Real-Time Recurrent Reinforcement Learning

Julian Lemmel, Radu Grosu

Technical University Vienna Karlsplatz 13, 1040 Wien AT julian.lemmel@tuwien.ac.at

Abstract

In this paper we propose real-time recurrent reinforcement learning (RTRRL), a biologically plausible approach to solving discrete and continuous control tasks in partiallyobservable markov decision processes (POMDPs). RTRRL consists of three parts: (1) a Meta-RL RNN architecture, implementing on its own an actor-critic algorithm; (2) an outer reinforcement learning algorithm, exploiting temporal difference learning and dutch eligibility traces to train the Meta-RL network; and (3) random-feedback local-online (RFLO) learning, an online automatic differentiation algorithm for computing the gradients with respect to parameters of the network. Our experimental results show that by replacing the optimization algorithm in RTRRL with the biologically implausible back propagation through time (BPTT), or real-time recurrent learning (RTRL), one does not improve returns, while matching the computational complexity for BPTT, and even increasing complexity for RTRL. RTRRL thus serves as a model of learning in biological neural networks, mimicking reward pathways in the basal ganglia.

Introduction

Model-free deep reinforcement-learning algorithms, leveraging recurrent neural network architectures (RNNs), have recently been shown to serve as a strong baseline for a wide range of partially-observable Markov decision processes (POMDPs) (Ni, Eysenbach, and Salakhutdinov 2022). Gradient-based reinforcement-learning (RL) algorithms, such as temporal-difference (TD) methods, have been shown to be sample efficient, and come with formal convergence guarantees when using linear function approximation (Sutton and Barto 2018).

However, linear functions are not able to infer hidden state variables that are required for solving POMDPs. RNNs, on the other hand, can compensate for the partial observability in POMDPs, since they are capable of aggregating information about the entire sequence of preceding observations.

Neural networks were originally inspired by biological neurons, which are in general recurrently connected, and subject to synaptic plasticity. These long-term changes of synaptic efficacy are mediated by locally accumulated proteins, and a scalar-valued reward signal represented by some neurotransmitter (e.g. dopamine) concentration (Wise 2004).



Figure 1: The RTRRL recurrent network architecture implements the Meta-RL framework, as defined in Wang et al. (2018), by feeding the past action a_{t-1} to the CT-RNN alongside the current reward r_t and the current observation o_t . The latent state h_t which is computed by the CT-RNN, is then used to produce the next action a_t and the next value estimate \hat{v}_t through linear mappings.

The ubiquitous backpropagation through time algorithm (BPTT) (Werbos 1990), which is used for training RNNs in practice, appears to be biologically implausible, due to the need of weight transport (Bartunov et al. 2018), and distinct forward and backward phases. Thus, BPTT requires individual neurons to alternate between computation and precisely communicating gradients of their synaptic parameters to each other. With BPTT, RNN-based RL algorithms renounce any claim to biological interpretation.

Biologically plausible methods for computing gradients in RNNs however exist, such as random-feedback online learning (RFLO) (Murray 2019), a simplified versions of real-time recurrent learning (RTRL) (Williams and Zipser 1989).

The main question we asked in this paper was if combining biologically-plausible methods for learning and gradient computation in RNNs, such as TD and RFLO, respectively, would be able to learn and solve POMDP RL tasks?

Taking advantage of the previous work, we are able here to answer the above question in a positive fashion, with a novel approach consisting of three basic building blocks:

- 1. *A Meta-RL RNN architecture*, whose interface and recurrent connections are such that when trained, they become on their own an actor-critic algorithm,
- 2. A $TD(\lambda)$ Actor-Critic algorithm exploiting temporal dif-

ference and dutch eligibility traces, to train the weights of the Meta-RL network, and

3. *The Biologically Plausible RFLO* online automatic differentiation algorithm, for computing the gradients of the parameters in the Meta-RL network.

We call our new, biologically plausible, online RL approach, real-time recurrent reinforcement learning (RTRRL).

The use of $TD(\lambda)$ in RTRRL, the most basic version of an actor-critic algorithm with no other enhancements, such as the batched experience replay, is able to solve discrete as well as continuous POMDP tasks, in a biologically-plausible fashion. Our experimental results show that, when replacing the RFLO optimization of RTRRL with the BPTT optimization, one achieves a similar accuracy, but in many cases at a much slower convergence time. Moreover, replacing RFLO with RTRL optimization considerably increases time complexity, without significant improvements in accuracy.

Motivation

The main objections for BPTT-plausibility in biological neural networks are the reliance on shared weights, for forward and backward connections, and reciprocal error-transport, a form of propagating back the errors without interfering with neural activity (Bartunov et al. 2018). Another concern is that it requires storing long sequences of exact the exact activation for each cell. (Lillicrap and Santoro 2019)

We slightly rephrase the premises of Bartunov et al. (2018) and Lillicrap and Santoro (2019) in the following fashion. A biologically plausible learning algorithm has to be:

- 1. Local, up to some low-dimensional reward signal.
- 2. Online, jointly computing updates and outputs.
- 3. Without weight transport, in its error computation.

By *weight transport*, we refer to synapses propagating back error signals mirroring the strength of forward synapses. We will further elaborate on the above requirements for obtaining biological plausibility in Appendix .

Previous work on feedback alignment (Lillicrap et al. 2014) demonstrated that weight transport is not strictly necessary for training deep neural networks. In particular, they showed that randomly initialized feedback matrices, used for propagating back gradients to previous layers in place of the forward weights, are also leading to acceptable function approximators. Moreover, the forward weights appear to align with the fixed backward weights during training.

Online supervised learning for RNNs was first described by Williams and Zipser (1989), who introduced RTRL as an alternative to backpropagation. More recently, a range of more computationally efficient variants were introduced. (Tallec and Ollivier 2017; Roth, Kanitscheider, and Fiete 2018; Mujika, Meier, and Steger 2018; Murray 2019). Marschall, Cho, and Savin (2020) proposed a unifying framework for all these variants and assessed their biological plausibility while showing that many of them can perform on-par with BPTT. RFLO (Murray 2019), stands out due to its biological update rule, that still performed well in the experiments of Marschall, Cho, and Savin (2020).

Another motivation for investigating online training algorithms of neural networks is the promise of energy-efficient neuromorphic hardware(Zenke and Neftci 2021). The aim is to produce bionic integrated circuits, mimicking biological neurons. However, they also require biologically plausible update rules to enable efficient training.

Our work shows that online reinforcement learning with RNNs is possible with RTRRL-RFLO, which fulfills all our premises for biologically plausible learning. Our algorithm combines online RNN gradient computation with $TD(\lambda)$. We create a fully online learning neural controller that does not require multi-step unrolling, weight transport or communicating gradients over long distances. Our algorithm succeeds in learning a policy from one continuing stream of experiences alone, meaning that no replay buffer is used.

To the best of our knowledge, the proposed RTRRL, is the first fully biologically plausible reinforcement learning algorithm for rate-based recurrent neural networks.

Real-Time Recurrent RL

In this section we present RTRRL, our novel real-time recurrent RL approach to POMDPs. As mentioned, this approach consists of: 1) A Meta-RL actor-critic RNN architecture, which after training becomes on its own an actor-critic algorithm, 2) A TD(λ)-Actor-Critic learning algorithm, exploiting temporal-difference and dutch eligibility traces, to train the weights of the Meta-RL network, and 3) A biologicallyplausible real-time RFLO optimization technique, computing the gradients of the parameters in the MetaRL network.

In the following subsections, we provide a gentle and selfcontained introduction to each of the constituent parts of our RTRRL framework. We then put all pieces together, and discuss the pseudocode of the overall RTRRL approach.

The Meta-RL RNN Architecture

The Meta-RL actor-critic RNN architecture used by RTRRL is shown in Figure 1. It features a RNN with linear output layers for the actor and critic functions. At each step, the RNN computes an estimated latent state and the two linear layers compute the next action and the next value, from the latent state, respectively. Since the synaptic weights of the network are trained (slowly) to choose the actions with most value, and the network states are also updated during computation (fast) towards the same goal, this architecture is also called a Meta-RL architecture (Wang et al. 2018).

As shown by Wang et al. (2018), a Meta-RL RNN can be trained to solve a general family of parameterized tasks, where all instances follow a common structure, and that the activations in trained RNNs mimic dopaminergic reward prediction errors (RPEs) measured in primates. They also show that the network is capable of inferring the underlying parameters of each new task and subsequently solve unseen instances after training. In particular, they use LSTM cells (Hochreiter and Schmidhuber 1997), receiving the observation, past action, and reward as inputs, and computing a latent state from which the next action and value estimate are inferred by a linear mapping. In Figure 1 we use the architecture of that work, but replace the LSTM with a continuoustime recurrent neural network (CT-RNN).

Continuous-Time RNN (Funahashi and Nakamura 1993) This type of RNN can be interpreted as rate-based model of biological neurons. They have the following state dynamics:

$$\dot{h} = \tau^{-1} \left(-h + \varphi (W_I x + W_R h + b) \right)$$

where factor τ is the time constant, φ , a non-linear activation function, x, the input, W_I, W_R , the synaptic weights for input and states, respectively, and b is the bias.

Wang et al. (2018) argue that, the Meta-RL RNN effectively learns to implement an inner RL algorithm, that empowers it to learn about, and solve tasks on its own. Behrens et al. (2018) argue that, the timescale of plasticity expressed through dopaminergic synapses in the prefrontal cortex, is magnitudes below the time needed for primates to solve tasks, hinting at a hierarchically layered RL algorithms reminiscent of the framework of Wang et al. (2018).

The Actor-Critic Learning Algorithm

As a learning algorithm, RTRRL also uses an Actor-Critic architecture with eligibility traces. Here, the actor and the critic (sharing the common recurrent CT-RNN backbone), are used to compute the synaptic-weights updates, ensuring that for the current observation and reward, and previous action, the network outputs the most valuable action.

In particular, the value produced by the critic, is used to computing the TD error and the forward dutch eligibility traces, which are thereafter used to properly update the synaptic weights. Hence, the RL algorithm used by RTRRL ia an Actor-Critic-TD(λ) algorithm. In order to understand it, we first discuss TD learning, and then eligibility traces.

Temporal-Difference Learning (TD). TD Learning relies only on local information by bootstrapping (Sutton and Barto 2018). It is on-line, which makes it applicable to a wider range of problems, as it does not rely on completing an entire episode prior to computing the updates. After each action, the reward r_t , and past and current experience s_t and s_{t+1} are used to compute the TD error δ of value estimates $\hat{v}_{\theta}(s)$. The following Equation defines the TD-error:

$$\delta_t = r_t + \gamma \hat{v}_{\theta_t}(s_{t+1}) - \hat{v}_{\theta_t}(s_t) \tag{1}$$

The value function is learned by regression towards the bootstrapped target. Updates can be computed by taking the gradient of the value approximator $\theta_{t+1} \leftarrow \theta_t + \eta \delta_t \nabla_{\theta} \hat{v}_{\theta_t}(s_t)$, where η is a small step size. In order to also learn behavior, we use an Actor-Critic (AC) policy gradient method. In AC algorithms, the actor computes the actions, and the critic evaluates the actor's decisions. In particular, the value $\hat{v}_{\theta_t}(s_t)$ of the current state s_t is compared to its value $r + \gamma \hat{v}_{\theta_t}(s_{t+1})$ after taking an action a_t , that is, after improving its knowledge by getting a reward r and moving to state s_{t+1} . This leads to the the TD-error discussed above.

The policy (actor) is in this case a parameterized function π_{φ} that maps state s to a distribution of actions p(a). Parameters φ are in most cases the weights of a neural network.

The policy π_{φ} is trained using gradient ascent. By taking the gradient of the policy with respect to the log action probability, and multiplying with the TD-error, we can compute the updates $\varphi \leftarrow \varphi + \alpha \delta \nabla_{\varphi} \log \pi_{\varphi}(a)$. Intuitively, this tells the actor to increase the probability for the chosen action whenever the RPE is positive, that is, when the reward was better than predicted. Likewise, when the reward was worse than predicted, the action probability is lowered.

The TD-error is a measure for the accuracy of the rewardprediction, acting as RPE. Given its importance, it is used to update both the actor and the critic, by acting as a reinforcement signal (Sutton and Barto 2018).

Eligibility Traces (ET). The algorithm just described is known as TD(0). It is impractical when dealing with delayed rewards in an online setting, since value estimates need to be updated backwards in order to factor in temporal dependencies. ETs are a way of factoring-in future rewards. The idea is to keep a running average of each synapse's contribution to the network's output. This can be thought of as a short-term memory, paralleling the long-term one represented by the synapse weights. ETs unify and generalize Monte-Carlo and TD methods (Sutton and Barto 2018).

Particularly, $TD(\lambda)$ makes use of ETs. Weight updates are computed by multiplying the trace with the TD-error δ . By accumulating the gradients of the state-value function, each weight's contribution is recorded. The trace e^{θ} decays with factor $\gamma \lambda$ where γ is the discount factor:

$$e_t^{\theta} = \gamma \lambda \ e_{t-1}^{\theta} + \nabla_{\theta} \hat{v}_{\theta_t}(s_t) \theta_{t+1} \leftarrow \theta_t + \eta^{\theta} \ \delta_t \ e_t^{\theta}$$
(2)

Since we use a parameterized linear-function approximation $\hat{v}_{\theta}(s_t) = w^{\top} s_t$, with parameters w as in the original TD(λ), the gradient of the loss with respect to w, is simply $\nabla_w \hat{v}_{\theta} = s_t$. A later refinement of linear TD(λ) called True online TD(λ) (Seijen and Sutton 2014; Sutton and Barto 2018), improves the accuracy of updates done to the value function. The resulting trace is called a "dutch" trace. Here, calculating the TD-error can be done without using the same set of weights twice. Hyperparameter α couples the rate of change of the ET with the optimization step size.

$$\delta_{t} = r_{t} + \gamma \hat{v}_{\theta_{t}}(s_{t+1}) - \hat{v}_{\theta_{t-1}}(s_{t})$$

$$e_{t}^{\theta} = \gamma \lambda e_{t-1}^{\theta} + \alpha s_{t-1} - \alpha \gamma \lambda \left[e_{t-1}^{\theta \top} s_{t-1} \right] s_{t-1} \qquad (3)$$

$$\theta_{t+1} = \theta_{t} + \delta_{t} e_{t}^{\theta} + \alpha \left[\theta_{t-1}^{\top} s_{t-1} - \theta_{t}^{\top} s_{t-1} \right] s_{t-1}$$

Linear TD methods come with provable convergence guarantees (Sutton and Barto 2018). However, the simplicity of the function approximator fails at accurately representing complex functions needed for solving harder tasks. Replacing the linear functions with non-linear DNNs can introduce inaccuracies in the optimization. However in practice, multilayer perceptrons (MLPs) can lead to satisfactory results, e.g. on the Atari benchmarks (Daley and Amato 2020).

The Biological-Optimization Algorithm

The optimization algorithm in RTRRL computes the gradients of the synaptic weights, in a biologically plausible fashion. This has to be done for the Meta-RL RNN, in an efficient online fashion as discussed in Section . For this purpose, RTRRL is using the RFLO algorithm. However, to asses the loss in precision of RTRRL while using RFLO, we also plug RTRL and BPTT in our algorithm, instead of RFLO. In Figure 2, we show how the gradients are passed back to each state of the Meta-RL RNN for RFLO.

Real-Time Recurrent Learning (RTRL). RTRL was proposed as an RNN online-optimization algorithm for infinite horizons (Williams and Zipser 1989). This is made possible by estimating the gradient of network parameters in each step, and computing an estimate of the error-vector during the feedforward computation. As in stochastic gradient descent, using a small learning rate η can mitigate the noise introduced by estimating the true gradient. One big advantage of RTRL-based algorithms, is that the computation time of an update-step is constant in the number of task steps. Nonetheless, RTRL has a much worse computation complexity than BPTT, and is therefore not used in practice.

The update rule used in RTRL is derived as follows. Given a dataset consisting of the multivariate time-series $x_t \in \mathbb{R}^I$ of inputs and $y_t \in \mathbb{R}^O$ of labels, we want to minimize some loss function $\mathcal{L}_{\theta} = \sum_{t=0}^{T} L_{\theta}(x_t, y_t)$ by gradient descent. This is achieved by taking small steps in the direction of the negative gradient of the total loss:

$$\Delta \theta = -\eta \nabla_{\theta} \mathcal{L}_{\theta} = -\eta \sum_{t=0}^{T} \nabla_{\theta} L_{\theta}(t)$$
(4)

We can compute the gradient of the loss at each individual step as $\Delta\theta(t) = \nabla_{\theta}L_{\theta}(t) = \nabla_{\theta}\hat{y}_t\nabla_{\hat{y}_t}L_{\theta}(t)$ with \hat{y}_t being the output of the RNN at timestep t. A common choice in supervised learning is the MSE $L_{\theta}(t) = \frac{1}{2}\langle \varepsilon_t, \varepsilon_t \rangle$ where $\varepsilon_t = \hat{y}(t) - y(t) \in \mathbb{R}^O$ and $\nabla_{\hat{y}_t}L_{\theta}(t) = \varepsilon_t$. When employing an RNN with linear output mapping the gradient of the model output can be further expanded into $\nabla_{\theta_R}\hat{y}_t = \nabla_{\theta_R}h_t\nabla_{h_z}\hat{y}_t$. RTRL is defined for the special class of Neural ODEs discussed in the previous section.

We calculate the gradient of the Neural ODE unit's activation with respect to the parameters θ , recursively. To this end, we define the immediate Jacobian $\bar{J}_t = \nabla_{\theta} f(x_t, h_t)$, with ∇_{θ} being the partial derivative with respect to θ , and the approximate Jacobian trace \hat{J}_t . We will denote the total derivative of f wrt. θ as $\frac{d}{d\theta} f(x_t, h_t)$.

$$\hat{J}_{t+1} := \frac{d}{d\theta} h_{t+1} = \frac{d}{d\theta} (h_t + f(x_t, h_t))$$

$$= \frac{d}{d\theta} h_t + \frac{d}{d\theta} f(x_t, h_t)$$

$$= \hat{J}_t + \nabla_{\theta} f(x_t, h_t) + \frac{d}{d\theta} h_t \nabla_{h_t} f(x_t, h_t)$$

$$= \hat{J}_t + \bar{J}_t + \hat{J}_t \nabla_{h_t} f(x_t, h_t)$$

$$= \hat{J}_t (\mathbb{I} + \nabla_{h_t} f(x_t, h_t)) + \bar{J}_t$$
(5)

The last row defines the Jacobian trace recurrently, in terms of the immediate Jacobian and a linear combination of the past trace $\hat{J}_t \nabla_{h_t} f(x_t, h_t)$, allowing to calculate it in parallel

to the forward computation of the RNN. When taking an optimization step we can calculate the actual gradients as:

$$\Delta \theta(t) = \hat{J}_t \nabla_{h_t} L_{\theta(t)} = \hat{J}_t W_{out}^\top \varepsilon_t \tag{6}$$

Being online, RTRL is biologically plausible in the time domain. However, the error signal $W^{\top} \varepsilon_t$ still assumes weight transport and so do the gradients communicated within the Neural ODE $\hat{J}_t \nabla_{h,t} f(x_t, h_t)$.

In general, RTRL has complexity $\mathcal{O}(n^4)$ in the number of neurons *n* compared to $\mathcal{O}(n^2)$ for BPTT (considering fixed horizon). However, a computationally more efficient version for sparse RNNs was proposed by Menick et al. (2020).

Random-Feedback Local Online Learning (RFLO). A biologically plausible descendant of RTRL is RFLO (Murray 2019). This leverages the Neural ODE of simple CT-RNNs in order to simplify the RTRL update substantially, dropping all parts that are biologically implausible. RFLO improves biological plausibility of RTRL in two ways:

- 1. Weight transport during error backpropagation is avoided by using feedback alignment.
- 2. Locality of gradient information is ensured by dropping all non-local terms from the gradient computation.

The CT-RNN above with N hidden states, I inputs, and activation φ , can be simplified by considering a combined weight matrix $W \in \mathbb{R}^{N \times X}$ and a vector $\xi \in \mathbb{R}^{Z}$, where Z = I + N + 1. Each neuron has a time-constant $\tau \in \mathbb{R}^{N}$ and the next state $h_{t+1} \in \mathbb{R}^{N}$ is computed as follows:

$$h_{t+1} = h_t + \frac{1}{\tau} \left(-h_t + \varphi(W\xi_t) \right) \quad \xi_t = \begin{bmatrix} x_t \\ h_t \\ 1 \end{bmatrix} \in \mathbb{R}^Z \quad (7)$$

where x_t is the input at time t, and the 1 concatenated to ξ_t accounts for the bias. The output $\hat{y}_t \in \mathbb{R}^O$ is given by a linear mapping $\hat{y}_t = W_{out}h_t$. The latent space follows the Neural ODE $\tau \dot{h} = -h_t + \varphi(W\xi_t)$.

RFLO leverages the state-update expression in order to simplify the RTRL update. For brevity, we only show the resulting update rule. A derivation is in the appendix.

$$\hat{J}_{t+1}^{W} \approx (1 - \frac{1}{\tau}) \, \hat{J}_{t}^{W} + \frac{1}{\tau} \varphi'(W\xi_{t})^{\top} \xi_{t}$$
 (8)

Weight transport is avoided by using feedback alignment for propagating gradients. Parameter updates are computed as $\Delta W(t) = \hat{J}_t^W B \varepsilon_t$ using a fixed random matrix *B*. Effective learning is still achievable with this simplified version as shown in (Murray 2019; Marschall, Cho, and Savin 2020).

While RTRL complexity is $\mathcal{O}(n^4)$ in the number of neurons, RFLO has the same complexity as backpropagation $\mathcal{O}(n^2)$. The reader is referred to (Murray 2019) and Marschall, Cho, and Savin (2020) for a detailed comparison between RTRL and RFLO.

Putting All Pieces Together

Having described each piece of our RTRRL approach, we now proceed discussing how RTRRL puts them together. Algorithm 1 shows the outline of the RTRRL approach. Intuitively, the approximate Jacobian \hat{J}_t is computed at each



Figure 2: Schematics showing how gradients are passed back to each state of the RNN. The network consists of Linear $TD(\lambda)$ and a shared RNN body. Gradients with respect to actor and critic losses are propagated back towards h_t and h_{t-1} respectively.

RNN step, and combined with the TD-error of True online TD(λ) to update the RNN weights. We follow the Meta-RL approach delineated in (Wang et al. 2018) feeding in the previous action and reward to the RNN body alongside the observation at each step. π and \hat{v} are the Actor and Critic functions parameterized by θ_A and θ_C .

We train the Actor and Critic using $\text{TD}(\lambda)$ and take small steps in direction of the action probability and value estimate respectively. The gradients for each function are accumulated by means of additive eligibility traces $e_{A,C}$ with λ decay. Additionally, the gradients are passed back to the single layer RNN, parameterized by θ_R , through random feedback matrices B_A and B_C respectively. A third, combined eligibility trace e_R is maintained for the RNN, summarizing the combined gradient. For improved exploration, we also compute the gradient of the action distribution's entropy $H(\pi)$ and add it to the gradient of the policy objective. When using RFLO the passing back of the gradients is done via randomly initialized fixed backwards matrices $B_{A,C}$. This way we arrive at a biologically-plausible algorithm.

The Jacobian in Algorithm 1 can also be computed using RTRL. This reduces the variance in the gradients at a much higher computational cost. Similarly, we can choose to propagate back to the RNN by using the forward weights as in backpropagation. However, we found that the biologically plausible feedback alignment works just as well.

RTRRL Neuron Model. We use the simplified CT-RNN of the previous section for the RNN body, where we compute the Jacobian \hat{J}_t by using RTRL or RFLO as explained above. Extending RFLO, we derive an update rule for the time-constant parameter τ . Again, the full derivation of this Jacobian can be found in Appendix :

$$\hat{J}_{t+1}^{\tau} \approx \hat{J}_{t}^{\tau} (1 - \frac{1}{\tau}) + \frac{1}{\tau^2} (h_t - \varphi(W\xi_t))$$
 (9)

RTRRL Hyperparameters. The hyperparameters that can be tuned for individual environments are γ , $\lambda_{A,C,R}$, $\alpha_{A,C}$ and α . Our approach does not introduce any new ones

over True online $TD(\lambda)$ other than lambda and learning rate for the RNN. A detailed list of all hyperparameters and RTRRL implementation details are given in the Appendix.

Experiments

We evaluate the feasibility of our RTRRL approach by testing it on the popular RL benchmarks, provided by the gymnax (Lange 2022) and brax (Freeman et al. 2021) python packages. The tasks comprise fully and partially observable MDPs, with discrete and continuous actions.

The gymnax package implements a range of popular classic and modern control tasks used to benchmark neuralnetwork policies. As baselines we consider $TD(\lambda)$ with Linear Function Approximation, and Proximal Policy Optimization (PPO) (Schulman et al. 2017), with BPTT for the LSTM and CT-RNN models. Our implementation of PPO is based on purejaxrl (Lu et al. 2022). For each environment, we trained a Meta-RL network with 32 recurrent neurons for either a maximum of 50 million steps or until 20 subsequent epochs showed no improvement.

Our results show that RTRRL with RFLO, as discussed in previous sections, as the default, can often outperform PPO that relies on BPTT. In Table 1 we summarize all our results for the experiments explained in the remainder of this section. The column "RFLO CT-RNN" corresponds to RTRRL as defined above, but we have also included for comparison, results for RTRRL, where RTRL replaces RFLO. The values presented in the table are the median and standard-deviation of the best evaluation episodes, experienced throughout the training of 5 runs per environment and algorithm.

Discrete Actions. For discrete actions, the outputs of the actor are the log probabilities for each action, and past actions fed to the RNN are represented in one-hot encoding. In Table 1 we report results for experiments on the classical Acrobot and CartPole benchmarks but also for the bsuite environments (Osband et al. 2020) including DeepSea, MemoryChain and UmbrellaChain, that are specifically tailored to test individual aspects of RL tasks

Model	Linear TD(λ)	RTRL CTRNN	RFLO CTRNN	PPO CTRNN	PPO LSTM
Environment					
MemoryChain-4 *	$0.07 {\pm} 0.01$	1.00 ±0.00	1.00 ±0.00	1.00 ±0.00	1.00 ±0.00
MemoryChain-8 *	$0.07 {\pm} 0.03$	1.00 ±0.00	1.00±0.28	1.00 ±0.00	1.00 ±0.00
MemoryChain-16 *	$0.12{\pm}0.04$	1.00±0.36	$0.87 {\pm} 0.42$	1.00 ±0.00	1.00 ±0.00
DeepSea-4	0.99 ±0.00	0.99 ±0.42	0.99 ±0.31	0.99 ±0.00	0.99 ±0.00
DeepSea-8	0.99 ±0.48	0.99 ±0.31	0.99 ±0.42	0.99 ±0.00	$0.00 {\pm} 0.54$
DeepSea-16	-0.00 ± 0.42	0.99 ±0.46	0.99 ±0.51	$0.00{\pm}0.00$	$0.00{\pm}0.00$
Acrobot	-500.00 ± 0.00	-80.97 ±187.54	-86.72 ± 227.18	-471.40 ± 221.84	-84.14 ± 180.65
CartPole-vel *	$454.55 {\pm} 26.52$	500.00 ±0.00	500.00 ±0.00	112.11 ± 40.38	138.37 ± 31.41
CartPole-pos *	$55.87 {\pm} 1.18$	137.36 ± 177.76	500.00±122.30	70.76 ± 11.42	56.15 ± 82.74
MetaMaze	$15.30{\pm}40.00$	$17.20{\pm}12.94$	19.60 ±13.49	$10.80{\pm}1.34$	$10.40{\pm}1.04$
DiscountingChain	1.10±0.05	$1.04{\pm}0.04$	1.10 ±0.04	1.10 ±0.00	1.10 ±0.00
UmbrellaChain	$0.83 {\pm} 0.31$	1.19 ±0.59	$1.09 {\pm} 0.60$	$1.04{\pm}0.12$	1.03 ± 0.12
PointRobot	$0.82{\pm}0.51$	$1.20{\pm}0.42$	1.40±0.27	$0.44{\pm}0.05$	$0.42{\pm}0.07$
Reacher	$115.47 {\pm} 17.68$	118.70 ± 8.59	118.77±10.66	$16.72 {\pm} 0.91$	17.27 ± 1.01
Swimmer	$96.56 {\pm} 4.91$	$87.69 {\pm} 10.96$	110.72 ±62.50	30.76 ± 1.74	29.05 ± 1.97
MountCarCont-vel *	-41.88 ± 14.85	-1.31 ±13.16	-20.00 ± 35.18	-1079.70 ± 167.95	-1126.11±212.34
MountCarCont-pos *	-46.91 ± 13.71	1.23 ±7.27	-59.54 ± 38.86	-1071.67 ± 36.35	-1077.49 ± 114.22

Table 1: Summary of RTRRL experiments: Numbers reported are the median reward and the standard deviation for 5 runs (larger is better), a " * " marks POMDPs. "RFLO CTRNN" denotes our biologically plausible instantiation of RTRRL. In each case, the best validation reward achieved throughout training was used. On average, RTRRL arguably achieves the best and most balanced results.

such as exploration and memory capacity. Additionally we included the MetaMaze environments that was introduced by Miconi, Clune, and Stanley (2018) and to test Meta-RL algorithms. As one can see, RTRRL performs best in average, followed by RTRRL with RTRL, as the second best algorithm.

Continuous Actions. To obtain a stochastic policy, a common trick is to use a probability distribution parameterized by the model output. A natural choice is to use a Gaussian whose mean and standard-deviation are outputs of the neural network. Assuming i.i.d. action components, the network output has to be in $\mathbb{\bar{R}}^{2U}$ where U is the size of the action vector. An action is obtained by sampling from the distribution. Reparametrization allows for obtaining the probability of the selected continuous action and subsequently computing the gradient of $\log \pi[a]$ in policy-gradient algorithms. In practice, clipping the magnitude of policy gradients to \leq 1, in order to avoid exploding gradients that occur when the variance of the action distribution becomes very small, worked best. The continuous action environments included in our experiments, and whose results are illustrated in Table 1 were MountCarCont, PointRobot, Reacher and Swimmer. Moreover, they also included environments implemented in brax, such as, mujoco.

Masked Observations. We took the classical CartPole MDP and transformed it into a POMDP by masking out parts of the observation. We consider observations that only contain the positions (x, θ) or the velocities $(\dot{x}, \dot{\theta})$, called CartPole-pos and CartPole-vel respectively. The same was done with MountainCarContinuous. Our results show that, when observing positions only, RFLO is the only method able to properly solve this task.

In a similar fashion, we masked observations for the continuous action brax environments. We only kept even entries of the observation, discarding odd ones, to create a POMDP. These physical simulations are computationally demanding which is why we only report the best evaluation epoch for RTRRL after tuning compared to the tuned PPO baselines that were provided by the package authors. The best rewards are 2720.03 (2386.09) for Ant, 2030.10 (2317.57) for Halfcheetah and 1026.78 (1072.85) for Hopper with the tuned baselines in parenthesis, respectively. A plot of the learning curves for the runs in question can be found in Appendix . As one can see form these results, RTRRL performs on-par with PPO with BPTT, which used batched experiences and handpicked best hyperparameters.

Deep Exploration. Exploration versus exploitation is a trade-off, central to all agents learning a task online. The DeepSea environment included in bsuite (Osband et al. 2020) is tailored to benchmark the ability of RL algorithms to explore in unfavourable environments. At each step, the agent has to decide on one of two actions which randomly map to left and right in a way that is fixed at each position. A small negative reward is given when the action leading right was chosen and a large positive reward is given at the end of each episode, if the agent chose right each time. The agent is required to explore to reach the rightmost position albeit receiving negative rewards when moving towards it. We give results for exponentially increasing task length.

The classic control environment Acrobot, too, requires extensive exploration. A double pendulum starting in hanging position is set into motion by controlling the middle

Algorithm 1: Real-Time Recurrent Reinforcement Learning

Require: Linear policy: $\pi_{\theta_A}(a|h)$ **Require:** Linear value function: $\hat{v}_{\theta_C}(h)$ **Require:** CT-RNN body: $\text{RNN}_{\theta_R}([o, a, r], h, J)$ $\theta_A, \theta_C, \theta_R \leftarrow \text{Randomly initialize parameters}$ $B_A, B_C \leftarrow$ Randomly initialize feedback matrices $h, e_A, e_C, e_R \leftarrow \mathbf{0}$ $o \leftarrow \text{Reset Environment}$ $h, J \leftarrow \text{RNN}_{\theta_R}([o, \mathbf{0}, 0], h, \mathbf{0})$ \triangleright RFLO to compute \hat{J} $v \leftarrow \theta_C^+ h$ while not done do $\pi \leftarrow \pi_{\theta_A}(h)$ $a \leftarrow \text{Sample(dist)}$ $o, R \leftarrow \text{Take action } a \text{ in Environment}$ $h', \hat{J}' \leftarrow \text{RNN}_{\theta_R}[o, a, r], h, \hat{J})$ $g_C \leftarrow B_C \mathbf{1}$ b feedback alignment $g_A \leftarrow B_A \nabla_{\mathbf{h}} [\ln \pi[a] + \eta_H H(\pi)]$ $e_{C} \leftarrow \gamma \lambda_{C} e_{C} + \alpha h - \alpha \gamma \lambda \left[e_{C}^{\top} h \right] h$ $e_A \leftarrow \gamma \lambda_A e_A + \nabla_{\theta_A} [\ln \pi[a] + \eta_H H(\pi)]$ $e_R \leftarrow \gamma \lambda_R e_R + \hat{J}(g_A + g_C)$ $v' \leftarrow \theta_C^\top h'$ $\delta \leftarrow R + \gamma v' - v$ ⊳ TD error $\theta_{C} \leftarrow \theta_{C} + \delta e_{C} + \alpha \left[v - \theta_{C}^{\top} h \right] h$ $\theta_A \leftarrow \theta_A + \alpha_A \delta e_A$ $\begin{array}{l} \theta_R \leftarrow \theta_R + \alpha_R \delta e_R \\ v \leftarrow v', \quad h \leftarrow h', \quad \hat{J} \leftarrow \hat{J'} \end{array}$ end while

joint. The outer segment has to be elevated above a certain height, located above the anchor of the inner segment. Figure 3 shows the median rewards of 5 runs over the number of environment steps. We find that RTRRL with RFLO finds a solution significantly faster than the other methods hinting at superior exploration of our RTRRL approach.

Memory Length. We tested the memory capacity of our RTRRL models by learning to remember the state of a bit for an extended number of steps. The MemoryChain environment can be thought of as a T-maze, known from behavioural experiments with mice. Here, a long corridor leads up to a T-section. Each mouse is given a queue at the entrance to the corridor, indicating on which path leaving the T-section it will find food (Osband et al. 2020).

The experiment tests if the model can remember the state of a bit for a fixed number of time steps. Increasing the number of steps increases difficulty. We conducted MemoryChain experiments with 32 neurons, again with exponentially increasing task length. A boxplot of the results is shown in figure 4. With this experiments we wanted to answer the question of how RTRL, RFLO, and BPTT compare to each-other. The results show that using approximate gradients hampers somewhat memory capacity. Quite surprisingly, RTRRL trained with RTRL, outperformed PPO with LSTMs of the same size that were trained with BPTT.

Related Work

Ladosz et al. (2022) extended DQNs by adding a neural net-



Figure 3: Median reward of 5 runs on Acrobot. The key to solving this environment is exploration. RNNs trained with RFLO solve it quicker than when trained with BPTT.

work head called MOHN, that is subject to Hebbian learning using eligibility traces. Their algorithm significantly outperforms vanilla DQN on confounding POMDPs although slowing down execution time by 66%. However, this work uses back-propagation in its optimization engine.

Ororbia and Mali (2022) introduced a biologically plausible model-based RL framework called active predicitve coding that enables online learning of tasks with extremely sparse feedback. Their algorithm combines predictive coding and active inference, two concepts grounded in neuroscience. Network parameters are trained with Hebbian updates and surrogate losses. This is quite different from our work.

The work of Johard and Ruffaldi (2014) uses a cascadecorrelation algorithm with two eligibility traces, that starts from a small neural network and grows it sequentially. This is a very interesting alternative to backpropagation, and it is claimed to be biologically plausible. The two eligibility traces approximate the gradient of the discounted reward, in a connectionist actor-critic setting. The method outperforms regular policy-gradient methods on the CartPole environment. However, it only considers feed-forward networks.

In his Master's Thesis, Chung (2019) introduces a network architecture similar to RTRRL that consists of a recurrent backbone and linear TD heads. Convergence for the RL algorithm is proven assuming the learning rate of the RNN being magnitudes below those of the heads. Albeit the similarity in network structure, gradients were nonetheless computed using biologically implausible BPTT.

An interesting approach to reduce the complexity of RTRL was proposed by Javed et al. (2023). Similar to the cascade-correlation approach used in Johard and Ruffaldi (2014), they train a recurrent neural network constructively, that is one neuron at a time, keeping preceding neuron's parameters fixed. This way the RTRL complexity is reduced to the one of BPTT. However, this work does not consider RL.

Recently, Irie, Gopalakrishnan, and Schmidhuber (2023) investigated the practical performance of RTRL-based recurrent RL on a set of memory tasks in 3D. They used LSTMs modified to allow for efficient RTRL updates, and were able to show an improvement over training with BPTT when used in the framework of IMPALA (Espeholt et al. 2018). The



Figure 4: Boxplots of the Rewards of 10 runs on MemoryChain per type of plasticity for increasing memory lengths. BPTT refers to PPO with LSTM, RFLO and RTRL denote the variants of RTRRL and LocalMSE is a naive approximation to RTRL.

main aim of this work was thus to improving RTRL in the context of LSTMs.

A similar idea was explored by Zucchet et al. (2023), who show that independent modules of complex-valued linear recurrent units (LRUs) allow for efficient computation of RTRL updates. Their architecture can be generalized to multi-layer networks setting it apart from other related work. However, their model renounces biological plausibility by restricting the recurrent connections to be diagonal.

Finally, a great number of recent publications deal with training recurrent networks of spiking neurons (RSNNs). (Bellec et al. 2020; Taherkhani et al. 2020; Pan et al. 2023) The different approaches to train RSNNs in a biologically plausible manner do mostly rely on discrete spike events, for example in spike-time dependent plasticity (STDP). The e-prop algorithm introduced by Bellec et al. (2020) stands out as the most similar to RFLO. It features the same computational complexity and has been shown to be capable of solving RL tasks, albeit only for discrete action spaces.

Discussion

We introduced real-time recurrent reinforcement learning (RTRRL), a new approach to solving discrete and continuous control tasks for POMDPs, in a biologically plausible fashion. RTRRL consists of three parts: 1) A Meta-RL RNN architecture, implementing on its own a nested RL algorithm, 2) A TD(λ) actor-critic algorithm exploiting temporal difference and dutch eligibility traces, to train the weights of the Meta-RL network, and 3) A biologically plausible, random-feedback local-online (RFLO) real-time optimization algorithm, for computing the weight gradients.

We compared RTRRL with popular, but biologically implausible RL algorithms that use BPTT or RTRL for gradient-computation. Our results show, that using approximate gradients as in RFLO, still finds satisfactory solutions, and in some cases, it even improves on the state of the art. Particularly, RTRRL outperformed PPO with BPTT, when solving tasks requiring exploration in unfavourable environments.

Although the results presented in this paper are empirically convincing, some limitations have to be discussed. The algorithm generally suffers from a larger variance due to an effective batch size of 1. A possible improvement of RTRRL might be achieved, by using batched experience replay, which has become a standard in RL. However, this has to be done in a way, that preserves biological plausibility.

RTRRL is grounded in neuroscience and adequately explains how biological neural networks learn to act in unknown environments. The network structure resembles the interplay of dorsal and ventral striatum of the basal ganglia, featuring global RPEs found in dopaminergic pathways projecting from the ventral tegmental area and the substantia nigra zona compacta to the striatum and cortex (Wang et al. 2018). The role of dopamine as RPE was established experimentally by Wise (2004) who showed that dopamine is release upon receiving an unexpected reward, reinforcing the recent behavior. However, dopamine is also released in response to a conditioned stimulus, before receiving the actual reward. If then the expected reward is absent, dopamine levels will drop below baseline - a negative reinforcement signal. Dopaminergic synapses are usually located at the dendritic stems of glutamate synapses (Kandel 2013) and can therefore effectively mediate synaptic plasticity.

More specifically, the ventral striatum would correspond to the critic in RTRRL and the dorsal striatum to the actor, with dopamine axons targeting both the ventral and dorsal compartmens (Sutton and Barto 2018). The axonal tree of dopaminergic synapses is represented by the backward weights in RTRRL. Dopamine subsequently encodes the TD-error as RPE, which is used to update both the actor and the critic. RTRRL can therefore be seen as a model of reward-based learning taking place in the human brain.

References

Bartunov, S.; Santoro, A.; Richards, B. A.; Marris, L.; Hinton, G. E.; and Lillicrap, T. 2018. Assessing the Scalability of Biologically-Motivated Deep Learning Algorithms and Architectures. *arXiv:1807.04587 [cs, stat]*.

Behrens, T. E. J.; Muller, T. H.; Whittington, J. C. R.; Mark, S.; Baram, A. B.; Stachenfeld, K. L.; and Kurth-Nelson, Z. 2018. What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*, 100(2): 490–509.

Bellec, G.; Scherr, F.; Subramoney, A.; Hajek, E.; Salaj, D.;

Legenstein, R.; and Maass, W. 2020. A Solution to the Learning Dilemma for Recurrent Networks of Spiking Neurons. *Nature Communications*, 11(1): 3625.

Chung, W. 2019. Two-Timescale Networks for Nonlinear Value Function Approximation. https://era.library.ualberta.ca/items/d0665ad5-a222-4911a20e-72d4f7916821.

Daley, B.; and Amato, C. 2020. Reconciling \$\lambda\$-Returns with Experience Replay. 1810.09967.

Espeholt, L.; Soyer, H.; Munos, R.; Simonyan, K.; Mnih, V.; Ward, T.; Doron, Y.; Firoiu, V.; Harley, T.; Dunning, I.; Legg, S.; and Kavukcuoglu, K. 2018. IMPALA: Scalable Distributed Deep-RL with Importance Weighted Actor-Learner Architectures. 1802.01561.

Freeman, C. D.; Frey, E.; Raichuk, A.; Girgin, S.; Mordatch, I.; and Bachem, O. 2021. Brax – A Differentiable Physics Engine for Large Scale Rigid Body Simulation.

Funahashi, K.-i.; and Nakamura, Y. 1993. Approximation of Dynamical Systems by Continuous Time Recurrent Neural Networks. *Neural Networks*, 6(6): 801–806.

Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Comput.*, 9(8): 1735–1780.

Irie, K.; Gopalakrishnan, A.; and Schmidhuber, J. 2023. Exploring the Promise and Limits of Real-Time Recurrent Learning. 2305.19044.

Javed, K.; Shah, H.; Sutton, R.; and White, M. 2023. Scalable Real-Time Recurrent Learning Using Sparse Connections and Selective Learning. arxiv:2302.05326.

Johard, L.; and Ruffaldi, E. 2014. A Connectionist Actor-Critic Algorithm for Faster Learning and Biological Plausibility. In 2014 IEEE International Conference on Robotics and Automation (ICRA), 3903–3909. Hong Kong, China: IEEE. ISBN 978-1-4799-3685-4.

Kandel, E. R. 2013. *Principles of Neural Science*. New York; Toronto: McGraw-Hill Medical. ISBN 978-0-07-139011-8.

Ladosz, P.; Ben-Iwhiwhu, E.; Dick, J.; Ketz, N.; Kolouri, S.; Krichmar, J. L.; Pilly, P. K.; and Soltoggio, A. 2022. Deep Reinforcement Learning With Modulated Hebbian Plus Q-Network Architecture. *IEEE Transactions on Neural Networks and Learning Systems*, 33(5): 2045–2056.

Lange, R. T. 2022. gymnax: A JAX-based Reinforcement Learning Environment Library.

Lillicrap, T. P.; Cownden, D.; Tweed, D. B.; and Akerman, C. J. 2014. Random Feedback Weights Support Learning in Deep Neural Networks. *arXiv:1411.0247 [cs, q-bio]*.

Lillicrap, T. P.; and Santoro, A. 2019. Backpropagation through Time and the Brain. *Current Opinion in Neurobiology*, 55: 82–89.

Lu, C.; Kuba, J.; Letcher, A.; Metz, L.; Schroeder de Witt, C.; and Foerster, J. 2022. Discovered Policy Optimisation. *Advances in Neural Information Processing Systems*, 35: 16455–16468.

Marschall, O.; Cho, K.; and Savin, C. 2020. A Unified Framework of Online Learning Algorithms for Training Recurrent Neural Networks. *Journal of Machine Learning Research*, 21(135): 1–34.

Menick, J.; Elsen, E.; Evci, U.; Osindero, S.; Simonyan, K.; and Graves, A. 2020. Practical Real Time Recurrent Learning with a Sparse Approximation. In *International Conference on Learning Representations*.

Miconi, T.; Clune, J.; and Stanley, K. O. 2018. Differentiable Plasticity: Training Plastic Neural Networks with Backpropagation.

Mujika, A.; Meier, F.; and Steger, A. 2018. Approximating Real-Time Recurrent Learning with Random Kronecker Factors. arxiv:1805.10842.

Murray, J. M. 2019. Local Online Learning in Recurrent Networks with Random Feedback. *eLife*, 8: e43299.

Ni, T.; Eysenbach, B.; and Salakhutdinov, R. 2022. Recurrent Model-Free RL Can Be a Strong Baseline for Many POMDPs. arxiv:2110.05038.

Ororbia, A.; and Mali, A. 2022. Active Predicting Coding: Brain-Inspired Reinforcement Learning for Sparse Reward Robotic Control Problems. arxiv:2209.09174.

Osband, I.; Doron, Y.; Hessel, M.; Aslanides, J.; Sezener, E.; Saraiva, A.; McKinney, K.; Lattimore, T.; Szepesvari, C.; Singh, S.; Roy, B. V.; Sutton, R.; Silver, D.; and Hasselt, H. V. 2020. Behaviour Suite for Reinforcement Learning.

Pan, W.; Zhao, F.; Zeng, Y.; and Han, B. 2023. Adaptive Structure Evolution and Biologically Plausible Synaptic Plasticity for Recurrent Spiking Neural Networks. *Scientific Reports*, 13(1): 16924.

Roth, C.; Kanitscheider, I.; and Fiete, I. 2018. Kernel RNN Learning (KeRNL). In *International Conference on Learning Representations*.

Rusu, S. I.; and Pennartz, C. M. A. ???? Learning, Memory and Consolidation Mechanisms for Behavioral Control in Hierarchically Organized Cortico-Basal Ganglia Systems. 30.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347.

Seijen, H.; and Sutton, R. 2014. True Online TD(Lambda). In *Proceedings of the 31st International Conference on Machine Learning*, 692–700. PMLR.

Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. A Bradford Book. ISBN 0-262-03924-9.

Taherkhani, A.; Belatreche, A.; Li, Y.; Cosma, G.; Maguire, L. P.; and McGinnity, T. M. 2020. A Review of Learning in Biologically Plausible Spiking Neural Networks. *Neural Networks*, 122: 253–272.

Tallec, C.; and Ollivier, Y. 2017. Unbiased Online Recurrent Optimization. arxiv:1702.05043.

Wang, J. X.; Kurth-Nelson, Z.; Kumaran, D.; Tirumala, D.; Soyer, H.; Leibo, J. Z.; Hassabis, D.; and Botvinick, M. 2018. Prefrontal Cortex as a Meta-Reinforcement Learning System. *Nature Neuroscience*, 21(6): 860–868.

Werbos, P. 1990. Backpropagation through Time: What It Does and How to Do It. *Proceedings of the IEEE*, 78(10): 1550–1560.

Williams, R. J.; and Zipser, D. 1989. A Learning Algorithm for Continually Running Fully Recurrent Neural Networks. *Neural Computation*, 1(2): 270–280.

Wise, R. A. 2004. Dopamine, Learning and Motivation. *Nature Reviews Neuroscience*, 5(6): 483–494.

Zenke, F.; and Neftci, E. O. 2021. Brain-Inspired Learning on Neuromorphic Substrates. *Proceedings of the IEEE*, 109(5): 935–950.

Zucchet, N.; Meier, R.; Schug, S.; Mujika, A.; and Sacramento, J. 2023. Online Learning of Long-Range Dependencies. arxiv:2305.15947.

Derivation of update equations

Consider a CT-RNN that has N hidden states and I inputs, activation φ and a combined weight matrix $W \in \mathbb{R}^{N \times X}$ where Z = I + N + 1. Each neuron has a time-constant $\tau \in \mathbb{R}^N$ and the next state $h_{t+1} \in \mathbb{R}^N$ is computed as follows:

$$h_{t+1} = h_t + \frac{1}{\tau} \left(-h_t + \varphi(W\xi_t) \right) \qquad \xi_t = \begin{bmatrix} x_t \\ h_t \\ 1 \end{bmatrix} \in \mathbb{R}^Z$$

where x_t is the input at time t, and 1 concatenated to ξ_t accounts for the bias. The output $\hat{y}_t \in \mathbb{R}^O$ is given by a linear mapping $\hat{y}_t = W_{out}h_t$. The latent space follows the ODE $\tau \dot{h} = -h_t + \varphi(W\xi_t)$.

RFLO leverages the state-update expression in order to simplify the RTRL update. For this we expand the gradient of f in equation 5. Note that previous work on RFLO kept the time constant τ fixed and trained the recurrent weights W only, hence the restricted \hat{J}^W in the equation:

$$\hat{J}_{t+1}^{W} = \frac{d}{dW}h_t + \frac{d}{dW}f(x_t, h_t) \\
= \frac{d}{dW}h_t + \frac{d}{dW}\frac{1}{\tau}\left(-h_t + \varphi(W\xi_t)\right) \\
= \frac{d}{dW}h_t\left(1 - \frac{1}{\tau}\right) + \frac{d}{dW}\frac{1}{\tau}\varphi(W\xi_t) \qquad (10) \\
= \left(1 - \frac{1}{\tau}\right)\hat{J}_t^W + \frac{1}{\tau}\nabla_W\varphi(W\xi_t) \\
+ \frac{1}{\tau}\hat{J}_t^W\nabla_{h_t}\varphi(W\xi_t)$$

In order to achieve biological plausibility, RFLO boldly drops the last summand, since it requires horizontal gradient communication. The partial derivative of the activation is simply $\nabla_W \varphi(W\xi_t) = \varphi'(W\xi_t)^{\top} \xi_t$ where φ' is the pointwise derivative of the activation function φ :

$$\hat{J}_{t+1}^{W} \approx (1 - \frac{1}{\tau}) \, \hat{J}_{t}^{W} + \frac{1}{\tau} \varphi'(W\xi_{t})^{\top} \xi_{t}$$
 (11)

We analogously derive the RFLO update for τ . Again we drop the communicated gradients $\hat{J}_t^{\tau} \nabla_{h_t} \varphi(W\xi_t)$ and arrive at the expression for \hat{J}^{τ} .

$$\hat{J}_{t+1}^{\tau} = \hat{J}_{t}^{\tau} - \frac{d}{d\tau} \frac{1}{\tau} h_{t} + \frac{1}{\tau} \varphi(W\xi_{t}) \\
= \hat{J}_{t}^{\tau} - \nabla_{\tau} \frac{1}{\tau} h_{t} - \hat{J}_{t}^{\tau} \nabla_{h_{t}} \frac{1}{\tau} h_{t} + \frac{d}{d\tau} \frac{1}{\tau} \varphi(W\xi_{t}) \\
= \hat{J}_{t}^{\tau} (1 - \frac{1}{\tau}) + \frac{1}{\tau^{2}} h_{t} + \nabla_{\tau} \frac{1}{\tau} \varphi(W\xi_{t}) \\
+ \hat{J}_{t}^{\tau} \nabla_{h_{t}} \frac{1}{\tau} \varphi(W\xi_{t}) \\
= \hat{J}_{t}^{\tau} (1 - \frac{1}{\tau}) + \frac{1}{\tau^{2}} h_{t} - \frac{1}{\tau^{2}} \varphi(W\xi_{t}) \\
+ \hat{J}_{t}^{\tau} \nabla_{h_{t}} \frac{1}{\tau} \varphi(W\xi_{t})$$
(12)

Implementation Details

Upon acceptance, we will publish our code on GitHub. Logging of experiments is implemented for Aim¹ or Weights & Biases² as backend. Our implementation is highly configurable and allows for many tweaks to the base algorithm. Available options include gradient clipping, learning rate decay, epsilon greedy policy, delayed RNN parameter updates and many more. Please refer to table 2 and the Readme.md in the code folder for a list of configurables. Table 2 summarizes the hyperparameters of RTRRL. We kept them at the listed default values for all our experiments.

Algorithm 2 repeats RTRRL with if-cases for RTRL without feedback alignment for demonstrative purposes. The algorithm can be divided into 4 distinct steps that are depicted in figure 5. When developing the algorithm, we had to figure out the proper order of operations. Figure 6 is a flowchart that was created to help understand at what point the eligibility traces are combined with the TD-error and approximate Jacobian, to form the parameter updates.

Since not using batched experiences, our algorithm unsurprisingly suffers from large variance and in some cases catastrophic forgetting ensues. Using an exponentially decaying learning rate for the RNN can help ins such cases but for simplicity we chose to make our experiments without this fix. Since a biologically plausible agent should retain its capability of reacting to shifts in the environment, the adaptability of our algorithm would be tainted as the learning rate of the RNN approaches 0. Nonetheless, RTRRL most of the time converges to an optimal solution without the use of decaying learning rates.

Neuron Model. The simplified CT-RNN outlined in the paper is taken from (Murray 2019). Our code allows for increasing the number of steps $k = dt^{-1}$ when solving the underlying Ordinary Differential Equation with the forward Euler method. More steps lead to a more expressive model meaning you can get away with fewer neurons, but also to increased computational complexity. In our experiments we kept k = 1 for simplicity.

¹https://aimstack.readthedocs.io/

²httpbs://wandb.ai/



Figure 5: RTRRL can be devided into 4 parts that are repeated throughout training.



Figure 6: Flowchart depicting the computations done throughout one RTRRL step.

More experiments

Memory Length

We tested the memory capacity of our RTRRL models by learning to remember the state of a bit for an extended number of steps. The experiment is explained in section . Figure 4 shows the rewards of 10 runs each as boxplots. We also included a naive approximation to RTRL we call *LocalMSE* that only uses the immediate Jacobian for updates. Looking at the boxplot we can clearly see that RTRL outperforms the other methods and that RFLO is situated at a middle ground between LocalMSE and RTRL in terms of memory capacity.

Description	Symbol	Value
number of neurons	n	32
discount factor	γ	0.99
Actor learning rate	α_A	1e-2
Critic learning rate	α_C	1.0
RNN learning rate	α_R	1e-3
entropy rate	η_H	1e-5
Actor eligibility decay	λ_A	0.9
Critic eligibility decay	λ_C	0.9
RNN eligibility decay	λ_R	0.9
patience in epochs	20	
maximum environment	50 mil.	
optimizer	SGD	
batch size	1	
learning rate decay	0	
action epsilon	0	
update period	1	
gradient norm clip	1.0	
normalize observations	False	

Table 2: Hyperparameters of RTRRL and values used.



Figure 7: CT-RNNs are solving an Ordinary Differential Equation. In general, any solver may be used. When using forward Euler, $dt = k^{-1}$ is a hyperparameter that determines the number of solver steps and subsequently the accuracy of the solution.

Physics Simulation

For this set of experiments we created POMDPs with continuous actions by masking the observations of three different brax environments. For each environment, we ran hyperparameter tuning for at least 10 hours and picked the best performing run. Figure 8 shows the evaluation rewards achieved throughout the learning process for RTRRL and the PPO baseline that was provided by the package.

Biological implausibility of BPTT

Here, we elaborate further on the objections to biological plausibility of Backpropagation-Through Time.

Fully online. The computation of updates should not depend on alternating forward and backward phases. Conceivably, such phases could be implemented using pacemaker neurons. However, there must not be any freezing of values, which occurs in BPTT. Furthermore, parallel streams of ex-

Algorithm 2: Real-Time Recurrent Reinforcement Learning with accumulate eligibility traces and RTRL

Require: Linear policy: $\pi(a|h, \theta_A)$ **Require:** Linear value function: $\hat{v}(h, \theta_C)$ **Require:** CT-RNN body: RNN($[o, a, r], h, \hat{J}, \theta_R$) $\theta_A, \theta_C, \theta_R \leftarrow \text{Randomly initialize parameters}$ if RFLO then $B_A, B_C \leftarrow$ Randomly initialize feedback matrices end if $h, e_A, e_C, e_R \leftarrow \mathbf{0}$ $o \leftarrow \text{Reset Environment}$ $h, \hat{J} \leftarrow \text{RNN}([o, \mathbf{0}, 0], h, \theta_R)$ $v \leftarrow \hat{v}(h, \theta_C)$ while not done do logits $\leftarrow \pi(h, \theta_A)$ $a \leftarrow \text{Sample(logits)}$ $o, R \leftarrow$ Take action a in Environment $h', \hat{J}' \leftarrow \text{RNN}([o, a, r], h, \hat{J}, \theta_R)$ $v' \leftarrow \hat{v}(h', \theta_C)$ $\delta \leftarrow R + \gamma v' - v$ $e_C \leftarrow \gamma \lambda_C e_C + \nabla_{\theta_C} \hat{v}$ $e_A \leftarrow \gamma \lambda_A e_A + \nabla_{\theta_A} [\ln \pi[a] + \eta_H \operatorname{Ent}(\pi)]$ if RTRL then $e_R \leftarrow \gamma \lambda_R e_R + \nabla_h [\hat{v} + \ln \pi [a] + \eta_H \operatorname{Ent}(\pi)]$ else if RFLO then $g_C \leftarrow B_C \mathbf{1}$ $g_A \leftarrow B_A \nabla_{\text{logits}} [\ln \pi[a] + \eta_H \text{Ent}(\pi)]$ $e_R \leftarrow \gamma \lambda_R e_R + g_C + g_A$ end if $\theta_C \leftarrow \theta_C + \alpha_C \delta e_C$ $\theta_A \leftarrow \theta_A + \alpha_A \delta e_A$ $\begin{array}{l} \theta_{R} \leftarrow \theta_{R} + \alpha_{R} \hat{J} \delta e_{R} \\ v \leftarrow v', \quad h \leftarrow h', \quad \hat{J} \leftarrow \hat{J}' \end{array}$ end while

perience, such as in batched environments of modern DRLs, violate this constraint, as a biological agent can only interact with the singular environment in which it is situated.

No weight transport. Synapses that propagate back error signals cannot have their strength mirroring the strength of forward synapses. This is heavily violated by backpropagation since backward pathways need access to, and utilize, parameters that were used in the forward pass.

No horizontal gradient communication. Individual neurons x_k are very likely unable to communicate exact activation gradients $\nabla_{w_{ij}} x_k(t)$, with respect to the synaptic parameters w_{ij} , to other neurons that are not in their vicinity. In other words, intermediate (source) neurons should not be able to tell other (target) neurons how each of the synaptic weights of the source influence the dynamics of the target, nor vice-versa.

Biological Interpretation of RTRRL

Real-Time Recurrent Reinforcement Learning is a model of goal-directed behaviour learning and training of reflexes in animals, located in the human basal ganglia (Rusu and Pennartz). It comprises a scalar-valued global reward signal that



Figure 8: Evaluation rewards for environments from the brax package that where masked as described in section to make them POMDPs. Shown are the best run for RTRRL and the tuned baseline shipped with the package.

can be interpreted as dopaminergic synapses, a value estimator (critic) and motor output (actor).

Dopamine and Learning

Dopamine is a neurotransmitter found in the central nervous system of mammals. It is widely agreed upon that dopamine plays an important role in rewards and reinforcement. The molecule is released by specialized neurons located primarily in two areas of the brain: the substantia nigra zona compacta (SNc) and the ventral tegmental area (VTA). (Wise 2004) These neurons have large branching axonal arbors making synapses with many other neurons - mostly in the striatum and the pre-frontal cortex. Those synapses are usually located at the dendritic stems of glutamate synapses and can therefore effectively influence synaptic plasticity. (Sutton and Barto 2018)

Whenever receiving an unexpected reward, dopamine is released which subsequently reinforces the behavior that led to the reward. However, if the reward is preceded by a conditioned stimulus, it is released as soon as the stimulus occurs instead of when receiving the expected reward. If then the expected reward is absent, levels will drop below baseline representing a negative reinforcement signal. (Wise 2004) Various experiments have shown that dopamine is crucial for stamping in response-reward and stimulus-reward associations which in turn is needed for motivation when confronted with the same task in the future. Particularly, moderate doses of dopamine antagonists (neuroleptics) given to a live animal will reduce motivation to act. Habitual responses decline progressively in animals that are treated with neuroleptics. (Wise 2004)

Ternary Synapses

Hetero-synaptic plasticity is a type of synaptic plasticity that involves at least three different neurons. Usually, a sensory neuron is forming a synapse with a motor neuron. A third so called *facilitating* neuron also forms a synapse at the same spot and so is able to influence the signal transmission between the sensory and the motor neuron. A historic example is found in the gill-withdrawal reflex circuitry of Aplysia (Kandel 2013). The RPE signalling, facilitating inter-neuron strengthens the motor response (withdrawing the gill) when a negative reward (shock) is experienced or expected. RTRRL could be implemented in biological neural networks with ternary synapses where the facilitating synapses are projecting back from the TD-error computing inter-neuron.

Brain Structures

There is ample evidence that certain brain structures are implementing actor-critic methods. The striatum is involved in motor and action planning, decision-making, motivation, reinforcement and reward perception. It is also heavily innervated by dopamine axons coming from the VTA and the SNc. It is speculated that dopamine release in the ventral striatum "energizes the next response" while it acts by stamping in the procedural memory trace in the dorsal striatum "establishing and maintaining procedural habit structures" (Wise 2004). Subsequently, the ventral striatum would correspond to the critic and the dorsal striatum to the actor of an RL algorithm (Sutton and Barto 2018). Dopamine would then correspond to the TD-error which is used to update both the actor and the critic. As dopamine neuron axons target both the ventral and dorsal striatum, and dopamine appears to be critical for synaptic plasticity, the similarities are evident. Furthermore, the TD-error and dopamine levels are both encoding the RPE: they are high whenever an unexpected reward is received and they are low (or negative in case of the TD-error) when an expected reward does not occur. (Sutton and Barto 2018) These similarities could be beneficial for RL as well as for Neuroscience as advances in either field could lead to new insights that are beneficial to the other.



Figure 9: The basal ganglia are located at the base of the forebrain and play a major role in motivation and behavioural learning. Source: Wikimedia, Leevanjackson³

Timing comparison

In order to compare wall-clock time, we did a quick performance test for RTRRL with CT-RNN and PPO with LSTM. In both cases we trained for 5000 steps, with 32 units and a batch size of 1, on a machine with a single GeForce RTX 2070 GPU. For both cases, we repeated the test 3 times and calculated the average time per step of the algorithm. Our results were 7,58 ms / step for PPO-LSTM and 7,49 ms / step for RTRRL-CTRNN. Please note that these results may not carry over to larger model or batch sizes.