

---

# Incrementally Added GVFs are Learned Faster with the Successor Representation

---

Craig Sherstan<sup>1</sup> Marlos C. Machado<sup>1</sup> Patrick M. Pilarski<sup>1 2</sup>

## Abstract

We propose using the *successor representation* (SR) to accelerate learning in a constructive knowledge system based on *general value functions* (GVFs). We consider an agent which incrementally constructs a predictive model of the world as a collection of GVFs. Using the SR provides a way of separating the dynamics of the world from the prediction targets and thus captures regularities that can be reused across multiple GVFs. The learning problem then consists of two parts—learning the SR and learning one-step predictions about a signal of interest. SR-based predictions are then produced by combining the SR with different one-step predictions. As a primary contribution of this work, we show in both a grid world and a robotic task that using SR-based predictions can improve sample efficiency and learning speed in a lifelong learning setting where new predictions are incrementally added and learned over time.

## 1. Introduction

A long standing goal in the pursuit of artificial general intelligence is that of knowledge construction—incrementally modeling and explaining the world and the agent’s interaction with it directly from the agent’s own experience. This is particularly important in fields such as lifelong learning (Thrun, 1995) and developmental robotics (Oudeyer et al., 2007), where we expect agents to be capable of learning, dynamically and incrementally, to interact and succeed in complex environments.

One proposed approach for representing such world models is a collection of *general value functions* (GVFs) (Sutton

et al., 2011), which models the world as a set of predictive questions each defined by a policy of interest, a target signal of interest under that policy, and a timescale (discounting schedule) for accumulating the signal of interest. For example, a GVF on a mobile robot could pose the question “How much current will my wheels consume over the next second if I drive straight forward?”

GVF questions are typically answered using temporal-difference (TD) methods from the field of reinforcement learning (RL) (Sutton & Barto, 1998). A learned GVF approximates the expected future value of a signal of interest, directly representing the relationship between the environment, policy, timescale, and target signal as the output of a single predictive unit.

Nevertheless, despite the success RL algorithms have achieved recently, methods for answering multiple predictive questions from a single stream of experience (critical in a robotic setting) are known to exhibit sample inefficiency. In our setting of interest, where multitudes of GVFs are learned in an incremental, sample by sample way, this problem is multiplied. Ultimately, the faster an agent can learn to approximate a new GVF, the better.

In this paper we show how one can accelerate learning in a constructive knowledge system based on GVFs by sharing the environment dynamics across the different predictors. This is done with the successor representation (SR) (Dayan, 1993), which allows us to learn the world dynamics under a policy independently of any signal being predicted.

We empirically demonstrate the effectiveness of our approach on both a tabular representation and on a robot arm which uses function approximation. We evaluate our algorithm in the lifelong learning setting where it is not possible to specify all GVFs *a priori*, but, rather, GVFs are added incrementally during the course of learning. As a key result, we show that using a learned SR enables an agent to learn newly added GVFs faster than when learning the same GVFs in the standard fashion without the use of the SR.

---

<sup>1</sup>Department of Computing Science, University of Alberta, Edmonton, Canada <sup>2</sup>DeepMind. Correspondence to: Craig Sherstan <sherstan@ualberta.ca>, Patrick M. Pilarski <pilarski@ualberta.ca>.

## 2. Background

We consider an agent interacting with the environment sequentially. We use the standard notation in the reinforcement learning (RL) literature (Sutton & Barto, 1998), modeling the problem as a Markov Decision Process. Starting from state  $S_0 \in \mathcal{S}$ , at each timestep the agent chooses an action,  $A_t \in \mathcal{A}$ , according to the policy distribution  $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ , and transitions to state  $S_{t+1} \in \mathcal{S}$  according to the probability transition  $p(\cdot | S_t, A_t)$ . For each transition  $S_t \xrightarrow{A_t} S_{t+1}$  the agent receives a reward,  $R_t$ , from the reward function  $R(S_t, A_t, S_{t+1}) \in \mathbb{R}$ .

In this paper we focus on the prediction problem in RL, in which the agent's goal is to predict the value of a signal from its current state (e.g., the cumulative sum of future rewards). Note that throughout this paper we use upper case letters to indicate random variables.

### 2.1. General Value Functions (GVFs)

The most common prediction made in RL is about the expected return. The return is defined to be the sum of future discounted rewards under policy  $\pi$  starting from state  $s$ . Formally,  $G_t = \sum_{t=0}^T \gamma^t R_t$ , with  $\gamma \in [0, 1]$  being the discount factor and  $T$  being the final timestep, where  $T = \infty$  denotes a continuing task. The function encoding the prediction about the return is known as the value function  $v_\pi(s) = \mathbb{E}_\pi[G_t | S_0 = s]$ .

GVFs (Sutton et al., 2011) extend the notion of predictions to different signals in the environment. This is done by replacing the reward signal,  $R_t$ , by any other target signal, which we refer to as the cumulant,  $C_t$ , and by allowing a state-dependent discounting function,  $\gamma_t = \gamma(S_t)$ , instead of using a fixed discounting factor. The general value of state  $s$  under policy  $\pi$  is defined as:

$$v_{\pi:\gamma}(s) = \bar{c}_\pi + \bar{\gamma}_\pi \sum_{s'} p_\pi(s'|s) v_\pi(s'), \quad (1)$$

where  $\bar{c}_\pi$  is the average cumulant from state  $s$ ,  $\bar{\gamma}_\pi$  is the average  $\gamma$  from state  $s$ , and  $p_\pi(s'|s)$  is the probability of transitioning from state  $s$  to  $s'$  under policy  $\pi$ . This can also be written in matrix form:  $\mathbf{v}_{\pi:\gamma} = \bar{\mathbf{c}}_\pi + \bar{\boldsymbol{\gamma}}_\pi \odot P_\pi \mathbf{v}_{\pi:\gamma}$ , where  $\odot$  denotes element-wise multiplication. Such an equation, when solved, gives us

$$\mathbf{v}_{\pi:\gamma} = (I - \bar{\boldsymbol{\gamma}}_\pi \odot P_\pi)^{-1} \bar{\mathbf{c}}_\pi, \quad (2)$$

where  $I$  is the identity matrix,  $\mathbf{v}, \bar{\mathbf{c}}, \bar{\boldsymbol{\gamma}} \in \mathbb{R}^{|\mathcal{S}| \times 1}$ , and  $P_\pi \in [0, 1]^{|\mathcal{S}| \times |\mathcal{S}|}$  is a probability matrix such that  $[P_\pi]_{ij} = \mathbb{E}[p_\pi(S_{t+1} = s_j | S_t = s_i)]$ .

### 2.2. The Successor Representation (SR)

The successor representation (Dayan, 1993) was initially proposed as a representation capable of capturing state sim-

ilarity in terms of time. It is formally defined, for a fixed  $\gamma < 1$ , as:

$$\psi_\pi(s, s') = \mathbb{E}_\pi \left[ \sum_{t=0}^{\infty} \gamma^t \mathbb{1}_{\{S_t = s'\}} | S_0 = s \right].$$

In words, the SR encodes the expected number of times the agent will visit a particular state when the sum is discounted by  $\gamma$  over time<sup>1</sup>. It can be re-written in matrix form as

$$\Psi_\pi = \sum_{t=0}^{\infty} (\gamma P_\pi)^t = (I - \gamma P_\pi)^{-1}. \quad (3)$$

Importantly, the SR can be easily computed, incrementally, by standard RL algorithms such as TD learning (Sutton & Barto, 1998), since its primary modification is to replace the reward signal by a state visitation counter. Nevertheless, despite its simplicity, the SR holds important properties we leverage in this paper.

The SR, in the limit, for a constant  $\gamma$  (see Eq. 3), corresponds to the first factor of the solution in Eq. 2. Thus, the SR can be seen as encoding the dynamics of the Markov chain induced by the policy  $\pi$  and by the environment's transition probability  $p$ . If the agent has access to the SR, it can accurately predict the (discounted) accumulated value of any signal, from any state, by simply learning the expected *immediate* value,  $\bar{c}_\pi$ , of that signal in each state. On the other hand, if the agent does not use the SR, the agent must also deal with the problem of credit assignment, having to look at  $n$ -step returns to control for *delayed* consequences. Importantly, the dynamics encoded by the SR are the same for all signals learned under the same policy and discount function. This factorization of the solution is the main property we use in our work, described in the next section.

## 3. Methods

As aforementioned, we are interested in the problem of knowledge acquisition in the lifelong learning setting, where knowledge is encoded as predictive questions (GVFs). In this setting it is not possible to specify all GVFs ahead of time. Instead, GVFs must be added incrementally by some, as yet unknown, mechanism. The standard approach would be to learn each newly added prediction from scratch. In this section we discuss how we can use the SR to accelerate learning by taking advantage of the factorization shown in Eq. 2. Our method leverages the fact that the SR is independent of the target signal being predicted, learning the SR separately and re-using it when learning to predict new signals.

<sup>1</sup>Note that Dayan describes the SR as predicting future state visitation from time  $t$  onward. This is non-standard in RL as we typically describe the return as predicting the signal from  $t + 1$  onward.

In the previous section, for clarity, we discussed the main concepts in the tabular case. In real world applications, where the state space is too large, assuming states can be uniquely identified is not often feasible. Instead, we generally represent states as a set of features  $\phi(s) \in \mathbb{R}^d$  where  $d \ll |\mathcal{S}|$ . Because both  $\psi$  and  $\bar{c}$  are a function of feature vector  $\phi(S)$ , they can easily be represented using function approximation and learned using TD algorithms. In order to present a more general version of our algorithm, we introduce it here using the function approximation notation.

The first step in our algorithm is to compute the one-step average cumulant, which we do with TD error:

$$\delta_t = C_{t+1} - \bar{c}(\phi(S_t)). \quad (4)$$

If we use linear function approximation to estimate  $\bar{c}$  then  $\bar{c}(\phi(s)) = \phi(s)^\top \mathbf{w}$ . The TD error for  $\psi$  is given as (note that  $\delta$  is a vector of length  $n$ )

$$\delta_t = \phi(S_t) + \gamma_{t+1}\psi(\phi(S_{t+1})) - \psi(\phi(S_t)). \quad (5)$$

This generalization of the SR to the function approximation case is known as successor features (Barreto et al., 2017).

If we use linear function approximation to estimate  $\psi$  then  $\psi(\phi(s)) = M^\top \phi(s)$ , where  $M \in \mathbb{R}^{d \times d}$ . Using the usual semi-gradient method,<sup>2</sup> often used with TD, we derive a TD(0) stochastic gradient descent update as:

$$\begin{aligned} M_{t+1} &= M_t - \frac{1}{2}\alpha \nabla_M (\delta_t^2) \\ &= M_t + \alpha \nabla_M (\psi(S_t)) \otimes \delta_t \\ &= M_t + \alpha \phi(S_t) \otimes \delta_t, \end{aligned}$$

where  $\nabla_M$  is the gradient with respect to  $M$  and  $\otimes$  is the outer product. Based on this derivation, as well as Eq. 4 and 5, we obtain Algorithm 1. Note that the last two lines of Algorithm 1, which update the SR for the state  $S'$ , are only required for the episodic case.

This algorithm allows us to predict the cumulant  $C_i$ , in state  $S$ , using the current estimate of the matrix  $M$  and the weights  $\mathbf{w}_i$ . We can then obtain the final prediction by simply computing  $\psi(\phi(S))^\top \mathbf{w} = (M^\top \phi(s))^\top \mathbf{w} = \phi(s)^\top M \mathbf{w}$ .

This algorithm accelerates learning because, generally, learning to estimate  $\bar{c}$  is faster than learning the GVF directly. This is exactly what our algorithm does. When predicting a new signal, it starts with its current estimate for the SR,  $\Psi$ . At the end, the multiplication  $\Psi \bar{c}$  is simply a weighted average of the one-step predictions across all states, weighted by the likelihood they will be visited. We

<sup>2</sup>In semi-gradient methods, the effect of  $M$  on the prediction target  $\phi(S_{t+1}) + \gamma\psi(\phi(S_{t+1}))$  is ignored when computing the gradient.

---

**Algorithm 1** GVF prediction with the SR
 

---

**Input:** Feature representation  $\phi$ , policy  $\pi$ , discount function  $\gamma$ , and step-sizes  $\alpha_C, \alpha_{\text{SR}}$

**Output:** Matrix  $M$  and vectors  $\bar{c}_i$  as predictors of  $C_i$

Initialize  $\mathbf{w}$  and  $M$  arbitrarily

**while**  $S'$  is not terminal **do**

    Observe state  $S$ , take action  $A$  selected according to  $\pi(S)$ , and observe a next state  $S'$  and the cumulants  $C_i$

$\delta_{\text{SR}} = \phi(S) + \gamma(S')M^\top \phi(S') - M^\top \phi(S)$

$M \leftarrow M + \alpha_{\text{SR}} \phi(S) \otimes \delta_{\text{SR}}$

**for each** cumulant  $C_i$  **do**

$\delta_{C_i} = C_i - \phi(S)^\top \mathbf{w}_i$

$\mathbf{w}_i \leftarrow \mathbf{w}_i + \alpha_{C_i} \phi(S) \delta_{C_i}$

**end for**

**end while**

$\delta_{\text{SR}} = \phi(S') - M^\top \phi(S')$

$M \leftarrow M + \alpha_{\text{SR}} \phi(S') \otimes \delta_{\text{SR}}$

---

provide empirical evidence supporting this claim in the next sections.

## 4. Evaluation in Dayan's Grid World

We first evaluated our algorithm in a tabular grid world. Its simplicity allowed us to analyze our method more thoroughly since we were not bounded by the speed and complexity of physical robots. The grid world we used was inspired by Dayan's (Dayan, 1993) (see Figure 1a). Four actions are available in this environment: up, down, left, right. Taking an action into a wall (blue) results in no change in position. Transitions are deterministic, i.e., a move in any direction moves the agent to the next cell in the given direction, except when moving into a wall. For each episode the agent spawns at location  $\mathbf{S}$  and the episode terminates when the agent reaches the goal  $\mathbf{G}$ .

We generated fifty different signals for the agent to predict. They were generated randomly from a collection of primitives enumerated in Table 1. They are composed of two different primitives, one for each axis, like so:  $(\text{sig}_x(x + \text{offset}_x) + \text{bias}_x) * (\text{sig}_y(y + \text{offset}_y) + \text{bias}_y)$ . The bias and offset were drawn from  $[-2.0, 2.0]$  and  $[0, 10]$ , respectively. Offset and bias were not applied to either the unit or shortest path primitives. Further, the shortest path primitive was not combined with a second signal, but was used on its own. Gaussian noise with a standard deviation of 0.3 was applied on top of each signal. The shortest path signal is inspired by a common reward function used in RL where each transition has a cost (a negative reward) meant to push the agent to completing a task in a timely manner

Table 1. Signal Primitives

Primitive	Parameters
Fixed value	value $\in [-2.0, 2.0)$
Square wave	period $\in [2, 40)$ , invert $\in [True, False]$
Sin wave	period $\in [2, 40)$
Random binary	Fixed random binary string generated over the length of an axis.
Random float	Fixed random floats $\in [0, 1.0)$ generated over the length of an axis.
Unit	Fixed value of 1.0
Shortest path	transition_cost $\in [-10.0, -1.0)$ , goal_reward $\in [1.0, 10.0)$

and reaching the goal produces a positive reward signal.

Our agent selects actions using  $\epsilon$ -greedy action selection where, at each timestep, with probability  $1 - \epsilon$ , it uses the action specified by a hand-coded policy (see Figure 1a), and otherwise chooses randomly from all four actions ( $\epsilon = 0.3$  in our experiments). A tabular representation is used with each grid cell uniquely represented by a one-hot encoding.

In this set of experiments we can compute the ground-truth predictors for the SR and the signal predictors. This is done by taking the average return observed from each state. The SR reference was averaged over 30,000 episodes and the signal predictor references were averaged over 10,000 episodes. Each episode started at the start state and followed the  $\epsilon$ -greedy policy already described.

We first evaluated the predictive performance of our SR learning algorithm with respect to the step-size for a variety of  $\gamma$  values. We report the average over 30 trials of 10,000 episodes. We initialize the SR weights to 0.0. The squared Euclidean distance was calculated between the predicted SR and the reference SR for each timestep. These values were summed over the run and the average was taken across the runs. These averages are shown in Figure 1b.

Using the results in Figure 1b we evaluated the performance of the two signal prediction approaches by sweeping across step-sizes for various  $\gamma$ . For each experimental run, learning of a new signal was enabled incrementally every 50 episodes. This produced runs with a total length of 2,500 episodes where the first pair of GVs added (the direct and one-step predictors) were trained for 2,500 episodes and the last added GVs were trained for 50 episodes. Further, for each run, the order in which the signals were added was randomized. Thirty runs were performed. The weights of the predictors and of the SR were initialized to 0.0. Notice that the SR was being learned at the same time as the direct and one-step predictions.

For each run a cumulative MSE for each signal  $i$  was calcu-

lated according to Eq. 6. This equation computes the total squared error between the predictor’s estimate,  $V$  and the reference predictor’s estimate,  $V^*$ . For each episode,  $E$ , the error of the current and previous episodes is averaged. Then, for each signal, the maximum error for a given  $\gamma$ , either in the direct or SR-based predictions, is found and used to normalize the errors in the signal across that particular value of  $\gamma$  (see Eq. 7). In this way we attempt to treat the error of each signal equally. If this is not done the errors of large magnitude signals dominate the results. These normalized values are then summed across the signals and the averages across all 30 runs are plotted in Figure 1c.

$$MSE_i = \frac{1}{E} \sum_{e_0}^E \sum_t^T (V_{i:t} - V_{i:t}^*)^2 \tag{6}$$

$$NMSE_{i:\alpha:\gamma} = \frac{MSE_{i:\alpha:\gamma}}{\max_{\alpha} (MSE_{i:\gamma:direct}, MSE_{i:\gamma:SR})} \tag{7}$$

The advantage of the SR-based method is clear as  $\gamma$  increases. This is to be expected since for  $\gamma = 0.0$  both methods are making one-step predictions. In the experiment of Figure 1c, the SR performs better in the vast majority of the signals, as shown in Table 2. For all step-sizes not listed the SR-based method was better on all signals. Analysis of these cases where the direct method did better reveal that some of the target signals have very small magnitudes, suggesting the SR-based approach may be more susceptible to signal-to-noise ratio. Further analysis remains to be done. Finally, we analyzed how the prediction error of our systems evolve with time. This is demonstrated in Figure 2 where we selected the best step-sizes for  $\gamma = 0.9$  and plotted the performance over time across 30 different runs. In this case the order of the signals remained fixed so that sensible averages could be plotted for each signal. Signal performance was normalized as before and summed across all active signals. As expected, we clearly see that the SR-based predictions (green) start with much higher error than the direct (blue), but as the error of the SR (red) drops low the newly added SR-based predictors are able to learn quicker, with less peak and overall error than the direct predictor.

In the lifelong learning setting we never have the opportunity to tune for optimal step-sizes as we did in our evaluation. Practically, fixed step-sizes are used for many robotics settings in RL, but, in order to ensure stable learning, small step-sizes are chosen. As we saw in Figure 1c, the advantage of using the SR-based predictions is enhanced with

Table 2. Signal performance for  $\gamma = 0.9$  of Figure 1c.

$\alpha$	Direct Better	SR-Based Better
0.25	3	47
0.5	5	45
0.75	4	46
1.0	1	49

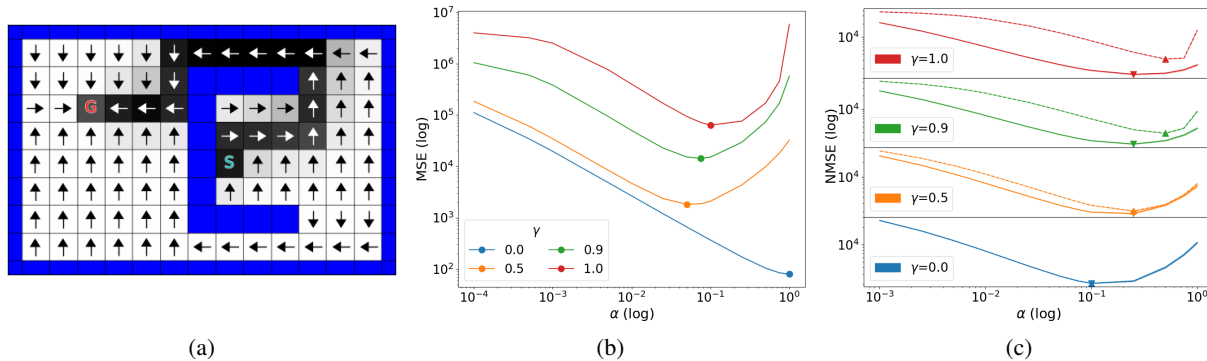


Figure 1. **a)** Dayan’s grid world. Arrows indicate the hand-coded policy (from the start state the policy is to go up). Black squares indicate the SR prediction given from the starting state,  $S$ , for  $\gamma = 1.0$ ; the darker the square the higher the expected visitation. Notice the graying around the central path caused by the  $\epsilon$ -greedy action selection. **b)** MSE of the SR as a function of step-size  $\alpha$  for different values of  $\gamma$ . Lowest error is indicated by the markers. **c)** A comparison of the NMSE of the direct (dashed lines) and SR-based (solid lines) predictions as a function of fixed step-size  $\alpha$  and discount factor  $\gamma$  summed across all signals. Lowest cumulative error is indicated by up arrows (direct predictions) and down arrows (SR-based predictions). Note, that although difficult to see, confidence intervals of 95% are included in both b and c.




Figure 2. All predictors learn from scratch with new predictors added in every 50 episodes. As the SR error (red, right axis) goes low the SR-based predictors (green) are able to learn faster than their direct (blue) counterparts. Shading indicates a 95% confidence interval.

smaller step-sizes. Ideally, however, we would imagine that a fully developed system would use some method of adapting step-sizes.

### 5. Evaluation on a Robot Arm

Tabular settings like Dayan’s grid world are useful for enabling analysis and providing insight into the behavior of our method. However, our goal is to accelerate learning on a real robot where states are not fully observed and cannot be represented exactly; instead we must use function approximation. Here we demonstrate our approach using a robot arm and learning sensorimotor predictions with respect to a human-generated policy. In our task a user controls a robot arm via joystick to trace a counter-clockwise circuit through the inside of the wire maze (see Figure 3) with a rod held in the robot’s gripper. The user performed this task for




Figure 3. The user controls the robot arm using a joystick to trace the inside of the wire maze in a counter-clockwise direction. Circuit path shown in blue.

approximately 12 minutes completing around 50 circuits.

In this experiment we used six different prediction targets: the current, position and speed of both the shoulder rotation and elbow flexion joints. A new predictor was activated every 2,000 timesteps (Note that the robot reports sensor updates at 30 timesteps/s). For this demonstration a discount factor of  $\gamma = 0.95$  was used. Four signals were used as input to our function approximator: the current position and a decaying trace of the position for the shoulder and elbow joints. The decaying trace for joint  $j$  was calculated as  $tr_{j:t+1} = 0.8 * tr_{j:t} + 0.2 * pos_{j:t+1}$ . These inputs were normalized over the joint ranges observed in the experiment and passed into a 4-dimensional tilecoding (Sutton & Barto, 1998) with 100 tilings of width 1.0 and a total memory size of 2048. Additionally a bias unit was added, resulting in a binary feature vector of length 2049 with a maximum of 101 active features on each timestep (hashing collisions can reduce this number). We use a decaying step-size for all the predictors where the step-size starts at 0.1 and decays




Figure 4. A 12 minute run tracing the maze circuit. A new predictor is added every 2,000 timesteps. NMSE errors are summed across all predictors.

linearly to zero over the entire dataset. At each timestep this step-size is further divided by the number of active features in  $\phi(S_t)$ . Finally, for each predictor, this step-size is offset such that the step-size starts at 0.1 when it is first activated and decays at the same rate as all the other predictors.

To compare the prediction error we compute a running MSE for each signal according to Eq. 8, where at each timestep  $t$  the sum is taken over all previous timesteps. Unlike the previous tabular domain, we do not have the ideal estimator to compare against and instead compare the predictions,  $V$ , against the actual return,  $G$ . In order to treat each signal equally we further normalize these errors according to Eq. 9. Note that the NMSE allows us to compare the predictions of a single signal between the two methods, but does not tell us how accurate the predictions are nor does it allow comparison between signals.

$$MSE_{i:t} = \frac{1}{T} \sum_{t=t_0}^T (V_{i:t} - G_{i:t})^2 \quad (8)$$

$$NMSE_i = \frac{MSE_i}{\max(MSE_{i:direct}, MSE_{i:SR})} \quad (9)$$

Figures 4 and 5 show a single run, approximately 12 minutes in length. A single ordering of the predictors was used. Figure 4 shows the error across all predictors while Figure 5 separates out each predictor. Here we see a clear advantage to using the SR-based predictions for most of the signals. Unlike the previous tabular results, there is little difference on the performance of the first predictor (shoulder current), even while the SR is being learned. To investigate, we ran experiments where each signal was learned from the beginning of the run. We observed that performance was rarely worse and sometimes even better when using the SR-based method. This suggests the SR-based approach is more robust than expected, but further experimentation is needed.




Figure 5. The same results as Figure 4, but with the NMSE for the individual predictors. Each is normalized from 0 to 1.

## 6. Further Advantages When Scaling

While this paper analyzed single policies and discount functions, this is not the setting in which the GVF framework is proposed to be used. Rather, it is imagined that massive numbers of GVFS over many policies and timescales will be used represent complex models of the world (Sutton et al., 2011; Modayil et al., 2014). In this setting we note that using SR-based predictions can offer additional benefits, allowing the robot to do more with less. Consider, for a single policy  $\pi$ , a collection of SRs learned for  $f$  discount functions and  $h$  one-step predictors. We can then represent  $fh$  predictions using  $f + h$  predictors. A first advantage is that far fewer GVFs need to be updated on each timestep, saving computational costs. As a second benefit, there is potential to reduce the number of weights used by the system. For example, consider learning in a tabular setting, with  $|\mathcal{S}|$  states, using linear estimators. For  $fh$  predictions the number of weights needed is:  $|w|_{direct} = fh|\mathcal{S}|$ ,  $|w|_{SR-based} = f|\mathcal{S}|^2 + h|\mathcal{S}|$ . It can be shown that for a fixed  $f$  and  $\mathcal{S}$  the total number of weights used by the direct prediction approach is greater when  $h > \frac{f|\mathcal{S}|}{f-1}$ .

## 7. Related Work

The SR was originally introduced as a function approximation method (Dayan, 1993). However, in recent years it has been applied to other settings. For instance, it has been used in transfer learning problems. Using the SR enables good transfer to new tasks when the reward function changes but the policy and dynamics of the environment are very similar to the previously learned task (Kulkarni et al., 2016; Barreto

et al., 2017). The work of Lehnert et al. (2017) shows that transfer with a single SR is limited and can even be detrimental when there is significant difference in the reward functions and therefore in the optimal policies. Barreto et al. (2017) acknowledge this limitation and introduce Generalized Policy Improvement, in which an agent learns multiple SRs, one for each task. These learned SRs and their corresponding policies then serve as a basis set and the agent can evaluate the value of following each known policy under the new reward function, thus enabling it to learn a policy over policies. This allows an agent to learn new policies faster. Interestingly, Gershman (2017) points out that the Google PageRank algorithm can be seen as predicting need probability using the SR with a random walk through links.

Since the SR predicts future visitation there is a natural relationship to planning and this was explored by Russek et al. (2017). They developed several algorithms based on the Dyna algorithm (Sutton & Barto, 1998) and showed that some of these variants were better able to respond to changes in the environment than Dyna.

The SR has also been used to produce exploratory behavior in reinforcement learning agents. Machado et al. (2018) produce a collection of options by performing eigendecomposition of the SR. The resulting eigenvectors are then used to define the options as potential based reward functions. By choosing over these options, which operate at many different timescales, the agent is able to rapidly explore its environment. In fact, this is well connected to another branch of interest in the SR—the relationship between the SR and the hippocampus and entorhinal cortex, which are understood to be used in memory formation and navigation. Stachenfeld et al. (2017) argue that the place cells of the hippocampus form an SR-like representation and that the grid cells of the entorhinal cortex encode a low-dimensional basis of the hippocampus.

A small number of improvements have been suggested for the basic SR algorithm. Gehring (2015) introduced a method for using a low-rank estimation of the SR, which may allow for computationally faster learning. Pitis (2018) proposed a connection between the SR and traces, observing that not only does the SR tell the agent what the successor states are, but also what the predecessor states are. This backward view can be used like traces to assign credit. They combine both the successor (forward) and predecessor (backward) views to improve the learning speed and accuracy of SR-based prediction. Both of these algorithmic improvements might serve to improve our own results.

GVF's were originally proposed as a method for building an agent's overall knowledge in a modular and hierarchical way, representing state as a collection of predictions (Sutton et al., 2011). To date they have been used with hand-coded fixed policies (Modayil et al., 2014; Sherstan et al., 2015)

and as state for learned policies (Pilarski et al., 2013). GVF's have also been used to track internal statistics such as the variance of the return and state visitation (Sherstan et al., 2016; 2018).

Another use of GVF's is as auxiliary tasks which drive representation learning. The UNREAL agent (Jaderberg et al., 2017) is a powerful demonstration of this approach in which a number of pre-defined auxiliary prediction tasks (GVF's) are used to drive representation learning of a convolutional neural network, resulting in accelerated learning and better policies. The work of Bellemare et al. (2017) estimated the distribution of the return with what can be seen as a collection of GVF's. Although this distributional approach and standard methods both use an estimate of the expected return the distributional approach learned better policies. This is surprising and the cause of this improvement is unknown, but may be do to the same auxiliary task effect driving improved representation learning.

Finally, the idea closest to our work is the concept of Universal Value Functions (UVFAs) (Schaul et al., 2015). UVFAs are a generalization of value functions over goals in a parametrized way. Ma et al. (2018) introduced the universal successor representation (USR), which combined the UVFA with the SR allowing SR-based predictions to generalize across goals. We believe our approach is mutually compatible with the USR and could be beneficially combined to further improve learning efficiency and generality.

## 8. Conclusions

In this paper we showed how the successor representation (SR), although originally introduced for another purpose, can be used to accelerate learning in a lifelong learning setting in which an agent incrementally constructs models of its world as a collection of predictions known as general value functions (GVF's). The SR enables a given prediction to be modularized into two components, one representing the dynamics of the environment (the SR) and the other representing the target signal (one-step signal prediction). This allows an agent to reuse its existing knowledge when adding a new prediction target, speeding up learning of the new predictor. We demonstrated this behaviour in both a tabular grid world and on a robot arm. These results suggest an effective method for improving the learning rate and sample efficiency for robots learning in the real world.

There are several clear opportunities for further research on this topic. The first is to provide greater understanding into why, for a given fixed step-size, some (few) signals are better predicted directly rather than through the SR. Further, the work in using the SR with function approximation is preliminary and more insight can yet be gained in this setting. Another opportunity for research is to explore using

SR-based predictions with state-dependent discount functions. Finally, we suggest that SR-based predictions with deep feature learning (Kulkarni et al., 2016) and an incrementally constructed architecture would be a very powerful tool to support lifelong or developmental learning in robotic domains with widespread real-world applications.

## References

- Barreto, A., Dabney, W., Munos, R., Hunt, J. J., Schaul, T., Silver, D., and van Hasselt, H. P. Successor Features for Transfer in Reinforcement Learning. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 4055–4065, 2017.
- Bellemare, M. G., Dabney, W., and Munos, R. A Distributional Perspective on Reinforcement Learning. In *International Conference on Machine Learning (ICML)*, Sydney, Australia, 2017.
- Dayan, P. Improving Generalization for Temporal Difference Learning: The Successor Representation. *Neural Computation*, 5(4):613–624, 1993.
- Gehring, C. A. Approximate Linear Successor Representation. In *Reinforcement Learning Decision Making (RLDM)*, 2015.
- Gershman, S. J. Predicting the Past, Remembering the Future. *Current Opinion in Behavioral Sciences*, 17, 2017.
- Jaderberg, M., Mnih, V., Czarnecki, W. M., Schaul, T., Leibo, J. Z., Silver, D., and Kavukcuoglu, K. Reinforcement Learning with Unsupervised Auxiliary Tasks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Kulkarni, T. D., Saedi, A., Gautam, S., and Gershman, S. J. Deep Successor Reinforcement Learning. *arXiv*, 1606.02396, 2016.
- Lehnert, L., Tellex, S., and Littman, Michael L. Advantages and Limitations of using Successor Features for Transfer in Reinforcement Learning. *arXiv*, 1708.00102, 2017.
- Ma, C., Wen, J., and Bengio, Y. Universal Successor Representations for Transfer Reinforcement Learning. *arXiv*, 1804.03758, 2018.
- Machado, M. C., Rosenbaum, C., Guo, X., Liu, M., Tesauro, G., and Campbell, M. Eigenoption Discovery through the Deep Successor Representation. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- Modayil, J., White, A., and Sutton, R. S. Multi-Timescale Nexting in a Reinforcement Learning Robot. *Adaptive Behavior*, 22(2):146–160, 2014.
- Oudeyer, P., Kaplan, F., and Hafner, V. V. Intrinsic Motivation Systems for Autonomous Mental Development. *Transactions on Evolutionary Computation*, 11(2):265–286, 2007.
- Pilarski, P. M., Dick, T. B., and Sutton, R. S. Real-Time Prediction Learning for the Simultaneous Actuation of Multiple Prosthetic Joints. In *International Conference on Rehabilitation Robotics (ICORR)*, pp. 1–8, 2013.
- Pitis, S. Source Traces for Temporal Difference Learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, New Orleans, Louisiana, USA, 2018.
- Russek, E. M., Momennejad, I., Botvinick, M. M., Gershman, S. J., and Daw, N. D. Predictive Representations Can Link Model-based Reinforcement Learning to Model-free Mechanisms. *PLoS Computational Biology*, 13(9): 1–42, 2017.
- Schaul, T., Horgan, D., Gregor, K., and Silver, D. Universal Value Function Approximators. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 1312–1320, 2015.
- Sherstan, C., Modayil, J., and Pilarski, P. M. A Collaborative Approach to the Simultaneous Multi-joint Control of a Prosthetic Arm. In *International Conference on Rehabilitation Robotics (ICORR)*, pp. 13–18, Singapore, Singapore, 2015.
- Sherstan, C., Machado, M. C., White, A., and Pilarski, P. M. Introspective Agents: Confidence Measures for General Value Functions. In *International Conference on Artificial General Intelligence (AGI)*, pp. 258–261, New York, New York, USA, 2016.
- Sherstan, C., Bennett, B., Young, K., Ashley, D. R., and Sutton, R. S. Directly Estimating the Variance of the  $\lambda$ -Return Using Temporal-Difference Methods. *arXiv*, 1801.08287, 2018.
- Stachenfeld, K. L., Botvinick, M. M., and Gershman, S. J. The Hippocampus as a Predictive Map. *Nature Neuroscience*, 20:1643–1653, 2017.
- Sutton, R. S. and Barto, A. G. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- Sutton, R. S., Modayil, J., Delp, M., Degris, T., Pilarski, P. M., White, A., and Precup, D. Horde: a Scalable Real-time Architecture for Learning Knowledge from Unsupervised Sensorimotor Interaction. In *International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 761–768, 2011.
- Thrun, S. A Lifelong Learning Perspective for Mobile Robot Control. In *Intelligent Robots and Systems*, pp. 201–214, 1995.