Fixing weak language grounding in a robot policy's weights

Three robot-policy architectures, one released model of each, on the same RoboLab DROID tasks. The VLA with an action expert is OpenPI π0.5, the VLA on a reasoning backbone is GR00T N1.7 DROID with Cosmos-Reason2-2B, and the world-action model is Cosmos3 Edge Policy. The tasks are Cube then banana, Eyeglasses, Cube left of bowl, Yellow on red, Mustard in left bin, Three-cube tower, Flat box, Mustard right of raisin box and a retention task Cube in bowl.

Method. When a clearer wording of the same goal makes a robot succeed, copy that wording's internal effect into the weights. It takes one least-squares calculation per layer and no training, and the robot still gets the original instruction.

VLA with an action expert, π0.5. The edit fixed its two grounding failures, Cube then banana from 2 to 40 of 64 unseen starts and Mustard right of raisin box from 0 to 28 of 64. A random edit of the same size did nothing.

Across architectures. The fix. A closed-form weight edit makes a world-action model, Cosmos3 Edge, succeed from 56 of 64 new starts, as often as rewording does. **When it works.** Only in the VLM or text pathway. **The cost.** The edit overgeneralises the spatial relation it fixes.



Cosmos3 Edge Policy, Mustard right of raisin box: the wording gap; the clip plays on the live page.

What the edit did on each task, on unseen starts unless marked practice

	Cosmos3 Edge Policy world-action model	**GR00T N1.7** reasoning-backbone VLA	**π0.5** OpenPI VLA, control
Cube then banana	Mustard right of raisin box edit: 18 → 18 of 32 9 / 9; p = 1	9 → 0 of 32 practice starts 0 / 9, p = 0.0039	2 → 40 of 64 38 / 0, p = 7.3×10^{-12}; paper's evaluation: fixed: 27 → 106 of 128
Eyeglasses	Mustard right of raisin box edit: 4 → 1 of 16 0 / 3; p = 0.25	unedited 0 of 32 not edited	Cube then banana edit: 12 → 21 of 32 12 / 3; p = 0.035; paper's evaluation: Cube then banana edit: 22 → 44 of 64
Cube left of bowl	Mustard right of raisin box edit: 13 → 0 of 32 0 / 13; p = 2.4×10^{-4}	unedited 1 of 32 not edited	Cube then banana edit: 1 → 6 of 32 6 / 1; p = 0.12; paper's evaluation: Cube then banana edit: 8 → 36 of 64
Yellow on red	Mustard right of raisin box edit: 24 → 22 of 32 4 / 6; p = 0.75	unedited 0 of 16 not edited	no gain lookalike perception; no clear wording
Mustard in left bin	Mustard right of raisin box edit: 9 → 0 of 32 0 / 9; p = 0.0039	unedited 0 of 32 not edited	no gain own edit 8 → 11 of 64; the grasp slips
Three-cube tower	Mustard right of raisin box edit: 1 → 1 of 32 1 / 1; p = 1	unedited 0 of 16 not edited	no gain own edit 35 → 39 of 64
Flat box	Mustard right of raisin box edit: 0 → 0 of 16 0 / 0; p = 1	unedited 0 of 32 not edited	no gain a missing skill; 0 of 64 unedited
Mustard right of raisin box	0 → 56 of 64 56 / 0, p = 2.8×10^{-17}	unedited 0 of 32 not edited	0 → 28 of 64 28 / 0, p = 7.5×10^{-9}; paper's evaluation: fixed: 0 → 27 of 64
Cube in bowl	24 → 28 of 32	unedited 28 of 32 not edited	31 → 31 of 32 paper's evaluation: kept: 32 → 32 of 32
Rendering: performance → default preset	—	5 → 16 of 16 Banana onto plate	—
Random control, same size	control 0, unedited 0, edit 56 of 64 Mustard right of raisin box	control 11, unedited 9, edit 0 of 32 practice starts Cube then banana	control 2, unedited 2, edit 40 of 64 Cube then banana
Where to edit	text pathway 13, video-action (generation) pathway 0 of 16 Mustard right of raisin box, 16 practice starts; unedited 0	Cosmos-Reason2 VLM backbone 0, flow-matching DiT action head 0 of 32 Cube then banana, 32 practice starts; unedited 9	VLM 55, action expert 20, output projection 0 of 64 Cube then banana; unedited 5; only the VLM edit transfers

	Cosmos3 Edge	**GR00T N1.7**	**π0.5**
Cube then banana	21/32 · 66%	9/32 · 28%	3/16 · 19%
Eyeglasses	2/32 · 6%	0/32 · 0%	7/16 · 44%
Cube left of bowl	13/32 · 41%	1/32 · 3%	0/16 · 0%
Yellow on red	22/32 · 69%	0/16 · 0%	9/16 · 56%
Mustard in left bin	11/32 · 34%	0/32 · 0%	3/16 · 19%
Three-cube tower	2/32 · 6%	0/16 · 0%	6/16 · 38%
Flat box	0/32 · 0%	0/32 · 0%	0/16 · 0%
Mustard right of raisin box	0/32 · 0%	0/32 · 0%	0/16 · 0%
Cube in bowl	21/32 · 66%	28/32 · 88%	16/16 · 100%

Each cell: the starts the unedited robot completed, of the starts tried, and that share; stronger shading is a higher share.

	Cosmos3 Edge	**π0.5**
Cube then banana	±0 / 32	**+38 / 64***
Eyeglasses	−3 / 16	**+9 / 32***
Cube left of bowl	**−13 / 32***	+5 / 32
Yellow on red	−2 / 32	+6 / 32
Mustard in left bin	**−9 / 32***	−1 / 32
Three-cube tower	±0 / 32	±0 / 32
Flat box	±0 / 16	±0 / 32
Mustard right of raisin box	**+56 / 64***	±0 / 32
Cube in bowl	+4 / 32	±0 / 32

gain, $p < 0.05$ loss, $p < 0.05$ no detectable change

Cell: change in starts completed / starts tried. * $p < 0.05$; outlined: the edit's own task

Each model's main edit, unchanged, on every suite task: the change in starts completed, edited minus unedited, from the same starts. Teal: a gain; near-black (near-white in dark mode): a loss (exact McNemar $p < 0.05$); grey: no detectable change; the outlined cell is the task the edit was built for.

Live page with every table, figure and record: ali-adeeb.com/weightsmith/cosmos · Paper: ali-adeeb.com/weightsmith/paper.pdf · Code and records: github.com/adeeb10abbas/weightsmith

Draft for discussion. Simulation only (RoboLab, Isaac Sim), one released checkpoint per architecture, shared slots 30000–30031 and 96000–96063 with identical initial scenes; π0.5 cells from the paper's evaluation name their own slots. Exact McNemar tests on paired starts.