Pi-R-Squared: Robot Policies That React Every 40 Milliseconds

Kabui, Charles

2026-08-01

Pi-R-squared is a retrofit for robot policies that predict several movements at once, then spend so long computing that the robot runs part of the plan blind. Carnegie Mellon researchers built it on NVIDIA's GR00T-N1.7 and split the control loop into a fast proprioception channel and a slower vision-language channel. Joint positions, torque, and fingertip force refresh every tick, while camera and instruction features update asynchronously. A latency-adaptive flow

schedule treats actions already in flight as context and emits a new action after one denoising step. On an xArm6 arm with an XHand, the authors report closed-loop replanning at about 25 Hz on an A5000 GPU, or one fresh observation every 40 ms, roughly four times faster than the base policy. They report gains of up to 23% in simulation and 30% on the physical robot.

That split targets a real failure mode: vision can be slightly stale, but a robot's joint and force readings cannot. The project page shows four physical tasks, and its rollout viewer exposes 20 trials per method. The result is not a universal real-time robot brain. Vision and language still lag, the method mainly supports local recovery, and reactive teleoperation data is hard to collect. The preprint is under review, while the public implementation provides deployment and training code for the xArm6 plus XHand setup.

Read More: TurboVLA runs a smaller robot policy at 32 Hz without an LLM.

Sources:

- Pi-R-squared paper
- Pi-R-squared project page
- Pi-R-squared rollout viewer
- Pi-R-squared implementation
