

---

# Representation Learning on Graphs: A Reinforcement Learning Application

---

**Author 1**  
Institution 1

**Author 2**  
Institution 2

**Author 3**  
Institution 3

## Abstract

In this work, we study value function approximation in reinforcement learning (RL) problems with high dimensional state or action spaces via a generalized version of representation policy iteration (RPI). We consider the limitations of proto-value functions (PVFs) at accurately approximating the value function in low dimensions and we highlight the importance of features learning for an improved low-dimensional value function approximation. Then, we adopt different representation learning algorithm on graphs to learn the basis functions that best represent the value function. We empirically show that *node2vec*, an algorithm for scalable feature learning in networks, and the *Variational Graph Auto-Encoder* constantly outperform the commonly used smooth proto-value functions in low-dimensional feat space.

## 1 INTRODUCTION

In reinforcement learning (RL), an agent, or decision maker, takes sequential actions and observes the consequent rewards and states, which are unknown *a priori*. These sequent observations improve the agent’s knowledge of the environment with the final goal of learning the optimal policy that maximizes the long term reward. The learning control problem is usually formulated as Markov decision process (MDP), where each state has an associated value function, which estimates the expected long term reward for some policy (usually the optimal one). Classical MDPs represent the value function by a lookup table, with one entry for each state. However, this does not scale with the state

(and implicitly also action) space dimension, leading to slow learning processes in high-dimensional reinforcement learning problems. Approximated reinforcement learning addresses this problem by learning a function to properly approximate the true value function. In the literature, many types of functions have been studied [Kaelbling et al., 1996, Sutton and Barto, 1998].

In this work, we study linear value function approximation, where the value function is represented as a weighted linear sum of a set of features (called *basis function*). Linear function approximation allows to represent complex value functions by choosing arbitrarily complex basis functions. Under this framework, one of the main challenges is to identify the right set of basis functions. Typical linear approximation architectures such as polynomial basis functions (where each basis function is a polynomial term) and radial basis functions (where each basis function is a Gaussian with fixed mean and variance) have been studied in the case of reinforcement learning [Lagoudakis, 2003]. These architectures make the assumption that the underlying state space has Euclidean geometry. However, in realistic scenarios, the MDP’s state space is likely to exhibit irregularities. For instance, let consider the environment depicted in Figure 1. As it can be seen in Figure 1(b), neighboring states can have values that are far apart (such as states on opposite sides of a wall). In such cases, these traditional parametric functions may not be able to accurately approximate value functions.

Consequently, other basis functions have been studied to address this issue. Example of such methods include Fourier basis [Konidaris et al., 2011], diffusion wavelets, [Mahadevan and Maggioni, 2006], Krylov bases [Petrik et al., 2010] and Bellman Error Basis Function [Parr et al., 2007, Parr et al., 2008]. In particular, [Mahadevan, 2007] proposes a graph spectral framework for solving Markov decision processes by jointly learning representations and optimal policies called *representation policy iteration* (RPI). They first note that MDPs can be intuitively represented using graphs, with states being the nodes and the transi-

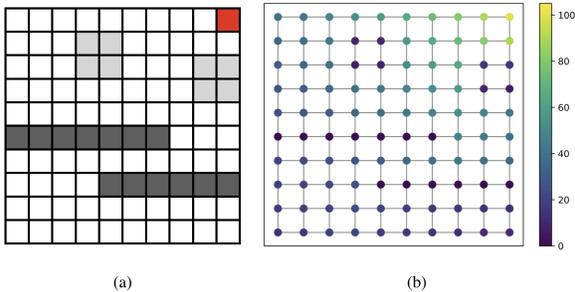


Figure 1: (a) Maze environment. The dark grey squares are strict walls, while the light grey square are difficult access rooms. The red square is the goal room. (b) Optimal value function computed using value iteration [Montague, 1999].

tion probability being the similarity matrix. Then, under the assumption that the value function is usually modelled as a diffusion process over the state-graph (and therefore it is a smooth function), they approximate the value function (smooth signal on the graph) as a linear combination of the first Laplacian eigenmaps on graphs. These features, known as *proto-value functions* (PVFs), preserve the smoothness of the value function. In this paper, we argue that constructing the graph that perfectly models the MDP such that the value function is indeed smooth on the graph is not trivial. Therefore, there is a need to automatically learn the basis functions that capture the geometry of the underlying state space from limited data to further improve the performance. Hence, given the success of recent node embedding models [Grover and Leskovec, 2016, Kipf and Welling, 2016b, Ribeiro et al., 2017, Donnat et al., 2018], we propose to investigate representation learning on graphs algorithms, to learn basis functions in the linear value function approximation.

The idea behind recent successful representation learning approaches is to learn a mapping that embeds nodes on a graph as low-dimensional vectors. They aim to optimize the representations so that geometric relationships in the embedding space preserve the structure of the original graph. [Hamilton et al., 2017] surveys recent representation learning on graph methods. Therefore, in this work we generalize the RPI algorithm [Mahadevan, 2007] algorithm to allow different basis functions, and analyze the performance of several representation learning methods for value function approximation.

The rest of this paper is structured as follows: in Section 2 we provide background material, providing details on Markov decision processes and value function approximation and describing the representation

learning algorithm used in this work. The General Representation Policy Iteration is described in Section 3. In Section 4 we run experiments and discuss the results. We proceed to summarize the main findings and give direction for future work. We finally conclude in Section 5.

## 2 BACKGROUND

### 2.1 Markov decision process (MDP)

Markov decision processes are discrete time stochastic control processes that provide a widely-used mathematical framework for modeling decision making strategies under uncertainty. Specifically, at each time step, the process is in some state  $s$ , and the agent can choose any action  $a$  that is available in state  $s$ . In part as a consequence of the action taken, the agent finds itself in a new state  $s'$  and observes an instantaneous reward  $r$ . We define a discrete MDP by the tuple  $M = (S, A, P, R)$ , where  $S$  is a finite set of discrete states,  $A$  a finite set of actions,  $P$  describes the transition model –with  $P(s, a, s')$  giving the probability of moving from state  $s$  to  $s'$  given action  $a$  – and  $R$  describes the reward function – with  $R(s, a)$  giving the immediate reward from taking action  $a$  in state  $s$ . Given a policy  $\pi : S \mapsto A$ , a value function  $V$  is a mapping  $S \mapsto \mathbb{R}$  that describes the expected long-term discounted sum of rewards observed by the agent in any given state  $s$  when following policy  $\pi$ . Solving a MDP requires to find a policy that defines the optimal value function  $V^*$ , which satisfies the following constraints:

$$V^*(s) = \max_a \left( R(s, a) + \gamma \sum_{s' \in S} P(s, a, s') V^*(s') \right).$$

This recursive equation is known as the *standard form* of Bellman’s equation. The optimal policy is a unique solution to the Bellman’s equation and can be found by dynamic programming, iteratively evaluating the value functions for all states.

### 2.2 Value Function Approximation

In large state spaces, computing exact value functions can be computationally intractable. A possible solution is to estimate the value function with function approximation (value function approximation method) [Bertsekas and Tsitsiklis, 1996]. Commonly, the value function is approximated as a weighted linear sum of a set of features (called *basis function*) [Montague, 1999, Mahadevan, 2007, Konidaris et al., 2011, Lagoudakis, 2003]

$$\phi_1, \phi_2, \dots, \phi_d : \tilde{V}(s|\theta) = \sum_{i=1}^d \theta_i \phi_i(s) \approx V(s),$$

where  $d$  is the dimension of the features space.

The basis functions  $\phi_i$  can be hand-crafted [Sutton and Barto, 1998] or automatically constructed [Mahadevan, 2007], and the model parameters  $\theta = [\theta_1, \theta_2, \dots, \theta_d]$  are typically learned via standard parameter estimation methods such as least-square policy iteration (LSPI) [Lagoudakis, 2003]. However, how to properly design the set of basis for a data-efficient function approximation framework is still an open question. The main question is how to find the set of basis  $\phi$  that is a *low-dimensional* (to ensure a data-efficient learning) and yet *meaningful representation* of the MPD (to reduce the suboptimality due to the value function approximation).

The representation policy iteration algorithm (RPI) was introduced in [Mahadevan, 2007] to address this problem. It is a three steps algorithm consisting of (1) a sample collection phase, (2) a representation learning phase and (3) a parameter estimation phase. RPI is described in further details in Section 3. In this work, we propose to generalize RPI to allow different representation learning method. In particular, we first observe that state space topology of MDPs can be intuitively modeled as (un)directed weighted graphs, with the nodes being the states and the transition probability matrix being the similarity matrix. When the transition probabilities are unknown, we can construct a graph from collected samples by connecting temporally consecutive states with a unit cost edge. Therefore, similarly to [Mahadevan, 2007], we propose to construct the graph from collected samples of an agent acting in the environment given by the MDP. We then learn representations on the graph induced by the MDP using node embedding models. Finally, we use the learned representation to linearly approximate the value function. In the next section, we describe the node embedding models that we exploit within this framework.

### 2.3 Representation Learning on Graph

We propose to use the following learned node embedding models as basis functions for the value function approximation in order to automatically learn to encode the graph structure - hence the MDP - into low-dimensional embeddings.

**Node2Vec** *Node2vec* [Grover and Leskovec, 2016] is an algorithmic framework for learning continuous feature representations for nodes in networks. It is inspired by the powerful language model *Skip-gram* [Mikolov et al., 2013] which is based on the hypothesis that words that appear in the same context share semantic meaning. In networks, the same hypothesis can be made for nodes, where the context of a node

is derived by considering the nodes that appear in the same random walk on the graph. Therefore, *node2vec* learns the node embeddings based on random walk statistics. The key is to optimize the node embeddings so that nodes have similar embeddings if they tend to co-occur on short (biased) random walks over the graph. Moreover, it allows for a flexible definition of random walks by introducing parameters that allow to interpolate between walks that are more breadth-first search or depth-first search.

Specifically, for a graph  $G = (\mathcal{V}, \mathcal{E}, W)$  (where  $\mathcal{V}$  is a set of nodes,  $\mathcal{E}$  is a set of edges and  $W$  the weight matrix) and a set  $\mathcal{W}_S$  of biased random walks  $w = u_1, u_2, \dots, u_T$  collected under sampling strategy  $S$  on the graph  $G$ , *node2vec* seeks to maximize the log-probability of observing a network neighborhood  $N_S(u)$  for a node  $u \in \mathcal{V}$  conditioned on its feature representation, given by  $f$  (a matrix of size  $|\mathcal{V}| \times d$  parameters, where  $d$  is the dimension of the feature space):

$$\max_f \sum_{w \in \mathcal{W}} \sum_{t=1}^T \log Pr(N_{S,w}(u_i) | f(u_i))$$

where  $N_{S,w}(u_i)$  describes the neighborhood of the  $i$ th node in the walk  $w$ .

**Struc2Vec** By introducing a bias in the sampling strategy, *node2vec* allows to learn representations that do not only focus on optimizing node embeddings so that nearby nodes in the graph have similar embeddings, but also consider representations that capture the structural roles of the nodes, independently of their global location on the graph. The recent node embedding approach, *struc2vec*, proposed by [Ribeiro et al., 2017] addresses the problem of specifically embedding nodes such that their structural roles are preserved. The model generates a series of weighted auxiliary graphs  $G'_k, k = 1, 2, \dots$  from the original graph  $G$ , where the auxiliary graph  $G_k$  captures structural similarities between nodes  $k$ -hop neighborhoods. Formally, let  $R_k(u_i)$  denotes the ordered sequence of degrees of the nodes that are exactly  $k$ -hops away from  $u_i$ , the edge-weights,  $w_k(u_i, v_j)$ , in the auxiliary graph  $G'_k$  are recursively represented by the *structural distance* between nodes  $u_i$  and  $v_j$  defined as

$$w_k(u_i, v_j) = w_{k-1}(u_i, v_j) + d(R_k(u_i), R_k(u_j)),$$

where  $w_0(u_i, v_j) = 0$  and  $d(R_k(u_i), R_k(u_j))$  is the *distance* between the ordered degree sequences  $R_k(u_i)$  and  $R_k(u_j)$  computed via dynamic time warping [Ribeiro et al., 2017].

Once the weighted auxiliary graphs  $G'_k$  are computed, *struc2vec* runs biased random walks over them and

proceeds as node2vec, optimising the log-probability of observing a network neighborhood based on these random walks.

**GraphWave** The *GraphWave* algorithm as proposed by [Donnat et al., 2018] takes a different approach to learning structural node embeddings. It learns node representations based on the diffusion of a spectral graph wavelet centered at each node. For a graph  $G$ ,  $\mathcal{L} = D - A$  denotes the graph Laplacian, where  $A$  is the adjacency matrix and  $D$  is a diagonal matrix, whose entries are row sums of the adjacency matrix. Let  $U$  denote the eigenvector decomposition of the graph Laplacian  $\mathcal{L} = U\Lambda U^T$  and  $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_{|V|})$  denote the eigenvalues of  $\mathcal{L}$ . Given a heat kernel  $g_s(\lambda) = e^{-s\lambda}$  for a given scale  $s$ , GraphWave uses  $U$  and  $g_s$  to compute a vector  $\psi_u$  representing diffusion patterns for node  $u$  as follows:

$$\psi_u = U \text{diag}(g_s(\lambda_1), g_s(\lambda_2), \dots, g_s(\lambda_{|V|})) U^T \delta_u$$

where  $\delta_u$  is the one-hot indicator vector for node  $u$ . Then, the characteristic function for each node’s coefficients  $\psi_u$  is computed as

$$\phi_u(t) = \frac{1}{|V|} \sum_{m=1}^{|V|} e^{it\Psi_{m,u}}$$

Finally, to obtain the structural node embedding  $f(u)$  for node  $u$ , the parametric function  $\phi_u(t)$  is sampled at  $d$  evenly spaced points  $t_1, \dots, t_d$ :

$$f(u) = [\text{Re}(\phi_u(t_i)), \text{Im}(\phi_u(t_i))]_{t_1, \dots, t_d}.$$

**Variational Graph Auto-Encoder** As opposed to directly encoding each node, auto-encoders aim at directly incorporating the graph structure into the encoder algorithm. The key idea is to compress information about a node’s local neighborhood. The *Variational Graph Auto-Encoder* proposed by [Kipf and Welling, 2016b] is a latent variable model for graph-structure data capable of learning interpretable latent representations for undirected graphs. The Graph Auto-Encoder uses a Graph Convolutional Neural Network (GCN) [Kipf and Welling, 2016a] to encode graphs and another GCN to reconstruct the graph. The Variational Graph Auto-Encoder makes use of latent variables.

### 3 GENERAL REPRESENTATION POLICY ITERATION

Within the context of approximated value function, the representation policy iteration algorithm (RPI) was introduced in [Mahadevan, 2007] to properly learn

the approximating function. Namely, RPI is a three steps algorithm consisting of (1) a sample collection phase, to build a training dataset with the quadruples  $\{(s_i, a_i, s_{i+1}, r_i)\}$ ; (2) a representation learning phase, that defines the proper set of basis functions; and (3) a parameter estimation phase, in which the coefficients  $\theta$  are learned. A generalized version of the RPI algorithm [Mahadevan, 2007] is described in Algorithm 1.

In the original RPI, the representation learning phase is predefined. Namely, an undirected weighted graph  $G$  is built from the available data set  $\mathcal{D}$ . Then a diffusion operator  $O$ , such as the normalized Laplacian is computed on graph  $G$  and the  $d$ -dimensional basis functions  $\phi = [\phi_1, \dots, \phi_d]$  are constructed from spectral analysis of the diffusion operator. Specifically,  $\phi_i$  are the smoothest eigenvectors of the graph Laplacian and are known as *proto-value functions* (PVFs). Under the assumption that the value function is actually smooth on the considered graph, it can be efficiently represented as a linear combination of the PVFs. However, it is not guaranteed that we can construct a graph from a limited number samples such that its derived proto-value functions reflect the underlying state space.

To overcome this limitation, we propose to use the node embedding methods described in 2.3 automatically learn the basis functions from the geometry of the underlying state space to further improve the performance.

---

#### Algorithm 1 General Representation Policy Iteration

---

**Input:**  $\pi_0$ : sampling strategy,  $N$ : number of random walks to sample,  $T$ : length of each walk,  $d$  dimension of the basis functions, *model* representation learning model,  $\epsilon$  convergence condition for LSPI.

**Output:**  $\epsilon$ -optimal policy  $\pi$

##### 1. Sample Collection Phase

Collect a data set  $\mathcal{D}$  of  $T$  successive samples  $\{(s_i, a_i, s_{i+1}, r_i), (s_{i+1}, a_{i+1}, s_{i+2}, r_{i+1}), \dots\}$  by following sampling strategy  $\pi_m$  for maximum  $T$  steps (terminating earlier if it results in an absorbing goal state).

##### 2. Representation Learning Phase

Build basis function matrix  $\phi = \text{model}(\mathcal{D}, d)$ .

##### 3. Control Learning Phase

Using a parameter estimation algorithm such as LSPI or Q-learning, find an  $\epsilon$ -optimal policy  $\pi$  that maximizes the action value function  $Q^\pi = \phi\theta^\pi$  within the linear span of the bases  $\phi$ .

---

In the following, we describe how to apply these features learning methodologies within reinforcement learning strategies.

In the original RPI, the representation learning phase is predefined. Namely, an undirected weighted graph  $G$  is built from  $\mathcal{D}$ . Then a diffusion operator  $O$ , such as the normalized Laplacian is computed on graph  $G$  and the basis functions are constructed from spectral analysis of the diffusion operator. Specifically,  $\phi_i$  are the smoothest eigenvectors of the graph Laplacian and are known as *proto-value functions* (PVFs). The key is that given a state-graph that perfectly represents the MDS, the value function is modelled as a diffusion process over the graph (and therefore it is a smooth function). Hence, given the spectral properties of the Laplacian operator, the proto-value functions are a good choice of basis functions for preserving the smoothness of the value function.

However, it is not guarantee that we can construct a graph from a limited number samples such that its derived proto-value functions reflect the underlying state space.

We propose to use the node embedding methods described in 2.3 automatically learn the basis functions from the geometry of the underlying state space to further improve the performance.

## 4 EXPERIMENTS

### 4.1 Set up

We consider the two-room environment used in [Mahadevan, 2007]. It is shown in Figure 2(a). It consists of 100 states in total, divided into 57 accessible states and 43 inaccessible states representing walls. There is one goal state, marked in red and the agent is rewarded by +100 for reaching the goal state.

We also consider an other environment Figure 2(b)) consisting of a single room with obstacles. There are also 100 states in total, some of which are inaccessible since they represent exterior walls and 14 of which are accessible from neighboring states with a fixed probability  $p = 0.2$  (they represent a moving obstacle or difficult access space). All other state are reachable with probability 0.9. The agent is rewarded by +100 for reaching the state located at the upper-right corner.

We construct the corresponding graphs where each location is a node, and the transitions (4 possible actions: left, right, up and down) are represented by the edges.

We run and evaluate the General Representation Policy Iteration (GRPI) algorithm using embedding models from Section 2.3 to compute the basis functions in the second phase of the algorithm.

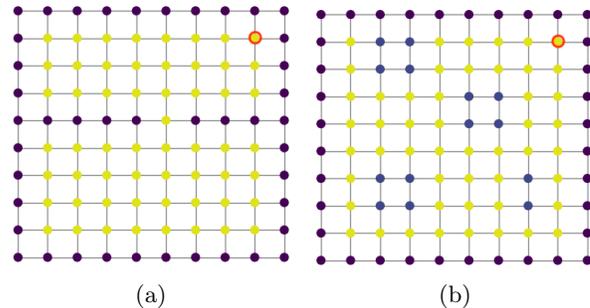
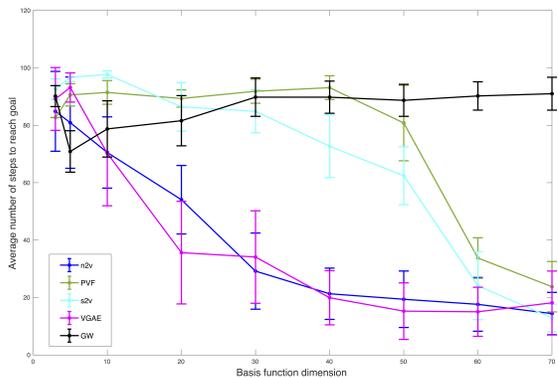
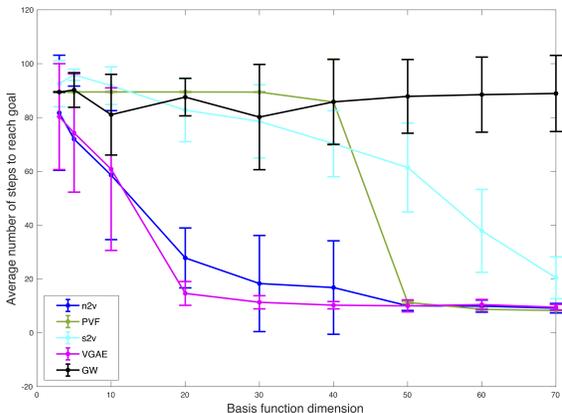


Figure 2: Two different maze environments. The purple nodes represent strict walls, while the blue nodes are difficult access rooms. All other nodes represent accessible rooms. The node shown in red is the goal room. Best seen in color.

1. We collected a set  $\mathcal{D}$  of samples from a 100 random walks, each of length 100 (or terminating early when the goal state was reached). Four actions (move up, down, left or right) are possible from each state. If a movement is possible, it succeeds with probability 0.9. Otherwise, the agent remains in the same state. When the agent reaches the goal state, it receives a reward of 100, and is randomly reset to an accessible interior state. We use off-policy sampling ( $\pi_0 =$  random policy) to collect the samples, except in the case of node2vec, where the samples are generated under a biased random walk. We use grid search to find the optimal hyperparameters  $p = 1$  and  $q = 4$  that guide the walk according to [Grover and Leskovec, 2016].
2. We used sample transitions in  $\mathcal{D}$  to build an undirected graph where the weight matrix  $W$  is the adjacency matrix and run  $model(\mathcal{D}, d)$  with  $model \in (\text{node2vec}(\text{n2v}), \text{struc2vec}(\text{s2v}), \text{variational graph auto-encoder (VGAE)}, \text{GraphWave (GW)})$  for different choices of  $d$ . In the case of node2vec, we directly use the collected samples to derive the node neighborhoods used in the objective function.
3. We learned the parameters of the linear value approximation using the parameter estimation method LSPI [Lagoudakis, 2003] using the samples set  $\mathcal{D}$ .
4. We used the policies learned by GRPI for each model to run simulations starting from each accessible states. We compare the performance of each models in terms of the average number of steps required to reach the goal. We also compare to the traditional PVF basis functions. The results for the two environments, averaged over 20



(a)



(b)

Figure 3: Average number of steps required to reach the goal steps using the various basis functions. On the  $x$  axis we make the dimension of the basis functions vary. Best seen in color. 3(a) corresponds to the two-room environment, while 3(b) corresponds to the obstacle-room environment.

independent runs, are shown in Figures 3(a) and 3(b). Each run consists of episodes of a maximum of 100 steps, where each episode is terminated earlier if the agent reached the goal state.

## 4.2 Discussion

Figures 3(a) and 3(b) provide the average number of steps to reach the goal as a function of the dimension of the basis function. We first observe that the policy learned via the GraphWave basis function lead to very poor performances regardless of the dimension size. We investigate this phenomenon by looking at the approximate value function learned under these basis. The approximate state values are depicted in

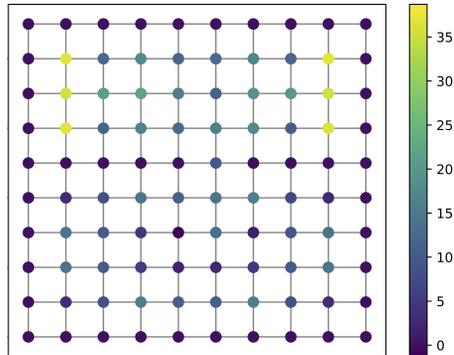


Figure 4: Approximate value function via GRPI using GraphWave basis function of dimension 70 on the two-room environment.

Figure 4. Because GraphWave aims at learning embeddings that are exclusively structural, we hypothesize that they fail at capturing global network properties. In fact, the embeddings learned by GraphWave for the corner states in the two-room environment are equal, making it obviously impossible to learn different state value with linear approximation. This suggests that although the GraphWave is a powerful model for capturing structural information in networks, it is not a good choice of basis function for approximating smooth value function on a graph.

On the other hand, we notice that although struc2vec was also designed to capture structural similarities between nodes, it also preserves the local properties of the graph by considering neighborhoods of different sizes [Ribeiro et al., 2017]. Hence, struc2vec is able to accurately approximate the value function even in graphs that have a perfectly symmetrical structure such as the two-room environment.

Finally, the result show that VGAE and node2vec are good choices of basis functions for approximating the value function in low dimension. Indeed, they lead to good performances in terms of number of steps to reach the goal states with basis functions of dimension as low as 20 for VGAE and 30 for node2vec. On the contrary, observe that the PVFs require dimension of at least 70 to reach comparable performances on the two-room domain and dimension of 50 on the obstacle-room domain.

The observed that the sampling strategy used in node2vec has a significant impact on the performance of the learned policy. Using grid search, we find that the optimal value of the parameters  $p$  and  $q$  that guide the random are 1 and 4 respectively. We show the performances of node2vec with selected values of  $p$  and  $q$  in 5. When  $p < q$  and  $q > 1$ , the strategy is bi-

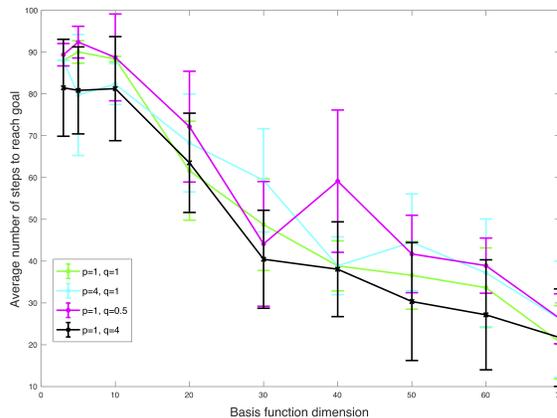


Figure 5: Average number of steps required to reach the goal steps using node2vec with varying parameters  $p$  and  $q$ . On the  $x$  axis we make the dimension of the basis functions vary.

used to encourage walks to backtrack a step and to visit nodes that are close to the current node in the walk. Therefore, it leads to walks that approximate a BFS behavior, gathering a local view of the underlying graph with respect to the starting node. On the other hand, when  $p > q$  and  $q < 1$ , the walk approximate a DFS behavior and lead to more outward exploration. [Grover and Leskovec, 2016] show that the first type sampling strategy allow to reflect structural equivalences of nodes whereas the second type allows to capture homophily within the network. Figure 5 suggests that for approximating value functions, structural equivalence plays a more important role than homophily.

### 4.3 Main Findings and Future Work

We summarize below the main findings of our work.

- Using basis functions that automatically learn to embed the geometry of the graph induced by the MPD can lead to improved performance over the proto-value function.
- Such embedding models need to capture the structural equivalence of the nodes while preserving the local properties of the graph.
- Under a sampling strategy that satisfy the requirements of the previous point, Node2vec [Grover and Leskovec, 2016] outperforms the commonly used proto-value functions.
- The Variational Graph Auto-Encoder, which is more complex than node2vec and requires more

training, leads to minor performance improvement compared to node2vec.

These findings encourage the further study of representation learning on graphs for achieving efficient and accurate policy learning for reinforcement learning. In particular, the question of scalability in large or continuous state space arises. Future work include analyzing to what extend one can efficiently learn good embeddings with limited samples in very large state spaces. Another interesting open question in this direction, is to investigate whether good representations can be inferred for states that have never been visited.

Naturally, future work should also aim at further improving the quality of the embeddings for the problem of reinforcement learning. A possibility would be to make use of the reward observed during the sample collection phase to build features that are not only based on state transitions, but capture reward information as well.

## 5 CONCLUSION

In this work, we study the representation policy iteration algorithm with a modified representation learning phase that allows to use any model for computing the basis functions in the linear value approximation. We investigate several models for learning high quality node embeddings that preserve the geometry of the graph induced by the Markov decision process. We compare the performance of several representation learning model in the context of value function approximation. Finally, we observe that models that are designed to capture the global structural geometry of the graph while preserving local properties do well at approximating the value function in low feature space dimensions, significantly outperforming the commonly considered proto-value functions for this task.

## References

- [Bertsekas and Tsitsiklis, 1996] Bertsekas, D. P. and Tsitsiklis, J. N. (1996). *Neuro-Dynamic Programming*. Athena Scientific, 1st edition.
- [Donnat et al., 2018] Donnat, C., Zitnik, M., Hallac, D., and Leskovec, J. (2018). Learning structural node embeddings via diffusion wavelets. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '18*, pages 1320–1329, New York, NY, USA. ACM.
- [Grover and Leskovec, 2016] Grover, A. and Leskovec, J. (2016). node2vec. In *Proceedings of the 22nd*

- ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*, pages 855–864.
- [Hamilton et al., 2017] Hamilton, W. L., Ying, R., and Leskovec, J. (2017). Representation learning on graphs: Methods and applications. *CoRR*, abs/1709.05584.
- [Kaelbling et al., 1996] Kaelbling, L. P., Littman, M. L., and Moore, A. P. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285.
- [Kipf and Welling, 2016a] Kipf, T. N. and Welling, M. (2016a). Semi-supervised classification with graph convolutional networks. *CoRR*, abs/1609.02907.
- [Kipf and Welling, 2016b] Kipf, T. N. and Welling, M. (2016b). Variational graph auto-encoders. *CoRR*, abs/1611.07308.
- [Konidaris et al., 2011] Konidaris, G., Osentoski, S., and Thomas, P. (2011). Value Function Approximation in Reinforcement Learning using the Fourier Basis. *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence*, pages 380–385.
- [Lagoudakis, 2003] Lagoudakis, M. (2003). Least-squares policy iteration. *The Journal of Machine Learning Research*, 4:1107–1149.
- [Mahadevan, 2007] Mahadevan, S. (2007). Learning Representation and Control in Markov Decision Processes: New Frontiers. *Foundations and Trends® in Machine Learning*, 1(4):403–565.
- [Mahadevan and Maggioni, 2006] Mahadevan, S. and Maggioni, M. (2006). Value function approximation with diffusion wavelets and laplacian eigenfunctions. In Weiss, Y., Schölkopf, B., and Platt, J. C., editors, *Advances in Neural Information Processing Systems 18*, pages 843–850. MIT Press.
- [Mikolov et al., 2013] Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781.
- [Montague, 1999] Montague, P. (1999). Reinforcement Learning: An Introduction, by Sutton, R.S. and Barto, A.G. *Trends in Cognitive Sciences*, 3(9):360.
- [Parr et al., 2008] Parr, R., Li, L., Taylor, G., Painter-Wakefield, C., and Littman, M. L. (2008). An analysis of linear models, linear value-function approximation, and feature selection for reinforcement learning. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, pages 752–759, New York, NY, USA. ACM.
- [Parr et al., 2007] Parr, R., Painter-Wakefield, C., Li, L., and Littman, M. (2007). Analyzing feature generation for value-function approximation. In *Proceedings of the 24th International Conference on Machine Learning, ICML '07*, pages 737–744, New York, NY, USA. ACM.
- [Petrik et al., 2010] Petrik, M., Taylor, G., Parr, R., and Zilberstein, S. (2010). Feature selection using regularization in approximate linear programs for markov decision processes. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pages 871–878.
- [Ribeiro et al., 2017] Ribeiro, L. F., Saverese, P. H., and Figueiredo, D. R. (2017). Struc2vec: Learning node representations from structural identity. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '17*, pages 385–394, New York, NY, USA. ACM.
- [Sutton and Barto, 1998] Sutton, R. S. and Barto, A. G. (1998). Reinforcement learning: An introduction. *IEEE Trans. Neural Networks*, 9(5):1054–1054.