

Gamma-Nets: Generalizing Value Estimation over Timescale

Craig Sherstan,¹ Shibhansh Dohare,¹ James MacGlashan,²
 Johannes Günther,¹ Patrick M. Pilarski¹

¹Department of Computing Science, University of Alberta
 Edmonton, Alberta, Canada

²Cogitai, USA

{sherstan, pilarski}@ualberta.ca

Abstract

Temporal abstraction is a key requirement for agents making decisions over long time horizons—a fundamental challenge in reinforcement learning. There are many reasons why value estimates at multiple timescales might be useful; recent work has shown that value estimates at different time scales can be the basis for creating more advanced discounting functions and for driving representation learning. Further, predictions at many different timescales serve to broaden an agent’s model of its environment. One predictive approach of interest within an online learning setting is general value function (GVFs), which represent models of an agent’s world as a collection of predictive questions each defined by a policy, a signal to be predicted, and a prediction timescale. In this paper we present Γ -nets, a method for generalizing value function estimation over timescale, allowing a given GVF to be trained and queried for arbitrary timescales so as to greatly increase the predictive ability and scalability of a GVF-based model. The key to our approach is to use timescale as one of the value estimator’s inputs. As a result, the prediction target for any timescale is available at every timestep and we are free to train on any number of timescales. We first provide two demonstrations by 1) predicting a square wave and 2) predicting sensorimotor signals on a robot arm using a linear function approximator. Next, we empirically evaluate Γ -nets in the deep reinforcement learning setting using policy evaluation on a set of Atari video games. Our results show that Γ -nets can be effective for predicting arbitrary timescales, with only a small cost in accuracy as compared to learning estimators for fixed timescales. Γ -nets provide a method for accurately and compactly making predictions at many timescales without requiring a priori knowledge of the task, making it a valuable contribution to ongoing work on model-based planning, representation learning, and lifelong learning algorithms.

Value Functions and Timescale

Reinforcement learning (RL) studies algorithms in which an agent learns to maximize the amount of reward it receives over its lifetime. A key method in RL is the estimation of *value*—the expected cumulative sum of discounted future rewards (called the *return*). In loose terms this tells an agent how good it is to be in a particular state. The agent can then

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

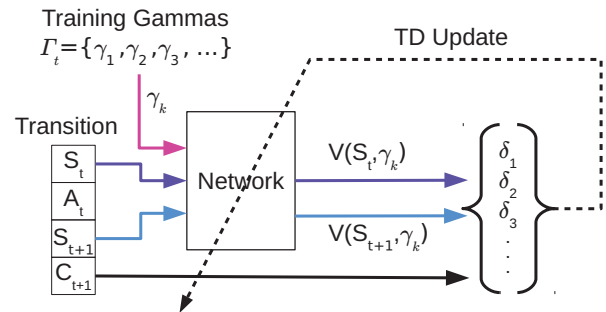


Figure 1: Training γ -nets. Values are estimated by providing state and timescale, γ , as inputs to the network parameterized by weights \mathbf{w} . An agent in state S takes action A and transitions to state S' receiving the new target signal C . The agent selects a set of timescales Γ_t on which to train and for each $\gamma_k \in \Gamma_t$ computes values $V(S, \gamma_k; \mathbf{w})$ and $V(S', \gamma_k; \mathbf{w})$. For each γ_k , the TD error is calculated according to $\delta_k = C + \gamma_k V(S', \gamma_k; \mathbf{w}) - V(S, \gamma_k; \mathbf{w})$. The TD errors are then collected and used to update \mathbf{w} using a chosen TD learning algorithm, such as TD(λ) or GTD.

use value estimates to learn a *policy*—a way of behaving—which maximizes the amount of reward received.

Sutton et al. (2011) broadened the use of value estimation by introducing general value functions (GVFs), in which value estimates are made of other sensorimotor signals, not just reward. GVFs can be thought of as representing an agent’s model of itself and its environment as a collection of questions about future sensorimotor returns; a predictive representation of state (Dayan 1993). A GVF is defined by three elements: 1) the policy, 2) the *cumulant* (the sensorimotor signal to be predicted), and 3) the prediction timescale, γ . Considering a simple mobile robot, examples of GVF questions include “How much current will my motors consume over the next 3 seconds if I spin clockwise?” or “How long until my bump sensor goes high if I drive forward?”

Modeling the world at many timescales is seen as a key problem in artificial intelligence (Sutton 1995; Sutton, Precup, and Singh 1999). Further, there is evidence that humans and other animals make estimates of reward and other sig-

nals at numerous timescales (Tanaka et al. 2016). This paper focuses on generalizing value estimation over timescale. Our work can be seen as directly connected to the concept of *nexting*, in which animals and people make large numbers of predictions of sensory input at many, short-term, timescales (Gilbert 2006). Modayil, White, and Sutton (2014) demonstrated the concept of nexting using GVF’s on a mobile robot. Until now, value estimation has generally been limited to a single fixed timescale. That is, for each desired timescale, a discrete and unique predictor was learned. However, there are situations where we may desire to have value estimates of the same cumulant over many different timescales. For example, consider an agent driving a car. Such an agent may make numerous predictions about the likelihood of colliding with various objects in its vicinity. The agent needs to consider the risk of collisions in both the near term and far term and the relevance of each may change with the speed of the car. If the engineer knew which timescales would be needed ahead of time they could design them into the system, but this is not the case for complex settings.

Here we present a class of algorithms which enables the explicit learning and inference of value estimates for any valid fixed discount. The key insights to our approach are: 1) the timescale can be treated as an input parameter for inference and learning and 2) the estimated bootstrapped prediction target for any fixed timescale is available at every timestep. We demonstrate Γ -nets in the policy evaluation setting: 1) predicting a square wave, 2) predicting sensorimotor signals on a robot arm, 3) predicting reward in Atari video games.

The ideas behind our approach are based on work by Schaul et al. (2015) which generalized value estimation across goals by providing a goal embedding vector as input to the value network. In contrast, our approach provides the discount, γ as input. Xu, van Hasselt, and Silver (2018) also provide γ as input to their value and policy networks. They present a meta-learning approach which learns the best γ to provide to an inner policy. Here we focus on determining what is necessary to effectively train a value network to train over timescale. Additionally, our algorithm trains on multiple timescales simultaneously.

Background

We model the environment as a Markov Decision Process. At each timestep t the agent, in state $S_t \in \mathcal{S}$, takes action $A_t \in \mathcal{A}$ according to policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ and transitions to state $S_{t+1} \in \mathcal{S}$ according to the transition probability $p(\cdot | S_t, A_t)$. In the traditional RL setting the agent receives a reward $R_{t+1} \equiv R(S_t, A_t, S_{t+1}) \in \mathbb{R}$. The agent tries to learn a policy which maximizes the cumulative reward it receives in the future, which is defined as the return: $G_t = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots$. In the case of GVF’s we simply substitute our signal of interest, the cumulant, C for reward, R . The term $\gamma \in [0, 1]$ is referred to by several names including the timescale, the continuation function and the discount; it represents the amount of emphasis applied to future rewards and is the focus of this paper.

A value estimate is simply the expectation of the return: $V_\pi(s) = \mathbb{E}_\pi[G_t | S_t = s]$. Temporal difference (TD) learning is a common class of algorithms used in RL for learning an

approximation of value (Sutton and Barto 1998). Estimation weights are typically trained by semi-gradient descent using the TD error: $\delta_t = C_{t+1} + \gamma V(S_{t+1}) - V(S_t)$.

While simple domains can be represented using tabular lookup, complex settings in which the state space is very large or infinite must use function approximation (FA) methods to estimate the value as $V(s; \mathbf{w})$, where \mathbf{w} is a set of weights parameterizing the network. Function approximation has the advantage that states are not treated independently, but rather, a learning step updates related states as well, allowing for generalization across state-space.

Generalizing over Timescale

Our goal is to predict the value function for any discount factor γ . While the GVF specification allows for γ that are a function of the transition, here we focus solely on the case of fixed timescale. To achieve that goal, we propose Γ -nets: an architecture for value functions that operates not only on the state, but also the desired target discount factor γ_k (see Figure 1). On each transition the network is trained on many $\gamma_k \in \Gamma_t$ values. Thus, the Γ -net learns to generalize over arbitrary γ_k values. Generating the error function for a γ -net is straightforward. For any single $\gamma_k \in \Gamma_t$, the TD error is:

$$\delta_{t;\gamma_k} = C_{t+1} + \gamma_k V(S_{t+1}, \gamma_k) - V(S_t, \gamma_k). \quad (1)$$

The total gradient can then be summed over all $\gamma_k \in \Gamma$ and applied to update the network.

Choosing Γ_t must be done with care. A naive approach might uniformly sample $\gamma_k \in [0, 1]$. However, value functions change non-linearly with γ . To illustrate this property, consider that γ can be viewed as the probability of continuation, allowing us to derive the expected number of timesteps (ts) until termination of the return as:

$$\tau = \frac{1}{1 - \gamma}. \quad (2)$$

This relationship is non-linear for large values of γ . Thus, naively drawing γ_k from a uniform distribution would tend to favor very short timescales. Conversely, drawing uniformly from τ would put little emphasis on short timescales. While the best method for selecting γ_k for training is outside the scope of this paper, we provide some comparisons in our experiments. Note that throughout this paper we will refer broadly to the word *timescale* for which we will use the parameters γ or τ as appropriate. It should be assumed that these terms can be used interchangeably using Eq. (2).

The representation of timescale used for input to the network may affect the network’s ability to represent different timescales. The γ scale compresses long timescales but spreads short ones and in the τ scale we have the opposite effect. Thus, providing both γ and τ as input may allow for good discrimination at all timescales.

Finally, the magnitude of returns at different timescales can be very different. To prevent large magnitude returns from dominating the network weights we need to scale the returns in some way. A general approach is given by van Hasselt et al. (2016), in which they continually normalize the target to have a mean of 0 and variance of

1. This allows them to handle rewards of varying magnitude. Here, we instead focus on keeping the magnitude of the returns, as a function of timescale, in the same ballpark. To do this we learn the value of the scaled cumulant, $f(s, \gamma_k; \mathbf{w}) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma_k^t (1 - \gamma_k) C_{t+1} | S_0 = s]$ and then define the unscaled value as $V(s, \gamma_k; \mathbf{w}) = \frac{f(s, \gamma_k; \mathbf{w})}{(1 - \gamma_k)}$. With these definitions we can then simply scale the TD error as:

$$\delta_{t; \gamma_k} = (1 - \gamma_k)(C_{t+1} + \gamma_k V(S_{t+1}, \gamma_k) - V(S_t, \gamma_k)).$$

Experiments

We first provide two demonstrations using linear function approximation (LFA): 1) predicting a square wave signal, 2) predicting joint current on a robot arm. Finally, we empirically evaluate Γ -nets in a deep learning setting by looking at performance on Atari games. Additional results and experimental details are available from Sherstan et al. (2019).

Square Wave

Our target signal was a repeating square wave 100 timesteps in length with a magnitude of $\{-1, 1\}$ (Figure 2). Inputs were normalized and then tilecoded (Sutton and Barto 1998) with 20 tilings of width 1.0, 20 tilings of width 0.5 and 30 tilings of width 0.1. Tiling positions were randomly shifted by small amounts at the time of initialization for each run. Value estimates were computed using LFA on the output of the tilecoding and the final layer of weights was updated using TD(0) (Sutton and Barto 1998). We also evaluated the impact of loss scaling. Results are shown in Figure 3. Unless otherwise stated: 1) timescale inputs were given on both the γ and τ scales simultaneously, 2) Γ_t was 6 elements long, with $\tau \in \{1, 100\}$ always included and two additional timescales drawn uniformly from each of the γ and τ timescales, 3) loss scaling was used.

Predictions on a Robot Arm

In this experiment a human operated the shoulder rotation and elbow flexion joints of a robot arm by joystick. The task was to maintain contact between a rod held by the robot and the inside of a wire maze while moving in a counter-clockwise direction (Figure 4a). Fifty circuits of the maze were completed in approximately 12 minutes. Network inputs were the normalized positions of the shoulder and elbow servos as well as both γ and normalized τ . Inputs were tilecoded (Sutton and Barto 1998) with 100 tilings of width 1.0 into a space of 2048 bits and a bias unit was added giving a feature vector of 2049 bits. Value estimates were computed by LFA and trained by TD(0). On each timestep Γ_t was generated from $\tau \in [1, 100]$ ts. The upper and lower bounds were included in the set and one γ and 29 τ were sampled uniformly from their respective scales for a total of 32 timescales. We receive updates from the robot at 30 Hz (1 ts \sim 0.03), because of this we expect longer timescales to be more important, thus we focus sampling on the τ scale. Loss prescaling was used. A baseline predictor was also trained on a fixed timescale using the same configuration excepting the inclusion of timescale input. Figure 4b shows predictions for shoulder joint speed. The Γ -net matches the true return well. In general, we found the Γ -net performed better than the baseline using this configuration.

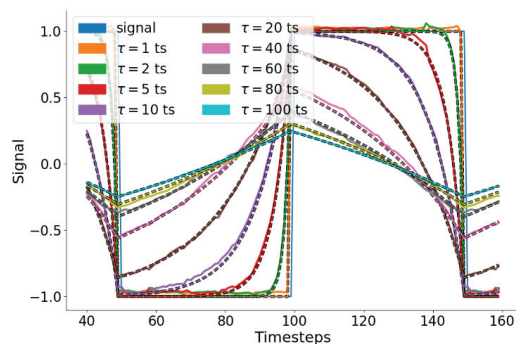


Figure 2: Square-wave Predictions. Predictions (solid) against the true return (dashed) after 50k ts of training using both γ and τ as input, scaling the loss and drawing two γ_k each from γ and τ scales plus $\tau = 1, 100$. For display purposes all predictions are normalized by $(1 - \gamma)$. We see good accuracy across all timescales.

Atari Environment

We examined the performance of Γ -nets under policy evaluation in the Arcade Learning Environment (ALE) (Bellemare et al. 2015). The agent’s policy was trained using the Dopamine project’s (Castro et al. 2018) implementation of the Rainbow agent (Hessel et al. 2018), which uses the same network architecture as the DQN agent (Mnih et al. 2015), but adds prioritized replay (Schaul et al. 2016), n-step returns, and distributional representation of the value estimates (Bellemare, Dabney, and Munos 2017). The results presented in the main body of the paper are for the game Centipede with a Rainbow agent trained for 25 million frames, which we will refer to as *Centipede@25M*. Additional Atari games can be found in Sherstan et al. (2019).

Figure 5 shows predictions on the early transitions of a single episode. For this episode the expected return was estimated by running 2000 Monte Carlo rollouts from each state visited along the way. The solid lines indicate the Γ -net predictions after training for 20 M frames (using the *direct* configuration which will be described in following sections).

Training The prediction networks were trained using transitions generated by pretrained policies. Agents select actions using ϵ -greedy over their Q-values. During policy training $\epsilon = 0.001$, but for generating the samples used for training the Γ -nets we use an evaluation mode where $\epsilon = 0.0001$. Transitions were generated sequentially and the environment was reset at the end of each episode or 27,000 steps, whichever came first. Transitions were saved to file in sequence and for each experiment they were reloaded in the same order. For each transition, we saved the reward and the activation of the final core layer of the agent’s network ϕ , which serves as the input to the Γ -nets. The Γ -net network consisted of five fully-connected layers of sizes [512, 256, 128, 16, 1], with all but the final layer using ReLU activation. Training of the Γ -nets proceeded as if the data was generated in an online fashion. That is, the agent would read in transition samples from the file, add them to a prioritized replay buffer, and then train by sampling from the replay buffer.

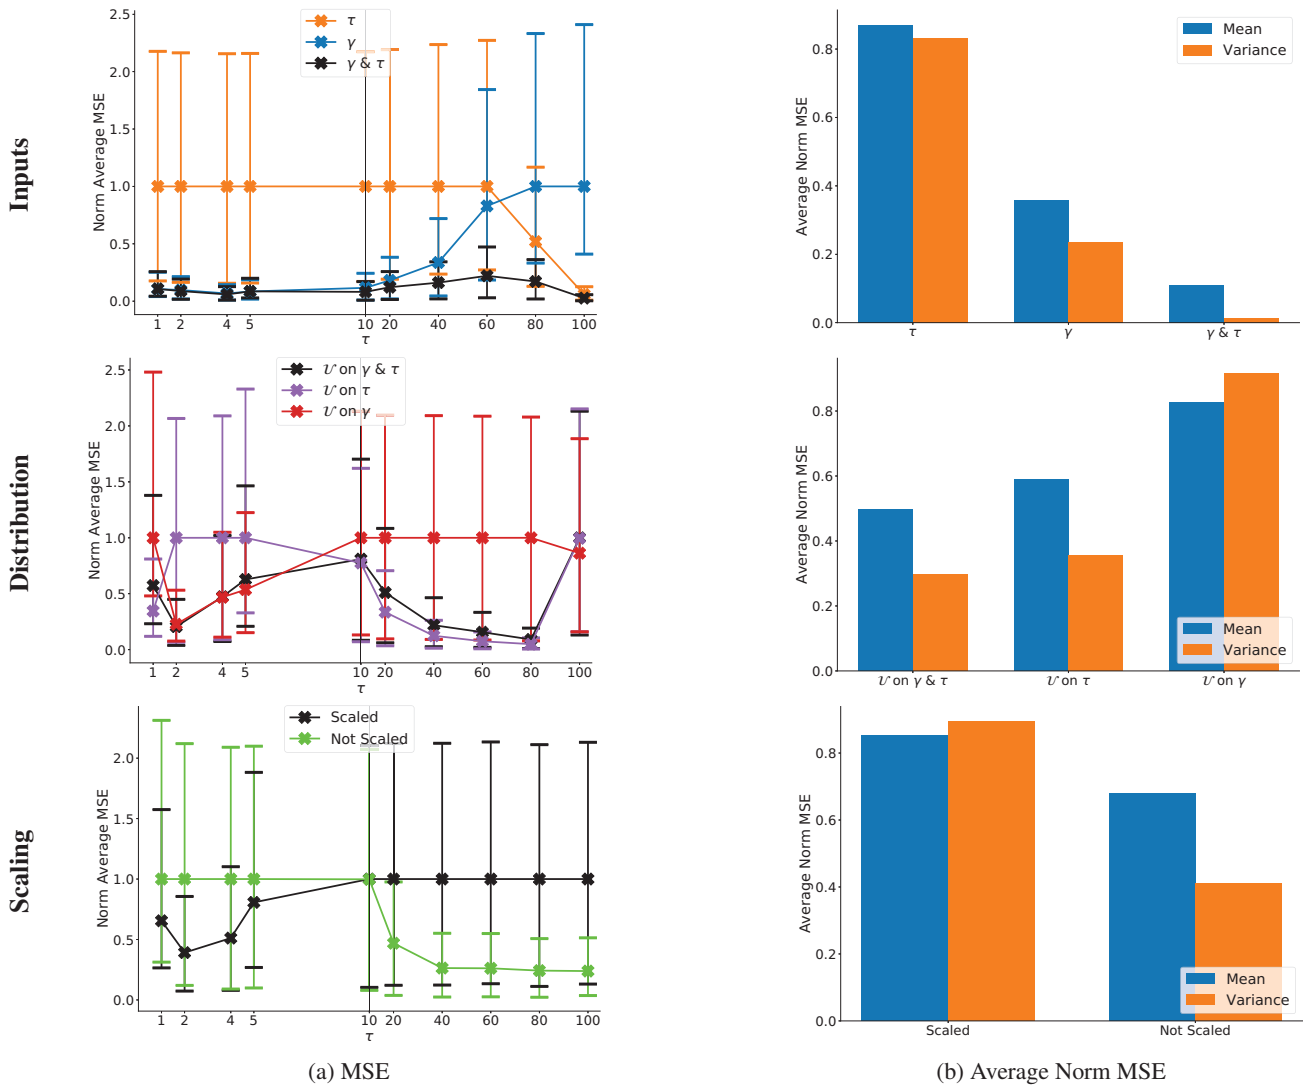



Figure 3: Square Wave. Each training run lasted for 50k timesteps and for each series 100 different runs were made. We show the normalized errors as a function of the prediction timescale, given on the τ -scale. Results are averaged over the last 5k timesteps. For each τ we normalize by the maximum mean error across the series in the plot. **Inputs)** Comparing the effect of using different timescale representations as input. As expected, providing timescale as γ did better than τ on the short timescales, but worse on the longer timescales, although this cross over occurred at a much longer timescale than expected. Providing both γ and τ did the best of all, producing the lowest errors across all probe timescales as well as providing the lowest variability. **Distribution)** Here we compare the effects of drawing γ_k from different distributions. Γ_t was 6 elements long, always including $\tau \in \{0, 100\}$, and sampling the 4 additional timescales. Excluding $\tau \in \{0, 100\}$ we see that drawing all γ_k from the γ scale performs better than drawing all from τ scale at shorter time scales, but does worse at longer timescales. Drawing half from each tends to follow the lower errors at all the timescales. **Scaling)** We compare the effects of scaling the cumulant. Here we see that scaling does improve performance on the shorter timescales, but causes worse performance on the longer ones.

Newly added samples were given the highest level of priority so that its probability of being sampled was high. Like the policy training we train on a batch of sampled transitions, using n-step returns. To update the priorities for a given sample in the batch we use the maximum squared loss across Γ_t .


A Γ_t of size 8 was used, which always included lower and upper bounds of $\tau = [1, 100]$. An additional 6 γ_k were drawn on each timestep. Unless otherwise stated the sampling was

done by drawing 3 timescales uniformly each from the γ scale on $[0, 0.99)$ and the τ scale on $[1, 100]$ (for τ we drew from the integer scales). Each network was trained for 20 M frames with network weights saved every 500k frames. Additional details can be found in Sherstan et al. (2019).

Evaluation To evaluate predictive accuracy we created a set of evaluation points for each game. These were gener-



(a)



(b)

Figure 4: Robot Arm. **a)** A user controls the shoulder and elbow joints of a robot arm via joystick to move a rod counter-clockwise around a wire maze. **b)** Predictions of the speed of the shoulder joint. In this snapshot the Γ -net (orange) matches the true return (dashed blue) very well, surprisingly better than a baseline predictor (green) trained just for this timescale. Predictions and returns are scaled by $(1 - \gamma)$ for display purposes.




Figure 5: Predictions on Centipede@25M for different γ from the start of a single episode. Γ -net predictions are shown in solid lines and the expected return, produced by Monte Carlo rollout, is shown by the dashed lines.

ated by running the agent in evaluation mode over multiple episodes. At the start of each episode an offset was randomly chosen between $[10, 100)$ steps. Then, starting at the offset, the state of the environment and agent were saved every 30 steps. For Centipede@25M a total of 269 evaluation points were created including the episode start state. From each evaluation points we ran 1000 episodes till termination and then computed the average return. These were used as the baseline against which we computed our prediction error. To compute the prediction error for a given evaluation point we restored the agent’s and environment’s state and recorded the network’s predictions for the probe timescales $[1, 2, 5, 10, 20, 40, 60, 80, 100]$ ts. For comparison, we trained fixed timescale networks for each probe timescale (plotted in fuschia). These networks used the same architecture as the Γ -net, but did not provide timescale as input to the network and only trained on the single fixed timescale. These probe networks also used loss scaling.

We use a reference configuration of the Γ -net across the different plots. We plot this series in black and refer to it as

direct although the legends may give it a different label to call out the significance of its configuration for a given comparison. For this configuration both γ and τ were provided as inputs to the network. Additionally, Γ_t was populated by drawing samples from both γ and τ scales and loss scaling was used. For each of the other configurations only a single setting was modified from this reference.

Plotting We focus our evaluation on the steady-state performance of the network, computing averages over the last 5 M frames of the 20 M frame runs (with evaluation at every 500k frames). Mean-squared error (MSE) for each experiment is presented as a function of the evaluation probe timescale given in τ (Ex. Figure 6a). For each τ we normalize across the different series by the largest mean error. Thus, the largest mean error for each τ is shown as 1.0. We do this to be able to clearly show results for all the different timescales in a single plot despite the large differences in magnitude. As a result, series can only be directly compared within a plot, not across plots. To rank each for comparison we provide a bar chart (Ex. Figure 6b) which averages the normalized means and normalized variances of the MSE. That is, we take the normalized mean MSE for each τ and average across all τ . Likewise we take the variance at each τ , normalize it by the maximum variance for each τ and take the average across all τ . Note that averaging this way is a biased approach in that it is dependent on what probe τ are used. For example, if we took many large τ and few small ones then our results would give more weight to the large τ . In practice, the weighting of errors for different timescales will be task dependent.

While conducting parameter sweeps it was observed that a particular network configuration might produce the lowest value of MSE but not actually be predictive. In this case the network would learn to always output a fixed value which captured the mean of the expected returns. Thus, we adopted a two step evaluation process. First, we took the evaluation points and concatenated them in sequence. We then computed the correlation between their expected returns and the predicted returns made by the network. If a configuration

had a positive correlation then it would be considered for comparison with other architectures. We have also included the plots of correlation by probe timescale (Ex. Figure 6c). Correlation values are easily interpreted with the maximum (best) value of 1. This tells us how closely the shapes of the target sequence and the prediction sequence match.

All series are an average over 6 seeds and the shading indicates max and min values. Note, that due to the high degree of overlap in many of the figures, color printing is required to discern individual series. Plots taken with respect to τ are produced by combining two different x-axes, allowing us to make both short and long timescales discernible. This split occurs at $\tau = 10$ and is indicated by the vertical black line.

While our evaluation method seeks to discern differences in performance due to the various configuration, in reality most configurations perform similarly. In order to rank configurations we first considered the MSE and then variance.

Embedding Comparison. We compare methods for combining the timescale inputs with the agent’s features, ϕ , using an embedding vector $\nu = \nu(\phi, \gamma)$ (Figure 6). The *direct* embedding performs a concatenation, $\nu = [\phi, \gamma]$. Xu, van Hasselt, and Silver (2018) learned a vector, $\xi(\gamma)$, of size 16 which was concatenated with ϕ , which we refer to as *L_embed*. We also considered a Hadamard embedding in which a learned vector, $\xi(\gamma)$, the same length as ϕ , was combined using element-wise multiplication with ϕ , that is $\nu = \phi \odot \xi$ (*h.L_embed*). Finally, we considered a matrix multiplication approach in which the timescales were given as inputs to a fully connected layer whose output was a square matrix, $\Xi(\gamma)$, with dimensions the same size as ϕ . The embedding was then formed by matrix multiplication: $\nu = \phi^T \Xi$. We found little difference between the approaches in terms of their MSE or correlation. Overall the linear embedding appears the best choice based on its lower variance, but this did not hold universally for all games evaluated. Learning and computation were both slower with the matrix multiplication approach and linear activations were generally slightly better than ReLU. See Sherstan et al. (2019) for additional results.

Timescale Input Comparison We examine how the input timescale representation affects prediction performance (Figure 7). We consider whether to use γ or τ inputs or both. The γ input values are naturally scaled between $[0, 1)$ and the τ input values were normalized by dividing by the max τ , which in these experiments was 100. We consistently see that using only γ produced the worst performance (Asteroids, is an exception (Sherstan et al. 2019)). Providing τ or both τ and γ performed very similarly, but we consistently observed that providing both representations performed best for very short timescales and had lower variance.

Distribution Comparison We look at the effect of drawing Γ_t from different distributions (Figure 8). We use a Γ_t of size 8, two of which are always the lower and upper bounds $\tau = [1, 100]$. Six additional γ_k are drawn from a given distribution. We either draw all six uniformly from the γ or τ scale or draw half from each. We see that drawing solely from γ performs worst overall, particularly at longer timescales, as is expected. Surprisingly, γ did not consistently outperform τ at

very short timescales. If we consider all timescales and games evaluated there is no clear winner between drawing solely from τ or from τ and γ . However, at very short timescales drawing from both tended to produce better results. Thus, we recommend drawing from both scales as a default.

Loss Scaling We examined the effect of loss scaling. Figure 9 shows that on Centipede@25M there is a clear benefit, with clearly lower MSE and variance. Scaling the loss was expected to improve short timescale performance. Surprisingly, in terms of MSE, the greatest impact was on longer timescales. However, such a pronounced difference was not seen in other Atari games (Sherstan et al. 2019). Instead we saw a general trend in which scaling did improve performance at short scales at the cost of performance at mid and long timescales, which was in line with our expectations (again, Asteroids was an exception).

Estimation by Interpolation An alternative approach to estimating value at arbitrary timescales is to have multiple prediction heads, each at a fixed timescale, and then linearly interpolate between the nearest bracketing timescales. In Figure 10 we show results for such an interpolation. Here we took the previously trained probe networks (with scaled loss and the taper network architecture) and performed linear interpolation for $\tau = [1.5, 3.5, 7.5, 15, 30, 50, 70, 90]$. Because of the non-linear relationship between τ and γ the linear interpolation gives different weighting depending on whether the interpolation is done on the τ or γ scale. Interpolating in these spaces is also compared. Results show that performance was fairly similar between the interpolation scales, but that the Γ -net did not perform as well. While it might have been expected that the ability of the neural network used by Γ -net to capture the non-linearity of the timescales would give it an advantage, this was not shown in this experiment. Rather, we suspect that the increased accuracy of the probe networks allowed the interpolation approach to win out.

Discussion

We have empirically evaluated various approaches to constructing Γ -nets and compared their predictive accuracy to baseline predictors. While we sought to separate the impacts of the various approaches, in reality all of the variants we explored performed similarly. We have considered several different Atari games with deep learning architectures as well as a simulation signal and robotics demonstration using a shallow architecture. Overall we found that Γ -nets worked reliably both for reward and sensorimotor prediction.

Despite the relatively minor differences in performance across the variants we do make some recommendations for implementation. Since there was no universal difference between the *direct* or *L_embed* embedding approaches we recommend just using the simplest, *direct*. If looking for a general approach that is not specifically adapted to the task then we recommend using both γ and τ as inputs to the network as well as drawing samples from both scales in order to populate Γ_t . On the other hand if longer timescales are preferred then it seems sufficient to use only τ for both input and sampling distribution. With regards to scaling the loss a clear universal




Figure 6: Embedding Comparison. Several approaches for adding the timescale dependency to the network were investigated. **direct**) Concatenates the timescale with ϕ . **l_embed**) Timescale is input to a fully connected layer of length 16 with linear activation whose output is concatenated with ϕ . **h_l_embed**) Timescale is input to a fully connected layer the same length as ϕ with linear activation and then combined with ϕ by element-wise multiplication. The *l_embed* approach appears to be slightly better due to its lower variance.




Figure 7: Input Comparison. We compare performance of the network when providing γ , τ or both to the network inputs. We see that providing only γ as the input timescale does the worst. Providing both γ and τ or just τ perform similarly, but providing both does better at very short timescales.




Figure 8: Distribution Comparison. We compare different distributions used to generate Γ_t . At lower timescales sampling from the γ scale does better than sampling from the τ scale and the opposite holds at longer timescales. Sampling from both provides a compromise in performance.

benefit has not been observed and we suggest that further investigation is required to determine the best way to balance the losses resulting from different timescales. Such an investigation is a clear opportunity for future work.

Our method is thus far limited to the fixed discounting case. However, one of the key generalizations of GVF is to support transition-dependent discounting functions: $\gamma_{t+1} \equiv \gamma(S_t, A_t, S_{t+1})$ (White 2017). This allows GVFs to be more expressive in terms of what the types of returns they can

estimate. Extending our method to support such discounting is clearly an important next-step in this work.

There are several ways in which our work and that of Fedus et al. (2019) are complementary. First, they demonstrated that using value predictions at many different timescales could serve as useful auxiliary tasks for driving representation learning. A clear next step is to investigate whether or not Γ -nets could also serve as a useful auxiliary task. Additionally, they demonstrated that alternative returns, such as hyperbolic




Figure 9: **Scaling Comparison.** We examined the effects that scaling the loss by $(1 - \gamma_k)$ has. We see that scaling results in lower overall error and variance. Note that such a clear separation was not observed over other games tested.




Figure 10: Interpolated. Predictions are made between the probe timescales by taking the weighted average of the predictions made by the bracketing probe networks. Due to the non-linear relationships of γ and τ different weightings are produced if the weighting is done in either scale. We see that either interpolation produces better results than the Γ -nets.

discounting, could be formed by using value estimates at multiple timescales as a basis function; Γ -nets could provide such a basis function using a single network.

Long timescale predictions can be difficult to learn due to the higher variance of the returns. Romoff et al. (2019) presented an algorithm which computes values for an ordered set of timescales by predicting the differences between the values using separate network heads. Value estimates are constructed in a cascade where each timescale prediction adds to the one that came before it. They showed their method could improve estimation accuracy for longer timescales by leveraging the accuracy of the easier to learn shorter timescales. We might expect a similar effect using Γ -nets where long timescale predictions could benefit from the short timescales being learned directly in the network. Our current evaluation approach is not fine grained enough to discern such a benefit. Thus, this area warrants further exploration.

Γ -nets is related to other works which seek to learn many different predictions simultaneously and tractably. The UVFA (Schaul et al. 2015), on which this work is based, generalizes over goals. The successor representation (SR) (Dayan 1993) separates environment dynamics from reward, providing a way to transfer learning across tasks (Barreto et al. 2017; Sherstan, Machado, and Pilarski 2018). These ideas have been combined (Mankowitz et al. 2018; Ma, Wen, and Bengio 2018) to enable transfer learning over multiple goals using off-policy learning. However, these methods still use fixed timescales, thus, a natural extension of Γ -nets is to combine

them with these approaches.

The original motivation for this work was to use Γ -nets to create GVFs which form a predictive representations of state for use by the agent’s policy. It now seems that the best approach would be to use multiple heads with predictions at fixed timescales and let the policy network learn to generalize over those predictions as it needed. Such an approach could be costly in terms of network weights and Γ -nets might accomplish the same thing with a smaller network.

Conclusion

We presented Γ -nets, a simple technique for generalizing value estimation across timescale. This technique allows a system to make predictions for values of any timescale within the training regime of the network. We expect that this ability will be useful in areas such as predictive representations of state—i.e., modeling the world as a collection of predictions about future sensorimotor signals. In complex environments complete models are not feasible, thus, being able to query for predicted outcomes at any timescale makes a model potentially more compact and expressive. An investigation of Γ -nets in different control learning scenarios is an important area for future work, and we believe they may be of benefit to ongoing research in planning and lifelong learning. In particular Γ -nets are complementary to approaches which seek to learn many things about the world simultaneously such as the successor representation and universal value functions,

suggesting that Γ -nets may provide us with a functional new tool for the pursuit of knowledgeable intelligent systems.

Acknowledgements

The authors would like to thank the following colleagues for providing thoughtful suggestions to this work: Alex Kearney, Marlos C. Machado, and Matt Schlegel. Additionally, Brendan Bennett, Jesse Farebrother and Vivek Veeriah provided helpful technical assistance. Initial stages of this work were funded by Cogitai and additional support was provided by the Natural Sciences and Engineering Research Council of Canada (NSERC), Compute Canada, the Canada Research Chair’s program, Alberta Innovates, and the Alberta Machine Intelligence Institute (Amii).

References

- Barreto, A.; Dabney, W.; Munos, R.; Hunt, J.; Schaul, T.; Silver, D.; and van Hasselt, H. 2017. Successor Features for Transfer in Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 4055–4065.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2015. The Arcade Learning Environment: An Evaluation Platform for General Agents. In *International Joint Conference on Artificial Intelligence*, 4148–4152.
- Bellemare, M. G.; Dabney, W.; and Munos, R. 2017. A Distributional Perspective on Reinforcement Learning. In *International Conference on Machine Learning (ICML)*, 449–458.
- Castro, P. S.; Moitra, S.; Gelada, C.; Kumar, S.; and Bellemare, M. G. 2018. Dopamine: A Research Framework for Deep Reinforcement Learning. *arXiv* 1812.06110.
- Dayan, P. 1993. Improving Generalization for Temporal Difference Learning: The Successor Representation. *Neural Computation* 5(4):613–624.
- Fedus, W.; Gelada, C.; Bengio, Y.; Bellemare, M. G.; and Larochelle, H. 2019. Hyperbolic Discounting and Learning over Multiple Horizons. *arXiv* 1902.06865.
- Gilbert, D. 2006. *Stumbling on Happiness*. Knopf.
- Hessel, M.; Modayil, J.; van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2018. Rainbow: Combining Improvements in Deep Reinforcement Learning. In *AAAI Conference on Artificial Intelligence*.
- Ma, C.; Wen, J.; and Bengio, Y. 2018. Universal Successor Representations for Transfer Reinforcement Learning. In *International Conference on Learning Representations (ICLR)*.
- Mankowitz, D. J.; Žídek, A.; Barreto, A.; Horgan, D.; Hessel, M.; Quan, J.; Oh, J.; van Hasselt, H.; Silver, D.; and Schaul, T. 2018. Unicorn: Continual Learning with a Universal, Off-policy Agent. *arXiv* 1802.08294.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. a.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; Petersen, S.; Beattie, C.; Sadik, A.; Antonoglou, I.; King, H.; Kumaran, D.; Wierstra, D.; Legg, S.; and Hassabis, D. 2015. Human-level Control through Deep Reinforcement Learning. *Nature* 518(7540):529–533.
- Modayil, J.; White, A.; and Sutton, R. S. 2014. Multi-Timescale Nexting in a Reinforcement Learning Robot. *Adaptive Behavior* 22(2):146–160.
- Romoff, J.; Henderson, P.; Touati, A.; Ollivier, Y.; Brunskill, E.; and Pineau, J. 2019. Separating Value Functions Across Time-scales. *arXiv* 1902.01883.
- Schaul, T.; Horgan, D.; Gregor, K.; and Silver, D. 2015. Universal Value Function Approximators. *International Conference on Machine Learning (ICML)* 1312–1320.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2016. Prioritized Experience Replay. *arXiv* 1511.05952.
- Sherstan, C.; Dohare, S.; MacGlashan, J.; Günther, J.; and Pilarski, P. M. 2019. Gamma-Nets: Generalizing Value Estimation Over Timescale. *arXiv* 1911.07794.
- Sherstan, C.; Machado, M. C.; and Pilarski, P. M. 2018. Accelerating Learning in Constructive Predictive Frameworks with the Successor Representation. In *IEEE International Conference on Robots and Systems (IROS)*, 2997–3003.
- Sutton, R. S., and Barto, A. G. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
- Sutton, R. S.; Modayil, J.; Delp, M.; Degris, T.; Pilarski, P. M.; White, A.; and Precup, D. 2011. Horde: A Scalable Real-time Architecture for Learning Knowledge from Unsupervised Sensorimotor Interaction. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 2, 761–768.
- Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between MDPs and Semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. *Artificial Intelligence* 112(1):181–211.
- Sutton, R. S. 1995. TD Models: Modeling the World at a Mixture of Time Scales. In *International Conference on Machine Learning (ICML)*, 531–539.
- Tanaka, S. C.; Doya, K.; Okada, G.; Ueda, K.; Okamoto, Y.; and Yamawaki, S. 2016. Prediction of Immediate and Future Rewards Differentially Recruits Cortico-basal Ganglia Loops. *Behavioral Economics of Preferences, Choices, and Happiness* 7(8):593–616.
- van Hasselt, H.; Guez, A.; Hessel, M.; Mnih, V.; and Silver, D. 2016. Learning Values Across Many Orders of Magnitude. In *Advances in Neural Information Processing Systems (NeurIPS)*, 4287–4295.
- White, M. 2017. Unifying Task Specification in Reinforcement Learning. In *International Conference on Machine Learning (ICML)*, 3742–3750.
- Xu, Z.; van Hasselt, H.; and Silver, D. 2018. Meta-Gradient Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2396–2407.