

Dexterous Humanoid Hands: Landscape, Benchmarks & How to Reproduce SOTA

Oct 2, 2026 · @julian gilyadov

The status quo to beat is an open, low-cost hand (ORCA or LEAP) trained by glove teleop, GPU sim-to-real RL, or pretraining on human video. The best real results now come from human-data pretraining, while real in-hand reorientation on cheap hands still manages only a few consecutive goals.

To know whether our work improves on this, reproduce four anchors on one hand in about eight weeks: simulated reorientation, real reorientation, glove-teleop imitation learning, and POMDAR plus a human-data baseline. Freeze them in the baseline sheet, then test our data, simulator or environment on the metric that matches it.

The biggest openings: nobody has published autonomous POMDAR scores or ORCA reorientation numbers, and real reorientation on low-cost hands lags simulation badly.

Hardware landscape

Labs train on cheap open hands (LEAP, RUKA, ORCA: about \$200–2,000 in parts), while 20–22 DoF tactile hands (Sharpa, Wuji, Tesla) set the 2026 frontier but publish few comparable benchmarks.

Research hands you can buy or build (sorted by price)

Hand	Released	Actuated DoF	Actuation	Tactile	Price (USD)	Known for
LEAP Hand v2 (CMU, 2025)	Jun 2025 (RSS paper)	8	Hybrid rigid-soft, servos	None reported	~200 parts	Under 2 h assembly
RUKA (NYU, 2025)	Apr 2025 (arXiv)	11 (15 joints)	Tendon, forearm motors	None	~1,300 parts	Learned controllers from MANUS-glove data; Kapandji 10/10; 29 of 33 GRASP-taxonomy grasps

Hand	Released	Actuated DoF	Actuation	Tactile	Price (USD)	Known for
LEAP Hand v1 (CMU, 2023)	Jul 2023 (RSS paper)	16 (4 fingers)	Servos in joints	None	~2,000 parts; 2,066 kit	Most-used low-cost research hand
ORCA (ETH Zurich, 2025)	v1: Apr 2025 (arXiv); v2 line (lite, standard, Touch): Mar 15, 2026; v2.1: no public date found	17 incl. wrist	Tendon, forearm motors	Fingertip FSR; Touch: 351 taxels	~2,000 parts; 3,500+ assembled	Published reliability data, LeRobot stack
Wuji Hand	Sep 18, 2025 (Hand 2 teased Jun 2026)	20	Micro-actuator per joint	None standard	16,000 US list	580 g, 15 N fingertip, 300k+ grasp-cycle rating
Inspire RH56DFX / BFX	2022 (DFX); BFX date not published	6 active (12 joints)	Linear servos	Optional fingertip force	20,599–21,599 US list	Common on Unitree humanoids
Allegro Hand	2012; V4 2018, V5 2024	16 (4 fingers)	Motors in joints, CAN bus	None standard	21,450 US list	Most-cited hand; HORA and DexPoint baselines
Sharpa Wave	Oct 2025 (first shipments)	22	Proprietary	1,000+ taxels per fingertip at 180 Hz	Not disclosed	30 N fingertip force; mass production claimed
Shadow Dexterous Hand	2005	20 (24 joints)	Tendon	Extensive	> 100,000 CHF	OpenAI Dactyl and Rubik's cube

US list prices for LEAP v1, Wuji, Inspire and Allegro are one reseller's (Robotics Center, July 2026); Shadow's cost is from the [ORCA paper](#).

Release dates: research hands use the first paper date; commercial hands use first sale or first shipment ([Wuji](#), [Inspire](#), [Allegro](#), [Sharpa](#), [Shadow](#), [ORCA v2 line](#)). ORCA v2.1 appears in no public release note or the orca_core repository, so its date is open.

Humanoid OEM hands (not sold separately; reported specs as of April 2026, [Wikipedia](#))

Robot hand	Released	DoF	Actuation	Sensing
Tesla Optimus Gen 3	22-DoF hand shown Oct 2024; full Gen 3 robot not revealed as of Apr 2026	22	Tendon; 25 actuators in the forearm	Vision-first; tactile little disclosed
Sanctuary Phoenix	21-DoF hand Dec 2024; Phoenix Gen 7 Apr 2024	20–21	Hydraulic	Micro-barometer pads, 5 mN sensitivity
Figure 02 (Helix 02)	Aug 6, 2024	16	Electric unit per finger	Fingertip tactile (~3 g), palm cameras
Fourier GR-2	Sep 30, 2024	12	Electric	6 tactile arrays per hand
Unitree H2 (Dex5 option)	Oct 20, 2025	10–12	Electric	Optional tactile
Appttronik Apollo (PSYONIC Ability)	Apollo Aug 2023; Ability Hand 2021	Low; prosthesis-derived	Electric	Multi-touch

DoF count is a weak proxy for dexterity. A May 2026 kinematic metric across 16 hands ([KaRMA](#)) ranked LEAP highest and Ability lowest, and found DoF alone does not predict reachable workspace.

Release dates are first public unveilings ([Optimus](#), [Sanctuary hand](#), [Phoenix Gen 7](#), [Figure 02](#), [GR-2](#), [H2](#), [Apollo](#), [Ability Hand](#)). Benchmark and dataset dates elsewhere are first paper or release dates ([Elliott & Connolly](#), [Isaac Lab](#), [EgoDex](#), [ActionNet](#), [RoboMind](#), [RealDex](#), [BrainCo Revo2](#)).

ORCA deep dive

ORCA is the best-documented low-cost anthropomorphic research hand: 17 DoF, tendon-driven, under 2,000 CHF in parts, with published reliability, accuracy and learning results ([ORCA paper](#), [IROS 2025](#)). Since June 2026 it also has an open learning stack that plugs into Hugging Face's LeRobot ([ORCA platform paper](#)).

Design (v1)

- Kinematics: 16 finger joints plus 1 wrist joint. Fingers 2–5 have MCP, PIP and abduction joints (no DIP); the thumb adds a CMC joint and sits at 15° supination for opposition.
- Actuation: each joint is pulled by a flexor/extensor pair of 0.4 mm braided nylon tendons over metal pins. Motors sit in a forearm "tower" with fans; the wrist is belt-driven with 60° of flexion and extension.
- Reliability features: pin joints that pop out instead of breaking, ratchet spools for re-tensioning in seconds, and auto-calibration that drives each joint to its stops to fit a motor-to-joint ratio without joint encoders.
- Sensing: binary force-sensing resistors under silicone skin on all five fingertips.

Published data points

Metric	Result	Conditions
Build	< 2,000 CHF materials, < 8 h assembly by one person	DIY, 3D-printed PLA
Continuous grasping	2,250 grasp cycles in 2.5 h, no failure	Grasp every 4 s, wrist cycle every 16 s; stopped by choice
Durability claim	> 10,000 cycles (~20 h) without hardware failure	Abstract figure
Joint tracking	Accuracy similar to LEAP Hand; mean latency < 0.2 s	2 Hz and 5 Hz sine, AprilTags filmed at 60 fps; latency mostly software
Sim-to-real RL	Zero-shot tennis-ball reorientation about a given axis	IsaacGymEnvs, 4,096 parallel envs, A2C, ~1 h training with domain randomization
Imitation learning	214 demos (~2.5 h) gave a policy that ran 7 h 17 min (~2,000 grasps) with no hardware intervention	Rokoko-glove teleop, Franka arm, 3 cameras, diffusion transformer, 500 epochs ≈ 4 h on one RTX 4090
IL ablation	Masked-cube input beat raw RGB	60 trials, 10 per table sub-area
Teleop skills	Cube stacking, jar-cap twist, fidget-spinner spin, writing "Hello", pouring 50 ml	Qualitative, gloves

Metric	Result	Conditions
Tactile limits	Detects 0.05 N; skin wear after ~2,000–4,000 cycles; sensor wires snapped after ~4,500–7,000	v1 binary FSR design

Buying options (Jul–Oct 2026)

Option	Price (USD)	Contents
BOM kit	643	Tendons, connectors; no motors, prints or silicone
Assembly kit with motors	4,234 per hand	Dynamixel motors, tested prints, cast skin, fans
Fully assembled v1 (2025)	5,929 per hand	Tested, hard case, spares
Assembled ORCA Hand, current listing	from 3,500 (Feetech) or 4,500 (Dynamixel)	Manufacturer-confirmed Jul 31, 2026
ORCA Hand Touch (Mar 2026)	6,100–7,100	351 Hall-effect taxels, 6D force per taxel, ~1.25 kg
9-DoF lite (Mar 2026)	BOM < 900	Coupled tendons, self-build only

The seller is ORCA Dexterity, Inc., a Delaware company with its team in Zurich; no grip-force or payload figure is published ([RoboZaps record](#)).

Software

- [orca_core](#): MIT-licensed Python controller with tension, calibrate and neutral scripts. It auto-detects Dynamixel or Feetech motors and ships v1 and v2 hand configs (~590 stars, 150 forks).
- Teleop: Rokoko gloves and Apple Vision Pro are supported out of the box; the retargeting code is open ([shop page](#)).
- [ORCA learning stack](#) (June 2026): one interface for control, simulation, consumer-device teleop and retargeting. Its reference workflow is VR-headset teleop of in-hand reorientation, a LeRobot-trained policy, and a reproducible evaluation setup.

What ORCA has not published: grip force, payload, YCB-style grasp success rates, or consecutive-success counts for reorientation. Those gaps are exactly where a new data or simulation product can show a measurable gain.

How dexterous hands get their skills

Five training routes dominate in 2026, and the fastest-moving one is human data: egocentric video or wearable capture, converted into robot trajectories and co-trained with a few robot demos.

Route	How it works	Typical cost	Main limitation
Teleop + imitation learning	An operator drives the hand with gloves or a VR headset; a policy (Diffusion Policy, ACT, π 0-style VLA) clones the demos, e.g. ORCA (Apr 2025)	50–200+ robot demos per task	Slow and tied to one hand; tracking quality caps data quality (gloves beat headsets under occlusion)
Sim-to-real RL	Train in a GPU simulator with domain randomization, deploy zero-shot, e.g. MuJoCo Playground (Feb 2025)	Minutes to days of GPU time	Needs object-pose tracking and reward design; real results trail simulation
Human video to robot	Retarget hand poses from egocentric video into robot trajectories, then pretrain, e.g. UniDex (Mar 2026), EgoScale (Feb 2026)	Thousands of hours of video	Kinematic and visual gaps; still needs some robot data
Human–robot co-training	Wearable rigs (gloves, head cameras) record humans; a few robot demos anchor the embodiment, e.g. DexWild (May 2025)	Human demos collect several times faster than robot demos	Success collapses without any robot demos

Route	How it works	Typical cost	Main limitation
World models and dexterous VLAs	Pretrain dynamics or policies on mixed human and robot data, then plan or fine-tune per hand, e.g. DexWM (Dec 2025)	Large pretraining, small fine-tune	Cross-hand transfer is still weak

Two caveats recur across routes: human data alone does not finish the job, and policies rarely transfer across hands without retraining. The State of the art section below gives the numbers.

Where hands are used today

- Research benchmarks: in-hand reorientation, grasping, tool use (scissors, spray bottles, kettles) and in-hand assembly.
- Humanoid demos: Sharpa’s robot assembled a paper windmill in 30+ autonomous steps at CES 2026 ([report](#)). OEMs publish videos far more often than success rates.
- Data collection: teleop rigs (ORCA supports Rokoko gloves and Apple Vision Pro) feed imitation-learning datasets.
- Prosthetics-derived hands: PSYONIC Ability (2021) and BrainCo Revo2 (Sep 2025) hands carry prosthetic designs onto humanoids.

Benchmarks and metrics

There is no single dexterity leaderboard. A credible claim stacks three levels: a hand-level benchmark (POMDAR), a simulation suite (MuJoCo Playground), and a real-world policy protocol with fixed trial counts and reported variance.

Hand-level benchmarks (what the hand plus controller can do)

Benchmark	Measures	Protocol and score	Status
POMDAR (ETH Zurich, Apr 2026)	12 in-hand manipulation and 6 grasp tasks from the Elliott & Connolly and GRASP taxonomies	3D-printed scaffolds; score = $0.8 \times \text{correctness} + 0.2 \times \text{speed}$ vs a human baseline; 20 trials per task; MuJoCo twin	ORCA scored by teleop only (1,140 trajectories, ~25 h); gloves beat Apple Vision Pro on occluded tasks. No autonomous-policy scores yet
Elliott & Connolly benchmark (CMU, Jul 2021)	13 in-hand patterns using the digits only	YCB objects, visual tracking	No grasp tasks
HD-marks (2020)	50 tasks: GRASP grasps, Kapandji thumb postures, in-hand axes	Mostly binary success	Broad, weak cross-lab comparability
Kapandji test (1986) + GRASP taxonomy (2016)	Thumb opposition (0-10) and 33 grasp types	Static postures	RUKA: 10/10 and 29 of 33
KaRMA (May 2026)	Kinematic workspace for fine manipulation	Computed from hand models	16 hands ranked
DexBench (RLWRLD, Jun 2026)	18 atomic industrial tasks in ~80 cases, placed on object-complexity and dexterity-regime axes	Real-world, customer-derived	Endorsed by Lotte, SK Telecom, CJ Logistics and others

The Elliott & Connolly and HD-marks rows are as summarized in the POMDAR paper.

Simulation suites

Suite	Contents	Headline metric
MuJoCo Playground (Feb 2025)	LeapCubeReorient and LeapCubeRotateZAxis; open source, pip-installable	Consecutive successes; ~35 min to train on one RTX 4090
Isaac Gym (2021) / Isaac Lab (Jun 2024)	GPU RL used by DeXtreme and by ORCA's ball-rotation policy	Consecutive successes
Bench2Dex (Sep 2026)	26 bimanual visuo-tactile tasks, 12 hands, ~1.3K teleop demos	Success, with mandatory seeds, episode counts and termination rules
DexVerse (Jul 2026)	Modular multi-task, multi-embodiment suite	Not reviewed here
LabDex (Aug 2026)	Hierarchical chemistry-lab skills, simulation and real robot	Skill-to-long-horizon scaling
In-hand assembly (Sep 2026)	Two-part assembly inside one hand, Isaac Sim	Goal-reaching error by hand morphology

Real-world policy metrics worth reporting

Metric	Used by	Definition
Consecutive successes	DeXtreme, MuJoCo Playground	Goals reached before a drop; median and mean over ≥ 10 trials
Task progress + final success	UniDex	Mean stage completion and full-task success over 20 trials
Unseen-scene success	DexWild	Success in environments absent from training
Zero-shot cross-hand transfer	UniDex	Task progress on a hand never trained on
Human-robot exchange rate	UniDex	Human demos needed to replace one robot demo ($\approx 2:1$)
Autonomy hours	ORCA	Hours or cycles without hardware intervention

Metric	Used by	Definition
Contact fidelity	TactiDex (Jul 2026)	Agreement of contacts and forces with human demos

Datasets

Dataset	Size	Content
EgoDex dataset (Apple, May 2025)	829 h of 1080p video	Egocentric human manipulation with hand-pose annotations
UniDex-Dataset (Mar 2026)	52K trajectories, 9M frames, 8 robot hands	Retargeted from H2O, HOI4D, HOT3D and TACO
ActionNet (Fourier, Mar 2025)	30K trajectories, 2 hands	Dexterous bimanual teleop
RoboMind (Dec 2024)	19K trajectories, 1 hand	Multi-embodiment teleop
RealDex (Feb 2024)	2K trajectories, 2 hands	Human-like grasping

ActionNet, RoboMind and RealDex sizes are from the UniDex paper's comparison table.

State of the art

The strongest 2026 results come from human-data pretraining: UniDex-VLA doubled π_0 's real tool-use success, and EgoScale showed policy quality scales predictably with hours of egocentric video. Real in-hand reorientation on low-cost hands remains weak.

Date	Work	Hand	Result
Sep 2026	In-hand assembly	Sharpa Wave; Wuji, Allegro, XHand in sim	Sharpa and Wuji performed best; Allegro and XHand failed deep insertion
Jul 2026	UHAS	LEAP	Real cube reorientation averaged 2.0 consecutive goals; multi-hand training scored lower than single-hand

Date	Work	Hand	Result
Jun 2026	ORCA learning stack	ORCA	First open, LeRobot-native loop: VR teleop, policy training, reproducible evaluation (no headline number)
Mar 2026	UniDex-VLA	Inspire, Wuji	81% task progress and 76% final success vs $\pi 0$ at 38% and 35%, with 50 demos per task; zero-shot transfer reached 40% on Wuji and 60% on Oymotion
Feb 2026	EgoScale	22-DoF hand	Pretrained on 20,854 h of human video with a log-linear scaling law; average success rate up 54% over no pretraining after ~50 h aligned human and 4 h robot data
Jan 2026	Sharpa North	Sharpa Wave	Paper windmill assembled in 30+ autonomous steps (demo, no success rate)
Dec 2025	DexWM	Allegro	10 of 12 real grasps (~83%) without robot fine-tuning
May 2025	DexWild	Multiple	68.5% success in unseen environments, ~4× robot-only; 5.8× better cross-embodiment generalization
Feb 2025	MuJoCo Playground	LEAP	Zero-shot cube reorientation: median 3.5, mean 7.1 consecutive goals over 10 trials (best 27)
Oct 2022	DeXtreme	Allegro	Vision-based reorientation: best run 112 consecutive goals; mean 23.1 when capped at 50; 60 h of training


Real in-hand reorientation varies about 10× by setup: DeXtreme averaged 23 goals on Allegro, while recent low-cost LEAP results average 2–7. Humanoid OEMs (Tesla, Figure, Sanctuary, Sharpa) publish demos rather than success rates, so their capability cannot be benchmarked from outside.

Reproduction plan

Reproduce four anchors in about eight weeks with one ORCA or LEAP hand, one arm and one RTX 4090-class GPU: simulated reorientation, real reorientation, glove-teleop imitation learning, and a human-data baseline. Freeze the protocol before testing our own work.

Four gated phases come before testing our own work

Reproduction plan, about eight weeks



reproduction roadmap · 4 phases, 4 gates

Each gate must pass before the next phase starts; the frozen sheet is what our own work is measured against.

Phase 0: simulation only (week 1, no hardware)

1. `pip install playground`, then train `LeapCubeReorient` with Brax PPO using the published settings (100M steps, 8,192 envs). Expect ~35 min on one RTX 4090; run 3 seeds ([MuJoCo Playground](#)).
2. Log simulated consecutive goals at 0.1 rad tolerance and wall-clock time to a fixed reward.
3. Load ORCA in MuJoCo through the [ORCA learning stack](#) and run the same task, so later comparisons are hand-matched.
4. Install POMDAR's MuJoCo version ([project page](#)) and teleoperate a few tasks to validate the pipeline.

Gate 0: reward curves and wall-clock match the paper within about 20%.

Phase 1: real in-hand reorientation (weeks 2–3, no arm needed)

5. Build the hand. ORCA assembles in under 8 h; then run `scripts/tension.py`, `scripts/calibrate.py` and `scripts/neutral.py` from [orca_core](#). LEAP v1 takes about 3 h.
6. Copy the Playground rig: palm tilted 20° down on an 80/20 frame, one RealSense D415 overhead, a 7 cm cube, policy at 20 Hz, and DeXtreme's cube-pose detector or AprilTags.
7. Deploy zero-shot. Run 10 trials, ending each when the cube drops or stalls for 30 s. Report every trial plus median and mean consecutive goals at 0.4 rad tolerance.

Gate 1: LEAP matches or beats the published LEAP numbers in State of the art.

Phase 2: teleop and imitation learning (weeks 3–6, arm required)

8. Mount the hand on a 7-DoF arm (Franka-class; ORCA uses an ISO 9409-1 flange) with two external cameras and one wrist camera.
9. Teleoperate with Rokoko gloves, calibrating retargeting each session. Apple Vision Pro also works but [scores lower](#) when fingers are occluded.
10. Collect ~200 demos of ORCA's self-resetting cube pick-and-place (about 2.5 h), plus 50 demos each for two UniDex-style tool tasks, such as kettle pouring and spray-bottle pressing.
11. Train with LeRobot: Diffusion Policy and ACT as baselines, plus a $\pi 0$ -class VLA fine-tune. Budget ~4 h per policy on one RTX 4090.
12. Evaluate pick-and-place over 60 trials across 6 table zones ([ORCA protocol](#)), and each tool task over 20 trials with stage-wise progress ([UniDex protocol](#)). Add one multi-hour autonomy run.

Gate 2: multi-hour autonomy without hardware intervention, and tool-task results within about 10 points of the published baselines.

Phase 3: hand-level score and human-data baseline (weeks 6–8)

13. Print the POMDAR rig. Score teleop first (20 trials per task) to calibrate against ETH's published ORCA plots, then score autonomous policies, which no one has published yet.
14. Build a UniDex-Cap-style capture rig (Apple Vision Pro plus a RealSense L515 on a printed mount), or start from public EgoDex or UniDex-Dataset data.
15. Co-train human and robot demos on one task. Sweep the human-demo count at fixed robot-demo counts to measure the exchange rate.

Gate 3: a frozen baseline sheet in which every published number is reproduced or its gap explained.

What you need

- One ORCA (17 DoF, tendon) or LEAP v1 (16 DoF, servos-in-joints) hand; buying both gives a morphology check. Prices are in the sections above.
- A 7-DoF arm with an ISO 9409-1 flange for Phases 2–3; Phase 1 needs only a frame.
- One RTX 4090-class GPU; every baseline above trains on one.
- RealSense D415 or D435 for in-hand tracking; L515 for UniDex-style RGB-D capture.
- Rokoko or MANUS gloves; Apple Vision Pro optional.
- Open question: arm, glove and camera prices are not covered here.

Evaluation protocol: proving our work beats the status quo

A claim holds only if, on the frozen protocol above, our data, simulator or environment moves a reproduced baseline number at equal or lower cost, with enough trials to separate signal from noise.

Pick the metric that matches the product

If our work is	Baseline to beat	Primary metric	Win condition
Glove or wearable data	DexWild co-training; UniDex-Cap	Human demos needed per robot demo; unseen-scene success	Fewer human demos per robot demo than UniDex's ratio, or higher unseen-scene success at equal robot demos
Egocentric video data	EgoScale scaling curve; EgoDex; UniDex-Dataset	Real task success after pretraining on matched hours; loss-vs-hours slope	Equal success from fewer hours, or a steeper scaling curve
Simulator or sim-to-real pipeline	MuJoCo Playground LEAP; DeXtreme	Real consecutive goals; sim-to-real drop; GPU-hours	Higher real median and mean at equal or fewer GPU-hours
Training or evaluation environments	POMDAR-sim; Bench2Dex	Rank agreement between sim and real scores across at least 5 policies	Higher rank correlation with real outcomes than existing suites

Rules for a fair comparison

- Pre-register tasks, objects, success definitions and trial counts before any run.
- Change one variable: same hand, arm, cameras, policy architecture and compute; only the data or environment differs.
- Randomize and blind: interleave policies A and B in random order, and keep the person resetting scenes unaware of which is running.
- Size trials for the effect: 20 trials leave about ± 22 points of 95% uncertainty at 50% success. Separating 50% from 70% with 80% power takes about 90 trials per arm.
- Report distributions: per-trial values plus median and mean, because consecutive-goal metrics are heavy-tailed.
- Split in-distribution results from held-out objects, scenes and hands.
- Normalize by cost: report success per hour of data collection and per GPU-hour, since human demos collect faster than robot demos.
- Train at least 3 seeds, and publish seeds, episode counts, termination rules and horizons, as Bench2Dex requires.

Baseline sheet

Measure	Published reference	Reproduced	Ours	Δ (95% CI)	Status
LEAP real reorientation, median / mean consecutive goals	MuJoCo Playground (2025)				Not started
ORCA real reorientation, median / mean consecutive goals	None published				Not started
ORCA pick-and-place success, 60 trials	ORCA (2025)				Not started
Autonomy hours without hardware intervention	ORCA (2025)				Not started
Tool-task progress and final success, 20 trials per task	UniDex (2026)				Not started
POMDAR score, teleop	POMDAR (2026)				Not started
POMDAR score, autonomous policy	None published				Not started
Human demos per robot demo	UniDex (2026)				Not started
GPU-hours to train reorientation	MuJoCo Playground (2025)				Not started

The two "None published" rows are open territory: the first credible autonomous POMDAR and ORCA reorientation numbers would themselves define the status quo.

Sources

Opened in full:

- [ORCA paper \(arXiv 2504.04259v2\)](#)
- [ORCA platform paper \(arXiv 2606.14561\)](#) and its [Hugging Face page](#)
- [orca_core repository](#)
- [RoboZaps ORCA Hand record](#)
- [Robotics Center dexterous hand comparison, July 2026](#)
- [Wikipedia: Humanoid hand](#)
- [Sharpa North at CES 2026](#)
- [UniDex \(arXiv 2603.22264\)](#)

- [MuJoCo Playground \(arXiv 2502.08844\)](#)
- [POMDAR \(arXiv 2604.09294\)](#)
- [WEF on DexBench, June 2026](#)

Cited from search excerpts of the source pages:

- ORCA shop: [fully assembled](#), [assembly kit](#), [BOM kit](#)
- [LEAP Hand v2 \(RSS 2025\)](#), [RUKA \(arXiv 2504.13165\)](#), [KaRMA \(arXiv 2605.15548\)](#)
- [DexWild \(arXiv 2505.07813\)](#), [DexWM \(arXiv 2512.13644\)](#), [EgoScale \(arXiv 2602.16710\)](#), [DeXtreme \(arXiv 2210.13702\)](#)
- [UHAS \(arXiv 2607.03570\)](#), [In-hand assembly \(arXiv 2609.10137\)](#)
- [Bench2Dex \(arXiv 2609.15726\)](#), [DexVerse \(arXiv 2607.08751\)](#), [LabDex \(arXiv 2608.18618\)](#), [TactiDex \(arXiv 2607.09190\)](#)