

HumanPlus: Humanoid Shadowing and Imitation from Humans

Zipeng Fu^{*1}, Qingqing Zhao^{*1}, Qi Wu^{*1}, Gordon Wetzstein¹, Chelsea Finn¹

^{*}project co-leads, ¹Stanford University

<https://humanoid-ai.github.io>



Figure 1: Stanford HumanPlus Robot. We present a full-stack system for humanoid robots to learn motion and autonomous skills from human data. Our system enables robots to shadow fast, diverse motions from a human operator, including boxing and playing table tennis, and to learn autonomous skills like wearing a shoe, folding clothes, and jumping high.

Abstract

One of the key arguments for building robots that have similar form factors to human beings is that we can leverage the massive human data for training. Yet, doing so has remained challenging in practice due to the complexities in humanoid perception and control, lingering physical gaps between humanoids and humans in morphologies and actuation,

and lack of a data pipeline for humanoids to learn autonomous skills from egocentric vision. In this paper, we introduce a full-stack system for humanoids to learn motion and autonomous skills from human data. We first train a low-level policy in simulation via reinforcement learning using existing 40-hour human motion datasets. This policy transfers to the real world and allows humanoid robots to follow hu-

man body and hand motion in real time using only a RGB camera, i.e. shadowing. Through shadowing, human operators can teleoperate humanoids to collect whole-body data for learning different tasks in the real world. Using the data collected, we then perform supervised behavior cloning to train skill policies using egocentric vision, allowing humanoids to complete different tasks autonomously by imitating human skills. We demonstrate the system on our customized 33-DoF 180cm humanoid, autonomously completing tasks such as wearing a shoe to stand up and walk, unloading objects from warehouse racks, folding a sweatshirt, rearranging objects, typing, and greeting another robot with 60-100% success rates using up to 40 demonstrations.

1. Introduction

Humanoid robots have long been of interest in the robotics community due to their human-like form factors. Since our surrounding environments, tasks and tools are structured and designed based on the human morphology, human-sized humanoids are the nature hardware platforms of general-purpose robots for potentially solving all tasks that people can complete. The human-like morphology of humanoids also presents a unique opportunity to leverage the vast amounts of human motion and skill data available for training, bypassing the scarcity of robot data. By mimicking humans, humanoids can potentially tap into the rich repertoire of skills and motion exhibited by humans, offering a promising avenue towards achieving general robot intelligence.

However, it has remained challenging in practice for humanoids to learn from human data. The complex dynamics and high-dimensional state and action spaces of humanoids pose difficulties in both perception and control. Traditional approaches, such as decoupling the problem into perception, planning and tracking, and separate modularization of control for arms and legs [10, 10, 23, 40], can be time-consuming to be designed and limited in scope, making them difficult to scale to the diverse range of tasks and environments that humanoids are expected to operate in. Moreover, although humanoids closely resemble humans compared to other forms of robots, physical differences between humanoids and humans in morphology and actuation still exist, including number of degrees of freedom, link length, height, weight, vision parameters and mechanisms, and actuation strength and responsiveness, presenting barriers for humanoids to effectively use and learn from human data. This problem is further exacerbated by the lack of off-the-shelf and integrated

hardware platforms. Additionally, we lack an accessible data pipeline for whole-body teleoperation of humanoids, preventing researchers from leveraging imitation learning as a tool to teach humanoids arbitrary skills. Humanoids developed by multiple companies have demonstrated the potential of this data pipeline and subsequent imitation learning from the data collected, but details aren't publicly available, and autonomous demonstrations of their systems are limited to a couple tasks. Prior works use motion capture systems, first-person-view (FPV) virtual reality (VR) headsets and exoskeletons to teleoperate humanoids [17, 20, 38, 59], which are expensive and restricted in operation locations.

In this paper, we present a full-stack system for humanoids to learn motion and autonomous skills from human data. To tackle the control complexity of humanoids, we follow the recent success in legged robotics using large-scale reinforcement learning in simulation and sim-to-real transfer [41, 51] to train a low-level policy for whole-body control. Typically, learning-based low-level policies are designed to be task-specific due to time-consuming reward engineering [19, 68], enabling the humanoid hardware to demonstrate only one skill at a time, such as walking. This limitation restricts the diverse range of tasks that the humanoid platform is capable of performing. At the same time, we have a 40-hour human motion dataset, AMASS [49], that covers a wide range of skills. We leverage this dataset by first retargeting human poses to humanoid poses and then training a task-agnostic low-level policy called Humanoid Shadowing Transformer conditioning on the retargeted humanoid poses. Our pose-conditioned low-level policy transfers to the real world zero-shot.

After deploying the low-level policy that controls the humanoid given target poses, we can shadow human motion to our customized 33-DoF 180cm humanoid in real time using a single RGB camera. Using state-of-the-art human body and hand pose estimation algorithms [58, 81], we can estimate real-time human motion and retarget it to humanoid motion, which is passed as input to the low-level policy. This process is traditionally done by using motion capture systems, which are expensive and restricted in operation locations. Using line of sight, human operators standing nearby can teleoperate humanoids to collect whole-body data for various tasks in the real world, like boxing, playing the piano, playing table tennis and opening cabinets to store a heavy pot. While being teleoperated, the humanoid collects egocentric vision data through binocular RGB cam-




Figure 2: Hardware Details. Our HumanPlus robot has two egocentric RGB cameras mounted on the head, two 6-DoF dexterous hands, and 33 degrees of freedom in total.

eras. Shadowing provides an efficient data collection pipeline for diverse real-world tasks, bypassing the sim-to-real gap of RGB perception.

Using the data collected through shadowing, we perform supervised behavior cloning to train vision-based skill policies. A skill policy takes in humanoid binocular egocentric RGB vision as inputs and predicts the desired humanoid body and hand poses. We build upon the recent success of imitation learning from human-provided demonstrations [11, 104], and introduce a transformer-based architecture that blends action prediction and forward dynamics prediction. Using forward dynamics prediction on image features, our method shows improved performance by regularizing on image feature spaces and preventing the vision-based skill policy from ignoring image features and overfitting to proprioception. Using up to 40 demonstrations, our humanoid can autonomously complete tasks such as wearing a shoe to stand up and walk, unloading objects from warehouse racks, folding a sweatshirt, rearranging objects, typing, and greeting with another robot with 60-100% success rates.

The main contribution of this paper is a full-stack humanoid system for learning complex autonomous skills from human data, named HumanPlus. Core to this system is both (1) a real-time shadowing system that allows human operators to whole-body control humanoids using a single RGB camera and Humanoid Shadowing Transformer, a low-level policy that is trained on massive human motion data in simulation, and (2) Humanoid Imitation Transformer, an imitation learning algorithm that enables efficient learning from 40 demonstrations for binocular perception and high-DoF control. The synergy between our shadowing system and imitation learning algorithm allows learning of whole-body manipulation

and locomotion skills directly in the real-world, such as wearing a shoe to stand up and walk, using only up to 40 demonstrations with 60-100% success.

2. Related Work

Reinforcement Learning for Humanoids. Reinforcement learning for humanoids has been predominantly focusing on locomotion. While model-based control [16, 36, 40, 55, 69, 92, 98] has made tremendous progress on a wide variety of humanoid robots [13, 32, 37, 54, 65, 89], learning-based methods can achieve robust locomotion performance for humanoids [39, 67, 68, 85, 91, 94, 103] and biped robots [4, 42, 44, 45, 83, 84, 93, 100] due to their training on highly randomized environments in simulation and their ability to adapt. Although locomanipulation and mobile manipulation using humanoids are mostly approached via model-predictive control [1, 24, 29, 30, 73, 77], there has been some recent success on applying reinforcement learning and sim-to-real to humanoids for box relocation by explicitly modeling the scene and task in simulation [19], and for generating diverse upper-body motion [9]. In contrast, we use reinforcement learning to train a low-level policy for task-agnostic whole-body control requiring no explicit modeling of real-world scenes and tasks in simulation.

Teleoperation of Humanoids. Prior works develop humanoid and dexterous teleoperation by using human motion capture suits [14, 17, 20, 21], exoskeletons [35, 35, 63, 70, 75], haptic feedback devices [6, 61, 71], and VR devices for visual feedbacks [7, 38, 59, 90] and for end-effector control [2, 46, 62, 86]. For example, Purushottam et al. develop whole-body teleoperation of a wheeled humanoid using an exoskeleton suit attached to a force plate for human motion recording. In terms of control space,




Figure 3: *Shadowing and Retargeting*. Our system uses one RGB camera for body and hand pose estimation.

prior works have done teleoperation in operation spaces [17, 76], upper-body teleoperation [7, 22], and whole-body teleoperation [31, 33, 34, 52, 56, 60, 90]. For example, He et al. use an RGB camera to capture human motion to whole-body teleoperate a humanoid. Seo et al. use VR controllers to teleoperate bimanual end-effectors and perform imitation learning on the collected data to learn static manipulation skills. In contrast, our work provides a full-stack system that consists of a low-cost whole-body teleoperation system using a single RGB camera for controlling every joint of a humanoid, enabling manipulation, squatting and walking, and an efficient imitation pipeline for learning autonomous manipulation and locomotion skills, enabling complex skills like wearing a shoe to stand up and walk.

Robot Learning from Human Data. Human data has been used extensively for robot learning, including for pre-training visual or intermediate representations or tasks [8, 28, 50, 53, 66, 79, 86, 99] leveraging Internet-scale data [18, 26, 27], and for imitation learning on in-domain human data [3, 12, 43, 47, 64, 78, 80, 82, 87, 88, 95–97, 101, 102]. For example, Qin et al. use in-domain human hand data for dexterous robotic hands to imitate. Recently, human data has also been used for training humanoids [9, 31]. Cheng et al. use offline human data for training humanoids to generate diverse upper-body motion, and He et al. use offline human data for training a whole-body teleoperation interface. In contrast, we use both offline human data for learning a low-level whole-body policy for real-time shadowing, and online human data through shadowing for humanoids to imitate human skills, enabling autonomous humanoid skills.

3. HumanPlus Hardware

Our humanoid features 33 degrees of freedom, including two 6-DoF hands, two 1-DoF wrists, and a 19-DoF body (two 4-DoF arms, two 5-DoF legs, and

a 1-DoF waist), as shown on the left of Figure 2. The system is built upon the Unitree H1 robot. Each arm is integrated with an Inspire-Robots RH56DFX hand, connected via a customized wrist. Each wrist has a Dynamixel servo and two thrust bearings. Both the hands and wrists are controlled via serial communication. Our robot has two RGB webcams (Razer Kiyo Pro) mounted on its head, angled 50 degrees downward, with a pupillary distance of 160mm. The fingers can exert forces up to 10N, while the arms can hold items up to 7.5kg. The motors on legs can generate instant torques of up to 360Nm during operation. Additional technical specifications of our robot are provided on the right of Figure 2.

4. Human Body and Hand Data

Offline Human Data. We use a public optical marker-based human motion dataset, AMASS [49] to train our low-level Humanoid Shadowing Transformer. The AMASS dataset aggregates data from several human motion datasets, containing 40 hours of human motion data on a diverse ranges of tasks, and consisting of over 11,000 unique motion sequences. To ensure the quality of the motion data, we apply a filtering process based on an approach outlined in [48]. Human body and hand motions are parameterized using the SMPL-X [57] model, which includes 22 body and 30 hand 3-DoF spherical joints, 3-dimensional global translational transformation, and 3-dimensional global rotational transformation.

Retargeting. Our humanoid body has a subset of the degrees of freedom of SMPL-X body, consisting of only 19 revolute joints. To retarget the body poses, we copy the corresponding Euler angles from SMPL-X to our humanoid model, namely for hips, knees, ankles, torso, shoulders and elbows. Each of the humanoid hip and shoulder joints consists of 3 orthogonal revolute joints, so can be viewed as one spherical joints. Our humanoid hand has 6 degrees

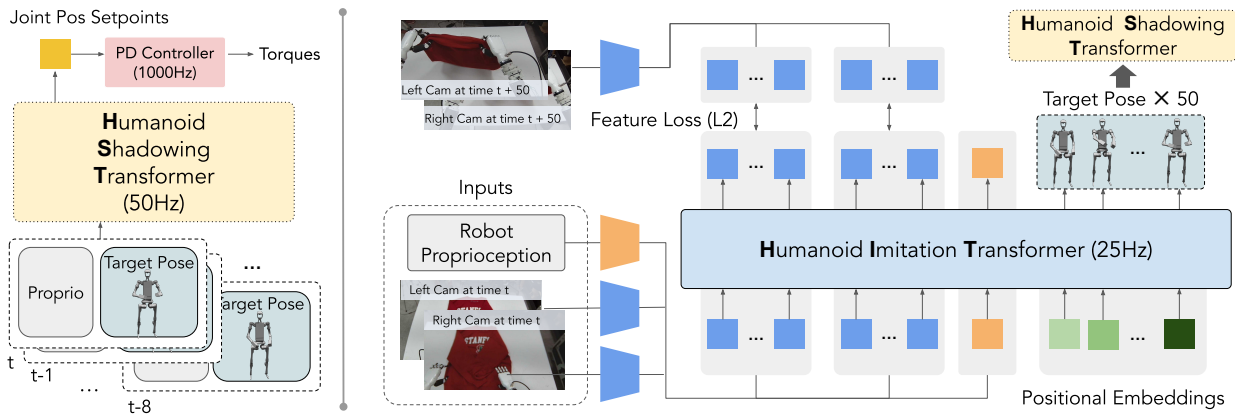


Figure 4: Model Architectures. Our system consists of a decoder-only transformer for low-level control, Humanoid Shadowing Transformer, and a decoder-only transformer for imitation learning, Humanoid Imitation Transformer.

Reward Teams	Expressions
target xy velocities	$\exp(- [v_x, v_y] - [v_x^{tg}, v_y^{tg}])$
target yaw velocities	$\exp(- v_{yaw} - v_{yaw}^{tg})$
target joint positions	$- q - q^{tg} _2^2$
target roll & pitch	$- [r, p] - [r^{tg}, p^{tg}] _2^2$
energy	$- \tau \dot{q} _2^2$
feet contact	$c == c^{tg}$
feet slipping	$- v_{feet} \cdot \mathbb{1}[F_{feet} > 1] _2$
alive	1

Table 1: Rewards in Simulation. We denote v_x as linear x velocity, v_y as linear y velocity, v_{yaw} as angular yaw velocity, q as joint positions, \dot{q} as joint velocities, r as roll, p as pitch, v_{feet} as feet velocities, c as feet contact indicator, F_{feet} as forces on feet, and \cdot^{tg} as targets.

of freedom: 1 DoF for each of the index, middle, ring, and little fingers, and 2 DoFs for the thumb. To re-target hand poses, we map the corresponding Euler angle of each finger using the rotation of the middle joint. To compute the 1-DoF wrist angle, we use the relative rotation between the forearm and hand global orientations.

Real-Time Body Pose Estimation and Retargeting. To estimate human motion in the real world for shadowing, we use World-Grounded Humans with Accurate Motion (WHAM) [81] to jointly estimate the human poses and global transformations in real time using a single RGB camera. WHAM uses SMPL-X for human pose parameterization. Shown in 3, we perform real-time human-to-humanoid body re-targeting using the approach described above. The body pose estimation and re-targeting runs at 25 fps on an NVIDIA RTX4090 GPU.

Real-Time Hand Pose Estimation and Retargeting. We use HaMeR [58], a transformer-based hand

Environment Params	Ranges
base payload	[-3.0, 3.0]kg
end-effector payload	[0, 0.5]kg
center of base mass	[-0.1, 0.1] ³ m
motor strength	[0.8, 1.1]
friction	[0.3, 0.9]
control delay	[0.02, 0.04]s

Table 2: Randomization in Simulation. We uniformly sample from these randomization ranges during training in simulation.

pose estimator using a single RGB camera, for real-time hand pose estimation. HaMeR predicts hand poses, camera parameters, and shape parameters using the MANO [72] hand model. We perform real-time human-to-humanoid hand re-targeting using the approach described above. Our hand pose estimation and re-targeting runs at 10 fps on an NVIDIA RTX4090 GPU.

5. Shadowing of Human Motion

We formulate our low-level policy, Humanoid Shadowing Transformer, as a decoder-only transformer, shown on the left side of Figure 4. At each time step, the input to the policy is humanoid proprioception and a humanoid target pose. The humanoid proprioception contains root state (roll, pitch, and base angular velocities), joint positions, joint velocities and last action. The humanoid target pose consists of target forward and lateral velocities, target roll and pitch, target yaw velocity and target joint angles, and is re-targeted from a human pose sampled from the processed AMASS dataset mentioned in Section 4. The output of the policy is 19-dimensional joint position setpoints for humanoid body joints, which are subsequently converted to torques using

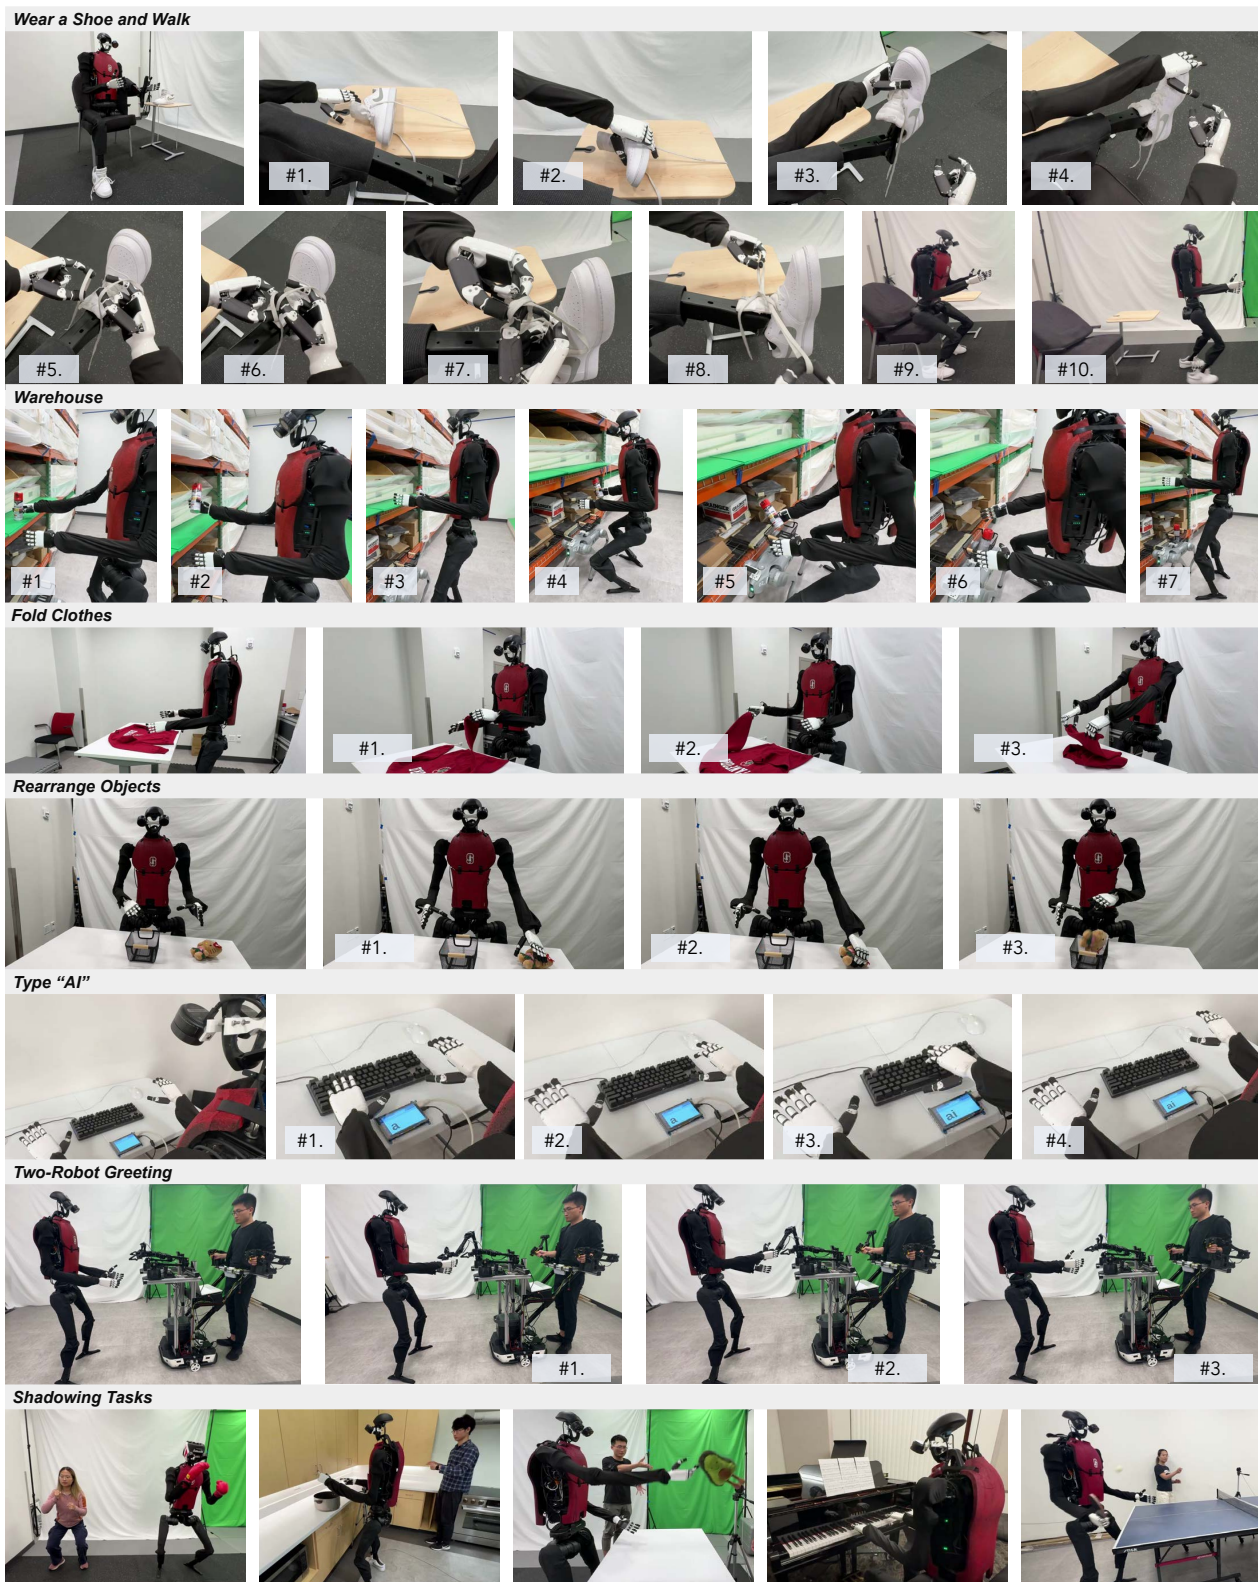


Figure 5: Task Definitions. We illustrate 5 autonomous tasks through imitation learning, and 5 shadowing tasks. Details are in Section 7.

a 1000Hz PD controller. The target hand joint angles are directly passed to the PD controller. Our low-level policy operates at 50Hz and has a context length of 8, so it can adapt to different environments given the observation history [67].

We use PPO [74] to train our Humanoid Shad-

owing Transformer in simulation by maximizing discounted expected return $\mathbb{E} \left[\sum_{t=0}^{T-1} \gamma^t r_t \right]$, where r_t is the reward at time step t , T is the maximum episode length, and γ is the discount factor. The reward r is the sum of terms encouraging matching

	Cost	# of Operators	Whole-Body Approach	Rearrange Objects			Whole Task (s)	Stand (%)	Rearrange Lower Objects (s)
				Object (s)	Pick Object (s)	Place Object (s)			
Kinesthetic	\$50	3	✗	2.10	1.38	3.12	6.60	90.5	-
ALOHA	\$7050	2-3	✗	2.70	1.30	3.15	7.15	100	-
Meta Quest	\$250	2	✗	3.57	1.63	3.67	8.87	95.3	-
Ours	\$50	1	✓	1.76	0.95	2.59	5.20	100	15.34

Table 3: Teleop Comparisons & User Studies. We report averaged completion time for 6 participants on 2 tasks.

target poses while saving energy and avoiding foot slipping. We list all the reward terms in the Table 1. We randomize the physical parameters of the simulated environment and humanoids with details in Table 2.

After training Humanoid Shadowing Transformer in simulation, we deploy it zero-shot to our humanoid in the real world for real-time shadowing. The proprioceptive observations are measured using only onboard sensors including an IMU and joint encoders. Following Section 4 and shown in Figure 3, we estimate the human body and hand poses in real time using a single RGB camera, and retarget the human poses to humanoid target poses. Illustrated in Figure 1, human operators stand near the humanoid to shadow their real-time whole-body motion to our humanoid, and use line of sight to observe the environment and behaviors of humanoid, ensuring a responsive teleoperation system. When the humanoid sits, we directly send the target poses to the PD controller, since we don’t need the policy to compensate gravity, and simulating sitting with rich contacts is challenging. While being teleoperated, the humanoid collects egocentric vision data through binocular RGB cameras. Through shadowing, we provide an efficient data collection pipeline for various real-world tasks, circumventing the challenges of having realist RGB rendering, accurate soft object simulation, and diverse task specifications in simulation.

6. Imitation of Human Skills

Imitation learning has shown great success on learning autonomous robot skills given demonstrations on a wide range of tasks [5, 11, 15, 25, 104]. Given the real-world data collected through shadowing, we apply the same recipe to humanoids to train skill policies. We make several modifications to enable faster inference using limited onboard compute, and efficient imitation learning given binocular perception and high-DoF control.

In this work, we modify the Action Chunking

Transformer [104] by removing its encoder-decoder architecture to develop a decoder-only Humanoid Imitation Transformer (HIT) for skill policies, as depicted on the right side of Figure 4. HIT processes the current image features from two egocentric RGB cameras, proprioception, and fixed positional embeddings as inputs. These image features are encoded using a pretrained ResNet encoder. Due to its decoder-only design, HIT operates by predicting a chunk of 50 target poses based on the fixed positional embeddings at the input, and it can predict tokens corresponding to the image features at their respective input positions. We incorporate an L2 feature loss on these predicted image features, compelling the transformer to predict corresponding image feature tokens for future states after execution of ground truth target pose sequences. This approach allows HIT to blend target pose prediction with forward dynamics prediction effectively. By using forward dynamics prediction on image features, our method enhances performance by regularizing image feature spaces, preventing the vision-based skill policies from ignoring image features and overfitting to proprioception. During deployment, HIT operates at 25Hz onboard, sending predicted target positions to the low-level Humanoid Shadowing Transformer asynchronously, while discarding the predicted future image feature tokens.

7. Tasks

We select six imitation tasks and five shadowing tasks that need bimanual dexterity and whole-body control. Shown in Figure 5, these tasks span a diverse array of capabilities and objects relevant to practical applications.

In the *Wear a Shoe and Walk* task, the robot (1) flips a shoe, (2) picks it up, (3) puts it on, (4) presses it down to secure the fit on the left foot, (5) tangles the shoelaces with both hands, (6) grasps the right one, (7) grasps left one, (8) ties them, (9) stands up and (10) walks forward. This task demonstrates the robot’s ability in complex bimanual manipulation

	Maximum Force Thresholds					Whole-Body Skills		
	forward (↑)	backward (↑)	leftward (↑)	rightward (↑)	Recovery Time (↓)	Squat(↓)	Jump(↑)	Stand Up
Ours	32N	44N	70N	100N	1.2s	0.44m	0.35m	✓
H1 Default	24N	36N	40N	40N	15s	0.85m	0m	✗

Table 4: Robustness Evaluation. Our low-level policy (Ours) can withstand large disturbance forces, has a shorter recovery time, and enables more whole-body skills than the manufacturer controller (H1 Default).




Figure 6: Baseline Teleoperation Systems.

with dexterous hands and capability in agile locomotion like standing up and walk while wearing shoes. The shoe is placed on the table uniformly randomly on a 2cm line along of robot’s forward facing. Each demonstration has 1250 steps or 50 seconds.

In the **Warehouse** task, the robot (1) approaches the paint spray on warehouse shelves with its right hand, (2) grasps the sprayer, (3) retracts the right hand, (4) squats, (5) approaches the cart on the back of the quadruped, (6) release the spray, and (7) stands up. This tasks tests the whole-body manipulation and coordination of the robot. The standing location of the robot is randomized along a 10cm line. Each demonstration has 500 steps or 20 seconds.

In the **Fold Clothes** task, while maintaining balance, the robot (1) folds left sleeve, (2) folds right sleeve, (3) folds the bottom of the sweatshirt, requiring both dexterity to manipulate fabric with complex dynamics and maintaining an upright pose. The robot starts in a standing position, with a uniformly randomly sampled root yaw deviation from +10 to -10 degrees. The sweatshirt is uniformly randomly placed with a deviation of 10cm x 10cm on the table and -30 degrees to 30 degrees in rotation. Each demonstration has 500 steps or 20 seconds.

In the **Rearrange Objects** task, while maintaining balance, the robot (1) approaches the object, (2) picks up the object, and (3) places the object into a basket. The complexity arises from the diverse shapes, colors, and orientations of the objects, requiring the robot to choose the appropriate hand based on the object’s location and to plan its actions accordingly. In total, we uniformly sample from 4 soft objects, including stuffed toys and a ice bag, where the object is placed uniformly randomly along a 10cm line on




Figure 7: Maximum Force Thresholds. Our low-level policy can withstand larger forces compared to H1 Default controller.

either left or right of the basket. Each demonstration has 250 steps or 10s.

In the **Type “AI”** task, the robot (1) types the letter ‘A’, (2) releases the key, (3) types the letter ‘I’, and (4) releases the key. Despite being seated, the robot needs high precision in manipulation. Each demonstration has 200 steps or 8 seconds.

In the **Two-Robot Greeting** task, the robot (1) approaches the other robot with the correct hand after observing the other bimanual robot starts to extend one hand/arm, (2) touches the hand with the other robot, and (3) releases the hand. The other robot uniformly samples which hand to extend and stops in an end-effector region of 5cm x 5cm x 5cm. The robot needs to quickly and accurately recognize which hand to use and approaches the other robot with the correct hand while maintaining balance. Each demonstration has 125 steps or 5 seconds.

For **Shadowing Tasks**, we demonstrate five tasks: boxing, opening a two-door cabinet to store a pot, tossing, playing the piano, playing table tennis, and typing “Hello World”, showcasing the mobility and stability in shadowing fast, diverse motions and manipulating heavy objects. Videos of qualitative shadowing results can be found on the project website: <https://humanoid-ai.github.io>.

8. Experiments on Shadowing

	Fold Clothes (40 demos)					Rearrange Objects (30 demos)				
	Fold Left Sleeve	Fold Right Sleeve	Fold Bottom	Stand	Whole Task	Approach Object	Pick up Object	Place Object	Stand	Whole Task
HIT (Ours)	100	100	100	100	100	100	90	100	100	90
Monocular	80	50	100	100	40	80	88	100	100	70
ACT	100	100	100	100	100	70	86	83	100	50
Open-loop	20	50	0	-	0	0	-	-	-	0

	Type "AI" (30 demos)				Two-Robot Greeting (30 demos)					
	Type "A"	Leave "A"	Type "I"	Leave "I"	Whole Task	Approach Hand	Touch Hand	Release Hand	Stand	Whole Task
HIT (Ours)	90	100	89	100	80	100	90	100	100	90
Monocular	90	100	44	100	40	100	80	100	100	80
ACT	30	20	0	-	0	100	90	100	100	90
Open-loop	82	100	79	100	60	50	0	-	-	0

	Warehouse (25 demos)								
	Approach	Grasp	Retract	Squat	Approach	Release	Stand Up	Whole Task	
HIT (Ours)	100	90	100	100	100	100	100	90	
Monocular	100	80	100	100	100	100	100	80	
ACT	100	90	100	100	100	100	100	90	
Open-loop	60	0	-	-	-	-	-	0	

	Wear a Shoe and Walk (40 demos)										
	Flip Shoe	Pick Up Shoe	Put on Shoe	Press Shoe	Tangle Shoelaces	Grasp Right Shoelaces	Grasp Left Shoelace	Tie Shoelace	Stand Up	Walk Forward	Whole Task
HIT (Ours)	100	100	80	100	100	100	75	100	100	100	60
Monocular	0	-	-	-	-	-	-	-	-	-	0
ACT	50	60	0	-	-	-	-	-	-	-	0
Open-loop	20	0	-	-	-	-	-	-	-	-	0

Table 5: Comparisons on Imitation. We show success rates of Humanoid Imitation Transformer (Ours), HIT with monocular input, ACT and open-loop trajectory replay across all tasks. Overall HIT (Ours) outperforms others.

8.1. Comparisons with Other Teleoperation Methods

We compare our teleoperation system with three baselines: Kinesthetic Teaching, ALOHA [104], and Meta Quest, shown in Figure 6. For *Kinesthetic Teaching*, both arms are in passive mode and manually positioned. For *ALOHA*, we build a pair of bimanual arms for puppeteering from two WidowX 250 robots with similar kinematic structure as our humanoid arms. For *Meta Quest*, we use positions of the controllers for operational space control through inverse kinematics with gravity compensation. Shown in Table 3, all baselines do not support whole-body control and require at least two human operators for hand pose estimation. In contrast, our shadowing system simultaneously control humanoid body and hands, requiring only one human operator. Also, both *ALOHA* and *Meta Quest* are more

expensive. In contrast, our system and *Kinesthetic Teaching* only require a single RGB camera.

Shown in Table 3, we conduct user studies on 6 participants to compare our shadowing system with three baselines in terms of teleoperation efficiency. Two participants have no prior teleoperation experience, while the remaining four have varying levels of expertise. None of the participants has prior experience with our shadowing system. The participants are tasked to perform the *Rearrange Objects* task and its variant, *Rearrange Lower Objects*, where an object is placed on a lower table of height 0.55m, requiring the robot to squat and thus necessitating whole-body control.

We record the average task completion time over six participants, with three trials each and three unrecorded practice rounds. We also record the average success rates of stable standing during teleoperation

using our low-level policy. While *ALOHA* enables precise control of robot joint angles, its fixed hardware setup makes it harder to adapt to people with different heights and body shapes, and it does not support whole-body control of humanoids by default. *Meta Quest* often results in singularities and mismatches between target and actual poses in Cartesian space due to the limited 5 degrees of freedom of each humanoid arm plus a wrist, resulting to the longest completion time and destabilized standing at arm singularities. While *Kinesthetic Teaching* is intuitive and has a low time-to-completion, it requires multiple operators, and sometimes external forces on arms during teaching causes the humanoid to stumble. In contrast, our system has the lowest time-to-completion, has the highest success rate of stable standing, and is the only method that can be used for whole-body teleoperation, solving the *Rearrange Lower Objects* task.

8.2. Robustness Evaluation

Shown in Table 4, We evaluate our low-level policy by comparing it to the manufacturer default controller (*H1 Default*). The robot must maintain balance while manipulating with objects, so we assess robustness by applying forces to the pelvis and record the minimum forces causing instability. Shown in Figure 7, our policy withstands significantly larger forces and has a shorter recovery time. When robot is unbalanced, the manufacturer default controller takes several steps and up to 20 seconds to stabilize the robot, while ours typically recovers within one or two steps and below 3 seconds. More recovery steps result in jittery behavior and compromise manipulation performance. We also show that our policy enables more whole-body skills that are not possible by the default controller like squatting, high jumping, standing up from sitting on a chair.

9. Experiments on Imitation

Shown in Table 5, we compare our imitation learning method Humanoid Imitation Transformer with three baseline methods: HIT policies with monocular inputs (*Monocular*), *ACT* [104], and *Open-loop* trajectory replay, across all tasks: *Fold Clothes*, *Rearrange Objects*, *Type "AI"*, *Two-Robot Greeting*, *Warehouse*, and *Wear a Shoes and Walk*, detailed in Section 7 and Figure 5. Although each skill policy solves its task continuously autonomously without stopping, we document the success rates of consecutive sub-tasks within each task for better analysis. We conduct 10 trials per task. We calculate the success rate for a sub-task by dividing the number of suc-

cessful attempts by the number of total attempts. For example in the case of *Put on Shoe* sub-task, the number of total attempts equals the number of success from the previous sub-task *Pick up Shoe*, as the robot could fail and stop at any sub-task.

Our HIT achieves higher success rates than other baselines across all tasks. Specifically, our method is only method solves the *Wear a Shoe and Walk* task, achieving success rates of 60% given 40 demonstrations, where all other methods fail. This is because our method uses binocular perception, and avoids overfitting to proprioception. *ACT* fails in *Wear a Shoe and Walk* and *Typing "AI"* tasks where it overfits to proprioception, where the robot repeatedly attempts and sticks at *Pick up Shoe* and *Leave "A"* respectively after successfully completing them, avoiding uses visual feedback. *Monocular* shows lower success rates due to its lack of depth information from a single RGB camera, yielding rough interaction with the table in *Fold Clothes*. It fails the *Wear a Shoe and Walk* task completely, where depth perception is crucial. However, due to its narrower field of view, it completes some sub-tasks more successfully than other methods in the *Typing "AI"* task. *Open-loop* only works in *Typing "AI"* with no randomization, and fails in all other tasks that require reactive control.

10. Conclusion, Limitations and Future Directions

In this work, we present HumanPlus, a full-stack system for humanoids to learn motions and autonomous skills from human data. Throughout the development of our system, we encountered several limitations. Firstly, our hardware platform offers fewer degrees of freedom compared to human anatomy. For instance, it uses feet with 1-DoF ankles, which restricts the humanoid’s ability to perform agile movements such as lifting and shaking one leg while the other remains stationary. Each arm has only 5 DoFs, including a wrist, which limits the application of 6-DoF operational space control and may result in unreachable regions during shadowing. Furthermore, the egocentric cameras are fixed on the humanoid’s head and are not active, leading to a constant risk of the hands and interactions falling out of view. In addition, we currently use a fixed retargeting mapping from human poses to humanoid poses, omitting many human joints that do not exist on our humanoid hardware. This may limit the humanoids to learn from a small subset of diverse human motions. Currently, pose estimation methods do not work well given large areas of occlusion, limiting the operating

regions of the human operators. Lastly, we focus on manipulation tasks with some locomotion tasks like squatting, standing up and walking in this work, as dealing with long-horizon navigation requires a much larger size of human demonstrations and accurate velocity tracking in the real world. We hope to address these limitations in future, and to enable more autonomous and robust humanoid skills that can be applied in various real-world tasks.

Acknowledgements

We thank Steve Cousins and Oussama Khatib at Stanford Robotics Center for providing facility support for our experiments. We also thank Inspire-Robots and Unitree Robotics for providing extensive supports on hardware and low-level firmware. We thank Huy Ha, Yihuai Gao, Chong Zhang, Ziwen Zhuang, Jiaman Li, Yifeng Jiang, Yuxiang Zhang, Xingxing Wang, Tony Yang, Walter Wen, Yunguo Cui, Rosy Wang, Zhiqiang Ma, Wei Yu, Xi Chen, Mengda Xu, Peizhuo Li, Tony Z. Zhao, Lucy X. Shi and Bartie for helps on experiments, valuable discussions and supports. This project is supported by The AI Institute and ONR grant N00014-21-1-2685. Zipeng Fu is supported by Pierre and Christine Lamond Fellowship.

References

- [1] Hitoshi Arisumi, Jean-Rémy Chardonnet, Abderrahmane Kheddar, and Kazuhito Yokoi. Dynamic lifting motion of humanoid robots. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, 2007. 3
- [2] Sridhar Pandian Arunachalam, Irmak Güzey, Soumith Chintala, and Lerrel Pinto. Holo-dex: Teaching dexterity with immersive mixed reality. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5962–5969. IEEE, 2023. 3
- [3] Shikhar Bahl, Abhinav Gupta, and Deepak Pathak. Human-to-robot imitation in the wild. *arXiv preprint arXiv:2207.09450*, 2022. 4
- [4] Hamid Benbrahim and Judy A Franklin. Biped dynamic walking using reinforcement learning. *Robotics and Autonomous systems*, 1997. 3
- [5] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla, Deeksha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez, Karl Pertsch, Jornell Quiambao, Kanishka Rao, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Kevin Sayed, Jaspiar Singh, Sumedh Sontakke, Austin Stone, Clayton Tan, Huong Tran, Vincent Vanhoucke, Steve Vega, Quan Vuong, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. Rt-1: Robotics transformer for real-world control at scale. In *arXiv preprint arXiv:2212.06817*, 2022. 7
- [6] Anais Brygo, Ioannis Sarakoglou, Nadia Garcia-Hernandez, and Nikolaos Tsagarakis. Humanoid robot teleoperation with vibrotactile based balancing feedback. In *Haptics: Neuroscience, Devices, Modeling, and Applications: 9th International Conference, EuroHaptics 2014, Versailles, France, June 24-26, 2014, Proceedings, Part II 9*, 2014. 3
- [7] Jean Chagas Vaz, Dylan Wallace, and Paul Y Oh. Humanoid loco-manipulation of pushed carts utilizing virtual reality teleoperation. In *ASME International Mechanical Engineering Congress and Exposition*, 2021. 3, 4
- [8] Annie S Chen, Suraj Nair, and Chelsea Finn. Learning generalizable robotic reward functions from "in-the-wild" human videos. *arXiv preprint arXiv:2103.16817*, 2021. 4
- [9] Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive whole-body control for humanoid robots. *arXiv preprint arXiv:2402.16796*, 2024. 3, 4
- [10] Joel Chestnutt, Manfred Lau, German Cheung, James Kuffner, Jessica Hodgins, and Takeo Kanade. Footstep planning for the honda asimo humanoid. In *ICRA*, 2005. 2
- [11] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023. 3, 7
- [12] Cheng Chi, Zhenjia Xu, Chuer Pan, Eric Cousineau, Benjamin Burchfiel, Siyuan Feng, Russ Tedrake, and Shuran Song. Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024. 4
- [13] Matthew Chignoli, Donghyun Kim, Elijah Stanger-Jones, and Sangbae Kim. The mit humanoid robot: Design, motion planning,

- and control for acrobatic behaviors. In *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*, 2021. 3
- [14] R Cisneros, M Benallegue, K Kaneko, H Kamina, G Caron, A Tanguy, R Singh, L Sun, A Dallard, C Fournier, et al. Team janus humanoid avatar: A cybernetic avatar to embody human telepresence. In *Toward Robot Avatars: Perspectives on the ANA Avatar XPRIZE Competition, RSS Workshop*, 2022. 3
- [15] Open X-Embodiment Collaboration, Abhishek Padalkar, Acorn Pooley, Ajinkya Jain, Alex Bewley, Alex Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anikait Singh, Anthony Brohan, Antonin Raffin, Ayzaan Wahid, Ben Burgess-Limerick, Beomjoon Kim, Bernhard Schölkopf, Brian Ichter, Cewu Lu, Charles Xu, Chelsea Finn, Chenfeng Xu, Cheng Chi, Chenguang Huang, Christine Chan, Chuer Pan, Chuyuan Fu, Coline Devin, Danny Driess, Deepak Pathak, Dhruv Shah, Dieter Büchler, Dmitry Kalashnikov, Dorsa Sadigh, Edward Johns, Federico Ceola, Fei Xia, Freek Stulp, Gaoyue Zhou, Gaurav S. Sukhatme, Gautam Salhotra, Ge Yan, Giulio Schiavi, Hao Su, Hao-Shu Fang, Haochen Shi, Heni Ben Amor, Henrik I Christensen, Hiroki Furuta, Homer Walke, Hongjie Fang, Igor Mordatch, Ilija Radosavovic, Isabel Leal, Jacky Liang, Jaehyung Kim, Jan Schneider, Jasmine Hsu, Jeannette Bohg, Jeffrey Bingham, Jiajun Wu, Jialin Wu, Jianlan Luo, Jiayuan Gu, Jie Tan, Jihoon Oh, Jitendra Malik, Jonathan Tompson, Jonathan Yang, Joseph J. Lim, João Silvério, Junhyek Han, Kanishka Rao, Karl Pertsch, Karol Hausman, Keegan Go, Keerthana Gopalakrishnan, Ken Goldberg, Kendra Byrne, Kenneth Oslund, Kento Kawaharazuka, Kevin Zhang, Keyvan Majd, Krishan Rana, Krishnan Srinivasan, Lawrence Yunliang Chen, Lerrel Pinto, Liam Tan, Lionel Ott, Lisa Lee, Masayoshi Tomizuka, Maximilian Du, Michael Ahn, Mingtong Zhang, Mingyu Ding, Mohan Kumar Srirama, Mohit Sharma, Moo Jin Kim, Naoaki Kanazawa, Nicklas Hansen, Nicolas Heess, Nikhil J Joshi, Niko Suenderhauf, Norman Di Palo, Nur Muhammad Mahi Shafiullah, Oier Mees, Oliver Kroemer, Pannag R Sanketi, Paul Wohlhart, Peng Xu, Pierre Sermanet, Priya Sundaesan, Quan Vuong, Rafael Rafailov, Ran Tian, Ria Doshi, Roberto Martín-Martín, Russell Mendonca, Rutav Shah, Ryan Hoque, Ryan Julian, Samuel Bustamante, Sean Kirmani, Sergey Levine, Sherry Moore, Shikhar Bahl, Shivin Dass, Shuran Song, Sichun Xu, Siddhant Halder, Simeon Adebola, Simon Guist, Soroush Nasiriany, Stefan Schaal, Stefan Welker, Stephen Tian, Sudeep Dasari, Suneel Belkhale, Takayuki Osa, Tatsuya Harada, Tatsuya Matsushima, Ted Xiao, Tianhe Yu, Tianli Ding, Todor Davchev, Tony Z. Zhao, Travis Armstrong, Trevor Darrell, Vidhi Jain, Vincent Vanhoucke, Wei Zhan, Wenxuan Zhou, Wolfram Burgard, Xi Chen, Xiaolong Wang, Xinghao Zhu, Xuanlin Li, Yao Lu, Yevgen Chebotar, Yifan Zhou, Yifeng Zhu, Ying Xu, Yixuan Wang, Yonatan Bisk, Yoonyoung Cho, Youngwoon Lee, Yuchen Cui, Yueh hua Wu, Yujin Tang, Yuke Zhu, Yunzhu Li, Yusuke Iwasawa, Yutaka Matsuo, Zhuo Xu, and Zichen Jeff Cui. Open X-Embodiment: Robotic learning datasets and RT-X models. <https://arxiv.org/abs/2310.08864>, 2023. 7
- [16] Steve Collins, Andy Ruina, Russ Tedrake, and Martijn Wisse. Efficient bipedal robots based on passive-dynamic walkers. *Science*, 2005. 3
- [17] Stefano Dafarra, Kouros Darvish, Riccardo Grieco, Gianluca Milani, Ugo Pattacini, Lorenzo Rapetti, Giulio Romualdi, Mattia Salvi, Alessandro Scalzo, Ines Sorrentino, et al. icub3 avatar system. *arXiv preprint arXiv:2203.06972*, 2022. 2, 3, 4
- [18] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European conference on computer vision (ECCV)*, 2018. 4
- [19] Jeremy Dao, Helei Duan, and Alan Fern. Sim-to-real learning for humanoid box locomanipulation. *arXiv preprint arXiv:2310.03191*, 2023. 2, 3
- [20] Kouros Darvish, Yeshasvi Tirupachuri, Giulio Romualdi, Lorenzo Rapetti, Diego Ferigo, Francisco Javier Andrade Chavez, and Daniele Pucci. Whole-body geometric retargeting for humanoid robots. In *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*, 2019. 2, 3
- [21] Anca D Dragan, Kenton CT Lee, and Siddhartha S Srinivasa. Legibility and predictability of robot motion. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2013. 3

- [22] Mohamed Elobaid, Yue Hu, Giulio Romualdi, Stefano Dafarra, Jan Babic, and Daniele Pucci. Telexistence and teleoperation for walking humanoid robots. In *Intelligent Systems and Applications: Proceedings of the 2019 Intelligent Systems Conference (IntelliSys) Volume 2*, pages 1106–1121. Springer, 2020. 4
- [23] Siyuan Feng, Eric Whitman, X Xinjilefu, and Christopher G Atkeson. Optimization based full body control for the atlas robot. In *International Conference on Humanoid Robots*, 2014. 2
- [24] Paolo Ferrari, Marco Cognetti, and Giuseppe Oriolo. Humanoid whole-body planning for loco-manipulation tasks. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017. 3
- [25] Zipeng Fu, Tony Z Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. *arXiv preprint arXiv:2401.02117*, 2024. 7
- [26] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The "something something" video database for learning and evaluating visual common sense. In *Proceedings of the IEEE international conference on computer vision*, pages 5842–5850, 2017. 4
- [27] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 4
- [28] Nicklas Hansen, Zhecheng Yuan, Yanjie Ze, Tongzhou Mu, Aravind Rajeswaran, Hao Su, Huazhe Xu, and Xiaolong Wang. On pre-training for visuo-motor control: Revisiting a learning-from-scratch baseline. *arXiv preprint arXiv:2212.05749*, 2022. 4
- [29] Kensuke Harada, Shuuji Kajita, Hajime Saito, Mitsuharu Morisawa, Fumio Kanehiro, Kiyoshi Fujiwara, Kenji Kaneko, and Hirohisa Hirukawa. A humanoid robot carrying a heavy object. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, 2005. 3
- [30] Kensuke Harada, Shuuji Kajita, Fumio Kanehiro, Kiyoshi Fujiwara, Kenji Kaneko, Kazuhito Yokoi, and Hirohisa Hirukawa. Real-time planning of humanoid robot's gait for force-controlled manipulation. *IEEE/ASME Transactions on Mechatronics*, 2007. 3
- [31] Tairan He, Zhengyi Luo, Wenli Xiao, Chong Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Learning human-to-humanoid real-time whole-body teleoperation. *arXiv preprint arXiv:2403.04436*, 2024. 4
- [32] Kazuo Hirai, Masato Hirose, Yuji Haikawa, and Toru Takenaka. The development of honda humanoid robot. In *Proceedings. 1998 IEEE international conference on robotics and automation*, 1998. 3
- [33] Kai Hu, Christian Ott, and Dongheui Lee. Online human walking imitation in task and joint space based on quadratic programming. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3458–3464. IEEE, 2014. 4
- [34] Yasuhiro Ishiguro, Kunio Kojima, Fumihito Sugai, Shunichi Nozawa, Yohei Kakiuchi, Kei Okada, and Masayuki Inaba. High speed whole body dynamic motion experiment with real time master-slave humanoid robot system. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5835–5841. IEEE, 2018. 4
- [35] Yasuhiro Ishiguro, Tasuku Makabe, Yuya Nagamatsu, Yuta Kojio, Kunio Kojima, Fumihito Sugai, Yohei Kakiuchi, Kei Okada, and Masayuki Inaba. Bilateral humanoid teleoperation system using whole-body exoskeleton cockpit tablis. *IEEE Robotics and Automation Letters*, 2020. 3
- [36] Shuuji Kajita, Fumio Kanehiro, Kenji Kaneko, Kazuhito Yokoi, and Hirohisa Hirukawa. The 3d linear inverted pendulum mode: A simple modeling for a biped walking pattern generation. In *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No. 01CH37180)*, 2001. 3
- [37] Ichiro Kato. Development of wabot 1. *Biomechanism*, 1973. 3
- [38] Doik Kim, Bum-Jae You, and Sang-Rok Oh. Whole body motion control framework for arbitrarily and simultaneously assigned upper-body tasks and walking motion. *Modeling, Simulation and Optimization of Bipedal Walking*, 2013. 2, 3

- [39] Lokesh Krishna, Guillermo A Castillo, Utkarsh A Mishra, Ayonga Hereid, and Shishir Kolathaya. Linear policies are sufficient to realize robust bipedal walking on challenging terrains. *IEEE Robotics and Automation Letters*, 2022. 3
- [40] Scott Kuindersma, Robin Deits, Maurice Fallon, Andrés Valenzuela, Hongkai Dai, Frank Permenter, Twan Koolen, Pat Marion, and Russ Tedrake. Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot. *Autonomous robots*, 2016. 2, 3
- [41] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots. *arXiv preprint arXiv:2107.04034*, 2021. 2
- [42] Ashish Kumar, Zhongyu Li, Jun Zeng, Deepak Pathak, Koushil Sreenath, and Jitendra Malik. Adapting rapid motor adaptation for bipedal robots. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022. 3
- [43] Sateesh Kumar, Jonathan Zamora, Nicklas Hansen, Rishabh Jangir, and Xiaolong Wang. Graph inverse reinforcement learning from diverse videos. In *Conference on Robot Learning*, 2023. 4
- [44] Zhongyu Li, Xuxin Cheng, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Reinforcement learning for robust parameterized locomotion control of bipedal robots. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021. 3
- [45] Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control. *arXiv preprint arXiv:2401.16889*, 2024. 3
- [46] Toru Lin, Yu Zhang, Qiyang Li, Haozhi Qi, Brent Yi, Sergey Levine, and Jitendra Malik. Learning visuotactile skills with two multifingered hands. *arXiv:2404.16823*, 2024. 3
- [47] YuXuan Liu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Imitation from observation: Learning to imitate behaviors from raw video via context translation. In *2018 IEEE international conference on robotics and automation (ICRA)*, 2018. 4
- [48] Zhengyi Luo, Jinkun Cao, Alexander W. Winkler, Kris Kitani, and Weipeng Xu. Perpetual humanoid control for real-time simulated avatars. In *International Conference on Computer Vision (ICCV)*, 2023. 4
- [49] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *International Conference on Computer Vision*, 2019. 2, 4
- [50] Arjun Majumdar, Karmesh Yadav, Sergio Arnaud, Jason Ma, Claire Chen, Sneha Silwal, Aryan Jain, Vincent-Pierre Berges, Tingfan Wu, Jay Vakil, et al. Where are we in the search for an artificial visual cortex for embodied intelligence? *Advances in Neural Information Processing Systems*, 2024. 4
- [51] Takahiro Miki, Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning robust perceptive locomotion for quadrupedal robots in the wild. *Science Robotics*, 2022. 2
- [52] Francisco-Javier Montecillo-Puente, Manish Sreenivasa, and Jean-Paul Laumond. On real-time whole-body human to humanoid motion transfer. 2010. 4
- [53] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3m: A universal visual representation for robot manipulation. *arXiv preprint arXiv:2203.12601*, 2022. 4
- [54] Gabe Nelson, Aaron Saunders, Neil Neville, Ben Swilling, Joe Bondaryk, Devin Billings, Chris Lee, Robert Playter, and Marc Raibert. Petman: A humanoid robot for testing chemical protective clothing. *Journal of the Robotics Society of Japan*, 2012. 3
- [55] Quan Nguyen, Ayush Agrawal, Xingye Da, William C Martin, Hartmut Geyer, Jessy W Grizzle, and Koushil Sreenath. Dynamic walking on randomly-varying discrete terrain with one-step preview. In *Robotics: Science and Systems*, 2017. 3
- [56] Kazuya Otani and Karim Bouyarmane. Adaptive whole-body manipulation in human-to-humanoid multi-contact motion retargeting. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, pages 446–453. IEEE, 2017. 4
- [57] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings IEEE*

- Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019. 4
- [58] Georgios Pavlakos, Dandan Shan, Ilija Radosavovic, Angjoo Kanazawa, David Fouhey, and Jitendra Malik. Reconstructing hands in 3D with transformers. In *CVPR*, 2024. 2, 5
- [59] Luigi Penco, Nicola Scianca, Valerio Modugno, Leonardo Lanari, Giuseppe Oriolo, and Serena Ivaldi. A multimode teleoperation framework for humanoid loco-manipulation: An application for the icub robot. *IEEE Robotics & Automation Magazine*, 2019. 2, 3
- [60] Luigo Penco, Brice Clément, Valerio Modugno, E Mingo Hoffman, Gabriele Nava, Daniele Pucci, Nikos G Tsagarakis, J-B Mouret, and Serena Ivaldi. Robust real-time whole-body motion retargeting from human to humanoid. In *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*, pages 425–432. IEEE, 2018. 4
- [61] Luka Peternel and Jan Babič. Learning of compliant human-robot interaction using full-body haptic interface. *Advanced Robotics*, 2013. 3
- [62] Tiffany Portela, Gabriel B Margolis, Yandong Ji, and Pulkit Agrawal. Learning force control for legged manipulation. *arXiv preprint arXiv:2405.01402*, 2024. 3
- [63] Amartya Purushottam, Christopher Xu, Yeongtae Jung, and Joao Ramos. Dynamic mobile manipulation via whole-body bilateral teleoperation of a wheeled humanoid. *IEEE Robotics and Automation Letters*, 2023. 3
- [64] Yuzhe Qin, Yueh-Hua Wu, Shaowei Liu, Hanwen Jiang, Ruihan Yang, Yang Fu, and Xiaolong Wang. Dexmv: Imitation learning for dexterous manipulation from human videos. In *European Conference on Computer Vision*, 2022. 4
- [65] Nicolaus A Radford, Philip Strawser, Kimberly Hambuchen, Joshua S Mehling, William K Verdeyen, A Stuart Donnan, James Holley, Jairo Sanchez, Vienny Nguyen, Lyndon Bridgwater, et al. Valkyrie: Nasa’s first bipedal humanoid robot. *Journal of Field Robotics*, 2015. 3
- [66] Ilija Radosavovic, Tete Xiao, Stephen James, Pieter Abbeel, Jitendra Malik, and Trevor Darrell. Real-world robot learning with masked visual pre-training. In *Conference on Robot Learning*, 2023. 4
- [67] Ilija Radosavovic, Tete Xiao, Bike Zhang, Trevor Darrell, Jitendra Malik, and Koushil Sreenath. Real-world humanoid locomotion with reinforcement learning. *Science Robotics*, 2024. 3, 6
- [68] Ilija Radosavovic, Bike Zhang, Baifeng Shi, Jathushan Rajasegaran, Sarthak Kamat, Trevor Darrell, Koushil Sreenath, and Jitendra Malik. Humanoid locomotion as next token prediction. *arXiv preprint arXiv:2402.19469*, 2024. 2, 3
- [69] Marc H Raibert. *Legged robots that balance*. MIT press, 1986. 3
- [70] Joao Ramos and Sangbae Kim. Humanoid dynamic synchronization through whole-body bilateral feedback teleoperation. *IEEE Transactions on Robotics*, 2018. 3
- [71] Joao Ramos and Sangbae Kim. Dynamic locomotion synchronization of bipedal robot and human operator via bilateral feedback teleoperation. *Science Robotics*, 2019. 3
- [72] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 2017. 5
- [73] Shimpei Sato, Yuta Kojio, Kunio Kojima, Fumihito Sugai, Yohei Kakiuchi, Kei Okada, and Masayuki Inaba. Drop prevention control for humanoid robots carrying stacked boxes. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021. 3
- [74] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 6
- [75] Max Schwarz, Christian Lenz, Andre Rochow, Michael Schreiber, and Sven Behnke. Nimbro avatar: Interactive immersive telepresence with force-feedback telemanipulation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5312–5319, 2021. 3
- [76] Mingyo Seo, Steve Han, Kyutae Sim, Seung Hyeon Bang, Carlos Gonzalez, Luis Sentis, and Yuke Zhu. Deep imitation learning for humanoid loco-manipulation through human teleoperation. In *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*, 2023. 4
- [77] Alessandro Settimi, Danilo Caporale, Przemyslaw Kryczka, Mirko Ferrati, and Lucia Pallottino. Motion primitive based random planning for loco-manipulation tasks. In *2016*

- IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, 2016. 3
- [78] Nur Muhammad Mahi Shafiullah, Anant Rai, Haritheja Etukuru, Yiqian Liu, Ishan Misra, Soumith Chintala, and Lerrel Pinto. On bringing robots home. *arXiv preprint arXiv:2311.16098*, 2023. 4
- [79] Lin Shao, Toki Migimatsu, Qiang Zhang, Karen Yang, and Jeannette Bohg. Concept2robot: Learning manipulation concepts from instructions and human demonstrations. *The International Journal of Robotics Research*, 40(12-14):1419–1434, 2021. 4
- [80] Pratyusha Sharma, Deepak Pathak, and Abhinav Gupta. Third-person visual imitation learning via decoupled hierarchical controller. *Advances in Neural Information Processing Systems*, 2019. 4
- [81] Soyong Shin, Juyong Kim, Eni Halilaj, and Michael J. Black. Wham: Reconstructing world-grounded humans with accurate 3d motion. In *Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 5
- [82] Maximilian Sieb, Zhou Xian, Audrey Huang, Oliver Kroemer, and Katerina Fragkiadaki. Graph-structured visual imitation. In *Conference on Robot Learning*, 2020. 4
- [83] Jonah Siekmann, Yesh Godse, Alan Fern, and Jonathan Hurst. Sim-to-real learning of all common bipedal gaits via periodic reward composition. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021. 3
- [84] Jonah Siekmann, Kevin Green, John Warila, Alan Fern, and Jonathan Hurst. Blind bipedal stair traversal via sim-to-real reinforcement learning. *arXiv preprint arXiv:2105.08328*, 2021. 3
- [85] Rohan P Singh, Zhaoming Xie, Pierre Gergondet, and Fumio Kanehiro. Learning bipedal walking for humanoids with current feedback. *IEEE Access*, 2023. 3
- [86] Aravind Sivakumar, Kenneth Shaw, and Deepak Pathak. Robotic telekinesis: Learning a robotic hand imitator by watching humans on youtube. *arXiv preprint arXiv:2202.10448*, 2022. 3, 4
- [87] Laura Smith, Nikita Dhawan, Marvin Zhang, Pieter Abbeel, and Sergey Levine. Avid: Learning multi-stage tasks via pixel-level translation of human videos. *arXiv preprint arXiv:1912.04443*, 2019. 4
- [88] Shuran Song, Andy Zeng, Johnny Lee, and Thomas Funkhouser. Grasping in the wild: Learning 6dof closed-loop grasping from low-cost demonstrations. *IEEE Robotics and Automation Letters*, 2020. 4
- [89] Olivier Stasse, Thomas Flayols, Rohan Budhiraja, Kevin Giraud-Esclasse, Justin Carpentier, Joseph Mirabel, Andrea Del Prete, Philippe Souères, Nicolas Mansard, Florent Lamiroux, et al. Talos: A new humanoid research platform targeted for industrial applications. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, 2017. 3
- [90] Susumu Tachi, Yasuyuki Inoue, and Fumihiro Kato. Telesar vi: Telexistence surrogate anthropomorphic robot vi. *International Journal of Humanoid Robotics*. 3, 4
- [91] Annan Tang, Takuma Hiraoka, Naoki Hiraoka, Fan Shi, Kento Kawaharazuka, Kunio Kojima, Kei Okada, and Masayuki Inaba. Humanmimic: Learning natural locomotion and transitions for humanoid robot via wasserstein adversarial imitation. *arXiv preprint arXiv:2309.14225*, 2023. 3
- [92] Yuval Tassa, Tom Erez, and Emanuel Todorov. Synthesis and stabilization of complex behaviors through online trajectory optimization. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012. 3
- [93] Russ Tedrake, Teresa Weirui Zhang, and H Sebastian Seung. Stochastic policy gradient reinforcement learning on a simple 3d biped. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, 2004. 3
- [94] Bart van Marum, Aayam Shrestha, Helei Duan, Pranay Dugar, Jeremy Dao, and Alan Fern. Revisiting reward design and evaluation for robust humanoid standing and walking. *arXiv preprint arXiv:2404.19173*, 2024. 3
- [95] Chen Wang, Linxi Fan, Jiankai Sun, Ruohan Zhang, Li Fei-Fei, Danfei Xu, Yuke Zhu, and Anima Anandkumar. Mimicplay: Long-horizon imitation learning by watching human play. *arXiv preprint arXiv:2302.12422*, 2023. 4
- [96] Chen Wang, Haochen Shi, Weizhuo Wang, Ruohan Zhang, Li Fei-Fei, and C. Karen Liu. Dexcap: Scalable and portable mocap data collection system for dexterous manipulation. *arXiv preprint arXiv:2403.07788*, 2024.
- [97] Jianren Wang, Sudeep Dasari, Mohan Kumar Srirama, Shubham Tulsiani, and Abhinav Gupta. Manipulate by seeing: Creating ma-

- nipulation controllers from pre-trained representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3859–3868, 2023. 4
- [98] Eric R Westervelt, Jessy W Grizzle, and Daniel E Koditschek. Hybrid zero dynamics of planar biped walkers. *IEEE transactions on automatic control*, 2003. 3
- [99] Tete Xiao, Ilija Radosavovic, Trevor Darrell, and Jitendra Malik. Masked visual pre-training for motor control. *arXiv preprint arXiv:2203.06173*, 2022. 4
- [100] Zhaoming Xie, Glen Berseth, Patrick Clary, Jonathan Hurst, and Michiel Van de Panne. Feedback control for cassie with deep reinforcement learning. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018. 3
- [101] Haoyu Xiong, Quanzhou Li, Yun-Chun Chen, Homanga Bharadhwaj, Samarth Sinha, and Animesh Garg. Learning by watching: Physical imitation of manipulation skills from human videos. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021. 4
- [102] Sarah Young, Dhiraj Gandhi, Shubham Tulsiani, Abhinav Gupta, Pieter Abbeel, and Lerrel Pinto. Visual imitation made easy. In *Conference on Robot Learning*. PMLR, 2021. 4
- [103] Qiang Zhang, Peter Cui, David Yan, Jingkai Sun, Yiqun Duan, Arthur Zhang, and Renjing Xu. Whole-body humanoid robot locomotion with human reference. *arXiv preprint arXiv:2402.18294*, 2024. 3
- [104] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bi-manual manipulation with low-cost hardware. *RSS*, 2023. 3, 7, 9, 10