BraintankAI

DISPATCH · ROBOT FORESIGHT · 3 MIN

Robot model forecasts three futures and weighs which parts matter for each motion

HiFOR predicts the next few moments, the next task milestone and the end goal, then lets its action generator weight what looks relevant, a preprint reports.

10 November 2026 · arXiv preprint, not yet peer reviewed

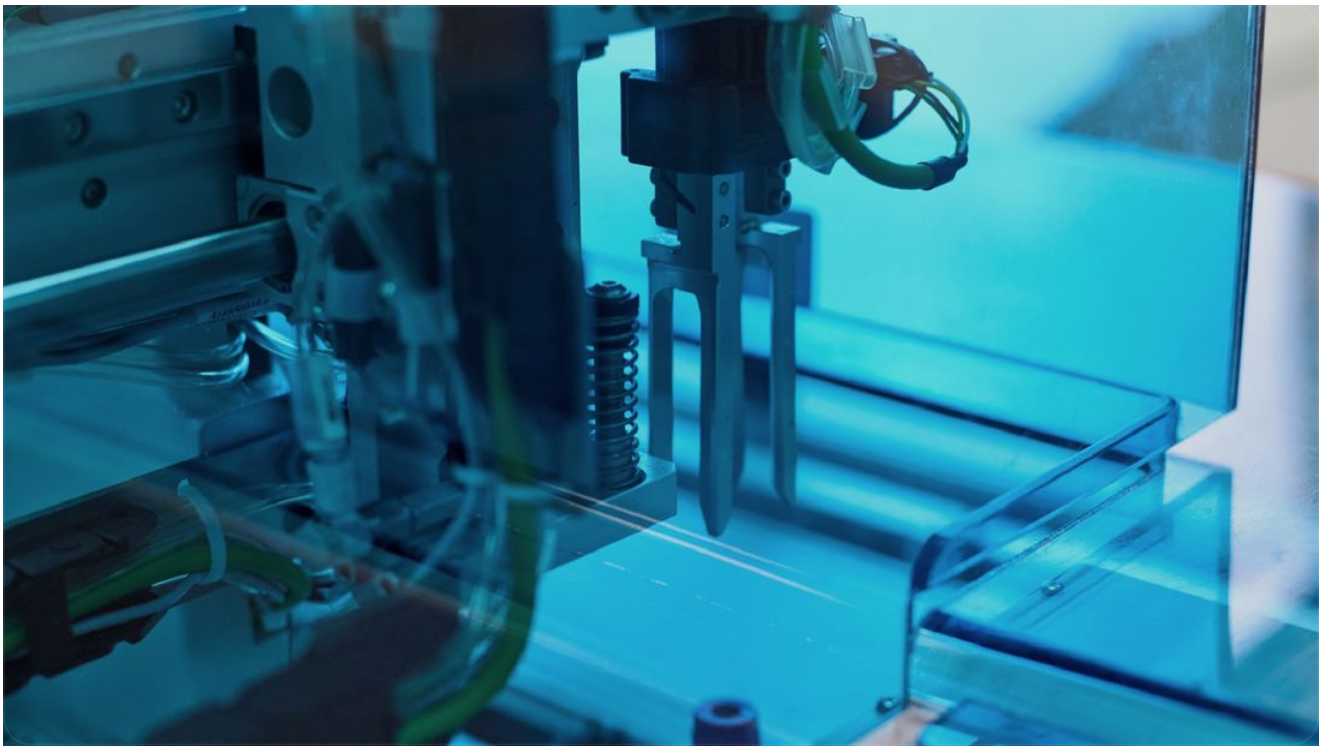


Photo: Testalize Me

Robot policies known as vision-language-action (VLA) models turn camera images and a typed instruction directly into motor commands. Some also predict how the scene will look later and use that foresight as a hint. HiFOR, described in a 2026 arXiv preprint that has not yet been peer reviewed, forecasts at three distances at once (a few steps ahead, the next checkpoint in the task and the finished goal). Its action generator then weights those predictions by their apparent relevance to the current motion, rather than receiving every prediction wholesale.

Preparing the training data takes work. A planning agent, which the paper does not name, splits each instruction into ordered subgoals, such as grasping a mug, each with a completion criterion. The moment each subgoal was reached in successful demonstrations is located automatically from gripper transitions, motion phases and robot state, and low-confidence cases are flagged for manual review. HiFOR does not forecast pixels. Each camera frame passes through DINOv2, a pretrained image encoder whose features summarise what a picture contains, and is pooled into a

grid of 16 patches. Three groups of learned queries predict those features for the short-term, milestone and goal frames. A Hierarchical Dependency Mask, a rule about which parts of the network may attend to which, lets the milestone forecast see the short-term one and the goal forecast see both, but not the reverse.

Actions are produced by flow matching, a method that starts from random noise and refines it over several denoising steps into a short sequence of motor commands. At every step the action generator weighs the forecast at three nested levels: which horizon, then which camera given that horizon, then which image patch. The combined weights decide what flows back into the action.

Success rates were lowest on the Hard setting of RoboTwin 2.0, a simulated manipulation benchmark that randomises clutter, lighting, textures and object heights. Across 20 tasks with 100 attempts each, the authors report an average success rate of 49.40%. The figure for π0.5, an existing flow-matching VLA whose results the authors reproduced and that HiFOR appears to build on, judging by the paper's tables, is 40.35%. The strongest of the remaining baselines were Spatial Forcing at 29.40% and X-VLA at 23.25%. HiFOR did not win everywhere: X-VLA beat it on a bell-clicking task, 63% to 48%. On the LIBERO benchmark, where scores are already high, HiFOR averaged 98.3 against 97.3 for the strongest other method, VLA-Adapter. Its largest LIBERO gain over π0.5 came on the long-horizon suite, from 92.4 to 95.8. The authors also tested a real AgileX Piper arm with wrist and third-person cameras on six grasping, placing and stacking tasks, each run 50 times. HiFOR averaged 81.7%, 11.0 points above π0.5, the only model it is compared with in the text.

Adding the components one at a time lifted RoboTwin success from 40.35 to 44.15, 45.60, 46.95 and finally 49.40 once the retrieval step was included. Predicting DINOv2 features rather than raw colour pixels scored 95.8 on LIBERO's long tasks, against 94.2. An attention visualisation showed the model leaning on the milestone forecast and the wrist camera while turning a stove knob, and on the goal forecast and the third-person camera during final placement. That evidence covers a few representative moments chosen by the authors.

The evidence has gaps. Every result is a single number, with no repeated seeds or error bars. Only ten of the 20 RoboTwin tasks appear in the main text. Only π0.5 is marked as reproduced; the paper does not say how the other baseline figures were obtained. There is no head-to-head comparison with the foresight methods the paper criticises. Because components were added in one fixed order, the mask and the three-step lookup are never tested in isolation. The extra inference cost is not reported. The data pipeline needs successful expert demonstrations, and the amount of manual review is not quantified.

braintankAI has previously covered a VLA that anticipates motion through compact tokens describing how a scene will change, rather than through generated video (earlier story). HiFOR likewise avoids pixel forecasts, and its own comparison favours encoder features over raw images. ThinkWorld used imagined outcomes as a training reward so that robot models could learn from human footage (earlier story). HiFOR instead consults its forecasts while generating each action, and relies on annotated robot demonstrations.

PAPER	*HiFOR: Hierarchical Foresight with Action-Guided World Grounding for Vision-Language-Action Model*
AUTHORS	Rongxing Ding, Hongyu Qu, Kehong Wu, Xudong Wang, Fan Zhang, Xinguang Xiang et al.
VENUE	arXiv preprint — not yet peer reviewed · 29 October 2026
LINK	arxiv.org/abs/2610.40373
INDEX	stacks.braintank.ai