Dense Policy: Bidirectional Autoregressive Learning of Actions

— Supplementary Document —

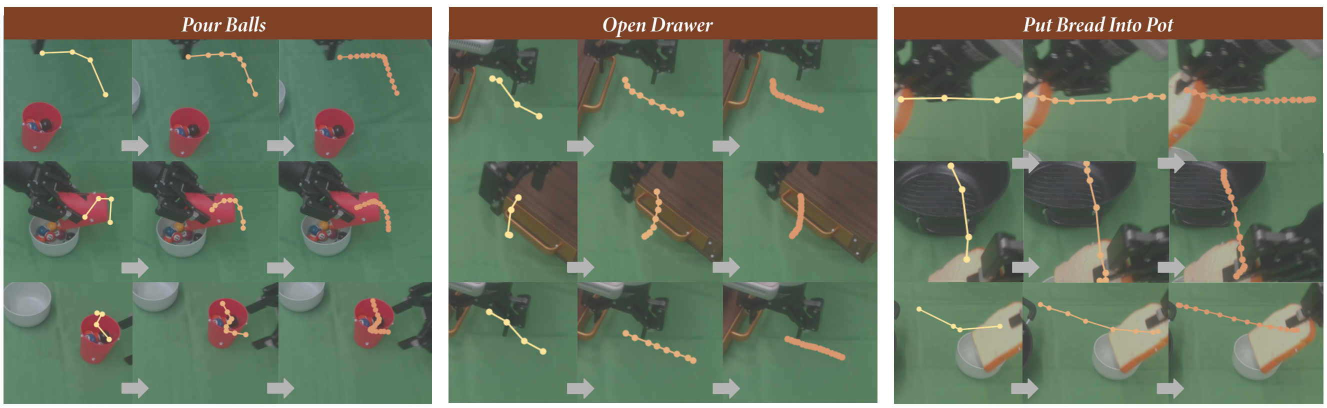

Figure 1. Visualizations of different action levels in Dense Process.

1. Dense Process

In this section, we provide visualizations during the different levels of Dense Process as shown in Fig. 1. Although supervision is only applied at the densest action level, the visualization reveals a clear coarse-to-fine and latent-to-explicit process within the dense process.

We also found that introducing supervision in each dense process level may let policy performs worse. Factually, early supervision channels policy's attention in high-frequency actions representations. Although this makes the model's early predictions more consistent with the corresponding time domain results of the ground truth, it is difficult to learn the overall trend of the action in the frequency domain, which makes it difficult for the early dense process reasoning to play a good exploratory role in global reasoning. As shown in Fig. 1, the early sparse action representations tend not to strictly correspond to their actions at the final sequence timestep, but rather tend to model the key transition actions in the sequence.

This is an interesting finding, which clearly shows that the dense policy is slightly different from the diffusion-based policies [1, 4]: in the inference phase, it explicitly models the interdependence of the sequence instead of making an overall distance prediction for each inference step [2].

What makes us even more excited is that this kind of sequence modeling not only reflects the time domain correlation of the sequence, but also reflects the frequency domain characteristics from low-frequency prediction to high-frequency prediction and from sparse key transition actions to complete trajectories. This makes us look forward to exploring whether dense policy can draw on the experience of frequency domain analysis [3] to complete more accurate sequence modeling in the future.

References

[1] C. Chi and et al. Diffusion policy: Visuomotor policy learning via action diffusion. In *RSS*, 2023. 1

[2] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 1

[3] Karl Pertsch, Kyle Stachowicz, Brian Ichter, Danny Driess, Suraj Nair, Quan Vuong, Oier Mees, Chelsea Finn, and Sergey Levine. Fast: Efficient action tokenization for vision-language-action models, 2025. 1

[4] Y. Ze and et al. 3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations. In *RSS*, 2024. 1