
Re-understanding Finite-State Representations of Recurrent Policy Networks

Mohamad H. Danesh¹ Anurag Koul¹ Alan Fern¹ Saeed Khorram¹

Abstract

We introduce an approach for understanding control policies represented as recurrent neural networks. Recent work has approached this problem by transforming such recurrent policy networks into finite-state machines (FSM) and then analyzing the equivalent minimized FSM. While this led to interesting insights, the minimization process can obscure a deeper understanding of a machine’s operation by merging states that are semantically distinct. To address this issue, we introduce an analysis approach that starts with an unminimized FSM and applies more-interpretable reductions that preserve the key decision points of the policy. We also contribute an attention tool to attain a deeper understanding of the role of observations in the decisions. Our case studies on 7 Atari games and 3 control benchmarks demonstrate that the approach can reveal insights that have not been previously noticed.

1. Introduction

What roles do observations and memory play in the decision making of deep policy networks for complex control tasks? While such networks have yielded state-of-the-art performance in reinforcement and imitation learning (e.g. (Mnih et al., 2013; Hausknecht & Stone, 2015; Wang et al., 2015; Van Hasselt et al., 2016; Ho & Ermon, 2016)), there are limited tools and approaches for giving insight into this question. This is particularly the case for policies represented as recurrent networks, e.g. LSTMs and GRUs (Hochreiter & Schmidhuber, 1997; Chung et al., 2014; Cho et al., 2014), which condition on high-dimensional memory vectors with no preconceived semantics. Prior works have attempted to gain insight via attention maps over the network input, e.g. (Greydanus et al., 2018; Iyer et al., 2018; Gupta et al., 2020; Atrey et al., 2020), however, this ultimately involves subjective

human interpretation of the underlying “strategic role” of the attended-to elements. In this paper, we develop an analysis approach that reveals such human interpretations can sometimes be highly misleading.

Our work builds on a recent approach for understanding recurrent policy networks by quantizing memory and observations within an RNN (Koul et al., 2019). A quantized RNN is a type of FSM, known as a Moore Machine (MM), which can be visualized and analyzed. In particular, the originally large MMs were algorithmically minimized, resulting in small machines for a variety of domains, which yielded high-level insights. For example, the minimal MM for Atari Pong showed that memory was not actually needed (i.e. a state-action mapping), while for Atari Bowling the MM was an open-loop strategy that ignored observations.

While analyzing minimal MMs allows for determining global properties, such as memory/observation use or not, we have found it difficult to gain more in-depth insight from minimal MMs. To see this, Figure 4 (a) and (b) show minimal MMs for Pong and Bowling and Figure 5 shows a set of frames/observations for each MM that all map to a single state in the corresponding MM. For a human, the frames are semantically distinct, whereas the minimization process was able to merge them all into a single state. This is due to the minimization focusing on maintaining logical equivalence to the unminimized MM policy, rather than maintaining any semantic meaning of states. Thus, we have found it very difficult to understand the strategic role of different states in such MMs to gain deeper insight into their decision logic. This experience has led us to the view that minimal MMs are *unlikely* to be a good starting point for gaining a deeper understanding of recurrent policies.

Our main contribution is to develop a new approach for RL researchers and advanced practitioners to analyze MM policy representations. We demonstrate that this approach can yield significantly more insight into the decision making of recurrent policies. Rather than start with a minimal MM after quantization, we start with the unminimized MM, which is often quite large, and apply “interpretable reductions” in order to reduce the visual complexity. The reductions are intentionally simple in order to preserve the inherent decision structure in the machine while effective at compressing the visual representation. For example, one reduction operation

¹Department of EECS, Oregon State University, Corvallis, OR, USA. Correspondence to: Mohamad H. Danesh <daneshm@oregonstate.edu>.




Figure 1. Overall approach for Pong. a) *QBN Insertion* (Section 3) discretizes observations and memory of the original RNN (top). b) Resulting *MM* with finite states and observations. c) *Interpretable reductions* (Section 4.1) are applied to the *MM*, yielding a single decision point s_{197} conditioning on observations. *Differential attention* (Section 4.2) is used to understand decisions in terms of observations. d) *Functional pruning* (Section 4.3) removes unnecessary branches, leaving an open-loop policy.

replaces a fixed series of non-branching memory states (i.e. a macro) with a single abstract transition. The result is a simplified machine with a fixed set of “branching states” where the flow of the machine depends on observations.

In order to help understand the decisions made at branching states, we further develop a new differential attention tool, aiming to identify parts of observations that are most important for selecting one branch over another. Using the tool, we found that often, especially in Atari, the attention was unintuitive from a strategic point of view, which led us to question whether observations were used for strategically important reasons, or whether they were “arbitrary” branches. This led us to consider the “functional pruning” reduction to assess that issue. In particular, this operation eliminates all but one branch at a decision point to test whether the observation-conditioned decision among multiple branches was strategically important, or just an artifact of learning a non-compact policy.

We explore this approach by studying of 7 deterministic Atari games and 3 continuous control environments. For the control tasks, the approach identifies interpretable machines, whose decision points are understandable and strategically meaningful. In contrast, for Atari, the analysis reveals new and surprising insights. Prior works have attempted to understand the most salient pixels for policy decisions in Atari games (Greydanus et al., 2018; Iyer et al., 2018; Gupta et al., 2020), however, little insight was gained into how that information was used. Our approach reveals that for the Atari policies, observations were not used for “strategically useful” purposes. In particular, at each state that branches on observations, it was possible to remove all branches, except for one, resulting in an open-loop policy that ignores observations, while maintaining performance. We call such policies, *pruned open-loop policies*, and observe that all of our Atari policies are of this type. Finally, we identify limits of the current approach as problems become more stochastic, which suggests important avenues of future work.

2. Related work

Attention Maps. Attention maps have been used as a tool to identify the most relevant parts of the input with respect to the agent’s decision (Greydanus et al., 2018; Iyer et al., 2018; Gupta et al., 2020; Atrey et al., 2020). Perturbation-based attention methods have been investigated to gain insight into learned Atari policies (Greydanus et al., 2018; Iyer et al., 2018; Gupta et al., 2020), but have been criticized for relying on the application of networks to potentially non-sensical perturbed states (Atrey et al., 2020). This has been partly addressed by using more advanced counterfactual state generation (Olson et al., 2021). In general, attention approaches produce a “local explanation” for the decisions made at specific states, which is in contrast to “global explanations” we primarily focus on in this work. Attention methods are also limited in the type of insights they can provide. For example, while they are applied to recurrent policies, they provide no insight into how memory is used and the strategic role of salient inputs. As we will show, our results suggest that a human’s intuition about the role of salient inputs can be highly misleading.

Hybrid Architectures. There have been efforts to induce explanations components in the architecture to make agents implicitly explainable (Yang et al., 2018; Mott et al., 2019). For example, soft attention modules have been used in a recurrent architecture to gain insight into the “attention” of an agent (Mott et al., 2019). While attention has been used in several context for explanation, the soundness of this approach is not clear. In particular, the attention weights are computed from the raw observations and memory, which is ignored in the explanation process. Thus, it is difficult to determine whether the attention patterns carry key strategic information based on other parts of the raw observation, or whether, indeed, it is only the information highlighted by the attention that is key to decisions. Investigating ways to distinguish these two cases is interesting future work.

Extracting Finite-State Machine. There has been significant prior work on extracting FSMs from RNNs in the context of formal language learning of fixed finite alpha-




Figure 2. a) An example of an MM with 7 memory states and 9 observations. The initial memory state is S_0 , and based on the input observation, it goes to the next memory state until it reaches to the memory state S_5 , then it loops back to S_2 . In this example, MM has a decision point at S_2 . b) Interpretable reductions applied to the MM.

bets (Tiño et al., 1998; Cechin et al., 2003; Weiss et al., 2017; 2019). Only recently have there been attempts to learn FSMs for complex RL problems (Koul et al., 2019) where machine minimization was used to extract high-level insights. Our work builds on (Koul et al., 2019) and aims to significantly advance the depth of understanding that such methods can provide.

3. Recurrent Networks to Moore Machines

In this section we review *recurrent policy networks (RPNs)* and how they can be converted to *MMs* as illustrated in Figure 1a. This work is agnostic about how a policy is learned, which, for example, could be via RL, imitation learning, or other training approaches.

Recurrent Policy Networks (RPNs). An RPN is an RNN policy that, at each time step, is given an observation o_t and outputs an action a_t . As illustrated in Figure 1a top, during execution, an RPN maintains a continuous-valued hidden memory state h_t , which is updated on each transition and influences the action choice a_t . Specifically, given current observation o_t and current state h_t , an RPN outputs an action $a_t = \pi(h_t)$ where π may be a feed-forward network. Then, it updates the state according to $h_{t+1} = \delta(f_t, h_t)$, where f_t is a set of features extracted from o_t , for example, using a CNN when observations are images. δ is the transition function, which is often implemented via different types of gating networks such as LSTMs or GRUs.

Moore Machines. Understanding action choices of an RPN is complicated by the memory’s high-dimensionality and lack of predefined semantics. Recent work (Koul et al., 2019) has attempted to address this issue by transforming RPNs to MMs, which allows for visualization and analysis of a finite system. An MM is a finite-state machine defined by a finite set of labeled hidden states H , a distinguished initial state $h_0 \in H$, a finite set of observation symbols O , and a transition function $\Delta : H \times O \rightarrow H$, which returns the next state $h_{t+1} = \Delta(h_t, o_t)$, given the current state and observation symbol. The label associated with each state corresponds to an action. An MM policy initializes the state to h_0 and then updates the state as observations arrive and

outputs the action associated with each state.

Quantized Bottleneck Insertion (QBN Insertion). We now overview the approach of *Quantized Bottleneck Insertion* (Koul et al., 2019) for transforming an RPN to an MM policy, which is illustrated in Figure 1a bottom. Full details of this approach can be found in the original paper and are not critical to the contributions of this paper. The key components of the approach are *Quantized Bottleneck Networks (QBNs)*, which are simply auto-encoders, for which the encoder produces a quantized latent representation. In this work, bottleneck representation is composed of discrete units with output values in $\{-1, 0, 1\}$. Given a trained RPN, a representative set of RPN trajectories is produced and the resulting sets of hidden states $\{h_t\}$ and observation features $\{f_t\}$ are collected. Next, a hidden-state QBN Q_h and observation QBN Q_o are trained to minimize reconstruction error on the data sets. The encoders of Q_h and Q_o can be viewed as discretizing the state and observation spaces. The trained QBNs are then inserted into the RPN in place of the “wires” that propagate the continuous memory vector h_t and observation features f_t (see Figure 1a). This creates a discrete representation of the hidden states \hat{h}_t and observation features \hat{f}_t in the RPN. In practice, QBNs have reconstruction errors, which may impact the RPN performance. Supervised fine-tuning of the discretized RPN can be used to improve performance via imitation learning with respect to the original RPN.

Trajectories of the discretized RPN are then run to collect a representative transition set $\{(\hat{h}_t, a_t, \hat{f}_t, \hat{h}_{t+1})\}$, indicating that action a_t was taken in discrete state \hat{h}_t and a transition to \hat{h}_{t+1} was observed when the discrete observation was \hat{f}_t . The MM is constructed by creating a transition graph edge for each data tuple. The key parameters relevant to this paper are sizes of the bottlenecks of Q_h and Q_o , denoted N_h and N_o respectively. Larger values give more potential to produce finer grained quantization and in turn more states.

4. Analyzing non-minimal Moore Machines

RPNs learned for complex problems can result in MMs with large numbers of discrete states and observations. In order




Figure 3. Differential Attention Pipeline. The pair of images under comparison result in discrete representations that differ on the highlighted features. We produce an attention map for each of those features using the Integrated Gradient (IG) approach and average the magnitude of the maps for an overall differential attention map.

to aid understanding, previous work (Koul et al., 2019) used a standard minimization algorithm (Paull & Unger, 1959) to produce equivalent minimal state MMs. As described in Section 1, the minimal machines allowed for interesting global insights into the general use of states and observations. However, as also described in Section 1, we have found that minimal MMs obscure the decision making behavior, since they compress the original MM structure with no regard for interpretability. For example, the observations mapped to single states often appear to be semantically very different (Figure 5), making it difficult to understand the role of states and how observations influence their choices. Thus, in this paper, we start with the unminimized MMs with the aim to gain a more detailed understanding. Below, we describe the steps of our analysis approach.

4.1. Interpretable Reductions

Figure 2a shows an example MM, which illustrates the main structures we observed in the learned MMs we analyzed. These structures provide several opportunities for simple *interpretable reductions*, which simplify the visualization of an MM without obscuring decision structure. *The reductions are very simple by design and one of our contributions is to notice that this simple set can be effective for interpreting very large MMs.* Note that some of the following reductions may hide some steps, for example long macros with no decision points, to make the visualization manageable. But such simplifications do not essentially influence the completeness of the decision process in MM. The first stage of our approach applies the following interpretable reductions.

Sequence Reduction. It is common to see long sequences of states with a single observation between consecutive states (e.g. S_0 to S_2 in Figure 2a). These sequences are open open-loop macros that simply execute a fixed sequence of actions whenever encountered. Examples are S_0 and S_1 in the sequence of state-transitions from state S_0 to S_2 that have no branching states. We reduce these sequences to a single “macro arc” represented as a dotted line with a number indicating its length.

Loop Unrolling Reduction. There are many loops attached

to sequences that are only traversed once when the sequence is visited. These loops increase the visual complexity of the MMs and make it appear as if there is a decision that controls loop exit, when there is not. Thus, we simply unroll such loops before applying sequence reduction. To identify such loops, once the MM is generated, we run the policy once more and count the number of times each node is visited; thus we exactly know how many times a loop is covered.

Parallel Reduction. There are often multiple transition arcs between two states with different observation labels (e.g. between S_5 and S_6 in Figure 2a). This can also occur for self-transitions. We merge these arcs into an abstract arc labeled by the number of observations.

Startup and Termination Reduction. In some MMs, there is a period of state transitions that corresponds to a warm-up/termination period in the environment where actions have no impact. If desired, these parts of an MM can be replaced with a macro arc represented by two parallel lines with the number of transitions in that period. This reduction requires minimal human annotation of the steps where episodes “meaningfully” begin and end, which is usually straightforward. For example, in Atari games, there is usually a warm-up period before actions impact the game and the machines can have arbitrary structure that is not important to the game playing behavior. It is straightforward for a human to mark the time when this warm-up period ends.

The result of these reductions is shown in Figure 2b and often result in orders of magnitude smaller visualizations, e.g., going from Figure 1b to Figure 1c. Note that interpretable reductions do not change the behavior of the agent, nor the control flow, but are only for improved visualization. The states remaining in the reduced diagram are *decision points* (e.g. state S_2 is the single decision point in Figure 2b), where the next state, and hence future behavior, depends on the observation. The decision points are the key states in the machine that dictate how its behavior is influenced by observations. It is these points where we can gain the most insight about an MM.

Table 1. MM results for control tasks and Atari environments. “DP” refers to decision points, “Obs.” is short for observations, and “Perf.” is short for performance.

Game	QBN Sizes		Original MM				Functional pruning				Minimal MM			
	N_h	N_o	DP	States	Obs.	Perf.	DP	States	Obs.	Perf.	DP	States	Obs.	Perf.
Acrobot	4	4	9	12	38	-77.1	2	4	3	-80.7	3	3	11	-80
	8	8	79	42	200	-77.1	2	4	5	-79.9	3	3	7	-86
CartPole	4	4	6	7	18	500	2	4	5	500	2	3	8	500
	8	8	7	8	58	500	2	4	5	500	3	3	11	500
Lunar Lander	32	32	195	1426	1150	180.4	184	1387	980	172.2	41	41	92	165
	32	64	249	1389	1437	204.9	230	1321	1327	197.3	19	19	73	147
Bowling	32	50	5	608	525	60	0	530	437	60	5	5	19	60
	64	100	5	630	552	60	0	546	488	60	3	4	12	60
Boxing	32	50	0	1274	1270	100	0	1274	1270	100	19	19	109	100
	64	100	0	1097	1095	100	0	1097	1095	100	14	14	101	100
Breakout	32	50	4	2479	2466	404	0	2365	2345	404	8	8	28	404
	64	100	5	1659	1608	404	0	1598	1540	404	11	11	43	404
Pacman	32	50	43	728	716	3060	0	612	597	3060	21	21	70	3060
	64	100	34	895	876	3060	0	776	758	3080	9	9	45	3060
Pong	32	50	0	383	369	21	0	383	369	21	3	3	12	21
	64	100	2	384	369	21	0	271	268	21	4	4	10	21
SeaQuest	32	50	46	2167	2233	2580	0	1679	1577	2580	16	16	140	2580
	64	100	18	2244	2261	2580	0	1834	1883	2580	17	17	135	2580
Space Invaders	32	50	102	1700	1709	1820	0	1314	1350	1820	30	30	40	1820
	64	100	35	1914	1852	1820	0	1832	1802	1820	11	11	27	1820

4.2. Differential Attention for Decision Points

Given an MM decision point, we are interested in understanding how the raw observations (e.g. image pixels) influence its decisions. In particular, for a pair of outgoing branches labeled by discrete observation \hat{f}_1 and \hat{f}_2 , we would like to answer the question: “*What features of the raw observations are most influential to selecting the \hat{f}_1 branch versus \hat{f}_2 ?*”. To help answer this, we consider pairs of raw observations o_1 and o_2 that occur at the decision point during MM execution (e.g. o_{342} and o_{267} in Figure 1c), such that $\hat{f}_i = E_o(o_i)$ where E_o is the encoder of the observation QBN Q_o which discretizes the input features. That is, o_1 and o_2 are example observations that cause the machine to differentiate between the branches. We then produce a *differential attention map* $S(o_1, o_2)$ that highlights the parts of o_1 and o_2 that are most responsible for preferring branch \hat{f}_1 over \hat{f}_2 . To compute $S(o_1, o_2)$, we focus on the set $F(o_1, o_2)$ of discrete features produced by E_o that differ between o_1 and o_2 . As described below, for each $f \in F(o_1, o_2)$, we first produce an attention map $S[f](o_1, o_2)$ that highlights the parts of o_1 and o_2 that “explain” the difference in value of f . $S(o_1, o_2)$ is then just the average of the individual maps.

We compute each $S[f](o_1, o_2)$ via a straightforward, but novel, adaptation of the *Integrated Gradient (IG)* attention

approach (Sundararajan et al., 2017). Originally, IG was used to compute attention maps that explain the decision of a classifier on a single image/observation o . Fundamental to the approach is the notion of a *baseline* image o_b , which is used as a reference that is assumed to not excite the classifier. In an image domain, the baseline is often a constant or noise image. Let $f(o)$ be the classifier response for observation o (usually the largest class-specific input to the softmax layer), noting that $f(o_b)$ will be small, and let $IG_i[f](o, o_b)$ be the corresponding attention value produced by IG for feature/pixel i . The key property of IG, which makes it a meaningful notion of attention, is the relation $\sum_i IG_i[f](o, o_b) = f(o) - f(o_b)$. Thus, the attention value of pixel i can be viewed as its additive contribution to the difference in classification responses for O over O_b . For space reasons we refer the reader to the original paper (Sundararajan et al., 2017) for details of the IG computation.

Figure 3 shows how we adapt IG to compute a differential attention map $S[f](o_1, o_2)$ by treating o_2 as the baseline and letting the response function $f \in F(o_1, o_2)$ be the continuous features computed by the encoder just before the discretization step. That is, differential attention is given by $S[f](o_1, o_2) = IG[f](o_1, o_2)$. This means that differential attention satisfies $\sum_i S_i[f](o_1, o_2) = f(o_1) - f(o_2)$, which has the interpretation that each attention value $S_j[f]$ can




Figure 4. a) Pong minimal MM, b) Bowling minimal MM, c) Bowling pruned MM, d) Acrobot minimal MM.

be viewed as an additive contribution to the difference in response for o_1 and o_2 , i.e. the preference of o_1 over o_2 .¹

4.3. Functional Pruning

After exploring the reduced MM graphs and attentions in Atari, we made two high-level observations. First, as shown in our experiments the differential attention results often indicated that observations were not being used in a strategically meaningful way, e.g., branching on observations that were extremely similar. Second, the graphs beyond different branches at a decision point were often very similar and appeared to address similar situations. We hypothesized that many decision points may not be strategically meaningful, but rather just an artifact of learning. That is, even if a branching decision is not required at a certain point in a game, the inclusion of a decision point that conditions on an arbitrary part of the observation will not hurt the performance as long as good behavior is learned for each of the resulting branches. In such situations, the choice of which branch to traverse may be arbitrary and any one of them may work. Note that the usual attention-based tools (Greydanus et al., 2018; Iyer et al., 2018; Gupta et al., 2020) will simply indicate the agent’s attention and can lead a human to incorrectly infer strategic relevance.

Detecting unnecessary decision points cannot be done by just graph analysis—rather, empirical analysis of modified machines is required. To do this, we conduct a simple form of *functional pruning* (details and pseudo-code in appendix), to identify and prune unnecessary branches at decision points. Our approach considers each MM decision point, and prunes each of the branches one at a time, in order of the least to most frequently visited based on multiple MM runs. When a branch is pruned, we empirically estimate the performance of the resulting MM, noting that when the machine would have previously taken the pruned branch, we force it to take the most frequent branch. If the empirical performance does not degrade beyond a threshold, we permanently remove the branch and move on to the

next pruning step. *The intent of function pruning is not to preserve logical equivalence.* Rather, the intent is to test whether observations were strategically important or just an artifact of learning. After gaining the insights, one could decide to use the original machine or attempt to improve it.

We found greedy pruning to be effective for the MMs considered in this paper. For example, in Figure 1 we see that functional pruning resulted in pruning all but one branch from the single decision point in the MM, leaving an open-loop policy. In such cases, when an MM can be functionally pruned to the point of removing all decision points and leaving an open loop policy, we say that the MM is a *pruned open-loop policy*, which generalizes the notion of open-loop policy to include machines that condition on observations, but in ways that are not strategically necessary. As an analogy, consider a pruned open-loop policy for a human driving to a store along one of two equally good routes. The human may measure the temperature and decide between the routes depending on whether the temperature was an even or odd value. The observation had no real impact on the quality of the policy, since a fixed route could have been selected, but behavior is still seen to depend on observations.

5. Experiments

The only prior approach to compare against is the recent MM minimization approach for analysis of MMs (Koul et al., 2019). As described earlier, we have found it very difficult to gain insights from minimized MMs. To further illustrate this point, in the appendix we provide a quantitative and qualitative comparison of the MM approach to our approach in Atari games. Overall, the minimization approach does not reveal the insights of our new approach.

We consider 7 deterministic Atari environments: Bowling, Boxing, Breakout, MsPacman, Pong, SeaQuest and SpaceInvaders, and 3 stochastic discrete-action classic control tasks: Acrobot, CartPole, and LunarLander. For each experiment, we follow the choices of prior discretization work (Koul et al., 2019) for pre-processing, RPN architecture, QBN architectures, and training via A3C (Mnih et al., 2016) rein-

¹Source code is available at: github.com/modanesh/Differential_IG

enforcement learning. Detailed information on these choices along with hyperparameter choices are in the appendix. In addition to A3C, we investigate and interpret another policy learned by the R2D2 algorithm (Kapturowski et al., 2018) which results are provided in the appendix. We considered two sets of QBN sizes for each environment, shown in Table 1. From the table we see that the agent performance remains the same after reductions for all domains, except for Acrobot and LunarLander, where functional pruning reduced performance within the specified tolerance. The table also gives the number of states and observations, N_h and N_o , after the reductions. Interestingly, in Atari games, no decision points are left after functional pruning, i.e. all of the policies were “pruned open-loop” policies. The following case studies illustrate how our approach is useful for gaining insights and revealing unexpected properties of the decision logic. Additional examples are in the Appendix.

5.1. Atari: Case Studies

Example Insights: Pong and Bowling. The RPN for Pong achieved a maximum score of 21. The policy displays repetitive behavior by performing a “kill shot” against the opponent to win each point, though the behavior is not exactly the same across all shots. Figure 1b shows a view of the original large MM. Figure 1c shows the graph after the interpretable reductions, which is quite small with only one decision point at state S_{197} . This key decision point is the starting point of 3 possible loops, depending on the branch taken, and is entered once per round (each round is one point). This raised the questions of “*What basis is the machine using to decide which loop to enter?*” and “*Is there a strategic reason for the branching decision?*”.

To investigate, we computed the differential attention for the decision of choosing S_{275} over S_{352} , as shown in Figure 3. It is striking that the sample observations associated with the branches are almost identical. The differential attention indicates that the ball region and tips of the paddles are the most critical factors in deciding between the branches. Close inspection reveals that the appearance and location of the ball in the two observations are subtly different. To understand this, we observed that these differences are due to the fact that at the beginning of each round the starting position of the ball is minutely different for even versus odd rounds, which translated to the small difference observed at the decision point. Intuitively, this difference did not appear to have a strategic value. Indeed, after functional pruning (Figure 1d), we see that all branches were removed except for the one through S_{275} , leaving an open-loop strategy with no loss in performance. The even versus odd branching, was an unnecessary artifact of the RNN learning process.

For Bowling, the original large MM had 630 discrete states and 552 discrete observations. Our interpretable reductions




Figure 5. Four observations that enter a state in the corresponding minimal MM. a) Frames of Pong for in-going observations to S_2 , b) Frames of Bowling for in-going observations to S_0 .

revealed only 5 of those states corresponded to true decision points. Further, Figure 4c shows the result of functional pruning to get an MM with no loss in performance. Similar to Pong we end up with a pruned open-loop policy. Again the observations used at decision points were not strategically relevant and rather artifacts of learning.

Comparison to Prior Work. We now compare to the minimization approach of (Koul et al., 2019). For Pong, our approach was able to isolate a single decision point (Figure 1d) where behavior depended on observations in an understandable way, which was ultimately determined to be non-strategic. Meanwhile, Figure 4a shows the minimal MM produced by prior work for the same initial MM. This minimal MM merges the key decision point with semantically unrelated states (e.g. from multiple macros), obscures insights, and gives no understanding of how memory and observations are used. To highlight this, Figure 5a demonstrates four frames of Pong which enter S_2 in the corresponding minimal MM. This is an evidence of minimal MMs merging states that are semantically distinct. Further, it is unclear how the equivalent of functional pruning could be done using the minimal MM, due to the merging of decision points. Also, it is unclear how to uncover the key insight identified by our approach by starting with the minimal MM.

A similar comparison holds for Bowling where our approach resulted in the open-loop policy (Figure 4c). Rather, the minimal MM of the same initial policy is shown in Figure 4b, which resembles the tightly coupled minimal MM of Pong. Figure 5b shows frames that lead to decision point S_0 , which appear to be semantically very different from a human perspective. This is also the case for other states, which makes it very difficult to extract meaningful insights from the minimal MM. Again, it is completely unclear how we could start with the minimal machine and gain the realization that the observations play no significant strategic role, which our approach revealed.

Overall Results. Due to page limit, qualitative comparisons of other Atari games are provided in the appendix. Nevertheless, they provide similar distinctions which show our approach’s advantage over (Koul et al., 2019)’s. Table 1, gives information about the MMs before and after functional pruning for each game, and a comparison with the minimization approach. Note that 0 decision points in-

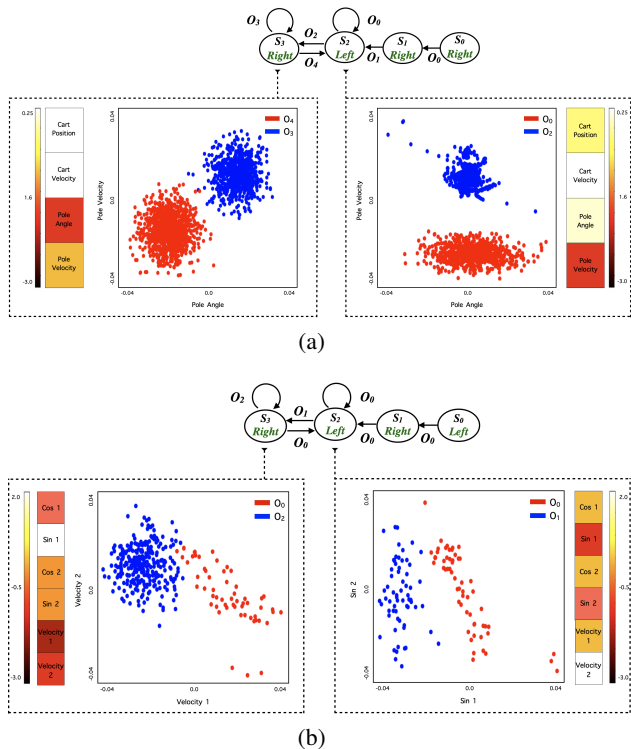


Figure 6. Pruned MMs for control tasks. a) CartPole, and b) Acrobot. In each case, we show attention of features for decision points. Also, we show scatter plots of continuous observations during an episode at each decision point for the two most salient features, where color indicates the discrete machine observation.

indicates an open-loop policy. We see that before functional pruning there is only one case of open-loop policies: Boxing. All other MMs have at least one decision point. This initially makes one to believe that observations are a key part of the overall MM strategy. However, in each such case, we found that it was rare to find a decision point where observations provide strategic values at a decision point. This was confirmed by our most striking finding. After functional pruning, each of the games resulted in open-loop MMs (i.e. zero decision points). Thus, in all cases, Atari RPNs produced MMs that were pruned open-loop policies.

5.2. Stochastic Classic Control Tasks

Acrobot. This control task includes two joints and two links, where the joint between the two links is actuated. Initially, the links are hanging downwards, and the goal is to swing the end of the lower link up to a given height. The state vector gives the sin and cos of the two rotational joint angles and the joint angular velocities. The actions involved are $-1, 0$ or $+1$ torque. We share pruned state machine for QBN sizes $(4, 4)$ in Figure 6b. We applied our differential attention approach to decision point S_2 and found that the most important observation features were ‘sin of joint angle 1’ and ‘joint 2 velocity’. The Figure 6b shows the scatter

plot of the decisions at S_2 versus these features. The plot reveals that for positive sin values, a torque of -1 is applied by starting in S_2 forcing link 1 and link 2 to rise against gravity. This torque is applied until it transitions to S_3 via O_1 which corresponds to a positive velocity for joint 2 as shown in the scatter plot. This happens when link 2 cannot go further up against gravity under the current momentum. In this state, it applies $+1$ torque to supplement the momentum provided by the pull of gravity. It transitions back to S_2 via O_0 only when the joint velocity of link 1 is positive, indicating that it cannot go further up. This loop between state S_2 and S_3 generates enough momentum to eventually reach the goal.

In contrast, we compare against the minimal MM (Figure 4d) extracted via prior work (Koul et al., 2019). This MM, is fully connected and merges many different types of observations. Through various previous attempts, we were unable to elicit a clear description of the machines operation. Rather, our reduction approach was able to lead to the above relatively clear understanding of the machine.

CartPole. We use the standard OpenAI CartPole environment with randomized initial states. For QBN pairs of $(4, 4)$, Table 1 shows that the number of decision points is 6 before functional pruning and reduces to 2 after. The pruned MM is shown in Figure 6a, which has the same simple structure as for Acrobot. The machine primarily transitions between S_2 (left movement) and S_3 (right movement) with self-loops at those states between transitions. The figure shows the differential attention computed over the features at S_2 and S_3 with the key features being ‘pole-velocity’ followed by ‘pole-angle’. The scatter plots for the decisions against these features show that at S_2 there is a clear threshold of positive pole-velocity that transitions to S_3 to take the ‘right’ action, and otherwise continues with ‘left’. This is an intuitive strategy for reducing the velocity. Similar insights are gained via the scatter plot at S_3 , but here both pole-velocity and pole-angle play a role in the decision. Again, our approach was able to produce an MM that has meaningful interpretation.

Lunar Lander. This task involves landing a rover using actions that fire one of three thrusters: main, left, right, and no-op. The MMs for both QBN sizes were significantly larger than for the above tasks, likely due to increased complexity and stochasticity. Functional pruning was ineffective, leaving 184 and 230 decision points for the two QBN sizes, which indicates the observations have a strong strategic role. We analyzed many of the decision points and found that the ‘distance to the ground’ is usually the most salient feature across decision points. One exception is the initial decision point. At this decision point there were three branches with different observations leading to states with different fire actions *main engine*, *right engine* and *left engine*. The attention comparison of these observations shows prime

differences in x-velocity and y-velocity, thereby opposite direction engines are fired to stabilize the rover. This decision point does not attend to rover coordinates (x,y position) and leg-positions, which are more relevant when the lander is closer to the ground. While it is difficult to articulate a concise description of such a machine, we see that this analysis approach is able to provide insights about the decisions that may help build confidence or identify concerns.

6. Discussion and Current Challenges

Our analysis is the first to provide such detailed insights into the decision making of RPNs for deterministic Atari games. Indeed, the observation that all of the considered policies resulted in pruned open-loop controllers was unexpected a priori and not apparent from prior work. For example, prior work on attention analysis of Atari policies, even for the deterministic setting, leaves one to believe that observations play key roles in decision making. It is tempting to discount the above insights, due to the deterministic setting, since 1) it is clear that there exist open-loop controllers for any deterministic domain, and 2) prior work has shown that search is able to find effective open-loop plans for some Atari games (Bellemare et al., 2013; Machado et al., 2018). However, these points do not imply that an RPN would necessarily learn an open-loop controller and indeed we observed that they do not. It is reasonable to expect that RPNs would meaningfully use observations to get a more general policy, but we instead saw the role for observations was very different. This demonstrates the importance of developing a variety of tools to reveal insights that rely less on human assumptions and interpretation.

Our preliminary experience with larger and more complex environments shows that sometimes our approach does not reveal easily analyzable MMs. This was apparent in stochastic Atari experiments, where we observed that the discrete sequences produced across different episodes have some overlap, but are dominated by large numbers of disjoint states. Some potential explanations for this are: 1) An inadequacy of our approach—e.g., the quantization process and/or reduction steps/analysis may need to be improved; 2) Our approach may be identifying the fundamental nature of the learned RNN policies. That is, the policies may effectively use large numbers of trajectories to encode large numbers of effectively (pruned) open-loop patterns. As new observations are encountered the machines then map to encoded patterns in a nearest neighbor style. This second possibility would suggest the need for improved RPN training approaches to support data efficiency and interpretability.

Acknowledgements

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under the XAI and CAML programs. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the DARPA, the Army Research Office, or the United States government.

References

- Atrey, A., Clary, K., and Jensen, D. Exploratory not explanatory: Counterfactual analysis of saliency maps for deep {rl}. In *International Conference on Learning Representations*, 2020.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Cechin, A. L., Regina, D., Simon, P., and Stertz, K. State automata extraction from recurrent neural nets using k-means and fuzzy clustering. In *23rd International Conference of the Chilean Computer Science Society, 2003. SCCC 2003. Proceedings.*, pp. 73–78, Nov 2003. doi: 10.1109/SCCC.2003.1245447.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using rnn encoder-decoder for statistical machine translation, 2014.
- Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. Empirical evaluation of gated recurrent neural networks on sequence modeling, 2014.
- Greydanus, S., Koul, A., Dodge, J., and Fern, A. Visualizing and understanding Atari agents. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1792–1801, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- Gupta, P., Puri, N., Verma, S., Kayastha, D., Deshmukh, S., Krishnamurthy, B., and Singh, S. Explain your move: Understanding agent actions using focused feature saliency. In *International Conference on Learning Representations*, 2020.
- Hausknecht, M. and Stone, P. Deep recurrent q-learning for partially observable mdps. In *2015 AAAI Fall Symposium Series*, 2015.
- Ho, J. and Ermon, S. Generative adversarial imitation learning, 2016.

- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Iyer, R., Li, Y., Li, H., Lewis, M., Sundar, R., and Sycara, K. P. Transparency and explanation in deep reinforcement learning neural networks. *CoRR*, abs/1809.06061, 2018.
- Kapturowski, S., Ostrovski, G., Quan, J., Munos, R., and Dabney, W. Recurrent experience replay in distributed reinforcement learning. In *International conference on learning representations*, 2018.
- Koul, A., Fern, A., and Greydanus, S. Learning finite state representations of recurrent policy networks. In *International Conference on Learning Representations*, 2019.
- Machado, M. C., Bellemare, M. G., Talvitie, E., Veness, J., Hausknecht, M., and Bowling, M. Revisiting the arcade learning environment: Evaluation protocols and open problems for general agents. *Journal of Artificial Intelligence Research*, 61:523–562, 2018.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937, 2016.
- Mott, A., Zoran, D., Chrzanowski, M., Wierstra, D., and Rezende, D. J. Towards interpretable reinforcement learning using attention augmented agents. *CoRR*, abs/1906.02500, 2019.
- Olson, M., Khanna, R., Neal, L., Li, F., and Wong, W.-K. Counterfactual state explanations for reinforcement learning agents via generative deep learning. *Artificial Intelligence*, In Press, 2021.
- Paull, M. C. and Unger, S. H. Minimizing the number of states in incompletely specified sequential switching functions. *IRE Transactions on Electronic Computers*, pp. 356–367, 1959.
- Sundararajan, M., Taly, A., and Yan, Q. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 3319–3328. JMLR. org, 2017.
- Tiño, P., Horne, B. G., Giles, C. L., and Collingwood, P. C. Finite state machines and recurrent neural networks—automata and dynamical systems approaches. In *Neural networks and pattern recognition*, pp. 171–219. Elsevier, 1998.
- Van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double q-learning. In *Thirtieth AAAI conference on artificial intelligence*, 2016.
- Wang, Z., Schaul, T., Hessel, M., Van Hasselt, H., Lanctot, M., and De Freitas, N. Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv:1511.06581*, 2015.
- Weiss, G., Goldberg, Y., and Yahav, E. Extracting automata from recurrent neural networks using queries and counterexamples, 2017.
- Weiss, G., Goldberg, Y., and Yahav, E. Learning deterministic weighted automata with queries and counterexamples, 2019.
- Yang, Z., Bai, S., Zhang, L., and Torr, P. H. S. Learn to interpret atari agents. *CoRR*, abs/1812.11276, 2018.