
Group Equivariant Deep Reinforcement Learning

Arnab Kumar Mondal^{1,2} Pratheeksha Nair^{1,3} Kaleem Siddiqi^{1,2}

Abstract

In Reinforcement Learning (RL), Convolutional Neural Networks (CNNs) have been successfully applied as function approximators in Deep Q-Learning algorithms, which seek to learn action-value functions and policies in various environments. However, to date, there has been little work on the learning of symmetry-transformation equivariant representations of the input environment state. In this paper, we propose the use of Equivariant CNNs to train RL agents and study their inductive bias for transformation equivariant Q-value approximation. We demonstrate that equivariant architectures can dramatically enhance the performance and sample efficiency of RL agents in a highly symmetric environment while requiring fewer parameters. Additionally, we show that they are robust to changes in the environment caused by affine transformations.

1. Introduction

Reinforcement Learning has always faced the challenge of handling high dimensional sensory input, such as that given by vision or speech. To this end, it was demonstrated that a convolutional neural network could directly learn control policies from raw video data, with success in various Atari game environments (Mnih et al., 2013). More recently, there has been work to improve both the feature extraction from raw images (Grattarola, 2017) as well as the underlying Deep Q-Learning algorithm (Schaul et al., 2015; Horgan et al., 2018; Van Hasselt et al., 2016; Wang et al., 2015). Following this, a variety of models focusing on short-term memory (Kapturowski et al., 2018), episodic memory (Badia et al., 2020b) and meta controlling (Badia et al., 2020a) have been introduced. Despite these advances, the generalization of trained agents to new environments and the improvement of sample efficiency has not been widely explored. One way

¹School of Computer Science, McGill University, Montreal, Canada ²Centre for Intelligent Machines, McGill University, Montreal, Canada ³McGill Centre for Bioinformatics, McGill University, Montreal, Quebec, Canada. Correspondence to: Arnab Mondal <arnabm@cim.mcgill.ca>.

to tackle this problem is to apply standard regularization techniques such as L2 regularization, dropout (Srivastava et al., 2014), data augmentation and batch normalization (Ioffe & Szegedy, 2015), as proposed in (Farebrother et al., 2018; Cobbe et al., 2018). Approaches rooted in Meta-RL have also been proposed to address the generalization problem (Wang et al., 2016; Dasgupta et al., 2019; Kirsch et al., 2019).

In this work, we exploit the intrinsic properties of an environment, such as its symmetry, to improve the performance of Deep RL algorithms. In particular, we consider the efficacy of using an E(2)-Equivariant CNN (Weiler & Cesa, 2019) architecture as a function approximator for training RL agents using an Equivariant Q-Learning algorithm. We show that in a game environment, with a high degree of symmetry, such an approach provides a significant performance gain and improves sample efficiency as it learns from fewer experience samples. We further show that the inherent inductive bias for the equivariance of symmetry transformation of our proposed approach, enables the effective transfer of knowledge across previously unseen transformations of the environment. Our proposed method is complementary to the other generalization ideas in RL mentioned earlier, and hence can be used in conjunction with them. Using the proposed method adds negligible computational overhead, improves generalization and facilitates a higher degree of parameter sharing. The ideas explored in this paper could be extended to more challenging RL tasks, such as path planning in dynamic environments, where the dynamics is given by symmetry transformation, using aerial views. In such tasks, the policy may be designed to be equivariant to symmetric transformations of the viewpoint.

The rest of the paper is organized as follows: Section 2 gives a brief overview of relevant background. In Section 3, we review the theory of E(2)-equivariant convolution and introduce our Equivariant DQN model. Finally, we present empirical results on two environments, Snake and Pacman, in Section 4, demonstrating the promise of equivariant Deep RL.

2. Background

Group equivariant CNNs (G-CNN) (Cohen & Welling, 2016) exploit the group of symmetries of input images to reduce sample complexity, learn faster and improve the

capacity of CNNs without increasing the number of parameters. This network architecture uses a new convolution layer whose output feature map changes equivariantly with the group action on the input feature map and promotes higher degrees of weight sharing. The theory of steerable CNNs (Cohen & Welling, 2017; Weiler & Cesa, 2019; Weiler et al., 2018) generalizes this idea to continuous groups and homogeneous spaces. In this work, we focus on using an E(2)-Equivariant Steerable CNN (Weiler & Cesa, 2019) architecture for deep RL.

Given an input signal, CNNs extract a hierarchy of feature maps. The weight-sharing of the convolution layers makes them inherently translation-equivariant so that a translated input signal results in a corresponding translation of the feature maps (Cohen & Welling, 2016). An E(2)-Equivariant Steerable CNN carries out translation, rotation and reflection equivariant convolution on the image plane. The feature spaces of such Equivariant CNNs are defined as spaces of feature fields and are characterized by a group representation that determines their transformation behaviour under transformations of the input, as discussed in Section 3.1.

The Deep Q-learning Network (DQN) (Mnih et al., 2013) has been widely used in RL since its inception. The DQN utilizes “experience replay” (Lin, 1993) where the agent’s experiences at each time-step are stored in a memory buffer, and the Q-learning updates are done on samples drawn from this buffer, which breaks the correlation between them. A variant of this strategy is the “prioritized replay buffer” (Schaul et al., 2015), where the experiences are sampled according to their importance. A second variant, the Double DQN or DDQN (Van Hasselt et al., 2016), addresses the problem of maximization bias, which occurs due to the usage of the same Q network for the off-policy bootstrapped target. An additional improvement is the use of an advantage function and the learning of a value function to determine the action-values using a common convolutional feature learning module, in a Dueling Network (Wang et al., 2015). We experiment with the above mentioned variants.

3. Method

3.1. E(2)-equivariant convolution

In this section, we briefly describe the theory behind E(2)-equivariant convolution. First, we define the group $T(2) \rtimes G$ where $G \leq O(2)$. $T(2)$ is a translational group on \mathbb{R}^2 and G is a subgroup of the orthogonal group $O(2)$, which are continuous rotations and reflections under which the origin is invariant. Intuitively, we are dealing with the subgroups of the group of isometries of a 2-D plane called E(2). In contrast to regular CNNs, which work with a stack of multiple channels of features $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, the steerable CNN defines a steerable feature space of feature fields $f : \mathbb{R}^2 \rightarrow \mathbb{R}^c$ which associates a c dimensional feature vector $f(x) \in \mathbb{R}^c$

to every $x \in \mathbb{R}^2$. The feature fields are linked to a transformation law that defines their transformations under the action of a group. The transformation law of a feature field is characterized by the group representation $\rho : G \mapsto \text{GL}(\mathbb{R}^c)$, where $\text{GL}(\mathbb{R}^c)$ represents the group of all invertible $c \times c$ matrices. This defines how each of these c channels mixes when the vector $f(x)$ is transformed. The operator for a transformation tg , where $t \in T(2)$ and $g \in G$, is given by:

$$\left([\text{Ind}_G^{T(2) \rtimes G} \rho] (tg) \cdot f \right) (x) := \rho(g) \cdot f(g^{-1}(x - t)) \quad (1)$$

where $[\text{Ind}_G^{T(2) \rtimes G} \rho]$ is called the induced representation. Analogous to the channels of a regular CNN, we can stack multiple feature fields f_i with their corresponding representation ρ_i and the stack $\bigoplus_i f_i$ then transforms under $\bigoplus_i \rho_i$, which is a block diagonal matrix. Notice that due to ρ being a block diagonal matrix each feature field transforms independently. Having described the feature fields, we will next give the equation for equivariance and the constraint it imposes on the convolution kernel. Consider two feature fields $f_{in} : \mathbb{R}^2 \rightarrow \mathbb{R}^{c_{in}}$ with representation ρ_{in} , $f_{out} : \mathbb{R}^2 \rightarrow \mathbb{R}^{c_{out}}$ with representation ρ_{out} and a convolution kernel $k : \mathbb{R}^2 \rightarrow \mathbb{R}^{c_{out} \times c_{in}}$ then the desired equivariance is given by:

$$k * \left([\text{Ind}_G^{T(2) \rtimes G} \rho_{in}] (tg) \cdot f_{in} \right) = [\text{Ind}_G^{T(2) \rtimes G} \rho_{out}] (tg) \cdot (k * f_{out}) \quad (2)$$

where convolution is defined as usual as:

$$f_{out}(x) := (k * f_{in})(x) = \int_{\mathbb{R}^2} k(y) f_{in}(x + y) dy \quad (3)$$

This can only be achieved if we restrict ourselves to G-steerable kernels which satisfy the kernel constraint:

$$k(gx) := \rho_{out}(g)k(x)\rho_{in}(g^{-1}) \quad \forall g \in G \& x \in \mathbb{R}^2 \quad (4)$$

Imposing this constraint on the kernels significantly reduces the number of parameters and promotes parameter sharing. Also, by obtaining equivariance in each convolution layer of the network, they can be composed to extract equivariant features from the input 2D image signal. Further details on the kernel basis are provided in (Weiler & Cesa, 2019).

3.2. Environment

In this work, we primarily experiment with two environments - the Snake game of the Pygame Learning Environment (Tasfi, 2016) and the Atari Pacman environment (Brockman et al., 2016)¹. In the Snake game², the agent is a snake which grows in length each time it feeds on a

¹<https://gym.openai.com/envs/MsPacman-v0/>

²<https://pygame-learning-environment.readthedocs.io/en/latest/user/games/snake.html>

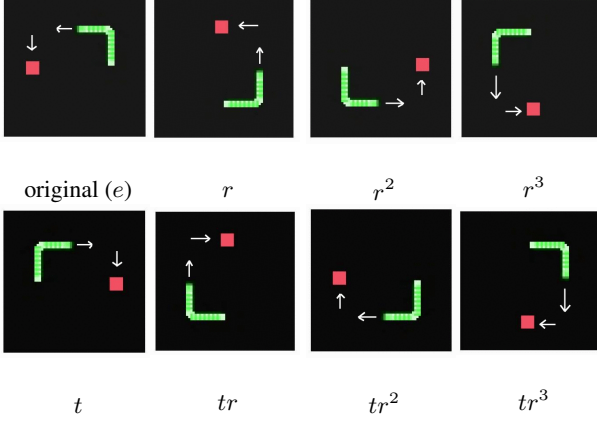


Figure 1. If we denote the 90 degree clockwise rotation by r and reflection over the vertical axis as t , then the group elements of D_4 are $\{e, r, r^2, r^3, t, tr, tr^2, tr^3\}$ where e is the identity action. These panels show the action of these group elements (transformations) on a game screen and how they affect the optimal policy (shown by white arrows).

food particle and gets a reward of +1. The food particle is randomly placed somewhere inside the valid area of a screen. The snake can choose four legal actions: move up, move down, move left, and move right. A terminal state is reached when the snake comes in contact with its body or the walls, and the agent then receives a score of -1. From Figure 1, we see that under the action of group elements of D_4 , the current optimal policy should change equivariantly, which suggests the possible benefits of learning the Q values for each action using equivariant features extracted from the game screen.

The Pacman game consists of a maze, a player agent and a few ghosts. Food particles are placed along the paths of the maze while the ghosts move freely around it. The player agent is also allowed four actions - move up, move down, move right and move left and it gets a positive reward for each particle it consumes without running into any of the ghosts. The game screen has a global D_1 symmetry and a degree of local D_4 symmetry.

3.3. Equivariant Deep Q-Network

Henceforth in this paper, “equivariant convolution” refers to $E(2)$ -equivariant steerable convolution. Suppose our pre-processed input is of dimension $m \times d \times d$ where m is the number of channels, and $d \times d$ is the size of the image. We convert it into a feature field represented by $s = \bigoplus_{i \in I} s_i$ where $I = \{1, \dots, m\}$ and s_i is an image of dimension $d \times d$. The transformation law of each channel is given by trivial representation (ρ_{triv}) of a chosen discrete group (G) for each channel. We further choose a regular representation (ρ_{reg}) for intermediate feature fields, which are permutation matrices given a group element $g \in G$, to derive the kernel basis of equivariant convolution. Using regular representa-

tion preserves the equivariance with point-wise nonlinear activation functions such as ReLU. We stack equivariant convolutions followed by ReLU to obtain an equivariant feature extractor $F_{eqv} : \mathbb{R}^{m \times d \times d} \rightarrow \mathbb{R}^n$ where n denotes the dimension of extracted equivariant features. The detailed architecture of this feature extractor and its relationship to the vanilla feature extractor we use in DDQN are in Appendix B. A discussion on how to choose the group and its representation for a feature field along with group restriction is included in Appendix A. Assuming that we do not restrict the group G along the depth of the network, our transformation rule of the extracted feature vector with respect to the transformation of input is given by:

$$F_{eqv} \left([\text{Ind}_G^{T(2) \times G} \rho_{triv}] (tg) s; \theta \right) = \rho(g) (F_{eqv} (s; \theta)) \quad (5)$$

where $\rho(g) = \bigoplus_{j \in J} \rho_{reg}(g)$ and

$$\left([\text{Ind}_G^{T(2) \times G} \rho_{triv}] (tg) s \right) (x) := \bigoplus_{i \in I} [s_i (g^{-1} (x - t))]$$

Note that Equation 5 gives the desired equivariance and $J = \{1, \dots, (n/N)\}$ where N is the order of the g . N divides n and n/N is the number of feature fields at the output. Intuitively, Equation 5 means that at every feature field the values permute along its dimension when we transform the input by some group element. Also note that if we restrict the group along the depth we will have $g_{res} \in G_{res} \leq G$ in the RHS of Equation 5 instead of g . Having obtained the feature vector which transforms equivariantly we can add a final linear layer to obtain the Q values:

$$Q_{eqv}(s, a, \phi) = W_a \cdot F_{eqv}(s; \theta) + b_a \quad (6)$$

where $a \in \mathcal{A} = \{1, \dots, |\mathcal{A}|\}$, $W_a = [w_{a1} \dots w_{an}]$ and $\phi = \{\theta, W_1, \dots, W_{|\mathcal{A}|}, b_1, \dots, b_{|\mathcal{A}|}\}$ (the set of all parameters). The linear layer learns whether or not to preserve the equivariance in output depending on the environment. We use DDQN as our baseline model throughout this work whose final loss at iteration l is given by:

$$L(\phi_l) = \mathbb{E}_{s,a,r,s'} \left[\left(y_l^{DDQN} - Q_{eqv}(s, a, \phi_l) \right)^2 \right] \quad (7)$$

with target:

$$y_l^{DDQN} = r + \gamma Q_{eqv}(s', \arg \max_{a'} Q_{eqv}(s', a', \phi_l), \phi^-) \quad (8)$$

where ϕ^- represents the parameters of the frozen network. The gradients computed through both the linear and the equivariant feature extractor networks are backpropogated to update their parameters.

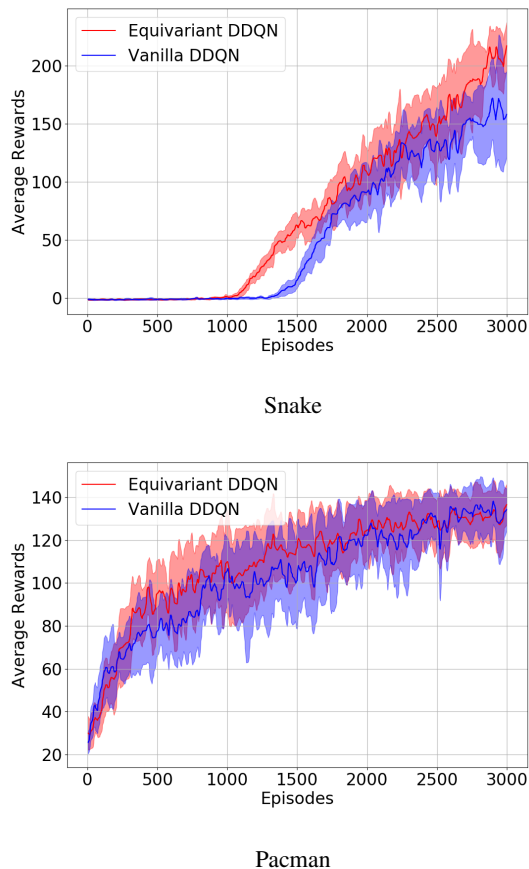


Figure 2. Plots of evolution of average rewards with the number of episodes. We also show the confidence intervals over 10 different seeds. Note that the plots are smoothed with a 1D Gaussian filter with $\sigma=3$ for better visualization.

4. Experiments

We first consider the performance of a carefully designed equivariant DDQN, keeping in mind the symmetry of the game (refer to Appendix B), compared to a vanilla DDQN. For a fair comparison, we keep the settings of the environment and hyperparameters the same for all the experiments³. We report in Figure 2 the evolution of rewards collected over the training episodes for both the models in the Snake and the Pacman environments. Our proposed model attains a 30% improvement in average reward collected after training for 3000 episodes in the highly symmetric Snake environment. It also learns faster with a 90% reduction in the number of parameters. This verifies our hypothesis that parameters required to learn policies of the identity transformation would be sufficient to generalize to optimal policies in other transformations for the Snake environment. In the case of Pacman, we notice our model performs slightly better in the initial episodes, with a 34% reduction in the number of parameters. But once both the models have seen

³Link to the code: <https://github.com/arnab39/EquivariantDQN>

enough samples, the margin of difference vanishes. In Appendix C, we show that the proposed method gives similar results with other subsequent improvements, such as using DDQN with priority replay and the Dueling architecture.

We further investigate the usefulness of the inherent inductive bias in the model in transfer learning with respect to the affine transformation of the environment screen. For this part, we remove the group restriction from Equivariant DDQN of the Pacman game and make the feature extractor D_4 equivariant. First we train both the Vanilla and Equivariant model. We then change the environment by rotating the input screen by 90 degrees clockwise (r). Leaving the rest of the network frozen, we retrain the final linear layer for this new environment. We show in Table 1 that while a regular CNN based feature extractor fails, the Equivariant feature extractor can still find a decent policy after learning the linear layers for certain epochs. The results of a simple path planning problem like Snake, indicate that our model can, in principle, be extended to more complex continuous path planning problems such as in UAVs (Zhang et al., 2015; Challita et al., 2018). Such scenarios would benefit both from faster learning due to increased sample efficiency and viewpoint transformation equivariant features for optimal policy learning, which can generalize to new transformations of the environment.

Transformation	Vanilla DDQN	Equivariant DDQN
e	129 ± 2.3	125 ± 4.5
r	48.9 ± 2	99 ± 4.7
r^2	53 ± 3.5	104 ± 3.4
r^3	51 ± 1.9	98 ± 3.9

Table 1. Average reward over 200 episodes of Pacman for 5 seeds reported with a confidence level of 95% for different environment transformations. e is the original screen.

5. Conclusions and future work

We have introduced an Equivariant Deep-Q learning algorithm, and have demonstrated that it provides a considerable boost to performance with parameter and sample efficiency when carefully designed for highly symmetric environments. We have also shown that this approach generalizes policies well to new unseen environments obtained by an affine transformation of the original environment. Although invariant models in supervised learning were shown to make the models robust, to the best of our knowledge, this is the first time equivariant learning has been proposed in a Deep RL framework. In follow-up work, we plan to implement continuous rotation and reflection group equivariance using an irreducible representation of the O_2 group, for more challenging path planning environments, with an extension to a continuous action space.

6. Acknowledgement

We would like to thank Gabriele Cesa for his valuable comments on the E(2)-equivariant convolution.

References

- Badia, A. P., Piot, B., Kapturowski, S., Sprechmann, P., Vitvitskiy, A., Guo, D., and Blundell, C. Agent57: Outperforming the atari human benchmark. *arXiv preprint arXiv:2003.13350*, 2020a.
- Badia, A. P., Sprechmann, P., Vitvitskiy, A., Guo, D., Piot, B., Kapturowski, S., Tieleman, O., Arjovsky, M., Pritzel, A., Bolt, A., et al. Never give up: Learning directed exploration strategies. *arXiv preprint arXiv:2002.06038*, 2020b.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Challita, U., Saad, W., and Bettstetter, C. Deep reinforcement learning for interference-aware path planning of cellular-connected uavs. In *2018 IEEE International Conference on Communications (ICC)*, pp. 1–7. IEEE, 2018.
- Cobbe, K., Klimov, O., Hesse, C., Kim, T., and Schulman, J. Quantifying generalization in reinforcement learning. *arXiv preprint arXiv:1812.02341*, 2018.
- Cohen, T. and Welling, M. Group equivariant convolutional networks. In *International conference on machine learning*, pp. 2990–2999, 2016.
- Cohen, T. S. and Welling, M. Steerable cnns. *International Conference on Learning Representations (ICLR)*, 2017.
- Dasgupta, I., Wang, J., Chiappa, S., Mitrovic, J., Ortega, P., Raposo, D., Hughes, E., Battaglia, P., Botvinick, M., and Kurth-Nelson, Z. Causal reasoning from meta-reinforcement learning. *arXiv preprint arXiv:1901.08162*, 2019.
- Farebrother, J., Machado, M. C., and Bowling, M. Generalization and regularization in dqn. *arXiv preprint arXiv:1810.00123*, 2018.
- Grattarola, D. Deep feature extraction for sample-efficient reinforcement learning. 2017.
- Horgan, D., Quan, J., Budden, D., Barth-Maron, G., Hessel, M., Van Hasselt, H., and Silver, D. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*, 2018.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Kapturowski, S., Ostrovski, G., Quan, J., Munos, R., and Dabney, W. Recurrent experience replay in distributed reinforcement learning. 2018.
- Kirsch, L., van Steenkiste, S., and Schmidhuber, J. Improving generalization in meta reinforcement learning using learned objectives. *arXiv preprint arXiv:1910.04098*, 2019.
- Lin, L.-J. Reinforcement learning for robots using neural networks. Technical report, Carnegie-Mellon Univ Pittsburgh PA School of Computer Science, 1993.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- Tasfi, N. Pygame learning environment. *GitHub repository*, 2016.
- Van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double q-learning. In *Thirtieth AAAI conference on artificial intelligence*, 2016.
- Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., Blundell, C., Kumaran, D., and Botvinick, M. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.
- Wang, Z., Schaul, T., Hessel, M., Van Hasselt, H., Lanctot, M., and De Freitas, N. Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv:1511.06581*, 2015.
- Weiler, M. and Cesa, G. General e(2)-equivariant steerable cnns. In *Advances in Neural Information Processing Systems*, pp. 14334–14345, 2019.
- Weiler, M., Hamprecht, F. A., and Storath, M. Learning steerable filters for rotation equivariant cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 849–858, 2018.
- Zhang, B., Mao, Z., Liu, W., and Liu, J. Geometric reinforcement learning for path planning of uavs. *Journal of Intelligent & Robotic Systems*, 77(2):391–409, 2015.

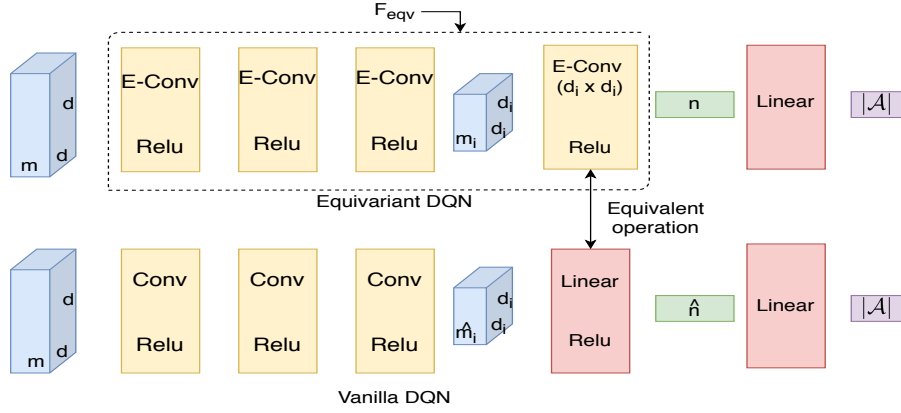


Figure 3 Juxtaposed Network architecture

A. Group representation and restriction in feature fields

We now discuss how one would choose a group (g) and its representation (ρ) to define a feature field. The group’s choice mainly depends on the problem we are tackling and to which kinds of transformation we wish the network to output equivariantly. We have several options for $E(2)$ -equivariant convolution, starting from discrete rotations and reflections (D_N) to continuous rotation and reflection ($O(2)$). Once a group is chosen, we need to choose its representation. The most common ones are trivial, irreducible, regular and quotient representations. The representation chosen determines the dimension c of a feature vector. While a trivial representation implies scalar features with dimension 1 the regular representation uses an N -dimensional feature field, where N denotes the order of the group we are using. Even though a regular representation was shown to perform the best (Weiler & Cesa, 2019), it is computationally infeasible to use it when using higher-order groups. In such a case, we use an irreducible representation, which takes the smallest dimension while leaving the representation of all the group elements unique.

Let us assume that we are working with a generic D_N group with its regular representation. The next thing we need to choose is the number of feature fields for each intermediate layer. Together the chosen representation and number of feature fields contribute to the dimension of the stack $\bigoplus_i f_i$ of intermediate feature fields, which further determines the depth of the convolution kernel we are using between two feature fields. Although increasing the number of feature fields increases the network’s capacity, this comes at the cost of increased computation during a single forward pass.

In an environment where we have a global D_N symmetry, where we want D_N equivariant features and $N > 1$, we can directly choose the group and keep it throughout. But in

most environments where there is usually a global D_1 symmetry and occasionally local D_N symmetry, using the same representation throughout would be futile as this is accompanied by order of N increase in feature field dimension. To alleviate this problem, we start with a higher-order group D_N where $N > 1$ and as we go deeper into our Q network, we restrict it to its subgroups ($\leq D_N$). This makes the network more computationally efficient while still extracting an equivariant feature vector.

B. Network Architecture

The baseline Vanilla DDQN used in this work is similar to the one used in (Mnih et al., 2013), which has an output dimension equal to the number of actions. As shown in Figure 3, our proposed Equivariant DDQN architecture mainly replaces the Vanilla convolutions ($Conv$) and the second last linear layer with equivariant convolutions ($E - Conv$). We call this an equivariant feature extractor. We want to emphasize on the last $E - Conv$ layer and point-out that its operation is similar to the second last linear layer in a Vanilla DDQN. As we use the filter size of the dimension of feature size before that layer, all the information is captured as a weighted sum into a 1-D vector. Although this is the same as the flattening of the feature and then applying a linear layer, using $E - Conv$ renders the output vector equivariant.

The final linear layer is the same for both and maps them to Q -values for each action. Notice, as mentioned in Appendix A, the group representation and the number of feature field will determine the sizes of intermediate features. We aim to make both networks similar with respect to computation time while not comprising the capacity of the Equivariant model. Below we provide the architecture of the Equivariant and Vanilla DDQN for both Snake and Pacman. We denote a basic convolution by: $Conv(filtersize \times filtersize, inchannels, outchannels, stride, padding)$ and equivariant one by: $E - Conv(filtersize \times$

$filtersize, infields, outfield, stride, padding)[Group]$
 The group restriction operation is denoted by: $GrpRes[Group \rightarrow Subgroup]$. In the Vanilla and Equivariant DDQN, we denote the size of the output of the third convolution by $\hat{m}_i \times d_i \times d_i$ and $m_i \times d_i \times d_i$ respectively. Using this, we give the exact architecture of both the networks below.

B.1. Snake

B.1.1. VANILLA DDQN

$Conv(7 \times 7, m, 32, 2, 2) - ReLU$
 $Conv(5 \times 5, 32, 64, 2, 1) - ReLU$
 $Conv(5 \times 5, 64, 64, 1, 1) - ReLU$
 $Linear(\hat{m}_i \times d_i \times d_i, 256) - ReLU$
 $Linear(256, |\mathcal{A}|)$

B.1.2. EQUIVARIANT DDQN

$E - Conv(7 \times 7, m, 8, 2, 2)[D_4] - ReLU$
 $E - Conv(5 \times 5, 8, 12, 2, 1)[D_4] - ReLU$
 $E - Conv(5 \times 5, 12, 12, 1, 1)[D_4] - ReLU$
 $E - Conv(d_i \times d_i, 12, 32, 1, 0)[D_4] - ReLU$
 $Linear(256, |\mathcal{A}|)$

B.2. Pacman

B.2.1. VANILLA DDQN

$Conv(7 \times 7, m, 32, 4, 2) - ReLU$
 $Conv(5 \times 5, 32, 64, 2, 2) - ReLU$
 $Conv(5 \times 5, 64, 64, 2, 1) - ReLU$
 $Linear(\hat{m}_i \times d_i \times d_i, 512) - ReLU$
 $Linear(512, |\mathcal{A}|)$

B.2.2. EQUIVARIANT DDQN

$E - Conv(7 \times 7, m, 8, 2, 2)[D_4] - ReLU$
 $E - Conv(5 \times 5, 8, 16, 2, 1)[D_4] - ReLU$
 $GrpRes[D_4 \rightarrow D_1]$
 $E - Conv(5 \times 5, 16, 64, 1, 1)[D_1] - ReLU$
 $E - Conv(d_i \times d_i, 64, 384, 1, 0)[D_1] - ReLU$
 $Linear(768, |\mathcal{A}|)$

Network Type	Vanilla DQN	Number of Parameters
Vanilla DDQN	Snake	583.46k
Equivariant DDQN	Snake	57.7k
Vanilla DDQN	Pacman	984.36k
Equivariant DDQN	Pacman	649.93k

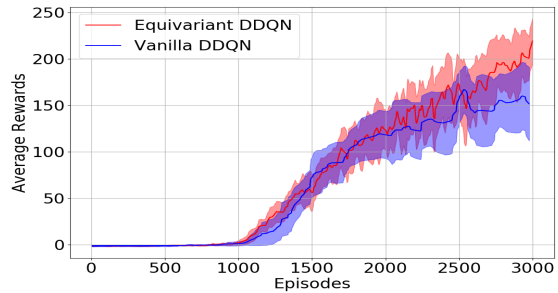
Table 2. This table gives the number of parameters of both the Networks.

Although there is a difference in the number of channels and feature fields, the overall runtime of the DDQN algorithms with both the networks are similar. The forward pass of the

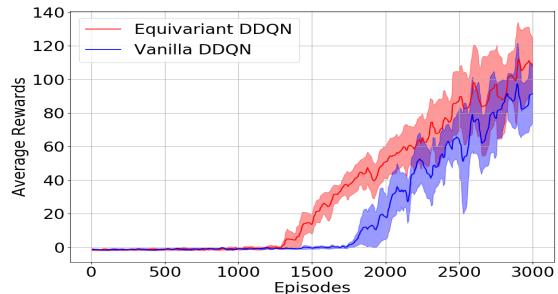
Equivariant network is more computationally expensive as the total dimension of the stack of feature fields in some layers is more than the number of channels in the Vanilla network. But this is partially compensated for during the backpropagation where we are updating fewer parameters in an Equivariant network. Note that, in general, adding feature fields increases the capacity at the cost of computation, but we keep the total cost with respect to the Vanilla model in mind while choosing them. Also, as the Pacman environment is globally symmetric to the D_1 group, we restrict the group once D_4 symmetric lower level features are extracted, which also reduces the dimension of representation and hence the computation cost significantly. It is interesting to note that higher-order symmetry in the environment leads to fewer parameters than the Vanilla DQN.

C. Additional Results

In this section, we provide some additional results of our proposed method applied to DDQN with priority replay and Dueling DDQN in the Snake Environment. We show that our proposed models outperform the Vanilla models in both the cases, which demonstrates that our approach scales to handle different algorithms. Note that we used a lower learning rate for Dueling DQN to stabilize training.



DDQN with priority replay in Snake



Dueling DQN in Snake

Figure 4. Plots of evolution of average rewards with the number of episodes over 10 different seeds and are smoothed with a 1D Gaussian filter with $\sigma=3$ for better visualization.