
Integrating Symmetry into Differentiable Planning

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 We study how group symmetry helps improve data efficiency and generalization
2 for end-to-end differentiable planning algorithms, specifically on 2D robotic path
3 planning problems: navigation and manipulation. We first formalize the idea from
4 Value Iteration Networks (VINs) on using convolutional networks for path plan-
5 ning, because it avoids explicitly constructing equivalence classes and enable end-
6 to-end planning. We then show that value iteration can always be represented as
7 some convolutional form for (2D) path planning, and name the resulting paradigm
8 Symmetric Planner (SymPlan). In implementation, we use steerable convolution
9 networks to incorporate symmetry. Our algorithms on navigation and manipula-
10 tion, with given or learned maps, improve training efficiency and generalization
11 performance by large margins over non-equivariant counterparts, VIN and GPPN.

1 Introduction

13 Model-based planning usually struggles in complex prob-
14 lems, and planning in more structured and abstract space
15 is a major solution [1, 2, 3, 4]. Symmetry is ubiquitous
16 in learning and decision-making problems and can effec-
17 tively reduce search space for planning. However, ex-
18 isting planning algorithms using symmetry assumes per-
19 fect dynamics knowledge, needs to explicitly build equiv-
20 alence classes, or does not consider problem structure
21 [5, 4, 6, 7, 8]. For example, if we use A* on path plan-
22 ning, we cannot specify visually obvious rotation sym-
23 metry in Figure 1, and need to detect in manually from
24 the provided dynamics model. This would be even more
25 challenging to detect in differentiable planning.

26 Nevertheless, symmetry in model-free deep reinforce-
27 ment learning (RL) has been studied recently [9, 10].
28 However, it can only *effectively* handle pixel-level
29 “element-wise” symmetry, such as flipping or rotating
30 state and action together. *Despite of this*, a critical bene-
31 fit of model-free RL agents that enables great asymptotic
32 performance is its end-to-end differentiability. This moti-
33 vates us to combine the spirit of both: *is it possible to de-
34 sign an end-to-end differentiable planning algorithm that
35 makes use of symmetry in environments?*

36 In this work, we propose to (1) avoid explicitly building equivalence classes for symmetric states
37 while (2) realize planning in an end-to-end differentiable manner. We are motivated by work in
38 the equivariant network and geometric deep learning community [11, 12, 13, 14, 15, 16], which




Figure 1: The *path planning* problem has symmetry, so we study how to *exploit* its symmetry in (differentiable) *planning*. Red dots are goal. The optimal actions $A = \text{SymPlan}(M)$ (bottom row) for the maps M (top row) are guaranteed to be equivariant $\text{SymPlan}(g.M) = g.\text{SymPlan}(M)$ under \odot rotations for (2D) path planning. For example, the action in the NW corner of A is the same as the action in the SW corner of $g.A$, after also rotating the arrow $\odot 90^\circ$.

39 treat an RGB image as a mapping $\mathbb{Z}^2 \rightarrow \mathbb{R}^3$ and apply equivariant convolutions between feature
 40 maps. It satisfies our desiderata: equivariant networks on images do not need to explicitly consider
 41 “symmetric pixels” while guarantee symmetry properties. Based on the intuition, we propose a
 42 framework, *Symmetric Planning* (SymPlan), to understand a straightforward but general problem,
 43 path planning, as operating like images, called steerable feature fields [14, 16]. We focus on 2D grid
 44 and prove that value iteration (VI) for 2D path planning is equivariant under the *isometries* of \mathbb{Z}^2 :
 45 translations, rotations, and reflections, and further show that VI here is a special form of steerable
 46 convolution network [14]. This provides us a foundation to equip Value Iteration Network (VIN,
 47 [17]) with steerable convolution. We implement the equivariant steerable version of VIN, named
 48 SymVIN, and use a variant, GPPN [18], to build SymGPPN. Both SymPlan methods achieve great
 49 improvement on training efficiency and generalization performance to unseen random maps, which
 50 showcases the advantage of exploiting symmetry from environments for planning.

51 Our contributions are as follows:

- 52 • Understand the inherent symmetry in path planning problems (on 2D grids), formulate value iter-
 53 ation in as steerable convolution network, and connect both to incorporate symmetry into VI.
- 54 • Based on the formulation, implement equivariant steerable version of VIN and GPPN.
- 55 • Show significant improvement in training and generalization on 2D navigation and manipulation.

56 2 Related work

57 **Planning with symmetries (Symmetric Planning).** Symmetries widely exist in various domains,
 58 and have been exploited in classic planning algorithms as well as model checking [5, 4, 6, 19, 20,
 59 21, 22, 23, 24, 25, 26, 27, 28, 29, 30]. Zinkevich and Balch [7] show the invariance of value function
 60 for an MDP with symmetry. Narayanamurthy and Ravindran [8] prove that finding exact symmetry
 61 in MDPs is graph isomorphism complete. However, they are based on classic planning algorithms,
 62 such as A*, and have a fundamental issue with exploitation of symmetries: they explicitly construct
 63 equivalence classes of symmetric states, which explicitly represents states and introduces symmetry
 64 breaking. Therefore, they are intractable (NP-hard) in maintaining symmetries in trajectory rollout
 65 and forward search (for large state space and symmetry group) and incompatible with differentiable
 66 pipelines for representation learning, hindering it from wider applications in RL and robotics.

67 **State abstraction for detecting symmetries.** Coarsest state abstraction aggregates all symmetric
 68 states into equivalence classes, studied in MDP homomorphisms and bisimulation [3, 31, 2]. How-
 69 ever, they usually require *perfect* MDP dynamics knowledge and do not scale up well, because of the
 70 complexity in maintaining abstraction mappings (homomorphisms) and abstracted MDPs. van der
 71 Pol et al. [32] integrate symmetry into model-free RL based on MDP homomorphisms [3], which
 72 avoids the challenges in handling symmetry in forward search. Park et al. [33] learn equivariant
 73 transition models, but do not consider planning. Additionally, the formulation in commonly defined
 74 symmetric MDPs [3, 9, 6, 7] is different from our symmetry formulation for path planning, since
 75 they study “element-wise” symmetry for every state-action pairs and require reward to be symmet-
 76 ric. Our reward is not symmetric and we mainly study symmetry of the underlying domain (2D
 77 grid), as further discussed in Section F.2.

78 **Symmetries and equivariance in deep learning.** Equivariant neural networks are used to incor-
 79 porate symmetry in supervised learning for different domains (e.g. grid and sphere), symmetry
 80 groups (e.g. translations and rotations), and group representation on feature spaces [34]. Cohen and
 81 Welling [15] introduce G-CNNs, followed by Steerable CNNs [14] which generalizes from scalar
 82 feature fields to vector fields with induced representations. Kondor and Trivedi [13], Cohen et al.
 83 [12] study theory on the relation between equivariant maps and convolutions. Weiler and Cesa [16]
 84 propose to solve kernel constraints under arbitrary representations for $E(2)$ and its subgroups by
 85 decomposing into irreducible representations, named $E(2)$ -CNN.

86 **Differentiable planning.** Our pipeline is based on learning to plan in a neural network in a differ-
 87 entiable manner. Value iteration network (VIN) [35] is a representative work that performs value
 88 iteration using convolution on lattice grids, and has been further extended [36, 18, 37, 38]. Other than
 89 using convolution network, works on integrating learning and planning into differentiable networks
 90 include [39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49]. In the theoretical side, Grimm et al. [50, 51]
 91 propose to understand the differentiable planning algorithms from value equivalence perspective,
 92 while Gehring et al. [52] study its gradient dynamics.

93 3 Preliminaries

94 **Markov decision processes.** We model the path planning problems as Markov decision processes
 95 (MDP) [1]. An MDP is a 5-tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$, with state space \mathcal{S} , action space \mathcal{A} , transition
 96 probability function $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}_+$, reward function $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, and discount factor
 97 $\gamma \in [0, 1]$. Value functions $V : \mathcal{S} \rightarrow \mathbb{R}$ and $Q : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ represent expected future returns [1].

98 **Symmetry groups and equivariance.** A symmetry *group* is defined as a set G together with a binary
 99 composition map satisfying the axioms of associativity, identity, and inverse. A (left) *group action*
 100 of G on a set \mathcal{X} is defined as the mapping $(g, x) \mapsto g.x$ which is compatible with composition.
 101 Given a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ and G acting on \mathcal{X} and \mathcal{Y} , then f is *G-equivariant* if it commutes
 102 with group actions: $g.f(x) = f(g.x), \forall g \in G, \forall x \in \mathcal{X}$. In the special case the action on \mathcal{Y} is trivial
 103 $g.y = y$, then $f(x) = f(g.x)$ holds, and we say f is *G-invariant*.

104 **Group representations.** We mainly use two groups: dihedral group D_4 and cyclic group C_4 . The
 105 cyclic group of 4 elements is $C_4 = \langle r \mid r^4 = 1 \rangle$, a symmetry group of rotating a square. The
 106 dihedral group $D_4 = \langle r, s \mid r^4 = s^2 = (sr)^2 = 1 \rangle$ includes both rotations r and reflections s , and
 107 has size $|D_4| = 8$. A group representation defines how a group action transforms a **vector space**
 108 $G \times \mathcal{S} \rightarrow \mathcal{S}$. These groups have three types of representations of our interest: *trivial*, *regular*, and
 109 *quotient* representations, see [16]. The *trivial representation* ρ_{triv} maps each $g \in G$ to 1 and hence
 110 fixes all $s \in \mathcal{S}$. The *regular representation* ρ_{reg} of C_4 group sends each $g \in C_4$ to a 4×4 permutation
 111 matrix that cyclically permutes a 4-element vector, such as a one-hot 4-direction action. The regular
 112 representation of D_4 maps each element to an 8×8 permutation matrix which does not act on 4-
 113 direction actions, which requires the *quotient representations* (quotienting out *sr² reflection part*)
 114 and forming a 4×4 permutation matrix. It is worth mentioning the *standard representation* of the
 115 cyclic groups, which are 2×2 rotation matrices, only used for visualization (Figure 2 middle).

116 **Steerable feature fields and Steerable CNNs.** The concept of *feature fields* is used in (equivariant)
 117 CNNs [11, 12, 13, 14, 15, 16]. The pixels of an 2D RGB image $x : \mathbb{Z}^2 \rightarrow \mathbb{R}^3$ on a domain $\Omega = \mathbb{Z}^2$
 118 is a feature field. In steerable CNNs for 2D grid, features are formed as *steerable feature fields*
 119 $f : \mathbb{Z}^2 \rightarrow \mathbb{R}^C$ that associate a C -dimensional feature vector $f(x) \in \mathbb{R}^C$ to each element on a
 120 base space, such as \mathbb{Z}^2 . Defined like this, we know how to transform a steerable feature field and
 121 also the feature field after applying CNN on it, using some group [14]. The type of CNNs that
 122 operates on steerable feature fields is called Steerable CNN [14], which is equivariant to groups
 123 including *translations* as subgroup $(\mathbb{Z}^2, +)$, extending [15]. It needs to satisfy a *kernel steerability*
 124 constraint, where the \mathbb{R}^2 and \mathbb{Z}^2 cases are considered in [16]. We consider the 2D grid as our domain
 125 $\Omega = \mathcal{S} = \mathbb{Z}^2$ and use $G = p4m$ group as the running example. The group $p4m = (\mathbb{Z}^2, +) \rtimes D_4$
 126 (wallpaper group) is semi-direct product of discrete translation group \mathbb{Z}^2 and dihedral group D_4 , see
 127 [15, 14]. We visualize the *transformation law* of $p4m$ on a feature field on $\Omega = \mathbb{Z}^2$ in **Figure 2**
 128 **(Middle)**, usually referred as *induced representation* [14, 16]. Additional details in Section G.

129 **Planning as convolution.** The core component behind dynamic programming (DP) based algo-
 130 rithms in planning or reinforcement learning is *Bellman (optimality) equation* [53, 1]: $V(s) =$
 131 $\max_a R(s, a) + \gamma \sum_{s'} P(s'|s, a) V(s')$. Value Iteration (VI) iteratively applies Bellman operator
 132 and converges to fixed points [1, 53]. The key component of our interest is $\sum_{s'} P(s'|s, a) V(s')$
 133 that aggregates values $V(s')$ from adjacent states by expectation using transition probabilities, here
 134 referred as **expected value operation**. Tamar et al. [17] propose Value Iteration Network (VIN) that
 135 uses convolution (networks) for planning, as an instance of differentiable planning, by recursively
 136 applying planar convolutions and max-pooling over feature spaces on 2D grid \mathbb{Z}^2 .

137 4 Symmetric Planning Framework

138 This section formulates the notion of Symmetric Planning (SymPlan). We expand the understanding
 139 of path planning in neural networks by planning as convolution on steerable feature fields (*steerable*
 140 *planning*). We use that to build *steerable value iteration* and show it is equivariant.

141 4.1 Steerable Planning: planning on steerable feature fields

142 We start the discussion based on Value Iteration Networks (VINs, [17]) and use a running example
 143 of planning on the 2D grid \mathbb{Z}^2 . We aim to understand (1) how VIN-style networks embed planning




Figure 2: **(Left)** Construction of spatial MDPs from path planning problems, enabling G -invariant transition. **(Middle)** The group acts on a feature field (MDP actions). We need to find the element in the original field by $f(r^{-1}x)$, and also rotate the arrow by $\rho(r)$, where $r \in D_4$. We represent one-hot actions as arrows (vector field, using ρ_{std}) for visualization. **(Right)** Equivariance of $V \mapsto Q$ in Bellman operator on feature fields, under $\circlearrowleft 90^\circ \in C_4$ rotation, which visually explains Theorem 4.1. The example simulates VI for one step (see red circles; minus signs omitted) with true transition P using \circlearrowleft N-W-S-E actions. The Q -value field are for 4 actions and can be viewed as either $\mathbb{Z}^2 \rightarrow \mathbb{R}^4$ ([14, 16]) or $\mathbb{Z}^2 \times C_4 \rightarrow \mathbb{R}$ (on $p4$ group, [15]). See Appendix H for more details.

144 and how its idea generalizes, (2) how is symmetry structure defined in path planning and how could
 145 it be injected into such planning networks.

146 **Constructing G -invariant transition: spatial MDP.** Intuitively, the embedded MDP in a VIN
 147 is different from the original path planning problem, since (planar) convolutions are translation
 148 equivariant but there are different obstacles in different regions.

149 We found the key insight in VINs is that it implicitly uses an MDP that has translation equivariance.
 150 The core idea behind the construction is that it converts *obstacles* (encoded in transition probability
 151 P , by *blocking*) into “traps” (encoded in reward \bar{R} , by $-\infty$ reward). This allows to use planar con-
 152 volutions with translation equivariance, and also enables use to further use steerable convolutions.

153 The demonstration of the idea is shown in **Figure 2 (Left)**. We call it *spatial MDP*, with different
 154 transition and reward function $\bar{\mathcal{M}} = \langle \mathcal{S}, \mathcal{A}, \bar{P}, \bar{R}_m, \gamma \rangle$, which converts the “complexity” in the
 155 transition function P in \mathcal{M} to the reward function \bar{R}_m in $\bar{\mathcal{M}}$. The state and action space are kept
 156 the same: state $\mathcal{S} = \mathbb{Z}^2$ and action $\mathcal{A} \subset \mathbb{Z}^2$ to move Δs in four directions in a 2D grid. We provide
 157 the detailed construction of the spatial MDP in Appendix H.1.

158 **Steerable features fields.** We generalize the idea from VIN, by viewing functions (in RL and
 159 planning) as *steerable feature fields*, motivated by [11, 12, 14]. This is analogous to pixels on images
 160 $\Omega \rightarrow [255]^3$, and would allow us to apply convolution on it. The state value function is expressed
 161 as a field $V : \mathcal{S} \rightarrow \mathbb{R}$, while the Q -value function needs a field with $|\mathcal{A}|$ channels: $Q : \mathcal{S} \rightarrow \mathbb{R}^{|\mathcal{A}|}$.
 162 Similarly, a policy field¹ has probability logits of selecting $|\mathcal{A}|$ actions. For the transition probability
 163 $P(s'|s, a)$, we can use action to index it as $P^a(s'|s)$, similarly for reward $R^a(s)$. The next section
 164 will show that we can convert the transition function to field and even convolutional filter. Additional
 165 details are in Appendix H.

166 4.2 Symmetric Planning: symmetry by equivariance

167 The seemingly slight change in the construction of spatial MDPs brings important symmetry struc-
 168 ture. The general idea in exploiting symmetry in path planning is to use *equivariance* to avoid
 169 explicitly constructing equivalence classes of symmetric states. To this end, we construct value
 170 iteration over steerable feature fields, and show it is *equivariant* for path planning.

171 In VIN, the convolution is over 2D grid \mathbb{Z}^2 , which is symmetric under D_4 (rotations and reflections).
 172 However, we also know that VIN is already equivariant under translations. To consider all symme-
 173 tries, as in [14, 16], we understand the group $p4m = G = B \times H$ as constructed by a *base space*

¹We avoid the symbol π for policy since it is used for induced representation in [14, 16].

174 $B = G/H = (\mathbb{Z}^2, +)$ and a fiber group $H = D_4$, which is a *stabilizer subgroup* that fixes the origin
 175 $\mathbf{0} \in \mathbb{Z}^2$. We could then formally study such symmetry in the spatial MDP, since we construct it to
 176 ensure that the transition probability function in \mathcal{M} is G -invariant. Specifically, we can uniquely
 177 decompose any $g \in \mathbb{Z}^2 \rtimes D_4$ as $t \in \mathbb{Z}^2$ and $r \in D_4$ (and translations act "trivially" on action), so

$$\bar{P}(s' | s, a) = \bar{P}(g.s' | g.s, g.a) \equiv \bar{P}((tr).s' | (tr).s, r.a), \quad \forall g = tr \in \mathbb{Z}^2 \rtimes D_4, \forall s, a, s'. \quad (1)$$

178 **Expected value operator as steerable convolution.** The equivariance property can be shown step-
 179 by-step: (1) *expected value operation*, (2) *Bellman operator*, and (3) *full value iteration*. First, we
 180 use G -invariance to prove that the expected value operator $\sum_{s'} P(s'|s, a)V(s')$ is equivariant.

181 **Theorem 4.1.** *If transition is G -invariant, the expected value operator E over \mathbb{Z}^2 is G -equivariant.*

182 The proof is in Appendix I.1 and visual understanding is in Figure 2 middle. However, this provides
 183 intuition but is inadequate since we do not know: (1) how to implement it with CNNs, (2) how to use
 184 multiple feature channels like VINs, since it shows for scalar-valued transition probability and value
 185 function (corresponding to trivial representation). To this end, we next prove that we can implement
 186 value iteration using steerable convolution with general steerable kernels.

187 **Theorem 4.2.** *If transition is G -invariant, there exists a (one-argument, isotropic) matrix-valued
 188 steerable kernel $P^a(s - s')$ (for every action), such that the expected value operator can be written
 189 as a steerable convolution and is G -equivariant:*

$$E^a[V] = P^a \star V, \quad [g.[P^a \star V]](s) = [P^{g.a} \star [g.V]](s), \quad \forall s \in \mathbb{Z}^2, \forall g \in \mathbb{Z}^2 \rtimes D_4. \quad (2)$$

190 The full derivation is provided in Appendix I. We write the transition probability as $P^a(s, s')$, and
 191 we show it only depends on *state difference* $P^a(s - s')$ (or *one-argument kernel* [12]) using G -
 192 invariance, which is the key step to show it is some *convolution*. Note that we use one kernel P^a
 193 for each action (four directions), and when the group acts on E , it also acts on the action $P^{g.a}$ (and
 194 state, so technically acting on $\mathcal{S} \times \mathcal{A}$). Additionally, if the steerable kernel also satisfies the D_4 -
 195 *steerability constraint* [16, 54], the steerable convolution is *equivariant* under $p4m = \mathbb{Z}^2 \rtimes D_4$. We
 196 can then extend VINs from \mathbb{Z}^2 translation equivariance to $p4m$ -equivariance (translations, rotations,
 197 reflections). The derivation follows the existing work on steerable CNNs [15, 14, 16, 12], while this
 198 is our goal: to justify the close connection between path planning and steerable convolutions.

199 **Steerable Bellman operator and value iteration.** We can now represent all operations in Bellman
 200 (optimality) operator on steerable feature fields over \mathbb{Z}^2 (or *steerable Bellman operator*) as follows:
 201

$$V_{k+1}(s) = \max_a R^a(s) + \gamma \times [P^a \star V_k](s), \quad (3)$$

202 where V, R^a, \bar{P}^a are steerable feature fields over \mathbb{Z}^2 . As for the operations, \max_a is (max) pooling
 203 (over group channel), $+, \times$ are point-wise operations, and \star is convolution. As the second step,
 204 the main idea is to prove every operation in Bellman (optimality) operator on steerable fields is
 205 equivariant, including the nonlinear \max_a operator and $+, \times$. Then, iteratively applying Bellman
 206 operator forms value iteration and is also equivariant, as shown below and proved in Appendix I.4.

207 **Proposition 4.3.** *For a spatial MDP with G -invariant transition, the optimal value function can be
 208 found through G -steerable value iteration.*

209 **Remark.** Our framework generalizes the idea behind VINs and enables us to understand its appli-
 210 cability and restrictions. More importantly, this allows us to integrate symmetry but avoid explicitly
 211 building equivalence classes and enables planning with symmetry in end-to-end fashion. We em-
 212 phasize that the *symmetry in spatial MDPs* is different from *symmetric MDPs* [7, 3, 9], since our
 213 reward function is *not G -invariant (if not conditioning on reward)*. Although we focus on \mathbb{Z}^2 , we
 214 can generalize to path planning on higher-dimensional or even continuous Euclidean spaces (like \mathbb{R}^3
 215 space [54] or spatial graphs in \mathbb{R}^3 [55]), and use *equivariant operations on steerable feature fields*
 216 (such as steerable convolutions, pooling, and point-wise non-linearities) from steerable CNNs. We
 217 refer the readers to Appendix H and to [15, 14, 56, 16] for more details.

218 5 Symmetric Planning in Practice

219 In this section, we discuss how to achieve Symmetric Planning on 2D grids with E(2)-steerable
 220 CNNs [16]. We focus on implementing symmetric version of value iteration, SymVIN, and gener-
 221 alize the methodology to make a symmetric version of a popular follow-up of VIN, GPPN [18].




Figure 3: Commutative diagram for the full pipeline of SymVIN on steerable feature fields over \mathbb{Z}^2 (every grid). If rotating the input map M by $\pi_M(g)$ of any g , the output action $A = \text{SymVIN}(M)$ is guaranteed to be transformed by $\pi_A(g)$, i.e. the entire steerable SymVIN is equivariant under induced representations π_M and π_A : $\text{SymVIN}(\pi_M(g)M) = \pi_A(g)\text{SymVIN}(M)$. We use stacked feature fields to emphasize that SymVIN supports direct-sum of representations beyond scalar-valued.

222 **Steerable value iteration.** We have showed that, value iteration for path planning problems on \mathbb{Z}^2
 223 consists of equivariant maps between steerable feature fields. It can be implemented as an equivari-
 224 ant steerable CNN, with recursively applying two alternating (equivariant) layers:

$$Q_k^a(s) = R_m^a(s) + \gamma \times [P_\theta^a \star V_k](s), \quad V_{k+1}(s) = \max_a Q_k^a(s), \quad s \in \mathbb{Z}^2, \quad (4)$$

225 where $k \in [K]$ indexes iteration, V_k, Q_k^a, R_m^a are steerable feature fields over \mathbb{Z}^2 output by equiv-
 226 ariant layers, P_θ^a is a learned kernel in neural network, and $+, \times$ are element-wise operations.

227 **Pipeline.** We follow the pipeline in VIN [17]. The commutative diagram for the full pipeline is
 228 shown in Figure 3. The path planning task is given by a $m \times m$ spatial binary obstacle occupancy
 229 map and one-hot goal map, represented as a feature field $M : \mathbb{Z}^2 \rightarrow \{0, 1\}^2$. For the iterative
 230 process $Q_k^a \mapsto V_k \mapsto Q_{k+1}^a$, the reward field R_M is predicted from map M (by a 1×1 convolution
 231 layer) and the value field V_0 is initialized as zeros. The network output is (logits of) planned actions
 232 for all locations², represented as $A : \mathbb{Z}^2 \rightarrow \mathbb{R}^{|\mathcal{A}|}$, predicted from the final Q-value field Q_K (by
 233 another 1×1 convolution layer). The number of iterations K and the convolutional kernel size F
 234 of P_θ^a are set based on map size M , and the spatial dimension $m \times m$ is kept consistent.

235 **Building Symmetric Value Iteration Networks.** Given the pipeline of VIN fully on steerable
 236 feature fields, we are ready to build equivariant version with $E(2)$ -steerable CNNs [16]. The idea
 237 is to replace every Conv2d with a steerable convolution layer between steerable feature fields, and
 238 associate the fields with proper fiber representations $\rho(h)$.

239 VINs use ordinary CNNs and can choose the size of intermediate feature maps. The design choices
 240 in steerable CNNs is the feature fields and fiber representations (or *type*) for every layer [14, 16].
 241 The main difference³ in steerable CNNs is that we also need to tell the network how to *transform*
 242 every *feature field*, by specifying *fiber representations*, as shown in Figure 3.

243 **Specification of input map and output action.** We first specify *fiber representations* for the input
 244 and output field of the network: map M and action A . For input **occupancy map and goal** $M : \mathbb{Z}^2 \rightarrow \{0, 1\}^2$,
 245 it does not D_4 to act on the 2 channels, so we use two copies of trivial representations
 246 $\rho_M = \rho_{\text{triv}} \oplus \rho_{\text{triv}}$. For **action**, the final action output $A : \mathbb{Z}^2 \rightarrow \mathbb{R}^{|\mathcal{A}|}$ is for logits of four actions
 247 $\mathcal{A} = (\text{north}, \text{west}, \text{south}, \text{east})$ for every location. If we use $H = C_4$, it naturally acts on
 248 the four actions (ordered \odot) by *cyclically* \odot *permuting* the \mathbb{R}^4 channels. However, since the D_4
 249 group has 8 elements, we need a *quotient representation*, see [16] and Appendix J.

250 **Specification of intermediate fields: value and reward.** Then, for the intermediate feature fields:
 251 Q-values Q_k , state value V_k , and reward R_m , we are free to choose fiber representations, as well as
 252 the width (number of copies). For example, if we want 2 copies of regular representation of D_4 , the
 253 feature field has $2 \times 8 = 16$ channels and the stacked representation is 16×16 (by direct-sum).

²Technically, it also includes values or actions for obstacles, since the network needs to learn to approximate the reward $R_M(s, \Delta s) = -\infty$ with enough small reward and avoid obstacles.

254 For the **Q -value field** $Q_k^a(s)$, we use representation ρ_Q and its size as C_Q . We need at least $C_A \geq |\mathcal{A}|$
 255 channels for all actions of $Q(s, a)$ as in VIN and GPPN, then stacked together and denoted as
 256 $Q_k \triangleq \bigoplus_a Q_k^a$ with dimension $Q_k : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_Q * C_A}$. Therefore, the representation is direct-sum
 257 $\bigoplus \rho_Q$ for C_A copies. The **reward** is implemented similarly as $R_M \triangleq \bigoplus_a R_M^a$ and must have same
 258 dimension and representation to add element-wisely. For **state value** field, we denote the choose as
 259 fiber representation as ρ_V and its size C_V . It has size $V_k : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_V}$. Thus, the steerable kernel
 260 is *matrix-valued* with dimension $P_\theta : \mathbb{Z}^2 \rightarrow \mathbb{R}^{(C_Q * C_A) \times C_V}$. In practice, we found using *regular*
 261 *representations* for all three works the best. It can be viewed as "augmented" state and is related to
 262 group convolution, detailed in Appendix J.

263 **Other operations.** We now visit the remained (equivariant) operations. (1) The max **operation** in
 264 $Q_k \mapsto V_{k+1}$. While we have showed the max operation in $V_{k+1}(s) = \max_a Q_k^a(s)$ is equivariant
 265 in Theorem 4.3, we need to apply max(-pooling) for all actions along the "representation channel"
 266 from stacked representations $C_A * C_Q$ to one C_Q . More details are in Appendix J.5. (2) The **final**
 267 **output layer** $Q_K \mapsto A$. After the final iteration, the Q -value field Q_k is fed into the policy layer
 268 with 1×1 convolution to convert the action logit field $\mathbb{Z}^2 \rightarrow \mathbb{R}^{|\mathcal{A}|}$.

269 **Extended method: Symmetric GPPN.** Gated path planning network (GPPN [18]) proposes to use
 270 LSTM to alleviate the issue of unstable gradient in VINs. Although it does not strictly follow value
 271 iteration, it still follows the spirit of steerable planning. Thus, we first obtained a fully convolutional
 272 variant of GPPN from [Redacted for anonymous review], called ConvGPPN. It replaces the MLPs
 273 in the original LSTM cell with convolutional layers, and then replaces convolutions with equivariant
 274 steerable convolutions, resulting in a fully equivariant SymGPPN. See Appendix J.3 for details.

275 **Extended tasks: planning on learned maps with mapper networks.** We consider two planning
 276 tasks on 2D grids: 2D navigation and 2-DOF manipulation. To demonstrate the ability of handling
 277 symmetry in differentiable planning, we consider more complicated state space input: visual nav-
 278 igation and workspace manipulation, and discuss how to use mapper networks to convert the state
 279 input and use end-to-end learned maps, as in [18, 37]. See Appendix J.2 for details.

280 6 Experiments

281 We experiment VIN, GPPN and our SymPlan methods on four path planning tasks, including using
 282 *given* or *learned* maps. The additional experiments and ablation studies are in Appendix E.

283 **Environments and datasets.** We demonstrate the idea in two major robotics tasks: *navigation*
 284 and *manipulation*. We focus on the 2D regular grid setting for path planning, as adopted in
 285 prior work [17, 18, 37]. For each task, we consider using either *given* (2D navigation and 2-
 286 DOF configuration-space manipulation) or *learned* maps (visual navigation and 2-DOF workspace
 287 manipulation). In the latter case, the planner needs to jointly learn a mapper that converts ego-
 288 centric panoramic images (visual navigation) or workspace states (workspace manipulation) into
 289 plannable loss, as in [18, 37]. In both cases, we randomly generate training, validation and test
 290 data of $10K/2K/2K$ maps for all map sizes, to demonstrate data efficiency and generalization
 291 ability of symmetric planning. Note that the test maps are unlikely to be symmetric to the train-
 292 ing maps by any transformation from the symmetry groups G . For all environments, the planning
 293 domain is the 2D regular grid $\mathcal{S} = \Omega = \mathbb{Z}^2$, and the action space is to move in $4 \odot$ directions⁴:
 294 $\mathcal{A} = (\text{north, west, south, east})$.

295 **Methods: planner networks.** We compare five planner methods, where two are our SymPlan
 296 version of their non-equivariant counterparts. Our equivariant implementation is based on *Value*
 297 *Iteration Networks* (VIN, [17]) and *Gated Path Planning Networks* (GPPN, [18]). We implement
 298 the equivariant version of VIN, named **SymVIN**. For GPPN, we first obtained a *fully convolutional*
 299 version, named **ConvGPPN** [Redacted for anonymous review], and furthermore **SymGPPN** with
 300 steerable CNNs. All methods use (equivariant) convolutions with *circular padding* in planning
 301 in configuration spaces for the manipulation tasks, except GPPN that is not fully convolutional.
 302 Chaplot et al. [37] propose SPT based on Transformers, while integrating symmetry to Transformers
 303 is beyond steerable convolutions, thus we do not consider it but still adopt some useful setup.

⁴Note that the MDP action space \mathcal{A} needs to be *compatible* with the group action $G \times \mathcal{A} \rightarrow \mathcal{A}$. Since the E2CNN package [16] uses *counterclockwise* rotations \odot as generators for rotation groups C_n , the action space needs to be *counterclockwise* \odot .




Figure 4: **(Left)** A visual navigation environment rendered from a randomly generated 7×7 maze **(Middle)**, where the hover is the visualization of four views at position $(5, 3)$. **(Right)** A 2-joint manipulation task in workspace (topdown) and configuration space (2 DOFs) in 18×18 resolution.




Figure 5: Training curves on **(Left)** 2D navigation with 10K of 15×15 maps and on **(Right)** 2DoFs manipulation with 10K of 18×18 maps in configuration space. Faded areas indicate standard error.

304 **Training and evaluation.** We report success rate and training curves over 3 seeds. The training
 305 process (on given maps) follows [17, 18], where we train 30 epochs with batch size 32, and use
 306 kernel size $F = 3$ by default. The gradient clip threshold is set to 5. The default batch size is 32,
 307 while we need to reduce for some GPPN variants, since LSTM consumes much more memory.

308 6.1 Planning on given maps

309 **Environmental setup.** In the **2D navigation** task, the map and goal are randomly generated, where
 310 the map size is $\{15, 28, 50\}$. In **2-DOF manipulation** in configuration space, we adopt the setting
 311 in [37] and train networks to take as input of configuration space, represented by two joints. We
 312 randomly generate 0 to 5 obstacles in the manipulator workspace. Then the 2 degree-of-freedom
 313 (DOF) configuration space is constructed from workspace and discretized into 2D grid with sizes
 314 $\{18, 36\}$, corresponding to bins of 20° and 10° , respectively. All methods are trained using the same
 315 network size, where for equivariant versions, we use *regular* representations for all layers, which has
 316 size $|D_4| = 8$. We keep the same parameters for all methods, so all equivariant convolution layers
 317 with *regular* representations will have higher embedding sizes. Due to memory constraint, we use
 318 $K = 30$ iterations for 2D maze navigation, and $K = 27$ for manipulation. We use kernel sizes
 319 $F = \{3, 5, 5\}$ for $m = \{15, 28, 50\}$ navigation, and $F = \{3, 5\}$ for $m = \{18, 36\}$ manipulation.

320 **Results.** We show the averaged test results for both 2D navigation and C-space manipulation tasks
 321 on generalizing to unseen maps (Table 1) and the training curves for all methods (Figure 5). For VIN
 322 series, our SymVIN is much better than the vanilla VIN in terms of generalization and training per-
 323 formance in both environments, which learns much faster and achieves almost perfect asymptotic
 324 performance. As for GPPN, we found the fully convolutional variant ConvGPPN actually works
 325 better than the original one in [18], especially in learning speed. However, SymVIN does fluctu-
 326 ate in some runs, which seems to come from initialization and label, further studied in Appendix.
 327 SymGPPN further boosts ConvGPPN and outperforms all other methods. One exception is GPPN
 328 learns poorly in C-space manipulation. For GPPN, the added circular padding in the convolution
 329 encoder leads to gradient vanishing problem.

330 Additionally, we found using regular representations (for D_4 or C_4) for state value $V : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_V}$
 331 (and for Q -value) works better than trivial representations. This is counterintuitive since we expect
 332 the V value to be scalar $\mathbb{Z}^2 \rightarrow \mathbb{R}$. One reason is that switching between regular (for Q) and trivial
 333 (for V) representation introduces unnecessary bottleneck. Depending on the choice of representa-
 334 tions, we implement different max-pooling, with details in Appendix J.5. We also empirically found
 335 using FC only in the final layer $Q_K \mapsto A$ helps stabilize the training a bit. The ablation study on
 336 this and more are in Appendix E.

Table 1: Averaged test success rate (%) for using 10K/2K/2K dataset for all four types of tasks.

Method (10K Data)	Navigation			Visual	Manipulation		
	15 × 15	28 × 28	50 × 50		18 × 18	36 × 36	Workspace
VIN	66.97	67.57	57.92	50.83	77.82	84.32	80.44
SymVIN	98.99	98.14	86.20	95.50	99.98	99.36	91.10
GPPN	96.36	95.77	91.84	93.13	2.62	1.68	3.67
ConvGPPN	99.75	99.09	97.21	98.55	99.98	99.95	89.88
SymGPPN	99.98	99.86	99.49	99.78	100.00	99.99	90.50

337 **Remark.** Two symmetric planners are both significantly better than their counterparts. Notably,
 338 we did not include any symmetric maps to the test data that symmetric planners would perform
 339 much better. There are several potential sources of advantages: (1) SymPlan allows parameter
 340 sharing across positions and maps and implicitly enables planning in a reduced space: every (s, a, s')
 341 seamlessly generalizes to $(g.s, g.a, g.s')$ for any $g \in G$, (2) thus it uses training data more efficiently,
 342 (3) it reduces the space of hypothesis class and facilitate generalization to unseen maps.

343 6.2 Planning on learned maps: simultaneously planning and mapping

344 **Environmental setup.** For **visual navigation**, we randomly generate maps using the same strategy
 345 as before, and then render four egocentric panoramic views for each location from produced 3D
 346 environments with *Gym-MiniWorld* [57], since it allows to generate 3D mazes with any layout. For
 347 $m \times m$ maps, all egocentric views for a map is represented by $m \times m \times 4$ RGB images. For
 348 **workspace manipulation**, we randomly generate 0 to 5 obstacles in workspace as before. We use a
 349 mapper network to convert the 96×96 workspace (image of obstacles) to the $m \times m$ 2 degree-of-
 350 freedom (DOF) configuration space (2D occupancy grid). In both environments, the setup is similar
 351 to Section 6.1, while we only use $m = 15$ maps but longer 100 epochs for visual navigation and
 352 $m = 18$ maps still with 30 epochs for workspace manipulation.

353 **Methods: mapper networks and setup.** For **visual navigation**, we follow the mapper network
 354 setup in [18]. A mapper network converts every image into a 256-dimensional embedding $m \times m \times$
 355 4×256 and then predicts map layout $m \times m \times 1$. For **workspace manipulation**, we use U-net [58]
 356 with residual-connection [59] as a mapper. See Section E for details.

357 **Results.** The results are also shown in Table 1, denoted as Visual (navigation, 15×15) and
 358 Workspace (manipulation, 18×18). In visual navigation, the trends are similar to 2D case: two
 359 symmetric planners both train much faster. Besides vanilla VIN, all approaches finally converge to
 360 near-optimal successful rate (around 95%), while the validation and test results show large gaps.
 361 SymGPPN has almost no generalization gap, while VIN does not generalize well to new 3D visual
 362 navigation environments. Our SymVIN improves test successful rate from less than 50% to 90%
 363 and is comparable with GPPN. Since the input is raw images and a mapper is used to learn end-to-
 364 end, it potentially causes one major source of generalization gap for some approaches. In workspace
 365 manipulation, the results are also analogous to C-space, while ours advantages over baselines are
 366 smaller. In our inspection, we found the mapper network is the bottleneck, since the mapping for
 367 obstacles from workspace to C-space is nontrivial to learn.

368 **Remark.** The SymPlan models demonstrate end-to-end planning and learning ability, potentially
 369 enabling further applications to other tasks as a differentiable component for planning. The addi-
 370 tional results and ablation studies are provided in Appendix E.

371 7 Discussion

372 In this work, we study the symmetry in 2D path planning problem, and build a framework using
 373 the theory of steerable CNNs to prove that value iteration in path planning is actually a form of
 374 steerable CNN (on 2D grids). Although we focus on \mathbb{Z}^2 , we can generalize to path planning on
 375 higher-dimensional or even continuous Euclidean spaces [54, 55], and use *equivariant operations* on
 376 *steerable feature fields* (such as steerable convolutions, pooling, and point-wise non-linearities) from
 377 steerable CNNs. We practically show that the SymPlan algorithms exactly motivated by the theory
 378 provide great improvement. We hope the framework along with the design of practical algorithms
 379 can provide a new pathway to exploiting symmetry structure in differentiable planning.

References

- 380
- 381 [1] Richard S. Sutton and Andrew G. Barto. *Reinforcement learning: an introduction*. Adaptive
382 computation and machine learning series. The MIT Press, Cambridge, Massachusetts, second
383 edition edition, 2018. ISBN 978-0-262-03924-6.
- 384 [2] Lihong Li, Thomas J. Walsh, and M. Littman. Towards a Unified Theory of State Abstraction
385 for MDPs. In *AI&M*, 2006.
- 386 [3] Balaraman Ravindran and Andrew G Barto. *An algebraic approach to abstraction in rein-*
387 *forcement learning*. PhD thesis, University of Massachusetts at Amherst, 2004.
- 388 [4] Maria Fox and Derek Long. Extending the exploitation of symmetries in planning. In *In*
389 *Proceedings of AIPS'02*, pages 83–91, 2002.
- 390 [5] Maria Fox and Derek Long. The Detection and Exploitation of Symmetry in Planning Prob-
391 lems. In *In IJCAI*, pages 956–961. Morgan Kaufmann, 1999.
- 392 [6] Nir Pochter, Aviv Zohar, and Jeffrey S. Rosenschein. Exploiting Problem Symmetries in State-
393 Based Planners. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*, August 2011.
394 URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI11/paper/view/3732>.
- 395 [7] Martin Zinkevich and Tucker Balch. Symmetry in Markov decision processes and its impli-
396 cations for single agent and multi agent learning. In *In Proceedings of the 18th International*
397 *Conference on Machine Learning*, pages 632–640. Morgan Kaufmann, 2001.
- 398 [8] Shravan Matthur Narayanamurthy and Balaraman Ravindran. On the hardness of finding
399 symmetries in Markov decision processes. In *Proceedings of the 25th international confer-*
400 *ence on Machine learning - ICML '08*, pages 688–695, Helsinki, Finland, 2008. ACM Press.
401 ISBN 978-1-60558-205-4. doi: 10/bkswc2. URL [http://portal.acm.org/citation.](http://portal.acm.org/citation.cfm?doid=1390156.1390243)
402 [cfm?doid=1390156.1390243](http://portal.acm.org/citation.cfm?doid=1390156.1390243).
- 403 [9] Elise van der Pol, Daniel E. Worrall, Herke van Hoof, Frans A. Oliehoek, and Max
404 Welling. MDP Homomorphic Networks: Group Symmetries in Reinforcement Learning.
405 *arXiv:2006.16908 [cs, stat]*, June 2020. URL <http://arxiv.org/abs/2006.16908>. arXiv:
406 2006.16908.
- 407 [10] Dian Wang, Robin Walters, and Robert Platt. $\mathrm{SO}(2)$ -Equivariant Reinforcement
408 Learning. September 2021. URL https://openreview.net/forum?id=7F9c0hdvfk_.
- 409 [11] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learn-
410 ing: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- 411 [12] Taco Cohen, Mario Geiger, and Maurice Weiler. A General Theory of Equivariant CNNs on
412 Homogeneous Spaces. *arXiv:1811.02017 [cs, stat]*, January 2020. URL [http://arxiv.](http://arxiv.org/abs/1811.02017)
413 [org/abs/1811.02017](http://arxiv.org/abs/1811.02017). arXiv: 1811.02017.
- 414 [13] Risi Kondor and Shubhendu Trivedi. On the Generalization of Equivariance and Convolution
415 in Neural Networks to the Action of Compact Groups. *arXiv:1802.03690 [cs, stat]*, November
416 2018. URL <http://arxiv.org/abs/1802.03690>. arXiv: 1802.03690.
- 417 [14] Taco S. Cohen and Max Welling. Steerable CNNs. November 2016. URL [https:](https://openreview.net/forum?id=rJQKYt511)
418 [//openreview.net/forum?id=rJQKYt511](https://openreview.net/forum?id=rJQKYt511).
- 419 [15] Taco S. Cohen and Max Welling. Group Equivariant Convolutional Networks.
420 *arXiv:1602.07576 [cs, stat]*, June 2016. URL <http://arxiv.org/abs/1602.07576>. arXiv:
421 1602.07576.
- 422 [16] Maurice Weiler and Gabriele Cesa. General $\mathrm{E}(2)$ -Equivariant Steerable CNNs.
423 *arXiv:1911.08251 [cs, eess]*, April 2021. URL <http://arxiv.org/abs/1911.08251>.
424 arXiv: 1911.08251.

- 425 [17] Aviv Tamar, Yi Wu, Garrett Thomas, Sergey Levine, and Pieter Abbeel. Value Iteration Net-
426 works. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial In-*
427 *telligence*, pages 4949–4953, Melbourne, Australia, August 2017. International Joint Confer-
428 ences on Artificial Intelligence Organization. ISBN 978-0-9992411-0-3. doi: 10/ggjfst. URL
429 <https://www.ijcai.org/proceedings/2017/700>.
- 430 [18] Lisa Lee, Emilio Parisotto, Devendra Singh Chaplot, Eric Xing, and Ruslan Salakhutdi-
431 nov. Gated Path Planning Networks. *arXiv:1806.06408 [cs, stat]*, June 2018. URL
432 <http://arxiv.org/abs/1806.06408>. arXiv: 1806.06408.
- 433 [19] Carmel Domshlak, Michael Katz, and Alexander Shleyfman. Enhanced Symmetry Breaking
434 in Cost-Optimal Planning as Forward Search. page 5.
- 435 [20] Alexander Shleyfman. Symmetry Breaking: Satisficing Planning and Landmark Heuristics.
436 page 5.
- 437 [21] Alexander Shleyfman, Michael Katz, Malte Helmert, Silvan Sievers, and Martin Wehrle.
438 Heuristics and Symmetries in Classical Planning. *Proceedings of the AAAI Conference on*
439 *Artificial Intelligence*, 29(1), March 2015. ISSN 2374-3468. URL [https://ojs.aaai.org/](https://ojs.aaai.org/index.php/AAAI/article/view/9649)
440 [index.php/AAAI/article/view/9649](https://ojs.aaai.org/index.php/AAAI/article/view/9649). Number: 1.
- 441 [22] Silvan Sievers, Martin Wehrle, Malte Helmert, Alexander Shleyfman, and Michael Katz. Fac-
442 tored Symmetries for Merge-and-Shrink Abstractions. page 8.
- 443 [23] Martin Wehrle, Malte Helmert, Alexander Shleyfman, and Michael Katz. Integrating Partial
444 Order Reduction and Symmetry Elimination for Cost-Optimal Classical Planning. page 7.
- 445 [24] Mohammad Abdulaziz, Michael Norrish, and Charles Gretton. Exploiting Symmetries by
446 Planning for a Descriptive Quotient. page 8.
- 447 [25] Silvan Sievers, Martin Wehrle, Malte Helmert, and Michael Katz. An Empirical Case Study
448 on Symmetry Handling in Cost-Optimal Planning as Heuristic Search. In Steffen Hölldobler,
449 Rafael Peñaloza, and Sebastian Rudolph, editors, *KI 2015: Advances in Artificial Intelligence*,
450 volume 9324, pages 166–180. Springer International Publishing, Cham, 2015. ISBN 978-
451 3-319-24488-4 978-3-319-24489-1. doi: 10.1007/978-3-319-24489-1_13. URL [http://](http://link.springer.com/10.1007/978-3-319-24489-1_13)
452 link.springer.com/10.1007/978-3-319-24489-1_13. Series Title: Lecture Notes in
453 Computer Science.
- 454 [26] Silvan Sievers. Structural Symmetries of the Lifted Representation of Classical Planning Tasks.
455 page 8.
- 456 [27] Dominik Winterer, Martin Wehrle, and Michael Katz. Structural Symmetries for Fully Ob-
457 servable Nondeterministic Planning. page 7.
- 458 [28] Gabriele Röger, Silvan Sievers, and Michael Katz. Symmetry-Based Task Reduction for
459 Relaxed Reachability Analysis. In *Twenty-Eighth International Conference on Automated*
460 *Planning and Scheduling*, June 2018. URL [https://aaai.org/ocs/index.php/ICAPS/](https://aaai.org/ocs/index.php/ICAPS/ICAPS18/paper/view/17772)
461 [ICAPS18/paper/view/17772](https://aaai.org/ocs/index.php/ICAPS/ICAPS18/paper/view/17772).
- 462 [29] Silvan Sievers, Gabriele Röger, Martin Wehrle, and Michael Katz. Theoretical Foundations
463 for Structural Symmetries of Lifted PDDL Tasks. *Proceedings of the International Conference*
464 *on Automated Planning and Scheduling*, 29:446–454, 2019. ISSN 2334-0843. URL <https://ojs.aaai.org/index.php/ICAPS/article/view/3509>.
- 465
- 466 [30] Daniel Fišer, Álvaro Torralba, and Alexander Shleyfman. Operator Mutexes and Symmetries
467 for Simplifying Planning Tasks. *Proceedings of the AAAI Conference on Artificial Intelligence*,
468 33(01):7586–7593, July 2019. ISSN 2374-3468. doi: 10.1609/aaai.v33i01.33017586. URL
469 <https://ojs.aaai.org/index.php/AAAI/article/view/4751>. Number: 01.
- 470 [31] N. Ferns, P. Panangaden, and Doina Precup. Metrics for Finite Markov Decision Processes. In
471 *AAAI*, 2004.

- 472 [32] Elise van der Pol, Daniel Worrall, Herke van Hoof, Frans Oliehoek, and Max Welling. Mdp
473 homomorphic networks: Group symmetries in reinforcement learning. *Advances in Neural*
474 *Information Processing Systems*, 33, 2020.
- 475 [33] Jung Yeon Park, Ondrej Biza, Linfeng Zhao, Jan Willem van de Meent, and Robin Walters.
476 Learning Symmetric Embeddings for Equivariant World Models. *arXiv:2204.11371 [cs]*, April
477 2022. URL <http://arxiv.org/abs/2204.11371>. arXiv: 2204.11371.
- 478 [34] Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric Deep Learn-
479 ing: Grids, Groups, Graphs, Geodesics, and Gauges. *arXiv:2104.13478 [cs, stat]*, April 2021.
480 URL <http://arxiv.org/abs/2104.13478>. arXiv: 2104.13478.
- 481 [35] Aviv Tamar, Yi Wu, Garrett Thomas, Sergey Levine, and Pieter Abbeel. Value iteration net-
482 works. *arXiv preprint arXiv:1602.02867*, 2016.
- 483 [36] Sufeng Niu, Siheng Chen, Hanyu Guo, Colin Targonski, Melissa C. Smith, and Jelena Ko-
484 vačević. Generalized Value Iteration Networks: Life Beyond Lattices. *arXiv:1706.02416 [cs]*,
485 October 2017. URL <http://arxiv.org/abs/1706.02416>. arXiv: 1706.02416.
- 486 [37] Devendra Singh Chaplot, Deepak Pathak, and Jitendra Malik. Differentiable Spatial Planning
487 using Transformers. *arXiv:2112.01010 [cs]*, December 2021. URL <http://arxiv.org/abs/2112.01010>. arXiv: 2112.01010.
- 489 [38] Andreea Deac, Petar Veličković, Ognjen Milinković, Pierre-Luc Bacon, Jian Tang, and Mladen
490 Nikolić. Neural Algorithmic Reasoners are Implicit Planners. October 2021. URL <https://arxiv.org/abs/2110.05442v1>.
- 492 [39] Junhyuk Oh, Satinder Singh, and Honglak Lee. Value Prediction Network. *arXiv:1707.03497*
493 *[cs]*, November 2017. URL <http://arxiv.org/abs/1707.03497>. arXiv: 1707.03497.
- 494 [40] Peter Karkus, David Hsu, and Wee Sun Lee. QMDP-Net: Deep Learning for Planning under
495 Partial Observability. *arXiv:1703.06692 [cs, stat]*, November 2017. URL <http://arxiv.org/abs/1703.06692>. arXiv: 1703.06692.
- 497 [41] Théophane Weber, Sébastien Racanière, David P. Reichert, Lars Buesing, Arthur Guez,
498 Danilo Jimenez Rezende, Adria Puigdomènech Badia, Oriol Vinyals, Nicolas Heess, Yu-
499 jia Li, Razvan Pascanu, Peter Battaglia, Demis Hassabis, David Silver, and Daan Wierstra.
500 Imagination-Augmented Agents for Deep Reinforcement Learning. *arXiv:1707.06203 [cs,*
501 *stat]*, February 2018. URL <http://arxiv.org/abs/1707.06203>. arXiv: 1707.06203.
- 502 [42] Aravind Srinivas, Allan Jabri, Pieter Abbeel, Sergey Levine, and Chelsea Finn. Universal
503 Planning Networks. *arXiv:1804.00645 [cs, stat]*, April 2018. URL <http://arxiv.org/abs/1804.00645>. arXiv: 1804.00645.
- 505 [43] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre,
506 Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, Timothy
507 Lillicrap, and David Silver. Mastering Atari, Go, Chess and Shogi by Planning with a Learned
508 Model. *arXiv:1911.08265 [cs, stat]*, November 2019. URL <http://arxiv.org/abs/1911.08265>. arXiv: 1911.08265.
- 510 [44] Brandon Amos, Ivan Dario Jimenez Rodriguez, Jacob Sacks, Byron Boots, and J. Zico Kolter.
511 Differentiable MPC for End-to-end Planning and Control. *arXiv:1810.13400 [cs, math, stat]*,
512 October 2019. URL <http://arxiv.org/abs/1810.13400>. arXiv: 1810.13400.
- 513 [45] Tingwu Wang and Jimmy Ba. Exploring Model-based Planning with Policy Networks. June
514 2019. URL <https://arxiv.org/abs/1906.08649v1>.
- 515 [46] Arthur Guez, Mehdi Mirza, Karol Gregor, Rishabh Kabra, Sébastien Racanière, Théophane
516 Weber, David Raposo, Adam Santoro, Laurent Orseau, Tom Eccles, Greg Wayne, David Silver,
517 and Timothy Lillicrap. An investigation of model-free planning. *arXiv:1901.03559 [cs, stat]*,
518 May 2019. URL <http://arxiv.org/abs/1901.03559>. arXiv: 1901.03559.

- 519 [47] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to Control:
520 Learning Behaviors by Latent Imagination. *arXiv:1912.01603 [cs]*, March 2020. URL <http://arxiv.org/abs/1912.01603>. arXiv: 1912.01603.
- 522 [48] Vitchyr Pong, Shixiang Gu, Murtaza Dalal, and Sergey Levine. Temporal Difference Models:
523 Model-Free Deep RL for Model-Based Control. *arXiv:1802.09081 [cs]*, February 2018. URL <http://arxiv.org/abs/1802.09081>. arXiv: 1802.09081.
- 525 [49] Ignasi Clavera, Violet Fu, and Pieter Abbeel. Model-Augmented Actor-Critic: Backpropagating
526 through Paths. *arXiv:2005.08068 [cs, stat]*, May 2020. URL <http://arxiv.org/abs/2005.08068>. arXiv: 2005.08068.
- 528 [50] Christopher Grimm, André Barreto, Satinder Singh, and David Silver. The Value Equivalence
529 Principle for Model-Based Reinforcement Learning. *arXiv:2011.03506 [cs]*, November 2020.
530 URL <http://arxiv.org/abs/2011.03506>. arXiv: 2011.03506.
- 531 [51] Christopher Grimm, André Barreto, Gregory Farquhar, David Silver, and Satinder Singh.
532 Proper Value Equivalence. *arXiv:2106.10316 [cs]*, December 2021. URL <http://arxiv.org/abs/2106.10316>. arXiv: 2106.10316.
- 534 [52] Clement Gehring, Kenji Kawaguchi, Jiaoyang Huang, and Leslie Pack Kaelbling. Understanding
535 End-to-End Model-Based Reinforcement Learning Methods as Implicit Parameterization.
536 May 2021. URL <https://openreview.net/forum?id=xj2sE--Q90e>.
- 537 [53] Steven M. LaValle. *Planning Algorithms*. Cambridge University Press, May 2006. ISBN
538 978-1-139-45517-6.
- 539 [54] Maurice Weiler, M. Geiger, M. Welling, Wouter Boomsma, and Taco Cohen. 3D Steerable
540 CNNs: Learning Rotationally Equivariant Features in Volumetric Data. In *NeurIPS*, 2018.
- 541 [55] Johannes Brandstetter, Rob Hesselink, Elise van der Pol, Erik J. Bekkers, and Max
542 Welling. Geometric and Physical Quantities Improve E(3) Equivariant Message Passing.
543 *arXiv:2110.02905 [cs, stat]*, December 2021. URL <http://arxiv.org/abs/2110.02905>.
544 arXiv: 2110.02905.
- 545 [56] T. S. Cohen. Equivariant convolutional networks. 2021. URL <https://dare.uva.nl/search?identifier=0f7014ae-ee94-430e-a5d8-37d03d8d10e6>.
- 547 [57] Maxime Chevalier-Boisvert. Miniworld: Minimalistic 3d environment for rl & robotics re-
548 search. <https://github.com/maximecb/gym-miniworld>, 2018.
- 549 [58] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for
550 biomedical image segmentation. *CoRR*, abs/1505.04597, 2015. URL <http://arxiv.org/abs/1505.04597>.
- 552 [59] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
553 recognition. *CoRR*, abs/1512.03385, 2015. URL <http://arxiv.org/abs/1512.03385>.

554 Checklist

- 555 1. For all authors...
- 556 (a) Do the main claims made in the abstract and introduction accurately reflect the paper's
557 contributions and scope? [Yes]
- 558 (b) Did you describe the limitations of your work? [Yes]
- 559 (c) Did you discuss any potential negative societal impacts of your work? [No]
- 560 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
561 them? [Yes]
- 562 2. If you are including theoretical results...
- 563 (a) Did you state the full set of assumptions of all theoretical results? [Yes] Briefly in
564 Section 4, and in full in the supplementary material.

- 565 (b) Did you include complete proofs of all theoretical results? [Yes] See supplementary
566 material.
- 567 3. If you ran experiments...
- 568 (a) Did you include the code, data, and instructions needed to reproduce the main exper-
569 imental results (either in the supplemental material or as a URL)? [No] The code and
570 data will be cleaned and released prior to final publication.
- 571 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
572 were chosen)? [Yes] Briefly in Section 6, and in full in the supplementary material.
- 573 (c) Did you report error bars (e.g., with respect to the random seed after running exper-
574 iments multiple times)? [Yes] Briefly in Section 6, and in full in the supplementary
575 material.
- 576 (d) Did you include the total amount of compute and the type of resources used (e.g., type
577 of GPUs, internal cluster, or cloud provider)? [Yes] See supplementary material.
- 578 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 579 (a) If your work uses existing assets, did you cite the creators? [Yes]
- 580 (b) Did you mention the license of the assets? [No]
- 581 (c) Did you include any new assets either in the supplemental material or as a URL? [No]
- 582
- 583 (d) Did you discuss whether and how consent was obtained from people whose data
584 you're using/curating? [N/A]
- 585 (e) Did you discuss whether the data you are using/curating contains personally identifi-
586 able information or offensive content? [N/A]
- 587 5. If you used crowdsourcing or conducted research with human subjects...
- 588 (a) Did you include the full text of instructions given to participants and screenshots, if
589 applicable? [N/A]
- 590 (b) Did you describe any potential participant risks, with links to Institutional Review
591 Board (IRB) approvals, if applicable? [N/A]
- 592 (c) Did you include the estimated hourly wage paid to participants and the total amount
593 spent on participant compensation? [N/A]

594 A Outline

595 The appendix is organized as follows. Blue text highlights new content in revision for rebuttal.

- 596 1. A temporary section for new figures and results for rebuttal.
- 597 2. A preliminary version of a section on Symmetric Planning with less prerequisites on equiv-
598 ariant networks.
- 599 3. Additional experimental setup and empirical results. (This section is moved here previously
600 at the end of appendix.)
- 601 4. Discussion on the considered symmetry, as well as limitations and extensions.
- 602 5. Additional technical background and concepts on steerable CNNs and group CNNs, useful
603 for understanding how to apply our setup to other problems and setup.
- 604 6. More details and interpretation on the steerable planning framework.
- 605 7. Full derivation and proofs.
- 606 8. Other practice implementation details.

607 B A temporary section for new figures and results

608 This temporary section is a collection of all new figures and results for the rebuttal purpose. All the
609 content will be merged into the corresponding sections in the future version.

610 B.1 Updated environment figures

611 To emphasize the tasks, we update the figures for the environments in our experiments, along with
612 demonstration of the learned model.

613 We show the figures for **Configuration-space and Workspace manipulation** in Figure 6, and the
614 figures for **2D and Visual Navigation** in Figure 7.




Figure 6: A set of visualization for a 2-joint manipulation task. The obstacles are randomly generated. (1) The 2-joint manipulation task shown in top-down workspace with 96×96 resolution. This is used as the input to the **Workspace Manipulation** task. (2) The predicted configuration space in resolution 18×18 from a mapper module, which is jointly optimized with a planner network. (3) The ground truth configuration space from a handcraft algorithm in resolution 18×18 . This is used as input to the **Configuration-space (C-space) Manipulation** task and as target in the auxiliary loss for the Workspace Manipulation task (as done in SPT [37]). (4) The predicted policy (overlaid with C-space obstacle for visualization) from an end-to-end trained SymVIN model that uses a mapper to take the top-down workspace image and plans on a learned map. The red block is the goal position.

615 B.2 Results on generalization to larger maps

616 To better demonstrate the empirical difference, we conduct new experiment on generalization to
617 larger maps. We hope this can alleviate some concern on (1) scalability and (2) performance gap
618 between SymGPPN and ConvGPPN.

619 We experiment all methods on map size 15×15 through 99×99 , averaging over 3 seeds (3 model
620 checkpoints) for each method and 1000 maps for each size. Note that all models are trained on
621 15×15 with $K = 30$.

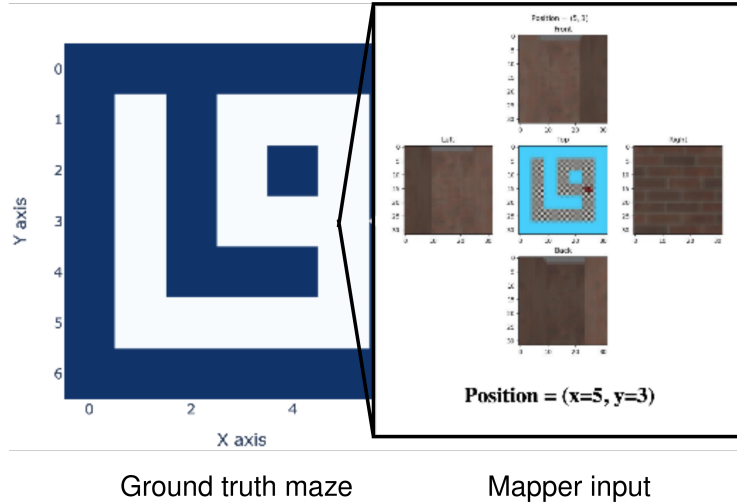


Figure 7: A set of visualization for 2D navigation and visual navigation. The maze is randomly generated. (1, top) The 3D visual navigation environment generated by an illustrative 7×7 map, where we highlight the panoramic view at a position $(5, 3)$ with four RGB images (resolution $32 \times 32 \times 3$). The entire observation tensor for this 7×7 example visual navigation environment is $7 \times 7 \times 4 \times 32 \times 32 \times 3$. This is used as the input to the **Visual Navigation** task. (2) Another predicted map in resolution 15×15 from a mapper module, which is jointly optimized with a planner network. We show the visualization a different map used in actual training. (3) The ground truth map in resolution 15×15 . This is also used as input to the **2D Navigation** task and as target in the auxiliary loss for the Visual Navigation task (as done in GPPN). (4) The predicted policy from an end-to-end trained SymVIN model that uses a mapper to take the observation images (formed as a tensor) and plans on a learned map. The red block is the goal position.

622 Between 15×15 and 49×49 we use all odd-size maps, and between 51×51 and 99×99 we
 623 use interval of 4 ($51 \times 51 \rightarrow 55 \times 55 \dots$). We only use odd size maps because for the even size
 624 maps the maze generation algorithm would cause non-symmetric pattern (missing right and bottom
 625 boundary).

626 Note that we disable backward pass during evaluation. However, we observe that GPPN variants do
 627 use much more memory if backward pass is enabled because (1) they rely on the costly computation
 628 of LSTM, and (2) the number of parameters is also significantly larger. The training and inference
 629 time used by GPPN variants is also significantly longer. We omit the consideration of resource and
 630 time issue and focus on the final generalization results.

631 We focus on comparing SymPlan methods with non-equivariant baselines, by grouping them based
 632 on VIN or GPPN. The results are shown in Figure 8.

633 **Fixed K .** For fixed K setup in Figure 8 (left), we keep number of iterations to be $K = 30$ and
 634 kernel size $F = 3$ for all methods.

635 For SymVIN, it far surpasses VIN for all sizes and preserves the gap throughout the evaluation. Ad-
 636 ditionally, SymVIN has slightly higher variance across three random seeds (three separately trained
 637 models).

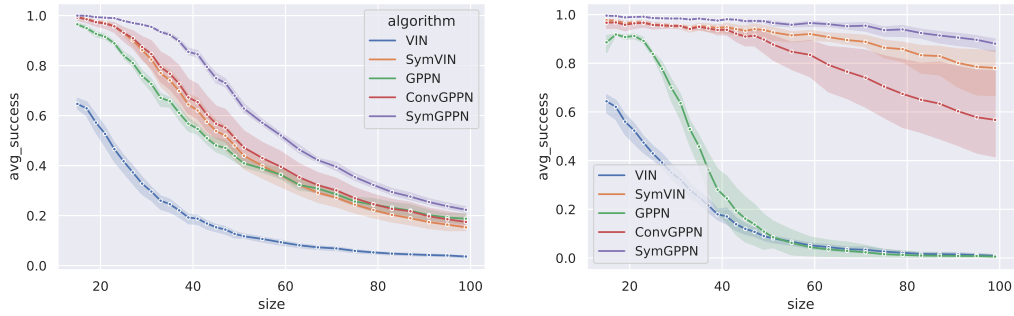


Figure 8: Results for generalization on larger maps for all methods. **(Left)** Fixed $K = 30$ iterations. **(Right)** Variable K iterations, where $K = \sqrt{2} \cdot M$ and M is the generalization map size (x-axis).

638 Among GPPN and its variants, SymGPPN significantly outperforms both GPPN and ConvGPPN.
 639 Interestingly, ConvGPPN has sharper drop with map size than both SymGPPN and ConvGPPN and
 640 thus has increasingly larger gap with SymGPPN and finally even got surpassed by GPPN. Across
 641 random seeds, the three trained models of ConvGPPN give unexpectedly high variance compared to
 642 GPPN and SymGPPN.

643 **Variable K .** We also experiment all methods in the same setting but with variable number of
 644 iterations $K = \sqrt{2} \cdot M$, where M is the generalization map size (x-axis). The trend is very different
 645 from fixed K setup.

646 SymVIN generalizes extremely well compared to VIN, although the variance is greater. GPPN
 647 seems to diverge for larger variable K since it is even worse than fixed $K = 30$ in all map sizes.
 648 ConvGPPN somehow helps convergence, while it fluctuates for different seeds, and SymGPPN
 649 is even better and more stable. Surprisingly, SymVIN is even better than ConvGPPN, although
 650 injecting symmetry (into SymVIN) does not explicitly deal with convergence.

651 C A Guide to Symmetric Planning

652 To address the common concern on the accessibility issue for technical section, as a step to solve
 653 this, we write a section on explaining the SymVIN method with PyTorch-style pseudocode, since it
 654 directly corresponds to what we propose in Section 4 and 5. We try to relate (1) existing concepts
 655 with VIN, (2) what we propose in Section 4 and 5 for SymVIN, and (3) actual PyTorch implementa-
 656 tion of VIN and SymVIN aligned line-by-line based on semantic correspondence.

657 We will consider to have another short section on intuitively explaining our Symmetric Planning
 658 framework and practical considerations in the next few days.

659 C.1 PyTorch-style pseudocode

660 We provide the key Python code snippets to demonstrate how easy it is to implement SymVIN, our
 661 symmetric version of VIN [17].

662 In the current Section 5 (SymPlan practice), we heavily use the concepts from Steerable CNNs.
 663 Thanks to the equivariant network community and the `e2cnn` package, the actual implementation
 664 is compact and closely corresponds to their non-equivariant counterpart, VIN, line-by-line. Thus,
 665 the ultimate goal here is to illustrate that, whatever concepts we have in regular CNNs (e.g., have
 666 whatever channels we want), we can use steerable CNNs that incorporate desired extra symmetry
 667 (of D_4 rotation+reflection or C_4 rotation).

668 We highlight the implementation of the value iteration procedure in VIN and SymVIN:

$$V := \max_a R^a + \gamma \times P^a * V. \quad (5)$$

669 Note that we use actual code snippets to avoid hiding any details.

```

1 import torch
2
3
4
5
6
7
8
9
10
11
12 # Define regular 2D convolution
13 q_conv = torch.nn.Conv2d(
14     in_channels=1,
15     out_channels=2 * q_size,
16     kernel_size=F, stride=1, bias=False
17 )

```

Listing 1: Define ‘expected value’ convolution layer for VIN.

```

1 import torch
2 import e2cnn
3
4 # Define the symmetry group to be D4
5 gspace = e2cnn.gspaces.FlipRot2d0nR2(N=4)
6 # Define feature (fiber) representations
7 field_type_q_in = e2cnn.nn.FieldType(
8     gspace=gspace,
9     representations=2 * q_size * [gspace.
10     regular_repr]
11 )
12 # Define steerable convolution
13 q_r2conv = e2cnn.nn.R2Conv(
14     in_type=field_type_q_in,
15     out_type=field_type_q_out,
16     kernel_size=F, stride=1, bias=False
17 )

```

Listing 2: Define ‘expected value’ (steerable) convolution layer for SymVIN.

670 **Defining (steerable) convolution layer.** First, we show the definition of the key convolution layer
671 for a key operation in VIN and SymVIN: expected value operator, in Listing 1 and 2.

672 As proved in Theorem 4.2, the expected value operator can be executed by a steerable convolution
673 layer for (2D) path planning. This serves as the theoretical foundation on how we should use a
674 steerable layer here.

675 For the left side, a regular 2D convolution is defined for VIN. The right side defines a steerable
676 convolution layer, using the library `e2cnn` from [16]. It provides high-level abstraction for building
677 equivariant 2D steerable convolution networks. As a user, we only need to specify how the feature
678 fields transform (as shown in Figure 3), and it will solve the G -steerability constraints, process what
679 needs to be trained for equivariant layers, etc. We use name `q_r2conv` to highlight the difference.

680 **Value iteration procedure.** Second, we compare the for loop for value iteration updates in VIN
681 and SymVIN, where the former one has regular 2D convolution `Conv2D` (Listing 3), and the latter
682 one uses steerable convolution [16] (Listing 4).

683 The lines are aligned based on semantic correspondence. The `e2cnn` layers, including steerable con-
684 volution layers, operate on its `GeometricTensor` data structure, which is to wrap a PyTorch tensor.
685 We denote them with `_geo` suffix. It only additionally needs to specify how this tensor (feature field)
686 transforms under a group (e.g., D_4), i.e. the user needs to specify a group representation for it.

687 `tensor_directsum` is used to concatenate two `GeometricTensor`’s (feature fields) and compute
688 their associated representations (by direct-sum).

689 Thus, the `e2cnn` steerable convolution layer on the right side `q_r2conv` can be used as a regular
690 PyTorch layer, while the input and output are `GeometricTensor`.

691 We also define the max operation as a customized max-pooling layer, named `q_max_pool`. The
692 implementation is similar to the left side of VIN and needs to additionally guarantee equivariance,
693 and the detail is omitted.

694 Note that for readability, we assume we use regular representations for the Q-value field Q and the
695 state-value field V . They are empirically found to work the best. This corresponds to the definition
696 in `field_type_q_in` in line 9 in the SymVIN definition listing and the comments in line 16-17 in
697 the steerable VI procedure listing for SymVIN.

698 Other components are omitted.

699 D Simplified Version: Symmetric Planning

700 This is a new preliminary version during the rebuttal period that aims to introduce symmetric plan-
701 ning in a more intuitive way, with minimum prerequisites of equivariant networks. This section is
702 intended to be an alternative and more intuitive version to Section 4 (SymPlan framework) and Sec-

```

1 # Input: maze and goal map, #iterations K
2
3
4
5
6 x = torch.cat([maze_map, goal_map], dim=1)
7
8 r = r_conv(x)
9
10 # Init value function V
11 v = torch.zeros(r.size())
12
13
14 for _ in range(K):
15     # Concat and convolve V with P
16     rv = torch.cat([r, v], dim=1)
17     q = q_conv(rv)
18
19     # Max over action channel
20     # > Q: batch_size x q_size x W x H
21     # > V: batch_size x 1 x W x H
22     q = q.view(-1, q_size, W, H)
23     v, _ = torch.max(q, dim=1)
24     v = v.view(-1, W, H)
25
26 # Output: 'q' (to produce policy map)

```

Listing 3: The central value iteration procedure for VIN. Some variable names are adjusted accordingly for readability. W and H are width and height for 2D map.

```

1 # Input: maze and goal map, #iterations K
2
3 from e2cnn.nn import GeometricTensor
4 from e2cnn.nn import tensor_directsum
5
6 x = torch.cat([maze_map, goal_map], dim=1)
7 x_geo = GeometricTensor(x, type=field_type_x)
8 r_geo = r_r2conv(x_geo)
9
10 # Init V and wrap V in e2cnn 'geometric tensor'
11 v_raw = torch.zeros(r_geo.size())
12 v_geo = GeometricTensor(v_raw, field_type_v)
13
14 for _ in range(K):
15     # Concat (direct-sum) and convolve V with P
16     rv_geo = tensor_directsum([r_geo, v_geo])
17     q_geo = q_r2conv(rv_geo)
18
19     # Max over group channel
20     # > Q: batch_size x (|G| * q_size) x W x H
21     # > V: batch_size x (|G| * 1) x W x H
22     v_geo = q_max_pool(q_geo)
23
24
25
26 # Output: 'q_geo' (to produce policy map)

```

Listing 4: The equivariant steerable value iteration procedure for SymVIN. Lines are aligned by semantic correspondence. Definition of other field types are similar and thus omitted.

703 tion 5 (SymPlan in practice). We would appreciate feedback and consider to make further revision
704 to this section and the organization of the entire paper.

705 D.1 Overview

706 In this work, we aim to exploit the inherent symmetry in a broadly existed problem: path planning.
707 Intuitively, since a rotated or reflected 2D map are still another instance of 2D map, such as in
708 Figure 1, their policies and optimal paths are related. This unveils an inherent symmetry property of
709 the path planning problem on the 2D grid that we could exploit.

710 In our work, we provide a rigorous algorithmic framework that can *provably* make use of symmetry
711 in an *efficient* manner. In this section, we will first introduce the algorithm we are based on: Value
712 Iteration Networks (VINs) [17], and use it as foundation to build our algorithm: Symmetric VIN.
713 Finally, we provide intuition to the theoretical guarantees on how we make use of symmetry.

714 D.2 Value Iteration Network: Background and Interpretation

715 Value Iteration Network (VIN) [17] is an example of a differentiable planning algorithm. It empir-
716 ically found that, for 2D path planning, value iteration can be implemented by a deep convolution
717 network.

718 **Background: VIN.** Value iteration is an instance of a dynamic programming (DP) method to
719 solve Markov decision processes (MDPs). It iteratively applies the Bellman (optimality) operator
720 until convergence, which is based on the following Bellman (optimality) equation:

$$Q(s, a) = R(s, a) + \gamma \sum_{s'} P(s'|s, a) V(s'), \quad V(s) = \max_a Q(s, a) \quad (6)$$

721 Tamar et al. [17] used a convolution network to parameterize value iteration. It jointly learns in a
722 latent MDP on 2D grid, which has the latent reward function $\bar{R} : \mathbb{Z}^2 \rightarrow \mathbb{R}^{|\mathcal{A}|}$ and value function
723 $\bar{V} : \mathbb{Z}^2 \rightarrow \mathbb{R}$, and applies value iteration on that MDP:

$$\bar{Q}_{\bar{a}, i', j'}^{(k)} = \bar{R}_{\bar{a}, i, j} + \sum_{i, j} W_{\bar{a}, i, j}^V \bar{V}_{i'-i, j'-j}^{(k-1)} = \bar{R}_{\bar{a}, i, j} + \text{Conv2D}(\bar{V}; W^V), \quad \bar{V}_{i, j}^{(k)} = \max_{\bar{a}} \bar{Q}_{\bar{a}, i', j'}^{(k)} \quad (7)$$

724 Later, we generalize the idea of VIN that (1) represents reward and value functions as fields on
725 2D grid, and (2) realizes value iteration by operations on the fields. Our final goal is to use VIN




Figure 9: (This is a copy of Figure 3.) The commutative diagram of Symmetric Value Iteration Network (SymVIN). Every *row* is a full computation graph of VIN. Every *column* is to rotate field by $\circlearrowleft 90^\circ$. The key message: if we *rotate* the map (from M to $g.M$), to guarantee the final policy function to also be *equivalently rotated* (from A to $g.A$), we shall guarantee every *transformation* (e.g., $Q_k \mapsto V_k$ and $V_k \mapsto Q_{k+1}$) in value iteration to also be *equivariant* ($g.f(x) = f(g.x)$, for every *pair of columns*).

726 to demonstrate a principled method for incorporating symmetry in differentiable planning. After
 727 reviewing the basics of VIN, we will next summarize the reasoning of choosing VIN.

728 D.3 Symmetric Value Iteration Network: A Practical Symmetric Planning Algorithm

729 **Why do we choose VIN?** There are two reasons behind the choice of VIN.

- 730 1. The expected value operator in value iteration $\sum_{s'} P(s'|s, a)V(s')$ is *linear* in value func-
 731 tion. As we show in Theorem 4.1, it is also *equivariant* for (2D) path planning, this means
 732 that it is a *linear equivariant operator*. According to Cohen et al. [12], any linear equiv-
 733 ariant operator (on spaces such as 2D grid) has one-to-one correspondence to a (group
 734 equivariant) convolution operator.
- 735 2. Value iteration, or Bellman (optimality) operator, is fully convolutional, i.e. only relies
 736 on operating on functions (“fields”) over \mathbb{Z}^2 , such as value function, reward function, and
 737 transition functions: $V_{k+1}(s) = \max_a R^a(s) + \gamma \times [P^a \star V_k](s)$. This enables us to inject
 738 symmetry by enforcing equivariance within convolution. For example, for a 2D map in
 739 Figure 1, the 4 corner states are symmetric under any one of the eight transformations in
 740 D_4 , and we can enforce those 4 states to have the same value if we rotate or flip the map
 741 (D_4 -equivariance). This avoids the need to find if a new state is symmetric to any existing
 742 state, which is shown to be NP-hard [8].

743 In summary, VIN satisfies both desiderata: (1) it uses convolution as the backbone, and (2) it operates
 744 on fields. Furthermore, we find VIN is empirically and conceptually the *simplest* differentiable
 745 planning algorithm that satisfies them, which leads to our decision.

746 **How to inject symmetry?** VIN uses a regular 2D convolutional network (Equation 7), which has
 747 *translation equivariance* [15, 13]. More concretely, a VIN will output the same value function for
 748 the same map patches that only differ by 2D translation. We omit how to characterize translation
 749 equivariance here, since it requires a different mechanism to handle and does not *decrease* the search
 750 space nor *reduce* a path planning MDP to an easier problem.

751 Beyond translation, we are more interested in *rotation* and *reflection* symmetries. Intuitively, as in
 752 Figure 1, if we find the optimal solution to a map, it automatically **generalizes** the solution to all 8
 753 transformed maps (4 rotations times 2 reflections, including identity transformation). This can be
 754 characterized by *equivariance* of a planning algorithm Plan , such as value iteration VI, visualized
 755 in Figure 9: $g.\text{Plan}(M) = \text{Plan}(g.M)$, where M is a maze map, and g is the symmetry group D_4
 756 under which 2D grid is invariant.

757 More importantly, symmetry also helps **training** of differentiable planning algorithms. Intuitively,
 758 symmetry in path planning poses additional constraint to its search space: if the goal is in the north,

759 go up; if in the east, go right. In other words, the knowledge can be shared between symmetric cases,
 760 or the path planning is effectively reduced by symmetry to a smaller one. This property can also be
 761 depicted by equivariance of Bellman operators \mathcal{T} , or a step of value iteration: $g.\mathcal{T}[V_0] = \mathcal{T}[g.V_0]$.
 762 If we use $\text{VI}(M)$ to denote applying Bellman operators on arbitrary initialization until convergence
 763 $\mathcal{T}^\infty[V_0]$, value iteration is also equivariant, as demonstrated in Figure 9:

$$g.\text{VI}(M) \equiv g.\mathcal{T}^\infty[V_0] = \mathcal{T}^\infty[g.V_0] \equiv \text{VI}(g.M). \quad (8)$$

764 Thus, we inject equivariance into value iteration w.r.t. *rotation* and *reflection*, in addition to *trans-*
 765 *lation*, through **steerable convolution (network)** from Cohen and Welling [14], which exactly
 766 matches our criteria. Cohen et al. [12] prove that steerable convolution is the most general linear
 767 equivariant map under some conditions, which value iteration satisfies. Weiler and Cesa [16] build
 768 $E(2)$ -Steerable CNNs for 2D space, and we use their package `e2cnn` in our implementation. In
 769 practice, to inject symmetry into VIN, we simply need to replace the translation-equivariant `Conv2D`
 770 with `SteerableConv`:

$$\bar{Q}_{\bar{a},i',j'}^{(k)} = \bar{R}_{\bar{a},i,j} + \text{SteerableConv}(\bar{V}; W^V), \quad \bar{V}_{i,j}^{(k)} = \max_{\bar{a}} \bar{Q}_{\bar{a},i',j'}^{(k)}. \quad (9)$$

771 We formally justify our design in Section D.4 below and provide more technical details in Section 4.

772 D.4 Theoretical Justification: Why does it work?

773 In the Section D.3, we show how to exploit symmetry in path planning by equivariance from con-
 774 volution via intuition. The goal of this new section is to (1) connect the theoretical justification with
 775 the algorithmic design, and (2) provide intuition for the justification. Even though we focus on
 776 a specific task, we hope that the underlying guidelines on integrating symmetry into planning are
 777 useful for broader planning algorithms and problems as well.

778 This version is for the purposes of rebuttal and preview, so we may refer some details to the original
 779 Section 4. We will consider further revision depending on how to form the method section.

780 **Overview.** There are numerous types of symmetry in various planning tasks. We study symmetry
 781 in **path planning** as an example, because it is a straightforward planning problem, and its solutions
 782 have been intensively studied in robotics and artificial intelligence [53, 1]. However, even for this
 783 problem, the symmetry has *not* been *effectively* exploited in its planning algorithms, such as Dijk-
 784 stra’s algorithm, A*, or RRT, because of NP-hard orbit finding [8]. Additionally, we focus on **value**
 785 **iteration** because it is both widely use and connects closely with convolution [14].

786 **Theory: symmetry in planning.** If we want to exploit symmetry in a task to improve planning,
 787 there are two major steps: (1) characterize the symmetry in the task, and (2) incorporate correspond-
 788 ing symmetry into the planning algorithm. The theoretical results in Section 4.2 mainly characterize
 789 the symmetry and direct us to a feasible planning algorithm.

790 The **symmetry in tasks** or MDPs can be specified by the equivariance property of the transition and
 791 reward function, studied in Ravindran and Barto [3], van der Pol et al. [32]:

$$\bar{P}(s' | s, a) = \bar{P}(g.s' | g.s, g.a), \quad \forall g \in G, \forall s, a, s' \quad (10)$$

$$\bar{R}_M(s, a) = \bar{R}_{g.M}(g.s, g.a), \quad \forall g \in G, \forall s, a \quad (11)$$

792 Note that how the group G acts on states and actions is decided by the space \mathcal{S} or \mathcal{A} , which has been
 793 discussed in Equation 1 in Section 4.2. We emphasize that the equivariance property of the reward
 794 function is different from prior work [3, 32]: in our case, the reward function encodes obstacles as
 795 well, and thus depends on map input M . Intuitively, using Figure 1 as an example, if a position s is
 796 rotated $g.s$, to find how the correct original reward R , the input map M must also be rotated $g.M$.
 797 More details in Section 4.2 and Section H.

798 As for exploiting the **symmetry in planning algorithms**, we focus on value iteration and the VIN
 799 algorithm. We first prove in Theorem 4.1 that value iteration for path planning respects the *equiv-*
 800 *ariance* property. This confirms that value iteration is a feasible method to incorporate symmetry.
 801 The next result in Theorem 4.2 further proves that value iteration is a general form of convolution
 802 (*steerable convolution*), motivating the use of steerable CNNs by Cohen and Welling [14] to replace
 803 regular CNNs in VIN.

805 **Retrospect.** We study how to inject symmetry into VIN for (2D) path planning, and expect the
806 task-specific technical details are useful for two types of readers. (i) *Using VIN.* If one uses VIN for
807 differentiable planning, the resulting algorithms SymVIN or SymGPPN can be a plug-in alternative,
808 as a part in a larger end-to-end system. (ii) *Studying path planning.* The proposed framework
809 characterizes the symmetry in (2D) path planning, so it is possible to apply the underlying ideas to
810 other domains. For example, it is possible to extend to higher-dimensional continuous Euclidean
811 spaces.

812 This concludes the section. We appreciate any feedback for this new simplified section on imple-
813 mentation and theory of symmetric planning. We will keep improving it and better integrate with
814 the current "detailed" version in the future iterations.

815 E Experiments: Details and Additional Results (moved)

816 E.1 Details: Setup

817 **Action space.** Note that the MDP action space \mathcal{A} needs to be *compatible* with the group action
818 $G \times \mathcal{A} \rightarrow \mathcal{A}$. Since the E2CNN package [16] uses *counterclockwise* rotations as generators for
819 rotation groups C_n , the action space needs to be *counterclockwise*.

820 **Mapper training: manipulation.** During training, we pre-train the mapper and the planner sep-
821 arately for 15 epochs. Where the mapper takes manipulator workspace and outputs configuration
822 space. The mapper is trained to minimize the binary cross entropy between output and ground truth
823 configurations space. The planner is trained in the same way as described in Section 6.1. After
824 pre-training, we switch the input to the planner from ground truth configuration space to the one
825 from the mapper. During testing, we follow the pipeline in [37] that the mapper-planner only have
826 access to the manipulator workspace.

827 E.2 Details: Environments.

828 **Manipulation.** For planning in configuration space, the configuration space of the 2 DoFs ma-
829 nipulator has no constraints in the $\{0, \pi\}$ boundaries, i.e., no joint limits. To reflect this nature
830 of the configuration space in manipulation tasks, we use circular padding before convolution op-
831 eration. The circular padding is applied to convolution layers in VIN, SymVIN, ConvGPPN, and
832 SymGPPN. Moreover, in GPPN, there is a convolution encoder before the LSTM layer. We add the
833 circular padding in the convolution layers in GPPN as well.

834 In **2-DOF manipulation** in configuration space, we adopt the setting in [37] and train networks
835 to take as input of configuration space, represented by two joints. We randomly generate 0 to 5
836 obstacles in the manipulator workspace. Then the 2 degree-of-freedom (DOF) configuration space
837 is constructed from workspace and discretized into 2D grid with sizes $\{18, 36\}$, corresponding to
838 bins of 20° and 10° , respectively.

839 We allow each joint to rotate over 2π , so the configuration space of 2-DOF manipulation forms a
840 torus \mathbb{T}^2 . Thus, the both boundaries need to be connected when generating action demonstrations,
841 and (equivariant) convolutions need to be circular (with padding mode) to wrap around for all meth-
842 ods. We allow each joint to rotate over 2π , so the both boundaries in configuration space need
843 to be connected when generating action demonstrations, and (equivariant) convolutions need to be
844 circular (with padding mode) to wrap around for all methods.

845 E.3 Details: Model Architecture

846 We try to mimic the setup in VIN and GPPN [18].

847 For non-SymPlan related parameters, we use learning rate of 10^{-3} , batch size of 32 if possible
848 (GPPN variants need smaller), RMSprop optimizer.

849 For SymPlan parameters, we use 150 hidden channels (or 150 *trivial* representations for SymPlan
850 methods) to process the input map. We use 100 hidden channels for Q-value for VIN (or 100 *regular*
851 representations for SymVIN), and use 40 hidden channels for Q-value for GPPN and ConvGPPN

852 (or 40 *regular* representations for SymGPPN on 15×15 , and 20 for larger maps because of memory
853 constraint).

854 E.4 Visualization of learned models

855 We visualize a trained VIN and a SymVIN, evaluated on a 15×15 map and its rotated version. For
856 non-symmetric VIN in Figure 10, the learned policy is obviously not equivariant under rotation.

857 We also visualize SymVIN on larger map sizes: 28×28 and 50×50 , to demonstrate its performance
858 and equivariance.




Figure 10: A trained VIN evaluated on a 15×15 map and its rotated version. It is obvious that the learned policy is not equivariant under rotation.




Figure 11: A trained SymVIN evaluated on a 15×15 map and its rotated version.

859 E.5 Further Analysis

860 **Additional training curves.** We also provide other training curves that we only show test numbers
861 in the main text.

862 **Training efficiency with less data.** Since the supervision is still dense, we experiment on training
863 with even smaller dataset to experiment in more extreme setup. We experiment how symmetry may
864 affect the training efficiency of Symmetric Planners by further reducing the size of training dataset.
865 We compare on two environments: 2D navigation and visual navigation, with training/validation/test
866 size of 1K/200/200, for all methods.




Figure 12: A fully trained SymVIN evaluated on a 28×28 map and its rotated version.




Figure 13: A fully trained SymVIN evaluated on a 50×50 map and its rotated version.

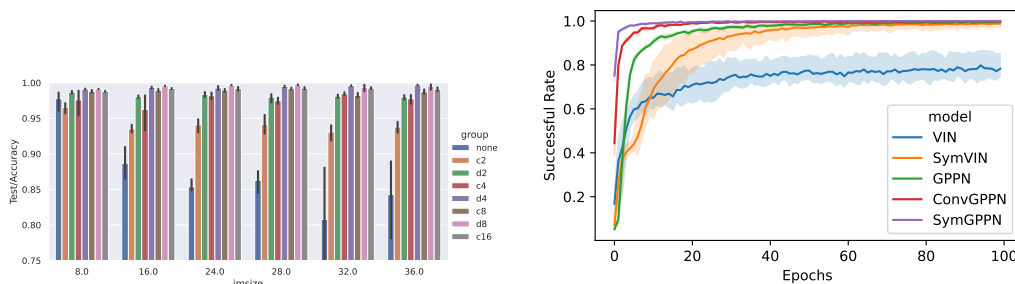


Figure 14: **(Left)** Accuracy evaluated on unseen test maps. The x-axis is the width of the map, and the y-axis is the accuracy, reported on every map size and every size and every chose symmetry group G . **(Right)** Visual navigation 15×15 with 10K data.

867 **Choose of symmetry groups for navigation.** One important benefit of partially equivariant network is that, we do not need to design the group representation of MDP action space $\rho_A(g)$ for different group or action space. Thus, we experiment several G -equivariant variants with different group equivariance: (discrete rotation group) C_2, C_4, C_8, C_{16} , and (dihedral group) D_2, D_4, D_8 , all based on $E(2)$ -steerable CNN [16]. For all intermediate layers, we use regular representations $\rho_{\text{reg}}(g)$ of each group, followed by a final policy layer with non-equivariant 1×1 convolution.

873 The results are reported in the Figure 14 (left). We only compare VIN (denoted as "none" symmetry) against our $E(2)$ -VIN (other symmetry group option) on 2D navigation with 15×15 maps.

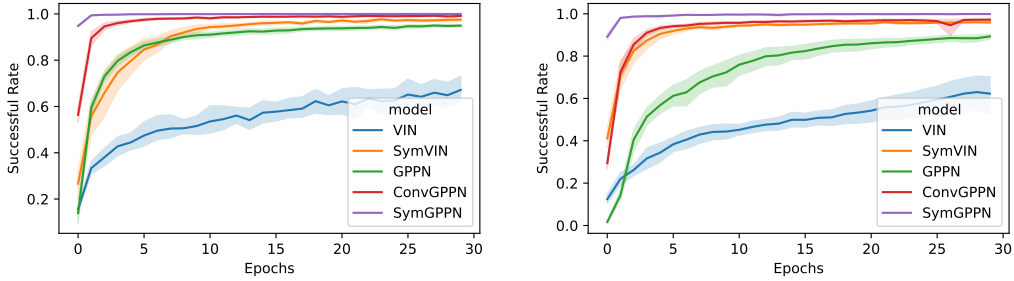


Figure 15: Training curves for (Left) 28×28 and (Right) 50×50 .

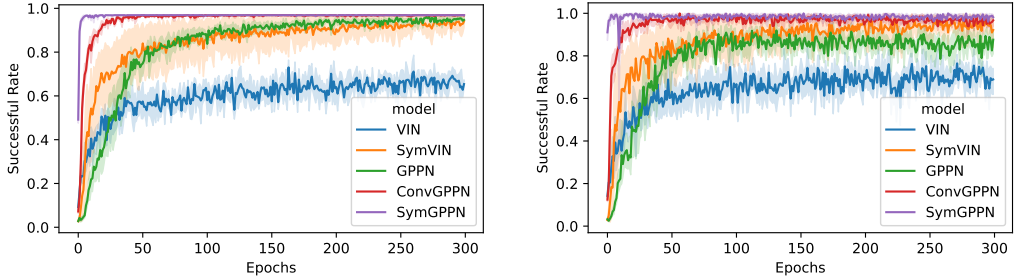


Figure 16: Training curves for 15×15 2D navigation 1K data (Left) training and (Right) validation successful rate.

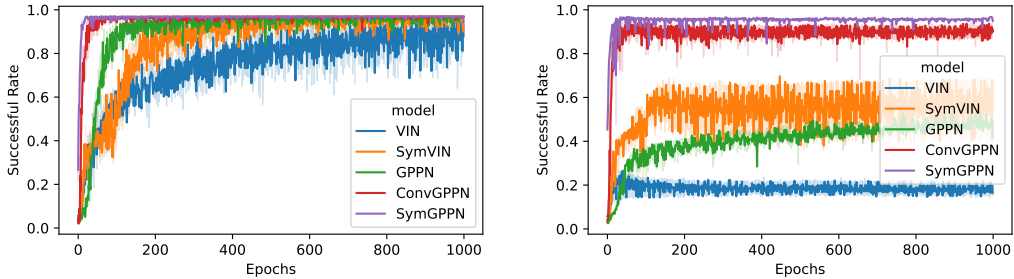


Figure 17: Training curves for 15×15 visual navigation 1K data (Left) training and (Right) validation successful rate.

Table 2: Fiber representations

(Fiber representation)	SymVIN
Default	98.45
Hidden: trivial to regular	99.07
State-value ρ_V : regular to trivial	63.08
Q-value ρ_Q : regular to trivial	21.30
ρ_Q and ρ_V : both trivial	2.814

875 In general, the planners equipped with any G group equivariance outperform the vanilla non-
 876 equivariant VIN, and D_4 -equivariant steerable CNN performs the best on most map sizes. Additionally,
 877 since the environment has actions in 8 directions (4 diagonals), C_8 or D_8 groups seem to
 878 take advantage of that and have slightly higher accuracy on some map sizes, while C_{16} is over-
 879 constrained compared to the true symmetry $G = D_4$ and be detrimental to performance. The
 880 non-equivariant VIN also experiences higher variance on large maps.

881 **Choosing fiber representations.** As we use steerable convolutions [16] to build symmetric plan-
 882 ners, we are free to choose the representations for feature fields, where intermediate equivariant con-
 883 volutional layers will be equivariant between them $f(\rho_{in}(g)x) = \rho_{out}(g)f(x)$. We found represen-
 884 tations for some feature fields are critical to the performance: mainly $V : \mathcal{S} \rightarrow \mathbb{R}$ and $Q : \mathcal{S} \rightarrow \mathbb{R}^{|A|}$.

885 We use the best setting as default, and ablate every option. As shown in Table 2, changing ρ_V or ρ_Q
 886 to trivial representation would result in much worse results.

887 **Fully vs. Partially equivariance for symmetric planners.** One seemingly minor but critical
 888 design choice in our SymPlan networks is the choice of the final policy layer, which maps Q-values
 889 $\mathcal{S} \rightarrow \mathbb{R}^{|\mathcal{A}|}$ to policy logits $\mathcal{S} \rightarrow \mathbb{R}^{|\mathcal{A}|}$. Fully equivariant is expected to perform better, but it has some
 890 points worth to mention. (1) We experience unstable training at the beginning, where the loss can
 891 go up to 10^6 in the first epoch, while we did not observe it in non-equivariant or partially equivariant
 892 counterparts. However, this only slightly affects training.

893 In summary, we found even though fully equivariant version can perform slightly better in the best
 894 tuned setting, on average setting, partially equivariant version is more robust and the gap is much
 895 larger, as shown in the follow table, which an example of averaging over three choices of represen-
 896 tations introduced in the last paragraph. On average partially equivariant version is much better. In
 897 our experiments, partially equivariant version also is easier to tune.

Table 3: Fully vs. Partially equivariance

	(Equivariance) SymVIN
<i>Partially</i> equivariant averaged over all representations	91.04
<i>Fully</i> equivariant averaged over all representations	42.61

898 F Additional Discussion

899 F.1 Limitations and Extensions

900 **Assumption on known domain structure.** As in VIN, although the framework of steerable plan-
 901 ning can potentially handle different domains, one important hidden assumption is that the under-
 902 lying domain Ω (state space), is known. In other words, we fix the structure of learned transition
 903 kernels $p(s' | s, a)$ and estimate coefficients of it. One potential method is to use Transformers that
 904 learn attention weights to all states in \mathcal{S} , which has been partially explored in SPT [37]. Additionally,
 905 it is also possible to treat unknown MDPs as learned transition graphs, as explored in XLVIN [38].
 906 We leave the consideration of symmetry in unknown underlying domains for future work.

907 **The curse of dimensionality.** The paradigm of steerable planning still requires full expansion
 908 in computing value iteration (opposite to *sampling-based*), since we realize the symmetric planner
 909 using group equivariant convolutions (essentially summation or integral). Convolutions on high-
 910 dimensional space could suffer from the curse of dimensionality for higher dimensional domains,
 911 and are vastly under-explored. This is a primary reason why we need sampling-based planning
 912 algorithms. If the domain (state-action transition graph) is sparsely connected, value iteration can
 913 still scale up to higher dimensions. It is also unclear either when steerable planning would fail, or
 914 how sampling-based algorithms could be integrated with the symmetric planning paradigm.

915 F.2 The considered symmetry in spatial MDPs

916 We need to differentiate between two types of symmetry in MDPs. Let’s take spatial graph as
 917 illustrative example to understand the potential symmetry from a higher level, which means that the
 918 nodes \mathcal{V} in the graph have spatial coordinates \mathbb{Z}^n or \mathbb{R}^n . Our 2D path planning is a special case of
 919 spatial graph, where the actions can only move to adjacent spatial nodes.

920 Let the graph denoted as $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$. \mathcal{E} is the set of edges connecting two states with an action.
 921 One type of symmetry is the symmetry of the graph itself. For the grid case, it means that after D_4
 922 rotation or reflection, the map is unchanged.

923 Another type of symmetry comes from the isometries of the space. For a spatial graph, we can rotate
 924 it freely in a space, while the relative positions are unchanged. For our grid case, it is shown in the
 925 Figure 1 that rotating a map resulting in the rotated policy. However, the map or policy itself can
 926 never be equal under any transformation in D_4 .

927 In other words, the first type is symmetry within a MDP (rely on the property of the MDP itself
 928 \mathcal{M} , or $\text{Aut}(\mathcal{M})$), and the second type is symmetry between MDPs (only rely on the property of the
 929 underlying spatial space \mathbb{Z}^2 , or $\text{Aut}(\mathbb{Z}^2)$).

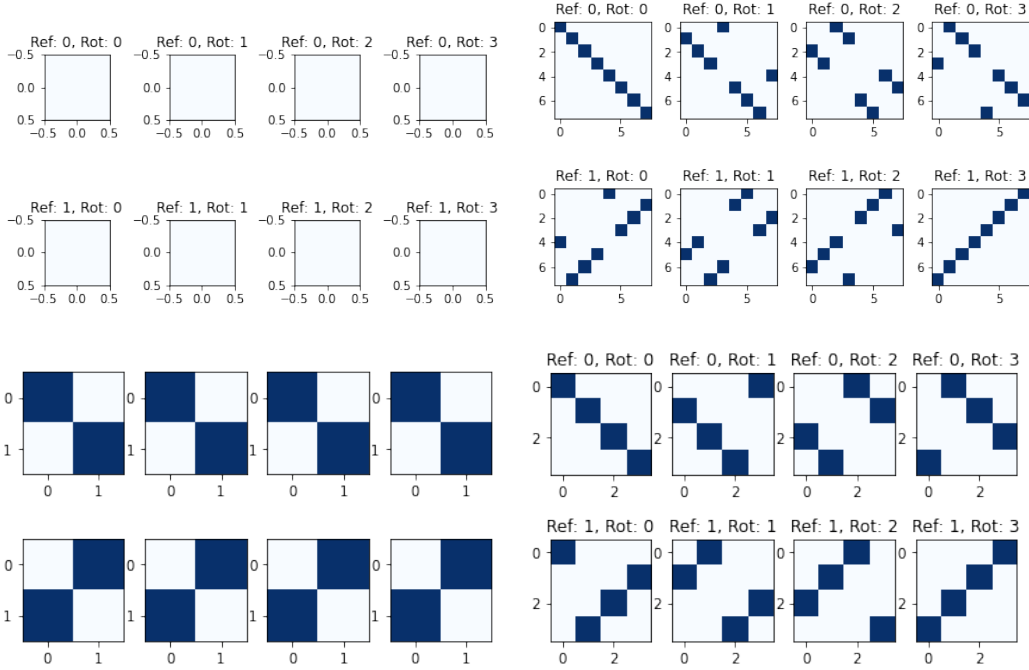


Figure 18: Visualization of the permutation representations of D_4 group for every element $g \in D_4$ (4 rotations each row and 2 reflections each column). They are (1) the trivial representation, (2) the regular representation, (3) the quotient representation (quotienting out *rotations*), (4) the quotient representation (quotienting out *reflections*).

930 Nevertheless, we could input map M and somehow treat symmetric states between MDPs as one
 931 state. See the proofs section for more details.

932 G Additional Background and Concepts

933 G.1 Group representations: visual understanding

934 A group representation is a (linear) group action that defines how a group acts on some space. Cohen
 935 and Welling [15, 14], Weiler and Cesa [16] provide more formal introduction to them in the context
 936 of equivariant neural networks. We provide visual understanding and refer the readers to them for
 937 comprehensive account.

938 To visually understand how the group D_4 acts on some vector space, we visualize the trivial, regular,
 939 and quotient (quotienting out reflections sr^2) representations, which are *permutation matrices*. If
 940 we apply such a representation $\rho(g)(g \in D_4)$ to a vector, the elements get *cyclically permuted*. See
 941 Figure 18.

942 The quotient representation that quotients out reflections and has dimension 4×4 is what we need
 943 to use on the 4-direction action space.

944 G.2 Geometric Deep Learning

945 We review another set of important concepts that motivate our formulation of steerable planning:
 946 geometric deep learning and the theories on connecting equivariance and convolution [11, 12, 13].
 947 Bronstein et al. [11] use x for feature fields while Cohen and Welling [14], Cohen et al. [12], Weiler
 948 and Cesa [16] use f .

949 **Convolutional feature fields.** The signals are taken from set $\mathcal{C} = \mathbb{R}^D$ on some structured *domain*
 950 Ω , and all mappings from the domain to signals forms the space of \mathcal{C} -valued signals $\mathcal{X}(\Omega, \mathcal{C}) = \{f : \Omega \rightarrow \mathcal{C}\}$,
 951 or $\mathcal{X}(\Omega)$ for abbreviation. For instance, for RGB images, the domain is the 2D $n \times n$

952 grid $\Omega = \mathbb{Z}_n \times \mathbb{Z}_n$, and every pixel can take RGB values $\mathcal{C} = \mathbb{R}^3$ at each point in the domain
 953 $u \in \Omega$, represented by a mapping $x : \mathbb{Z}_n \times \mathbb{Z}_n \rightarrow \mathbb{R}^3$. A function on images thus operates on
 954 $3n^2$ -dimensional inputs.

955 It is argued that the underlying geometric structure of domains Ω plays key role in alleviating the
 956 curse of dimensionality, such as convolution networks in computer vision, and this framework is
 957 named *Geometric Deep Learning*. We refer the readers to Geometric Deep Learning [11] for more
 958 details, and to more rigorous theories on the relation between equivariant maps and convolutions
 959 in [12] (vector fields through induced representations) and [13] (scalar fields through trivial repre-
 960 sentations).

961 **Group convolution.** Convolutions are shift-equivariant operations, and vice versa. This is the
 962 special case for $\Omega = \mathbb{R}$, which can be generalized to any group G (that we can integrate or sum
 963 over). The *group convolution* for signals on Ω is then defined⁵ as

$$(f \star \psi)(g) = \langle f, \rho(g)\psi \rangle = \int_{\Omega} f(u)\psi(g^{-1}u)du, \quad (12)$$

964 where $\psi(u)$ is shifted copies of a filter, usually locally supported on a subset of Ω and padded
 965 outside. Note that although x takes $u \in \Omega$, the feature map $(x \star \psi)$ takes as input the elements
 966 $g \in G$ instead of points on the domain $u \in \Omega$. All following group convolution layers take G :
 967 $\mathcal{X}(G) \rightarrow \mathcal{X}(G)$. In the grid case, the domain Ω is *homogeneous* space of the group G , i.e. the group
 968 G acts transitively: for any two points $u, v \in \Omega$ there exists a symmetry $g \in G$ to reach $u = gv$.

969 Analogous to classic shift-equivariant convolutions, the generalized group convolution is G -
 970 equivariant [12]. It is observed that $\langle x, \rho(g)\theta \rangle = \langle \rho(g^{-1})x, \theta \rangle$, and from the defining property
 971 of group representations $\rho(h^{-1})\rho(g) = \rho(h^{-1}g)$, the G -equivariance of group convolution fol-
 972 lows [11]:

$$(\rho(h)x \star \theta)(g) = \langle \rho(h)x, \rho(g)\theta \rangle = \langle x, \rho(h^{-1}g)\theta \rangle = \rho(h)(x \star \theta)(g) \quad (13)$$

973 **Steerable convolution kernels.** Steerable convolutions extend group convolutions to more general
 974 setup and decouple the computation cost with the group size [14, 56]. For example, $E(2)$ -steerable
 975 CNNs [16] apply it for $E(2)$ group, which is semi-direct product of translations \mathbb{R}^2 and a fiber
 976 group H , where H is a group of transformations that fixes the origin and is $O(2)$ or its subgroups.
 977 The representation on the signals/fields is induced from a representation of the fiber group H . Use
 978 \mathbb{R}^2 as example, a steerable kernel only needs to be H -equivariant by satisfying the following con-
 979 straint [16]:

$$\psi(hx) = \rho_{\text{out}}(h)\psi(x)\rho_{\text{in}}(h^{-1}) \quad \forall h \in H, x \in \mathbb{R}^2. \quad (14)$$

980 G.3 Steerable CNNs

981 We still use the running example on \mathbb{Z}^2 and group $p4m = \mathbb{Z}^2 \rtimes D_4$.

982 **Induced representations.** We follow [14, 12] to use π for *induced* representations. We still use
 983 feature fields over \mathbb{Z}^2 as example.

984 As shown in **Figure 2 middle**, to transform a feature field $f : \mathbb{Z}^2 \rightarrow \mathbb{R}^C$ on base \mathbb{Z}^2 with group
 985 $p4m = \mathbb{Z}^2 \rtimes D_4$, we need the *induced representation* [14, 12]. The induced representation in this
 986 case is denoted as $\pi(g) \triangleq \text{ind}_{D_4}^{\mathbb{Z}^2 \rtimes D_4} \rho(g)$ (for all g), which means how the group action of D_4
 987 transforms a feature field on $\mathbb{Z}^2 \rtimes D_4$.

988 It acts on the feature field with two parts: (1) on the base space \mathbb{Z}^2 and (2) on the fibers (feature
 989 channels \mathbb{R}^C) by fiber group $H = D_4$ [14, 16]. More specifically, applying a translation $t \in \mathbb{Z}^2$ and
 990 a transformation $r \in D_4$ to some field f , we get $\pi(tr)f$ [14, 16]:

$$f(x) \mapsto [\pi(tr)f](x) \triangleq \rho(r) \cdot [f((tr)^{-1}x)]. \quad (15)$$

⁵The definition of group convolution needs to assume that (1) signals $\mathcal{X}(\Omega)$ are in a Hilbert space (to define an inner product $\langle x, \theta \rangle = \int_{\Omega} x(u)\theta(u)du$) and (2) the group G is locally compact (so a Haar measure exists and "shift" of filter can be defined).

991 $\rho(r)$ is the fiber representation that transforms the fibers \mathbb{R}^C , and $(tr)^{-1}x$ finds the element before
 992 group action (or equivalently transforming the base space \mathbb{Z}^2). Thus, π only depends on the fiber
 993 representation ρ but not the latter part, thus named *induced representation* by ρ .

994 **Steerable convolution vs. group convolution.** The steerable convolution on \mathbb{Z}^2 The understand-
 995 ing of this point helps to understand how a group acts on various feature fields and the design of
 996 state space for path planning problems. We use the discrete group $p4 = \mathbb{Z}^2 \rtimes C_4$ as example, which
 997 consists of \mathbb{Z}^2 translations and 90° rotations. The only difference with $p4m$ is $p4$ does not have
 998 reflections.

999 The group convolution with filter ψ and signal x on grid (or $\mathbf{p} \in \mathbb{Z}^2$), which outputs signals (a
 1000 function) on group $p4$

$$[\psi \star x](\mathbf{t}, r) := \sum_{\mathbf{p} \in \mathbb{Z}^2} \psi((\mathbf{t}, r)^{-1}\mathbf{p}) x(\mathbf{p}). \quad (16)$$

1001 A group G has a natural action on the functions over its elements; if $x : G \rightarrow \mathbb{R}$ and $g \in G$, the
 1002 function $g.x$ is defined as $[g.x](h) := x(g^{-1} \cdot h)$.

1003 For example: The group action of a rotation $r \in C_4$ on the space of functions over $p4$ is

$$[r.y](\mathbf{p}, s) := y(r^{-1}(\mathbf{p}, s)) = y(r^{-1}\mathbf{p}, r^{-1}s), \quad (17)$$

1004 where $r^{-1}\mathbf{p}$ spatially rotates the pixels, $r^{-1}s$ cyclically permutes the 4 channels.

1005 The G-space (functions over $p4$) with a natural action of $p4$ on it:

$$[(\mathbf{t}, r).y](\mathbf{p}, s) := y((\mathbf{t}, r)^{-1} \cdot (\mathbf{p}, s)) = y(r^{-1}(\mathbf{p} - \mathbf{t}), r^{-1}s) \quad (18)$$

1006 The group convolution in discrete case is defined as

$$[\psi \star x](g) := \sum_{h \in H} \psi(g^{-1} \cdot h) x(h). \quad (19)$$

1007 The group convolution with filter ψ and signal x on $p4$ group is given by:

$$[\psi \star x](\mathbf{t}, r) := \sum_{s \in C_4} \sum_{\mathbf{p} \in \mathbb{Z}^2} \psi((\mathbf{t}, r)^{-1}(\mathbf{p}, s)) x(\mathbf{p}, s). \quad (20)$$

1008 Using the fact

$$\psi((\mathbf{t}, r)^{-1}(\mathbf{p}, s)) = \psi(r^{-1}(\mathbf{p} - \mathbf{t}, s)) = [r.\psi](\mathbf{p} - \mathbf{t}, s), \quad (21)$$

1009 the convolution can be equivalently written into

$$[\psi \star x](\mathbf{t}, r) := \sum_{s \in C_4} \left(\sum_{\mathbf{p} \in \mathbb{Z}^2} [r.\psi](\mathbf{p} - \mathbf{t}, s) x(\mathbf{p}, s) \right). \quad (22)$$

1010 So $\left(\sum_{\mathbf{p} \in \mathbb{Z}^2} [r.\psi](\mathbf{p} - \mathbf{t}, s) x(\mathbf{p}, s) \right)$ can be implemented in usual shift-equivariant convolution
 1011 CONV2D.

1012 The inner sum $\sum_{\mathbf{p} \in \mathbb{Z}^2}$ is equivalently for the sum in steerable convolution, and the outer sum $\sum_{s \in C_4}$
 1013 implement rotation-equivariant convolution that satisfies H -steerability kernel constraint. Here, the
 1014 outer sum is essentially using the *regular* fiber representation of C_4 .

1015 In other words, group convolution on $p4 = \mathbb{Z}^2 \rtimes C_4$ group is equivalent to steerable convolution on
 1016 base space \mathbb{Z}^2 with the fiber group of C_4 with regular representation.

1017 **Stack of feature fields.** Analogous to ordinary CNNs, a feature space in steerable CNNs can consist
 1018 of multiple feature fields $f_i : \mathbb{Z}^2 \rightarrow \mathbb{R}^{c_i}$. The feature fields are stacked $f = \bigoplus_i f_i$ together
 1019 by concatenating the individual feature fields f_i (along the fiber channel), which transforms under
 1020 the directly sum $\rho = \bigoplus_i \rho_i$ of individual (fiber) representations. Every layer will be equivariant
 1021 between input and output field $f_{\text{in}}, f_{\text{out}}$ under induced representations $\pi_{\text{in}}, \pi_{\text{out}}$. For a steerable con-
 1022 volution between more than one-dimensional feature fields, the kernel is matrix-valued [12, 16].

⁵Technically, we still need to solve the linear equivariance constraint in Eq. 35 to enable weight-sharing for equivariance, while Weiler and Cesa [16] have implemented it for 2D case.

1023 H Symmetric Planning Framework: Additional Details

1024 H.1 Path planning in neural networks

1025 We provide the detailed construction of doing path planning in neural networks in the Section 4.
1026 This further explains the visualization in Figure 2 left.

1027 We use the running example of planning on the 2D grid \mathbb{Z}^2 . We aim to understand (1) how VIN-
1028 style networks embed planning and how its idea generalizes, (2) how is symmetry structure defined
1029 in path planning and how could it be injected into such planning networks. Recall that we aim
1030 to understand (1) how VIN-style networks embed planning and how its idea generalizes, (2) how
1031 is symmetry structure defined in path planning and how could it be injected into such planning
1032 networks.

1033 **Path planning as MDPs.** To answer the above two questions, we first need to understand how
1034 a VIN embeds a path planning problem into a convolutional network as some embedded MDP.
1035 Intuitively, the embedded MDP in a VIN is different from the original path planning problem, since
1036 (planar) convolutions are translation equivariant but there are different obstacles in different regions.

1037 For path planning on the 2D grid $\mathcal{S} = \mathbb{Z}^2$, the objective is to avoid some obstacle region $\mathcal{C}_{\text{obs}} \subset \mathbb{Z}^2$
1038 and navigate to the goal region $\mathcal{C}_{\text{goal}}$ through free space $\mathcal{C} \setminus \mathcal{C}_{\text{obs}}$. An action $a = \Delta s \in \mathcal{A}$ is to
1039 move from the current state s to a next *free* state $s' = s + \Delta s$, where for now we limit it to be in
1040 four directions: $\mathcal{A} = \{ \uparrow, \downarrow, \leftarrow, \rightarrow \}$. Assuming deterministic transition, the agent moves to s' with probability 1 if
1041 $s + \Delta s \in \mathcal{C} \setminus \mathcal{C}_{\text{obs}}$. If it hits an obstacle, it stays at s if $s + \Delta s \in \mathcal{C}_{\text{obs}}$: $P(s + \Delta s | s, \Delta s) = 0$ and
1042 $P(s | s, \Delta s) = 1$. Every move has a constant negative reward $R(s, a) = -1$ to encourage shortest
1043 path. We call this *ground* path planning MDP, a 5-tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$.

1044 **Constructing embedded MDPs.** However, such transition function is not translation-invariant,
1045 i.e. at different position, the transition probabilities are not related by any symmetry: $P(s' | s, a) \neq$
1046 $P(g.s' | g.s, g.a)$. Instead, we could always construct a "symmetric" MDP that has equivalent optimal
1047 value and policy for path planning problems, which is implicitly realized in VINs. The idea is to
1048 move the information of obstacles from transition function to reward function: when we hit some
1049 action $s + \Delta s \in \mathcal{C}_{\text{obs}}$, we instead allow transition $\bar{P}(s + \Delta s | s, \Delta s) = 1$ (with all other s' as 0
1050 probability) while set a "trap" with negative infinity reward $\bar{R}_m(s, \Delta s) = -\infty$. The reward function
1051 needs the information from the occupancy map M , indicating obstacles \mathcal{C}_{obs} and free space. For the
1052 free region, the reward is still a constant $\bar{R}_M(s, \Delta s) = -1$, indicating the cost of movement.

1053 We call it the *embedded* MDP, with different transition and reward function $\bar{\mathcal{M}} = \langle \mathcal{S}, \mathcal{A}, \bar{P}, \bar{R}_M, \gamma \rangle$,
1054 which converts the "complexity" in the transition function P in \mathcal{M} to the reward function \bar{R}_m
1055 in $\bar{\mathcal{M}}$. Here, map M shall also be treated as an "input", thus later we will derive how the group acts
1056 on the map $g.M$. It has the same optimal policy and value as the ground MDP \mathcal{M} , since the optimal
1057 policies in both MDPs will avoid obstacles in \mathcal{M} or trap cells in $\bar{\mathcal{M}}$. It could be easily verified by
1058 simulating value iteration backward in time from the goal position.

1059 The transition probability \bar{P} of the embedded MDP $\bar{\mathcal{M}}$ is for an "empty" maze and thus translation-
1060 invariant. Note that the reward function \bar{R} is not not necessarily invariant. This construction is not
1061 limited to 2D grid and generalizes to continuous state space or even higher dimensional space, such
1062 as \mathbb{R}^6 configuration space for 6-DOF manipulation.

1063 Note, all of this is what we use to conceptually understand how a VIN is possible to learn. The
1064 reward cannot be negative infinity, but the network will learn it to be smaller than all desired Q-
1065 values.

1066 H.2 Understanding steerable planning

1067 How do we deal with potential symmetry in path planning? how do we characterize it? We try to
1068 understand symmetric planning (**steerable** planning after integrating symmetry with equivariance)
1069 and how it is difference classic planning algorithms, such as A*, for planning under *symmetry*.

⁵We avoid the symbol π for policy since it is used for induced representation in [14, 16].

1070 **Steerable planning.** Recall that we generalize the idea of VIN by considering it as a planning
 1071 network that composes of mappings between steerable feature fields.

1072 The critical point is that, convolutions directly operate on local patches of pixels and never directly
 1073 touch coordinates of pixels. In analogy, this avoids a critical drawback in other *explicit* planning
 1074 algorithms: in sampling-based planning, a trajectory $(s_1, a_1, s_2, a_2, \dots)$ is sampled and inevitable
 1075 represented by states $\Omega = \mathcal{S}$. However, to find another symmetric state $g.s$, we potentially need to
 1076 compare it against all known states $\mathcal{S}' \subset \mathcal{S}$ with all symmetries $g \in G$. On high level, an implicit
 1077 planner can avoid such symmetry breaking and is more easily compatible with symmetry by using
 1078 equivariant constraints.

1079 We can use MDP homomorphism to understand this [3, 32].

1080 **MDP homomorphisms.** An *MDP homomorphism* $h : \mathcal{M} \rightarrow \overline{\mathcal{M}}$ is a mapping from one MDP
 1081 $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$ to another $\overline{\mathcal{M}} = \langle \overline{\mathcal{S}}, \overline{\mathcal{A}}, \overline{P}, \overline{R}, \gamma \rangle$ [3, 32]. h consists of a tuple of surjective
 1082 maps $h = \langle \phi, \{\alpha_s \mid s \in \mathcal{S}\} \rangle$, where $\phi : \mathcal{S} \rightarrow \overline{\mathcal{S}}$ is the state mapping and $\alpha_s : \mathcal{A} \rightarrow \overline{\mathcal{A}}$ is the
 1083 *state-dependent* action mapping. The mappings are constructed to satisfy the following conditions:

$$\begin{aligned} \overline{R}(\phi(s), \alpha_s(a)) &\triangleq R(s, a), \\ \overline{P}(\phi(s') \mid \phi(s), \alpha_s(a)) &\triangleq \sum_{s'' \in \phi^{-1}(\phi(s'))} P(s'' \mid s, a), \end{aligned} \quad (23)$$

1084 for all $s, s' \in \mathcal{S}$ and for all $a \in \mathcal{A}$.

1085 We call the *reduced* MDP $\overline{\mathcal{M}}$ the *homomorphic image* of \mathcal{M} under h . If $h = \langle \phi, \{\alpha_s \mid s \in \mathcal{S}\} \rangle$ has
 1086 *bijective* maps ϕ and $\{\alpha_s\}$, we call h an *MDP isomorphism*. Given MDP homomorphism h , (s, a)
 1087 and (s', a') are said to be *h-equivariant* if $\sigma(s) = \sigma(s')$ and $\alpha_s(a) = \alpha_{s'}(a')$.

1088 **Symmetry-induced MDP homomorphisms.** Given group G , an MDP homomorphism h is said
 1089 to be *group structured* if any state-action pair (s, a) and its transformed counterpart $g.(s, a)$ are
 1090 mapped to the same abstract state-action pair: $(\phi(s), \alpha_s(a)) = (\phi(g.s), \alpha_{g.s}(g.a))$, for all $s \in$
 1091 $\mathcal{S}, a \in \mathcal{A}, g \in G$. For convenience, we denote $g.(s, a)$ as $(g.s, g.a)$, where $g.a$ implicitly⁶ depends
 1092 on state s . Applied to the transition and reward functions, the transition function P is G -invariant
 1093 if P satisfies $P(g.s' \mid g.s, g.a) = P(s' \mid s, a)$, and reward function R is G -invariant if $R(g.s, g.a) =$
 1094 $R(s, a)$, for all $s \in \mathcal{S}, a \in \mathcal{A}, g \in G$.

1095 However, this only fits the type of symmetry in [9, 10]. And also, they cannot handle invariance to
 1096 translation \mathbb{Z}^2 . In our case, we need to augment the reward function with map M input:

$$R_{g.M}(g.s, g.a) = R_M(s, a), \quad (24)$$

1097 for all $s \in \mathcal{S}, a \in \mathcal{A}, g \in G = p4m$.

1098 This means that, at least for rotations and reflections D_4 , the MDPs constructed from transformed
 1099 maps $\{g.M\}$ are MDP *isomorphic* to each other.

1100 I Symmetric Planning Framework: Proofs

1101 We show the derivation and proofs for all theoretical results in this section.

1102 We follow the notation in [12] to use \star for (one-argument) convolution and \cdot for (two-argument)
 1103 multiplication:

$$E^a[V](s) = [P^a \cdot V](s) \equiv \sum_{s'} P^a(s' \mid s) \cdot V(s') \quad (25)$$

1104 I.1 Proof: equivariance of scalar-valued expected value operation

1105 We present the Theorem 4.1 here and its formal definition.

⁶The group operation acting on action space \mathcal{A} depends on state, since G actually acts on the *product space* $\mathcal{S} \times \mathcal{A}$: $(g, (s, a)) \mapsto g.(s, a)$, while we denote it as $(g.s, g.a)$ for consistency with $h = \langle \phi, \{\alpha_s \mid s \in \mathcal{S}\} \rangle$. As a bibliographical note, in van der Pol et al. [32], the group acting on state and action space is denoted as state transformation $L_g : \mathcal{S} \rightarrow \mathcal{S}$ and *state-dependent* action transformation $K_g^s : \mathcal{A} \rightarrow \mathcal{A}$.

Theorem I.1. *If transition is G -invariant, the expected value operator E over \mathbb{Z}^2 is G -equivariant:*

$$[g.E^a[V]](s) = [E^{g.a}[g.V]](s), \quad \text{for all } g = tr \in \mathbb{Z}^2 \rtimes D_4.$$

1106 *Proof.* E is the expected value operator. We also write the transition probability as

1107 Recall the G -invariance condition of transition probability, the group element g acts on s, a, s' :

$$\bar{P}(s' | s, a) = \bar{P}(g.s' | g.s, g.a) \equiv \bar{P}((tr).s' | (tr).s, r.a), \quad \forall g = tr \in \mathbb{Z}^2 \rtimes D_4, \forall s, a, s', \quad (26)$$

1108 where we can uniquely decompose any $g \in \mathbb{Z}^2 \rtimes D_4$ as $t \in \mathbb{Z}^2$ and $r \in D_4$ [14]. Note that, since
1109 the action is the difference between states $a = \Delta s = s' - s$, the translation part t acts trivially on it,
1110 so $g.a = (tr).a = r.a$ for all $r \in D_4$.

1111 We transform the feature field and show its equivariance:

$$[g.E^a[V]](s) \equiv [g.[P^a \cdot V]](s) \quad (27)$$

$$\equiv \sum_{s'} \rho_{\text{triv}}(r) P^a(s' | (tr)^{-1}.s) \cdot V(s') \quad (28)$$

$$= \sum_{s'} \rho_{\text{triv}}(r) P^{r.a}((tr).s' | s) \cdot V(s') \quad (29)$$

$$= \sum_{\tilde{s}'} \rho_{\text{triv}}(r) P^{r.a}(\tilde{s}' | s) \cdot V((tr)^{-1}\tilde{s}') \quad (30)$$

$$= \sum_{\tilde{s}'} P^{r.a}(\tilde{s}' | s) \cdot \rho_{\text{triv}}(r) V((tr)^{-1}\tilde{s}') \quad (31)$$

$$\equiv [P^{r.a} \cdot [g.V]](s) \quad (32)$$

$$\equiv [E^{r.a}[g.V]](s). \quad (33)$$

1112 We use the trivial representation $\rho_{\text{triv}}(g) = \text{Id}_{1 \times 1} = 1$ to emphasize that (1) the group element g acts
1113 on *feature fields* P^a and V , and (2) both feature fields P^a and V are scalar-valued and correspond
1114 to the one-dimensional trivial representation of $r \in D_4$.

1115 In the third line, we use the G -invariance of transition probability.

1116 The fourth line uses substitution $\tilde{s}' \triangleq (tr).s'$, for all $s' \in \mathbb{Z}^2$ and $tr \in \mathbb{Z}^2 \rtimes D_4$. This is an one-to-one
1117 mapping and the summation does not change.

1118 □

1119 I.2 Proof: expected value operator as steerable convolution

1120 In this section, we derive how to cast expected value operator as steerable convolution. The equiv-
1121 ariance proof is in the next section.

1122 In Theorem 4.1, we show equivariance of value iteration in 2D path planning, while it is only for
1123 the case that feature fields P^a and V are scalar-valued and correspond to one-dimensional trivial
1124 representation of $r \in D_4$.

1125 Here, we provide the derivation for Theorem 4.2 show that steerable CNNs [14] can achieve value
1126 iteration since we could construct the G -invariant transition probability as a steerable convolutional
1127 kernel. This generalizes Theorem 4.1 from scalar-valued kernel (for transition probability) with triv-
1128 ial representation to matrix-valued kernel with any combination of representations, enabling using
1129 stack (direct-sum) of feature fields and representations.

1130 We state Theorem 4.2 here for completeness:

1131 **Theorem I.2.** *If transition is G -invariant, there exists a (one-argument, isotropic) matrix-valued*
1132 *steerable kernel $P^a(s - s')$ (for every action), such that the expected value operator can be written*
1133 *as a steerable convolution and is G -equivariant:*

$$E^a[V] = P^a \star V, \quad [g.[P^a \star V]](s) = [P^{g.a} \star [g.V]](s), \quad \forall s \in \mathbb{Z}^2, \forall g \in \mathbb{Z}^2 \rtimes D_4. \quad (34)$$

1134 **Steerable kernels.** In our earlier definition, ψ^a and f_{in} are transition probability and value func-
 1135 tion, which are both real-valued $\psi^a : \mathbb{Z}^2 \rightarrow \mathbb{R}$, $f_{\text{in}} : \mathbb{Z}^2 \rightarrow \mathbb{R}$. However, this is a *special case*
 1136 which corresponds to use one-dimensional *trivial representation* of the fiber group D_4 . In the gen-
 1137 eral case in steerable CNNs [14, 16], we can choose the feature fields $\psi^a : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_{\text{out}} \times C_{\text{in}}}$ and
 1138 $f_{\text{in}} : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_{\text{in}}}$ and their fiber representations, which we will introduce the group representations
 1139 of D_4 and how to choose in practice in the next section.

1140 Weiler et al. [54] show that *convolutions* with *steerable kernels* $\psi^a : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_{\text{out}} \times C_{\text{in}}}$ is the most
 1141 general *equivariant linear map* between steerable feature space, transforming under ρ_{in} and ρ_{out} . In
 1142 analogy to the continuous version⁷ in [16], the convolution is equivariant *iff* the kernel satisfies a
 1143 H -steerability kernel constraint:

$$\psi^a(hs) = \rho_{\text{out}}(h)\psi^a(s)\rho_{\text{in}}(h^{-1}) \quad h \in H = D_4, s \in \mathbb{Z}^2. \quad (35)$$

1144 **Expected value operation as steerable convolution.** The foremost step is to show that the ex-
 1145 pected value operation is a form of convolution and is also G -equivariant. By definition, if we want
 1146 to write a (linear) operator as a form of convolution, we need one-argument kernel. Cohen et al. [12]
 1147 show that every linear equivariant operator is some convolution and provide more details. For our
 1148 case, this is formally shown as follows.

1149 **Proposition I.3.** *If the transition probability is G -invariant, it can be expressed as an (one-
 1150 argument) kernel $P^a(s'|s) = P^a(s' - s)$ that only depends on the difference $s' - s$.*

1151 *Proof.* The form of our proof is similar to [12], while its direction is different from us. We construct
 1152 a MDP such that the transition probability kernel is G -invariant, while Cohen et al. [12] assume the
 1153 linear operator $\psi \cdot f$ is linear *equivariant* operator on a homogeneous space, and then derive that the
 1154 kernel is G -invariant and expressible as one-argument kernel. Additionally, our kernel $\psi^a(s, s')$ and
 1155 $\psi^a(s - s')$ both live on the base space $B = \mathbb{Z}^2$ but not on the group $G = \mathbb{Z}^2 \rtimes D_4$.

1156 We show that the transition probability only depends on the difference $\Delta s = s' - s$, so we can define
 1157 the two-argument kernel $P^a(s'|s)$ on $\mathcal{S} \times \mathcal{S}$ by an one-argument kernel $P^a(s' - s)$ (for every action
 1158 a) on $\mathcal{S} = \mathbb{Z}^2$, without loss of generality:

$$P^a(s' - s) \equiv P^a(\mathbf{0}, s' - s) \quad (36)$$

$$= P^{g \cdot a}(g \cdot \mathbf{0}, g \cdot (s' - s)) \quad (37)$$

$$= P^{r \cdot a}((rs) \cdot \mathbf{0}, (rs) \cdot (s' - s)) \quad (38)$$

$$= P^{r \cdot a}(r \cdot s, r \cdot (s' - s + s)) \quad (39)$$

$$= P^{r \cdot a}(r \cdot s, r \cdot s') \quad (40)$$

$$= P^a(s, s'), \quad (41)$$

1159 where the second step uses G -invariance with $g = sr$, understood as the composition of a translation
 1160 $s \in \mathbb{Z}^2$ and a transformation in $r \in D_4$.

1161 □

1162 Additionally, we can also derive that, for the one-argument kernel, if we rotate state difference
 1163 $r \cdot (s' - s)$, the probability is the same for rotated action $r \cdot a$.

$$P^a(s' - s) = P^{r \cdot a}(r \cdot (s' - s)), \text{ for all } r \in D_4, s, s' \in \mathbb{Z}^2 \quad (42)$$

1164 The *expected value operator* with two-argument kernel can be then written as

$$E[V](s) \equiv [P^a \cdot V](s) = \sum_{s'} P^a(s'|s)V(s') = \sum_{s'} P^a(s' - s)V(s') \equiv [P^a \star V](s). \quad (43)$$

1165 Note that we do not differentiate between cross-correlation ($s' - s$) and convolution ($s - s'$).

⁷Weiler and Cesa [16] use letter G to denote the stabilizer subgroup $H \leq O(2)$ of $E(2)$.

1166 **I.3 Proof: equivariance of expected future value**

1167 Our derivation follows the existing work on group convolution and steerable convolution net-
 1168 works [15, 14, 16, 12]. However, the goal of providing the proof is not just for completeness,
 1169 but instead to emphasize the close connection between how we formulate our planning problem and
 1170 the literature of steerable CNNs, which explains and justifies our formulation.

1171 Additionally, there are several subtle differences worth to mention. (1) Throughout the paper, we
 1172 do not discuss kernels or fields that live on a group G to make it more approachable. Nevertheless,
 1173 group convolutions are a special case of steerable convolutions with fiber representation ρ as regular
 1174 representation. (2) We use \mathbb{Z}^2 as running example. Some prior work uses \mathbb{R}^2 or \mathbb{Z}^2 , but they are
 1175 merely just differ in integral and summation. (3) The definition of convolution and cross-correlation
 1176 might be defined and used interchangeably in the literature of (equivariant) CNNs.

1177 **Notation.** To keep notation clear and consistent with the literature [14, 12, 16], we denote the
 1178 transition probability $\bar{P}(s'|s, a) \triangleq \psi^a(s, s') \in \mathbb{R}$ (one kernel for an action) and value function as
 1179 $V(s') \triangleq f_{\text{in}}(s') \in \mathbb{R}$, and the resulting expected value as $f_{\text{out}}^a(s) = \sum_{s'} \psi^a(s, s') f_{\text{in}}(s')$ (given a
 1180 specific action a).

1181 **Transformation laws: induced representation.** For some group acting on the base space \mathbb{Z}^2 , the
 1182 signals $f : \mathbb{Z}^2 \rightarrow \mathbb{R}^c$ are transformed like [14]:

$$[\pi(g)f](x) = f(g^{-1}x) \quad (44)$$

1183 Apply a translation t and a transformation $r \in D_4$ to f , we get $\pi(tr)f$. The transformation law on
 1184 the input space f_{in} is [14, 16]:

$$f(x) \mapsto [\pi(tr)f](x) \triangleq \rho(r) \cdot [f((tr)^{-1}x)] \quad (45)$$

1185 The transformation law of the output space after applying π_{in} on input f_{in} is given by [14]:

$$[\psi \star f](x) \mapsto [\psi \star [\pi(tr)f]](x) \triangleq \rho(r) \cdot [[\psi \star f]((tr)^{-1}x)]. \quad (46)$$

1186 In our case, the output space is $f_{\text{out}}^a : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_{\text{out}}}$ and the input space is $f_{\text{in}} : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_{\text{in}}}$. Intuitively,
 1187 if we rotate a vector field (fibers represent arrows) by the induced representation $\pi(tr)$ of f , we also
 1188 need to rotate the direction of arrows by $\rho(r)$, $r \in D_4$.

1189 **Equivariance.** Now we prove the steerable convolution is equivariant:

$$[\psi^a \star [\pi_{\text{in}}(g)f_{\text{in}}]](s) = [\pi_{\text{out}}(g)f_{\text{out}}^a](s) \quad \forall s \in \mathcal{S}, \forall g \in G. \quad (47)$$

1190 The induced representation of input field f_{in} is induced by the fiber representation ρ_{in} , expressed by
 1191 $\pi_{\text{in}} \triangleq \text{ind}_H^G \rho_{\text{in}} = \text{ind}_{D_4}^{\mathbb{Z}^2 \times D_4} \rho_{\text{in}}$, where ρ_{in} is the fiber representation of group $H = D_4$. The induced
 1192 representation of output field π_{out} is analogously from ρ_{out} .

1193 Weiler and Cesa [16] proved equivariance of steerable convolutions for \mathbb{R}^2 case, while we include the
 1194 proof under our setup for completeness. The definition in [16] uses a form of *cross-correlation* and
 1195 we use *convolution*, while it is usually referred to interchangeably in the literature and is equivalent.
 1196 Cohen and Welling [14], Weiler et al. [54], Weiler and Cesa [16], Cohen et al. [12], Cohen [56]
 1197 provide more details and we refer the readers to them for more comprehensive account.

1198 The convolution on discrete grids \mathbb{Z}^2 with input field f_{in} transformed by the induced representation
 1199 π_{in} gives:




Figure 19: We attach a copy of the commutative diagram of SymVIN to show the equivariance of steerable value iteration. Commutative diagram for the full pipeline of SymVIN on steerable feature fields over \mathbb{Z}^2 (every grid). If rotating the input map M by $\pi_M(g)$ of any g , the output action $A = \text{SymVIN}(M)$ is guaranteed to be transformed by $\pi_A(g)$, i.e. the entire steerable SymVIN is equivariant under induced representations π_M and π_A : $\text{SymVIN}(\pi_M(g)M) = \pi_A(g)\text{SymVIN}(M)$. We use stacked feature fields to emphasize that SymVIN supports direct-sum of representations beyond scalar-valued.

$$\begin{aligned}
[\psi^a \star [\pi_{\text{in}}(rt)f_{\text{in}}]](s) &= \sum_{s' \in \mathbb{Z}^2} \psi^a(s - s') [\pi_{\text{in}}(rt)f_{\text{in}}](s') \\
&= \sum_{s' \in \mathbb{Z}^2} \psi^a(s - s') \rho_{\text{in}}(r) f_{\text{in}}(r^{-1}(s' - t)) \\
&= \sum_{s' \in \mathbb{Z}^2} \rho_{\text{out}}(r) \psi^a(r^{-1}(s - s')) \rho_{\text{in}}(r)^{-1} \rho_{\text{in}}(r) f_{\text{in}}(r^{-1}(s' - t)) \\
&= \rho_{\text{out}}(r) \sum_{s' \in \mathbb{Z}^2} \psi^a(r^{-1}(s - s')) f_{\text{in}}(r^{-1}(s' - t)) \\
&= \rho_{\text{out}}(r) \sum_{\tilde{s} \in \mathbb{Z}^2} \psi^a(r^{-1}(s - t) - \tilde{s}) f_{\text{in}}(\tilde{s}) \\
&= \rho_{\text{out}}(r) f_{\text{out}}(r^{-1}(s - t)) \\
&= [\pi_{\text{out}}(rt) f_{\text{out}}^a](s),
\end{aligned} \tag{48}$$

1200 where $s' \in \mathcal{S} = \mathbb{Z}^2$, and thus satisfies the equivariance condition:

$$[\psi^a \star [\pi_{\text{in}}(rt)f_{\text{in}}]](s) = [\pi_{\text{out}}(rt) f_{\text{out}}^a](s), \forall s \in \mathbb{Z}^2, \forall rt \in \mathbb{Z}^2 \rtimes D_4. \tag{49}$$

- 1201 1. Definition of \star
- 1202 2. Transformation law of the induced representation π_{in} [14, 16]
- 1203 3. Kernel steerability $\psi^a(s) = \rho_{\text{out}}(h)\psi^a(h^{-1}s)\rho_{\text{in}}(h^{-1})$ [16]
- 1204 4. Move and cancel
- 1205 5. Substitutes $\tilde{s} = r^{-1}(s' - t)$, $r^{-1}s' = r^{-1}t + \tilde{s}$, so $r^{-1}(s - s') = r^{-1}(s - t) - \tilde{s}$.
- 1206 Since $r \in D_4$ and $s - s' \in \mathbb{Z}^2$, the result is still in $p4m$, it is one-to-one correspondence
- 1207 $p4m \times \mathbb{Z}^2 \rightarrow \mathbb{Z}^2$, and the summation does not change. Weiler and Cesa [16] analogously
- 1208 considers the continuous case, where D_4 is orthogonal transformations so the Jacobian is
- 1209 always 1.
- 1210 6. Definition of \star
- 1211 7. Transform law of the induced representation π_{out}

1212 I.4 Proof: equivariance of steerable value iteration

1213 As the third and final step, we would like to show that the full steerable value iteration pipeline
1214 is equivariant under $G = \mathbb{Z}^2 \rtimes D_4$. We need to show that every operation in the steerable value
1215 iteration is equivariant.

1216 The key is to prove that \max_a is an equivariant non-linearity over feature fields, which follows
 1217 Section D.2 in [16].

1218 **Step 1:** $V \mapsto Q$. Here, we prove the equivariance of $Q_k^a(s) = \bar{R}_M^a(s) + \gamma \times [\bar{P}_\theta^a \star V_k](s)$. First,
 1219 let the group acts on both sides:

$$Q_k^a(s) = \bar{R}_M^a(s) + \gamma \times [\bar{P}_\theta^a \star V_k](s) \quad (50)$$

$$\iff [\pi_{\text{out}}(g)Q_k^a](s) = [\pi_{\text{out}}(g)\bar{R}_M^a](s) + \gamma \times [\pi_{\text{out}}(g)[\bar{P}_\theta^a \star V_k]](s) \quad (51)$$

$$\iff [\pi_{\text{out}}(g)Q_k^a](s) = [\pi_{\text{out}}(g)\bar{R}_M^a](s) + \gamma \times [\bar{P}_\theta^a \star [\pi_{\text{in}}(g)V_k]](s) \quad (52)$$

$$\iff Q_k^{g \cdot a}(g^{-1}s) = \bar{R}_{g \cdot M}^{g \cdot a}(g^{-1}s) + \gamma \times [\bar{P}_\theta^{g \cdot a} \star V_k](g^{-1}s) \quad (53)$$

$$\iff Q_k^{\tilde{a}}(\tilde{s}) = \bar{R}_{\pi_M(g)M}^{\tilde{a}}(\tilde{s}) + \gamma \times [\bar{P}_\theta^{\tilde{a}} \star V_k](\tilde{s}) \quad (54)$$

1220 The the last step we substitute $\tilde{s} = g^{-1}s$ and $\tilde{a} = g \cdot a$.

1221 $M : \mathbb{Z}^2 \rightarrow \{0, 1\}^2$ is the concatenation of maze occupancy map and goal map, which also lives on
 1222 \mathbb{Z}^2 . We use two copies of trivial representations as fiber representation ρ_M , and denote the induced
 1223 representation of the field M as π_M .

1224 Then, we prove the equivariance: if we transform the occupancy map (and goal map), the value
 1225 iteration should have both input V and output Q transformed. Since this is an iterative process, the
 1226 only input to the value iteration is actually the occupancy map $M : \mathbb{Z}^2 \rightarrow \{0, 1\}^2$.

1227 Before that, we observe that the reward also has G -invariance when we have map as input:

$$\bar{R}_M^a(s) = \bar{R}_{g \cdot M}^{g \cdot a}(g \cdot s). \quad (55)$$

1228 Additionally, since the reward $\bar{R}_M^a(s)$ means the reward at given position in map M **after executing**
 1229 **action** a , when we transform the map, we also need to transform the action: $\bar{R}_{g \cdot M}^{g \cdot a}(s)$.

1230 Since it is iterative process, let the Q -map being transformed by g :

$$[g \cdot Q_k^a](s) = Q_k^a(g^{-1}s) \quad (56)$$

$$= \bar{R}_M^a(g^{-1}s) + \gamma \times [\bar{P}_\theta^a \star V_k](g^{-1}s) \quad (57)$$

$$= \bar{R}_{g \cdot M}^{g \cdot a}(s) + \gamma \times [\bar{P}_\theta^a \star V_k](g^{-1}s) \quad (58)$$

$$= \bar{R}_{g \cdot M}^{g \cdot a}(s) + \gamma \times [\bar{P}_\theta^{g \cdot a} \star [g \cdot V_k]](s) \quad (59)$$

1231 The second last step uses the G -invariance condition $\bar{R}_M^a(s) = \bar{R}_{g \cdot M}^{g \cdot a}(g \cdot s)$. The last step uses the
 1232 equivariance of steerable convolution.

1233 It should be understood as: (1) transforming map $g \cdot M$ and action $g \cdot a$, is always equal to (2) trans-
 1234 forming values $[g \cdot Q_k^a]$ and $[g \cdot V_k]$. This proves the equivariance visually shown in Figure 19.

1235 **Step 2:** $Q \mapsto V$. The second step is to show for $V_{k+1}(s) = \max_a Q_k^a(s)$.

1236 Intuitively, we sum over every channel of each representation. For example, if we have N copies
 1237 of the regular representation with size $|D_4| = 8$, we transform the tensor $(N \times 8) \times m \times m$ to
 1238 $(1 \times 8) \times m \times m$ along the N channel. Thus, how we use the 8×8 regular representation to
 1239 transform the $N \times 8$ channels still holds for 1×8 , which implies equivariance. The $m \times m$ spatial
 1240 map channels form the base space \mathbb{Z}^2 and are transformed as usual (spatially rotated).

1241 Weiler and Cesa [16] provide detailed illustration and proofs for equivariance of different types of
 1242 non-linearities.

1243 **Step 3: multiple iterations.** Since each layer is equivariant (under induced representations), Co-
 1244 hen and Welling [15], Kondor and Trivedi [13], Cohen et al. [12] show that stacking multiple equiv-
 1245 ariant layers is also equivariant. Thus, we know iteratively applying step 1 and 2 (*equivariant*
 1246 *steerable Bellman operator*) is also *equivariant (steerable value iteration)*.

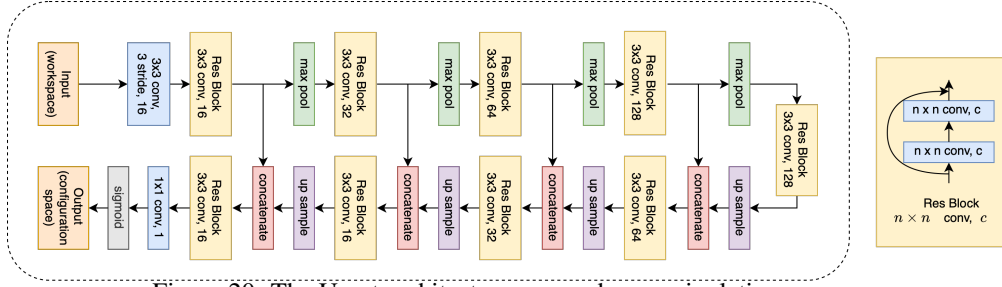


Figure 20: The U-net architecture we used as manipulation mapper.

1247 J Practice and Implementation Details

1248 J.1 Note

1249 We provide additional practical and implementation details, and leave results in the next section.

1250 J.2 Building Mapper Networks

1251 **For visual navigation.** For navigation, we follow the setting in GPPN [18]. The input is $m \times m$
 1252 panoramic egocentric RGB images in 4 directions of resolution $32 \times 32 \times 3$, which forms a tensor of
 1253 $m \times m \times 4 \times 32 \times 32 \times 3$. A mapper network converts every image into a 256-dimensional embedding
 1254 and results in a tensor in shape $m \times m \times 4 \times 256$ and then predicts map layout $m \times m \times 1$.

1255 For the first image encoding part, we use a CNN with first layer of 32 filters of size 8×8 and stride
 1256 of 4×4 , and second layer with 64 filters of size 4×4 and stride of 2×2 , with a final linear layer
 1257 of size 256.

1258 The second obstacle prediction part, the first layer has 64 filters and the second layer has 1 filter, all
 1259 with filter size 3×3 and stride 1×1 .

1260 **For workspace manipulation.** For **workspace manipulation**, we use U-net [58] with residual-
 1261 connection [59] as a mapper, see Figure.20. The input is 96×96 top-down occupancy grid of the
 1262 workspace with obstacles, and the target is to output 18×18 configuration space as the maps for
 1263 planning.

1264 During training, we pre-train the mapper and the planner separately for 15 epochs. Where the
 1265 mapper takes manipulator workspace and outputs configuration space. The mapper is trained to
 1266 minimize the binary cross entropy between output and ground truth configurations space. The plan-
 1267 ner is trained in the same way as described in Section 6.1. After pre-training, we switch the input to
 1268 the planner from ground truth configuration space to the one from the mapper. During testing, we
 1269 follow the pipeline in [37] that the mapper-planner only have access to the manipulator workspace.

1270 J.3 SymGPPN

1271 ConvGPPN [Redacted for anonymous review] is inspired by VIN and GPPN. To avoid the training
 1272 issues in VIN, GPPN proposes to use LSTM to alleviate them. In particular, it does not use max
 1273 pooling in the VIN. Instead, it uses a CNN and LSTM to mimic the value iteration process. Con-
 1274 vGPPN, on the other hand, integrates CNN into LSTM, resulting in a single component convLSTM
 1275 for value iteration. We found that ConvGPPN performs better than GPPN in most cases. Based on
 1276 ConvGPPN, SymGPPN replaces each convolutional layer with steerable convolutional layer.

1277 J.4 Understand group conv and "augmented state"

1278 We derive the relationship between group convolution and steerable convolution in Section G.3.

1279 The augmented state $\mathbb{Z}^2 \times D_4 \rightarrow \mathbb{R}$ can be similarly treated on the group $p4m = \mathbb{Z}^2 \rtimes D_4$. It is
 1280 equivalent to using regular representation on the base space \mathbb{Z}^2 as $\mathbb{Z}^2 \rightarrow \mathbb{R}^8$.

1281 **J.5 Implementation of max operation**

1282 Here, we consider how to implement the max operation in $V_{k+1}(s) = \max_a Q_k^a(s)$. The max is
1283 taken over every state, so the computation mainly depends on our choice of fiber representation.

1284 For example, if we use *trivial representations* for both input and output, the input would be $Q_k : \mathbb{Z}^2 \rightarrow \mathbb{R}^{1 \times C_A}$ and the output is state-value $V_k : \mathbb{Z}^2 \rightarrow \mathbb{R}$. This recovers the default value iteration
1285 since we take max over \mathbb{R}^{C_A} vector.
1286

1287 In steerable CNNs, we can use stack of fiber representations. We can choose from regular-regular,
1288 trivial-trivial, and regular-trivial (trivial-regular is not considered).

1289 We already covered *trivial* representations for both input and output, they would be $Q_k : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_Q \times C_A}$ and $V_k : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_V}$ with $C_Q = C_V = 1$, since every channel would need a trivial
1290 representation.
1291

1292 If we use *regular* representation for Q and *trivial* for V , they are $Q_k : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_Q \times C_A}$ and $V_k : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C_V}$ with $C_Q = |D_4| = 8$ and $C_V = 1$. It degenerates that we just take max over all
1293 $C_Q * C_A$ channels.
1294

1295 For both using regular representations, we need to make sure they use the same fiber group (such as
1296 D_4 or C_4), so $C_Q = C_V$. If using D_4 , we have $Q_k : \mathbb{Z}^2 \rightarrow \mathbb{R}^{8 \times C_A}$ and $V_k : \mathbb{Z}^2 \rightarrow \mathbb{R}^8$, and we
1297 take max over every C_A channels (for every location) and have 8 channels left, which are used as
1298 $\mathbb{Z}^2 \rightarrow \mathbb{R}^8$.

1299 Empirically, we found using regular representations for both works the best overall.