Explain My Surprise: Learning Efficient Long-Term Memory by Predicting Uncertain Outcomes

Anonymous Author(s)
Affiliation
Address
email

Abstract

In many sequential tasks, a model needs to remember relevant events from the distant past to make correct predictions. Unfortunately, a straightforward application of gradient based training requires intermediate computations to be stored for every element of a sequence. This requires prohibitively large compute memory if a sequence consists of thousands or even millions elements, and as a result, makes learning of very long-term dependencies infeasible. However, the majority of sequence elements can usually be predicted by taking into account only temporally local information. On the other hand, predictions affected by long-term dependencies are sparse and characterized by high uncertainty given only local information. We propose MemUP, a new training method that allows to learn long-term dependencies without backpropagating gradients through the whole sequence at a time. This method can be potentially applied to any gradient based sequence learning. MemUP implementation for recurrent architectures shows performances better or comparable to baselines while requiring significantly less compute memory.

1 Introduction

Two dominating approaches for memory augmentation in deep learning include recurrent neural networks [1; 2; 3; 4; 5; 6; 7; 8; 9]) and Transformer-based [10] architectures [11; 12; 13; 14; 15]. Historically, recurrent networks and Transformers have been primarily developed for applications in natural language processing, which is characterized by strong dependencies between closely located elements. On the other hand, in the reinforcement learning setting the agent often has to remember only a few bits of information but for a very long time [16].

Evolution of neural architectures tackling the problem of learning long-term dependencies follows the path of increasing size and complexity [6; 8; 17]. For example, Transformer-based [18] and memory augmented [7] neural networks perform much better than the classical LSTM in tasks with long-term temporal dependencies. Still, even gigantic language models like GPT-3 [19] process sequences only up to 2048 elements due to quadratic scaling complexity of self-attention. So, these architectures are much more demanding in terms of the memory capacity and computational power required for training.

To learn a long-term dependency between a distant observation in the past and agent's actions in the current state, both the recurrent networks trained with back propagation through time and Transformer-like architectures need to store gradients for all intermediate steps. Taking into account, that the minimum distance between useful information in the past and the moment of its utilization can be measured in thousands or millions of time steps, such solutions seem unrealistic given the current hardware capabilities. We propose a training method that allows an agent to find and store dependencies between temporarily distant events without the need to process intermediate steps.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

Our contributions are the following:

- We propose a new formulation for the problem of memory learning as maximisation of conditional mutual information between memory state and future outcomes.
- We introduce an original method `MemUP` (**Mem**ory for **U**ncertainty **P**rediction) for efficient long-term memory training by minimizing uncertainty of predictions for states where local information is not enough.
- We show that implementation of our method for recurrent neural networks is able to learn long-term dependencies better than baselines while propagating gradients only through a small fractions of a sequence in one optimization step. Notably, the applications of MemUP were successful in supervised learning tasks with different loss functions as well as reinforcement learning tasks.

2 Memory to explain surprise

The main idea of our approach is very simple. Find the most surprising events or elements of a sequence and train a memory to detect and store information that allows to explain out these events. Below we formulate this idea in the framework of information theory.

Lets consider a sequence τ with inputs $\{x_t\}_{t=0}^{T}$ and targets $\{y_t\}_{t=0}^{T}$. A model is trained to predict targets given inputs. An ideal memory could just store all previous inputs $\{x_i\}_{i=0}^{t}$ at each step t and predict y_t from them, but it is too computationally expensive. In practice a `memory model` g_θ has to selectively store some aggregation of this sequence in a memory state $m_t = g_\theta(x_t, m_{t-1})$. Then it should be decided what to write and what to remove from the memory.

Estimating past information importance. Consider an arbitrary time step k. The past element x_t has useful information about the target y_k, if y_k depends on x_t given current x_k, i.e. $p(y_k|x_k, x_t) \neq p(y_k|x_k)$. In this case, we can say that there is a temporal dependency with the length of $k - t$ steps, that starts at step t and ends at step k. The strength of temporal dependency or importance of remembering x_t can be measured as an amount of change in the distribution of y_k given a knowledge of x_t . Multiplication of the both parts of the inequality above by $p(x_t|x_k)$ and application of KL-divergence as a measure of the difference between distributions leads to the conditional mutual information as a measure of x_t importance:

$$I(y_k; x_t|x_k) = \mathbb{E}_{x_k}\mathcal{KL}[p(y_k, x_t|x_k)\|p(y_k|x_k)p(x_t|x_k)], \tag{1}$$

here x_k is taken from sequences stored in some dataset or generated by an agent over interaction with the environment.

Objective function for memory training. Then, the problem of learning what to store in the memory can be defined as maximization of the mutual information between memory states and prediction targets with respect to the parameters θ:

$$\max_{\theta} \sum_{k=t}^{T} I(y_k; m_t = g_\theta(x_t, m_{t-1})|x_k). \tag{2}$$

The sum in eq. 2 allows us to process each memory update in a separate gradient step. As accounting for all future steps prevents loss of information that can be helpful in the distant future. Unfortunately, this still requires the entire remaining (future) sequence to be processed for every update and has the complexity $O(T)$ with respect to the sequence length.

Reducing the cost of memory training. Fortunately, updating the memory for all future steps is not necessary for the majority of real-world tasks. As the distance between events increases, the number of dependencies also decreases. This phenomenon is sometimes called locality of reference [20]. In other words, information from the greater part of past inputs may be useless for y_k prediction given x_k. Accordingly, most of the elements in the sum (eq. 2) will be close to zero. Therefore, depending on the structure of the problem, we could vary the number of elements in the sum, leaving only those for which long-term memory is critical for correct predictions. The main problem is to find such steps in the sequence.

Lets assume that the ideal memory m_t^* that stores all possible information from the past is available. Then the conditional mutual information at a step t:

$$I(y_t, m_t^*|x_t) = H(y_t|x_t) - H(y_t|m_t^*, x_t), \tag{3}$$

shows the potential utility of memory for prediction at this step. When the number of elements in the sum in eq. 2 is minimized to reduce computations, the steps with the highest value of mutual information from eq. 3 should be preserved to assure good quality of prediction. Naturally, we cannot estimate $I(y_t, m_t^*|x_t)$ directly, because m_t^* is not available to us.

As seen from the eq.3, m_t^* participates only in the second entropy, which indicates how well one can predict the future in the presence of ideal memory and ideal model. On the other hand, the local entropy $H(y_t|x_t)$ can be easily estimated. A value of local entropy would be sufficient for our task if we could use it to determine the ordering for values of the potential memory utility $I(y_t, m_t^*|x_t)$. Specifically, for every pair of elements i, j if $H(y_i|x_i) > H(y_j|x_j)$ then $I(y_i, m_i^*|x_i) > I(y_j, m_j^*|x_j)$. This becomes true under one of two possible conditions: (1) $H(y_t|x_t) \gg H(y_t|m_t^*, x_t)$ or (2) $H(y_t|m_t^*, x_t) \sim \epsilon$. The first condition means that given a perfect memory and a model, one could make high-quality predictions for the task at hand. The second condition means that the real distribution of the target variable y_t has a similar amount of noise at each step.

If our task satisfies either one of these two conditions, then valuable information from the past could make the biggest contribution for the elements with the highest values of average "surprise" $H(y_t|x_t)$.

We should note that estimation of $I(y_j, m_j^*|x_j)$ by the local entropy makes MemUP vulnerable to the Noisy-TV problem[21], when the conditions we have described are not satisfied. Sensitivity tests for the Noisy-TV problem are discussed in Section 7.

Thus, instead of optimizing the memory for each future step, we introduce a new objective function that allows us to train it only for steps where the potential gain from the past information is maximal:

$$\max_{\theta} \sum_{k \in \mathcal{U}_t} I(y_k; m_t = g_\theta(x_t, m_{t-1})|x_k)\,, \tag{4}$$

here $\mathcal{U}_t$ denotes the set of top-K indices of steps from t to T with the highest estimated local entropy $H(y_k|x_k)$. The hyperparameter K controls fraction of a sequence to be processed in one gradient step. On the other hand, maximizing the mutual information between the memory m_t and an arbitrary distant event at step $k > t$ allows to learn long-term dependencies.

Optimization. Directly optimizing mutual information can be a challenging task. In our case, we use variational lower bound on mutual information proposed by [22]. Using the lower bound we can maximize mutual information by minimizing Cross-Entropy (CE) between the empirical and model distribution for all selected high entropy events:

$$\min_{\theta,\phi} \sum_{k \in \mathcal{U}_t} \mathbb{E}_{x,y} \left[-\log q_\phi(y_k|m_t^\theta, x_k) \right]\,, \tag{5}$$

where m_t^θ is a shortcut for $m_t = g_\theta(x_t, m_{t-1})$, q_ϕ is a `predictor` network with parameters ϕ, that estimates probability of y_k given x_k and m_t. For a detailed derivation of eq. 5 see Supplementary Materials Section A.

Collection of elements with high uncertainty $\mathcal{U}_t$ requires a special `uncertainty detector` d_ψ model. The main requirement for d_ψ is to produce uncertainty estimates $s_t = d_\psi(x_t)$. In the simplest case, the detector estimates "surprise" $-\log p(y_t|x_t)$ (prediction error), which can be seen as a single point estimate of uncertainty. Other models that directly estimate the uncertainty or a whole distribution [23; 24] can also be used.

3 MemUP for Recurrent Neural Architectures

Recurrent neural net implementation of `MemUP` consists of (1) training of uncertainty detector model d_ψ, (2) selection of elements with highest uncertainty $\mathcal{U}_t$ and (3) training memory g_θ and predictor q_ϕ models. A pseudocode for `RNN MemUP` training is shown in Supplementary Materials E.

Uncertainty detector training. Information maximization reasoning does not limit the choice of solutions for uncertainty detector d_ψ. In our experiments on algorithmic tasks we use a recurrent

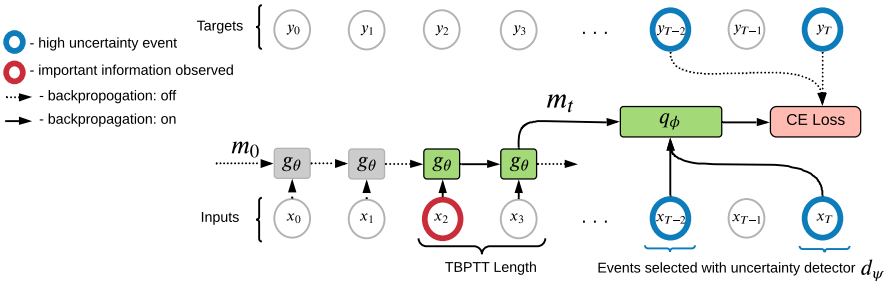

Figure 1: MemUP gradient update. The recurrent memory g_θ processes a sequence. The number of prediction targets $K = 2$ and TBPTT length $r = 2$. Blue circles denote states with the highest uncertainty estimates $d_\psi(y_k|x_k)$. Here x_2 marked by red circle contains information that can help to predict outcomes y_{T-2} and y_T. At the end of RNN TBPTT rollout these states are selected to form a set $U_3 = \{T-2, T\}$. Then, m_3, x_{T-2}, y_{T-2}, x_T, y_T are used to compute CE-Loss according to eq. 5. While states 2 and T could be separated by many millions steps, MemUP allows memory to capture the utility of information from x_2 by propagating gradients only through the sequence of $K + r = 4$ elements at a time.

neural classifier trained with Cross-Entropy Loss. In this case detector's "surprise" $-\log\, d_\psi(y_t|x_t)$ is used as an uncertainty estimate.

In Reinforcement Learning experiments we use distributional algorithm QRDQN [24] trained to predict discounted future returns R_t ($y_t = R_t$ in RL experiments). As QRDQN approximates a whole distribution of the future returns, estimating uncertainty becomes a simple task.

Processing a sequence with the memory network. MemUP allows to learn long term-dependencies without propagating gradients through the whole sequence. To demonstrate that in our experiments we use Truncated BPTT to train memory network g_θ. In the majority of supervised learning experiments we set rollout length r between 10 and 60 steps. In all Reinforcement Learning experiments $r = 1$, i.e. recurrent memory is trained without actually using backpropagation through time. In All Experiments g_θ is a simple LSTM network with additional input encoder.

Selecting states with high uncertainty. Given uncertainty estimates $\{s_i\}_{i=t}^{T}$ for each future element in the sequence, we use softmax sampling $p(k) = e^{s_k/\tau} / \sum_{i=t}^{T}(e^{s_i/\tau})$ without replacement with $\tau = 0.02$. To implement sampling without replacement Gumbel-Max Trick [25] is used. A scheme illustrating a single gradient update is shown in Figure 1.

4 Related work

The problem of learning long-term dependencies has long been studied in both supervised learning [26; 27] and reinforcement learning domains [28]. However, most of the research related to recurrent networks architectures is aimed at solving the problems of exploding and vanishing gradients [2; 29; 7; 30]. The common trend is that many of the existing solutions to deal with vanishing gradients simultaneously increase the cost associated with the calculation of these gradients [29; 7]. For example, Transformer-based [10] architectures require to store $O(N^2)$ of intermediate results for training on sequences of length N. The recently proposed Linformer [31] and BigBird [20] architectures allow training transformers with linear complexity in space. However, most of these new Transformer-based architectures have not yet found their way in the reinforcement learning setting.

In the field of deep reinforcement learning, the dominant approaches to implementation of memory are based on recurrent neural networks in combination with modern reinforcement learning algorithms [5; 32; 33; 34]. There are many studies that propose memory architectures for specific features of the reinforcement learning problems [9; 8]. In the paper by Parisotto and Salakhutdinov [35], a specifically modified DNC [7] architecture learned long-term dependencies of several hundred steps in the setting of 2D and 3D navigation tasks. The downside of this architecture is that the agent needs access to the information about its absolute or relative location in a 2D/3D environment. In another work, Ha and Schmidhuber [36] used a complex procedure with pre-training of memory and state embeddings. Their overall learning procedure is similar to ours. However, their memory/world model is trained to make predictions about the next step observations only, while we train memory to make predictions about events that can be arbitrarily far in the future.

Wayne at al. [17] demonstrated the state of the art results with the MERLIN algorithm. MERLIN's memory kept observation embeddings trained with a complex variational autoencoder in a DNC-like [7] list. The memory module was also trained separately from the policy. Hung et al. [37] followed up on the MERLIN architecture's success and used the soft attention mechanism over all past embeddings to encourage exploratory actions in the POMDP environment. Their experiments showed models learning temporal dependencies with a length of 500+ steps. In contrast to our work, both of these solutions still require processing an arbitrarily long sequence of state embeddings by the policy networks at each step.

The direct attempts in applying Transformer architecture in reinforcement learning setting resulted in conclusion that Transformers are too unstable to work properly in RL [38]. Parisotto et al. [18] were able to overcome the instability of Transformers, showing state of the art result in the memory dependent 3D tasks. Another recent work, by Loynd et al. [39], has also applied Transformers to the POMDP problems, but the WMG architecture relies on availability of factored observations to the agent which imposes additional requirements on the environment. Switching focus from Transformers and complex architectures, Beck et al. [16] proposed AMRL architecture: a simple modification to the LSTM architecture to combat the problem of learning long-term dependencies in environments with a significant amount of observational noise.

5 Supervised Learning Experiments

For evaluation and comparison of our method we use four tasks: Copy [40], Scattered copy (modification of the original Copy task), Add [2] and permuted sequential MNIST (pMNIST) [41]. All these tasks are benchmarks that are used for testing models with long-term memory. In the original Copy task a sequence of length l ($l = 10$ in out experiments) should be copied after a $T - l$ steps ($T \in \{120, 220, 520, 1020\}$ in our experiments), and in Scattered copy task a model has to make predictions in locations inside range $[l, T - 1]$ that are chosen at random. We add this task in order to make detection of high-uncertainty locations harder. For Copy, Scattered copy and Add tasks train and test datasets have sizes 10K and 1K, for pMNIST 60K and 10K. For detailed tasks description see Supplementary Materials Section C.

To encode inputs we use a recurrent encoder architecture E^{rnn} that consists of a single fully-connected layer followed by two LSTM-layers with 128 hidden units and dropout probability 0.1.

Detector d_ψ consists of encoder E^{rnn} followed by three-layered MLP. The detector is trained to predict y_t at each step with CE-loss. Memory model has the same architecture as E^{rnn} and does not share parameters with detector d_ψ.

Predictor is a three-layered MLP that is trained with the loss from eq. 5 given a concatenation of m_t and x_k. Here m_t is an output of the memory model. Instead of building a separate encoder for x_k we share the same recurrent encoder with the detector d_ψ.

In the Add Task we train detector and predictor with MSE loss, which is a special case of CE loss under the model-assumption that target distribution is a Gaussian with unit variance. Both memory and detector are trained with the same Truncated BPTT length. Truncation length equals 10 in all tasks with sequence length < 1000, 20 in tasks with length ≥ 1000 and 30, 60 in pMNIST 784, 3136 respectively.

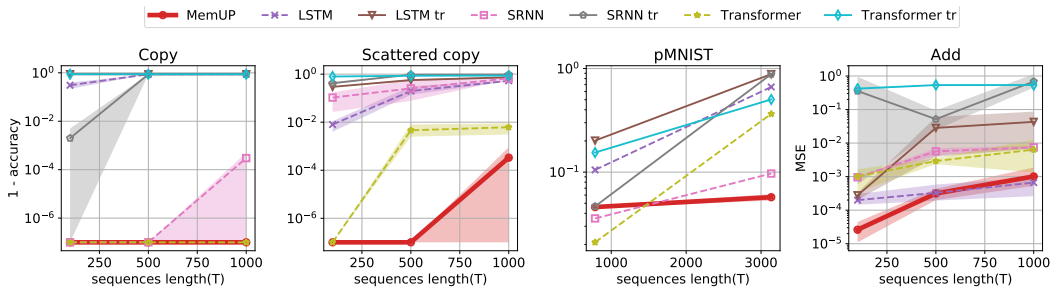

Figure 2: Final performanse on supervised learning tasks. X-Axis shows task scaling, while Y-axis correspond to metric score at the end of training. Metrics: Inverted Accuracy (1. - Accuracy) in Copy, Scattered copy and pMNIST tasks, MSE in Add task. All curves are averaged over 3 runs.

We compare MemUP with LSTM [2] and Transformer [10] models from pytorch library, as well as a recurrent network SRNN [30] which is designed for training on long sequences and has reduced saturation of gradients. LSTM* and SRNN* have the same architecture as E^{rnn} with 3 layered MLP on top. In the case of SRNN baselines, we have replaced all LSTM layers with SRNN layers. The Transformer model consists of a single FC layer followed by 4 self-attention layers and the same 3 layered MLP on top. Baselines were trained in two modes. In the first mode (columns with Transformer, SRNN, LSTM) all elements of input sequence of length T were fed into a model simultaneously, so gradients propagate through the whole sequence during the training. In the second mode (Transformer Tr., SRNN Tr., LSTM Tr.) BPTT rollout for recurrent architectures and attention window for Transformer were truncated to process the same number of steps as MemUP (r + K).

Results in Figure 2 show that only MemUP allows to achieve a good quality of prediction in all tasks. In truncated BPTT setting (solid lines) MemUP is significantly better than LSTM tr, SRNN tr or Transformer tr due to MemUP ability to learn temporal dependencies longer that the length of BPTT rollout. When full information about the sequence is available to LSTM, SRNN and Transformer (dashed lines) but not MemUP, the latter still shows the best results for Add as well as pMnist, and has better overall performance over all tasks. Moreover, when compared with Full BPTT baselines, MemUP requires to store much less intermediate computations (10-50 times) due to usage of short Truncated BPTT rollouts.

6 Reinforcement Learning Experiments

Performance of the MemUP algorithm was studied in two reinforcement learning tasks.

The first task is classic T-Maze environment where the main difficulty is a long-term dependency between a hint at the starting position and location of a reward at the exit (see Figure 3.1). The main advantage of this problem is that it allows to test the agent's memory mechanisms in isolation because of primitive policy and the observation space. We base our experiments on the noisy version introduced by [16](with length increased by 10 times).

The second task is the color dependent object collection in the Doom environment [42; 43] which requires long-term memory in combination with a more complex policy and rich observation space. In the task introduced by [44], the agent is placed in a room filled with acid (see Figure 3.2.) and constantly loses health. Objects of green and red colors are scattered throughout the environment. Object of one color replenish the agent's health and give a +1 reward, while others take away health and give a -1 reward. The correspondence between effects and items' colors is determined randomly at the beginning of each episode. This information is conveyed by a column whose color matches the color of health replenishing items. In our version the column disappear after first 45 steps (see Figure 3.3). The episode ends when the agent's health drops to zero or after 1050 time steps.

The most straightforward extension of the MemUP algorithm to a reinforcement learning problem is to use memory state m_t as an additional input to an RL agent. There are several possible solutions to combine MemUp's training with agent's policy training: (1) pre train MemUP and then train an agent, (2) use alternating memory and agent training phases, train memory and agent in parallel, (3) combining the MemUP and RL agent into a single into an end-to-end network. In these experiments, we test the simplest version: train MemUP on trajectories generated with a frozen policy, then train an agent with a frozen MemUP memory.

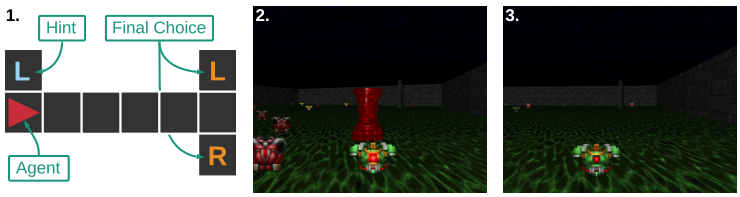

Figure 3: Environments for testing long-term time dependencies. **1.** *T-Maze environment.* The agent should reach the T-shaped junction and choose one of the arms (L or R). A hint about what arm to choose is provided at the very beginning. **2.** and **3.** *Vizdoom-Two-Colors environment*. The agent is in the room and constantly loses health. To replenish his health and recieve a reward the agent needs to collect items of the same color as the column. The color of the column is chosen randomly from red or green options at the beginning of the episode. After 45 steps from the beginning the column disappears. The episode lasts for a maximum of 1050 steps.

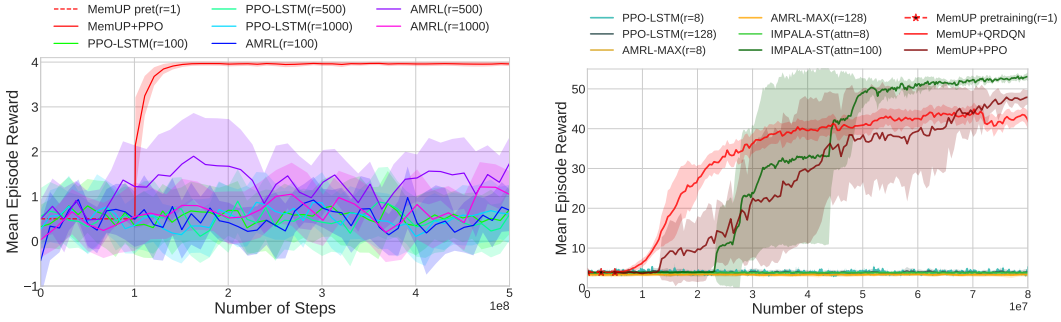

(a) (b)

Figure 4: **a)** Learning curves for `MemUP` and baselines in T-Maze-LNR-100. **b)** Learning plots in the Vizdoom-Two-Colors environment for all agents. All curves are averaged over 3 runs. Translucent areas around each curve show the standard deviation computed over all runs. The memory pre-training phase of `MemUP` agent is marked with the dashed line.

In RL experiments inputs x_t are formed from observations o_t and previous action a_{t-1}. Observations are encoded using convolutional networks or fully connected networks. We use discounted future returns $R_t = \sum_{i=t}^{T} \gamma^{i-t} r_i$. as targets y_t. For more details see Supplelemtary Materials Section F. In the experiments, `MemUP` is compared with the following baselines.

PPO-LSTM is a recurrent version of `PPO` [45] with a single `LSTM`-layer [2]. Our experiments were performed with the `PPO-LSTM` implementation from the RLPyt library [46].

IMPALA-ST is an IMPALA [47] agent with Transformer-based [10] architecture. The Stabilized Transformer architecture was presented by Parisotto et al. [18]. For experiments, we use the only available open implementation of this algorithm [48].

AMRL is proposed by [16] for the Noisy T-Maze Task. In the original study `AMRL` was compared with many different memory architectures including `DNC`, multi-layered `LSTM` and outperformed all of them on T-Maze-LN-100 task as well as other memory experiments. `AMRL` is similar to `PPO-LSTM` baseline, but extends `LSTM` with `AMRL` Layer.

In all experiments we use the same encoder networks for observation embeddings for all baselines and `MemUP`. Additionally the size of an `LSTM`-layer is the same for the memory module g_θ, `PPO-LSTM` and `PPO-AMRL` baselines.

Results for the T-Maze-LNR-1000 environment are presented in Figure 4a. The designation T-Maze-LNR-1000 means that the maze length (and corresponding tempral dependency) in each episode can be from 1000 to 1009 steps. In T-Maze-LNR-1000 environment (Figure 4a),we tested several rollout lengths for RNN-based baselines. We could not successfully train a single `PPO-LSTM` or `AMRL` agent even with TBPTT rollout of $r = 1000$[1].

During MemUP pre-training phase, episodes are generated with a random policy. We used the discounted future reward with $\gamma = 0$ as a prediction target. A policy training for the `MemUP` agent starts after the memory pre-training phase is completed. The time spent in the pre-training phase is marked with a dashed line in all figures. As shown in Figure4a the agent learns almost instantly, given a memory pre-trained with our method. For the `MemUP` agent we train all components including memory module and uncertainty detector for each individual run from scratch. In this experiment MemUP learned a temporal-dependency of 1000 steps by propogating gradients only through 2 steps for each gradient update.

Results for the Vizdoom-Two-Colors environment are shown in Figure 4b. IMPALA-ST agent was trained in two configurations with the attention span of 100 and 8 steps. IMPALA-ST(attn=8) with a short attention span could not learn any reasonable policy, as the knowledge about only the last 8 steps is not enough to store relevant information about the color of the healing items. On the other hand, after 45 million training steps, IMPALA-ST(attn=100) was able to solve the problem and

[1] We tested our AMRL implementation on T-Maze-LNR-100 with full rollout and get results aligning with the original paper. But the algorithm did not withstand 10 times increase in episode length and truncation constants

reliably survive in the environment for 1050 steps. Both PPO-LSTM and AMRL baselines with TBPTT rollouts of 128 and 8 steps could not learn to survive longer than a random agent. While the rollout of 128 steps is potentially enough to remember the color of the column or a previously collected items, PPO-LSTM and AMRL agents did not utilize long-term information from the past. They stuck with the "collect all items" policy.

For MemUP we used the discounted future reward with $\gamma = 0.8$ as a prediction target. The memory module g_θ was trained with the rollout length of 1 step. In Vizdoom-Two-Colors we set number of targets $K = 3$. Thus, to train the memory g_θ we used only 4 separate observations per episode for one gradient update. In ViZDoom, we trained PPO (MemUP+PPO) and QRDQN (MemUP+QRDQN) agents with MemUP pretraining.

Both MemUP+PPO and MemUP+QRDQN learn to survive in the environement for full 1050 steps by collecting healing items. Though the final performance of IMPALA-ST(attn = 100) is slightly better. We speculate that this is due to the fact that the pre-trained memory **M** has learned to store mostly information about the color of the column, while IMPALA-ST (attn=100) has slightly better spatial awareness. On the other hand, our algorithm used significantly fewer resources in terms of the size of processed sequences during training/pre-training and evaluation. As well, MemUP significantly outperforms baselines with comparable resource requirements, like IMPALA-ST (attn=8), PPO-LSTM(r=8) and AMRL-MAX(r=8).

7 Sensitivity to the Noisy-TV problem

In section 2 we discussed conditions on true uncertainty of targets y_t given a perfect memory m_t^*. The goal was to set conditions under which we can make a good guess about relative values of potential memory utility $I(y_i|m_i^*, c_i)$ by estimating local entropy $H(y_i|c_i)$. However, this may not be the case. For example, some targets y_i may be completely unpredictable and no information from the past can reduce the entropy for the distribution of y_i: $H(y_t|x_t) \sim H(y_t|m_t^*, x_t)$. In this case local entropy is large, while potential utility of memory for y_t predictions may be close to zero. Training memory to predict such targets is pointless, and, at the same time, it may prevent MemUP from focusing on events for which training long-term memory is essential. The problem of over-emphasis on unpredictable surprising events is also called Noisy-TV problem[21]. Noisy-TV problem is often encountered by algorithms using curiocity-based exploration in reinforcement learning[49].

To test the robustness of the algorithm to the Noisy-TV problem, we use modified version T-maze-LNR-100 environment.

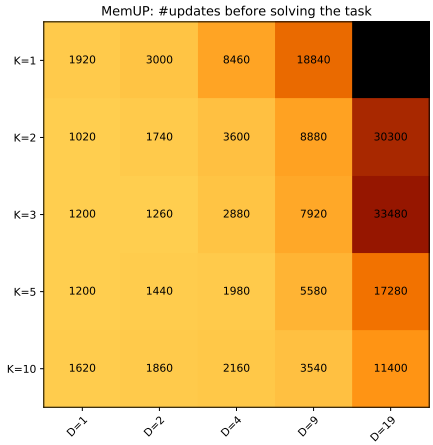

Figure 5: Each cell shows the mean number of updates required for the memory to correctly predict (achieving 0.1 RMSE) the reward after the decision point connected with the temporal dependency. The x-axis shows the number of unpredictable decision points in the environment. The Y-axis shows number of events predicted by MemUP. Darker colors means slower solution.

Noisy T-maze With Distractors. In this version, agent can receive +4 or -3 rewards in D+1 decision points. Their location in the corridor is chosen randomly at a distance of at least 50 steps from the hint. All of them can be detected when $o_t[1]! = 0$. element in observation vector o_t. Decision points are distinguishable from each other. Each decision point has a unique value of $o_t[1]$. Only in one decision point, the next reward depends on the agent's action and the value of the hint ($o_t[1] = 1$. in this case). In other D decision points, the next rewards are completely random (chosen with a probability of 0.5). Thus, Noisy T-maze With Distractors generate sequences with one long-term dependency and D events acting as Noisy-TV distractors.

We conducted 25 experiments with 6 independent runs in each (150 runs in total). The differences in the experiments come down to two parameters: number of distracting events D (X-axis in Fig. 5) and number of events sampled for prediction K (Y-axis in Fig. 5). More distracting events make the task of learning the temporal dependency harder. Increasing the size of $\mathcal{U}_t$ set(by changing K) improves the chance of selecting a future event that is a part of temporal dependency. On the other hand, with an increase in $\mathcal{U}_t$, the resource cost for memory training also increase. The TBPTT rollout length in these experiments is 1.

Results in Fig. 5 show that with an increase in the number of distractions, memory learning speed decreases. However, increasing the set $\mathcal{U}_t$ allows to alleviate this problem. For the experiment (D=19, K=1), none of the runs solved the task in the first 45000 updates. In experiments (D=19, K=2) and (D=19, K=3) only 4 out of 5 runs solved the task. The results show that the agent can learn the time dependency fairly quickly even if the uncertainty detector select noisy events 80% of the time.

8 Conclusions

In this paper, we proposed a new method for training long-term memory. The main idea is to train a memory network to predict future outcomes of high uncertainty and skip all others. Predicting a small number of arbitrarily distant future outcomes substantially saves computational resources required to backpropagate gradients. At the same time, the emphasis on high uncertainty outcomes allows not to miss long-term dependencies in the task.

Experimental results show that our training algorithm allows to learn temporal dependencies significantly longer than the number of steps processed for a single gradient update. None of the baseline architectures trained in a classical way demonstrate such ability. All these baselines with comparable results require to store at least 50 times more intermediate calculations for the Add, Copy, Scattered Copy, pMNIST tasks as well as at least 500 times more for the T-maze, and 625 times more for the Vizdoom-Two-Colors (Stabilized Transformer baseline has quadratic complexity while processing 25 times more steps).

Even using fewer resources, MemUP outperforms non-truncated baselines that has simultaneous access to all elements of a sequence (including Transformer) on Add and pMNIST(31) tasks. On Scattered Copy, T-maze, and Vizdoom-Two-Colors tasks, MemUP is better than all recurrent baselines. Another advantage is that MemUP can be combined with any recurrent architecture and applied both to supervised as well reinforcement learning settings.

We believe that the MemUP algorithm applied to recurrent networks has shown promising results, and the main idea demonstrates exciting avenues for future research.

References

[1] Paul J Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.

[2] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.

[3] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv:1410.5401*, 2014.

[4] Jason Weston, Sumit Chopra, and Antoine Bordes. Memory networks. *arXiv:1410.3916*, 2014.

[5] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep

reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016.

[6] Junhyuk Oh, Valliappa Chockalingam, Satinder P. Singh, and Honglak Lee. Control of memory, active perception, and action in minecraft. In Maria-Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 2790–2799. JMLR.org, 2016. URL `http://proceedings.mlr.press/v48/oh16.html`.

[7] Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538 (7626):471, 2016.

[8] Adam Santoro, Ryan Faulkner, David Raposo, Jack W. Rae, Mike Chrzanowski, Theophane Weber, Daan Wierstra, Oriol Vinyals, Razvan Pascanu, and Timothy P. Lillicrap. Relational recurrent neural networks. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 7310–7321, 2018. URL `https://proceedings.neurips.cc/paper/2018/hash/e2eabaf96372e20a9e3d4b5f83723a61-Abstract.html`.

[9] Steven Kapturowski, Georg Ostrovski, John Quan, Remi Munos, and Will Dabney. Recurrent experience replay in distributed reinforcement learning. *7th International Conference on Learning Representations (ICLR)*, 2019.

[10] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008, 2017. URL `http://papers.nips.cc/paper/7181-attention-is-all-you-need`.

[11] Mikhail S Burtsev, Yuri Kuratov, Anton Peganov, and Grigory V Sapunov. Memory transformer. *arXiv preprint arXiv:2006.11527*, 2020.

[12] Jie Lei, Liwei Wang, Yelong Shen, Dong Yu, Tamara Berg, and Mohit Bansal. MART: Memory-augmented recurrent transformer for coherent video paragraph captioning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2603–2614, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.233. URL `https://aclanthology.org/2020.acl-main.233`.

[13] Ankit Gupta and Jonathan Berant. Gmat: Global memory augmentation for transformers. *arXiv preprint arXiv:2006.03274*, 2020.

[14] Qingyang Wu, Zhenzhong Lan, Jing Gu, and Zhou Yu. Memformer: The memory-augmented transformer. *arXiv preprint arXiv:2010.06891*, 2020.

[15] Angela Fan, Thibaut Lavril, Edouard Grave, Armand Joulin, and Sainbayar Sukhbaatar. Addressing some limitations of transformers with feedback memory. *arXiv preprint arXiv:2002.09402*, 2020.

[16] Jacob Beck, Kamil Ciosek, Sam Devlin, Sebastian Tschiatschek, Cheng Zhang, and Katja Hofmann. AMRL: aggregated memory for reinforcement learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL `https://openreview.net/forum?id=Bkl7bREtDr`.

[17] Greg Wayne, Chia-Chun Hung, David Amos, Mehdi Mirza, Arun Ahuja, Agnieszka Grabska-Barwinska, Jack W. Rae, Piotr Mirowski, Joel Z. Leibo, Adam Santoro, Mevlana Gemici, Malcolm Reynolds, Tim Harley, Josh Abramson, Shakir Mohamed, Danilo Jimenez Rezende,

David Saxton, Adam Cain, Chloe Hillier, David Silver, Koray Kavukcuoglu, Matthew Botvinick, Demis Hassabis, and Timothy P. Lillicrap. Unsupervised predictive memory in a goal-directed agent. *CoRR*, abs/1803.10760, 2018. URL `http://arxiv.org/abs/1803.10760`.

[18] Emilio Parisotto, H. Francis Song, Jack W. Rae, Razvan Pascanu, Çaglar Gülçehre, Siddhant M. Jayakumar, Max Jaderberg, Raphael Lopez Kaufman, Aidan Clark, Seb Noury, Matthew M. Botvinick, Nicolas Heess, and Raia Hadsell. Stabilizing transformers for reinforcement learning. *CoRR*, abs/1910.06764, 2019. URL `http://arxiv.org/abs/1910.06764`.

[19] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.

[20] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. Big bird: Transformers for longer sequences. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL `https://proceedings.neurips.cc/paper/2020/hash/c8512d142a2d849725f31a9a7a361ab9-Abstract.html`.

[21] Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.

[22] David Barber Felix Agakov. The im algorithm: a variational approach to information maximization. *Advances in neural information processing systems*, 16(320):201, 2004.

[23] Evgenii Tsymbalov, K. Fedyanin, and Maxim Panov. Dropout strikes back: Improved uncertainty estimation via diversity sampled implicit ensembles. *ArXiv*, abs/2003.03274, 2020.

[24] Will Dabney, Mark Rowland, Marc G Bellemare, and Rémi Munos. Distributional reinforcement learning with quantile regression. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[25] Iris AM Huijben, Wouter Kool, Max B Paulus, and Ruud JG van Sloun. A review of the gumbel-max trick and its extensions for discrete stochasticity in machine learning. *arXiv preprint arXiv:2110.01515*, 2021.

[26] Herbert Jaeger. *Tutorial on training recurrent neural networks, covering BPPT, RTRL, EKF and the" echo state network" approach*, volume 5. GMD-Forschungszentrum Informationstechnik Bonn, 2002.

[27] Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, Jürgen Schmidhuber, et al. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. *A field guide to dynamical recurrent neural networks. IEEE Press*, 2001.

[28] Bram Bakker. Reinforcement learning with long short-term memory. In Thomas G. Dietterich, Suzanna Becker, and Zoubin Ghahramani, editors, *Advances in Neural Information Processing Systems 14 [Neural Information Processing Systems: Natural and Synthetic, NIPS 2001, December 3-8, 2001, Vancouver, British Columbia, Canada]*, pages 1475–1482. MIT Press, 2001.

[29] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *CoRR*, abs/1410.5401, 2014. URL `http://arxiv.org/abs/1410.5401`.

[30] Michael Rotman and Lior Wolf. Shuffling recurrent neural networks. *CoRR*, abs/2007.07324, 2020. URL `https://arxiv.org/abs/2007.07324`.

[31] Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *CoRR*, abs/2006.04768, 2020. URL `https://arxiv.org/abs/2006.04768`.

[32] Matthew Hausknecht and Peter Stone. Deep recurrent q-learning for partially observable mdps. *CoRR, abs/1507.06527*, 7(1), 2015.

[33] Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, and Koray Kavukcuoglu. Feudal networks for hierarchical reinforcement learning. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 3540–3549. PMLR, 2017. URL `http://proceedings.mlr.press/v70/vezhnevets17a.html`.

[34] Bowen Baker, Ingmar Kanitscheider, Todor M. Markov, Yi Wu, Glenn Powell, Bob McGrew, and Igor Mordatch. Emergent tool use from multi-agent autocurricula. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL `https://openreview.net/forum?id=SkxpxJBKwS`.

[35] Emilio Parisotto and Ruslan Salakhutdinov. Neural map: Structured memory for deep reinforcement learning. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL `https://openreview.net/forum?id=Bk9zbyZCZ`.

[36] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 2455–2467, 2018. URL `https://proceedings.neurips.cc/paper/2018/hash/2de5d16682c3c35007e4e92982f1a2ba-Abstract.html`.

[37] Chia-Chun Hung, Timothy Lillicrap, Josh Abramson, Yan Wu, Mehdi Mirza, Federico Carnevale, Arun Ahuja, and Greg Wayne. Optimizing agent behavior over long time scales by transporting value. *Nature communications*, 10(1):1–12, 2019.

[38] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL `https://openreview.net/forum?id=B1DmUzWAW`.

[39] Ricky Loynd, Roland Fernandez, Asli Çelikyilmaz, Adith Swaminathan, and Matthew J. Hausknecht. Working memory graphs. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 6404–6414. PMLR, 2020. URL `http://proceedings.mlr.press/v119/loynd20a.html`.

[40] Martin Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. *International Conference on Machine Learning*, pages 1120–1128, 2016.

[41] Quoc V. Le, Navdeep Jaitly, and Geoffrey E. Hinton. A simple way to initialize recurrent networks of rectified linear units. *ArXiv*, abs/1504.00941, 2015.

[42] Michał Kempka, Marek Wydmuch, Grzegorz Runc, Jakub Toczek, and Wojciech Jaśkowski. Vizdoom: A doom-based ai research platform for visual reinforcement learning. In *2016 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 1–8. IEEE, 2016.

[43] Guillaume Lample and Devendra Singh Chaplot. Playing FPS games with deep reinforcement learning. In Satinder P. Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 2140–2146. AAAI Press, 2017. URL `http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14456`.

[44] Edward Beeching, Christian Wolf, Jilles Dibangoye, and Olivier Simonin. Deep reinforcement learning on a budget: 3d control and reasoning without a supercomputer. *CoRR*, abs/1904.01806, 2019. URL `http://arxiv.org/abs/1904.01806`.

[45] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[46] Adam Stooke and Pieter Abbeel. rlpyt: A research code base for deep reinforcement learning in pytorch. *arXiv preprint arXiv:1909.01500*, 2019.

[47] Lasse Espeholt, Hubert Soyer, Rémi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. IMPALA: scalable distributed deep-rl with importance weighted actor-learner architectures. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 1406–1415. PMLR, 2018. URL `http://proceedings.mlr.press/v80/espeholt18a.html`.

[48] Shakti Kumar, Jerrod Parker, and Panteha Naderian. Adaptive transformers in rl. *arXiv preprint arXiv:2004.03761*, 2020.

[49] Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A Efros. Large-scale study of curiosity-driven learning. *arXiv preprint arXiv:1808.04355*, 2018.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] Our experimental results and ablation study align with our claims. See Section 6 and Ablation Study in Supplementary materials
   (b) Did you describe the limitations of your work? [Yes] See Section 2 and Noisy-TV experiments 7
   (c) Did you discuss any potential negative societal impacts of your work? [No] No, Negative consequences of practical application of our idea (that we could think of), can be coupled with any improvement in DL/RL methods rather than specifically with our proposal.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] We describe our assumptions in Section 2 (L.83-96)
   (b) Did you include complete proofs of all theoretical results? [Yes] See Supplementary Materials. Section A.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] We are working towards code publication, but not yet ready to do so.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We have included these details in *Supplementary Materials*
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] All graphs include mean and standard deviation computed over multiple runs with different random seed.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See *Supplementary Materials*.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We use RLPyt library and open implementation of Stabilized Transformer and SRNN for our baselines and cite authors in sections 5 and 6
   (b) Did you mention the license of the assets? [No]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No]