation learning for text based games with combinatorial action spaces. *The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM) 2019*.
- Watkins, C. J. C. H. 1989. *Learning from delayed rewards*. Ph.D. Dissertation, Cambridge University, Cambridge, England.
- Yuan, X.; Côté, M.-A.; Sordoni, A.; Laroché, R.; des Combes, R. T.; Hausknecht, M.; and Trischler, A. 2018. Counting to explore and generalize in text-based games.
- Zahavy, T.; Haroush, M.; Merlis, N.; Mankowitz, D. J.; and Mannor, S. 2018. Learn what not to learn: Action elimination with deep reinforcement learning. In *Proceedings of the 32Nd International Conference on Neural Information Processing Systems, NIPS’18*, 3566–3577. USA: Curran Associates Inc.