G1 Moves

A Motion Capture Dataset for Unitree G1 Humanoid Robot

Experiential Technologies

March 2026

60	29	30.3	60
Motion Clips	Degrees of Freedom	Minutes of Motion	Trained Policies

ABSTRACT

G1 Moves is a comprehensive motion capture dataset containing 60 clips spanning dance, martial arts, and acrobatic motions, captured from human performers and retargeted to the Unitree G1 humanoid robot (29 DOF). Each clip is processed through a four-stage pipeline: markerless motion capture via MOVIN TRACIN (or monocular video via PromptHMR), inverse-kinematics retargeting to robot joint space, MuJoCo forward-kinematics simulation for body-state generation, and PPO reinforcement learning across 8,192 parallel MuJoCo-Warp environments. All 60 clips have fully trained policies achieving <17 cm mean position error and <8 deg mean orientation error, exported as both PyTorch checkpoints and ONNX models for edge deployment. The dataset totals 106,434 frames (30.3 minutes) at 60 FPS and includes an interactive browser-based viewer using MuJoCo WASM and ONNX Runtime Web for real-time policy visualization.

`Dataset:` `huggingface.co/datasets/exptech/g1-moves`
`Code:` `github.com/experientialtech/g1-moves`
`Space:` `huggingface.co/spaces/exptech/g1-moves`

Generated with Claude Code (claude.ai/code)

1. INTRODUCTION

Teaching humanoid robots to perform expressive whole-body motions remains a central challenge in embodied AI. While reinforcement learning (RL) has shown promise for locomotion, the gap between simulated training and real-world deployment is compounded by the scarcity of high-quality, robot-specific motion datasets with corresponding trained policies.

G1 Moves addresses this gap by providing an end-to-end, open dataset and pipeline for the Unitree G1 humanoid robot. The dataset spans three motion categories--dance (28 clips), karate (27 clips), and bonus motions (5 clips)--captured from professional performers using markerless motion capture. Each clip flows through a reproducible four-stage pipeline: capture, retarget, train, and deploy.

Key contributions include: (1) a standardized pipeline from human mocap to robot-specific joint trajectories with automated self-collision correction; (2) detailed reward function specification enabling reproducible RL training; (3) quality-validated policies for all 60 clips with convergence analysis; (4) ONNX-exported models for browser-based and edge inference; and (5) an interactive web viewer for real-time policy visualization without installation.

2. DATASET OVERVIEW

2.1 Summary Statistics

Property	Value
Total clips	60 (59 unique + 1 remapping)
Total frames	106,434
Total duration	30.3 minutes
Frame rate	60 FPS (59 clips), 30 FPS (1 clip)
Robot platform	Unitree G1 (Edition EDU, 29 DOF)
Performers	5 (see Section 12)
Trained policies	60/60 (100%)
Export formats	BVH, PKL, CSV, NPZ, PT, ONNX

2.2 Category Breakdown

Category	Clips	Frames	Duration	Avg Length
Dance	28	59,344	16.5 min	35.5 s
Karate	27	42,339	11.8 min	26.2 s
Bonus	5	4,751	1.4 min	17.0 s
Total	60	106,434	30.3 min	30.4 s

2.3 Robot Platform

The Unitree G1 is a full-size humanoid robot standing approximately 1.3 m tall. The dataset targets Edition EDU with 29 actuated degrees of freedom: 6 per leg (hip pitch/roll/yaw, knee, ankle pitch/roll), 3 in the waist (yaw, roll, pitch), and 7 per arm (shoulder pitch/roll/yaw, elbow, wrist roll/pitch/yaw).

2.4 Joint Configuration (29 DOF)

Index	Group	Joints
0-5	Left leg	hip_pitch, hip_roll, hip_yaw, knee, ankle_pitch, ankle_roll
6-11	Right leg	hip_pitch, hip_roll, hip_yaw, knee, ankle_pitch, ankle_roll
12-14	Waist	yaw, roll, pitch
15-21	Left arm	shoulder_pitch/roll/yaw, elbow, wrist_roll/pitch/yaw
22-28	Right arm	shoulder_pitch/roll/yaw, elbow, wrist_roll/pitch/yaw

3. MOTION CAPTURE & RETARGETING

3.1 Capture Systems

Primary capture uses MOVIN TRACIN, a markerless motion capture system combining LiDAR and computer vision to track full-body motion at 60 FPS. The system outputs BVH files with a 51-joint humanoid skeleton. No markers, suits, or calibration targets are required--performers move freely in an open space.

An alternative pipeline, video2robot, extracts 3D human pose from monocular video using PromptHMR (SMPL-X body model) and retargets via GMR inverse kinematics. This enables motion extraction from smartphone video, YouTube content, or AI-generated footage.

3.2 Retargeting Pipeline

Human BVH skeleton (51 joints) is mapped to the G1's 29-DOF joint space via per-frame inverse kinematics using the MOVIN SDK Python library. The retargeting process assumes a standard human height of 1.75 m and solves for joint angles that best match the human pose within G1 joint limits.

Ground calibration is applied post-retargeting: MuJoCo forward kinematics computes the minimum ankle Z across all frames, and the root position is shifted downward by this offset to ensure consistent ground contact throughout the motion.

3.3 Self-Collision Correction

The video2robot pipeline does not enforce self-collision avoidance during IK. A post-processing step detects collision frames using MuJoCo contact detection and resolves them by interpolating arm joint angles (indices 15-28) between nearest clean boundary frames. If arm-only interpolation is insufficient, full-joint interpolation is applied iteratively (up to 3 passes). This typically corrects collisions with <5 deg mean deviation from the original motion.

3.4 Output Format: PKL

Key	Shape	Type	Description
fps	(1,)	int	Frame rate (60)
root_pos	(N, 3)	float32	Root pelvis position in world frame [m]
root_rot	(N, 4)	float32	Root pelvis rotation quaternion (xyzw)

dof_pos	(N, 29)	float32	Joint angles [radians]

4. TRAINING DATA GENERATION

4.1 CSV Intermediate Format

Retargeted PKL data is exported to a 36-column CSV (no header): 3 columns for root position (x, y, z), 4 columns for root quaternion (x, y, z, w), and 29 columns for joint angles. Values are stored at 10 decimal places of precision. This format enables manual inspection and frame-range selection.

4.2 NPZ Generation via MuJoCo Forward Kinematics

The csv_to_npz script loads joint trajectories, optionally resamples via cubic spline interpolation (positions) and SLERP (quaternions) to the target frame rate, then runs MuJoCo forward kinematics to compute full body state: positions, quaternions, linear velocities, and angular velocities for all 30 rigid bodies in the G1 model. Velocities are computed via torch.gradient() with the appropriate timestep.

4.3 NPZ Array Dictionary

Key	Shape	Type	Description
fps	(1,)	int	Output frame rate
joint_pos	(N, 29)	float32	Joint positions [radians]
joint_vel	(N, 29)	float32	Joint velocities [rad/s]
body_pos_w	(N, 30, 3)	float32	30 body positions, world frame [m]
body_quat_w	(N, 30, 4)	float32	30 body quaternions (wxyz), world frame
body_lin_vel_w	(N, 30, 3)	float32	30 body linear velocities [m/s]
body_ang_vel_w	(N, 30, 3)	float32	30 body angular velocities [rad/s]

5. REINFORCEMENT LEARNING

5.1 Framework

Policies are trained using Proximal Policy Optimization (PPO) via the RSL-RL library, within the mjlab framework. mjlab combines Isaac Lab's manager-based environment API with MuJoCo-Warp for GPU-accelerated physics simulation, enabling 8,192 parallel environments on a single GPU.

5.2 Observation Space (160-dimensional)

The actor observation vector consists of 160 dimensions, constructed each timestep from simulation state and reference motion data:

Component	Dim	Description
Reference joint positions	29	Target joint angles from NPZ frame t
Reference joint velocities	29	Target joint velocities from NPZ frame t

Motion anchor position (body frame)	3	Torso pos. delta, rotated to body frame
Motion anchor orientation (6D)	6	Relative rotation as 2 columns of rotation matrix
Base linear velocity (body frame)	3	Velocimeter reading at IMU location
Base angular velocity (body frame)	3	Gyroscope reading
Joint position (relative to default)	29	q_actual - q_default
Joint velocity	29	dq/dt from simulation
Last action	29	Previous control signal a_{t-1}
Total	160	

The critic receives an extended 430-dimensional observation that additionally includes 30 body positions (90 dims) and 30 body orientations as 6D rotation representations (180 dims), all in body frame.

5.3 Observation Noise (Training Only)

During training rollouts, uniform noise is injected into the actor observation to encourage robustness. The critic always receives clean (noise-free) observations.

Component	**Noise Range**	**Units**
Motion anchor position	U[-0.25, 0.25]	m
Motion anchor orientation	U[-0.05, 0.05]	rad
Base linear velocity	U[-0.5, 0.5]	m/s
Base angular velocity	U[-0.2, 0.2]	rad/s
Joint position	U[-0.01, 0.01]	rad
Joint velocity	U[-0.5, 0.5]	rad/s

5.4 Action Space (29-dimensional)

The policy outputs a 29-dimensional continuous action vector, representing desired joint position offsets from the default standing pose. Actions are clipped to [-10, 10] and scaled per-joint before application:

```
joint_target[i] = q_default[i] + clip(action[i], -10, 10) * scale[i]
```

Per-joint action scales range from 0.0745 rad (wrist joints, ~4.3 deg) to 0.5475 rad (hip/waist joints, ~31.4 deg). The policy runs at 50 Hz (decimation factor of 4 over the 200 Hz physics timestep), meaning each action is held constant for 4 MuJoCo substeps of 0.005 s each.

6. REWARD FUNCTION

The tracking task uses 9 reward terms combining exponential tracking rewards with regularization penalties. Tracking rewards use a Gaussian kernel: r = w * exp(-(error / std)^2), where higher fidelity yields reward closer to the weight w.

6.1 Tracking Rewards

Term	**Weight**	**Std**	**Error Metric**
Global root position	0.5	0.3 m	L2 distance of anchor pelvis to reference
Global root orientation	0.5	0.4 rad	Quaternion geodesic distance
Body position (14 bodies)	1.0	0.3 m	Mean L2 of body positions in body frame

Body orientation (14 bodies)	1.0	0.4 rad	Mean quaternion distance in body frame
Body linear velocity	1.0	1.0 m/s	Mean L2 of body linear velocities
Body angular velocity	1.0	3.14 rad/s	Mean L2 of body angular velocities

6.2 Regularization Penalties

Term	Weight	Formulation
Action rate L2	-0.1	||action_t - action_{t-1}||^2 (smoothness)
Joint limit violation	-10.0	Count of joints exceeding G1 limits
Self-collision	-10.0	Binary: 1 if any non-ground contact pair

The total reward is the sum of all terms. Converged policies typically achieve total episode rewards of ~35-40, with tracking terms contributing positively and penalties near zero.

7. NETWORK ARCHITECTURE

7.1 Actor Network (~142k parameters)

```
Input (160) -> Linear(512) + ELU -> Linear(256) + ELU -> Linear(128) + ELU -> Linear(29)

Output: action mean (29-dim)
Noise:  log_std parameter (29-dim), initialized at log(1.0)
Action: sampled from N(mean, exp(log_std)^2) during training
```

7.2 Critic Network (~298k parameters)

```
Input (430) -> Linear(512) + ELU -> Linear(256) + ELU -> Linear(128) + ELU -> Linear(1)

Output: scalar value estimate V(s)
```

7.3 Training Details

Hyperparameter	Value
Activation function	ELU (alpha=1.0)
Weight initialization	PyTorch default (Kaiming uniform)
Observation normalization	Running mean/std (actor and critic)
Optimizer	Adam (beta1=0.9, beta2=0.999)
Learning rate	5e-4, adaptive (KL-based annealing)
Target KL divergence	0.01
PPO clip parameter	0.2
Discount factor (gamma)	0.99
GAE lambda	0.95
Entropy coefficient	0.005
Max gradient norm	1.0

Rollout steps per env	24
Mini-batches	4
Epochs per update	5

8. TERMINATION CONDITIONS

Condition	Threshold	Type	Description
Time out	10.0 s (500 steps)	Natural	Episode ends after full duration
Anchor position	depth > 0.25 m	Early	Pelvis too far below reference
Anchor orientation	error > 0.8 rad (46 deg)	Early	Torso too rotated from reference
Body position	any limb Z < -0.25 m	Early	Limb penetrates ground plane

Converged policies achieve near-100% time-out rates, meaning the robot completes the full 10-second motion episode without triggering any early termination condition.

9. DOMAIN RANDOMIZATION

Domain randomization is applied once at environment startup (per-environment, not per-episode) to promote sim-to-real transfer robustness:

Parameter	Range	Description
Center of mass offset (x, y)	U[-2.5, 2.5] cm	Lateral COM perturbation
Center of mass offset (z)	U[-5.0, 5.0] cm	Vertical COM perturbation
Joint encoder bias	U[-0.01, 0.01] rad	Simulates encoder miscalibration
Foot friction coefficient	U[0.3, 1.2]	Surface friction variation

Additionally, the motion command applies per-timestep perturbations to the reference trajectory: position jitter of +/-5 cm (x,y) and +/-1 cm (z), rotation jitter up to +/-0.2 rad (yaw) and +/-0.1 rad (roll/pitch), and joint target perturbation of +/-0.1 rad.

10. TRAINING RESULTS

10.1 Training Configuration

Parameter	Value
Parallel environments	8,192
GPU	NVIDIA RTX PRO 6000 (96 GB VRAM)
Physics throughput	~174,000 FPS
Iterations (short clips <10s)	15,000 (1.5-2.5 hours)
Iterations (medium clips 10-25s)	20,000 (2.5-4 hours)
Iterations (long clips >25s)	30,000 (3-5 hours)
Early stopping criterion	time_out ratio >= 0.95 for 3 consecutive checks

10.2 Convergence Metrics (B_DadDance, Representative)

Metric	Value	Interpretation
Total episode reward	36.9	Excellent convergence
Mean episode length	481.3 steps (9.6s)	Near-maximum duration
Anchor position error	0.167 m	Pelvis tracking accuracy
Anchor rotation error	0.129 rad (7.4 deg)	Pelvis orientation accuracy
Mean body position error	0.042 m (4.2 cm)	Limb position tracking
Mean body rotation error	0.143 rad (8.2 deg)	Limb orientation tracking
Time-out ratio	1.0 (100%)	All episodes complete successfully
Action noise std	0.169	Well-controlled exploration
Final learning rate	0.00011	Adapted down from 0.0005
Self-collisions	0	No self-contact violations

10.3 Training Progression

Training follows a characteristic three-phase progression. In the first few hundred iterations, the untrained policy produces immediate falls (episode length ~1 step, reward ~0). Within 1,000 iterations, basic balance emerges and episode length increases rapidly. By 5,000-10,000 iterations, tracking accuracy converges and the policy maintains upright motion for the full episode. The remaining iterations refine motion quality, reduce jitter, and improve dynamic tracking of fast movements.

11. QUALITY ASSURANCE

All clips undergo automated validation covering five categories:

Check	Pass Rate	Details
Joint limit violations	57/59	B_Fence1 (17 frames), B_Fence2 (35 frames) minor ankle_roll
Ground penetration	59/59	Minimum foot Z >= 3.1 cm across all clips
Frame consistency	59/59	PKL frames match BVH; NPZ = BVH - 1 (expected)
NaN values	59/59	No NaN in any PKL or NPZ array
File completeness	59/59	All pipeline outputs present (.pkl, .csv, .npz, videos)

12. DEPLOYMENT

12.1 ONNX Export

All trained policies are exported to ONNX format for cross-platform inference. The ONNX model takes a 160-dimensional float32 observation vector as input and outputs a 29-dimensional action vector. Observation normalization (running mean/std) is baked into the model, requiring no external state.

12.2 Browser-Based Viewer

An interactive web viewer runs MuJoCo physics via WebAssembly (mujoco-js) and policy inference via ONNX

Runtime Web. The viewer constructs the full 160-dimensional observation from simulation state, runs the policy at 50 Hz, and renders via Three.js with orbit camera controls. No installation required.

12.3 Physical Robot Deployment

Sim-to-real transfer uses the RoboJuDo framework, which provides a modular architecture separating controller (joystick/motion sequence), environment (MuJoCo sim or real robot via Unitree SDK), and policy (ONNX/TorchScript). The real robot communicates over Ethernet (192.168.123.x subnet) at 50 Hz.

13. EQUIPMENT

System	Component	Specification
Training workstation	GPU	NVIDIA RTX PRO 6000 Blackwell (96 GB VRAM)
	CPU	Intel Core Ultra 9 285K
	RAM	128 GB DDR5
	Platform	Dell Pro Max Tower T2, Ubuntu 24.04
Motion capture	System	MOVIN TRACIN (markerless, LiDAR + vision)
	Software	MOVIN Studio
	Output	BVH (51-joint, 60 FPS)
Robot	Model	Unitree G1 (Edition EDU)
	DOF	29 actuated joints
	Interface	Ethernet (Unitree SDK)
	Onboard compute	Jetson Orin

14. DATA FORMATS REFERENCE

Format	Stage	Columns/Keys	Typical Size
BVH	Capture	51-joint skeleton hierarchy, N frames @ 60 FPS	~3 MB
PKL	Retarget	fps, root_pos(N,3), root_rot(N,4), dof_pos(N,29)	~360 KB
CSV	Intermediate	36 cols: 3 pos + 4 quat + 29 joints, no header	~850 KB
NPZ	Training	7 arrays: joint/body pos/vel/quat (see Sec 4.3)	~4.5 MB
PT	Policy	model_state_dict, optimizer, iter, infos	~1.1 MB
ONNX	Deploy	Input: (1,160) float32, Output: (1,29) float32	~600 KB

15. PERFORMERS & CREDITS

Role	Name	Contribution
Director + Development	Mitch Chaiet	Pipeline design, training, 5 bonus dance clips
DIT	Molly Maguire	Data management
Dance performer	Jasmine Coro	24 dance clips (J_Dance, J_ShortDance series)
Karate performer	Mike Gassaway	20 karate clips (M_Move, M_ShortMove series)

Fencing performer	Maya Coro	B_Fence1, B_Fence2
Deployment	Joe DiPrima	Physical robot deployment
Deployment	John DiPrima	Physical robot deployment

16. COMPUTATIONAL COST

16.1 Per-Clip Processing Time

Stage	Time per Clip	Hardware
Retarget (BVH to PKL)	~5 seconds	CPU (4 workers)
NPZ generation (FK sim)	~10 seconds	CPU + MuJoCo
RL training (short clip)	1.5-2.5 hours	RTX PRO 6000, 8192 envs
RL training (medium clip)	2.5-4 hours	RTX PRO 6000, 8192 envs
RL training (long clip)	3-5 hours	RTX PRO 6000, 8192 envs
ONNX export	~5 seconds	CPU

16.2 Total Dataset Cost

The complete dataset of 60 clips required approximately 180 GPU-hours on the RTX PRO 6000 for RL training alone. The batch training pipeline ran continuously over 7 days with early stopping, achieving 100% policy coverage. Total storage (all formats) is approximately 580 MB per clip across all pipeline stages, or ~34 GB for the complete dataset including videos.

17. CITATION

If you use G1 Moves in your research or project, please cite:

```
@misc{chaiet2026g1moves,
  title     = {G1 Moves: A Motion Capture Dataset for Unitree G1
               Humanoid Robot},
  author    = {Chaiet, Mitch},
  year      = {2026},
  publisher = {Hugging Face},
  url       = {https://huggingface.co/datasets/exptech/g1-moves},
  note      = {60 motion capture clips with trained RL policies
               for sim-to-real transfer}
}
```

18. LICENSE

G1 Moves is released under the Creative Commons Attribution 4.0 International License (CC BY 4.0). You are free to share and adapt the material for any purpose, including commercial use, provided you give appropriate credit. See the dataset repository for full license terms.

SAFETY WARNING: Deploying learned locomotion policies to physical robots is inherently dangerous and can cause serious injury, death, or property damage. The policies in this dataset are trained in simulation and have not been validated for safe real-world operation. Do not deploy without appropriate safety infrastructure, expertise, and risk assessment. See the dataset page for the full safety disclaimer and limitation of liability.