Maintaining Demonstration Quality in a 100-Robot Teleoperation Pipeline

Kunal Kapoor Shiraz Khan Joseph McCalmon Jesse Michel Arif Mohammed
Katerina Nikiforova Adam Oppenheimer Victor Szabo David Watkins
Tutor Intelligence, Watertown, MA • jesse@tutorintelligence.com

RSS 2026
"It's the Demos" Workshop
Sydney, Australia

100 robots in the fleet

up to **1,000** demo-hours / day

35.8% faster teleoperation

compensation for speed-based multimodality

fast edge-case discovery

Data Factory 1 — 100 *Sonny* bimanual manipulators running continuously in Watertown, MA.

A Data Factory for Robot Manipulation

- **Challenge:** at fleet scale, the many people defining tasks, teleoperating, and annotating strain data quality and consistency
- **DF1:** 100 *Sonny* bimanual robots collecting nonstop for VLA training; **Ti0** trained entirely on DF1 data
- **Operators:** an internationally distributed remote team ("Tutors") — **1:1 control** (one operator, one robot) + **1:N supervision** (one operator, many rollouts)
- **Why engineer quality:** direct communication with one operator held a few hundred demos consistent, but broke at 100 nonstop robots
- **Thesis:** demonstration quality is an *explicit objective at every stage* — a deficiency anywhere caps quality downstream

1. Collection interface (PTeleop) → 2. Low-latency curation → 3. Data preprocessing (VN) → 4. Quality modeling (reward) → 5. Failure discovery (rollouts)

1. The Collection Interface: PTeleop

- 2D + latency → move–wait–move → hesitant, hard-to-imitate motion
- **PTeleop (VR):** live 3D point cloud <50 ms, 1:1 hand↔gripper, haptic grasp
- *Ghost robot* previews the next state → operators stop waiting
- **35.8% faster**, **>10% fewer mistakes**, new operators ready in **6 h**

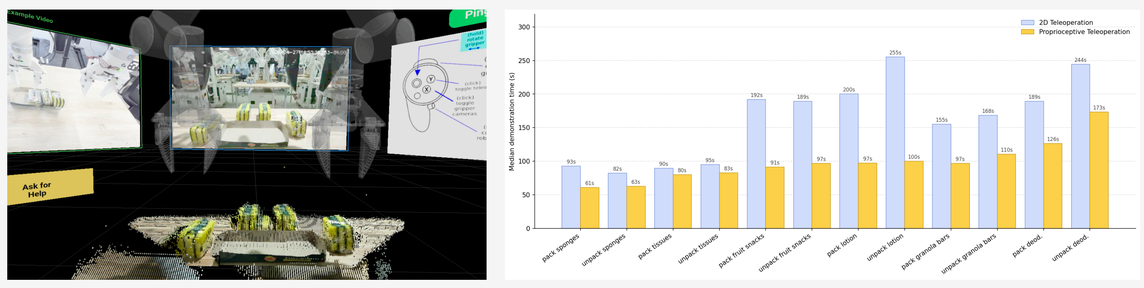

PTeleop VR interface, one 3D frame (left); demonstration time, 2D vs. VR, over 12 tasks (right).

2. Low-Latency Curation

- Flag off-spec demos; preserve in-spec diversity (randomized scenes)
- Continuous human + automated labeling — not an after-the-fact audit
- Every episode scored within **minutes**; **thousands of hours/day**

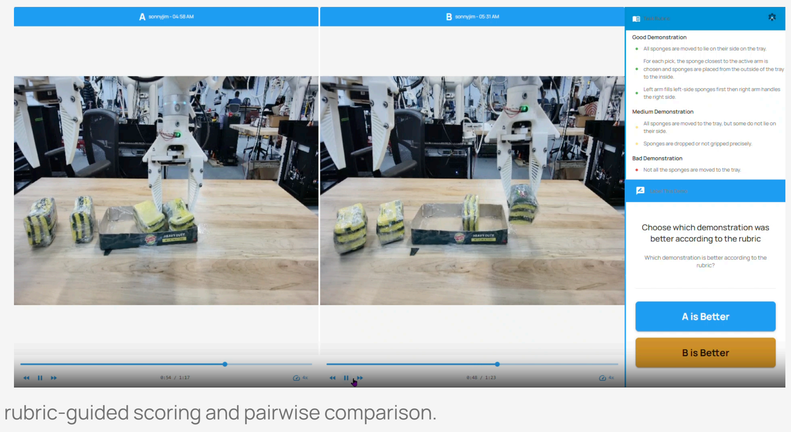

In-house labeling: rubric-guided scoring and pairwise comparison.

3. Data Preprocessing: Velocity Normalization

- **Speed-based multimodality:** same path at 2× speed → conflicting labels
- **Stage 1:** rescale to characteristic speed, clamp $[0.75,\ 1.5]$
- **Stage 2:** smooth monotonic speed map; protect grasps, drop idle
- **62% less** cross-episode speed variance

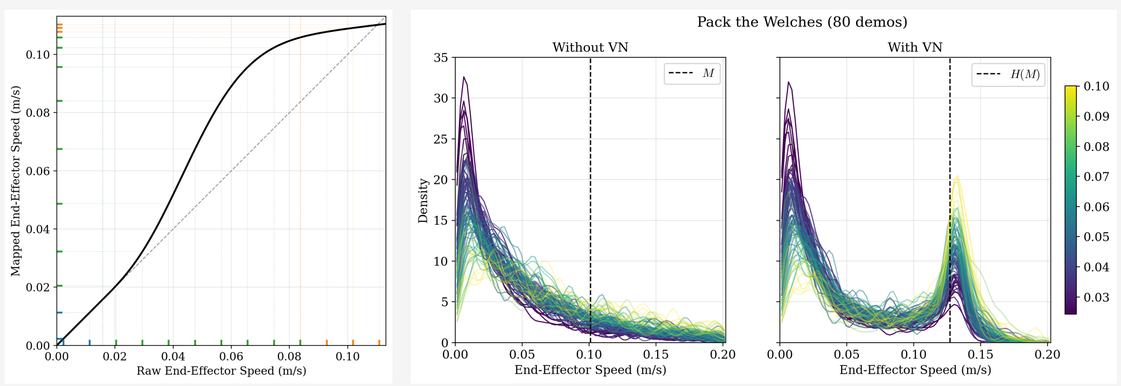

Stage-2 speed map $H(s)$ (left); per-episode speed distributions, before vs. after VN (right).

4. Modeling Demonstration Quality

- Mixed quality → model as a **scalar reward** → reward-weighted imitation
- 5 human label types → terminal $R \in [0, 10]$ → per-frame progress reward
- Applied to **13,282 demonstrations**
- Feedback loop: coach operators, sharpen task specs
- Robometer off-distribution → dense *in-domain* labels needed

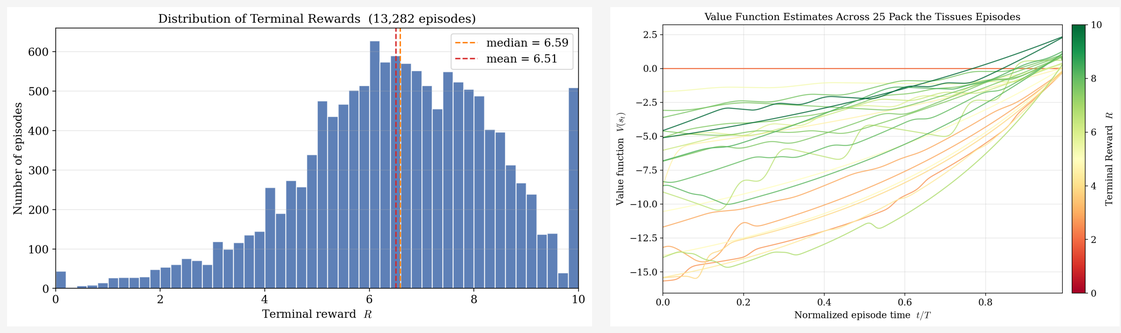

Terminal-reward distribution (left); learned value curves, high- vs. low-quality demos (right).

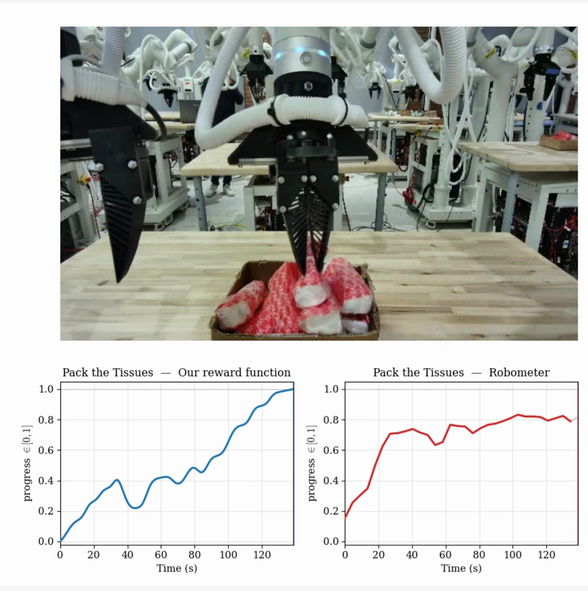

"Pack the tissues": our in-domain reward (left) reaches completion; Robometer (right) plateaus.

5. Discovering Failures at Scale

- Pipeline → >**80% success**; higher needs on-policy corrections
- Edge case: 8 h on 1 robot → ∼**5 min** on 100 (∼**100**× faster)
- 1:N supervision: operator recovers in seconds → corrections train the model

Discussion

- Quality is an **explicit objective at every stage** of the pipeline
- Impact order: PTeleop → spec monitoring → VN + reward → corrections → commodity HW
- Commodity hardware (Meta Quest 3S, NVIDIA 5060 Ti) → international scale
- Open: >1-DoF dexterity · sub-100 ms loops · in-domain progress modeling

Together, a blueprint for producing high-quality real-robot data at scale.

Full technical report: *Data Factory 1 Technical Report*, Tutor Intelligence (2026).

Authors listed in alphabetical order.

Contact: jesse@tutorintelligence.com