Better Safe than Sorry: Evidence Accumulation Allows for Safe Reinforcement Learning

Akshat Agarwal[*] [†]
Robotics Institute
Carnegie Mellon University

Abhinau Kumar V[*]
Department of Electrical Engineering
Indian Institute of Technology Hyderabad

Kyle Dunovan, Erik Peterson, Timothy Verstynen
Department of Psychology
Carnegie Mellon University

Katia Sycara
Robotics Institute
Carnegie Mellon University

Abstract

In the real world, agents often have to operate in situations with incomplete information, limited sensing capabilities, and inherently stochastic environments, making individual observations incomplete and unreliable. Moreover, in many situations it is preferable to delay a decision rather than run the risk of making a bad decision. In such situations it is necessary to aggregate information before taking an action; however, most state of the art reinforcement learning (RL) algorithms are biased towards taking actions *at every time step*, even if the agent is not particularly confident in its chosen action. This lack of caution can lead the agent to make critical mistakes, regardless of prior experience and acclimation to the environment. Motivated by theories of dynamic resolution of uncertainty during decision making in biological brains, we propose a simple accumulator module which accumulates evidence in favor of each possible decision, encodes uncertainty as a dynamic competition between actions, and acts on the environment only when it is sufficiently confident in the chosen action. The agent makes no decision by default, and the burden of proof to make a decision falls on the policy to accrue evidence strongly in favor of a single decision. Our results show that this accumulator module achieves near-optimal performance on a simple guessing game, far outperforming deep recurrent networks using traditional, forced action selection policies.

1 Introduction

Traditional reinforcement learning (RL) algorithms map the state of the world to an action so as to maximize a reward signal it receives from environmental feedback. With the success of deep RL, this action-value mapping is increasingly being approximated by deep neural networks to maximize success in complex tasks, such as Atari games (Mnih et al. 2015) and Go (Silver et al. 2016). In the real world, RL agents usually operate with incomplete information about their surroundings due to a range of issues issues such as limited sensor coverage, noisy data, occlusions, and the inherent randomness in the environment which comes from factors that can't be modeled. With individual observations being incomplete and/or unreliable, it

is imperative that agents accrue sufficient evidence in order to make the most task or environmentally appropriate decision. In current RL approaches, this is accomplished by using recurrent layers in the neural network to aggregate information (Lample and Chaplot 2017; Agarwal et al. 2018; Zoph and Le 2016); however, this pipeline is biased towards taking a decision at every time step, even if the agent is not particularly confident in any of the possible actions. This is highly undesirable, especially in situations where an incorrect action could be catastrophic (or very heavily penalized). The usual mechanism to allow the possibility of not taking an action at any time step is to have one possible action as a 'No-Op', which can be chosen by the agent when it does not want to act on the environment. This requires the policy to actively choose to not act, which is a counter-intuitive requirement for any real world scenario, where not taking an action should be the default setting.

In biological networks, the circuit-level computations of Cortico-Basal-Ganglia-Thalamic (CBGT) pathways (Mink 1996) are ideally suited for performing the multiple sequential probability ratio test (MSPRT) (Wald 1945; Bogacz and Gurney 2007; Bogacz 2007), a simple algorithm of information integration that optimally selects single actions from a competing set of alternatives based on differences in input evidence (Draglia, Tartakovsky, and Veeravalli 1999; Baum and Veeravalli 1994). Motivated by these theories of dynamic resolution of decision uncertainty in the CBGT pathways in mammalian brains (see also (Redgrave, Prescott, and Gurney 1999; Mink 1996; Dunovan and Verstynen 2016)), *we propose modifying existing RL architectures by replacing the policy/Q-value output layers with an accumulator module* that makes a decision by accumulating the evidence for alternative actions until a threshold is met. Each possible action is represented by a channel through which environmental input is sampled and accumulated as evidence at each time step, and an action is chosen only when the evidence in one of the channels crosses a certain threshold. This ensures that when the environment is stochastic and uncertainty is high, the agent can exercise greater caution by postponing the decision to act until sufficient evidence has been accumulated, thereby avoiding catastrophic outcomes. While evidence accumulation necessarily comes at a cost to decision speed, there are many real world scenarios in which longer decision times are consid-

---

[*]Denotes equal contribution

[†]Corresponding Author: aa7@cmu.edu

ered a perfectly acceptable price to pay for assurances that those decisions will be both safe and accurate.

The accumulator module can work with both tabular and deep reinforcement learning, with on-policy and off-policy RL algorithms, and can be trained via backpropagation. We present a simple guessing task where the environment is partially observable, and show that a state of the art RL algorithm (A2C-RNN (Mnih et al. 2016)) fails to learn the task when using traditional, forced action selection policies (even when equipped with a 'No-Op'), but achieves near-optimal performance when allowed to accumulate evidence before acting.

2 Related Work

Partially Observable Markov Decision Processes (POMDPs) (Kaelbling, Littman, and Cassandra 1998; Jaakkola, Singh, and Jordan 1995; Kimura, Miyazaki, and Kobayashi 1997) are the de-facto choice for modeling partially observable stochastic domains. Hausknecht and Stone (2015) first successfully used an LSTM layer in a DQN (Mnih et al. 2015). Since then, it has become a standard part of Deep RL architectures, including both on-policy and off-policy RL algorithms (Agarwal, Hope, and Sycara 2018; Lample and Chaplot 2017). Another strategy consists of using Hidden Markov Models (Monahan 1982) to learn a model of the environment, for domains where the environment is itself Markovian, but does not appear to be so to the agent because of partial observability.

The implementation of accumulation-to-threshold dynamics in single neurons, where inputs are accumulated over time as a sub-threshold change in potential until a threshold is reached, causing the cell to "fire", has been studied to a great extent as spiking neural networks (O'Connor and Welling 2016) for supervised learning using backpropagation (Lee, Delbruck, and Pfeiffer 2016). Each neuron in the neural network is replaced by a Stochastic/Leaky Integrate-and-Fire neuron, with Winner-Take-All (WTA) circuits. Florian (2007) presented a reinforcement learning algorithm for spiking networks through modulation of spike timing-dependent plasticity. Zambrano, Roelfsema, and Bohte (2015) also presented a continuous-time on-policy RL algorithm to learn task-specific working memory in order to decouple action duration from the internal time-steps of the RL model using a Winner-Take-All action selection mechanism. The approach taken here differs from these previous examples in three key ways: first, we consider simple additive accumulators without any leakage; second, the dynamic competition between neurons is modeled using center-surround inhibition, allowing between-channel dynamics to modulate the evidence criterion; and third, evidence accumulation is restricted to neurons in the last (e.g., output) layer of the network.

Prior safe reinforcement learning models (García and Fernández 2015) have primarily been divided into two lines of work - the first is based on modification of optimality criteria to incorporate worst-case criteria, or risk-sensitive criteria. The second line focuses on the modification of the exploration process through incorporation of external knowledge, teacher guidance or risk-directed exploration. Re-
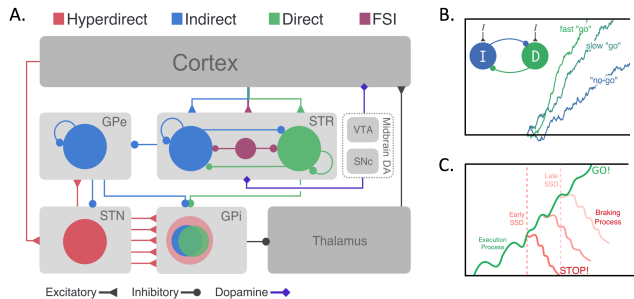

Figure 1: **Cortico-Basal-Ganglia-Thalamus (CBGT) networks & the dependent process accumulator model.** (A) CBGT circuit. Striatum, STR; fast-spiking interneurons, FSI; globus pallidus external, GPE; globus pallidus internal (GPi); subthalamic nucleus, STN; ventral tegmental area, VTA; substantia nigra pars compacta, SNc. (B) Competing pathways model of the direct (D) & indirect (I) pathways. (C) The dependent process model. Panels A and B were recreated with permission from (Dunovan and Verstynen 2016) and Panel C was recreated with permission from (Dunovan et al. 2015).

cently, Lipton et al. (2016) used intrinsic fear and reward shaping to learn and avoid a set of dangerous (catastrophic) states, in the context of lifelong RL. Chow et al. (2018) used Lyapunov functions to guarantee the safety of a behavior policy during training via a set of local, linear constraints. These works focus on a different aspect of safety in reinforcement learning, and are complementary to ours.

3 Decision Making in the Brain

We now briefly describe the dependent process model of decision making in the brain, which serves as the biological inspiration for evidence accumulation in reinforcement learning. Decision making in CBGT circuits can be modeled as an interaction of three parallel pathways: the direct pathway, the indirect pathway and the hyperdirect pathway (see Fig.1A). The direct and indirect pathways act as gates for action selection, with direct pathway facilitating and indirect pathway inhibiting action selection. These pathways converge on a common output nucleus (the GPi). From a computational perspective, this convergence of the direct and indirect pathways suggests that their competition encodes the rate of evidence accumulation in favor of a given action (Bogacz and Gurney 2007; Dunovan and Verstynen 2016), resulting in action execution if the direct pathway sufficiently overpowers the indirect pathway. The decision speed is thus modulated by the degree of response conflict across actions (see Fig. 1B). A second inhibitory pathway (the hyperdirect pathway) globally suppresses all action decisions when the system is going to make an inappropriate response, with the competition between all three major pathways formalized by the so-called dependent process model of CBGT computations (see Fig. 1C) (Dunovan et al. 2015). The center-surround architecture of the CBGT network (such that inputs to a direct pathway for one action also excite indi-

Figure 2: Flowchart describing agent-environment interaction

rect pathways for alternative actions) (Mink 1996), as well as the competitive nature of direct and indirect pathways within a single action channel (Bogacz and Gurney 2007; Dunovan and Verstynen 2016) allows for modulation of both the rate and threshold of the evidence accumulation process (Dunovan and Verstynen 2017). Moreover, this selection mechanism implicitly handles situations in which no action is required from the agent, as evidence simply remains at sub-threshold levels until a significant change is registered in the environment (Dunovan and Verstynen 2016).

This dynamic selection process runs in sharp contrast to standard Deep RL methods (Mnih et al. 2015; 2016). Deep RL operates at a fixed frame rate, mandating actions to be taken with a particular fixed frequency even in times of high uncertainty or times when actions might not be needed. Deep RL uses backpropagation to modulate representations of state-value and action-value (a process analogous, but not identical, to how the dopamine projections to cortex alter cortical representations), whereas the actual gating units are static units with no feedback-dependent plasticity. We posit that incorporating additional plasticity at the output layer holds significant promise to improving existing Deep RL algorithms, as adaption of the selection process will interact with the action representation process to facilitate complex action repertoires.

4 Methods

4.1 Mode Estimation (ME) Task

We propose a simple episodic task where the agent receives a sample from a discrete unimodal distribution at each time step, and has to estimate the mode of the distribution. Within each episode, the agent receives a sample from the environment distribution at each step, and has the choice to make a decision (estimate the mode) or not. If the agent chooses not to make a decision at that time step, it simply receives another sample from the environment. The episode ends either when the agent makes a guess, or if the maximum allowed length of an episode, T_{max}, is exceeded. The beginning of each new episode resets the environment and randomly changes the distribution from which the samples being observed by the agent are generated.

Conditioned on the spread of the hidden environmental distribution, the agent must learn to delay its decision and aggregate samples over multiple time steps to make an in-

formed estimate of the mode. The agent receives feedback from the environment in the form of a reward r_t at the end of each episode, such that the agent is rewarded for making the correct decision and penalized for tardiness and for making an incorrect decision (or no decision at all).

$$
r_t = \begin{cases} 0 & t \leq T_{max} \text{ and No guess} \\ R_1 - (t-1) & t \leq T_{max} \text{ and Correct guess} \\ R_2 & t \leq T_{max} \text{ and Incorrect guess} \\ R_3 & t > T_{max} \end{cases} \tag{1}
$$

where t is the number of time steps the agent waited and accumulated information before making a guess. R_1 is the reward the agent receives if it guesses correctly at the first time step, R_2 is the penalty the agent receives for making an incorrect guess and R_3 for not making a guess at all within T_{max} time steps. The reward received by the agent decays linearly with the number of time steps it waits before making a decision, requiring it to balance the trade-off between making a decision quickly and making a decision accurately. This becomes especially relevant at higher values of ϵ, when the observations are very noisy. The interaction between the agent and the environment is illustrated in Fig.2.

4.2 The Accumulator Module

At each time step, the agent receives an observation o_t from the environment, from which it extracts an evidence vector $\kappa_t = f(o_t), \kappa_t \in \mathbb{R}^n$, consisting of one component for each available action. The cumulative evidence received since the beginning of each episode ($t = 0$) is stored in accumulator channels ν

$$
\nu^i = \sum_{t=0} \kappa_t^i \tag{2}
$$

where κ_t^i is the i-th component of the evidence vector κ_t. The preference vector ρ over the action choices is given by a softmax over ν, such that

$$
\rho^i = \frac{\exp(\nu^i)}{\sum_{i=1}^{A} \exp(\nu^i)} \tag{3}
$$

ρ encodes the agent's confidence in actions - if the accumulated evidence values ν_i are high for multiple actions, then the preference values ρ_i for all of them will be relatively low, indicating that the agent is not very confident in any particular action choice. Since we do not want any decision to be made in such a situation, action i is taken only when ρ_i crosses some threshold τ, failing which no guess is made and the agent keeps observing more information from the environment, until the time when it becomes sufficiently confident to act on the environment. Note that the evidence accumulation process, as defined here, mirrors the hypothesis proposed in Bogacz and Gurney (2007) regarding how the basal ganglia and cortex implement optimal decision making between alternative actions.

4.3 Learning Algorithm

We use the Advantage Actor-Critic (A2C) algorithm (Mnih et al. 2016) to learn the RNN, the accumulator threshold

τ and the evidence mapping f in our experiments below. This is a policy gradient method, which performs an approximate gradient descent on the agent's discounted return $G_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k+1}$. The A2C gradient is as follows:

$$
\begin{aligned}
& (G_t - v_\theta(s_t)) \nabla_\theta \log \pi_\theta(a_t|s_t) \\
& + \eta(G_t - v_\theta(s_t)) \nabla_\theta v_\theta(s_t) \\
& + \beta \sum_a \pi_\theta(a|s) \log \pi_\theta(a|s)
\end{aligned} \quad (4)
$$

where s_t is the observation, a_t the action selected by the policy π_θ defined by a deep neural network with parameters θ and $v_\theta(s)$ is a value function estimate of the expected return $\mathbb{E}[G_t|s_t = s]$ produced by the same network. Instead of the full return, we use a 1-step return $G_t = r_{t+1} + \gamma v_\theta(s_{t+1})$ in the gradient above. The last term regularizes the policy towards larger entropy, which promotes exploration, and β is a hyper-parameter which controls the importance of entropy in the overall gradient. We keep the value loss coefficient η and discount factor γ fixed at 1 and 0.95, respectively, for all the experiments.

5 Experiments and Results

The simple mathematical structure of the Mode Estimation task allows us to find the optimal values of the accumulator threshold τ using Monte Carlo simulations, providing a good reference to compare the performance of our learning algorithms with. We first measure the performance of recurrent actor-critic policy gradient RL with a forced action-selection policy, verify that it is unable to learn anything meaningful when the environment stochasticity ϵ is high, and then demonstrate how using the accumulator module achieves near-optimal performance on a wide range of ϵ values. We train the accumulator threshold directly (with observations as evidence) with Advantage Actor-Critic (A2C) (Mnih et al. 2016), and then successfully jointly train deep networks to learn both the evidence mapping f and accumulator threshold τ values using A2C.

5.1 Task Instantiation

In the particular instance of the Mode Estimation task we use for running experiments, the environment chooses an integer, say n_0, uniformly at random from the set of integers $\mathbb{Z} = \{0, 1, 2, \cdots, 9\}$. Then, at each step during that episode, the agent receives an observation $n \in \mathbb{Z}$, with probability p such that

$$
p(n) = \begin{cases} 1 - \epsilon & \text{if } n = n_0 \\ \frac{\epsilon}{9} & \text{otherwise} \end{cases} \quad (5)
$$

where ϵ is an environment parameter encoding the amount of randomness/noise inherent in the environment. The agent's task is to correctly guess the mode n_0 for that particular episode, based on these noisy observations. As soon as the agent makes a guess, the episode resets (a new n_0 is chosen). The reward received by the agent follows Eqn. 1, with $R_1 = -R_2 = -R_3 = 30$ and $T_{max} = 30$.

5.2 Baseline Monte-Carlo Estimates of Accumulator Performance

The advantage of using a simple task for evaluation is that we can obtain Monte-Carlo (MC) estimates of the best expected performance of the accumulator model for various values of the environment's randomness parameters ϵ. The accumulator is parameterized by only one hyperparameter - the threshold τ. Since τ is compared with components of the preference vector ρ, which are the output of a softmax operation, $\tau \in (0, 1)$ covers the entire range of values which are useful. The agent receives observations in the form of one-hot vector representations of the distribution samples, which are directly treated as evidence to be accumulated. For each value of $\tau \in \{0, 0.1, 0.2, \ldots 0.9\}$, we complete 10,000 episode rollouts, tracking the reward received in each episode. The threshold with the highest expected reward is selected, and that reward is used as a near-optimal estimate of the accumulator's performance. This process is repeated for environments with varying levels of stochasticity, specifically, with $\epsilon \in \{0, 0.2, 0.4, 0.6, 0.8\}$. In Fig. 3, the accuracy, decision time and reward achieved by the optimal thresholds for each ϵ value are plotted with a dashed red line, and the expected reward received with the optimal thresholds is specified in Row 1 of Table 1. Note that since we use the same discretization of τ (or a subset of it) when learning the threshold in subsequent sections, these are the highest possible rewards that our learning agents could receive, which is also reflected in Fig. 3, where the rewards achieved by any learning agent never exceed the MC estimates.

5.3 Recurrent A2C with Forced Action Selection

Using traditional, forced action selection policies, we now train a recurrent policy for the Mode Estimation task. To allow a fair comparison with the evidence accumulators, the agent is given an additional 'No-Op' action output which allows the agent the possibility of choosing not to make a decision and wait for more samples. We call this agent the 'A2C-RNN' agent. Note that recurrent policies are the state of the art method of dealing with partial observability in deep reinforcement learning (Hausknecht and Stone 2015; Mnih et al. 2016).

The policy network takes as input observations from the environment, and outputs (a) a probability distribution over the actions, and (b) a value function estimate of the expected return. The observation is a 4-dimensional binary representation of the sample (e.g. 0110 for 6), and is passed through a linear layer with a ReLU non-linearity to get an output of size 25. This goes through an RNN cell (Elman 1990) of output size 25 with a ReLU non-linearity, and is then 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. The agent has 11 possible action choices (including a 'No-Op', and 10 choices for each of the possible modes).

The agent is trained for 50k episodes with entropy regularization with coefficient $\beta = 5.0$ to encourage exploration. Performance is evaluated after every 500 episodes.

Uncertainty ϵ / Agent	0	0.2	0.4	0.6	0.8
Monte-Carlo Estimate	30	27.6	25	18.4	-8.2
A2C-RNN	30	26.9	7.5	-23	-30
Learning τ	30	27.6	24.9	17.4	-13.7
Joint Training of τ and f	29.7	25.7	22.2	12.5	-25.2

Table 1: The expected rewards received by learning agents in environments with varying levels of uncertainty, after 50k episodes of training. The Monte-Carlo estimates provide a near-optimal baseline to compare the learning approaches with. It can be seen that the joint training method far outperforms A2C-RNN, while the agent learning only the accumulator threshold τ reaches near-optimal values close to the MC estimate.

The learning curves for expected accuracy, decision time and reward achieved by the learned network at each evaluation point are plotted in Fig. 3 using dotted blue lines. Row 2 of Table 1 presents results for the final expected rewards achieved by the policy trained with A2C-RNN. Using the Adam optimizer with learning rate 1×10^{-3}, we find that while the A2C-RNN agent achieves near-optimal performance for low values of environment randomness ($\epsilon = \{0, 0.2\}$), it's performance saturates at a lower reward level for $\epsilon = 0.4$ and it is unable to learn anything meaningful for $\epsilon = \{0.6, 0.8\}$, with the expected reward not increasing from its initial value of -30, which is the lowest reward possible. This clearly shows that the A2C-RNN agent is unable to learn that it should wait, and make a safe decision only when it is confident in its chosen action. We hypothesize that this poor performance in the absence of the accumulator module, especially in environments with high uncertainty, is because learning to wait for long periods of time without having a built-in default 'no-go' mechanism is difficult for any continuous parameterized function to learn, including an RNN. Intuitively, the agent would have to actively choose the 'No-Op' action for multiple time steps (say, the first 10 observations), and then, at the 11th observation, change it's neuron activations to choose the correct mode. In fact, the RNN would be required to have chosen 'No-Op' when it received the exact same observation previously (but it's 'confidence' was low). This sudden change in the output, which has to be precipitated only by the cell state of the RNN (since the input form *does not change*), is difficult for neural networks, which are continuous function approximators, to learn.

5.4 Learning Accumulator Threshold τ

Having established that a state of the art algorithm (A2C-RNN) with a forced action-selection policy is not able to do well at the Mode Estimation task, especially at high levels of environment randomness, we will show that replacing the traditional policy outputs of the actor-critic network with an accumulator module enables the agent to make safe decisions and achieve consistently high rewards in the ME task. In this section, we learn only the accumulator threshold τ, verifying its utility - and in the next, use it to functionally replace the traditional outputs of a policy network.

The agent learns (using RL) a separate **Accumulator Network**, which predicts the optimal threshold as a function of the observation, $\tau = g(o)$. The observation o received by the agent is a 10-dimensional one-hot vector representation of the sample, which is directly treated as the evidence to be accumulated (i.e., $\kappa = f(o) = o$). We note here that directly treating environment observations as evidence to be accumulated is possible only in simple environments as the one chosen here, but will not scale to more complex tasks, for which we jointly train both f and τ (see Section 5.5). The approach used in this subsection, however, is presented as a minimal working example of the accumulator module.

The accumulator network's action space consists of 10 possible values for the threshold $\tau \in \{0.1, 0.2, \ldots 0.9\}$, while the observations it takes in as input are a 10-dimensional one-hot vector representation of the sample. The observation is passed through a linear layer with a ReLU non-linearity to get an output of size 25, which is then passed as input to two linear layers that output the probability distribution over actions (choices of threshold) and the value estimate of the expected return. At every step, the agent's accumulated evidence is compared against the threshold decided by the accumulator network, and the agent makes a decision, or not, accordingly.

The agent trains its accumulator network to choose a threshold that maximizes the rewards r_t returned by the environment using the A2C algorithm and the Adam optimizer with learning rate 1×10^{-4}, and includes entropy regularization with coefficient $\beta = 0.5$ to encourage exploration. It is trained for 50k episodes, with performance evaluated after every 500 episodes. The learning curves are plotted using dash-dot orange lines in Fig. 3, while the final rewards are listed in Row 3 of Table 1. The agent is able to achieve optimal performance matching the MC simulations for all values of ϵ but 0.8. It clearly outperforms the performance of the A2C-RNN agent, both in final performance and sample efficiency.

5.5 Jointly Learning τ and Evidence Mapping f

We have now established the viability of the proposed evidence accumulation mechanism; however, this instantiation comes with meaningful evidence (in the form of one-hot vector representations of the observation) received directly from the environment. In real situations, an agent will first need to extract evidence from its environment. For example, an agent receiving visual observations of its surroundings needs to extract evidence like objects, faces etc. from those images in order to present as evidence to the accumulator. As a simplified version of that we force the agent to learn a meaningful evidence mapping f by providing it
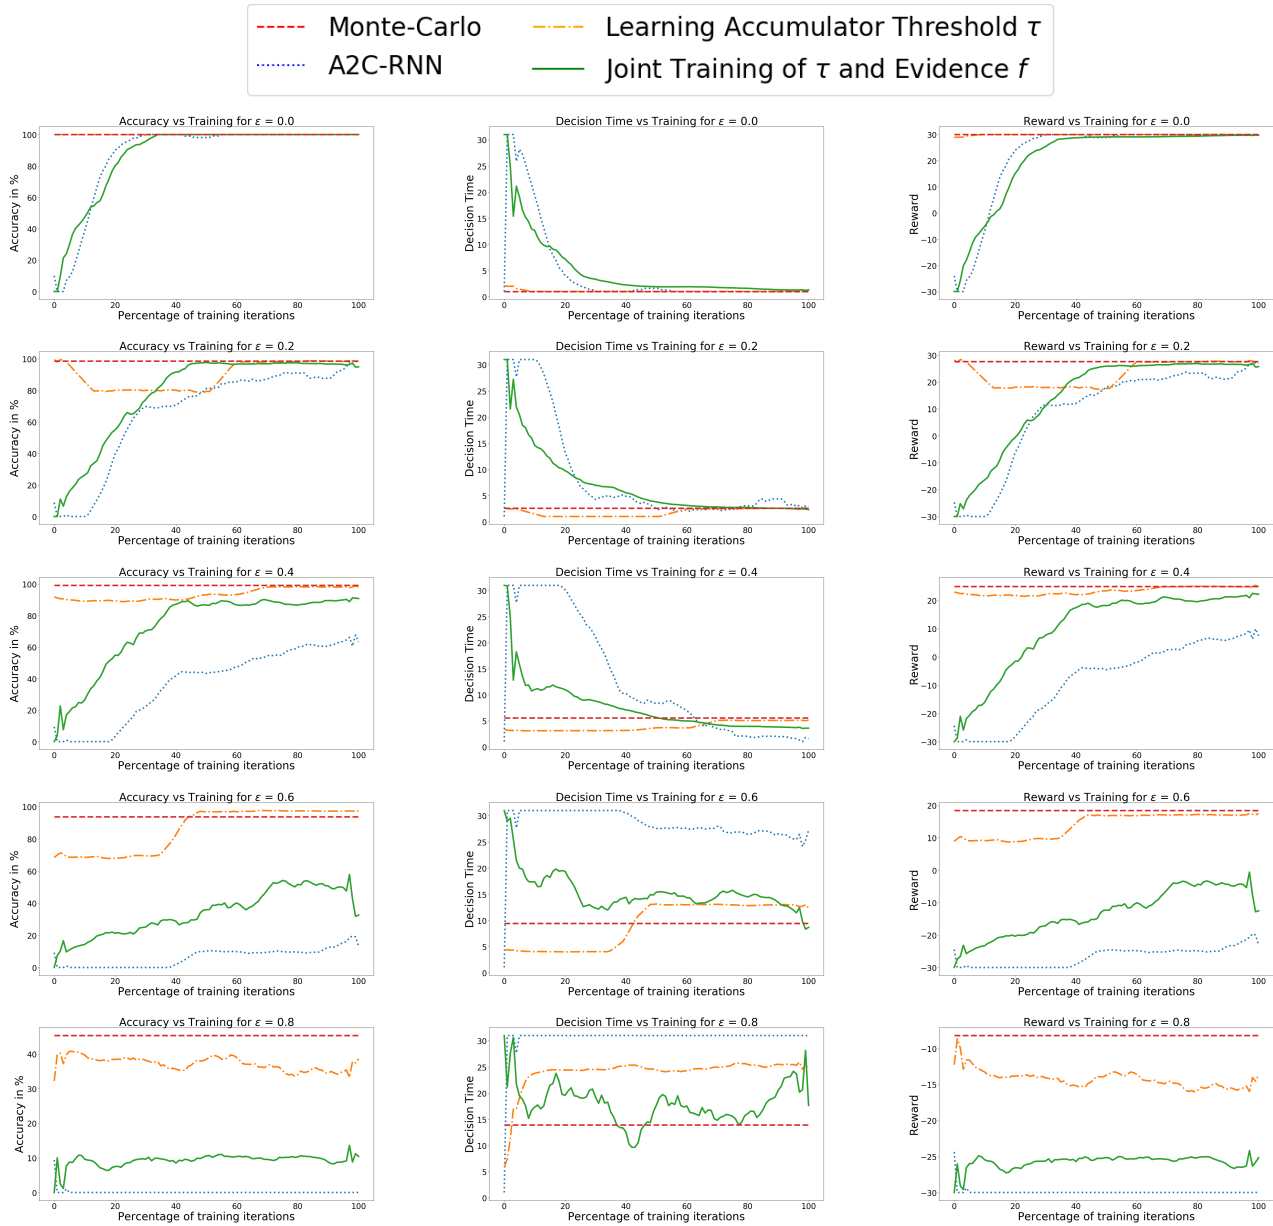
Figure 3: Plots of Accuracy, Decision Time and Reward vs Training Iterations for 5 different values of $\epsilon\{0, 0.2, 0.4, 0.6, 0.8\}$. Each row has 3 plots for Accuracy, Decision Time and Reward (from left to right) for one value of ϵ (increasing from top to bottom). The plots for accuracy and reward often mirror each other but are both shown for completeness. As ϵ increases and the environment becomes more stochastic, the A2C-RNN agent heavily underperforms the accumulator module.

with 4-dimensional binary representations of the samples - requiring the agent to learn to extract evidence κ for the 10 accumulator channels, while simultaneously learning the accumulator threshold τ - which is functionally analogous to replacing the traditional policy outputs in an actor-critic network with an accumulator module.

Here, we make two observations. First, the preference ρ is calculated by a softmax over accumulator channels ν, which means that ρ is only a function of the differences between values in ν. Hence, imposing a uniform lower bound on all the evidence vectors κ_t does not restrict the performance of the accumulator. Second, since the highest threshold value we consider is 0.9, allowing arbitrarily large values in the accumulator channels ν is redundant since any evidence value κ_t^i large enough such that ρ^i exceeds 0.9 would be sufficient. Consequently, we are able to impose a loose upper bound on the evidence values κ^i, restricting $\kappa_t^i \in [0, S] \ \forall i, t$. We can further simplify this by restricting $\kappa_t^i \in [0, 1] \ \forall i, t$ and instead accumulating $\nu^i = \sum(S_t \kappa_t^i)$, where $1/S_t$ can be interpreted as the global sensitivity of the accumulator *across channels*, and can be used to incorporate global action suppression mechanisms mirroring the hyper-direct pathway of the basal ganglia (section 3). For now, we treat S_t as a hyperparameter that does not vary with time.

We now sample each component of the evidence $\kappa_t^i \in [0, 1]$ from a Beta distribution $\mathbb{B}_t^i$, whose concentration parameters α_t^i, β_t^i are predicted conditioned on the environment observation. The 4-dimensional (binary representations of samples) observations are passed through a linear layer of size 20, with ReLU non-linearity. The output is then passed as input to two linear layers that output the α and β parameters, respectively, for all 10 components of the evidence vector, hence defining the distributions $\mathbb{B}_t^i$. We call this neural network, from which evidence κ_t^i is sampled and accumulated in ν, the **Evidence Network**.

Similar to the previous section, a separate **Accumulator Network** is responsible for deciding the accumulator threshold τ, which follows the same architecture and training process as described there, except that we restrict the choice to 5 possible values of the threshold $\tau \in \{0.5, 0.6, \ldots 0.9\}$. This ensures that there is always a single winner, since even with the lowest choice of $\tau = 0.5$, no two components of the preference ρ could exceed 0.5 at the same time.

After both the evidence κ_t and threshold τ are obtained from the Evidence and Accumulator Networks, respectively, the agent accumulates evidence in its accumulator channels ν, calculates the preference ρ, compares it with the threshold τ and accordingly decides whether or not to make a guess (act upon the environment). Both the evidence and accumulator networks are trained using the A2C algorithm, using the same reward (Eqn. 1). We use the Adam optimizer with learning rate of 5×10^{-4} for environments where $\epsilon = \{0, 0.2\}$ and 1×10^{-3} when $\epsilon = \{0.4, 0.6, 0.8\}$. Entropy regularization is used for both networks, with coefficients β 1.0 and 2.0, for the evidence and accumulator networks respectively. The agent trains for 50k episodes with evaluation every 500 episodes, and the learning curves have been plotted using solid green lines in Fig. 3, while the final rewards are listed in Row 4 of Table 1. The jointly trained agent easily outperforms the A2C-RNN agent, learning greater patience and consequently winning greater reward. In the environments with $\epsilon = \{0.6, 0.8\}$, where the A2C-RNN agent does not learn anything, the jointly trained agent learns even greater patience, and achieves significantly better performance than the A2C-RNN agent.

6 Conclusion

In this paper, we propose a modification to existing RL architectures by replacing the policy/Q-value outputs with an accumulator module which sequentially accumulates evidence for each possible action at each time step, acting only when the evidence for one of those actions crosses a certain threshold. This ensures that when the environment is stochastic and uncertainty is high, the agent can exercise greater caution by postponing the decision to act until sufficient evidence has been accumulated, thereby avoiding catastrophic outcomes. We first define a partially observable task where the agent must estimate the mode of a probability distribution from which it is observing samples, and show that a state-of-the-art RL agent (A2C-RNN) is unable to learn even this simple task without an accumulator module, even though it is allowed to choose a 'No-Op' action. We run Monte-Carlo simulations which provide baseline optimal estimates of the performance of the accumulator, and then learn the accumulator threshold as a function of the environment observations, showing that the accumulator module helps the agent achieve near-optimal performance on the task. Recognizing that in more complex real world tasks, the agent will have to extract meaningful evidence from the high-dimensional observations, we also jointly learn the evidence and the threshold, finding that this agent also easily outperforms the A2C-RNN agent, while being equally or more sample efficient.

These results make a strong case for adding an accumulator module to existing Deep RL architectures, especially in real-world scenarios where individual observations are incomplete and unreliable, the cost of making a bad decision is very high, and longer decision times are an acceptable price to pay for assurances that those decisions will be both safe and accurate.

7 Future Work

The Mode Estimation task as defined in this paper is, essentially, a partially observable multi-armed contextual bandit. While the context (the mode of the distribution) is unknown to the agent, it does not transition to different contexts within an episode, as is common in reinforcement learning tasks. We plan to test the accumulator module on tasks with state transitions, and then on more complex domains (such as the Atari games (Bellemare et al. 2013)). Another interesting line of work is to add a global suppression mechanism (similar to the hyperdirect pathway in the CBGT, see Section 3), by allowing the agent to change the sensitivity across accumulator channels based on environmental signals. Having a global stopping mechanism would be very useful for agents operating in very dynamic and reactive environments, such as self-driving vehicles on open roads.

References

Agarwal, A.; Gurumurthy, S.; Sharma, V.; Lewis, M.; and Sycara, K. 2018. Community regularization of visually-grounded dialog. *arXiv preprint arXiv:1808.04359*.

Agarwal, A.; Hope, R.; and Sycara, K. 2018. Challenges of context and time in reinforcement learning: Introducing space fortress as a benchmark. *arXiv preprint arXiv:1809.02206*.

Baum, C. W., and Veeravalli, V. V. 1994. A sequential procedure for multihypothesis testing. *IEEE Transactions on Information Theory* 40(6).

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.

Bogacz, R., and Gurney, K. 2007. The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural computation* 19(2):442–477.

Bogacz, R. 2007. Optimal decision-making theories: linking neurobiology with behaviour. *Trends in cognitive sciences* 11(3):118–125.

Chow, Y.; Nachum, O.; Duenez-Guzman, E.; and Ghavamzadeh, M. 2018. A lyapunov-based approach to safe reinforcement learning. *arXiv preprint arXiv:1805.07708*.

Draglia, V.; Tartakovsky, A. G.; and Veeravalli, V. V. 1999. Multihypothesis sequential probability ratio tests. i. asymptotic optimality. *IEEE Transactions on Information Theory* 45(7):2448–2461.

Dunovan, K., and Verstynen, T. 2016. Believer-skeptic meets actor-critic: Rethinking the role of basal ganglia pathways during decision-making and reinforcement learning. *Frontiers in neuroscience* 10:106.

Dunovan, K., and Verstynen, T. 2017. Errors in action timing and inhibition facilitate learning by tuning distinct mechanisms in the underlying decision process. *bioRxiv* 204867.

Dunovan, K.; Lynch, B.; Molesworth, T.; and Verstynen, T. 2015. Competing basal ganglia pathways determine the difference between stopping and deciding not to go. *Elife* 4:e08723.

Elman, J. L. 1990. Finding structure in time. *Cognitive science* 14(2):179–211.

Florian, R. V. 2007. Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity. *Neural Computation* 19(6):1468–1502.

Garcıa, J., and Fernández, F. 2015. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research* 16(1):1437–1480.

Hausknecht, M., and Stone, P. 2015. Deep recurrent q-learning for partially observable mdps. *CoRR, abs/1507.06527*.

Jaakkola, T.; Singh, S. P.; and Jordan, M. I. 1995. Reinforcement learning algorithm for partially observable markov decision problems. In *Advances in neural information processing systems*, 345–352.

Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence* 101(1-2):99–134.

Kimura, H.; Miyazaki, K.; and Kobayashi, S. 1997. Reinforcement learning in pomdps with function approximation. In *ICML*, volume 97, 152–160.

Lample, G., and Chaplot, D. S. 2017. Playing fps games with deep reinforcement learning. In *AAAI*, 2140–2146.

Lee, J. H.; Delbruck, T.; and Pfeiffer, M. 2016. Training deep spiking neural networks using backpropagation. *Frontiers in neuroscience* 10:508.

Lipton, Z. C.; Azizzadenesheli, K.; Kumar, A.; Li, L.; Gao, J.; and Deng, L. 2016. Combating reinforcement learning's sisyphean curse with intrinsic fear. *arXiv preprint arXiv:1611.01211*.

Mink, J. W. 1996. The basal ganglia: focused selection and inhibition of competing motor programs. *Progress in neurobiology* 50(4):381–425.

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.

Monahan, G. E. 1982. State of the arta survey of partially observable markov decision processes: theory, models, and algorithms. *Management Science* 28(1):1–16.

O'Connor, P., and Welling, M. 2016. Deep spiking networks. *arXiv preprint arXiv:1602.08323*.

Redgrave, P.; Prescott, T. J.; and Gurney, K. 1999. The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience* 89(4):1009–1023.

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.

Wald, A. 1945. Sequential tests of statistical hypotheses. *Annals of Mathematical Statistics* 16(2):117–186.

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.

Zoph, B., and Le, Q. V. 2016. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*.