
Gated Path Planning Networks

Lisa Lee^{*1} Emilio Parisotto^{*1} Devendra Singh Chaplot¹ Eric Xing¹ Ruslan Salakhutdinov¹

Abstract

Value Iteration Networks (VINs) are effective differentiable path planning modules that can be used by agents to perform navigation while still maintaining end-to-end differentiability of the entire architecture. Despite their effectiveness, they suffer from several disadvantages including training instability, random seed sensitivity, and other optimization problems. In this work, we reframe VINs as recurrent-convolutional networks which demonstrates that VINs couple recurrent convolutions with an unconventional max-pooling activation. From this perspective, we argue that standard gated recurrent update equations could potentially alleviate the optimization issues plaguing VIN. The resulting architecture, which we call the Gated Path Planning Network, is shown to empirically outperform VIN on a variety of metrics such as learning speed, hyperparameter sensitivity, iteration count, and even generalization. Furthermore, we show that this performance gap is consistent across different maze transition types, maze sizes and even show success on a challenging 3D environment, where the planner is only provided with first-person RGB images.

1. Introduction

A common type of sub-task that arises in various reinforcement learning domains is path finding: finding a shortest set of actions to reach a subgoal from some starting state. Path finding is a fundamental part of any application which requires navigating in an environment, such as robotics (Ackerman & Guizzo, 2015) and video game AI (Silver, 2005). Due to its ubiquity in these important applications, recent work (Tamar et al., 2017) has designed a differentiable submodule that performs path-finding as directed by the agent in some inner loop. These *Value Iteration Network (VIN)*

modules mimic the application of Value Iteration on a 2D grid world, but without a pre-specified model or reward function. VINs were shown to be capable of computing near-optimal paths in 2D mazes and 3D landscapes where the transition model $P(s'|s, a)$ was not provided a priori and had to be learned.

In this paper, we show that VINs are often plagued by training instability, oscillating between high and low performance between epochs; random seed sensitivity, often converging to different performances depending on the random seed that was used; and hyperparameter sensitivity, where relatively small changes in hyperparameters can cause diverging behaviour. Owing to these optimization difficulties, we reframe the VIN as a recurrent-convolutional network, which enables us to replace the unconventional recurrent VIN update (convolution & max-pooling) with well-established gated recurrent operators such as the LSTM update (Hochreiter & Schmidhuber, 1997). These *Gated Path Planning Networks (GPPNs)* are a more general model that relaxes the architectural inductive bias of VINs that was designed to perform a computation resembling value-iteration.

We then establish empirically that GPPNs perform better or equal to the performance of VINs on a wide variety of 2D maze experiments, including different transition models, maze sizes and different training dataset sizes. We further demonstrate that GPPNs exhibit fewer optimization issues than VINs, including reducing random seed and hyperparameter sensitivity and increasing training stability. GPPNs are also shown to work with larger kernel sizes, often outperforming VINs with significantly fewer recurrent iterations, and also learn faster on average and generalize better given less training samples. Finally, we present results for both VIN and GPPN on challenging 3D ViZDoom environments (Kempka et al., 2016), where the planner is only provided with first-person RGB images instead of the top-down 2D maze design.

2. Background

In reinforcement learning, the environment is formulated as a Markov decision process (MDP) consisting of states s , actions a , a reward function \mathcal{R} , and state transition kernels $P(s' | s, a)$. *Value iteration* is a method of

^{*}Equal contribution ¹Carnegie Mellon University, Machine Learning Department. Correspondence to: Lisa Lee <lsllee@cs.cmu.edu>, Emilio Parisotto <eparisot@cs.cmu.edu>.

computing an optimal policy π and its value $V^\pi(s) = \mathbb{E}^\pi [\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(s_t, a_t, s_{t+1}) \mid s_0 = s]$, where $\gamma \in [0, 1]$ is a discount factor and $\mathcal{R}(s_t, a_t, s_{t+1})$ is a reward function. More specifically, value iteration starts with an arbitrary function $V^{(0)}$ and iteratively computes:

$$Q^{(k)}(s, a) = \sum_{s'} P(s' \mid s, a) \left(\mathcal{R}(s, a, s') + \gamma V^{(k-1)}(s') \right),$$

$$V^{(k)}(s) = \max_a Q^{(k)}(s, a).$$

The value function $V^{(k)}$ converges to V^* in the limit as $k \rightarrow \infty$, and the optimal policy can be recovered as $\pi^*(s) := \arg \max_a Q^{(\infty)}(s, a)$ (Sutton & Barto, 2018).

Despite the theoretical guarantees, value iteration requires a pre-specified environment model. Tamar et al. (2017) introduced the *Value Iteration Network (VIN)*, which is capable of learning these MDP parameters from data automatically. The VIN reformulates value iteration as a recursive process of applying convolutions and max-pooling over the feature channels:

$$\bar{Q}_{\bar{a}, i', j'}^{(k)} = \sum_{i, j} \left(W_{\bar{a}, i, j}^R \bar{R}_{i' - i, j' - j} + W_{\bar{a}, i, j}^V \bar{V}_{i' - i, j' - j}^{(k-1)} \right),$$

$$\bar{V}_{i, j}^{(k)} = \max_{\bar{a}} \bar{Q}_{\bar{a}, i, j}^{(k)}, \quad (1)$$

where the indices $i, j \in [m]$ correspond to cells in the $m \times m$ maze, $\bar{R}, \bar{Q}, \bar{V}$ is the VIN estimated reward, action-value and value functions, respectively, \bar{a} is the action index of the \bar{Q} feature map, and W^R, W^V are the convolutional weights for the reward function and value function, respectively. In the following iteration, the previous value \bar{V} is stacked with \bar{R} for the convolution step.

Tamar et al. (2017) showed that VINs have much greater success at path planning than baseline CNN and feedforward architectures in a variety of 2D and graph-based navigation tasks. The demonstrated success of VIN has made it an important component of models designed to solve downstream tasks where navigation is crucial (Karkus et al., 2017; Gupta et al., 2017a;b). For example, Gupta et al. (2017a;b) designed a Deep RL agent to perform navigation within partially observable and noisy environments by combining a VIN module with a 2D-structured memory map.

3. Method

In this work, we explore whether the inductive biases provided by the VIN are even necessary: is it possible that using alternative, more general architectures might work better than those of the VIN? We can view the VIN update (1) within the perspective of a convolutional-recurrent network, updating a recurrent state $V_{i', j'}^{(k)}$ at every spatial position

(i', j') in each iteration:

$$\bar{V}_{i', j'}^{(k)} = \max_{\bar{a}} \left(\sum_{i, j} W_{\bar{a}, i, j}^R \bar{R}_{i' - i, j' - j} + W_{\bar{a}, i, j}^V \bar{V}_{i' - i, j' - j}^{(k-1)} \right)$$

$$= \max_{\bar{a}} \left(W_{\bar{a}}^R \bar{R}_{[i', j', 3]} + W_{\bar{a}}^V \bar{V}_{[i', j', 3]}^{(k-1)} \right), \quad (2)$$

where $X_{[i', j', F]}$ denotes the image patch centered at position (i', j') with kernel size F . From (2), it can be seen that VIN follows the standard recurrent neural network (RNN) update where the recurrent state is updated by taking a linear combination of the input \bar{R} and the previous recurrent state $\bar{V}^{(k-1)}$, and passing their sum through a nonlinearity $\max_{\bar{a}}$. The main differences from a standard RNN are the following: the non-conventional nonlinearity (channel-wise max-pooling) used in VIN; the hidden dimension of the recurrent network, which is essentially one; the sparse weight matrices, where the non-zero values of the weight matrices represent neighboring inputs and units which are local in space; and the restriction of kernel sizes to 3.

Under this perspective, it is easy to question whether the adherence to these strict architectural biases is even necessary, given the long history of demonstrations that standard non-gated recurrent operators are difficult to optimize due to effects such as vanishing and exploding gradients (Pascanu et al., 2013).

We can easily replace the recurrent VIN update in (2) with the well-established LSTM update (Hochreiter & Schmidhuber, 1997), whose gated update alleviates many of the problems with standard recurrent networks:

$$h_{i', j'}^{(k)}, c_{i', j'}^{(k)} = \text{LSTM} \left(\sum_{\bar{a}} \left(W_{\bar{a}}^R \bar{R}_{[i', j', F]} + W_{\bar{a}}^h h_{[i', j', F]}^{(k-1)} \right), c_{i', j'}^{(k-1)} \right), \quad (3)$$

where F is the convolution kernel size. This recurrent update (3) still maintains the convolutional properties of the input and recurrent weight matrix as in VIN. It involves taking as input the $F \times F$ convolution of the input vector \bar{R} and previous hidden states $h^{(k-1)}$, and the previous cell state $c_{i', j'}^{(k-1)}$ of the LSTM at the central position (i', j') . We call path planning modules which use these gated updates *Gated Path Planning Networks (GPPNs)*. The GPPN is an LSTM which uses convolution of previous spatially-contiguous hidden states for its input.

4. Environments and Maze Transition Types

We test VIN and GPPN on 2D maze environments and 3D ViZDoom environments (Figure 1) on a variety of settings




Figure 1. (a) A sample 2D maze. (b) A sample 3D Doom maze and examples of screenshots showing the first-person view of the environment at three locations.

such as training dataset size, maze size and maze transition kernel.

We used three different maze transition kernels: In **NEWS**, the agent can move North, East, West, or South; in **Differential Drive**, the agent can move forward along its current orientation, or turn left/right by 90 degrees; in **Moore**, the agent can move to any of the eight cells in its Moore neighborhood. In the NEWS and Moore transition types, the target is an x-y coordinate, while in Differential Drive the target contains an orientation along with the x-y coordinate. Consequently, the dimension of the goal map given as input to the models is $1 \times m \times m$ for NEWS and Moore, and $4 \times m \times m$ for Differential Drive, where m is the maze size.

4.1. 2D Maze Environment

The 2D maze environment is created with a maze generation process that uses Depth-First Search with the Recursive Backtracker algorithm (Maze Generation Algorithms, 2018) to construct the maze tree, resulting in a fully connected maze (see Figure 1a). For each maze, we sample a probability d uniformly from $[0,1]$. Then for each wall, we delete the wall with probability d .

For our experiments on the 2D mazes, the state vector consists of the maze and the goal location, each of which are represented by a binary $m \times m$ matrix, where $m \times m$ is the maze size. We use early stopping based on validation set metrics to choose the final models.

4.2. 3D ViZDoom Environment

We use the Doom Game Engine and the ViZDoom API (Kempka et al., 2016) to create mazes in a simulated 3D environment (see Figure 1b). The maze design for the 3D mazes are generated in exactly the same manner as the 2D mazes, using Depth-First Search with the Recursive Backtracker algorithm followed by wall pruning with a uniformly sampled probability d . For each Doom maze, we take RGB screenshots showing the first-person view of the environment at each position and orientation. A sample

3D Doom maze and example screenshot images are shown in Figure 1. For an $m \times m$ maze with 4 orientations, this results in a total of $4m^2$ images.

In the 3D ViZDoom experiments, these map images are given as input to the model (instead of the 2D map design). This setup is similar to the one used for localization experiments by Chaplot et al. (2018) who argue that these images are easier to obtain as compared to constructing an accurate map design of an environment in the real world. The model needs to learn to infer the map design from these images along with learning to plan, which makes the task more challenging in 3D environments.

5. Experiments & Discussion

In this section, we empirically compare VIN and GPPN using two metrics: **%Optimal** (%Opt) is the percentage of states whose predicted paths under the policy estimated by the model has optimal length, and **%Success** (%Suc) is the percentage of states whose predicted paths under the policy estimated by the model reach the goal state. The reported performance is on a held-out test split. In contrast with the metrics reported in (Tamar et al., 2017), we do not stochastically sample rollouts but instead evaluate and train the output policy of the models directly on all states simultaneously. This reduces optimization noise and makes it easier to tell whether difficulties with training are due to sampling noise or model architecture/capacity.

All analyses are based on 2D maze results, except in Section 5.8 where we discuss 3D ViZDoom results. In order to make comparison fair, we utilized a hidden dimension of 150 for GPPN and 600 for VIN, owing to the approximately $4\times$ increase in parameters a GPPN contains due to the 4 gates it computes. Unless otherwise noted, the results were obtained by doing a hyperparameter sweep of (K, F) over $K \in \{5, 10, 15, 20, 30\}$ and $F \in \{3, 5, 7, 9, 11\}$, and using a 25k/5k/5k train-val-test split. Other experimental details are deferred to the Appendix.

5.1. Varying Kernel Size F

One question that can be asked of the architectural choices of the VIN is whether the kernel size needs to be the same dimension as the true underlying transition model. The kernel size used in VIN was set to 3×3 with a stride of 1, which is sufficient to represent the true transition model when the agent can move anywhere in the Moore neighborhood, but it limits the rate at which information propagates spatially with each iteration. With a kernel size of 3×3 and stride of 1, the receptive field of a unit in the last iteration’s feature map increases with rate $(3 + 2K) \times (3 + 2K)$ where K is the iteration count, meaning that the maximum path length information travels scales directly with iteration count k .

Table 1. Test performance on 2D mazes of size 15×15 with **varying kernel sizes** F and best K setting for each F . Bold indicates best result across all F for each model and transition kernel. VIN performs worse with larger F , while GPPN is more robust when F is varied and actually works better with larger F .

Model	F	NEWS		Moore		Diff. Drive	
		%Opt	%Suc	%Opt	%Suc	%Opt	%Suc
VIN	3	93.4	93.5	90.5	91.3	98.4	99.1
VIN	5	93.9	94.1	96.3	96.6	96.4	98.6
VIN	7	92.7	93.0	95.1	95.6	92.2	96.2
VIN	9	86.8	87.8	92.0	93.0	91.2	95.2
VIN	11	87.6	88.3	92.7	93.8	87.9	93.8
GPPN	3	97.6	98.3	96.8	97.6	96.4	98.1
GPPN	5	98.6	99.0	98.4	99.1	98.7	99.5
GPPN	7	99.0	99.3	98.8	99.3	99.1	99.7
GPPN	9	99.0	99.4	98.8	99.3	99.3	99.7
GPPN	11	99.2	99.5	98.6	99.2	99.2	99.6

Therefore for long-term planning in larger environments, Tamar et al. (2017) designed a multi-scale variant called the Hierarchical VIN. Hierarchical VINs rely on downsampling the maps into multi-scale hierarchies, and then doing VIN planning and up-scaling, progressively growing the map until it regains its original, un-downsampled size.

Another potential method to do long-range planning without requiring a multi-scale hierarchy is to instead increase the kernel size. An increased kernel size would cause the receptive field to grow more rapidly, potentially allowing the models to require fewer iterations K before reaching well-performing policies. In this section, we sought to test out the feasibility of increasing the kernel size of VINs and GPPNs. These results are summarized in Table 1. All the models were trained with the best K setting for each F and transition kernel. From the results, we can clearly see that GPPN can handle training with larger F values, and moreover, GPPN often performs better than VIN with larger values of F . In contrast, we can observe that VIN’s performance drops significantly after its kernel size is increased more than 5, with its best performing settings being either 3 or 5 depending on the true transition model. These results show that GPPN can learn planning approximations that work with $F > 3$ much more stably than VIN, and could further suggest that GPPN can work as well as VIN with less iterations.

5.2. Varying Iteration Count K

Following the above results showing that GPPN benefits from increased F , we further evaluated the effect of varying both iteration count K and kernel size F on the VIN and GPPN models. Table 2 shows %Optimal and %Success results of VIN and GPPN on 15×15 2D mazes for different values of F and K . We can see from NEWS column in the table that GPPN with $F > 7$ can get results on par with the

best VIN model with only $K = 5$ iterations. This shows that GPPN can learn to more effectively propagate information spatially in a smaller number of iterations than VIN can, and outperforms VIN even when VIN is given a much larger number of iterations. Additionally, we can see that VIN has significant trouble learning when both K and F are large in the differential drive mazes and to a lesser extent in the NEWS mazes.

Table 3 shows the results of VIN and GPPN with varying iteration counts K and the best F setting for each K . Owing to the larger kernel size, GPPN with smaller number of iterations $K \leq 10$ can get results on par with the best VIN model. Generally, both models benefit from a larger K (assuming the best F setting is used).

5.3. Different Maze Transition Kernels

From Tables 1 and 3, we can observe the performance of VIN and GPPN across a variety of different underlying groundtruth transition kernels (NEWS, Moore, and Differential Drive). From these results, we can see that GPPN consistently outperforms VIN on all the transition kernel types. An interesting observation is that VIN does very well at Differential Drive, consistently obtaining high results, although GPPN still does better than or on par with VIN. The reasons why VIN is so well suited to Differential Drive are not clear, and a preliminary analysis of VIN’s feature weights and reward vectors did not reveal any intuition on why this is the case.

5.4. Effect of Dataset Size

A potential benefit of the stronger architectural biases of VIN might be that they can enable better generalization given less training data. In this section, we designed experiments that set out to test this hypothesis. We trained VINs and GPPNs on datasets with varying number of training samples for all three maze transition kernels, and the results are given in Table 4. We can see that GPPN consistently outperforms VIN across all dataset sizes and maze models. Interestingly, we can observe that the performance gap between VIN and GPPN is larger the less data there is, demonstrating the opposite effect to our hypothesis. This could suggest that the architectural biases do not in fact aid generalization performance, or that there is another problem, such as perhaps the difficulty of optimizing VIN, that overshadows the benefit that the inductive bias could potentially provide.

5.5. Random Seed and Hyperparameter Sensitivity

The hypothesis this section sought to verify was whether the particular recurrent-convolutional form of the VIN did indeed negatively affect its optimization, as many ungated

Table 2. Test performance on 2D mazes of size 15×15 with **varying kernel sizes F** and **iteration counts K** . “–” indicates the training diverged. GPPN outperforms VIN under best settings of (K, F) , indicated in bold. By utilizing a larger F , GPPN can learn to more effectively propagate information spatially in a smaller number of iterations ($K \leq 10$) than VIN can.

Model	K	%Opt for NEWS					%Opt for Moore					%Opt for Differential Drive				
		$F=3$	$F=5$	$F=7$	$F=9$	$F=11$	$F=3$	$F=5$	$F=7$	$F=9$	$F=11$	$F=3$	$F=5$	$F=7$	$F=9$	$F=11$
VIN	5	55.6	87.7	84.6	86.3	86.6	75.0	86.7	88.9	92.0	92.3	74.8	91.9	91.5	91.2	87.9
VIN	10	79.0	83.3	92.2	86.8	86.7	90.5	91.4	95.1	89.4	92.7	92.4	96.1	92.2	84.0	64.4
VIN	15	91.3	92.9	92.7	85.4	87.6	88.7	89.6	92.4	90.0	91.0	96.7	96.4	90.1	65.2	23.0
VIN	20	93.4	93.9	91.4	86.3	85.5	80.9	92.8	90.7	89.1	90.4	97.7	94.8	89.0	40.0	22.3
VIN	30	71.2	92.8	84.5	86.5	86.4	80.5	96.3	92.5	91.7	89.1	98.4	95.9	89.5	–	–
GPPN	5	66.2	86.5	90.8	92.4	93.0	75.9	90.4	93.4	93.9	94.1	62.4	82.3	88.6	90.1	91.2
GPPN	10	91.2	96.1	97.1	97.6	97.7	93.3	96.5	97.4	97.6	97.4	87.7	95.4	96.1	97.0	97.4
GPPN	15	95.3	98.1	98.5	98.3	98.8	96.1	97.7	98.1	98.1	98.3	93.5	97.1	97.8	97.7	99.0
GPPN	20	97.4	98.4	99.0	99.0	99.2	96.8	98.4	98.5	98.7	98.6	95.8	97.9	98.4	98.4	98.9
GPPN	30	97.6	98.6	99.0	98.6	98.8	98.0	98.4	98.8	98.8	98.4	96.4	98.7	99.1	99.3	99.2

Table 3. Test performance on 2D mazes of size 15×15 with **varying iteration counts K** and best F setting for each K . Bold indicates best result across all K for each model and transition kernel. Generally, increasing K improves performance.

Model	K	NEWS		Moore		Diff. Drive	
		%Opt	%Suc	%Opt	%Suc	%Opt	%Suc
VIN	5	87.7	88.4	92.3	93.3	91.9	95.8
VIN	10	92.2	92.5	95.1	95.6	96.1	97.9
VIN	15	92.9	93.0	92.4	93.9	96.7	98.3
VIN	20	93.9	94.1	92.8	94.0	97.7	98.8
VIN	30	92.8	93.2	96.3	96.6	98.4	99.1
GPPN	5	93.0	94.3	94.1	96.1	91.2	95.6
GPPN	10	97.7	98.4	97.6	98.4	97.4	98.8
GPPN	15	98.8	99.2	98.3	98.9	99.0	99.6
GPPN	20	99.2	99.5	98.7	99.2	98.9	99.5
GPPN	30	99.0	99.3	98.8	99.3	99.3	99.7

recurrent updates suffer from optimization problems including training instability and higher sensitivity to weight initialization and hyperparameters due to gradient scaling problems (Pascanu et al., 2013).

We test each architecture’s sensitivity to random seeds by running several experiments with the same hyperparameters but different random seeds, and measuring the variance in their final performance. These results are reported in Table 5. The results show that GPPN gets consistently lower variance than VIN over different random seed initializations, supporting the hypothesis that the LSTM update enables more training stability and easier optimization than the un-gated recurrent update in VIN.

We additionally test hyperparameter sensitivity in Figure 2. We take all the results obtained on a hyperparameter sweep over settings (K, F) where K was varied over $K \in \{5, 10, 15, 20, 30\}$ and F was varied over $F \in \{3, 5, 7, 9, 11\}$. We then rank these results, and the x-axis is the top- x ranked hyperparameter settings and the

Table 4. Test performance on 2D mazes of size 15×15 with **varying dataset sizes N** under best settings of (K, F) for each model. Both models improve with more training data (larger N). GPPN performs relatively better than VIN with less data, suggesting that the VIN architectural biases do not help generalization performance.

N	Model	NEWS		Moore		Diff. Drive	
		%Opt	%Suc	%Opt	%Suc	%Opt	%Suc
10k	VIN	90.3	90.6	88.1	90.5	97.5	98.4
10k	GPPN	97.8	98.6	97.6	98.4	98.0	99.4
25k	VIN	93.9	94.1	96.3	96.6	98.4	99.1
25k	GPPN	99.2	99.5	98.8	99.3	99.3	99.7
100k	VIN	97.3	97.3	97.1	97.5	98.9	99.4
100k	GPPN	99.9	99.9	99.7	99.8	99.9	99.9

corresponding y-axis is the average %Opt/%Suc of those x settings. This plot thus measures how stable the performance of the architecture is to hyperparameter changes as the number of hyperparameter settings we consider grows. Therefore, architectures whose average top- x ranked performance remains high and relatively flat demonstrates that good performance with the architecture can be obtained with many different hyperparameter settings. This suggests that these models are both easier to optimize and consistently better than alternatives, and higher performance was not due to a single lucky hyperparameter setting. We can see from the figures that the performance of GPPN is clearly both higher and more stable over hyperparameter settings than VIN.

In Figure 3, we plot the learning curves for VIN and GPPN on 2D mazes with varying K and F . These plots show that VIN’s performance often oscillates between epochs (especially for larger kernel sizes $F > 3$), while GPPN is much more stable. Learning curves for other experiments showing a similar result are included in the Appendix. The

Table 5. Mean and standard deviation %Opt after 30 epochs, taken over 3 runs, on 2D mazes of size 15×15 . Bold indicates best result across all K for each model and transition kernel. The results were obtained using the best setting of F for each K and dataset size 100k. GPPN exhibits lower variance between runs.

Model	K	NEWS %Opt				Diff. Drive %Opt			
		Train		Val.		Train		Val.	
		mean	std	mean	std	mean	std	mean	std
VIN	5	90.1	0.1	90.1	1.5	88.4	1.0	95.4	1.1
VIN	10	92.8	0.6	92.7	1.4	92.3	0.5	93.9	0.2
VIN	15	93.4	1.2	94.2	0.6	95.8	0.5	97.0	0.3
VIN	20	93.1	1.5	94.3	0.8	96.4	0.2	96.8	1.1
GPPN	5	95.5	0.2	95.2	<0.1	93.8	0.2	93.4	<0.1
GPPN	10	99.1	0.1	99.0	<0.1	98.7	0.1	98.2	0.2
GPPN	15	99.6	<0.1	99.6	<0.1	99.4	0.1	99.3	0.1
GPPN	20	99.7	<0.1	99.7	0.1	99.8	<0.1	99.7	0.1




Figure 2. The y-axis is the average Test %Opt (or %Suc) of the top- n hyperparameter settings (K, F) over $K \in \{5, 10, 15, 20, 30\}$ and $F \in \{3, 5, 7, 9, 11\}$. The results are on 2D mazes of size 15×15 . These plots measure how stable the performance of each model is to hyperparameter changes as we increase the number of hyperparameter settings considered. GPPN exhibits less hyperparameter sensitivity.

training stability of GPPN provides more evidence to the hypothesis that GPPNs are simpler to optimize than VINs and consistently outperform them.

5.6. Learning Speed

In this section, we examine whether VINs or GPPNs learn faster. To do this, we measure the number of training epochs (passes over the entire dataset) that it takes for each model to reach a specific %Opt for the first time. These results are reported in Table 6. We can see from this table that GPPN learns significantly faster, often reaching 95% within 5-6 epochs. Comparatively, VIN sometimes never reaches 95%, as is the case for the NEWS mazes, or it takes 2-5 times as many epochs. This is the case even on the Differential Drive mazes, where VIN takes 2-3 times longer to train despite also getting high final performance.

Table 6. The number of epochs it takes for each model to attain a certain %Opt (50%, 75%, 90%, 95%) on the validation set under best settings of (K, F). The results are on 2D mazes of size 15×15 . GPPN learns faster.

Model	NEWS				Moore				Diff. Drive			
	50	75	90	95	50	75	90	95	50	75	90	95
VIN	1	6	17	—	1	1	11	23	2	3	5	14
GPPN	1	1	3	5	1	1	3	5	1	2	3	6

Table 7. Test performance on 2D mazes with **varying maze sizes** $m \times m$ under best settings of (K, F) for each model. For the larger 28×28 maze, we train for 100 epochs and sweep over $K \in \{14, 28, 56\}$ to account for longer trajectories required to solve some mazes. GPPN performs better.

m	Model	NEWS		Moore		Diff. Drive	
		%Opt	%Suc	%Opt	%Suc	%Opt	%Suc
15	VIN	93.9	94.1	96.3	96.6	98.4	99.1
15	GPPN	99.2	99.5	98.8	99.3	99.3	99.7
28	VIN	93.0	93.2	95.0	95.8	93.8	96.8
28	GPPN	98.3	98.9	97.8	98.7	99.0	99.6

5.7. Larger Maze Size

To test whether the improved performance GPPN persists even on larger, more challenging mazes, we evaluated the models on a dataset of mazes of size 28×28 , and varied $K \in \{14, 28, 56\}$ (Table 7). We used a training dataset size of 25k. GPPN outperformed VIN by a significant margin (3-6% for %Opt and %Suc) for all cases except Diff. Drive 15×15 , where the gap was closer (GPPN 99.3 vs. VIN 98.3 for %Opt).

5.8. 3D ViZDoom Experiments

In the 3D ViZDoom experiments, the state vector consists of RGB images showing the first-person view of the environment at each position and orientation, instead of the top-down 2D maze design (represented by a binary $m \times m$ matrix) as in the 2D maze experiments. To process the map images, we use a Convolutional Neural Network (LeCun et al., 1989) consisting of two convolutional layers: first layer with 32 filters of size 8×8 and a stride of 4, and second layer with 64 filters of size 4×4 with a stride 2×2 , followed by a linear layer of size 256.¹ The 256-dimensional representation for all the 4 orientations at each location is concatenated to create a 1024-dimensional representation. These representations of each location are then stacked at the corresponding x-y coordinate to create a map representation of size $1024 \times m \times m$. The map representation is then

¹This architecture was adapted from a previous work which is shown to perform well at playing deathmatches in Doom (Lample & Chaplot, 2017).

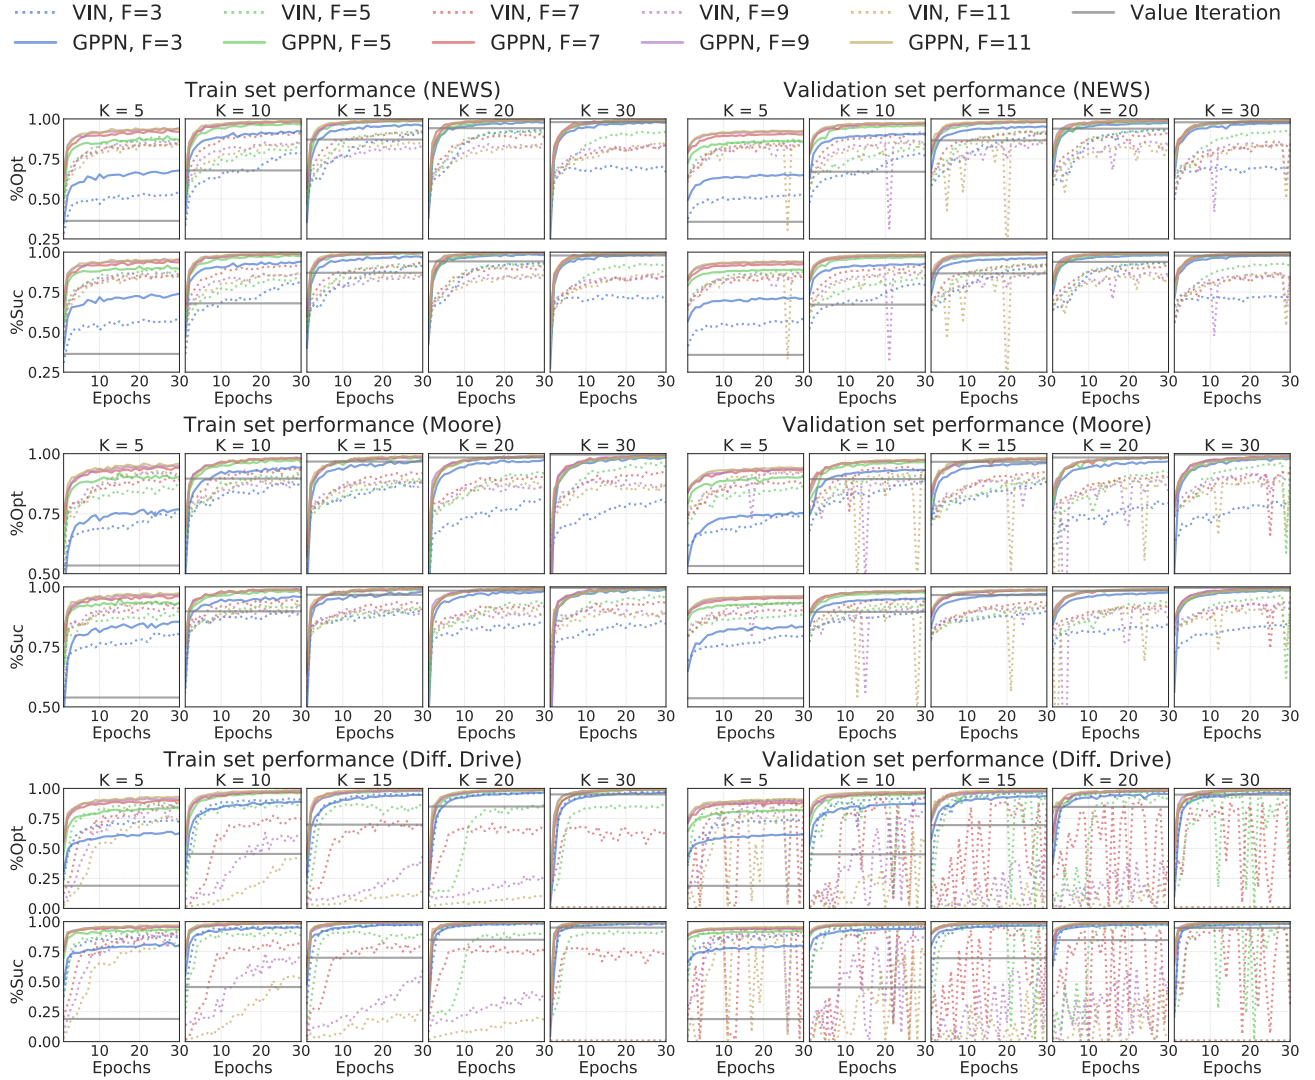


Figure 3. Performance on 2D mazes of size 15×15 with varying iteration counts K and kernel sizes F . All models are trained using dataset size 25k. VIN exhibits higher training instability, its performance often oscillating between epochs.

passed through two more convolutional layers (first layer with 64 filters and the second layer with 1 filter, both of size 3×3 and a stride of 1) to predict a maze design matrix of size $1 \times m \times m$, which is trained using an auxiliary binary cross-entropy loss. The predicted maze design is then stacked with the goal map and passed to the VIN or GPPN module in the same way as the 2D experiments.

The 3D ViZDoom results are summarized in Table 8. %Acc is the accuracy for predicting the top-down 2D maze design from first-person RGB images. Learning to plan in the 3D environments is more challenging due to the difficulty of simultaneously optimizing both the original planner loss and the auxiliary maze prediction loss. We can see that when %Acc is low, *i.e.*, the planner module must rely on a noisy

maze design, then the planner metrics %Opt and %Suc also suffer. We observe that VIN is more prone to overfitting on the training dataset: its validation %Acc is low ($< 91\%$) for all three transition kernels, whereas GPPN achieves higher validation %Acc on NEWS and Moore. However, GPPN also overfits on the Differential Drive.

6. Related Works

Karkus et al. (2017) looked at extending differentiable planning towards being able to plan in partially observable environments. In their setting, the agent is not provided a-priori with its position within the environment and thus needs to maintain a belief state over where it actually is. Similar to VIN’s differentiable extension of VI, the QMDP-Net archi-

Table 8. Performance on **3D ViZDoom mazes**. %Acc is accuracy for predicting the top-down 2D maze design from first-person RGB maze images. When %Acc is low, then the model must use a noisy maze design from which to plan, so %Opt and %Suc suffer as well. The results were obtained using $K = 30$, the best setting of F for each transition kernel, a smaller dataset size 10k (due to memory and time constraints), a smaller learning rate $5e-4$, and 100 training epochs. VIN is more prone to overfitting: its validation %Acc is low for all three transition kernels, while GPPN achieves higher validation %Acc on NEWS and Moore.

Kernel	Model	Train			Val			Test	
		%Acc	%Opt	%Suc	%Acc	%Opt	%Suc	%Opt	%Suc
NEWS	VIN	99.9	82.3	83.0	81.5	80.8	81.5	79.0	79.7
NEWS	GPPN	99.9	99.4	99.7	94.9	93.2	94.9	94.1	95.9
Moore	VIN	99.6	86.5	88.9	89.1	86.7	89.1	84.6	87.6
Moore	GPPN	99.6	98.1	99.4	97.4	95.3	97.4	94.5	97.2
Diff. Drive	VIN	100.0	99.4	99.7	90.5	89.0	90.5	96.9	97.9
Diff. Drive	GPPN	99.8	99.5	100.0	85.0	81.0	85.0	91.4	96.0

ecture was based on creating a differentiable analogue of the QMDP algorithm (Littman et al., 1995), an algorithm designed to approximate belief space planning in POMDPs. The architecture itself consisted of a filter module, which maintained the beliefs over which states the agent currently was in, and a planning module, which determined what action to take next. The planning module was essentially using a VIN to enable it to make more informed decisions on which parts of the environment to explore.

In recent work there has been a variety of deep reinforcement learning models that have examined combining an internal planning process with model-free methods. The Predictron (Silver et al., 2016) was a value function approximator which predicted a policy’s value by internally rolling out an LSTM forward predictive model of the agent’s future rewards, discounts and values. These future rewards, values and discounts were then accumulated together, with the idea that this would predict a more accurate value by forcing the architecture to model a multi-step rollout. A later extension, Value Predictive Networks (Oh et al., 2017), learnt a forward model that is used to predict the future rewards and values of executing a multi-step rollout. Although similar to the Predictron, they considered the control setting, where not only a value function had to be learnt but a policy as well. They demonstrated that their model, trained using model-free methods, was able to outperform existing methods on a 2D goal navigation task and outperformed DQN on Atari games.

Convolutional-recurrent networks similar to the VIN and GPPN have had a recent history of use within computer vision, particularly for applications which have both a spatial and temporal aspect. Convolutional LSTMs (ConvLSTMs) were first used in the application of precipitation nowcasting, where the goal was to predict rainfall intensity

within a region using past data (Shi et al., 2015). Recurrent-convolutional networks have also been used within computer vision applications where there is no explicit temporal aspect, such as object recognition. Feedback Networks (Zamir et al., 2017) utilized a ConvLSTM in order to allow information to feedback from higher layers to lower layers by unrolling the ConvLSTM over time. This enabled the Feedback Network to attain performance better than or on par with Residual Networks (ResNets) (He et al., 2016), one of the most commonly used feedforward architectures for object recognition.

A deeper connection has also been explored between residual and convolutional-recurrent networks. (Liao & Poggio, 2016) tested whether weight tying between layers in a ResNet significantly affects performance, finding that although performance slightly degrades, the change is not drastic. They provide some hypotheses on these results, suggesting that deep feedforward networks like ResNets are approximating recurrent networks in some capacity. While the GPPN can be seen as an instance of ConvLSTMs, our paper is the first to apply it to the domain of differentiable path planning and to show that, in general, structuring differentiable path planning within the context of convolutional-recurrent networks enables use of previous well-established recurrent architectures such as LSTM and GRUs.

7. Conclusion

In this work, we re-formulated VIN as a convolutional-recurrent network and designed a new planning module called the Gated Path Planning Network (GPPN) which replaced the unconventional recurrent update in VIN with a well-established gated LSTM recurrent operator. We presented experimental results comparing VIN and GPPN on 2D path-planning maze tasks and a 3D navigation task in the video game Doom, showing that the GPPN achieves results no worse and often better than VIN. The LSTM update alleviates many of the optimization issues including training instability and sensitivity to random seeds and hyperparameter settings. The GPPN is also able to utilize a larger kernel size, which the VIN is largely unable to do due to training instability, allowing the GPPN to work as well as VIN with fewer iterations. The GPPN also learns significantly faster, attaining high performance after only a few epochs, whereas the VIN takes longer to train. Finally, the relative performance improvement of GPPN over VIN increases with less training data. In conclusion, our analyses suggest that the inductive biases of VIN are not necessary in the design of a well-performing differentiable path planning module, and that the use of more general, gated recurrent architectures provides significant benefits over VINs.

Acknowledgements

LL is supported by a NSF GRFP Fellowship and by the CMU SEI under Contract FA8702-15-D-0002, Section H Clause, AFLCMC (H)-H001: 6-18014. EP, DC, and RS are supported in part by Apple, Nvidia NVAIL, DARPA D17AP00001, IARPA DIVA award D17PC00340, and ONR award N000141512791. The authors would also like to thank Nvidia for providing GPU support.

References

- Ackerman, E. and Guizzo, E. irobot brings visual mapping and navigation to the roomba 980, 2015.
- Chaplot, D. S., Parisotto, E., and Salakhutdinov, R. Active neural localization. In *Proceedings of the 6th International Conference on Learning Representations (ICLR 2018)*, 2018.
- Gupta, S., Davidson, J., Levine, S., Sukthankar, R., and Malik, J. Cognitive mapping and planning for visual navigation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2017)*, pp. 7272–7281, 2017a.
- Gupta, S., Fouhey, D. F., Levine, S., and Malik, J. Unifying map and landmark based representations for visual navigation. *CoRR*, abs/1712.08125, 2017b.
- Ha, D., Dai, A. M., and Le, Q. V. Hypernetworks. In *Proceedings of the 5th International Conference on Learning Representations (ICLR 2017)*, 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2016)*, pp. 770–778, 2016.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Karkus, P., Hsu, D., and Lee, W. S. QMDP-Net: Deep learning for planning under partial observability. In *Advances in Neural Information Processing Systems (NIPS 2017)*, pp. 4697–4707, 2017.
- Kempka, M., Wydmuch, M., Runc, G., Toczek, J., and Jaśkowski, W. ViZDoom: A Doom-based AI research platform for visual reinforcement learning. In *IEEE Conference on Computational Intelligence and Games*, pp. 341–348. IEEE, Sep 2016.
- Lample, G. and Chaplot, D. S. Playing fps games with deep reinforcement learning. In *AAAI*, pp. 2140–2146, 2017.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- Liao, Q. and Poggio, T. Bridging the gaps between residual learning, recurrent neural networks and visual cortex. *arXiv preprint arXiv:1604.03640*, 2016.
- Littman, M. L., Cassandra, A. R., and Kaelbling, L. P. Learning policies for partially observable environments: Scaling up. In *Machine Learning Proceedings 1995*, pp. 362–370. Elsevier, 1995.
- Maze Generation Algorithms. Maze generation algorithms — Wikipedia, the free encyclopedia, 2018. URL https://en.wikipedia.org/wiki/Maze_generation_algorithm#Recursive_backtracker. [Online; accessed 9-Feb-2018].
- Oh, J., Singh, S., and Lee, H. Value prediction network. In *Advances in Neural Information Processing Systems*, pp. 6120–6130, 2017.
- Pascanu, R., Mikolov, T., and Bengio, Y. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on Machine Learning (ICML 2013)*, pp. 1310–1318, 2013.
- Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., kin Wong, W., and chun Woo, W. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems (NIPS 2015)*, pp. 802–810, 2015.
- Silver, D. Cooperative pathfinding. In *AIIDE*, pp. 117–122, 2005.
- Silver, D., van Hasselt, H., Hessel, M., Schaul, T., Guez, A., Harley, T., Dulac-Arnold, G., Reichert, D., Rabinowitz, N., Barreto, A., et al. The predictron: End-to-end learning and planning. *arXiv preprint arXiv:1612.08810*, 2016.
- Sutton, R. and Barto, A. *Introduction to Reinforcement Learning*. MIT Press, 2nd edition, 2018.
- Tamar, A., Wu, Y., Thomas, G., Levine, S., and Abbeel, P. Value Iteration Networks. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI 2017)*, pp. 4949–4953, 2017.
- Zamir, A. R., Wu, T.-L., Sun, L., Shen, W. B., Shi, B. E., Malik, J., and Savarese, S. Feedback networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1808–1817. IEEE, 2017.