

What's the Move?

Hybrid Imitation Learning via Salient Points

Priya Sundaresan^{*1}, Hengyuan Hu^{*1},
Quan Vuong², Jeannette Bohg¹, Dorsa Sadigh¹

^{*}Equal contribution

¹Stanford University, ²Physical Intelligence




*How can we learn **sample-efficient** IL policies
that **generalize spatially** for **complex** tasks?*

Prior Work

Dense Policies


Actions on a *per-timestep* basis.



[1, 2, 3]

Hybrid Policies

Multiple action modes.



[6]

Waypoint Policies

Actions predicted *sparingly* as poses.



[4, 5]

[1] Chi, Cheng, et al. "Diffusion policy: Visuomotor policy learning via action diffusion." *IJRR*, 2023.

[2] Zhao, Tony Z., et al. "Learning fine-grained bimanual manipulation with low-cost hardware." *RSS*, 2023.

[3] Ze, Yanjie, et al. "3d diffusion policy." *arXiv preprint arXiv:2403.03954* (2024).

[4] Shridhar, Mohit, Lucas Manuelli, and Dieter Fox. "Perceiver-actor: A multi-task transformer for robotic manipulation." *CoRL*, 2023.

[5] Goyal, Ankit, et al. "Rvt: Robotic view transformer for 3d object manipulation." *CoRL*, 2023.

[6] Belkhale, Suneel, Yuchen Cui, and Dorsa Sadigh. "Hydra: Hybrid robot actions for imitation learning." *CoRL*, 2023.

Key Insight

- Different **input modalities** & **action spaces** are suited to different task phases:

Key Insight

- Different **input modalities** & **action spaces** are suited to different task phases:

Point Clouds + Waypoints → Long-Range Movements

Key Insight

- Different **input modalities** & **action spaces** are suited to different task phases:

Point Clouds + Waypoints → Long-Range Movements

Wrist Images + Dense Actions → Precise Movements

These mode switches typically occur around task-relevant **salient points** which can be learned.

SPHINX

Salient-Point based **H**ybrid **I**mitation and **e**Xecution

SPHINX Data Collection Interface


Record Waypoint

End Demo


[Translation] Key: Shift



SPHINX Data Collection



SPHINX



SPHINX: Waypoint Policy



SPHINX: Waypoint Policy



Classify salient (task-relevant) point,
predict waypoint as a relative offset to that point

SPHINX: Waypoint Policy




SPHINX: Waypoint Policy




SPHINX: Dense Policy


Diffusion Policy which outputs dense actions + mode.




SPHINX



Experiments: Waypoint-Only Tasks




Experiments: Waypoint-Only Tasks




■ DP ■ 3D DP ■ Vanilla Waypoint ■ Vanilla Waypoint + Aux. Salient Pts.

Experiments: Waypoint-Only Tasks




Experiments: Hybrid Tasks


Diffusion Policy




SPHINX




Experiments: Hybrid Tasks




SPHINX Results: Precise, Long-Horizon Tasks



SPHINX Results: Visual Generalization



Visual Distractors




Novel Viewpoint




Raised Drawer Height

SPHINX Results: Novel Execution Speeds



SPHINX
Vanilla Controller



SPHINX
2x Sped Up Controller

SPHINX Results: Spatial Generalization



What's the Move?

Hybrid Imitation Learning via Salient Points

Priya Sundaresan^{*1}, Hengyuan Hu^{*1},
Quan Vuong², Jeannette Bohg¹, Dorsa Sadigh¹

^{*}Equal contribution

¹Stanford University, ²Physical Intelligence



sphinx-manip.github.io

