
EDGI: Equivariant Diffusion for Planning with Embodied Agents

Supplementary Material

Anonymous Author(s)

Affiliation

Address

email

1 A Architecture details

2 On a high level, EDGI follows Diffuser [1]. In the following, we will describe the key difference: our
3 $SE(3) \times \mathbb{Z} \times S_n$ -equivariant architecture for the diffusion model.

4 **Overall architecture.** We illustrate the architecture in Fig. 2 in the main paper. After converting the
5 input data in our internal representation (see Sec. 3.2 in the main paper), the data is processed with
6 an equivariant U -net with four levels. At each level, we process the hidden state with two residual
7 standard blocks, before downsampling (in the downward pass) or upsampling (in the upward pass).

8 **Residual standard block.** The main processing unit of our architecture processes the current hidden
9 state with an equivariant block consisting of a temporal layer, an object layer, a normalization layer,
10 and a geometric layer. In parallel, the context information (an embedding of diffusion time and a
11 conditioning mask) is processed with a context block. The hidden state is added to the output of the
12 context block and processes with another equivariant block. Finally, we process the data with a linear
13 attention layer over time. This whole pipeline consists of an equivariant block, a context block, and
14 another equivariant block is residual (the inputs are added to the outputs).

15 **Temporal layers.** Temporal layers consist of one-dimensional convolutions without bias along the
16 time dimension. We use a kernel size of 5.

17 **Normalization layers.** We use a simple equivariant normalization layer that for each batch element
18 rescales the entire tensor w_{toc} to unit variance. This is essentially an equivariant version of LayerNorm.
19 The difference is that our normalization layer does not shift the inputs to zero means, as that would
20 break equivariance with respect to $SO(3)$.

21 **Geometric layers.** In the geometric layers, the input state is split into scalar and vector components.
22 The vector components are linearly transformed to reduce the number of channels to 16. We
23 then construct all $SO(3)$ invariants from these 16 vectors by taking pairwise inner products and
24 concatenating them with the scalar inputs. This set of scalars is processed with two MLPs, each
25 consisting of two hidden layers and ReLU nonlinearities. The MLPs output the scalar outputs and
26 coefficients for a linear map between the vector inputs and the vector outputs, respectively. Finally,
27 there is a residual connection that adds the scalar and vector inputs to the outputs.

28 **Linear attention over time.** To match the architecture used by Janner et al. [1] as closely as possible,
29 we follow their choice of adding another residual linear attention over time at the end of each level in
30 the U -net. We make the linear attention mechanism equivariant by computing the attention weights as

31 **Context blocks.** The embeddings of diffusion time and conditioning information are processed with
32 a Mish nonlinearity and a linear layer, like in Janner et al. [1]. Finally, we embed them in our internal
33 representation by zero-padding the resulting tensor.

34 **Upsampling and downsampling.** During the downsampling path, there is a final temporal layer
35 that implements temporal downsampling and increases the number of channels by a factor of two.
36 Conversely, during the upsampling path, we use a temporal layer for temporal upsampling and a
37 reduction of the number of channels.

38 B Navigation experiments

39 We introduce a new navigation environment. The scene consists of a spherical agent navigating a
40 plane populated with a goal state and $n = 10$ spherical obstacles. At the beginning of every episode,
41 the agent position, agent velocity, obstacle positions, and goal position are initialized randomly (in a
42 rotation-invariant way). We simulate the environment dynamics with PyBullet [2].

43 **Offline dataset.** To obtain expert trajectories, we train a TD3 [3] agent in the implementation by
44 Raffin et al. [4] for 10^7 steps with default hyperparameters on this environment. We generate 10^5
45 trajectories for our offline dataset.

46 **State.** The state contains the agent position, agent velocity, goal position, and obstacle positions.

47 **Actions.** The action space is two-dimensional and specifies a force acting on the agent.

48 **Rewards.** At each time step, the agent receives a reward equal to the negative Euclidean distance
49 to the goal state. In addition, a penalty of -0.1 is added to the reward if the agent touches any of
50 the obstacles. Finally, there is an additional control cost equal to -10^3 times the force acting on the
51 agent. We affinely normalize the rewards such that a normalized reward of 0 corresponds to that
52 achieved by a random policy and a normalized reward of 100 corresponds to the expert policy.

53 C Kuka experiments

54 We use the object manipulation environments and tasks from Janner et al. [1], please see that work
55 for details on the environment. In our experiments, we consider three tasks: unconditional stacking,
56 conditional stacking, and block rearrangement. For a fair comparison, we re-implement the Diffuser
57 algorithm while making bug fixes in the codebase of Janner et al. [1], which mainly included properly
58 resetting the environment.

59 **State.** We experiment with two parameterizations of the Kuka environment state. For the Diffuser
60 baseline, we use the original 39-dimensional parameterization from Janner et al. [1].

61 For our EDGI, we need to parameterize the system in terms of $SE(3) \times \mathbb{Z} \times S_n$ representations. We,
62 therefore, describe the robot and block orientations with $SO(3)$ vectors as follows. Originally, the
63 robot state is specified through a collection of joint angles. One of these encodes the rotation of the
64 base along the vertical z -axis. We choose to represent this angle as a ρ_1 vector in the xy -plane. In
65 addition, we add the gravity direction (the z -axis itself) as another ρ_1 vector, which is also the normal
66 direction of the table on which the objects rest. Combined, these vectors define the pose of the base of
67 the robot arm. Rotating gravity direction, and the robot and object pose by $SO(3)$ can be interpreted
68 as a passive coordinate transformation, or as an active rotation of the entire scene, including gravity.
69 As the laws of physics are invariant to this transformation, this is a valid symmetry of the problem.

70 The n objects can be translated and rotated. Their pose is thus given by a translation $t \in \mathbb{R}^3$ and
71 rotation in $r \in SO(3)$ relative to a reference pose. The translation transforms by a global rotation
72 $g \in SO(3)$ as a vector via representation ρ_1 . The rotational pose transforms by left multiplication
73 $r \mapsto gr$. The $SO(3)$ pose is not a Euclidean space, but a non-trivial manifold. Even though diffusion
74 on manifolds is possible [5, 6], we simplify the problem by embedding the pose in a Euclidean space.
75 This is done by picking the first two columns of the pose rotation matrix $r \in SO(3)$. These columns
76 each transform again as a vector with representation ρ_1 . This forms an equivariant embedding
77 $\iota : SO(3) \hookrightarrow \mathbb{R}^{2 \times 3}$, whose image is two orthogonal 3-vectors of unit norm. Via the Gram-Schmidt
78 procedure, we can define an equivariant map $\pi : \mathbb{R}^{2 \times 3} \rightarrow SO(3)$ (defined almost everywhere), that is
79 a left inverse to the embedding: $\pi \circ \iota = \text{id}_{SO(3)}$. Combining with the translation, the roto-translational
80 pose of each object is thus embedded as three ρ_1 vectors.

81 We also tested the performance of the baseline Diffuser method on this reparameterization of the state
82 but found worse results.

83 **Hyperparameters.** We also follow the choices of Janner et al. [1], except that we experiment with
84 a linear noise schedule as an alternative to the cosine schedule they use. For each model and each
85 dataset, we train the diffusion model with both noise schedules and report the better of the two results.

86 **References**

- 87 [1] Michael Janner, Yilun Du, Joshua B Tenenbaum, and Sergey Levine. Planning with diffusion for
88 flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022. (Cited on pages 1, 2, and 3)
- 89 [2] Erwin Coumans and Yunfei Bai. PyBullet, a Python module for physics simulation for games,
90 robotics and machine learning. <http://pybullet.org>, 2016–2019. (Cited on page 2)
- 91 [3] Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in
92 actor-critic methods. In *International conference on machine learning*, pages 1587–1596. PMLR,
93 2018. (Cited on page 2)
- 94 [4] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah
95 Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of*
96 *Machine Learning Research*, 22(268):1–8, 2021. URL [http://jmlr.org/papers/v22/](http://jmlr.org/papers/v22/20-1364.html)
97 [20-1364.html](http://jmlr.org/papers/v22/20-1364.html). (Cited on page 2)
- 98 [5] Valentin De Bortoli, Emile Mathieu, Michael John Hutchinson, James Thornton, Yee Whye
99 Teh, and Arnaud Doucet. Riemannian Score-Based generative modelling. October 2022. URL
100 <https://openreview.net/pdf?id=oDRQGo8I7P>. (Cited on page 2)
- 101 [6] Chin-Wei Huang, Milad Aghajohari, Avishek Joey Bose, Prakash Panangaden, and Aaron
102 Courville. Riemannian diffusion models. August 2022. URL [http://arxiv.org/abs/](http://arxiv.org/abs/2208.07949)
103 [2208.07949](http://arxiv.org/abs/2208.07949). (Cited on page 2)