

463 **A Illustration of Group and Group Representation in CarFlag-2D**

464 **Domain** We consider a small version of CarFlag-2D (see Figure 9) with a grid size of 3x3, where
 465 the agent (red) must navigate to an unknown target cell (green) in a grid world. The agent can always
 466 observe its current location but only observe the target cell when it visits the information cell (blue),
 467 which is also unknown to the agent.

468 **Observation** The observation is a two-channel image size 2x3x3, where the first channel encodes
 469 the agent’s location and the second encodes the target location. The values of the second channel
 470 are non-zero only when the agent is at the information cell (Figure 9).

471 **Actions** Movements in four directions (the location does not change if going out of the world).

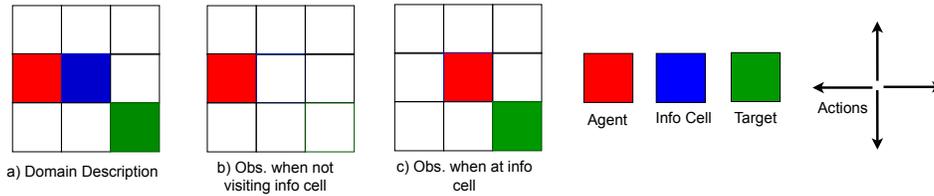


Figure 9: Illustration of the domain, action, and observation.

472 **Domain Symmetry** Consider Scenario 1 and Scenario 2 in Figure 10: Scenario 2 is the rotated
 473 version of Scenario 1 after a 90° counter-clockwise (CCW) rotation. Therefore, an optimal path
 474 (denoted with colored arrows) in Scenario 1 is equally optimal in Scenario 2 if we rotate the path
 475 similarly. The same happens if the rotation angle is 180° or 270°. We can capture the rotational
 symmetry using group $C_4 = \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$.

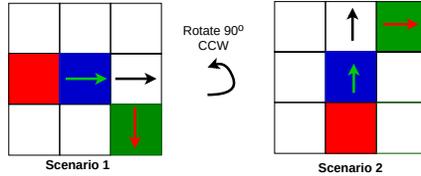


Figure 10: Illustration of domain symmetry. Scenario 2 is the rotated version of Scenario 1 after a 90° counter-clockwise rotation. An optimal path in Scenario 1 can be rotated similarly to become optimal in Scenario 2.

476

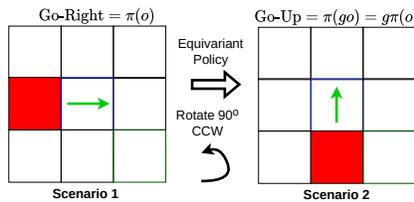


Figure 11: Illustration of the effect of an equivariant policy: the action is automatically rotated when the input observation is rotated.

477 **Equivariant Policy** We want our policy to automatically capture the domain symmetry above by
 478 making it equivariant. In Figure 11, we illustrate the property of an equivariant policy π with g
 479 being a 90° CCW rotation. In Scenario 1, given the first observation o , we assume that π already
 480 knows it should go to the right $\text{Go-Right} = \pi(o)$ towards the information cell. Now, moving to
 481 Scenario 2, when the first observation is the rotated version of o , denoted as go . An equivariant
 482 policy automatically calculates the next action in Scenario 2 as:

$$\pi(go) = g\pi(o) = g(\text{Go-Right}) = \text{Go-Up} \quad (4)$$

483 **Group Representation** From Equation (4), to construct an equivariant policy, we need to define
 484 how a rotation g acts on an observation at the input (i.e., define go) and on an action at the output
 485 (i.e., define $g\pi(o)$). For that purpose, besides defining the group, we need to specify the group rep-
 486 resentation, i.e., defining an *observation group representation* ρ_o and an *action group representation*
 487 ρ_a for π (see below).

488 **Example of Group Acting on Observation and Action** The effect of a 90° CCW rotation g
 489 on the observation via a *trivial* representation $\rho_o = \rho_t$ and the action via a *regular* representation
 490 $\rho_a = \rho_r$ is illustrated in Figure 12. A trivial representation ρ_t rotates the observation (like rotating
 491 the normal image) while keeping the pixel values unchanged (the value in cell 0 is still $(0_0, 0_1)$). In
 492 contrast, a regular representation ρ_r permutes the action distribution output, resulting in a different
 action, i.e., Go-Right \rightarrow Go-Up).

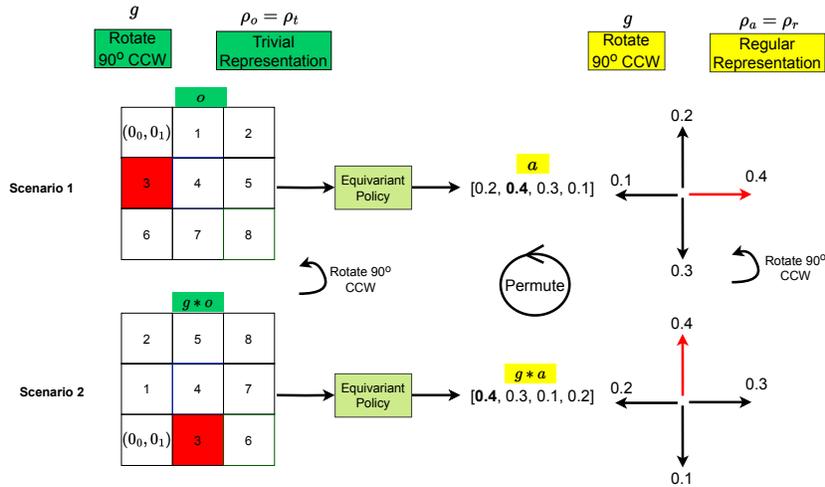


Figure 12: Illustration of the effect of a trivial representation acting on the observation ($\rho_o = \rho_t$) and a regular representation acting on the action ($\rho_a = \rho_r$) with g being a 90° CCW rotation. ρ_o rotates the observation and keeps the pixel values unchanged. ρ_a permutes the action distribution output, resulting in a different action (Go-Right \rightarrow Go-Up).

493

494 **B Proof of Theorem 1**

495 In this section, we introduce the framework of history representation MDP [63] and prove a support-
496 ing lemma before arriving at the proof.

497 **B.1 History Representation MDP**

498 A POMDP can be converted into a *history representation MDP* (HR-MDP) [63] whose state is a
499 sufficient statistic of the POMDP history for control, e.g., the well-known *Belief-MDP* [64] construct
500 is a special case of an HR-MDP based on the belief representation. Useful representations such
501 as the belief might require a known POMDP model; however, we adopt a model-free approach
502 with no such knowledge and therefore use the trivial identity representation whereby the history is
503 represented by itself. This effectively converts the POMDP into an equivalent *History-MDP*, which
504 is defined by the tuple $(\mathcal{H}, \mathcal{A}, \bar{T}, \bar{R})$, where:

$$\bar{T}(h, a, h') = \mathbb{E}_{o|h,a} [\mathbb{I}\{h' = hao\}] \quad \bar{R}(h, a) = \mathbb{E}_{s|h} [R(s, a)] , \quad (5)$$

505 where $\mathbb{I}\{\cdot\}$ is the indicator function, and

$$\Pr(o | h, a) = \mathbb{E}_{s|h} \left[\sum_{s'} T(s, a, s') O(a, s', o) \right] , \quad (6)$$

$$\Pr(s' | h') \propto \mathbb{E}_{s|h} [T(s, a, s') O(a, s', o)] . \quad (7)$$

506 **B.2 Supporting Lemma**

507 **Lemma 1.** *The belief function of a group-invariant POMDP (as defined by Definition 1) is group-*
508 *invariant,*

$$\Pr(gs | gh) = \Pr(s | h) . \quad (8)$$

509 *Proof By Induction.*

510 **Base Case.** We first prove that the belief after the first observation is invariant. We note here that
511 the observation function for the first timestep takes the form $O(s, o)$, with no preceding action.

$$\Pr(gs_0 | go_0) \propto b_0(gs_0) O(gs_0, go_0) = b_0(s_0) O(s_0, o_0) \propto \Pr(s_0 | o_0) . \quad (9)$$

512 Since $\Pr(gs_0 | go_0)$ and $\Pr(s_0 | o_0)$ are both proportional to the same quantity, and they are both
513 normalized to be distributions over states, then they are themselves equal.

514 **Inductive Step.** We then prove that if $\Pr(s_t | h_t)$ is invariant, then $\Pr(s_{t+1} | h_{t+1})$ is also invariant.
515 Per Equation (7),

$$\begin{aligned} \Pr(gs_{t+1} | gh_{t+1}) &\propto \Pr(gs_t | gh_t) T(gs_t, ga_t, gs_{t+1}) O(ga_t, gs_{t+1}, go_{t+1}) \\ &= \Pr(s_t | h_t) T(s_t, a_t, s_{t+1}) O(a_t, s_{t+1}, o_{t+1}) \propto \Pr(s_{t+1} | h_{t+1}) . \end{aligned} \quad (10)$$

516 Since $\Pr(gs_{t+1} | gh_{t+1})$ and $\Pr(s_{t+1} | h_{t+1})$ are both proportional to the same quantity, and they
517 are both normalized to be distributions over states, then they are themselves equal. By induction,
518 given the base case and the inductive step, the belief function $\Pr(s_t | h_t)$ is invariant for any t . \square

519 **B.3 Proof**

520 *Proof.* We begin by constructing the History-MDP associated with a group-invariant POMDP, and
521 showing that it is itself a group-invariant MDP. The transition and reward functions of the History-
522 MDP are shown in Equation (5) and satisfy the group invariance properties.

523 For this proof, it is simpler to express the history transition function as $\bar{T}(h, a, h') = \Pr(o | h, a)$,
524 where o is the observation (if any exists) s.t. $h' = hao$. If no such observation exists, then
525 $\bar{T}(h, a, h') = 0$ is trivially invariant. If it does exist, then it is necessarily the last observation
526 of h' ,

$$\bar{T}(gh, ga, gh') = \Pr(go | gh, ga) = \sum_{s,s'} \Pr(s | gh) \Pr(s' | s, ga) \Pr(go | ga, s')$$

527 since g permutes the elements of \mathcal{S} , we can re-index using $s = g\bar{s}$ and $s' = g\bar{s}'$,

$$\begin{aligned}
&= \sum_{\bar{s}, \bar{s}'} \Pr(g\bar{s} \mid gh) T(g\bar{s}' \mid g\bar{s}, ga) O(go \mid ga, g\bar{s}') \\
&= \sum_{s, s'} \Pr(s \mid h) T(s' \mid s, a) O(o \mid a, s') = \bar{T}(h, a, h'). \tag{11}
\end{aligned}$$

528 By using $s = g\bar{s}$, we proceed similarly for history rewards,

$$\begin{aligned}
\bar{R}(gh, ga) &= \sum_s \Pr(s \mid gh) R(s, ga) = \sum_{\bar{s}} \Pr(g\bar{s} \mid gh) R(g\bar{s}, ga) \\
&= \sum_s \Pr(s \mid h) R(s, a) = \bar{R}(h, a). \tag{12}
\end{aligned}$$

529 Therefore, $\bar{T}(h, a, h')$ and $\bar{R}(h, a)$ are invariant, and History-MDPs are group-invariant MDPs. By
530 the theory developed in [1], this implies that the optimal Q-value function $Q^*(h, a)$ is invariant and
531 that there exists at least one equivariant deterministic optimal policy $\pi^*(h)$. Moreover,

$$\begin{aligned}
V^*(gh) &= Q^*(gh, \pi^*(gh)) = Q^*(gh, g\pi^*(h)) \\
&= Q^*(h, \pi^*(h)) = V^*(h), \tag{13}
\end{aligned}$$

532 this ends our proof by showing that $V^*(h)$ is invariant. □

533 C Environment Details

534 C.1 Grid-world Domains

535 C.1.1 CarFlag-1D

- 536 • Action: Go-Left or Go-Right
- 537 • Observation (Discrete): The position of the car, the side of the green flag (-1 or 1 if the car
- 538 is at the blue flag, and 0 otherwise)
- 539 • Reward: step reward: -0.01, reaching the green flag: 1.0, and reaching the red flag: -1.0
- 540 • Episode Initialization: The car is randomized such that it is not at the information location
- 541 (blue flag). The goal (green flag) is always either at the leftmost or rightmost end. The red
- 542 flag is on the opposite end
- 543 • Episode Termination: Reaching either flags or an episode lasts more than 50 timesteps
- 544 • World size: The distance between the red and the green flag is 50

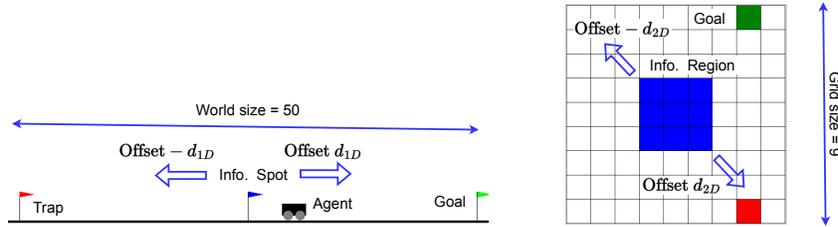


Figure 13: CarFlag-1D and CarFlag-2D domains. The information regions are not visible to the agent. These domains become asymmetric when the offsets from the information region to the world center, i.e., d_{1D} and d_{2D} , are non-zero.

545 C.1.2 CarFlag-2D

- 546 • Action: Right/Left/Up/Down
- 547 • Observation: The observation is encoded as an $N \times N \times 2$ image, where N is the grid
- 548 size, the first channel encodes the position of the car, and the second channel encodes the
- 549 position of the green cell. The second channel is only informative when the agent is inside
- 550 the information region (blue)
- 551 • Reward: Reaching the green cell: 1.0, otherwise 0.0
- 552 • Episode Initialization: The agent and the goal cell are randomized such that the minimum
- 553 distance between them is at least two steps. Moreover, both the agent and the goal are not
- 554 initialized inside the information region (blue)
- 555 • Episode Termination: Reached the goal or an episode lasts more than 50 timesteps

556 C.2 Robot Manipulation Domains

557 For these domains, an episode is terminated when it lasts over 50 timesteps or the task is achieved.
 558 Because all robot domains share the same observation and action, we only describe them below.

559 **Action.** An action $a = (\delta_w, \delta_x, \delta_y, \delta_z, \delta_r)$, where $\delta_w \in [0, 1]$ is the absolute openness of the gripper
 560 (0: fully open, 1: fully closed), $\delta_{x,y,z} \in [-0.05, 0.05]$ are the displacements of the gripper in the X,
 561 Y, and Z axis, and $\delta_r \in [-\pi/8, \pi/8]$ is the angular rotation around the Z axis (see Figure 14a)

562 **Observation.** An observation is a top-down depth image taken from a camera located at the end-
 563 effector. Specifically, an observation $o = (I, k)$, where $I \in \mathbb{R}^{84 \times 84}$ is the depth image and $k \in$
 564 $\{1, 0\}$ indicates the current holding status of the gripper. I and k are combined to create a unified
 565 depth observation $o \in \mathbb{R}^{2 \times 84 \times 84}$. Moreover, two fingers of the gripper are also projected on I (black
 566 squares in Figure 14b)

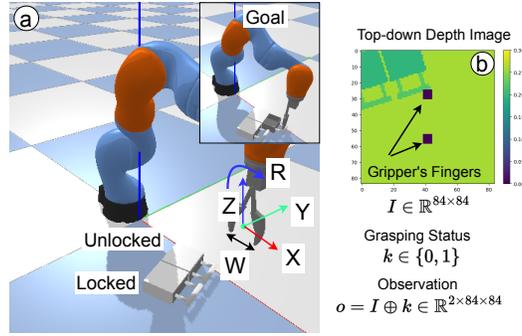


Figure 14: Visual description of Drawer-Opening.

567 **Partial Observability.** These domains characterize the natural partial observability when certain
 568 physical properties of objects, e.g., whether a drawer in Figure 14a is unlocked or not, are often
 569 unobservable using pixel observations alone

570 C.2.1 Block-Picking

- 571 • Reward: A reward of 1.0 only when the movable block is picked and brought higher than
 572 8cm
- 573 • Episode Initialization: The poses of the two blocks are randomized. The arm is initialized
 574 at a fixed pose
- 575 • Expert Generation: An expert (a planner with access to all object poses) randomly chooses
 576 one block to pick. If the expert picks the movable block, it will bring the block up to
 577 achieve the task. Otherwise, the expert keeps trying for several timesteps before switching
 578 to pick the movable block to achieve the task

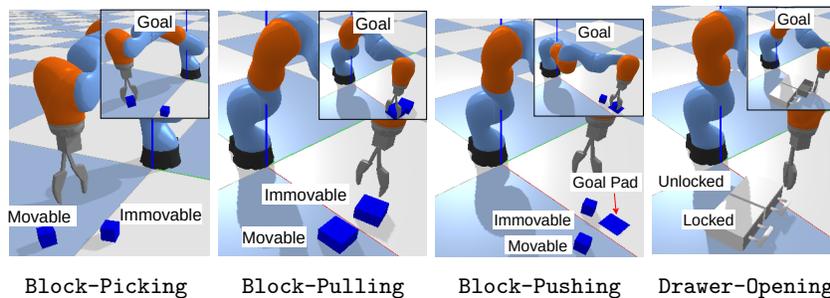


Figure 15: Robot manipulation domains.

579 C.2.2 Block-Pulling

- 580 • Reward: A reward of 1.0 only when two blocks are in contact
- 581 • Expert Generation: An expert randomly chooses one block to pull towards the other block.
 582 If the block is pullable, then it will be pulled towards the other block to achieve the task.
 583 Otherwise, the expert keeps trying for a while before pulling the other block

584 C.2.3 Block-Pushing

- 585 • Reward: A reward of 1.0 only when the pushable block is within 5cm from the center of
 586 the goal pad. The agent additionally receives a penalty of 0.1 per timestep if it changes the
 587 height of the movable block by 5mm to prevent picking the block instead of pushing it
- 588 • Episode Initialization: The poses of the two blocks and the goal pad are randomly initial-
 589 ized

590 • Expert Generation: An expert randomly chose one block to push towards the goal pad. If
591 the block is pushable, it will then continue pushing until reaching the goal pad. Otherwise,
592 the expert keeps trying for several timesteps before doing the same thing with the other
593 (pushable) block

594 **C.2.4** Drawer-Opening

595 • Reward: A reward of 1.0 only when the unlocked drawer is opened more than 5cm
596 • Episode Initialization: Two drawers are randomly placed next to each other with the same
597 heading angle
598 • Expert Generation: An expert randomly chooses one drawer to open. If it chooses the
599 unlocked drawer, it will then open the drawer to achieve the task. Otherwise, the expert
600 keeps opening the unlocked drawer several timesteps before opening the other drawer

601 D Implementation Details

602 D.1 Network Structure of Equivariant Recurrent A2C (Equi-RA2C)

603 Figure 16 shows the specific architecture of Equi-RA2C used in CarFlag domains. Because the
 604 actions can be inferred from the observations in these domains, we do not include the feature extrac-
 605 tor for the previous actions. We also omit the skip-connections. The input representation is some
 representation of the observation ρ_o , depending on the domains (see below).

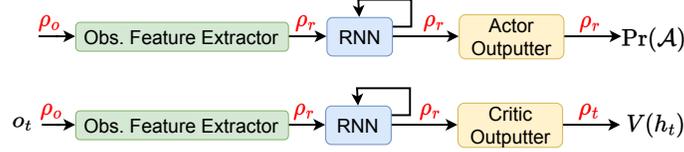


Figure 16: The architecture of Equi-RA2C used in CarFlag-1D and CarFlag-2D.

606

607 Figure 17 shows the details of Equi-RA2C used in CarFlag-1D for the `flip2d0nR2` group in the
 608 `escnn`¹ [45, 65] library. Notice that the input x_t for the LSTM cell using the *irreducible* repre-
 609 sentation of the `flip2d0nR2` group denoted as ρ_{irr} . For CarFlag-1D, using this representation in the
 610 input would negate the signs of every component in x_t , i.e., flipping the positions of the car,
 611 the sides of the green flag, and the previous actions in the history. Because the observation in this domain is
 612 feature-based, we remove the observation feature extractor and directly feed the observation to the
 equivariant LSTM.

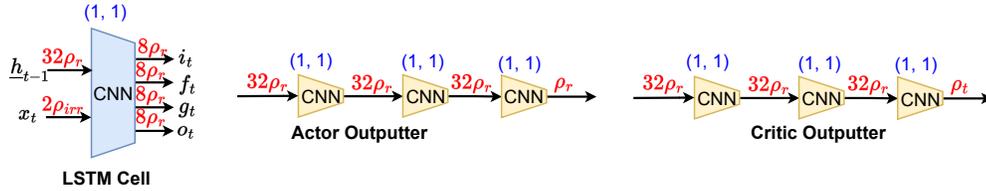


Figure 17: Details of Equi-RA2C used in CarFlag-1D for the `flip2d0nR2` group. Numbers inside
 brackets (blue - on top) denote the value of kernel sizes and strides used for the CNN modules on
 the bottom. The numbers next to the representations, e.g., $32\rho_r$, denote the number of feature fields.

613

Figure 18 shows the details of Equi-RA2C used in CarFlag-2D for the C_4 group.

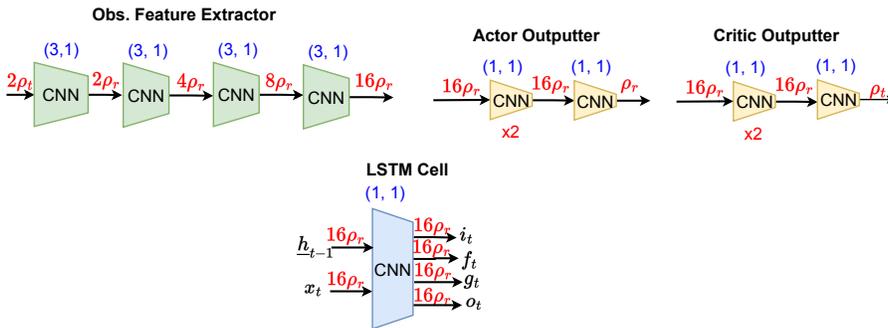


Figure 18: The details of Equi-RA2C used in CarFlag-2D for the C_4 group.

614

615 D.2 Network Structure of Equivariant Recurrent SAC (Equi-RSAC) with C_4 Group

616 Figure 19 shows the details of Equi-RSAC used in the robot manipulation domains with the C_4
 617 group. The input representation is *mixed* for the action feature extractor because the action in-

¹<https://github.com/QUVA-Lab/escnn>

618 put has components that transform differently under a rotation. Specifically, given an action
 619 $a = (\delta_w, \delta_x, \delta_y, \delta_z, \delta_r)$, the trivial representation ρ_t is chosen for the $\delta_w, \delta_z, \delta_r$ components (which
 620 should be unchanged under the rotation). In contrast, the standard representation ρ_s is chosen for
 621 the lateral components (δ_x, δ_y) , which should rotate. For the same reason, for the actor outputter,
 622 ρ_μ is mixed, i.e., the trivial representations ρ_t are used for the w, z, r components, and ρ_s is used for
 the x, y components.

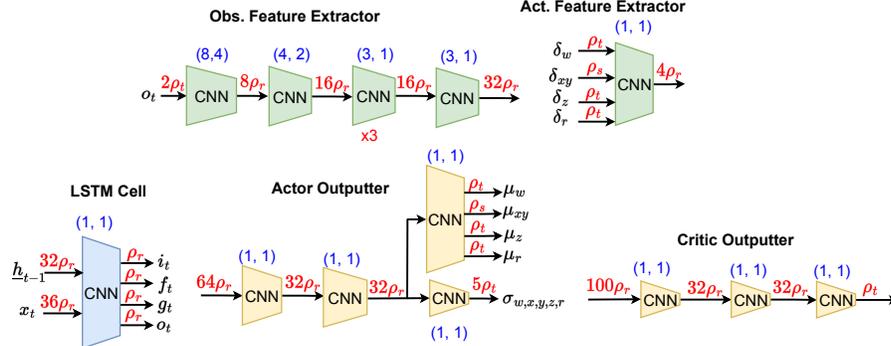


Figure 19: Details of Equi-RSAC with the robot manipulation domains and the C_4 group.

623

624 D.3 Implementation Using The ESCNN Library

625 Given the definition of each equivariant component above, we can easily implement it with `escnn`.
 626 For instance, the following PyTorch [66] code defines the observation feature extractor in Figure 18a
 627 with ReLU as a non-linearity component:

```

import escnn.nn as enn

# Define group C4
s = escnn.gspaces.rot2d0nR2(4)

# Define in/out representations
repr_i = enn.FieldType(s, 2*[s.trivial_repr])
repr_m0 = enn.FieldType(s, 2*[s.regular_repr])
repr_m1 = enn.FieldType(s, 4*[s.regular_repr])
repr_m2 = enn.FieldType(s, 8*[s.regular_repr])
repr_o = enn.FieldType(s, 16*[s.regular_repr])

obs_feature_extractor = enn.SequentialModule(
    enn.R2Conv(repr_i, repr_m0, 3, 1),
    enn.ReLU(repr_m0),
    enn.R2Conv(repr_m0, repr_m1, 3, 1),
    enn.ReLU(repr_m1),
    enn.R2Conv(repr_m1, repr_m2, 3, 1),
    enn.ReLU(repr_m2),
    enn.R2Conv(repr_m2, repr_o, 3, 1),
    enn.ReLU(repr_o),
)

```

628 Implementing the mixed representation is also straightforward by summing different field
 629 types. In order to create the actor and the critic, we simply chain components by using the
 630 `SequentialModule` as in native PyTorch.

631 D.4 Training Details

632 We implement using PyTorch. The batch size for all agents is 32 (episodes). The replay buffer
 633 has a capacity of 100,000 transitions. We use the Adam optimizer [67] with a learning rate of $3e$ -
 634 4 for actors and critics and $1e$ -3 for optimizing α for SAC-based agents. The target entropy \bar{H}

635 for SAC-based agents is $-\dim(\mathcal{A})$ followed the common practice, and α is initialized at 0.1. After
636 prepopulating the replay buffer with 80 expert episodes, the buffer is filled with 20 episodes with
637 random actions. We use the same 1:1 environment/gradient step ratio for all agents.

638 **D.5 Implementing Equivariant LSTM**

639 We implement the equivariant LSTM [47] based on a public code of ConvLSTM [48] at <https://github.com/Hzzone/Precipitation-Nowcasting> as the authors did not release the official
640 code.
641

642 **E Baseline Details**

643 **RA2C [51]** We modified the code at [https://github.com/ikostrikov/](https://github.com/ikostrikov/pytorch-a2c-ppo-acktr-gail)
644 [pytorch-a2c-ppo-acktr-gail](https://github.com/ikostrikov/pytorch-a2c-ppo-acktr-gail). We used 16 environments in parallel and used recurrent
645 policies. Other hyper-parameters are kept at default.

646 **DPFRL [14]** We used the authors' code at <https://github.com/Yusufma03/DPFRL>. We used
647 30 particles, MGF particle aggregation type, and the hidden dimension is 128.

648 **RAD [54]** We collected depth images of size 90x90 to perform random cropping to reduce the size
649 to 84x84. We perform the same type of random cropping for every depth image within an episode.

650 **DrQ [55]** We used random shift of ± 4 pixels as suggested by the original work. The same type of
651 shifting is used for every depth image within a sequence. We also followed the authors' suggestions
652 when using the numbers of augmentations for calculating the Q-targets and the Q-values are $K = 2$
653 and $M = 2$, respectively.

654 **SLAC [56]** We used a Pytorch implementation at [https://github.com/toshikwa/slac.](https://github.com/toshikwa/slac.pytorch)
655 [pytorch](https://github.com/toshikwa/slac.pytorch), which has been benchmarked against the performance reported in the original paper. We
656 pre-train the latent variable model for 2k steps before iterating between data collection, model up-
657 date, and evaluation. We also pre-fill the replay buffer with the same number of expert and random
658 episodes before training and use four extra augmented episodes for each episode during training
659 to ensure a fair comparison. The sequence length is extended from 8 (originally) to 50 (maximum
660 episode length). We varied the sequence length for better performance, but the performance did not
661 improve much. For any episode shorter than 50 steps, we zero-pad dummy transitions *in front*.

662 **DreamerV2 [52]** We used the official code at <https://github.com/danijar/dreamerv2>. For
663 CarFlag domains, we mainly keep the default hyper-parameters (suggested by the authors). In
664 CarFlag-2D, the observation image is extended to have the size of $64 \times 64 \times 3$ by zero-padding
665 around the original image and is added with a dummy channel (all zero).

666 **DreamerV3 [53]** We used the official code at <https://github.com/danijar/dreamerv3> and
667 performed similar steps like in the case of DreamerV2. We used the *small* world models with about
668 18M trainable parameters (predefined in the repo's configuration file) for our CarFlag domains.

669 F Visualization of SLAC Reconstructed Images

670 Figure 20 shows the comparison between the depth images produced by the trained latent model of
671 SLAC [56] (top row) and the ground-truth ones (bottom row) in Block-Pulling after 40k training
672 steps. It can be seen that small squares representing the gripper have been reconstructed quite well,
673 but the model fails to reconstruct the two blocks representing the gripper's position in the scene.

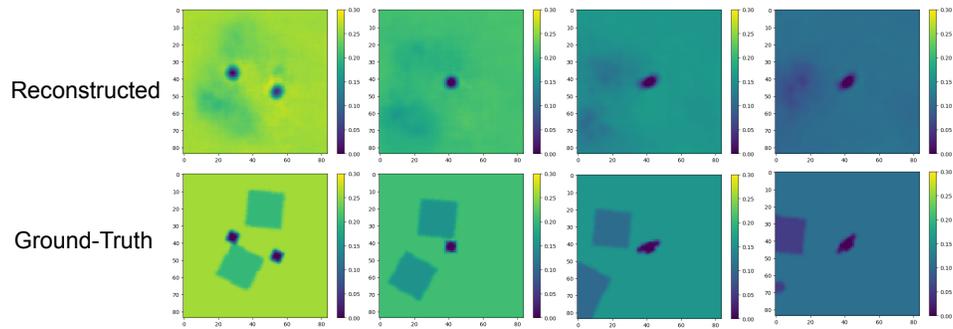


Figure 20: Images reconstructed by the latent model of SLAC [56] in Block-Pulling: reconstructed (top row), ground-truth (bottom row).

674 **G Visualization of Data Augmentations**

675 We show visualizations of different ways for augmenting the observations with a training se-
676 quence in Drawer-Opening: random rotation (Figure 21), random crop (Figure 22), and random
677 shift (Figure 23). Note that the same operation (rotation/crop/shift) is applied similarly to every
678 observation in an episode. For each training episode, we perform this augmentation four times to
generate four auxiliary episodes.

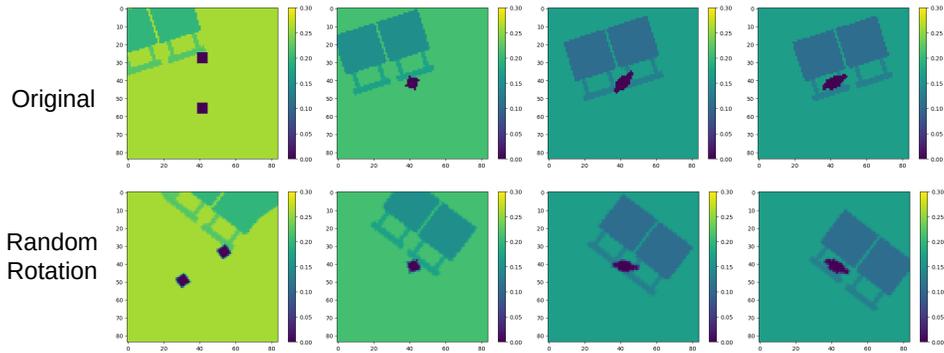


Figure 21: Visualization of randomly rotated augmentations in Drawer-Opening: original obser-
vations (top row), randomly rotated observations (bottom row).

679

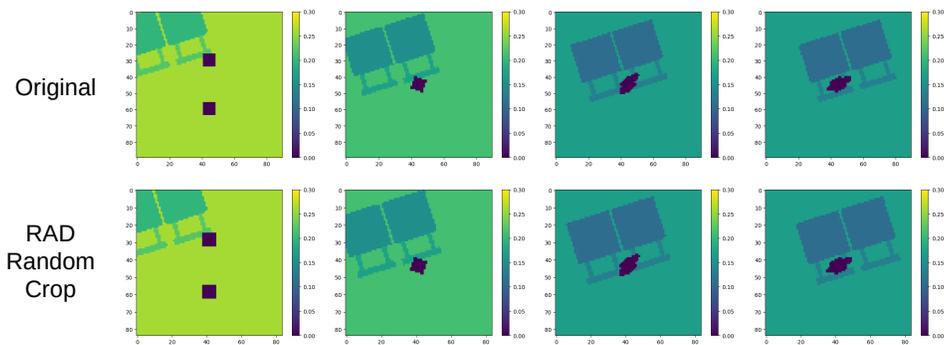


Figure 22: Visualization of randomly cropped augmentation for RAD [54] in Drawer-Opening:
original observations (top row), randomly cropped observations (bottom row).

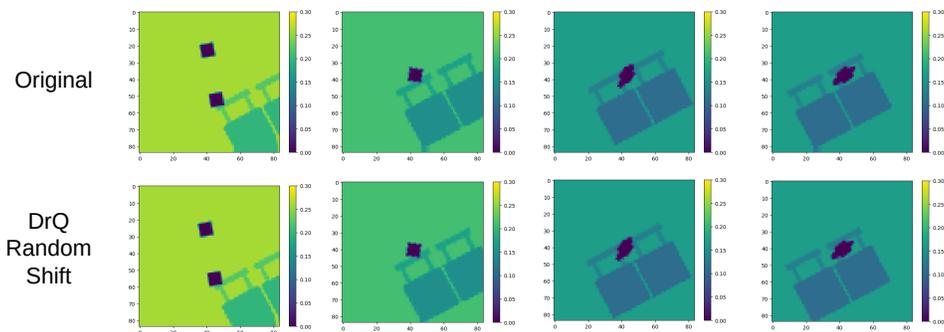


Figure 23: Visualization of randomly shifted augmentation for DrQ [55] in Drawer-Opening: ori-
ginal observations (top row), randomly shifted observations (bottom row).

680 H Ablation Studies

681 H.1 Equivariant Actor or Critic Only

682 In Figure 24, we additionally show the learning performance when only either actor or critic is
683 equivariant in Block-Pushing and Drawer-Opening. From the figure, having an equivariant critic
684 (purple) is more beneficial than having an equivariant actor (blue). However, having both being
685 equivariant (green) yields the best performance.

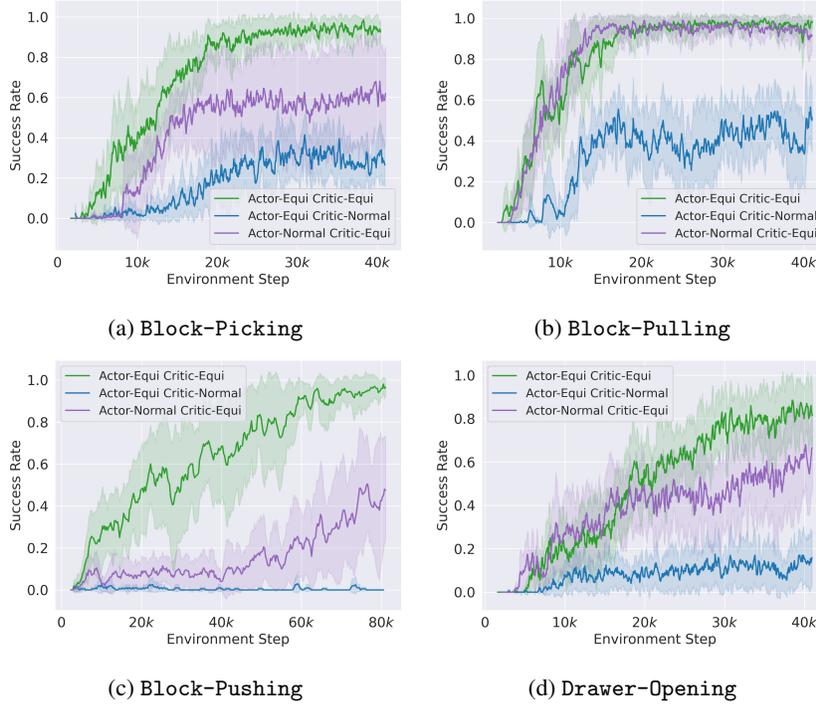


Figure 24: Comparing the effect of only using equivariant actor or critic.

686 **H.2 Different Symmetry Groups**

687 Figure 25 shows the performance when the C_4 and C_8 symmetry groups in the robot manipulation
688 domains. Using C_4 is much better than using C_8 in Block-Pushing, but the two groups perform
689 similarly in the remaining domains. Furthermore, it is possible to use other group symmetries which
690 extend C_n with reflection, such as the dihedral groups D_4 or D_8 .

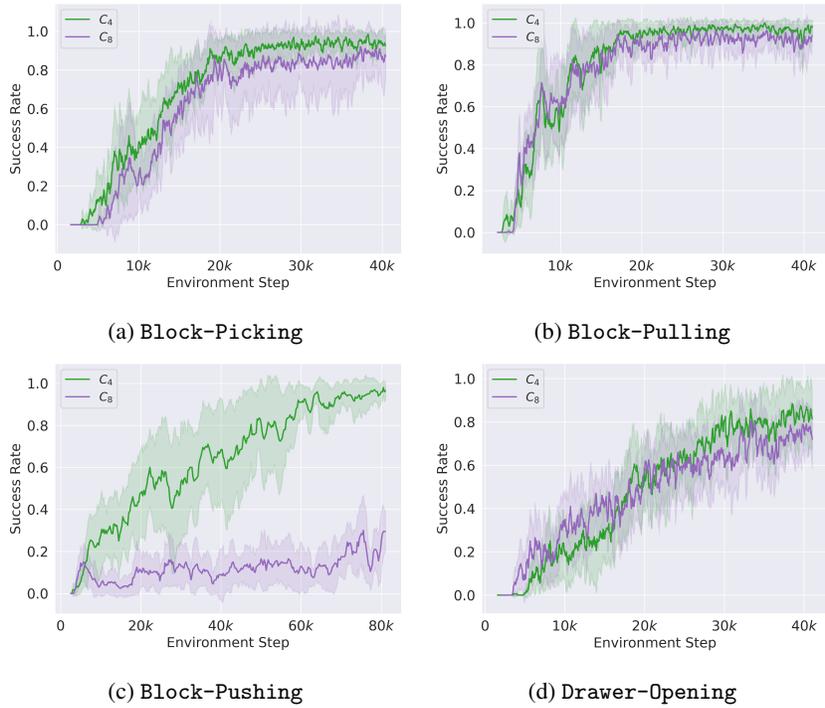


Figure 25: Comparing the effect of using symmetry groups C_4 and C_8 .

691 **H.3 Randomly Initialized Cell and Hidden States of Equivariant LSTM**

692 Figure 26 shows the performance when the equivariant LSTM is initialized with random instead
693 of zero cell and hidden states. Random initialization results in a worse performance because the
694 equivariance of the actor and the critic is broken. However, our method is generally robust to this
695 change when the performance is still better than the baselines.

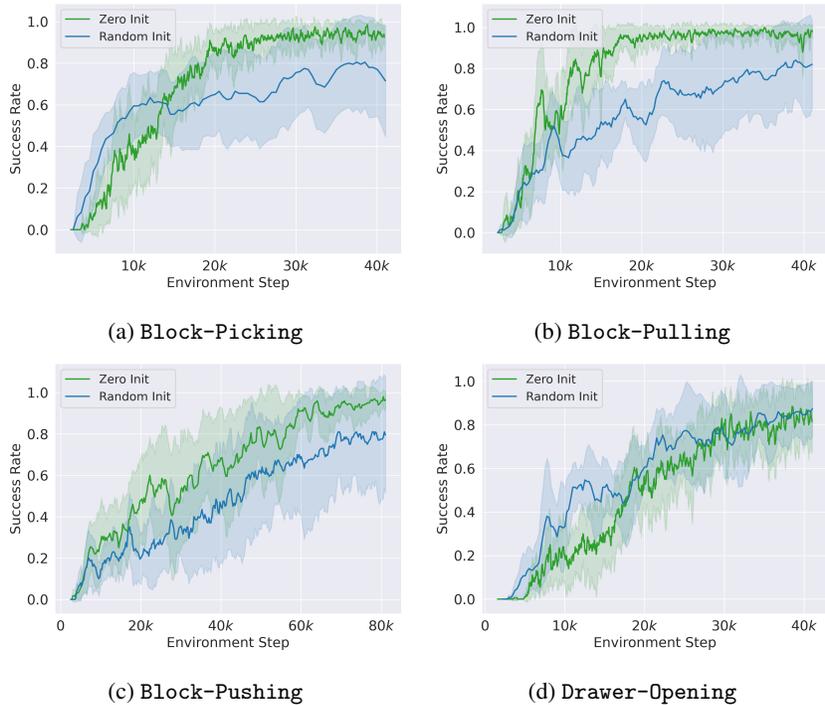


Figure 26: Comparing the performance when initializing the cell and hidden states of the equivariant LSTM with zero and random values. Random initialization results in a worse performance because the actor and the critic’s equivariance is broken.

696 **I Additional Experimental Results**

697 **I.1 Performance in Asymmetric CarFlag Domains**

698 Figure 27 shows the evaluation success rates in asymmetric variants of CarFlag domains with different offsets.

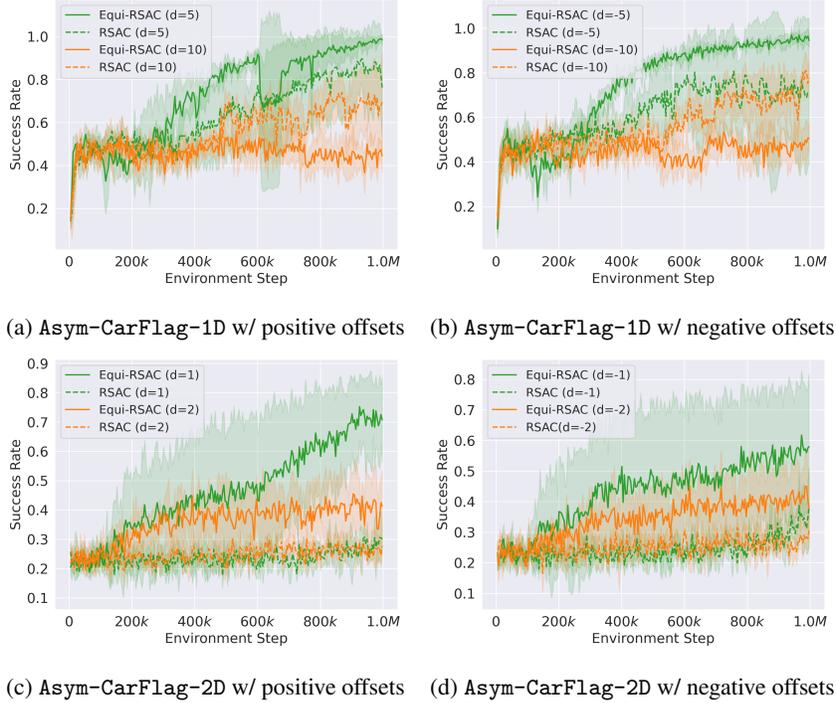


Figure 27: Learning performance with asymmetric version of CarFlag domains.

699

700 **I.2 Performance in Variants of CarFlag Domains**

701 Figure 28 show the evaluation success rates in different variants of CarFlag domains with a different world size and grid size. Our equivariant agent still outperforms other baselines.

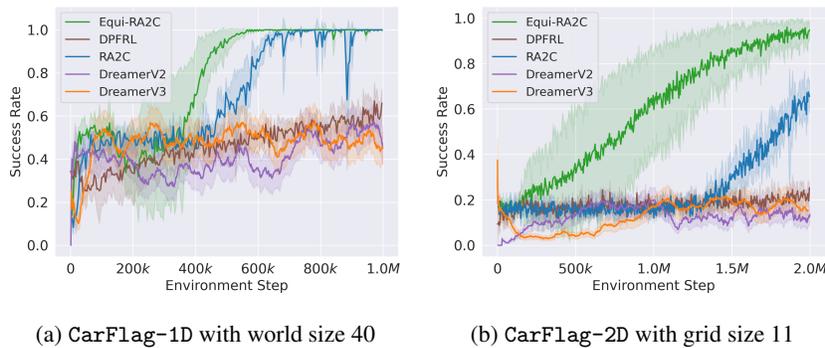


Figure 28: Learning performance in CarFlag domains with different sizes.

702

703 **I.3 Effect of Rotational Augmentation**

704 Figure 29 shows that including rotational augmented episodes significantly improves the learning
705 performance of equivariant agents. These rotational augmented episodes possibly help equivariant
agents distinguish different discrete rotations within a group, thus boosting performance.

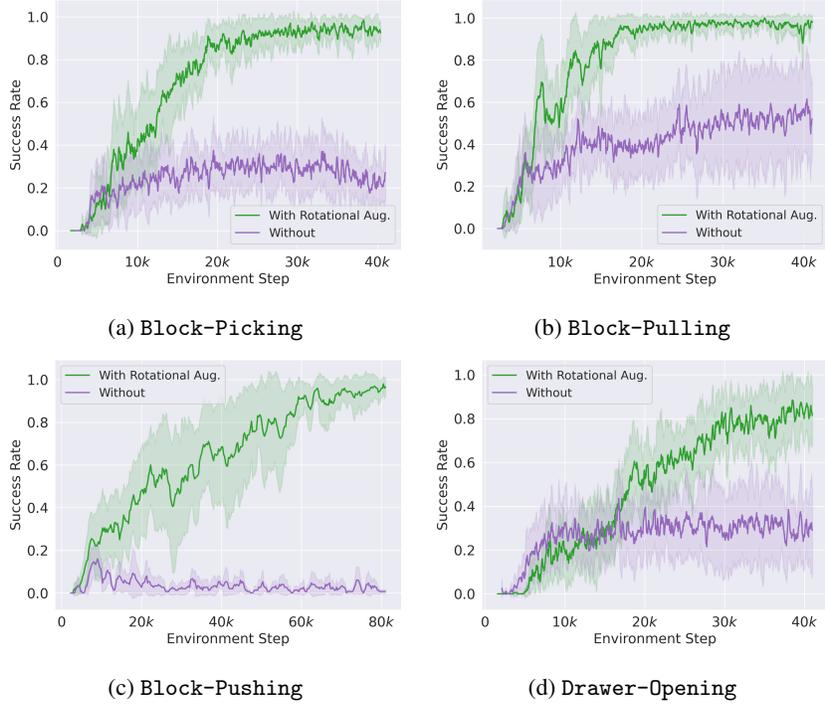


Figure 29: Comparing the performance of our equivariant agents when using/not using rotational augmentation episodes.

706

707 **I.4 Effect of Number of Demonstration Episodes**

708 Figure 30 shows that the performance improves when using more demonstrations in all domains, as expected.

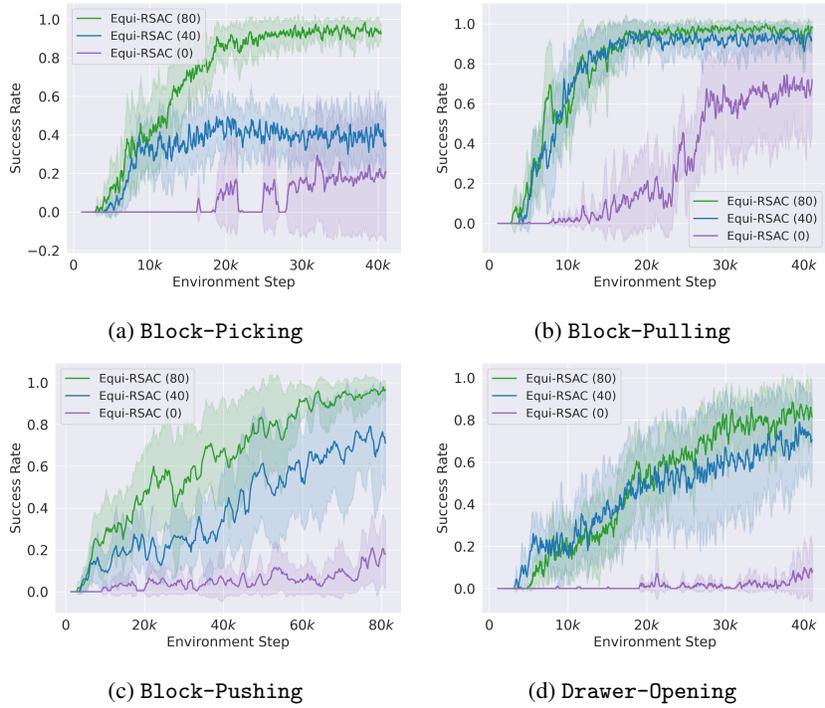


Figure 30: Using different numbers of demonstration episodes.

709