Tactile Sensors Are an Effective Substitute for Wrist Cameras in Contact-Rich or Occluded Manipulation

Aijia Yang **Arsh Tangri** **Sergey Zakharov** **Krishna Murthy Jatavallabhula**
Toyota Research Institute
any015@ucsd.edu
https://annieyaj.github.io/tactile-vs-vision/

Abstract:
Learned manipulation policies often rely on wrist-mounted cameras for close-range visual feedback that complements external views of the workspace. However, contact-rich manipulation requires establishing and maintaining contacts that can remain hidden even from wrist cameras. Tactile sensors mounted on fingers can reliably and directly measure contact, without needing visual access. In this work, we show that this feedback enables policies to recover the performance advantage of wrist cameras using scene and egocentric vision together with touch. Our approach incorporates two low-cost tactile arrays into a visuotactile ACT policy, without wrist images during training or deployment. We evaluate six real-world task families spanning placement, insertion, occluded retrieval, and sustained-contact tool use. On three contact-rich tasks, we independently vary wrist-camera and tactile availability while keeping the demonstrations, policy architecture, and optimization settings fixed. Adding wrist cameras improves success by 18.3–28.3 percentage points over vision without wrist images. Adding tactile sensing instead improves success by 31.7–60.0 percentage points and yields higher observed success than vision-only policies with wrist cameras on all three tasks, with significant advantages on two under Holm-corrected tests. These results support finger-mounted tactile sensing as an effective alternative to wrist cameras in the evaluated contact-rich settings.

Keywords: Tactile Sensing, Imitation Learning, Contact-Rich Manipulation

1 Introduction

Learned manipulation policies rely on external cameras for scene context and wrist-mounted cameras for close-range observations of the gripper and manipulated object. However, during manipulation, the fingers can occlude grasped objects, and whether contact is maintained may be difficult to observe even with wrist cameras. Tactile sensors provide direct measurements at the finger–object interface, supplying feedback without requiring visual access to the contact. This raises two research questions: (a) *Can tactile sensing substitute for, and improve on, wrist cameras?* and (b) *When?* We investigate these questions by comparing visuotactile and vision-only policies while controlling wrist-camera availability.

Using the Raiden toolkit [1], we train and deploy Action Chunking with Transformers (ACT) [2] policies on a YAM robot arm equipped with two low-cost FlexiTac [3] tactile sensor arrays. Our primary configuration combines scene and egocentric images with tactile observations, without wrist images during training or deployment. We evaluate multiple tasks spanning placement, fine alignment, occluded retrieval, and sustained contact. We compare vision-only and visuotactile policies while varying the availability of wrist cameras, under the same trajectory data distribution. Our contribution is a controlled empirical study of when finger-mounted tactile input can substitute for wrist images in an otherwise fixed ACT pipeline.

10th Conference on Robot Learning (CoRL 2026), Austin TX, USA.

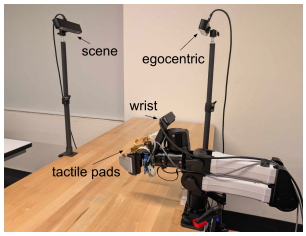

(a) station setup



(b) tactile sensor

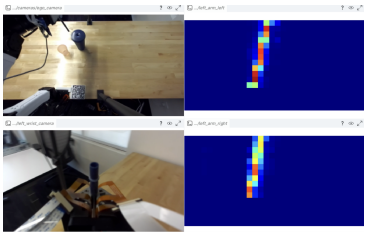

(c) T3 tactile heatmap



(d) (e) (f) (g) (h) (i)

Figure 1: **System and evaluation tasks.** Two FlexiTac pads are mounted on the inner gripper surfaces and recorded alongside scene, egocentric, and wrist RGB observations. Panels d–i show tasks T1–T6, spanning placement, fine alignment, occluded retrieval, and sustained contact.

We observe substantial success gains from adding wrist cameras for vision-only policies. However, we also find that adding tactile sensing results in similar (or better) performance, where vision+tactile (V+T) policies without wrist cameras achieve higher observed success than vision-only (V) policies with wrist cameras on contact-rich tasks. These results support tactile sensing as an effective substitute for wrist cameras in the evaluated contact-rich or occluded manipulations.

2 Related Work

Low-cost tactile sensors such as FlexiTac [3] and eFlesh [4] have made contact sensing increasingly accessible. TacO [5] shows that tactile benefits depend on both task and sensor modality, rather than providing uniform improvements. Earlier approaches such as Visuo-Tactile Transformers [6] and 3D-ViTac [7] demonstrate the benefits of learned visual-tactile representations for manipulation. We use FlexiTac to study how the benefit of tactile input changes with visual observations available to a learned manipulation policy.

Closest to our setting, LeFlexiTac [8] integrates FlexiTac into several manipulation policies and reports particularly large gains for occluded retrieval. VT-Refine [9] and T-Rex [10] also demonstrate the value of tactile feedback for assembly and contact-rich dexterity. Whereas these works establish that touch can help, we hold the sensor type and policy class fixed and vary tactile input and wrist-vision availability, examining whether touch provides benefits beyond wrist vision and whether it can substitute for wrist vision when scene and egocentric observations remain available.

3 Methodology

Experiments use a YAM robotic platform equipped with scene, egocentric, and wrist RGB cameras (Fig. 1a). Throughout this paper, without wrist vision means that scene and egocentric images are retained and only wrist images are omitted. Two FlexiTac [3] piezoresistive arrays are mounted on the inner surfaces of the TPU compliant gripper fingers (Fig. 1b, Appendix A). Tactile sensing is recorded at approximately 99 Hz and RGB observations at approximately 30 Hz using a shared clock.

During dataset construction, tactile measurements are aligned to each policy timestep by nearest timestamp. The full-rate tactile stream is retained in recordings rather than down-sampled to camera frequency, allowing future use of high-frequency contact information.

Table 1: **Vision-only vs. vision+tactile success.** Each policy entry reports successes/rollouts (success rate). T4–T6 are evaluated both with and without wrist images. Δ is the observed V+T-minus-V difference in percentage points.

	Task	Wrist	V	V+T	Δ
T1	Tight-clearance block	on	15/30 (50.0%)	19/30 (63.3%)	+13.3 pp
T2a	Rigid block	on	15/25 (60.0%)	15/25 (60.0%)	0.0 pp
T2b	Sponge	on	19/25 (76.0%)	22/25 (88.0%)	+12.0 pp
T3	Straw insertion	on	26/60 (43.3%)	39/60 (65.0%)	+21.7 pp
T4	Opaque-bag retrieval	off	13/60 (21.7%)	49/60 (81.7%)	+60.0 pp
		on	30/60 (50.0%)	43/60 (71.7%)	+21.7 pp
T5	Plug insertion	off	9/60 (15.0%)	28/60 (46.7%)	+31.7 pp
		on	20/60 (33.3%)	30/60 (50.0%)	+16.7 pp
T6	Whiteboard wiping	off	13/60 (21.7%)	43/60 (71.7%)	+50.0 pp
		on	28/60 (46.7%)	39/60 (65.0%)	+18.3 pp

All experiments use ACT (Action Chunking with Transformers) [2], predicting 19-step chunks of absolute 7-DoF left-arm joint position commands (Training details in Appendix B). We compare **vision-only (V)**, which receives RGB observations, against **vision+tactile (V+T)**, which additionally receives two 32×16 tactile maps. Each grayscale tactile map is replicated to three channels, resized to 224×224, and processed by the same shared ImageNet-pretrained ResNet-18 [11] used for camera views. Tactile sensing therefore enters ACT as two additional image-like observations. The policy class, demonstration data, action representation, optimization settings, and deployment procedures are otherwise fixed between observation conditions. Checkpoint selection is described separately in Appendix B.

We evaluate tight-clearance block placement (T1), rigid-block and deformable-sponge placement (T2a/b), straw insertion (T3), opaque-bag marker retrieval (T4), plug insertion (T5), and whiteboard wiping (T6). T1–T2 are largely visually observable placement tasks; T3 and T5 require precise alignment; T4 involves retrieving an object from inside an opaque bag; and T6 additionally requires maintaining contact during wiping.

For each task, V and V+T are evaluated from matched initial-condition sets using fixed success criteria and rollout counts. For T4–T6, both policies are trained and evaluated with and without wrist images. The evaluation comprises 1,000 rollouts across 20 task-and-observation conditions, counting T2a and T2b separately. Dataset construction, episode-level splits, and evaluation allocations are described in Appendix B.

4 Experimental Results

Table 1 reports observed success rates. Posterior summaries, tactile-gain intervals, and Holm-corrected tests are provided in Appendix C. Counts by initial condition are in Appendix D.

Smaller gains on placement tasks. Tactile input leaves the observed rigid-block success unchanged at 60.0% and improves tight-clearance and sponge placement by $13.3\,\mathrm{pp}$ and $12.0\,\mathrm{pp}$, respectively. None of these differences are statistically significant under the corrected tests.

Touch helps with fine alignment. Straw insertion success increases from 43.3% to 65.0% ($+21.7\,\mathrm{pp}$). This improvement occurs despite the target being visible, showing that tactile benefits are not limited to fully occluded manipulation.

Largest gains without wrist vision. Without wrist vision, tactile input increases success by $60.0\,\mathrm{pp}$ for opaque-bag retrieval, $31.7\,\mathrm{pp}$ for plug insertion, and $50.0\,\mathrm{pp}$ for whiteboard wiping. These tasks involve occluded grasping, fine insertion, and sustained contact, respectively.

4.1 Wrist-Camera Control Comparison

Adding wrist vision improves V success by $28.3\,\mathrm{pp}$, $18.3\,\mathrm{pp}$, and $25.0\,\mathrm{pp}$ on T4–T6, respectively. The corresponding tactile gains shrink to $21.7\,\mathrm{pp}$, $16.7\,\mathrm{pp}$, and $18.3\,\mathrm{pp}$.

Adding wrist images does not consistently improve V+T. Observed success changes from 81.7% to 71.7% on T4, from 46.7% to 50.0% on T5, and from 71.7% to 65.0% on T6. The wrist-on and wrist-off V+T conditions are not significantly different on T4 or T6 under the Holm-corrected tests (Appendix C), although this does not establish equivalence. Smaller marginal tactile gains are consistent with overlapping task-relevant information, but information overlap alone does not explain the lower absolute point estimates. Finite evaluation samples, training and checkpoint variability, and differences in how the learned policies use the available observations are possible explanations; the present experiment does not distinguish among them.

For T4, the lower V+T wrist-on success is concentrated in far-right trials ($9/20$ vs. $15/20$ without wrist images; Table 6). This post hoc comparison does not establish causality.

4.2 Can Touch Substitute for Wrist Vision?

We directly compare V+T without wrist images against V with wrist images. Observed success is higher with touch on all three contact-rich tasks: 81.7% vs. 50.0% for opaque-bag retrieval, 46.7% vs. 33.3% for plug insertion, and 71.7% vs. 46.7% for whiteboard wiping. The differences are statistically significant for T4 and T6 under the Holm-corrected tests.

These results support touch as an alternative to wrist vision on the evaluated contact-rich tasks, not as a replacement for vision overall: both conditions retain scene and egocentric images.

Stagewise outcomes in wiping. In T6, tactile input increases observed eraser-pickup success from $23/60$ to $48/60$ without wrist images and from $31/60$ to $42/60$ with wrist images. Task success given pickup increases from $13/23$ to $43/48$ without wrist images, but is similar with wrist images ($28/31$ vs. $39/42$).

Comparing the two V+T conditions, observed pickup success is $48/60$ without wrist images and $42/60$ with wrist images, while task success given pickup is $43/48$ (89.6%) and $39/42$ (92.9%), respectively. The pickup decrease is concentrated in the bottom initial position, with $15/20$ successful pickups without wrist images and $9/20$ with wrist images (Table 8). Thus, lower overall wrist-on success coincides with fewer successful pickups, while observed completion among successful pickups remains high. These comparisons are descriptive: the successful-pickup subsets differ across policies, and the position-level analysis does not establish a causal mechanism.

5 Discussion and Limitations

The benefit of touch depends strongly on the task and on how much close-range visual information is available. Observed gains are smaller on placement tasks and largest on T4–T6 without wrist images. Adding wrist vision narrows the tactile advantage but does not consistently improve V+T; the cause of the lower wrist-on success rates remains unresolved. Without wrist images, V+T achieves higher observed success than V with wrist images on all three contact-rich tasks, with significant advantages on T4 and T6. In these tasks, touch can recover—and in some cases exceed—the benefit of adding wrist vision, while scene and egocentric vision remain available.

Our study is limited to one robot platform, one tactile sensor type, one policy-learning method, and one training seed. Low-cost tactile pads also exhibit baseline variation after replacement, motivating improved normalization and evaluation across additional policies and sensors. Overall, the results suggest that vision guides the robot toward contact, while touch resolves interaction states that cameras cannot reliably observe.

References

[1] S. Iwase, P. Miller, J. Yao, K. Jatavallabhula, and S. Zakharov. RAIDEN: A toolkit for policy learning with YAM bimanual robot arms, 2026. URL https://tri-ml.github.io/raiden/.

[2] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning fine-grained bimanual manipulation with low-cost hardware, 2023. URL https://arxiv.org/abs/2304.13705.

[3] B. Huang and Y. Li. Flexitac: A low-cost, open-source, scalable tactile sensing solution for robotic systems, 2026. URL https://arxiv.org/abs/2604.28156.

[4] V. Pattabiraman, Z. Huang, D. Panozzo, D. Zorin, L. Pinto, and R. Bhirangi. eflesh: Highly customizable magnetic touch sensing using cut-cell microstructures, 2025. URL https://arxiv.org/abs/2506.09994.

[5] A. Zorin, Z. Si, M. Park, J. Park, A. Buynitsky, S. Bhadang, T. Park, S. J. Yoon, Y.-L. Park, O. Kroemer, Z. Temel, M. T. Tolley, S. Yi, and X. Wang. Taco: Benchmarking tactile sensors for object manipulation, 2026. URL https://arxiv.org/abs/2605.21976.

[6] Y. Chen, A. Sipos, M. V. der Merwe, and N. Fazeli. Visuo-tactile transformers for manipulation, 2022. URL https://arxiv.org/abs/2210.00121.

[7] B. Huang, Y. Wang, X. Yang, Y. Luo, and Y. Li. 3d-vitac: Learning fine-grained manipulation with visuo-tactile sensing, 2025. URL https://arxiv.org/abs/2410.24091.

[8] N. Tao, Y. He, W. Maa, B. Huang, and Y. Li. LeFlexiTac: Giving robots a sense of touch. *Columbia University RoboPIL Blog*, 2026. https://tna001-ai.github.io/LeFlexiTac/.

[9] B. Huang, J. Xu, I. Akinola, W. Yang, B. Sundaralingam, R. O'Flaherty, D. Fox, X. Wang, A. Mousavian, Y.-W. Chao, and Y. Li. VT-refine: Learning bimanual assembly with visuo-tactile feedback via simulation fine-tuning. In *9th Annual Conference on Robot Learning*, 2025. URL https://openreview.net/forum?id=bOVF8Rj33i.

[10] D. Niu, Z. Liu, Z. Wang, B. Shao, Z.-H. Yin, A. Pai, Y. Sharma, S. Saravalle, R. Zheng, J. Wang, R. Punamiya, M. Xu, Y. Xie, Y. Jiang, L. Fu, K. Kallidromitis, M. Gioia, J. Zhang, J. Ge, H. Feng, F. Galasso, W. Zhan, D. M. Chan, Y. Bai, R. Herzig, J. Lei, F.-F. Li, K. Goldberg, J. Malik, P. Abbeel, Y. Zhu, D. Xu, J. Fan, and T. Darrell. T-rex: Tactile-reactive dexterous manipulation, 2026. URL https://arxiv.org/abs/2606.17055.

[11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition, 2015. URL https://arxiv.org/abs/1512.03385.

[12] D. Snyder, A. J. Hancock, A. Badithela, E. Dixon, P. Miller, R. A. Ambrus, A. Majumdar, M. Itkina, and H. Nishimura. Is your imitation learning policy better than mine? policy comparison with near-optimal stopping, 2025. URL https://arxiv.org/abs/2503.10966.

A Hardware Designs

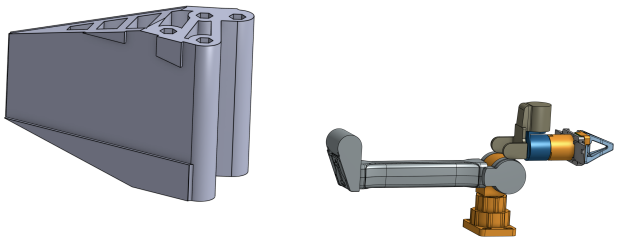

(a) TPU gripper design with sensor slots

(b) Single arm CAD with electronics mounts

Figure 2: Customized station design

B Training and Evaluation Details

Dataset construction. All demonstrations are collected on the YAM station described in Sec. 3. RGB observations are recorded at approximately 30 Hz and tactile observations at approximately 99 Hz using a shared clock. The full-rate tactile stream is retained during recording, and each policy timestep is paired with the tactile measurement having the nearest timestamp. Each demonstration is stored as one episode. Training and validation sets are split by episode, with 10% of episodes held out for validation using a fixed split by episode ID. Within each task, observation conditions use the same demonstration episodes and episode-level split. All experiments use random seed 42. ACT targets contain 19 consecutive future action steps; the data-loader stride changes the spacing between sampled training windows and does not subsample actions within the chunk.

Table 2: **Demonstration dataset sizes.** Total recorded episodes and successful episodes before the episode-level training/validation split.

Task	Recorded	Successful
T1: Block placement	267	229
T2: Block/sponge dataset	225	208
T3: Straw insertion	220	209
T4: Bag retrieval	235	203
T5: Plug insertion	231	202
T6: Whiteboard wiping	218	207
Total	1396	1258

Table 2 summarizes the recorded demonstration episodes and successful episodes for each task. Only successful demonstration episodes were retained for policy training before the episode-level training/validation split.

Preprocessing. RGB images are stored at 384×384 using JPEG quality 95 and processed at 224×224. During training, camera observations are resized, cropped to 224×224, normalized with ImageNet statistics, and augmented with color jitter $(0.2, 0.2, 0.2)$. Because the stored frames are square and the resize and crop sizes are identical, the nominal random crop does not introduce spatial augmentation in our setting. At evaluation, the corresponding deterministic preprocessing is used. Each tactile observation is a 32×16 grayscale map, replicated to three channels and resized directly to 224×224. Tactile inputs receive no crop or color augmentation and use the same ImageNet normalization.

ACT architecture and optimization. All experiments use ACT [2] with a shared ImageNet-pretrained ResNet-18 [11] for camera and tactile views. The backbone uses `FrozenBatchNorm2d` and produces a $7\times7\times512$ feature map per view. The transformer uses model dimension 512, 8 attention heads, a 3200-dimensional feed-forward layer, four encoder layers, two decoder layers, four CVAE encoder layers, and a 32-dimensional latent variable. The model contains approximately 56.9M parameters. Table 3 lists the optimization settings.

Table 3: **ACT training settings**

Setting	Value
Action space	7-D absolute joint positions
Action chunk	19 consecutive steps
KL weight	10.0
Batch size	48
Optimizer	AdamW
Learning rate	10^{-5}
Weight decay	10^{-4}
Gradient clipping	1.0 global norm
Warmup	500 steps
Schedule	cosine decay to 0
Training steps	100,000
Precision	bf16 autocast
Seed	42

Training minimizes masked L1 action reconstruction plus the KL term,

$$\mathcal{L}_{\text{train}} = \mathcal{L}_{\text{L1}} + 10\,\mathcal{L}_{\text{KL}}.$$

Validation and checkpoints. Validation is performed every 1,000 optimization steps on up to 50 batches from held-out episodes. In evaluation mode, ACT disables the CVAE branch, so the reported validation objective is masked L1 only rather than the full training objective. The training script saves `best.pt` whenever held-out L1 improves and numbered checkpoints every 10,000 steps. Reported evaluations use `best.pt` for all tasks except T5. For T5, we compared three checkpoints: `best.pt`, `last.pt`, and `80000.pt`, using 10 physical development rollouts per checkpoint for each of the four observation conditions. All four conditions followed the same checkpoint-selection procedure, and the 80,000-step checkpoint achieved higher observed success than the alternatives. We therefore selected `80000.pt` for the final T5 evaluations. All checkpoint-selection rollouts were conducted separately from and excluded from the 60 final evaluation rollouts per condition. Thus, T5 checkpoint selection was based on physical rollout performance rather than validation loss alone.

Experimental conditions and rollout protocol. The primary comparison is between V and V+T policies trained from the same task demonstrations. For T4–T6, this comparison is crossed with wrist-camera availability. When wrist images are disabled, scene and egocentric images remain available in both training and deployment.

At deployment, ACT predicts a new 19-step action chunk at each policy timestep, and overlapping predictions are combined using temporal ensembling. Each rollout is limited to 900 control steps. Within each task, policies are evaluated using matched initial-condition sets, identical rollout counts, and the same success criterion. The reported evaluation comprises 30 rollouts per condition for T1, 25 each for T2a and T2b, and 60 for T3–T6, totaling 1000 rollouts. These are evaluation executions, not demonstration episodes.

Table 4 lists the initial-condition allocation, which is shared across the observation conditions within each task. For T4, reported success rates pool the five positions and therefore reflect the nonuniform allocation of trials. Tables 5–8 give the corresponding outcome counts.

Success criteria. T1 requires successful placement in the target box. T2 requires successful grasping and placement in the box, with no sponge deformation for T2b. T3 requires the straw to be

inserted through the cup-lid opening. T4 requires retrieval of the marker from the opaque bag; a trial in which the marker is subsequently lost before completion is counted as a failure. T5 requires the plug to be inserted in the socket. T6 requires at least 50% of the marked whiteboard region to be erased. In T6, successful eraser pickup is also recorded separately to distinguish pickup failures from failures after pickup.

Table 4: **Evaluation allocation per observation condition.**

Task	Initial conditions	Trials per position	Total
T1	Middle, top left, top right, bottom left, bottom right, far right in front of the box	5 each	30
T2a/b	Middle, top left, top right, bottom left, bottom right	5 each	25 each
T3	Middle, top left, top right, bottom left, bottom right	12 each	60
T4	Far left, left, middle, right, far right	10, 5, 5, 20, 20	60
T5	Left, right, middle, top, bottom	12 each	60
T6	Top, center, bottom	20 each	60

C Posterior Success Probabilities and Pairwise Tests

Table 1 reports success-rate point estimates. For s successes in n rollouts, a uniform $\mathrm{Beta}(1,1)$ prior yields the posterior

$$p \mid s, n \sim \mathrm{Beta}(s+1, n-s+1).$$

Figure 4 shows the condition-wise posteriors together with compact letter displays [12]. Pairwise comparisons use Barnard tests, with Holm correction applied within each task panel. The four-condition panels for T4–T6 each contain six pairwise comparisons; each two-condition panel contains one comparison. Conditions sharing a letter are not distinguishable at the corrected $\alpha = 0.05$ level. Failure to reject a difference does not establish equivalence.

The tests and posterior intervals concern the evaluated policies and rollout allocation. They do not account for training-seed variation or checkpoint-selection uncertainty. The additional position-level comparisons in the main text and Appendix D are descriptive post hoc analyses; no subgroup significance claims are made.

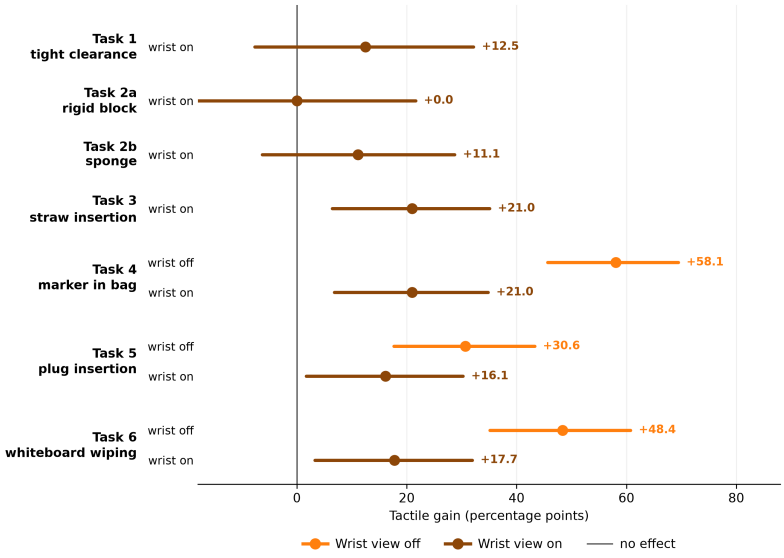

Figure 3: **Tactile gain by task and wrist-camera availability.** Posterior mean and 90% credible interval for the difference in success probability between V+T and V.

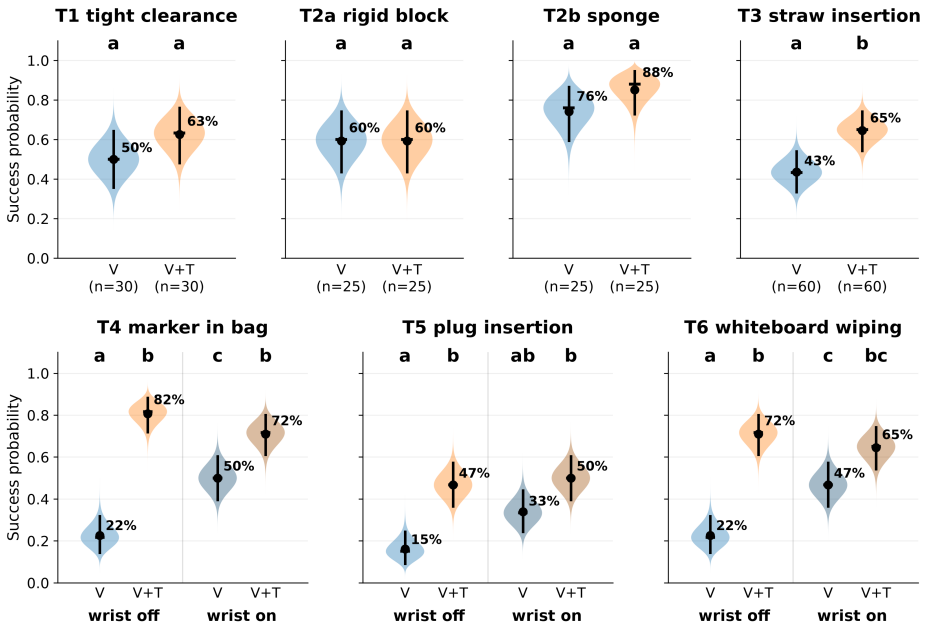

Posterior mean and 90% credible interval; short bar = observed rate. Blue = without tactile, orange = with tactile; muted shade = wrist view on. Shared letter = no significant difference (Barnard exact, Holm-corrected, alpha=0.05). Letters are comparable WITHIN a panel only.

Figure 4: **Posterior success probability for every task, with compact letter displays.** Beta posteriors under a uniform prior. Top row: the wrist-on tasks (T1–T3), two conditions each. Bottom row: the ablated tasks (T4–T6), four conditions each. Letters come from Holm-corrected pairwise Barnard tests computed within each panel; conditions sharing a letter could not be distinguished at α=0.05. Letters are not comparable across panels. In T4 and T6 the two V+T conditions share a letter while the two V conditions do not—no significant wrist-on versus wrist-off difference was detected for V+T in these tasks.

D Evaluation Outcomes by Initial Condition

The following tables provide the recorded numerators and denominators underlying Table 1. T2a and T2b are reported as separate evaluation variants. All counts are evaluation outcomes. The breakdowns support the descriptive analyses in Sec. 4.1 and Sec. 4.2; they do not constitute additional training runs or checkpoint evaluations.

Table 5: **T1–T3 success counts by initial condition.** Wrist images are enabled. Each entry is successes/rollouts.

Task	Initial condition	V	V+T
T1	Middle	4/5	5/5
	Top left	2/5	3/5
	Top right	0/5	2/5
	Bottom left	1/5	0/5
	Bottom right	5/5	5/5
	Far right in front of the box	3/5	4/5
	Total	15/30	19/30
T2a	Middle	2/5	2/5
	Top left	3/5	3/5
	Top right	4/5	3/5
	Bottom left	3/5	4/5
	Bottom right	3/5	3/5
	Total	15/25	15/25
T2b	Middle	4/5	4/5
	Top left	3/5	4/5
	Top right	4/5	5/5
	Bottom left	4/5	4/5
	Bottom right	4/5	5/5
	Total	19/25	22/25
T3	Middle	5/12	6/12
	Top left	8/12	8/12
	Top right	8/12	9/12
	Bottom left	2/12	7/12
	Bottom right	3/12	9/12
	Total	26/60	39/60

Table 6: **T4 retrieval success by initial position.** Each entry is successes/rollouts. The aggregate V+T wrist-on decrease is concentrated in the far-right group; the other positions yield $34/40$ successes in both V+T configurations.

	Wrist off		Wrist on	
Initial position	V	V+T	V	V+T
Far left	3/10	9/10	5/10	8/10
Left	4/5	5/5	5/5	4/5
Middle	2/5	4/5	5/5	5/5
Right	4/20	16/20	11/20	17/20
Far right	0/20	15/20	4/20	9/20
Total	13/60	49/60	30/60	43/60

Table 7: **T5 insertion success by initial position.** Each entry is successes/rollouts.

	Wrist off		Wrist on	
Initial position	V	V+T	V	V+T
Left	0/12	8/12	3/12	7/12
Right	1/12	3/12	5/12	3/12
Middle	3/12	8/12	6/12	7/12
Top	4/12	4/12	4/12	6/12
Bottom	1/12	5/12	2/12	7/12
Total	9/60	28/60	20/60	30/60

Table 8: **T6 pickup and task success by initial position.** Each policy entry is (P, S), where P is the number of successful pickups and S is the number of successful complete trials. The trial count n applies to every policy in that row.

		Wrist off		Wrist on	
Initial position	n	V	V+T	V	V+T
Top	20	$(5, 4)$	$(17, 13)$	$(8, 7)$	$(18, 15)$
Center	20	$(15, 6)$	$(16, 15)$	$(17, 15)$	$(15, 15)$
Bottom	20	$(3, 3)$	$(15, 15)$	$(6, 6)$	$(9, 9)$
Total	60	$(23, 13)$	$(48, 43)$	$(31, 28)$	$(42, 39)$