Catch & Carry: Reusable Neural Controllers for Vision-Guided Whole-Body Tasks

JOSH MEREL, SARAN TUNYASUVUNAKOOL, ARUN AHUJA, YUVAL TASSA, LEONARD HASENCLEVER, VU PHAM, TOM EREZ, GREG WAYNE, and NICOLAS HEESS, DeepMind

Fig. 1. A catch-carry-toss sequence (bottom) from first-person visual inputs (top). Note how the character's gaze and posture track the ball.

We address the longstanding challenge of producing flexible, realistic humanoid character controllers that can perform diverse whole-body tasks involving object interactions. This challenge is central to a variety of fields, from graphics and animation to robotics and motor neuroscience. Our physics-based environment uses realistic actuation and first-person perception – including touch sensors and egocentric vision – with a view to producing active-sensing behaviors (e.g. gaze direction), transferability to real robots, and comparisons to the biology. We develop an integrated neural-network based approach consisting of a motor primitive module, human demonstrations, and an instructed reinforcement learning regime with curricula and task variations. We demonstrate the utility of our approach for several tasks, including goal-conditioned box carrying and ball catching, and we characterize its behavioral robustness. The resulting controllers can be deployed in real-time on a standard PC.¹

CCS Concepts: • **Computing methodologies** → **Artificial intelligence**; **Control methods**; **Physical simulation**; **Motion capture**.

Additional Key Words and Phrases: reinforcement learning, physics-based character, motor control, object interaction

¹See overview video, [Video 1](#). Refer to Supplementary Section F for all video captions.

Authors' address: Josh Merel, jmerel@google.com; Saran Tunyasuvunakool; Arun Ahuja; Yuval Tassa; Leonard Hasenclever; Vu Pham; Tom Erez; Greg Wayne; Nicolas Heess, DeepMind.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2020 Copyright held by the owner/author(s).

0730-0301/2020/7-ART39

<https://doi.org/10.1145/3386569.3392474>

ACM Reference Format:

Josh Merel, Saran Tunyasuvunakool, Arun Ahuja, Yuval Tassa, Leonard Hasenclever, Vu Pham, Tom Erez, Greg Wayne, and Nicolas Heess. 2020. Catch & Carry: Reusable Neural Controllers for Vision-Guided Whole-Body Tasks. *ACM Trans. Graph.* 39, 4, Article 39 (July 2020), 14 pages. <https://doi.org/10.1145/3386569.3392474>

1 INTRODUCTION

Endowing embodied agents with the motor intelligence that is required for natural and flexible goal-directed interaction with their physical environment is a longstanding challenge in artificial intelligence [Pfeifer and Scheier 2001]. This is a problem of obvious practical relevance to a number of fields including robotics [Arkin et al. 1998] and motor neuroscience [Merel et al. 2019b]. But it is also a topic of longstanding interest in the graphics and animation setting [since e.g. Raibert and Hodgins 1991; Van de Panne and Fiume 1993] – the ability to control agents with physically simulated bodies and sensors that naturally behave in response to high-level instructions may reduce the effort in creating realistic animations of agent-environment interactions.

Conventional approaches such as manual editing and kinematic blending of motion capture data require specification of character movements in their entirety, including how movements conform to the surrounding scene and task context. This can be challenging, especially when fine-grained motion is required, e.g. for object interaction. However, when controlling a goal-directed agent with physically simulated body, many aspects of its movements will emerge from the physical interaction; while other features, such as gaze direction, emerge from the interaction between the agent's goals and the constraints imposed by its body and sensors.

Recent developments in Deep Reinforcement Learning (Deep RL) have enabled great strides in learning from scratch for many game

domains [Mnih et al. 2015; Silver et al. 2018]; and while complex motor control problems involving physical bodies remain difficult even in simulation, there have been some successes discovering locomotion from scratch for reasonably sophisticated bodies [Heess et al. 2017]. But generation of more complex behaviors, especially whole body movements that include object interactions, have remained largely out of reach. These settings require the coordination of a complex, high-dimensional body to achieve a task goal, and satisfactory performance is often a complicated intersection of task success and additional constraints (e.g. naturalness, robustness, or energy-efficiency of the movements). Unlike problems with clear single-task performance objectives, these criteria can be hard to formalize as a reward function. Even where this is possible, the discovery of good solutions through RL can be difficult.

Consequently, physics-based control algorithms often leverage prior knowledge, for instance, in the form of demonstrations or skills transferred from other tasks. These can help with the discovery of rewarding behavior [e.g. Heess et al. 2016] as well as constrain the solutions that emerge [e.g. Peng et al. 2018; Merel et al. 2019c]. These control settings force us to confront a fundamental trade-off: while narrow, stereotyped skills, e.g. from demonstrations, can serve as useful initializations in settings where the controller only needs to reproduce one movement pattern, novel compositions of movements may be required for other settings, which, while related to demonstrations, are not completely consistent with them. This is particularly pertinent to whole-body humanoid control that includes object interaction. While locomotion skills are only a function of the body’s pose and its relation to the ground, manipulation skills are inherently tied to objects; yet we want manipulation skills to be general enough to apply not just a single scene with a particular object, but also to novel objects and object configurations.

Here we develop an integrated learning approach for humanoid whole-body manipulation and locomotion in simulation, that allows us to strike a satisfactory balance between task-specificity and motor generality for object interaction behaviors. It consists of the following components: (1) a general purpose *low-level motor skill module* that is derived from motion capture demonstrations, yet is scene agnostic and can therefore be deployed in many scenarios; (2) a hierarchical control scheme, consisting of a high-level *task policy* that operates from egocentric vision, possesses memory, and interfaces with the motor module; (3) a training procedure involving a broad distribution of *task variations* to achieve generalization to a number of different environmental conditions; and lastly, (4) training using a *phased task*, in which the task policy is trained to solve task stages using simple rewards, which, together with the use of demonstrations, greatly facilitates exploration and allows us to learn complex multi-step tasks while minimizing the need for complicated shaping rewards. This approach builds on the *neural probabilistic motor primitives* (NPMP) approach of Merel et al. 2019c. While that work demonstrated reusable skills for locomotion behaviors without objects, the present work adapts the training procedure to produce a similarly structured motor skill module from demonstrations that can successfully be reused for varied object interactions. In particular, to support invariance and transfer of motor skills across type and quantity of objects, the low-level controller is

not provided direct access to object state, even while trained from demonstrations containing objects.

We apply our approach to two challenging tasks, both involving a humanoid interacting bimanually with large objects such as boxes and medicine balls. The two tasks are an instructed box manipulation task in which the simulated character needs to follow user-specified instructions and move boxes between shelves (a simplified “warehouse” setting) and a ball catching and tossing task (“toss”). Both tasks are solved either from task features or egocentric vision by the same motor module (albeit different task policies) demonstrating the possibility of general and reusable motor skills that can be deployed in diverse settings. The goal directed training leads to the robust deployment of locomotion and manipulation skills as well as active control of the head for gaze direction. The results demonstrate the flexibility and generality of the approach, which achieves significant generalization beyond the raw demonstrations that the system was bootstrapped from, and constitute another step towards general learning schemes for sophisticated whole-body control in simulated physical environments.

2 RELATED WORK

Research aiming to achieve coordinated movement of virtual humanoids occurs in artificial intelligence and physics-based character animation, and advances in these fields relate to developments in robotics and biological motor control.

Virtual humanoid control. Realistic animation of humanoid characters is a major, enduring aim in the graphics community. For generalized locomotion, various kinematic approaches sequentially set the pose of the body according to motion capture snippets [Lee et al. 2002; Arikan and Forsyth 2002; Kovar et al. 2008], with behavioral sequencing resulting from approaches ranging from high-level planning [Levine et al. 2011] to timestep-level control by neural network sequence models [Holden et al. 2017]. A parallel line of research, that we focus on here, employs physics-based simulation to control a virtual body and produce realistic movement [e.g. Van de Panne and Fiume 1993; Faloutsos et al. 2001], often also using motion capture as a reference for the controller [e.g. Yin et al. 2007; Liu et al. 2010; Peng et al. 2018]. For both kinematic and physics-based approaches, the aim is essentially to re-sequence or schedule the available movements to produce appropriate combinations; this is adequate for settings that require limited movement diversity. In the context of dynamic locomotion, impressive results have been obtained by sequencing a small set of movements [Liu et al. 2012]. Complementary to the use of motion capture, Deep RL approaches enable learning from scratch [Peng et al. 2016; Heess et al. 2017], with behavior emerging from the constraints of the physics and task. Task-driven learning can also occur in combination with motion capture which constrains behavior to be more humanlike [Peng et al. 2017, 2018; Merel et al. 2017]. Recent approaches aim to more tightly couple kinematic and physics-based approaches by using kinematic approaches to specify targets for physics-based controllers [Chentanez et al. 2018; Park et al. 2019; Bergamin et al. 2019].

While locomotion is increasingly tractable in graphics settings using either kinematic or physics-based approaches, coordinated locomotion and manipulation remains more challenging. Recent

kinematic-based approaches show impressive interactions between a humanoid and the environment [Starke et al. 2019], but this requires modeling the motion as an explicit function of a high-dimensional volumetric scene context. For highly dynamic movements the precise coupling between motion and scene elements can be difficult to model accurately [Liu and Hodgins 2018]. For many physics-based graphics approaches involving object interactions, a shortcut is often taken by forming fixed attachments between the hands of the body and the manipulated object in order to simplify the problem [Coros et al. 2010; Mordatch et al. 2012; Liu et al. 2012; Peng et al. 2019]. Nevertheless, in addition to robotics-oriented efforts [e.g. Sentis and Khatib 2005; Otani and Bouyarmane 2017], there are instances of physically simulated object interactions, such as a humanoid sitting in a chair [Chao et al. 2019]. Another particularly impressive exception involves dribbling a basketball simulated with physically plausible object interactions [Liu and Hodgins 2018]; however, this work, while impressive, produces a dribbling controller that follows a relatively narrow trajectory. Narrow behaviors may be appropriate for certain animation settings, or even medical applications in which musculo-tendon models can help characterize the impact of surgical interventions [Holzbaur et al. 2005; Lee et al. 2019]. But generality of behavior may be important for settings that prioritize autonomy, such as autonomous virtual characters or robotics settings.

Finally, we note that a necessary component of an integrated embodied system is control of head and gaze. The animation and graphics literature has put considerable effort into making the eye movements of virtual characters appear natural [Ruhland et al. 2015; Pejsa et al. 2016], including by recent data-driven approaches [Klein et al. 2019]. In the context of the present work, the active gaze control serves a functional role in a visuomotor task, and body movements are informed by inferred state [see also Terzopoulos and Rabie 1995; Sprague et al. 2007]. In graphics research, this kind of gaze control has been considered, for example, to enable catching behaviors requiring upper body movement [Yeo et al. 2012; Nakada et al. 2018; Eom et al. 2019] as well as visually guided locomotion [Eom et al. 2019], generally using specially engineered gaze-control systems. In our specific setting, we do not model eyes, but the agent can learn to control its gaze via head movements in an emergent fashion in order to support task performance.

Generalizing from demonstrations. Motion capture for virtual character control in graphics is one particular instance of the more general problem of leveraging demonstrations for control. Demonstrations can be readily obtained for many simple real or simulated robotics systems, for instance through *teleoperation* or via a human operator physically guiding the pose of the robot. The classical approach for learning from demonstrations amounts to using the demonstration to initialize the policy, and learning how to deviate from the demonstrate to solve the task at hand [Smart and Kaelbling 2002; Schaal et al. 2003]. It has long been recognized that given a small number of demonstrations, it is not sufficient to try to directly mimic the demonstration as there will be some discrepancies when recapitulating the movement which will compound and lead to failure [Atkeson and Schaal 1997; Schaal 1997]. A fairly direct approach involves fitting the demonstrations to a parametric form and using

RL to modulate the parameters of the fitted model [Guenther et al. 2007; Peters and Schaal 2008; Kober and Peters 2009; Pastor et al. 2011]. Slightly more indirect approaches consist of using the demonstrations to learn local models from which a policy can be derived [Coates et al. 2008] or using the demonstrations to infer the objective for the policy through inverse optimal control [Ng and Russell 2000; Ho and Ermon 2016; Englert and Toussaint 2018]. In Deep RL settings involving a replay buffer, and when the demonstrations include actions and reward on the task being solved, it is possible to fill the replay buffer with teleoperation demonstrations [Večerík et al. 2017]. Finally, as noted previously, there are approaches in which both matching the demonstrations and solving the task serve as rewards during training [Kumar et al. 2016; Peng et al. 2018; Merel et al. 2017; Zhu et al. 2018]. Note that insofar as the main role of demonstrations is to serve as a form of prior knowledge, a similarly motivated approach is to design controllers that incorporate domain knowledge for some tasks and to then use learning to refine the behavior around this initial, engineered policy – this scheme has been applied to efforts in robotics for tossing objects [Zeng et al. 2019] and catching objects [Kim et al. 2014].

The commonality across this broad class of existing approaches for learning from demonstrations, both in virtual and robotic settings, is that they are well suited primarily when there is a single variety of movement that needs to be reproduced and where the demonstrations are well aligned with the behavior required to solve the task. While these approaches have been successful in various cases, they have not yet been demonstrated for more complex tasks that require composition and arbitrary re-sequencing of motor skills. Ideally we wish for a skill space to serve both as a generic “initialization” of the policy as well as a set of constraints on the behavior; yet we also want the skill space to be multipotent, in the sense that it can be leveraged for multiple distinct classes of target tasks, rather than serve only for a narrow range of movements. While some work has aimed to build motor skill modules from unstructured demonstrations [Jenkins and Mataric 2003; Niekum et al. 2012], limited work to date has aimed to learn flexible skill modules of the sort suitable for Deep RL [Merel et al. 2019c; Peng et al. 2019]. It remains open how best to generalize beyond individual trajectories and strike a balance between realism of the movements and the degree to which new movements can be synthesized from finite demonstrations.

3 APPROACH

In this work, we develop an approach for skill transfer and learning from demonstrations in the setting of visually-guided humanoid control with object interactions. By “skill transfer”, we refer to a setting which involves establishing basic motor competency on a source distribution of demonstrations or tasks and applying this prior knowledge to improve learning on new tasks. Our approach employs an episodic reinforcement learning (RL) paradigm, in which tasks are defined for a physical scene by specifying initial conditions, termination conditions, and a task-specific per timestep reward (r_t). A learning agent maintains a policy π that produces actions (a_t) in response to partial observations that it receives of the underlying physical state (s_t) of the environment. The agent aims to learn parameters of the policy so that it maximizes the expected

Fig. 2. **Overview of producing and reusing skills.** Stage 1 involves training a large set of separate, single-behavior “expert” policies. For each motion capture clip trajectory (depicted by a curve), we produce a policy that tracks that trajectory. Stage 2 involves distilling the experts produced in stage 1 into a single inverse model architecture. The inverse model receives the state for a few future steps (from $t + 1$ to $t + k$ where $k = 5$ in our setting), an encoder embeds this into a latent intention (z_t), and the decoder produces the action that will achieve the transition from s_t to s_{t+1} . Stage 3 involves training only the task policy to reuse the frozen low-level skill module, using the learned embedding space to communicate what to do.

sum of discounted future rewards during rollouts from the policy, $\mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t r_t]$. We propose and evaluate an approach that leverages unlabelled demonstrations (without rewards) and creates a *motor module* or low-level controller, which can then be used for multiple object-interaction tasks. This approach can also be seen as incorporating additional, previously learned structure into the policy to facilitate subsequent learning. We emphasize that while training a policy from scratch on motor problems is difficult insofar as RL-based policy optimization only obtains local optima, approaches that constrain the search space can result in more tractable optimization and yield better solutions.

The general workflow for producing and reusing the skill module is depicted in figure 2. It consists of three stages. Firstly, expert neural-network policies are generated which are capable of robustly tracking individual motion capture clips in the presence of noise. The second stage consists of distilling these policies into a single conditional policy, or inverse model, which maps the state at the current timestep (s_t) and the desired state at timesteps in the near future ($s_{t+1} \dots t+k$) to the first action (a_t) of a sequence of actions that would result in that desired future. Providing a horizon into the future (i.e., $k > 1$, as opposed to only providing s_{t+1}) is potentially useful in disambiguating the instruction for movements that require multiple timesteps of preparation to perform. As explained in more detail in Section 3.3, this inverse model is separated into an encoder and decoder, which communicate via a multi-dimensional, continuous random variable that reflects short term motor intention. The decoder can also be interpreted as a conditional policy that is trained via a form of behavioral cloning. This training procedure generally follows the approach of [Merel et al. 2019c], and so we refer to this architecture as “Neural Probabilistic Motor Primitives” (NPMP). Finally, the third stage amounts to reusing the NPMP decoder as

a low-level controller in the context of new tasks, by treating the learned motor intention space as an action space for a new controller. Here, a high-level task policy receives observations that are appropriate for the target task, either vision or states of relevant objects relative to the body, and this policy is trained by model-free RL. The task policy outputs “actions” corresponding to latent variables that serve as “commands” to the now fixed low-level module. The low-level module thus transforms the initial noise distribution of an untrained task policy into “colored”-noise that reflects the coordinated movement statistics of the motion capture data. By acting through the low-level controller, movement exploration as well as resulting solutions are constrained to the manifold of human-like behavior that can be produced by the motor module.

The particular challenges of the manipulation tasks considered in this work mean that several additional elements of the training process are critical. Manipulation requires directed interaction with objects in environment, and these are difficult to discover even when exploration is restricted to the space of movements expressed by the skill module. We address this by employing a suitable distribution of initial configurations of the body and objects, along with variations for object masses and sizes. Taken together, these initializations and variations facilitate learning through exposure to scene configurations that vary in difficulty and distance from reward, offering an organic curriculum. Finally, while we find that the skill module is somewhat robust to variations in which expert demonstrations are included; there is a trade-off between specificity and generality of skills. We elaborate on these elements of our proposed training process when presenting the tasks, and we demonstrate their relevance through ablations (see Results, Section 4.3).

All simulations are performed using the MuJoCo physics simulator [Todorov et al. 2012]. The humanoid body has 56 actuated

degrees of freedom, and was adapted from the standard version of this body that is available as part of the DeepMind Control codebase [Tassa et al. 2018]. Similarly to previous work using this body, we use position-control actuators that are limited to produce reasonable maximum torques. The body also includes force sensors in the shoulder as well as multiple binary touch/contact sensors on each hand. For this work, all body lengths are scaled from the available body according to the measured dimensions of the motion capture subject who performed reference movements.

3.1 Demonstrations for skills

We collected motion capture data of a person performing bimanual, whole-body box manipulation movements, ball tossing, and various locomotor behaviors with and without objects in-hand. When objects were involved, we also collected motion capture for the objects. To go from point-cloud, raw motion capture data to body-specific movements, we implemented simultaneous tracking and calibration [Wu et al. 2013], which solves a joint optimization problem over body pose and marker position. See Supplementary Section A for implementation details. Figure 3 shows a visualization of the virtualized “props” and humanoid set to poses from the motion capture. As noted above, it was important that we measured and re-sized the lengths of body segments of the virtual character to correspond to the person whose motion data we collected, as this ensured that the positions of the hands relative to tracked objects are similar in the virtual environments relative to the real setting. This precision in body dimensions also made the inference of body poses from point-clouds more robust. Nevertheless, the proportions of the virtual humanoid still only approximately correspond to the human actor and the dynamic properties differ substantially.

The dataset collected for this work consists of a single subject interacting with 8 objects (or “props”). The objects are two “large”

balls, two “small” balls, two “large” boxes, and two “small” boxes. Small objects weighed 3kg and large objects 10kg. We considered interactions at 3 heights, “floor-height”, “torso-height”, and “head-height”. For each object, at each height, we collected two repeats of behavior consisting of the actor approaching a pedestal on which an object is resting, picking it up, walking around with the object in hand, returning to the pedestal, placing the object back on the pedestal, and then backing away from the pedestal. In total this amounts to 48 clips (8 objects $\times$ 2 repeats $\times$ 3 heights), each of which is generally no less than 10 seconds and no longer than just over 20 seconds. Other less structured behavior was captured, including walking around with no object (“walking”) as well as tossing a ball with a second person (“ball-tossing”; one person and the ball were tracked). In total, we use a little less than 20 min of data (~ 1130 sec). For representative examples, see videos of motion capture playback: box interaction [Video 2](#) and ball tossing [Video 3](#).

3.2 Single-clip tracking for object manipulation

To produce expert policies, we use a tracking objective and train time-indexed policies to reproduce the movements observed via motion capture [Peng et al. 2018], here including the position of the object. Similarly to [Merel et al. 2019a], we provide the agent a normalized tracking reward, $r_t \in (0, 1]$, that reflects how well the body and object in the virtual environment match the reference:

$$r_t = \exp(-\beta E_{\text{total}}/w_{\text{total}}) \quad (1)$$

where w_{total} is the sum of the per energy-term weights and β is a sharpness parameter ($\beta = 10$ throughout). The energy term is a sum of tracking terms, each of which corresponds to a distance between the pose of the physically simulated body relative to a reference trajectory derived from motion capture:

$$E_{\text{total}} = w_{\text{qpos}}E_{\text{qpos}} + w_{\text{qvel}}E_{\text{qvel}} + w_{\text{ori}}E_{\text{ori}} + w_{\text{app}}E_{\text{app}} + w_{\text{vel}}E_{\text{vel}} + w_{\text{gyro}}E_{\text{gyro}} + w_{\text{obj}}E_{\text{obj}} \quad (2)$$

with terms for tracking the reference joint angles (E_{qpos}), joint velocities (E_{qvel}), root quaternion (E_{ori}), body-frame vectors from the root to appendages (hands, feet, head; E_{app}), translational velocity (E_{vel}), root rotational velocities (E_{gyro}) and object position (E_{obj}). See Supplementary Section B for more specific details. Note that to encourage robustness of the controller, we train in the presence of moderate action noise – noise is sampled from a Gaussian independently per actuator with $\sigma = .1$, for actions $\in [-1, 1]$.

Using the objective described above, we produce expert policies for all reference motions (at similar sample complexity to previous work). Each expert is trained to track motion capture clips that are roughly 3-5s snippets of the raw reference clips, yielding a few hundred varied experts that tile the behaviors discussed in the previous section. To assess performance of the expert tracking controllers, we collect rollouts starting from different points along the trajectory and see how well these align with the motion capture reference. In general we find that tracking performance is good for “warehouse” behavior experts (Figure 4A) with only a small falloff as a function of duration of rollout (Figure 4B). For “toss” behavior experts, performance sometimes shows a sharp fall-off after tossing the ball (Figure 4C). However, this performance decline is primarily due to the object tracking term when the ball is no longer directly controlled

Fig. 3. **Motion capture.** (A) Virtual analogs of the objects that are tracked with motion capture. (B) & (C) Frames of motion capture for box interaction and ball tossing, with prop and humanoid body set to those poses in the physics simulator. The green dots correspond to markers.

Fig. 4. **Tracking with expert policies.** (A) Tracking performance and corresponding filmstrip is shown for a representative warehouse expert clip. We initialize the expert at timepoints throughout the clip and the policy controls the behavior to end of the clip. Time within a rollout since the initialization is depicted via intensity. The policy is robust in that it controls the body (interacting with the box) to remain on track. (B) Representative summary of the expert tracking performance for rollouts as a function of time since initialization for all medium height experts – rollouts lasting up to 3 seconds show only limited accumulated tracking error, indicating experts are well-tracked. (C) Performance and filmstrip for a ball-tossing expert indicating good tracking performance until the ball is released from the hands at which point the performance deteriorates due largely to the loss of control over the ball. Despite the inability to control the ball to match the reference ball trajectory perfectly, visually, the expert looks reasonable through the release of the ball.

and does not reflect failure of body tracking, as visually discernible from the filmstrip in Figure 4C. As a data augmentation, we also produced “mime” experts for which the expert was trained to track the human reference movements involving object interactions, but for which the object was not present in the virtual environment. Inclusion of these mime experts helped balance the data used for distillation, which otherwise overrepresented object carrying movements, resulting in controllers that are overly predisposed to bring the hands of the humanoid together.

3.3 Training the motor module for locomotion and manipulation

The single-behavior expert policies track individual motion capture trajectories but do not directly generalize to new tasks or even configurations of the environment. To enable reusability of the skills, we therefore follow [Merel et al. 2019c] and distill expert behaviors

into a single module with suitable architecture (the “Neural Probabilistic Motor Primitives”, or NPMP). Unlike [Merel et al. 2019c] we are interested in manipulation skills which depend strongly on the environment, not just the body controlled by the agent. As a critical design choice that ensures usability of the motor module across various environments, we employ the following factorization: during training, we give the encoder access to the state both of the humanoid as well as the object used in the expert trajectory; however, the decoder only directly receives egocentric humanoid proprioceptive information. By construction, the decoder will therefore be reusable as a policy that only requires egocentric observations of the humanoid body, which are inherently consistent across environments. When reusing the skill module, any awareness of the objects in the scene must be passed to the low-level controller via the latent variable produced by the task policy.

The training procedure for the motor module follows the approach presented in [Merel et al. 2019c]. We train the model in a supervised fashion to model state-action sequences (trajectories) generated by executing the various single-skill experts policies, while adding independent Gaussian noise to the actions. Specifically, we maximize the Evidence Lower Bound (ELBO):

$$\mathbb{E}_q \left[\sum_{t=1}^T \log \pi(a_t | s_t, z_t) + \beta (\log p_z(z_t | z_{t-1}) - \log q(z_t | z_{t-1}, s_{t+1 \dots t+k})) \right], \quad (3)$$

where $\pi(a_t | s_t, z_t)$ corresponds to the decoder, which is a policy conditioned on the latent variable z_t and the current state s_t . The distribution $q(z_t | z_{t-1}, s_{t+1 \dots t+k})$ corresponds to the encoder and produces latent embeddings based on short snippets into the future. The latent variable z_t carries the semantics of motor intention, inherited from training on the behavior distribution of the experts, and this space is therefore called the *skill embedding* space. This latent z will now serve as a task-relevant instruction to the decoder-policy. β controls the weight of the autoregressive prior $p_z(z_t | z_{t-1})$, which regularizes the skill embedding space by encouraging z ’s temporal continuity. For further analysis of the skills well-reflected in this space, see Supplementary Section C.

3.4 Training task policies that reuse low-level skills

To train the task policy, which reuses the low-level skill module, we use a model-free distributed RL setup with a single learner and many actors (here 1000). This training paradigm as well as our use of off-policy correction via V-trace to train the value-function are consistent with IMPALA [Espeholt et al. 2018]. Policy updates are performed using a version of V-MPO [Song et al. 2020], which is in the MPO family of approaches [Abdolmaleki et al. 2018]; however, while V-MPO proposes the use of on-policy updates, we perform updates from a replay buffer, which we found stabilized learning in our setting. Given these choices, the relevant learning parameters are the learning rate (1e-4), the MPO ϵ that controls the KL-divergence constraint on the policy (we swept .5 and 1.), and discount factor ($\gamma = 0.99$). In addition, the policy and value functions are updated with minibatches of size 128 and trajectory length of 50, from a replay buffer.

Fig. 5. **High-level task policy architecture.** When training a high-level task policy to reuse the low-level controller, three streams of input are available potentially. For whichever of the egocentric image input, task instruction input, and proprioception streams are available, each is passed through a preprocessor network. A value function and policy branch from a shared LSTM, and the policy also receives skip connections for the task and proprioception input streams. The policy output here refers to high-level actions that serve as inputs to the low-level controller.

The architecture we used for the task policies is close to the simplest architecture that is suited for the setting (see Figure 5 for a schematic). We separately encode the qualitatively different streams of inputs with separate preprocessing networks – small 1-2 hidden-layer MLPs for the task and proprioceptive features and a generic ResNet for image inputs [He et al. 2016]. The outputs of these three input channels are concatenated and used as inputs to an LSTM [Hochreiter and Schmidhuber 1997]. The shared LSTM branches into the value function output as well as a second LSTM which is used for the policy. By having the first shared LSTM, learned representations that are useful both for the value function and policy can be shared. We did not extensively tune the network structure.

We want the task policy to produce actions consistent with the values seen by the low-level skill module during its supervised training (the embedding space is regularized to have values close to zero). To prevent actions from the task policy from being too “out-of-distribution” for the pretrained low-level controller, we restrict the high-level actions to a limited range of values, in the range $(-2, 2)$.

4 RESULTS

4.1 Core tasks

In this work, we defined two challenging object interaction tasks, and we show that the low-level skill module can be used to solve either of these, when a high-level, task-specific policy is trained to reuse the skills on each task. Our two core tasks are a proto-warehouse task (“warehouse”) and a ball tossing task (“toss”). The warehouse task involves going to a box that is on a pedestal, picking up the box, bringing it to another pedestal, putting the box down, and repeating. To make the task unambiguous (e.g. whether the current goal is to pick up or put down), we provide the agent with a task “phase” or “instruction” that indicates which of these four phases of the task the agent is presently in. This phase also provides a natural way of providing sub-goals, insofar as sparse rewards are provided after each phase of the task has been completed. In addition, we provide the agent with the position (relative to itself) of the target pedestal (to which it must go in order to pick up a box, or to put down a box). This active pedestal is highlighted in the

Fig. 6. **Tasks.** (A) The “warehouse” task involving instructed movements of indicated boxes from one indicated pedestal to another. (B) The “toss” task involving catching a ball thrown towards the humanoid and tossing it into a bucket on the ground. Tasks can be performed either from state features or from egocentric vision.

videos, and this visual cue is also available to vision-based agents. See Supplementary Section D for further details about the task specification.

Our second task consists of catching a ball and then tossing it into a bucket. In this task, the ball is always initially thrown towards the humanoid. The task is terminated with a negative reward if the ball touches the ground, which incentivizes the agent to learn to catch the ball and avoid dropping it. A small shaping reward encourages the agent to bring the ball towards the bucket, and a sparse positive reward is provided if the ball is deposited into the bucket. See Supplementary Section E for details.

Both tasks are generated procedurally, with several task parameters sampled from a distribution on a per-episode basis. For the warehouse the pedestal heights, box dimensions, and box masses are each sampled from a distribution. For the tossing task, the ball size, mass, the trajectory of the ball thrown towards the humanoid, and the position of the bucket are each sampled from a distribution. In both tasks, mass variations are visualized by object color (darker is heavier). In the warehouse task, we also initialize episodes in the various phases of the task and sample initial poses of the body from the motion capture data. Again, see Supplementary Sections D and E for more details. These task variations and initializations are important for successful training of the task policies as we will show in Section 4.3. To observe the environment, agents are provided either

visual information (an egocentric camera mounted on the head of the humanoid) or state features which consist of the position of the prop relative to the humanoid as well as the orientation of the prop, and we compare performance using these different observations.

4.2 Performance on tasks

We train task policies that operate from state and visual observations on both tasks and found that successful reuse is possible using either observation type. However, note that comparable experiments from vision require longer walltime, since rendering image observations slows the simulation. On the warehouse task, visual information seems to improve learning (Figure 7A), whereas state information is better on the toss task (Figure 8A). This discrepancy between the tasks is perhaps explicable – box interaction from state may be performance limited by the available features. Indeed, given that state features only consist of center-of-box position and orientation, precise movements that require sensory access to the edges and faces of the box relative to body and hands may be more natural to learn from visual inputs (from which this information may be more apparent). However, in the toss task, the same state features may be more adequate for an optimal policy. Nevertheless, both policies using either feature set trained to a reasonable performance level. For representative performance and behavior of the vision based policies, see the representative “warehouse” task [Video 4](#) and “toss” task [Video 5](#). When task policies operate from vision, the agent must learn to coordinate the body to interact with the objects that it senses and interprets from the image inputs. In the warehouse task, the body and head movements lead to a somewhat jerky egocentric camera, but evidently the policy is still able to leverage this visual stream for successful object interaction. In the toss task, the camera control is more stable and it seems intuitive that the visual stream can support object tracking for interaction, though perhaps with some ambiguity of ball size vs ball distance. For all results in this work, any object state estimation or gaze control is learned implicitly, without any additional structure to encourage its emergence.

Note that without reusable motor skills, an alternative is to learn the task from scratch. This is difficult, as rewards in these tasks are sparse and therefore do not very strongly shape behavior. And critically, in these tasks it would be very difficult to design dense rewards that incentivize the right kind of structured behavior. This being said, it did turn out to be possible to learning from scratch on the toss task from state information. After training for an order of magnitude longer ($>100e9$ learning steps), the early terminations with penalty and shaping reward are sufficient to produce a policy that could solve the toss task, albeit without actually catching the ball, essentially using its back as a paddle. This same behavior was consistently learned across multiple seeds, see [Video 6](#). For the warehouse task, training from scratch for equivalently long did not yield behavior that could solve a whole cycle of the task, but some progress was made for certain initial conditions, see [Video 7](#). Note that for experiments from vision (slower than experiments from state), training a policy for $150e9$ steps took roughly 3 weeks of wall-clock time, so it was not feasible to systematically explore training from scratch for significantly longer intervals.

For the warehouse task, we provide an additional evaluative visualization of the final performance to provide a clearer sense of the quality of the learned solution. For any stage of the behavior, we can take a trained agent and assess how reliably it can perform a given behavior from different initial positions. We defined a 9×9 grid of initial x-y locations in the plane. For each location we initialized the humanoid there 10 times, randomizing over orientation, body configuration (sampled from motion capture), and initial velocity. We then computed the fraction of trials for which the humanoid was able to successfully pick up a prop. We visualize a top down view, with the agent aiming to pick up the prop located on the pedestal on the right side of the top down view, with the heatmap of success probability overlain (Figure 7B). The agent is generally robust to initial position of the humanoid; only a limited fraction of initializations are too close to the pedestal and lead to failures, presumably due to initial poses or velocities that make it especially difficult.

For the toss task, we similarly wanted to provide a statistical description of the core behavior of the trained agent. We discretized the space of initial ball velocities (both towards the humanoid and horizontally relative to the agent) – consistent with training, we computed an initial vertical velocity such that ball would be approximately shoulder height when near the initial position of the humanoid. We initialized the ball velocity for 10 repeats in each bin of the velocity, randomizing over other variations. The heatmap depicted in Figure 8B indicates the “strike zone” of parameters for which the agent is able to catch the ball. Naturally, for initial velocities that are too horizontal it is simply not possible to catch the ball and probability of success falls off to zero (indicated by episode return of -1 , corresponding to the ball hitting the ground).

We also remark that visual “quality” does not entirely align with performing the task optimally. Throughout the course of our research, we noticed that slightly worse optimizers or termination

Fig. 7. **Performance for the “warehouse” task:** (A) Representative learning curves (best of 3 seeds) comparing vision-based and state-based performance on the warehouse task, as a function of learner update steps. (B) For the trained vision-based policy, heatmap overlain on top-down view visualizes probability of successful pickup as a function of initial location. (C) Representative filmstrip of behavior in the warehouse task from egocentric and side view.

Fig. 8. **Performance for the “toss” task:** (A) Representative learning curves (best of 3 seeds) comparing vision-based and state-based performance on the toss task, as a function of learner update steps. (B) For the trained state-based policy, heatmap indicates the episode return as a function of initial ball velocity. (C) Representative filmstrip of behavior in the warehouse task from egocentric and side view.

partway through training resulted in policies that are more conservative, scored fewer points, but might subjectively be considered to look more humanlike (less hurried). This is consistent with the humanlikeness of the movements being determined largely by the low-level controller, but performance of the task becoming increasingly hectic as the task policy ultimately controls the body to move faster and with extreme movements to achieve more reward. For an example of a policy trained for less time, see [Video 8](#).

4.3 Task variations

In addition to demonstrating performance on the core tasks, a few themes emerged in developing our approach, for which we provide illustrative examples. In particular, some trends that we observed include: (1) the ratio of expert skills in the NPMP matter, (2) the initializations at different phases of the task matter for the warehouse task, but aren’t required for the toss task, & (3) more extreme variations benefit from a curriculum via variations (a similar result is reported in [Heess et al. 2017]).

First, we consider how important the relative ratios of different skills are in the NPMP. In extreme cases, this is trivially important. For example, an NPMP module that only contained locomotion skills, without object interactions, would intuitively offer limited utility for transfer to the warehouse task. A more nuanced question is how important the relative quantities of ball tossing behavior versus warehouse behavior affect the ability of the NPMP to learn the two tasks. For illustration, we trained three NPMPs, one that only had access to warehouse experts, one that had both warehouse experts and ball toss experts in proportion to how much was collected (more motion capture was warehouse relative to toss demonstrations), and one that trained on twice as much data from toss experts – that is, when training the NPMP, we recorded twice as many trajectories from ball toss experts as we did for other experts, thereby over-representing these experts in the NPMP training data. Note that in the toss upsampled NPMP, toss experts are over-represented

relative to our motion capture, but there was still more warehouse data relative to toss data even with this upsampling. We observed that while the upsampled toss NPMP learned an arguably slightly more aesthetically satisfying toss behavior (no meaningful change in performance), it was more difficult for the upsampled toss NPMP to learn the warehouse task. In figure 9A, we show comparisons of these different NPMPs on the warehouse task. While ultimately, the upsampled toss NPMP was able to learn the warehouse task, it was consistently lower and less robust for other hyperparameters.

In addition to the balance of expert data, we also examined the need to initialize the behavior in different phases of the warehouse task. In the warehouse task, the training episodes are initialized in all phases of the task in poses sampled from motion capture, forming a curriculum over variations. In the toss task, we did not initialize episodes in different task phases. To illustrate the need for starting the warehouse task in the various phases, we ran comparisons involving starting only in the pickup or walk phases of the task and found that neither of these are able to learn to solve the task (see figure 9B).

We also explored both decreasing and increasing the range of procedural variations across episodes. Based on previous work [Heess et al. 2017], it was our starting intuition to design the task with a sensible range of variations to facilitate learning – this meant that our initial distribution of variations basically worked. However, we also attempted to train the task policy to perform the task with only large boxes. We probed the original task policy trained on variable box size on variants of the warehouse task that only included larger boxes, and we see that training with variations improves performance on the probe task (figure 9C). Essentially, it was much more difficult to learn to solve this task without the variations in box size

Fig. 9. **Ablations and probes.** (A) Performance on the warehouse task as a function of using NPMPs trained with different ratios of expert data. “Mixed” uses natural proportions of all of our warehouse and ball tossing data, “no toss” omits all tossing data, and “toss++” uses ball toss experts upsampled by a factor of two. (B) Comparison of performance on the warehouse task under the default setting involving initializations at “all” phases of the task versus training with initializations restricted to either the “pickup” or “walk” phases. (C) We trained the task policy to perform the warehouse task with either the baseline variation in box sizes (blue) or only a smaller range of large boxes (orange) and then evaluated performance only interacting with large boxes. Variation on the wider distribution during training improved performance on the evaluation tasks.

– no policy fully solved the task. For representative failure mode behavior, see [Video 9](#).

Finally, we considered training on a wider distribution than our standard range of pedestal heights and this tended to work – this indicates that a broader, continuous task distribution could allow a policy to perform a wide range of movements, so long as exploration and learning are guided from some examples that are initially achievable. See a representative [Video 10](#) showing performance when trained on this broader range of pedestal heights, including pedestals that are quite low to the ground as well as higher up.

5 DISCUSSION

In this work, we demonstrated an approach for transfer of motor skills involving whole body humanoid movement and object interaction. We showed that a relatively small set of demonstration data can be used to provide a fairly generic low-level motor skill space that can be leveraged to improve exploration and learning on various tasks that can be solved via movements similar to those in the expert demonstrations. Importantly, a single skill module is multipotent, permitting reuse on multiple transfer tasks. The resulting task controllers can follow high-level goals and adapt naturally to changes or interactive perturbations of the scene and object details (see [Video 1](#)). These interactions with trained controllers are enabled by the ability to deploy trained controllers in real-time on a standard computer. Critically, a physics-based controller always responds to a scene in a manner that is consistent with the physics of the world, unlike kinematic approaches which can generalize in physically impossible ways. Furthermore, the use of egocentric vision induces natural behavioral phenomena such as directing the gaze towards objects of relevance; but perhaps more significantly, it makes task-specific scene representations unnecessary. The agent can understand what to do by looking at the scene.

What differentiates our approach from most preceding work that leverages demonstrations for physics-based control is that we synthesize a single, multipotent skill module. Instead of having to stay close to demonstrations of a single object interaction, we provide a large set of unlabeled demonstrations and automatically generalize skills from them. One open question that we believe will be important in future efforts involves how to best trade off the specificity of exploration, provided by staying close to a narrow set of demonstrations, versus generality obtained through leveraging more diverse demonstrations. As we showed, there is still some sensitivity to the relative ratios of the various skills in the space. It may be fundamental that there is some trade-off between exploration guidance and generality, and it would be interesting to better understand this in high-dimensional control settings.

A similar trade-off currently exists between generality and visual quality. We note that in this work, the task objectives are quite basic – essentially providing sparse indications of task progress. Where the controller generalizes beyond the motion capture, the behavior may look slightly less natural. We are optimistic that further refinement of the movement appearance could result from additional objectives (e.g. smoothness or energy regularization) when training the task policies, additional data that helps more completely cover the space of relevant behaviors, and adjustments to the physics of the body.

We have outlined a generic architectural scheme for motor reuse, within which it is possible to implement the encoder and decoder using various specific neural networks. We have elected to use the simple choice of shallow MLPs (both our encoder and decoder have two hidden layers). Conceptually, we view the specific choice of network as an important, but secondary, consideration in the context of the overall approach. Network engineering may further improve this scheme, for instance through either faster learning or in terms of motion quality. In particular, Peng et al. 2019 have proposed an alternative architecture that could be applied in this setting involving additional multiplicative structure and composable primitives. Although Peng et al. 2019 reported that their multiplicative compositional policies (MCP) network structure performed better than an MLP in their settings, our own exploratory investigation of the reported MCP architecture (using 8 primitives) did not reproduce this benefit. Instead, we found that while MCP reuse did produce generally similar movements in early stages of learning, it neither took off faster nor achieved effective performance in the warehouse setting. We speculate that this difference may have to do with the necessary degrees of freedom for the task – in the warehouse setting, limiting the reusable control to compositions of a discrete set of primitives may impair their reusability outside of the movements that are highly represented in the reference data. Regardless, as a comparison of decoder architectures is not a focus of this work, we did not pursue this deeply, and we leave more systematic comparison of architectural choices to future work.

There are a few additional caveats concerning the present approach that are related to the difficulty of exploration and learning in complicated tasks. We train the task policies by model-free RL, and this is not efficient with respect to data (requiring many environment interactions and learner updates). Substantially, this slowness arises from dithering exploration at the level of the task policy. While we use low-level skills to structure exploration, this alone is not sufficient to learn to control high-dimensional bodies interacting with objects, especially from only sparse rewards. We are optimistic that future improvements will accelerate learning, perhaps partly through more intelligent, goal-directed exploration strategies. For example, the motif of using skills for exploration could potentially be repeated hierarchically to encourage richer behavior. Another limitation is that, for the warehouse task, we leverage a curriculum via informative motion capture initializations, which expose the agent to favorable states that it may not have discovered on its own. It is interesting to note that the use of initializations is not required for the ball toss, where the combination of the ball being thrown towards the humanoid (forcing it to engage) and a weak shaping reward to induce movement towards the box are adequate.

Taken together, these limitations restrict the present approach to simulation settings; however there is a growing literature on approaches involving transfer of policies trained in simulation to real world systems (sim-to-real) [Rusu et al. 2017; Sadeghi and Levine 2017; Tobin et al. 2017; Andrychowicz et al. 2018; Zhu et al. 2018; Tan et al. 2018; Hwangbo et al. 2019; Xie et al. 2019]. While we believe that the study of sophisticated motor control problems in simulation is an important area of research in its own right, sim-to-real may offer a path to translate these results into real world applications.

ACKNOWLEDGMENTS

We thank Tim Lillicrap for constructive input at the outset of the project, Vicky Langston for help coordinating the motion capture acquisition, Thomas Rothörl for assistance during our studio visit, and Audiomotion Studios for services related to motion capture collection and clean up. We also thank others at DeepMind for input and support throughout the project. Finally, we thank Jaakko Lehtinen for comments on the manuscript.

REFERENCES

- Abbas Abdolmaleki, Jost Tobias Springenberg, Yuval Tassa, Remi Munos, Nicolas Heess, and Martin Riedmiller. 2018. Maximum a posteriori policy optimisation. In *International Conference on Learning Representations*.
- Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. 2018. Learning dexterous in-hand manipulation. *arXiv preprint arXiv:1808.00177* (2018).
- Okan Arikan and David A Forsyth. 2002. Interactive motion generation from examples. In *ACM Transactions on Graphics (TOG)*, Vol. 21. ACM, 483–490.
- Ronald C Arkin, Ronald C Arkin, et al. 1998. *Behavior-based robotics*. MIT press.
- Christopher G Atkeson and Stefan Schaal. 1997. Robot learning from demonstration. In *ICML*, Vol. 97. Citeseer, 12–20.
- Kevin Bergamin, Simon Clavet, Daniel Holden, and James Richard Forbes. 2019. DRCon: data-driven responsive control of physics-based characters. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 206.
- Yu-Wei Chao, Jimei Yang, Weifeng Chen, and Jia Deng. 2019. Learning to sit: Synthesizing human-chair interactions via hierarchical control. *arXiv preprint arXiv:1908.07423* (2019).
- Nuttapong Chentanez, Matthias Müller, Miles Macklin, Viktor Makoviyshuk, and Stefan Jeschke. 2018. Physics-based motion capture imitation with deep reinforcement learning. In *Proceedings of the 11th Annual International Conference on Motion, Interaction, and Games*. 1–10.
- Adam Coates, Pieter Abbeel, and Andrew Y Ng. 2008. Learning for control from multiple demonstrations. In *Proceedings of the 25th international conference on Machine learning*. ACM, 144–151.
- Stelian Coros, Philippe Beaudoin, and Michiel Van de Panne. 2010. Generalized biped walking control. In *ACM Transactions on Graphics (TOG)*, Vol. 29. ACM, 130.
- Peter Englert and Marc Toussaint. 2018. Learning manipulation skills from a single demonstration. *The International Journal of Robotics Research* 37, 1 (2018), 137–154.
- Haegwang Eom, Daseong Han, Joseph S Shin, and Junyong Noh. 2019. Model Predictive Control with a Visuomotor System for Physics-based Character Animation. *ACM Transactions on Graphics (TOG)* 39, 1 (2019), 1–11.
- Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. 2018. IMPALA: Scalable Distributed Deep-RL with Importance Weighted Actor-Learner Architectures. In *International Conference on Machine Learning*. 1406–1415.
- Petros Faloutsos, Michiel van de Panne, and Demetri Terzopoulos. 2001. Composable Controllers for Physics-Based Character Animation. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. Association for Computing Machinery.
- Florent Guenter, Micha Hersch, Sylvain Calinon, and Aude Billard. 2007. Reinforcement learning for imitating constrained reaching movements. *Advanced Robotics* 21, 13 (2007), 1521–1544.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- Nicolas Heess, TB Dhruva, Srinivasan Sriram, Jay Lemmon, Josh Merel, Greg Wayne, Yuval Tassa, Tom Erez, Ziyu Wang, Ali Eslami, et al. 2017. Emergence of Locomotion Behaviours in Rich Environments. *arXiv preprint arXiv:1707.02286* (2017).
- Nicolas Heess, Greg Wayne, Yuval Tassa, Timothy Lillicrap, Martin Riedmiller, and David Silver. 2016. Learning and transfer of modulated locomotor controllers. *arXiv preprint arXiv:1610.05182* (2016).
- Jonathan Ho and Stefano Ermon. 2016. Generative adversarial imitation learning. In *Advances in neural information processing systems*. 4565–4573.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- Daniel Holden, Taku Komura, and Jun Saito. 2017. Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 42.
- Katherine RS Holzbaur, Wendy M Murray, and Scott L Delp. 2005. A model of the upper extremity for simulating musculoskeletal surgery and analyzing neuromuscular control. *Annals of biomedical engineering* 33, 6 (2005), 829–840.
- Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladen Koltun, and Marco Hutter. 2019. Learning agile and dynamic motor skills for legged robots. *Science Robotics* 4, 26 (2019), eaau5872.
- Odest Chadwicke Jenkins and Maja J Mataric. 2003. Automated derivation of behavior vocabularies for autonomous humanoid motion. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*. ACM, 225–232.
- Seungsu Kim, Ashwini Shukla, and Aude Billard. 2014. Catching objects in flight. *IEEE Transactions on Robotics* 30, 5 (2014), 1049–1065.
- Alex Klein, Zerrin Yumak, Arjen Beij, and A Frank van der Stappen. 2019. Data-driven Gaze Animation using Recurrent Neural Networks. In *Motion, Interaction and Games*. 1–11.
- Jens Kober and Jan R Peters. 2009. Policy search for motor primitives in robotics. In *Advances in neural information processing systems*. 849–856.
- Lucas Kovar, Michael Gleicher, and Frédéric Pighin. 2008. Motion graphs. In *ACM SIGGRAPH 2008 classes*. ACM, 51.
- Vikash Kumar, Abhishek Gupta, Emanuel Todorov, and Sergey Levine. 2016. Learning dexterous manipulation policies from experience and imitation. *arXiv preprint arXiv:1611.05095* (2016).
- Jehee Lee, Jinxiang Chai, Paul SA Reitsma, Jessica K Hodgins, and Nancy S Pollard. 2002. Interactive control of avatars animated with human motion data. In *ACM Transactions on Graphics (TOG)*, Vol. 21. ACM, 491–500.
- Seunghwan Lee, Moonseok Park, Kyoungmin Lee, and Jehee Lee. 2019. Scalable muscle-actuated human simulation and control. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 73.
- Sergey Levine, Yongjoon Lee, Vladen Koltun, and Zoran Popović. 2011. Space-time planning with parameterized locomotion controllers. *ACM Transactions on Graphics (TOG)* 30, 3 (2011), 1–11.
- Libin Liu and Jessica Hodgins. 2018. Learning basketball dribbling skills using trajectory optimization and deep reinforcement learning. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 142.
- Libin Liu, KangKang Yin, Michiel van de Panne, and Baining Guo. 2012. Terrain runner: control, parameterization, composition, and planning for highly dynamic motions. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 154.
- Libin Liu, KangKang Yin, Michiel van de Panne, Tianjia Shao, and Weiwei Xu. 2010. Sampling-based contact-rich motion control. In *ACM SIGGRAPH 2010 papers*. 1–10.
- Josh Merel, Arun Ahuja, Vu Pham, Saran Tunyasuvunakool, Siqi Liu, Dhruva Tirumala, Nicolas Heess, and Greg Wayne. 2019a. Hierarchical visuomotor control of humanoids. In *International Conference on Learning Representations*.
- Josh Merel, Matthew Botvinick, and Greg Wayne. 2019b. Hierarchical motor control in mammals and machines. *Nature Communications* 10, 1 (2019), 1–12.
- Josh Merel, Leonard Hasenclever, Alexandre Galashov, Arun Ahuja, Vu Pham, Greg Wayne, Yee Whye Teh, and Nicolas Heess. 2019c. Neural probabilistic motor primitives for humanoid control. In *International Conference on Learning Representations*.
- Josh Merel, Yuval Tassa, Sriram Srinivasan, Jay Lemmon, Ziyu Wang, Greg Wayne, and Nicolas Heess. 2017. Learning human behaviors from motion capture by adversarial imitation. *arXiv preprint arXiv:1707.02201* (2017).
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529.
- Igor Mordatch, Emanuel Todorov, and Zoran Popović. 2012. Discovery of complex behaviors through contact-invariant optimization. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 43.
- Masaki Nakada, Tao Zhou, Honglin Chen, Tomer Weiss, and Demetri Terzopoulos. 2018. Deep learning of biomimetic sensorimotor control for biomechanical human animation. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–15.
- Andrew Y Ng and Stuart J Russell. 2000. Algorithms for Inverse Reinforcement Learning. In *Proceedings of the Seventeenth International Conference on Machine Learning*. Morgan Kaufmann Publishers Inc., 663–670.
- Scott Niekum, Sarah Osentoski, George Konidaris, and Andrew G Barto. 2012. Learning and generalization of complex tasks from unstructured demonstrations. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 5239–5246.
- Kazuya Otani and Karim Bouyarmane. 2017. Adaptive whole-body manipulation in human-to-humanoid multi-contact motion retargeting. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*. IEEE, 446–453.
- Soohwan Park, Hoseok Ryu, Seyoung Lee, Sunmin Lee, and Jehee Lee. 2019. Learning predict-and-simulate policies from unorganized human motion data. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 205.
- Peter Pastor, Mrinal Kalakrishnan, Sachin Chitta, Evangelos Theodorou, and Stefan Schaal. 2011. Skill learning and task outcome prediction for manipulation. In *2011 IEEE International Conference on Robotics and Automation*. IEEE, 3828–3834.
- Tomislav Pejsa, Daniel Rakita, Bilge Mutlu, and Michael Gleicher. 2016. Authoring directed gaze for full-body motion capture. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 1–11.

Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. 2018. DeepMimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 143.

Xue Bin Peng, Glen Berseth, and Michiel Van de Panne. 2016. Terrain-adaptive locomotion skills using deep reinforcement learning. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 81.

Xue Bin Peng, Glen Berseth, KangKang Yin, and Michiel Van De Panne. 2017. DeepLoco: Dynamic locomotion skills using hierarchical deep reinforcement learning. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 41.

Xue Bin Peng, Michael Chang, Grace Zhang, Pieter Abbeel, and Sergey Levine. 2019. MCP: Learning Composable Hierarchical Control with Multiplicative Compositional Policies. *arXiv preprint arXiv:1905.09808* (2019).

Jan Peters and Stefan Schaal. 2008. Reinforcement learning of motor skills with policy gradients. *Neural networks* 21, 4 (2008), 682–697.

Rolf Pfeifer and Christian Scheier. 2001. *Understanding intelligence*. MIT Press.

Marc H Raibert and Jessica K Hodgins. 1991. Animation of dynamic legged locomotion. In *Proceedings of the 18th annual conference on Computer graphics and interactive techniques*. 349–358.

Kerstin Ruhland, Christopher E Peters, Sean Andrist, Jeremy B Badler, Norman I Badler, Michael Gleicher, Bilge Mutlu, and Rachel McDonnell. 2015. A review of eye gaze in virtual agents, social robotics and hci: Behaviour generation, user interaction and perception. In *Computer graphics forum*, Vol. 34. Wiley Online Library, 299–326.

Andrei A Rusu, Mel Večerík, Thomas Rothörl, Nicolas Heess, Razvan Pascanu, and Raia Hadsell. 2017. Sim-to-Real Robot Learning from Pixels with Progressive Nets. In *Conference on Robot Learning*. 262–270.

Fereshteh Sadeghi and Sergey Levine. 2017. CAD2RL: Real Single-Image Flight Without a Single Real Image. In *Robotics: Science and Systems*.

Stefan Schaal. 1997. Learning from demonstration. In *Advances in neural information processing systems*. 1040–1046.

Stefan Schaal, Auke Ijspeert, and Aude Billard. 2003. Computational approaches to motor learning by imitation. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 358, 1431 (2003), 537–547.

Luis Sentis and Oussama Khatib. 2005. Synthesis of whole-body behaviors through hierarchical control of behavioral primitives. *International Journal of Humanoid Robotics* 2, 04 (2005), 505–518.

David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharmashan Kumaran, Thore Graepel, et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 362, 6419 (2018), 1140–1144.

William D Smart and L Pack Kaelbling. 2002. Effective reinforcement learning for mobile robots. In *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)*, Vol. 4. IEEE, 3404–3410.

H Francis Song, Abbas Abdolmaleki, Jost Tobias Springenberg, Aidan Clark, Hubert Soyer, Jack W Rae, Seb Noury, Arun Ahuja, Siqu Liu, Dhruva Tirumala, et al. 2020. V-MPO: On-Policy Maximum a Posteriori Policy Optimization for Discrete and Continuous Control. In *International Conference on Learning Representations*.

Nathan Sprague, Dana Ballard, and Al Robinson. 2007. Modeling embodied visual behaviors. *ACM Transactions on Applied Perception (TAP)* 4, 2 (2007), 11–es.

Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. 2019. Neural state machine for character-scene interactions. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 209.

Jie Tan, Tingnan Zhang, Erwin Coumans, Atil Iscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. 2018. Sim-to-real: Learning agile locomotion for quadruped robots. *arXiv preprint arXiv:1804.10332* (2018).

Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy Lillicrap, and Martin Riedmiller. 2018. DeepMind control suite. *arXiv preprint arXiv:1801.00690* (2018).

Demetri Terzopoulos and Tamer F Rabie. 1995. Animat vision: Active vision in artificial animals. In *Proceedings of IEEE International Conference on Computer Vision*. IEEE, 801–808.

Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. 2017. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RISJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 23–30.

Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RISJ International Conference on Intelligent Robots and Systems*. IEEE, 5026–5033.

Michiel Van de Panne and Eugene Fiume. 1993. Sensor-actuator networks. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*. ACM, 335–342.

Mel Večerík, Todd Hester, Jonathan Scholz, Fumin Wang, Olivier Pietquin, Bilal Piot, Nicolas Heess, Thomas Rothörl, Thomas Lampe, and Martin Riedmiller. 2017. Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards. *arXiv preprint arXiv:1707.08817* (2017).

Tingfan Wu, Yuval Tassa, Vikash Kumar, Javier Movellan, and Emanuel Todorov. 2013. STAC: Simultaneous tracking and calibration. In *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. IEEE, 469–476.

Zhaoming Xie, Patrick Clary, Jeremy Dao, Pedro Morais, Jonathan Hurst, and Michiel van de Panne. 2019. Iterative reinforcement learning based design of dynamic locomotion skills for cassie. *arXiv preprint arXiv:1903.09537* (2019).

Sang Hoon Yeo, Martin Lesmana, Debanga R Neog, and Dinesh K Pai. 2012. Eyecatch: Simulating visuomotor coordination for object interception. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–10.

KangKang Yin, Kevin Loken, and Michiel Van de Panne. 2007. Simbicon: Simple biped locomotion control. In *ACM Transactions on Graphics (TOG)*, Vol. 26. ACM, 105.

Andy Zeng, Shuran Song, Johnny Lee, Alberto Rodriguez, and Thomas Funkhouser. 2019. TossingBot: Learning to Throw Arbitrary Objects with Residual Physics. *Robotics: Science and Systems*.

Yuke Zhu, Ziyu Wang, Josh Merel, Andrei Rusu, Tom Erez, Serkan Cabi, Saran Tunyasuvunakool, János Kramár, Raia Hadsell, Nando de Freitas, and Nicolas Heess. 2018. Reinforcement and imitation learning for diverse visuomotor skills. In *Robotics: Science and Systems*.

APPENDICES

A SIMULTANEOUS TRACKING AND CALIBRATION

Simultaneous tracking and calibration (STAC) is an algorithm for inferring joint angles of a body from point-cloud data when it is not known in advance precisely where the markers are on the body [Wu et al. 2013]. The relevant variables include the body which has a pose ($\vec{q}$) as well as marker positions that are fixed to it ($\vec{x}_m$). We observe via motion capture the sensor readings ($\vec{s}^*$) which should be equal to the positions of the markers at each timestep, up to negligible noise. STAC makes the assumption that the markers are rigidly attached to the body with fixed offsets ($\vec{x}_m$) – if those offsets are known, a forward kinematics call ($f_k(\cdot)$) allows us to compute the positions at which we expect sensor readings.

For known marker offsets, the pose of the body ($\vec{q}$) can be inferred by optimizing (per frame):

$$\arg \min_{\vec{q}} \|f_k(\vec{q}, \vec{x}_m) - \vec{s}^*\|_2^2 \quad (4)$$

We additionally know that the marker offsets should be the same at every timestep (assuming rigid attachment of markers). So similarly, if the pose of the body is known, the marker offsets can be inferred by optimizing:

$$\arg \min_{\vec{x}_m} \sum_i \|f_k(\vec{q}_i, \vec{x}_m) - \vec{s}_i^*\|_2^2 \quad (5)$$

So overall, to perform joint optimization over unknown marker offsets and poses, we alternate between these optimization steps. We initialize the pose of the body to the null pose (approximately a t-pose) and roughly initialize the marker offsets by placing markers on the body part, without precise tuning. The first optimization is of the pose, using the initial, coarsely placed markers (per frame). We then optimize the marker positions using frames sampled at a regular interval throughout a range-of-motion [Video 11](#). We then re-optimize the joint angles per frame. We found that further alternation was not required and the marker offsets that are found using the range-of-motion clip worked well for all other clips.

In practice we also use a small regularization term, encouraging joints angles to be near the null pose, and we also warm start the per-frame optimization at the inferred pose from the preceding timestep.

B SINGLE-CLIP TRACKING OBJECTIVE

In the main text, we described that the tracking reward arises from a weighted sum of terms that score how well different features of the reference are being tracked. More specifically, these objectives are:

$$\begin{aligned}
 E_{\text{ori}} &= ||\log(\tilde{q}_{\text{ori}} \cdot \tilde{q}_{\text{ori}}^{\star-1})||_2 \\
 E_{\text{gyro}} &= 0.1 \cdot ||\tilde{q}_{\text{gyro}} - \tilde{q}_{\text{gyro}}^{\star}||_2 \\
 E_{\text{obj}} &= ||\tilde{x}_{\text{obj}} - \tilde{x}_{\text{obj}}^{\star}||_2 \\
 E_{\text{qpos}} &= \frac{1}{N_{\text{qpos}}} \sum |\tilde{q}_{\text{pos}} - \tilde{q}_{\text{pos}}^{\star}| \\
 E_{\text{qvel}} &= \frac{1}{N_{\text{qvel}}} \sum |\tilde{q}_{\text{vel}} - \tilde{q}_{\text{vel}}^{\star}| \\
 E_{\text{app}} &= \frac{1}{N_{\text{app}}} \sum ||\tilde{x}_{\text{app}} - \tilde{x}_{\text{app}}^{\star}||_2 \\
 E_{\text{vel}} &= 0.1 \cdot \frac{1}{N_{\text{vel}}} \sum |\tilde{x}_{\text{vel}} - \tilde{x}_{\text{vel}}^{\star}|
 \end{aligned}$$

where $\tilde{q}$ represents the pose or velocity and $\tilde{q}^{\star}$ represents the reference value. The $\tilde{x}_{\text{app}}$ and $\tilde{x}_{\text{obj}}$ are 3D Cartesian vectors from the root to the various appendages (head, hands, feet) or object (box or ball) in the root frame. The root is located in the pelvis of the humanoid. $\tilde{x}_{\text{vel}}$ is in the global reference frame. In this work, for the body terms, we used coefficients $w_{\text{qpos}} = 5$, $w_{\text{qvel}} = 1$, $w_{\text{ori}} = 20$, $w_{\text{app}} = 2$, $w_{\text{vel}} = 1$, $w_{\text{gyro}} = 1$. This has been used in previous work [Merel et al. 2019a]. The object tracking term coefficient, new to this work, was tuned to $w_{\text{obj}} = 10$ to relative strongly enforce object tracking. The same values are used for all clips, despite the diversity of behaviors, indicated relative robustness of this approach.

C ONE-SHOT IMITATION EVALUATION

One-shot imitation involves providing the trained NPMP with a state-sequence and asking it to generate a sequence of actions that would reproduce that movement. In asking the trained NPMP to perform one-shot imitation, we get a glimpse into which skills it is able to perform well, and we can assess this performance for overlapping subcategories of clips. Note that one-shot imitation is not actually the objective that the NPMP was trained to perform, and one-shot imitation is difficult due to object interactions (see figure A.1). Both walking behavior and ball toss behavior are better captured than the pickup and putdown interactions with boxes. This presumably reflects the fact that in terms of timesteps of data, there are fewer moments at which the difficult box interactions are being performed. As such, these quantifications may leave a misleading impression that one-shot behavior is worse than it is. To complement these quantitative metrics, we also provide a [Video 12](#) showing a representative assortment of one-shot behavior which show that while the object interactions can be difficult, movements are broadly sensible.

D INSTRUCTED WAREHOUSE TASK DETAILS

The warehouse task rewards moving a box from one pedestal to another, and repeating this process. The environment consists of a flat ground with four pedestals and two boxes that can be moved

Fig. A.1. **One-shot analyses:** Here we depict, by behavior category, the one-shot performance of the trained NPMP. Note that relative performance > 1 happens if an expert is imperfect and the one-shot imitation is effectively denoised.

freely (note that we varied the number of boxes and pedestals as well, not reported in this paper, and results are similar). The distance of each pedestal from the origin is individually drawn from a uniform distribution between 2.5 and 3.5 meters, and the pedestals are at equispaced angles around the origin. The height of each pedestal is set randomly from between 0.45 and 0.75 meters. The size of each box is taken from one of our motion capture trajectories, but with a random multiplicative variation of between 0.75 and 1.25 applied. The mass of each box is also individually drawn from a uniform distribution between 2kg and 7kg (the real boxes are either 3kg or 10kg). The size and mass of each box is not directly provided as an observation to the agent.

This task can be logically divided into four phases: walk empty-handed to a pedestal (GOTO), lifting the box from a pedestal (LIFT), carrying the box to a different pedestal (CARRY), and putting it down on the target pedestal (PUTDOWN). In our current work, we provide the agent with an observation that tells it which of these four phases it should be pursuing at a given timestep, as a one-hot vector. The position of the *focal pedestal* and *focal box* relative to the walker is also provided as observations, where the *focal box* is the box that needs to be moved, and the *focal pedestal* is dependent on the phase of the task: in GOTO and LIFT it is the pedestal on which the box is initially placed, while in CARRY and PUTDOWN it is the target pedestal. Each of the four phases has well-defined success criteria, as detailed in Table 1 (an empty cell indicates that a particular type of criterion is not used to determine success of a phase).

At each timestep, task logic determines whether the agent has successfully completed its current phase. If it has, a reward of 1.0 is given at that timestep only, and the task is advanced to the next phase. The phase transition is determined by a simple state machine:

GOTO → LIFT → CARRY → PUTDOWN → GOTO → ...,

and the task repeats indefinitely up to a final episode duration (15s simulated time), at which point the episode is terminated with bootstrapping. While there is no prespecified maximum score, obtaining an undiscounted return greater than 10 within a 15s episode requires moving through the phases rapidly. Note that the episode

Phase	Walker position	Walker/box contact	Pedestal/box contact
GOTO	within 0.65 meter of focal pedestal		
LIFT		at least one contact point with each hand	no contact points
CARRY	within 0.65 meter of focal pedestal	at least one contact point with each hand	
PUTDOWN		no contact points	at least 4 contact points

Table 1. The four phases of the warehouse task.

is terminated with a failure (no bootstrapping) if either the walker falls (contact between a non-foot geom of the walker and the ground plane) or if a box is dropped (contact between one of the boxes and the ground plane).

At the beginning of each episode, after randomly sampling the variations described above, one of the four phases is sampled uniformly as the initial phase for the episode. A motion capture trajectory is picked at random, and a random timestep from the clip consistent with the active phase in the episode is sampled. Note that the motion capture clips can be automatically segmented into phases by simply applying the task logic to the motion capture reference, so no manual segmentation is required (beyond defining the task). The joint configuration of the walker, the position of the box relative to the walker, and both the walker’s and box’s velocities, are synchronized to the state from this motion capture timestep. If the episode begins in either the LIFT or PUTDOWN phase, the displacement of the walker from the focal pedestal is also synchronized, otherwise we apply a random translation and rotation around the z-axis (i.e. yaw) to the walker and prop together as a rigid body.

E BALL TOSS TASK DETAILS

The toss task encourages catching a ball and subsequently throwing it into a bucket. The initial pose of the walker is randomly sampled from a range of motion capture poses related to ball tossing. The ball size and mass are procedurally randomized (radius randomly multiplied by a factor uniformly sampled between .95 and 1.5; mass sampled uniformly from the range of 2 to 4kg) and the angle and velocity of the ball are also procedurally randomized such that the ball is generally “thrown” towards the humanoid. More precisely, the ball is always initialized in midair, behind the bucket, at a distance of roughly 3m from the humanoid (d_x). To initialize the ball velocity in a way that ensures it is projected towards the humanoid in an appropriate *strike zone*, we must determine initial 3D velocity components of the ball, which define its trajectory. We first pick a random velocity towards the humanoid (v_x between 1.5 and 4.5m/s). We can also select a random horizontal velocity relative to the walker (v_y between .75m/s leftwards or rightwards). For the vertical

component, we compute a random target height (d_z between .1 and .4m from the ground), and then we analytically compute the time at which the ball should hit the humanoid ($t_{hit} = d_x/v_x$) as well as the initial vertical velocity required to hit the randomly selected target ($v_z = (4.9t_{hit}^2 + d_z)/t_{hit}$). For robustness, random angular velocities are also applied to the ball at the initial timestep.

The incentives of the task are specified through rewards and termination logic. The primary element of the task is that if the ball touches the ground or if the humanoid falls (contact between a non-foot geom of the walker and the ground plane), the episode terminates with a negative reward. This strongly disincentivizes letting the ball fall to the ground and encourages the humanoid to remain standing. Even reliably achieving this level of performance over the range of procedural ball trajectories is difficult. In addition, once the ball reaches the humanoid, a shaping reward is activated that corresponds to a small positive per-timestep reward inversely related to the distance between the ball and the bucket (in the x-y plane, neglecting vertical height). This reward encourages the humanoid, after catching the ball to walk towards the bucket. Finally, if the ball is in the bucket, there is a moderate per-timestep reward encouraging dropping the ball into the bucket – this final reward is sparse in the sense that it is achieved iff there is a contact between the bottom of the bucket and the ball. Once the agent has learned to drop the ball into the bucket, it learns to do this earlier (i.e. throw the ball) to achieve the reward as soon as possible.

F SUPPLEMENTARY VIDEO CAPTIONS

Video 1 Overview video summarizing highlights of the paper.

Video 2: Kinematic playback of a motion capture clip of a box interaction.

Video 3: Kinematic playback of a motion capture clip of ball tossing.

Video 4: A representative illustration of the behavior of a successfully trained vision-based policy on the “warehouse” task.

Video 5: A representative illustration of the behavior of a successfully trained vision-based policy on the “toss” task.

Video 6: A representative illustration of the behavior learned by a policy trained from scratch on the “toss” task (from state).

Video 7: A representative illustration of the behavior learned by a policy trained from scratch on the “warehouse” task (from vision).

Video 8: A representative illustration of the behavior of a partially trained vision-based policy on the “warehouse” task.

Video 9: A representative illustration of the behavior learned on the “warehouse” task when training only with large boxes.

Video 10: A representative illustration of the behavior of a successfully trained vision-based policy on the “warehouse” task when trained on a wider range of pedestal heights.

Video 11: Kinematic playback of a range-of-motion motion capture clip used for STAC calibration.

Video 12: Examples of one-shot imitation of object interaction behaviors.