



HAL
open science

Prescient Teleoperation of Humanoid Robots

Luigi Penco, Jean-Baptiste Mouret, Serena Ivaldi

► **To cite this version:**

Luigi Penco, Jean-Baptiste Mouret, Serena Ivaldi. Prescient Teleoperation of Humanoid Robots. 2023 IEEE RAS International Conference on Humanoid Robots, Dec 2023, Austin, United States. 10.1109/Humanoids57100.2023.10375166 . hal-04265528

HAL Id: hal-04265528

<https://hal.science/hal-04265528>

Submitted on 30 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Prescient Teleoperation of Humanoid Robots

Luigi Penco¹, Jean-Baptiste Mouret² and Serena Ivaldi²

Abstract—Teleoperated humanoid robots are ideally suited to act as human avatars in remote environments. Unfortunately, their deployment is hindered by the communication delays between the human input and the video feedback from the robot. Here, we introduce a direct teleoperation system in which the operator receives a synchronized video feed of real images, even when the communication channel imposes a 1 to 2-second delay. Our key idea is to leverage machine learning to allow the robot to execute commands before the operator performs them, so that the operator receives a delayed video stream that is almost indistinguishable from real-time feedback. In our experiments, the iCub humanoid robot (32 degrees of freedom) was successfully controlled to perform several whole-body manipulation tasks, including reaching different targets, picking up an object, and moving a box. This new technique may enable real-life avatars on long-range radio networks, from remote maintenance to space missions.

I. INTRODUCTION

Teleoperated robots have the potential to replace humans in numerous hazardous scenarios, ranging from contaminated environments to outer space. For such robots to be effective, they must possess both intuitive controls and high versatility, enabling operators to adapt to various situations. Humanoid robots are among the most promising solutions to meet these challenges, as they offer the capability for bimanual manipulation, walking, crawling, or climbing, while also being inherently intuitive for human teleoperation due to their similar physical structure.

Substantial progress in whole-body control and motion capture now makes it possible to map a whole-body motion of a human to a robot, despite the difference in physical constraints and dynamics [1]–[5]. This *whole-body direct teleoperation* approach allows the operator to control the robot creatively, quickly, and intuitively, and has been demonstrated in several lab experiments [1] and even competitions (e.g., ANA Avatar Xprize) [4].

Whole-body direct teleoperation becomes impractical in the presence of delays because operators rely on visual feedback from the robot to adapt and adjust their commands. Unfortunately, delays are inherent to many of the most promising applications of teleoperated humanoids. For

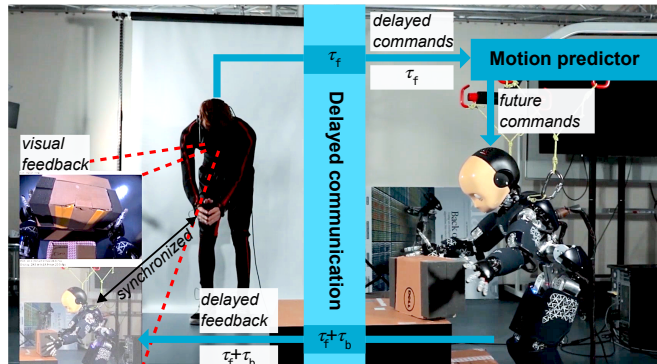


Fig. 1. **Concept of prescient teleoperation.** The robot executes the commands *before it receives them* thanks to a data-driven predictor. In this example, the robot drops a box onto a table before the human commands it to open its arms. When the predictions are accurate, by anticipating both the forward delay τ_f and the backward delay τ_b , the visual feedback appears to be synchronized with the operator. The robot sends two video streams, which are both displayed in the virtual reality headset: an internal view, from the head of the robot, and an external view. Video: https://youtu.be/_HnQw3ZZ3T8.

example, the roundtrip delay is approximately 0.81s when teleoperating a robot on Earth from the International Space Station [6], and it can be as long as 2.6 seconds with a laser communication link between the Earth and the Moon [7]. The DARPA Robotics Challenge [8] set the average delay to 1 second (with a maximum of 30 seconds), and the NASA Space Challenge [9] set it to 10 seconds (with a maximum of 20 seconds). These delays significantly impact the performance of the operator and limit the use of direct teleoperation in many real-life scenarios.

Here, we introduce a novel concept of direct teleoperation, where the operator receives a synchronized video feed of *real images* from the remote system (robot and environment/system cameras), even when the communication channel introduces a delay of 1 to 2 seconds. Our key idea is that if the robot executes the desired movement *before* the operator performs it, then the operator will observe a delayed video feed that will be almost indistinguishable from a real-time feed (Fig. 1). At each time-step, the robot analyses the data received from the operator, measures the communication time, estimates the communication time required to send the video feedback, and predicts the operator’s most likely motions in the upcoming seconds. This prediction makes it possible to execute the command with enough anticipation so that the user receives a video feed that corresponds to the past motion of the robot but that aligns with the present time for the operator. We refer to this prediction-based feedback scheme as *prescient teleoperation*.

*This work was supported by the European Union H2020 and HE Research and Innovation Programs under GA No. 731540, No. 101070596 (projects AnDy, euROBIN), the DGA grants “humanoïde résilient” and ATOR, and the Inria “ADT” wbCub/wbTorque. Experiments were performed in the Creativ’Lab facilities, supported by the CPER Sciarat and CyberEntreprise.

¹Luigi Penco is with the Florida Institute for Human and Machine Cognition. lpenco@ihmc.org

²Jean-Baptiste Mouret and Serena Ivaldi are with Inria Nancy — Grand Est, CNRS and Université of Lorraine. [jean-baptiste.mouret,serena.ivaldi}@inria.fr](mailto:{jean-baptiste.mouret,serena.ivaldi}@inria.fr)

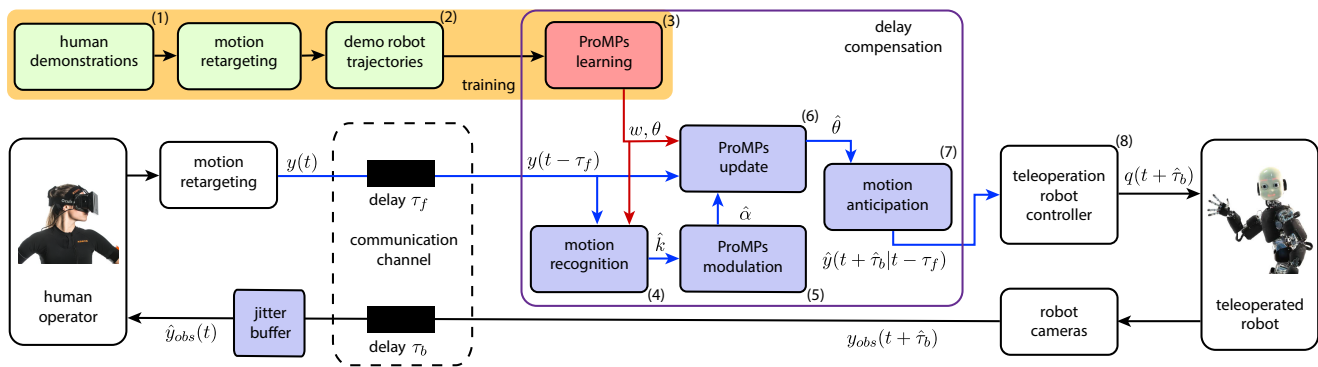


Fig. 2. **Flowchart of the prescient teleoperation system.** During the training phase, (1) the human operator teleoperates the robot without any significant delay (local network), and performs a variety of tasks. The retargeted human motions represent the robot trajectories demonstrations (2) that are later used to train the Probabilistic Movement Primitives (ProMPs). A ProMP is learned for every task (3). When teleoperating the robot in a non-ideal network, the ProMPs are used to predict the robot movement: (4) the system recognizes the current task, i.e. it identifies the most likely ProMP; (5) it estimates the speed of the teleoperated motion and (6) updates the selected ProMPs according to the delayed observed via-points; (7) it compensates for the delay by selecting the trajectory of the posterior ProMPs at the right timestep. This ensures synchronization between the user’s movements and what they observe from the remote cameras on the robot’s side. The resulting trajectories are then tracked by the whole-body controller (8), which calculates the joint commands for the robot.

This paper describes the system and presents several experiments to show the contribution of the different components, in terms of prediction, delay compensation, and tracking error. More details and additional figures, not included in the manuscript for space limits, are reported in the supplementary material available online¹. A video showing prescient teleoperation with the humanoid robot iCub is available online².

II. PREVIOUS WORK

The initial studies with teleoperated robot manipulators and delays [10] suggested that users often adopt a “move-and-wait” strategy to avoid overcompensating for delayed perceived errors. However, subsequent experiments demonstrated that this strategy is ineffective even with time delays of approximately 0.3 seconds [11].

Due to the challenges posed by delays, many teleoperation systems focused on supervised autonomy. In this approach, the operator provides high-level goals, such as selecting an object on a screen and pressing a button to grasp it, and the robot autonomously carries out the motions to fulfill the goals [3], [5], [12]–[14]. This method requires careful planning, making it difficult to spontaneously generate new behaviors for new goals or adapt to changes. Moreover, the division of goals into discrete operations often results in reduced teleoperation speed and fluidity: the operator must select an operation, wait for its completion, and then activate the next one, resulting in a sequential process.

In direct teleoperation, most research on delays focused on designing a stable control loop for systems that provide haptic feedback, specifically in multilateral teleoperation [6], [15]–[19]. This line of work does not compensate for the delays; instead, it mitigates the instability issues arising from delayed force feedback. The operator still experiences

delayed feedback and must move cautiously. To account for potential operator errors, impedance controllers are typically used by the robot [6], [20].

The challenge of delays in *visual* feedback during direct teleoperation is currently addressed through the use of *predictive displays* [21]–[24]. These displays present a non-delayed simulation of the robot and its environment. However, due to the inherent limitations of 3D modeling, disparities between the actual environment and the simulated one are unavoidable. This is particularly critical in unstructured scenarios, such as remote planets or damaged buildings, where mission-specific models are typically unavailable. Another related approach is the use of *predictor displays*, where the predicted trajectory of the system is overlaid on the visual feedback, simulated or not. This enables the operator to anticipate and better cope with delays [19], but human operators are charged with high cognitive load and can only tolerate moderate delays.

III. PRESCIENT TELEOPERATION

Fig. 2 describes the overall system. In prescient teleoperation, the primary concern revolves around determining how Cartesian and joint commands should be executed by the robot in order to effectively compensate for the delays.

A. Delay Estimation

To determine the amount of anticipation required for the commands, the robot needs to calculate the round-trip delay, which consists of a forward delay $\tau_f(t)$ in the communication from the operator to the robot and a backward delay $\tau_b(t)$ between the robot and the operator: $\tau(t) = \tau_f(t) + \tau_b(t)$. Each one-way delay consists of two components: a deterministic component primarily caused by transmission and propagation time, and a stochastic component [25], often referred to as “jitter”. We denote by $\tau_{f,D}$ the deterministic

¹<https://hal.science/hal-04166800/document>

²https://youtu.be/_HnQw3ZZ3T8

part of τ_f and by $\tau_{f,S}$ the stochastic part:

$$\tau_f(t) = \tau_{f,D} + \tau_{f,S} \quad (1)$$

$$\tau_b(t) = \tau_{b,D} + \tau_{b,S} \quad (2)$$

The robot computes the forward delay, including both the deterministic and stochastic components, by utilizing timestamps attached to the packets sent by the operator and clocks synchronized through NTP [26]. While the robot cannot predict the exact time required for a packet to reach the operator before transmission, it can obtain the average backward delay by querying the operator’s station. However, it remains unaware of the stochastic portion of the backward delay until the packet is sent. To overcome this challenge, we adopt an upper-bound approach: video streaming systems incorporate a “jitter buffer” [27], [28] that accumulates, reorders, and, if necessary, discards video packets within a specific time window. When the buffer is set to a constant size, it effectively converts the stochastic delay into a deterministic upper bound delay. As a result, the robot only needs to be aware of the jitter buffer size and the deterministic backward delay to determine when the current frame will be displayed on the operator’s computer.

B. Delay Compensation with Predictor Anticipation

To operate with delays, the system predicts the most likely future trajectories at each time step based on the commands received from the operator so far (for instance, the 3D trajectories of the hands). More formally, given a predictor of future trajectories $P(\cdot)$, the commands sent up to time step t over a history of length n , and a retargeting function³ $R(\cdot)$ that transforms the operator inputs x to robot inputs y (e.g., apply scaling factors), the predictor $P(\cdot)$ predicts the next k time steps:

$$y_{t+1}, \dots, y_{t+k} = P(R(x_{t-n}), \dots, R(x_t)) \quad (3)$$

At time-step t , the robot executes the command q_t that corresponds to the future input at time-step $\lfloor t + \tau_f + \hat{\tau}_b \rfloor$ using its whole-body controller $WB(\cdot)$:

$$q_t = WB(y_{\lfloor t + \tau_f + \hat{\tau}_b \rfloor}) \quad (4)$$

where $\lfloor \tau \rfloor$ denotes the nearest integer.

Numerous generic machine learning techniques have been proposed to predict the future of time-series, especially with neural networks [29], [30]. Nevertheless, the robotics community has long focused on regression techniques based on the concept of motion primitives, which are particularly suited for robot trajectories. Our system is based on Probabilistic Movement Primitives (ProMPs) [31]. They represent trajectories as probability distributions, making them well-suited for capturing the variability observed in human demonstrations as motion primitives. Another advantage of working with distributions is that the properties of motion primitives can be translated into operations from probability theory [31]. In particular, our system utilizes the conditioning

³In this work, we predict the retargeted values, but we could have equivalently predicted the input value and retargeted the prediction.

operator of ProMPs to adapt predictions based on incoming observations. This allows us to update the prediction, referred to as the posterior, for the ongoing movement at each time step, using the learned model of the associated primitive as the prior. In other words, ProMPs predict the mean trajectory of prior demonstrations when no conditioning data is available, but when data, e.g., as observations of the current movement, is present, they update the prediction to resemble the most likely trajectories from the training set.

At time-step t , a point ξ_t of a single trajectory is computed using a weight vector $\mathbf{w} \in \mathbb{R}^m$ and a vector of Gaussian basis functions Φ , assuming a trajectory noise variance ϵ_ξ :

$$\xi_t = \Phi_t \mathbf{w} + \epsilon_\xi \quad (5)$$

where the vector $\Phi_t \in \mathbb{R}^m$ corresponds to the m normalized radial basis functions evaluated at time t (see [31] or Appendix IV-A). Let us denote by Σ_ξ the observation noise variance. The probability of observing a trajectory \mathbf{y} given the weight vector \mathbf{w} is:

$$p(\mathbf{y}|\mathbf{w}) = \prod_t \mathcal{N}(\xi_t | \Phi_t \mathbf{w}, \Sigma_\xi), \quad (6)$$

To learn a ProMP for a given task (a set of demonstrations), we fit each demonstration (i.e., find \mathbf{w}) to this representation using ridge regression [31]. Then, we fit the distribution of weight vectors $p(\mathbf{w})$, assuming they follow a normal distribution $\mathcal{N}(\boldsymbol{\mu}_w, \Sigma_w)$.

The trajectory distribution $p(\mathbf{y})$ is obtained by marginalizing out the weight vector \mathbf{w} , i.e.

$$p(\mathbf{y}) = \int p(\mathbf{y}|\mathbf{w})p(\mathbf{w})d\mathbf{w}. \quad (7)$$

Since a whole-body trajectory is represented by N trajectories (x , y , z position of the center of mass, of the hands, etc.), we learn a ProMP for each of the N trajectories. These ProMPs all together encode a task, which means that each task is associated with a set of ProMPs. More details on our implementation are reported in the Supplementary Materials.

To enable online prediction, ProMPs are learned from demonstrations for each task (e.g., one to take a box, one to release it, one to push a bottle, etc.). During the offline training phase, an operator teleoperates the robot in a local network, which we consider to be an approximation of an ideal network without any delay. For each motion, we record the trajectories of the center of mass ground projection, the waist height, the hand Cartesian positions, the arms posture (shoulder rotation, elbow flexion, forearm rotation), the neck posture (flexion and rotation) and the torso posture (flexion, rotation and abduction). A ProMP is learned for each of these trajectories. Hence, the robot knows a set of ProMPs for each task.

For inference (prediction), we identify the task \hat{k} by computing the distance to the mean of each ProMP [32]:

$$\hat{k} = \arg \min_{k \in [1:K]} \left[\sum_{n=1}^N \sum_{t \in T_{obs}} |\mathbf{y}_n(t - \tau_f(t)) - \Phi_{n,t - \tau_f(t)} \boldsymbol{\mu}_{n,w_k}| \right], \quad (8)$$

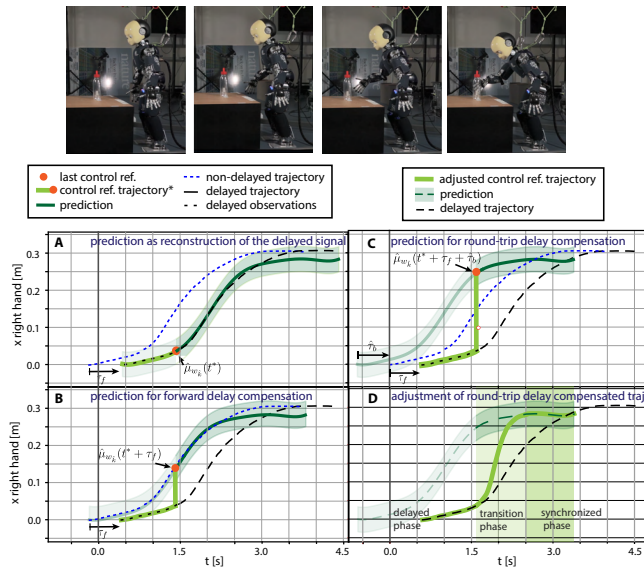


Fig. 3. **Round-trip delay compensation.** Given the past delayed observations, the robot produces at each time a prediction ($\hat{\mu}_{w_k}$) of the current command. To compensate for the delays, the right sample (orange dot) from the prediction has to be selected as a reference for the robot controller at each time. (A) The sample corresponding to the last received observation is an estimate of the delayed command. (B) By knowing the forward delay $\tau_f(t)$, a sample from the prediction can be selected to achieve a synchronization between the operator’s movement and the robot movements. (C) By knowing the forward $\tau_f(t)$ and backward delay $\tau_b(t)$, the robot can select the right sample from its prediction so as to achieve a synchronization between the operator’s movement and the feedback from the robot side. (D) A policy blending arbitrates the delayed observations with the samples selected from the prediction, which guarantees a smooth transition from delayed to compensated teleoperation.

where K is the number of tasks in the dataset and $T_{obs} = \{t_1, \dots, t_{n_{obs}}\}$ is the set of timesteps associated to the n_{obs} early observations. To account for different motion speeds, we additionally search for the best time-modulation by generating variants of each ProMP with different time-modulation values (see [32] or Appendix IV-D). This enables us to also predict the most likely duration of the current motion for the predicted task.

To predict the remaining of the trajectory, we condition the ProMP on the past trajectory using Bayes’ theorem [31], [32], which “fine-tunes” the prediction to the actual data (instead of simply following the mean of the ProMP):

$$p(\mathbf{w}_{\hat{k}} | \mathbf{x}_t^*) \propto \mathcal{N}(\mathbf{y}_t^* | \Phi_{\hat{a}t} \mathbf{w}_{\hat{k}}, \Sigma_{\hat{y}}^*) p(\mathbf{w}_{\hat{k}}). \quad (9)$$

Details on the operations can be found in the literature [31], [32].

When the robot has not received enough data to recognize the ProMP, it executes the delayed command. Once the ProMP is recognized, the robot transitions from delayed commands to compensated commands using a blending technique typical of shared autonomy approaches⁴ [33] (Fig. 3).

C. Whole-body controller

The whole-body motion of the robot is defined by the following trajectories: center of mass ground projection,

⁴More details can be found in the supplementary materials.

TABLE I

DATASETS USED TO TRAIN AND TEST PRESCIENT TELEOPERATION.

Dataset Multiple Tasks			
Bottle reaching			
	TOTAL	Bottle on table	Bottle on box
#training	12	6	6
#testing	20	10	10
Box pick & place			
	TOTAL	Picking	Placing
#training	42	18	24
#testing	21	9	12
Dataset Obstacles — Bottle on table with different obstacles			
#training	6		
#testing	9		
Dataset Goals — Bottle on table at different positions			
#training	7		
#testing	10		

See Supplementary Material, Fig. S1, S2, S3 and S4 for illustrations of the datasets.

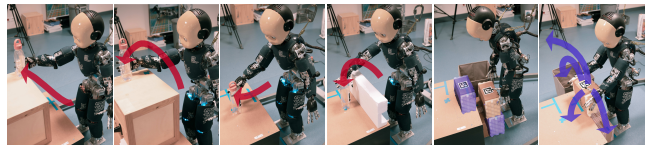


Fig. 4. **Tasks performed by the robot during the experiments (dataset “Multiple tasks”).** In the first scenario (4 left images), the robot is teleoperated to reach a bottle at different locations and in different ways; in the second scenario (2 right images), the robot has to pick up a box from different locations and then placing it in another location.

waist height, hands Cartesian positions, and the joint postures of the arms postures, neck, and torso. These trajectories, denoted by $\hat{\mathbf{y}}$, are either the delayed retargeted human motions or the delay-compensation trajectories generated by the prediction algorithm. $\hat{\mathbf{y}}$ are input to the whole-body controller, formulated as a constrained quadratic programming problem [38]:

$$\begin{aligned} \arg \min_{\hat{\mathbf{q}}} \quad & \sum_i w_i f_i + \sum_j w_j g_j + \mathbf{C} \hat{\mathbf{q}} \\ & f_i = \|\mathbf{J}_i \hat{\mathbf{q}} - \hat{\mathbf{x}}_i\|^2 \\ & g_j = \|\hat{\mathbf{q}}_j - \hat{\mathbf{q}}_j^r\|^2 \\ & \text{subject to } \mathbf{J} \hat{\mathbf{q}} = \hat{\mathbf{x}} \\ & \mathbf{A} \hat{\mathbf{q}} \leq \mathbf{b} \end{aligned} \quad (10)$$

where \mathbf{q} are the joints positions (and robot commands); w_i are the weights of the position tasks f_i ; $\hat{\mathbf{x}}_i = \dot{\hat{\mathbf{y}}}_i$ are the reference velocities for body link i ; \mathbf{J}_i is the Jacobian matrix for body link i ; w_j the weights for the postural tasks g_j ; $\hat{\mathbf{q}}_j^r = \dot{\hat{\mathbf{y}}}_j^r$ are the reference joint velocities for joints j ; $\mathbf{C} \hat{\mathbf{q}}$ is a regularization term to avoid singularities. The equality constraints include the center of mass x position and the feet poses. The inequality constraints contain the traditional robot joint velocity bounds and zero moment point bounds, that is the ZMP is constrained to stay inside the support polygon. More details on the controller are reported in [2].

IV. EXPERIMENTS

We implemented this system with the humanoid robot iCub [34], which has 32 degrees of freedom (excluding the hands and eyes) and is position-controlled. The whole-body motion of the operator is measured with a motion

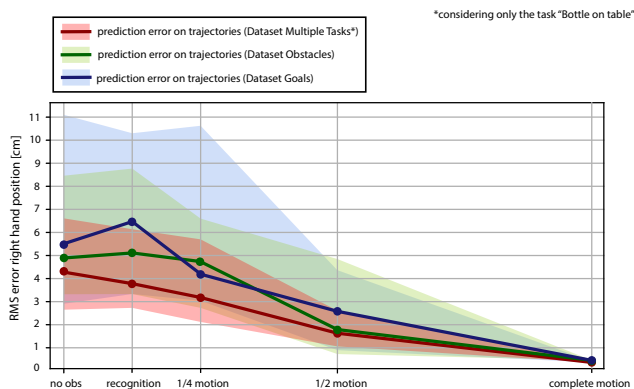


Fig. 5. **Comparison of the prediction error on different datasets.** The RMS of the error is computed based on the 10 testing trajectories of the task of reaching the bottle on the table in the dataset Multiple Tasks, the testing trajectories from the dataset Obstacles, and the dataset Goals (see Table I). The bold line represents the mean RMS error. The transparent region around it represents the maximum and minimum of the error. The prediction error is computed as the Euclidean distance between the predicted trajectory and the reference trajectory. The plot reports the error given by the mean trajectories of the ProMPs learned from the demonstrations, from the prediction updated after observing the first portion of motion used to infer the task, after observing a fourth of the motion, after observing half of the motion, and after observing the whole motion.

capture suit (Xsens MVN), which provides both joint angles and Cartesian positions of the operator’s body parts. The robot’s whole-body controller is implemented in the OpenSOT framework and uses the QP solver qpOASES [40]. It runs at 100 Hz, which is also the frequency of the motor commands. The values of the weights w_i, w_j were optimized using the procedure in [2].

We considered several teleoperation scenarios, which resulted in three different experiments and datasets (Table I), each divided into training and testing sets. All the datasets are variants of two tasks: reaching a bottle (scenario 1, Fig. 4) and manipulating a box with both hands (scenario 2). The test trajectories of the first dataset consist of various repetitions of these tasks, capturing the inherent movement variability of the operator. The second dataset, called Obstacles, features the same scenarios as in the first dataset, but the test trajectories involve different ways of reaching a bottle while avoiding obstacles. These obstacles were not considered during training, allowing the robot to adapt its task execution to different environments. The third dataset, called Goals, focuses on reaching a bottle placed at different positions. The test trajectories correspond to positions that were not used for training. All datasets and associated plots are available in the supplementary materials.

To facilitate result reproducibility and minimize time-consuming experiments, we used the real robot only for the Multiple Tasks dataset (Fig. 4). For the other two datasets, we utilized an accurate dynamic simulation. It is important to note that the real and simulated robots employ the same whole-body controller and follow the same joint trajectories. These trajectories are “played” in an open-loop manner, whether in simulation or on the robot.

Experiment 1: Evaluation of the predicted trajectories

Using testing data from the Multiple Tasks dataset, we initially evaluated the predictive capabilities of conditioned ProMPs in forecasting the future motion of the operator, independently of any consideration for teleoperation (Fig. 5). The results demonstrate that the prediction error decreases over time as more observations become available to update the prediction. For Cartesian trajectories (e.g., hands), the error reduces to approximately 0.5 cm after observing half of the motion (more detail in Supplementary Material, Table S1). Similarly, notable improvement is observed in the postural trajectories. These results highlight that continuously updating the prediction allows the operator to influence the predicted motion so that it matches better their intentions. Put differently, the system is not simply recognizing the motions and then executing the mean of the learned demonstrations: it is accurately predicting the future trajectories given the data received so far and is adapting to a specific trajectory. We obtained similar errors with the other datasets (Fig. 5).

Experiment 2: Prescient teleoperation

We conducted evaluations of the system on the iCub robot, considering a mean time-varying round-trip delay of 1.5 seconds. This delay consisted of a stochastic forward delay, following a normal distribution with a mean of 750ms, a standard deviation of 100ms, and a constant backward delay of 750ms (Fig. 6).

To assess the quality of compensation, we compared the compensated trajectory to the non-delayed trajectory for 20 testing motions in the bottle-reaching scenario of the Multiple Tasks dataset, as well as for 21 testing motions in the box handling scenario of the same dataset (see Supplementary Material, Table S2). In the box handling scenario, the results demonstrate that the error for all considered references (especially the hands) is approximately 1cm or less with compensation, whereas, without compensation, the error is roughly three times higher (around 3cm for the hands). Similarly, in the bottle reaching task, the compensated trajectory exhibits an error of approximately 1 to 1.4cm for the hands, compared to an error of about 4cm for the hands and 1cm for the center of mass when no compensation is applied. The angular errors exhibit a similar trend. While an error of around 1cm is often acceptable to accomplish a task, such as grasping an object, an error of 3 to 4cm significantly increases the likelihood of missing the object, in addition to causing the operator frustration and disorientation.

We also evaluated the performance of the compensation as the communication delay increases (Fig. 7). To do this, we compared the compensated trajectory to the non-delayed trajectory, for the right hand, in the task of reaching the bottle on the table of the dataset Multiple Tasks (Supplementary Material, Fig. S1). During the synchronization, the error is roughly proportional to the delay (Fig. 3), which adds up to the prediction errors. In this case, we observed a mean error of approximately 2.5cm for a 1-second delay, but over 10cm for a 3-second delay, primarily because the transition time constitutes a significant portion of a short trajectory (30%

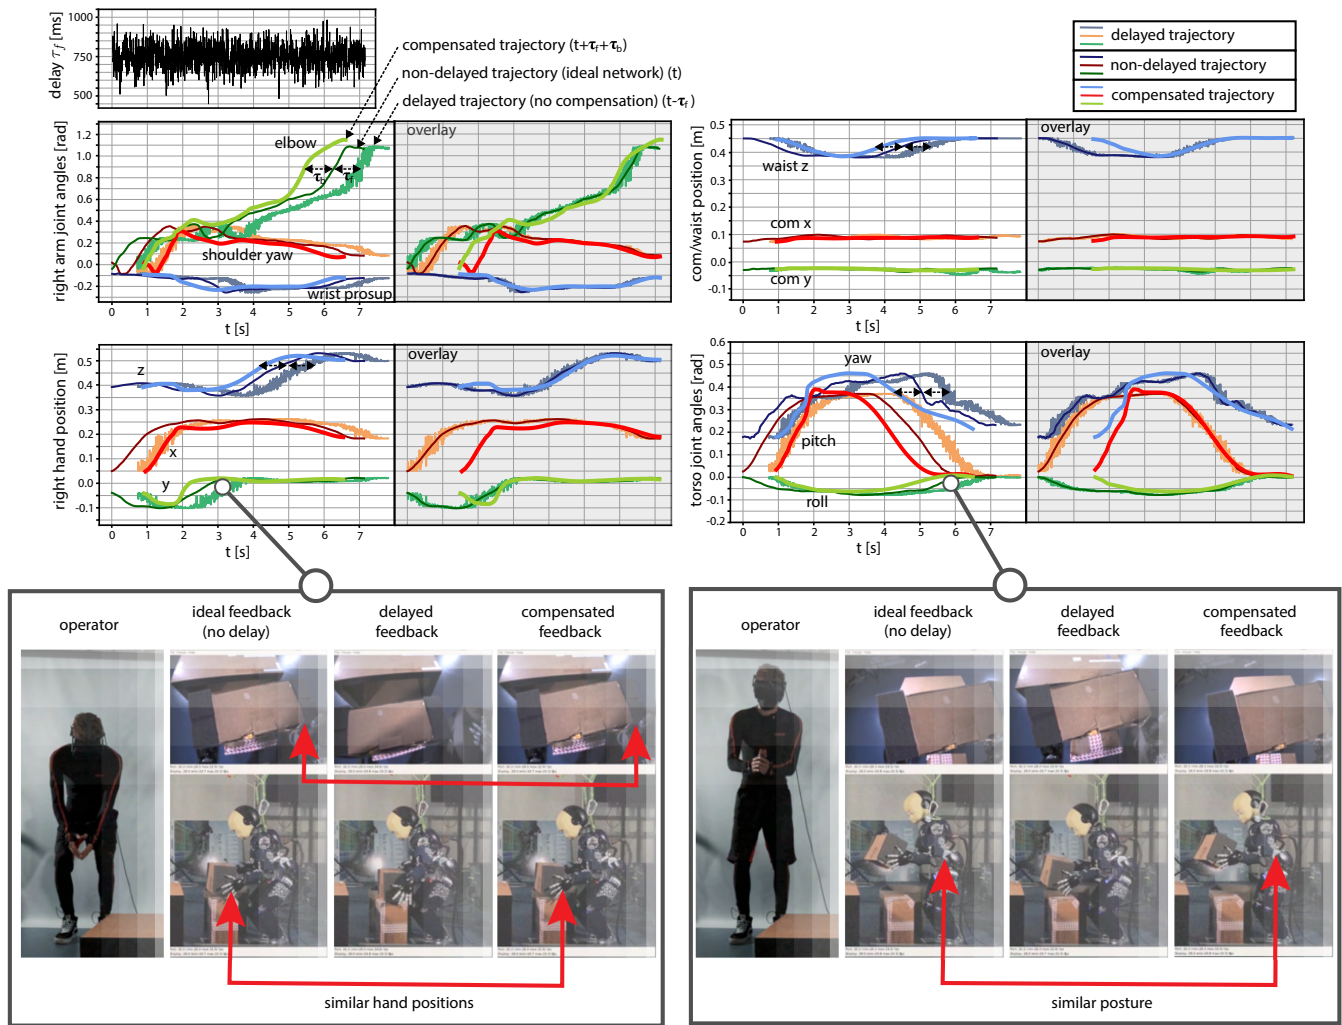


Fig. 6. **Teleoperation with compensation of a mean round-trip delay of 1.5s.** The robot is picking up a box in front of it at mid-height. The forward delay follows a normal distribution with 750ms as mean and 100ms as standard deviation, while the backward delay is 750ms (top). After the synchronization time (first 2 seconds), the compensated trajectory anticipates the delayed trajectory by $\tau_b + \tau_f$ to take into account both the delay from the operator to the robot and from the robot to the operator, this is why the robot anticipates *more* than the non-delayed trajectory. Once we shift all the trajectories to compare them (“overlay”), the commands (delayed trajectories) and the compensated trajectories align very well, which shows that the robot executes the right trajectory but shifted in time to anticipate the motion.

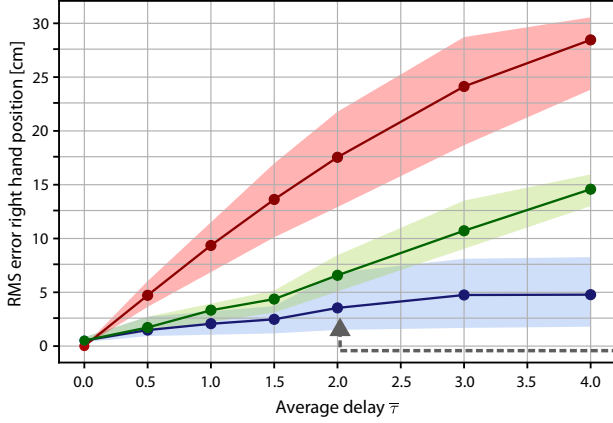
for a 3-second delay and a 10-second trajectory). Excluding the synchronization time, the results reveal that the compensated tracking error is below 2cm for delays around 0.5s, approximately 2.5cm for delays around 1.5s, and increases to about 5cm for delays of 3s and 4s. Qualitatively, the system becomes challenging to use for the operator with a delay exceeding 2 seconds, resulting in an average error of approximately 3cm after the transition and approximately 7cm when considering the transition. We recorded similar tracking errors with the other datasets (Fig. 7b).

Experiments 3: Conforming the teleoperation to the intended motion

Humans can often perform a task in many ways because they are highly redundant. For example, when picking up a box, an operator can choose to bend their back without utilizing their legs, bend their legs while minimizing back

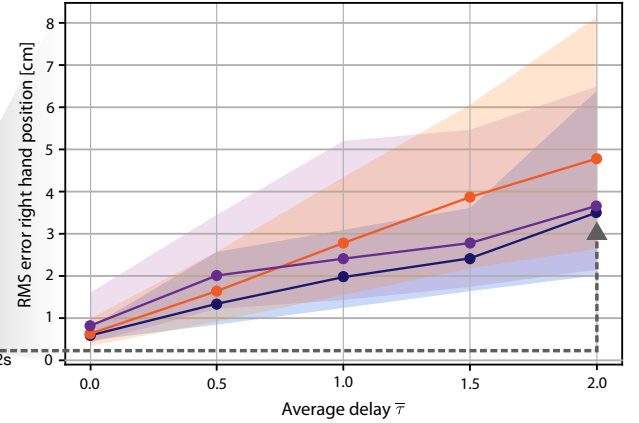
movements, or adopt any combination in between. Ideally, the robot should execute tasks in the preferred manner of the operator. To explicitly explore this motion adaptation, we teleoperated the robot in simulation to reach a bottle placed on a table using three different approaches during the training phase, dataset “Obstacles”). During testing, we assessed the effectiveness of our compensation approach for the same task but with different obstacles present. As with the previous dataset, the prediction steadily improves as more observations become available, increasingly aligning with the observed operator commands. The results show that the position error is comparable (if higher, only by about half a centimeter) to the error from the bottle-reaching scenario, where the same obstacles were used for both training and testing.

a. Tracking error for compensated and non-compensated trajectories (dataset Multiple tasks)



■ tracking error of non-compensated trajectories (Dataset Multiple Tasks*)
■ tracking error of compensated trajectories with transition (Dataset Multiple Tasks*)
■ tracking error of compensated trajectories after transition (Dataset Multiple Tasks*)
 *considering only the task "Bottle on table"

b. Tracking error for compensated trajectories: all test sets



■ tracking error of compensated trajectories after transition (Dataset Multiple Tasks*)
■ tracking error of compensated trajectories after transition (Dataset Obstacles)
■ tracking error of compensated trajectories after transition (Dataset Goals)

Fig. 7. Scalability of the delay compensation with respect to increasing time delays. a. Tracking error of the compensated trajectories for the right-hand position with respect to the non-delayed ones compared to the tracking error of the corresponding non-compensated (delayed) trajectories with respect to the same non-delayed ones. The tracking error of the compensated trajectories is considered both including the transition from the delayed phase to the synchronization phase (Fig. 3), which adds a non-compensable error, and without transition. The RMS of the error is computed from the 10 testing motions of the task of reaching the bottle on the table from dataset Multiple Tasks and its mean value is reported as a bold line. The transparent region around it, represents the maximum and minimum of the error. The tracking error is computed as the Euclidean distance between the evaluated trajectory and the reference trajectory. The compensated trajectories are temporally realigned with the non-delayed trajectories for computing the error, which is evaluated with different round-trip delays $\tau(t)$: 0s, around 0.5s, 1s, 1.5s, 2s, 3s and 4s. The time-varying forward delay follows a normal distribution with mean $\bar{\tau}_f = \bar{\tau}/2$ and standard deviation equal to $\frac{2}{15}\bar{\tau}_f$. The backward delay is set equal to $\bar{\tau}_f$. **b. Tracking error of the compensated trajectories for the right-hand position compared with the non-delayed ones (after the transition phase) for all the datasets.**

Experiment 4: Conforming the teleoperation to new goals

In many scenarios, the operator may need to perform motions that the system was not specifically trained for. To investigate this challenge, we assessed how the robot adapts to new object locations. For this purpose, we utilized the third dataset (dataset "Goals") and teleoperated the robot in simulation to reach a bottle positioned on a table at various locations. We then tested the approach by reaching the same bottle but at different positions. As one would expect, the prediction is less accurate compared to the previous datasets where the goal position remains constant. However, the error decreases to approximately 1cm after observing half of the motion (Fig. 5). When teleoperating the robot with a mean delay of 1.5s, the average tracking error on the hand position is 4cm (Fig. 7b), which is approximately a centimeter higher than the error with a 2-second delay in datasets where the bottle remains in the same position for both training and testing. The operator found it challenging to teleoperate the robot with such compensation errors. However, the approach can still be used effectively with novel goals for lower delays, typically between 0.5s and 1s, where similar accuracy to the other datasets is often achieved (Fig. 7b).

V. CONCLUSION AND DISCUSSION

Whole-body teleoperation offers an intuitive and flexible approach to operate humanoid robots exploiting their entire body, as long as the operator can rely on synchronized

feedback. Through our use of machine learning to anticipate operator commands, we have demonstrated the ability to compensate for delays ranging from 1 to 2 seconds, which are typically observed in round-trip communication times between Earth and space [6] and across continents on the Internet [35].

While we utilized ProMPs as our chosen method for predicting future motions, other data-driven techniques for trajectory or time-series prediction, such as LSTM [36], [37], could also be employed. However, these alternatives may require more training data and could potentially result in less smooth movements. In essence, the concept of *prescient teleoperation* can be implemented using any predictor, including neural networks, and the system's performance will improve with more accurate and advanced predictors.

All the experiments described in this study were conducted by an experienced user who is familiar with the teleoperation system and iCub. While novice users may not possess the same level of proficiency with the proposed system, it is unlikely that they would be entrusted with operating an expensive humanoid robot in a high-stakes mission, such as intervention in a damaged chemical plant or a remote lunar base. Just as drone pilots undergo rigorous training before their first mission, we expect that future humanoid operators will undergo extensive training as well. For this reason, in this work, we did not compare novice and expert user performance and we did not conduct a user study on usability

in general. In future work, we will conduct user studies to examine the impact of prediction and delay compensation on human cognition, behavior, and performance, to explore the requirements and limits of how experienced users can effectively utilize a prescient teleoperation system in real-world scenarios.

SUPPLEMENTARY MATERIAL

Supplementary Material are available online at <https://hal.science/hal-04166800/document>. They include more experiments and details about the methods.

REFERENCES

- [1] L. Penco *et al.*, “Robust real-time whole-body motion retargeting from human to humanoid,” in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2018, pp. 425–432.
- [2] L. Penco, E. M. Hoffman, V. Modugno, W. Gomes, J. Mouret, and S. Ivaldi, “Learning robust task priorities and gains for control of redundant robots,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2626–2633, 2020.
- [3] M. Zucker, S. Joo, M. X. Grey, C. Rasmussen, E. Huang, M. Stilman, and A. Bobick, “A general-purpose system for teleoperation of the dre-hubo humanoid robot,” *Journal of Field Robotics*, vol. 32, no. 3, pp. 336–351, 2015.
- [4] S. Dafarra, K. Darvish, R. Grieco, G. Milani, U. Pattacini, L. Rapetti, G. Romualdi, M. Salvi, A. Scalzo, I. Sorrentino, *et al.*, “icub3 avatar system,” *arXiv preprint arXiv:2203.06972*, 2022.
- [5] K. Darvish, L. Penco, J. Ramos, R. Cisneros, J. Pratt, E. Yoshida, S. Ivaldi, and D. Pucci, “Teleoperation of humanoid robots: A survey,” *IEEE Transactions on Robotics*, 2023.
- [6] M. Panzirsch *et al.*, “Exploring planet geology through force-feedback telemanipulation from orbit,” *Science Robotics*, vol. 7, no. 65, p. eabl6307, 2022.
- [7] D. M. Boroson, B. S. Robinson, D. V. Murphy, D. A. Burianek, F. Khatri, J. M. Kovalik, Z. Sodnik, and D. M. Cornwell, “Overview and results of the lunar laser communication demonstration,” in *Free-Space Laser Communication and Atmospheric Propagation XXVI*, vol. 8971. SPIE, 2014, pp. 213–223.
- [8] C. Atkeson *et al.*, *What Happened at the DARPA Robotics Challenge Finals*. Springer Tracts in Advanced Robotics, 2018, pp. 667–684.
- [9] “National aeronautics and space administration (nasa)-centennial challenges program-space robotics challenge phase 2,” <https://t.ly/t10-t>.
- [10] W. R. Ferrell, “Remote manipulation with transmission delay,” *IEEE Trans. on Human Factors in Electron.*, vol. 6, no. 1, pp. 24–32, 1965.
- [11] W. R. Ferrell and T. B. Sheridan, “Supervisory control of remote manipulation,” *IEEE Spectrum*, vol. 4, no. 10, pp. 81–88, 1967.
- [12] M. Stilman, K. Nishiwaki, and S. Kagami, “Humanoid teleoperation for whole body manipulation,” in *IEEE ICRA*, 2008.
- [13] B. Omarali, B. Denoun, K. Althoefer, L. Jamone, M. Valle, and I. Farkhatdinov, “Virtual reality based telerobotics framework with depth cameras,” in *IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2020.
- [14] A. Naceri, D. Mazzanti, J. Bimbo, Y. T. Tefera, D. Prattichizzo, D. G. Caldwell, L. S. Mattos, and N. Deshpande, “The Vicarios virtual reality interface for remote robotic teleoperation,” *Journal of Intelligent & Robotic Systems*, vol. 101, no. 4, pp. 1–16, 2021.
- [15] J.-H. Ryu, D.-S. Kwon, and B. Hannaford, “Stable teleoperation with time-domain passivity control,” *IEEE Transactions on robotics and automation*, vol. 20, no. 2, pp. 365–373, 2004.
- [16] M. Panzirsch, H. Singh, T. Krüger, C. Ott, and A. Albu-Schäffer, “Safe interactions and kinesthetic feedback in high performance Earth-to-Moon teleoperation,” in *IEEE Aerospace Conference*, 2020.
- [17] J. Artigas, R. Balachandran, C. Riecke, M. Stelzer, B. Weber, J.-H. Ryu, and A. Albu-Schaeffer, “Kontur-2: force-feedback teleoperation from the international space station,” in *IEEE ICRA*, 2016.
- [18] N. Lii *et al.*, “The robot as an avatar or co-worker? an investigation of the different teleoperation modalities through the KONTUR-2 and METERON SUPVIS Justin space telerobotic missions,” in *Int. Astronautical Cong.*, 2018.
- [19] T. B. Sheridan, “Space teleoperation through time delay: Review and prognosis,” *IEEE Trans. on Robotics and Automation*, vol. 9, no. 5, pp. 592–606, 1993.
- [20] P. Evrard, N. Mansard, O. Stasse, A. Kheddar, T. Schaub, C. Weber, A. Peer, and M. Buss, “Intercontinental, multimodal, wide-range tele-cooperation using a humanoid robot,” in *IEEE/RSJ IROS*, 2009, pp. 5635–5640.
- [21] L. Peñin, K. Matsumoto, and K. U. G. Kenkyūjo, *Teleoperation with Time Delay: A Survey and Its Use in Space Robotics*. National Aerospace Laboratory, 2002.
- [22] M. Hernando and E. Gambaio, *Advances in Telerobotics*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, ch. Teleprogramming: Capturing the Intention of the Human Operator, pp. 303–320.
- [23] A. K. Bejczy, W. S. Kim, and S. C. Venema, “The phantom robot: predictive displays for teleoperation with time delay,” in *IEEE ICRA*, 1990.
- [24] P. Mitra and G. Niemeyer, “Mediating time delayed teleoperation with user suggested models: Implications and comparative study,” in *Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 2008, pp. 343–350.
- [25] O. Gurewitz, I. Cidon, and M. Sidi, “One-way delay estimation using network-wide measurements,” *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2710–2724, 2006.
- [26] D. Mills, J. Martin, J. Burbank, and W. Kasch, “Network time protocol version 4: Protocol and algorithms specification,” Internet Requests for Comments, RFC 5905, June 2010.
- [27] B. Oklander and M. Sidi, “Jitter buffer analysis,” in *Proc. Int. Conf. on Comp. Comm. Networks, ICCCN*, 09 2008, pp. 1 – 6.
- [28] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, “RTP: A transport protocol for real-time applications,” Internet Requests for Comments, RFC Editor, STD 64, 2003.
- [29] A. R. S. Parmezan, V. M. Souza, and G. E. Batista, “Evaluation of statistical and machine learning models for time series prediction: Identifying the state-of-the-art and the best conditions for the use of each model,” *Information sciences*, vol. 484, pp. 302–337, 2019.
- [30] B. Lim and S. Zohren, “Time series forecasting with deep learning: A survey,” *Philosophical Transactions of the Royal Society A.*, vol. 379, pp. 20 200 209–20 200 209, 2021.
- [31] A. Paraschos, C. Daniel, J. Peters, and G. Neumann, “Using probabilistic movement primitives in robotics,” *Autonomous Robots*, vol. 42, p. 529–551, 07 2018.
- [32] O. Dermay, A. Paraschos, M. Ewerton, J. Peters, F. Charpillet, and S. Ivaldi, “Prediction of intention during interaction with iCub with probabilistic movement primitives,” *Frontiers in Robotics and AI*, vol. 4, pp. 45–45, 2017.
- [33] A. D. Dragan and S. S. Srinivasa, “A policy-blending formalism for shared control,” *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 790–805, 2013.
- [34] L. Natale, C. Bartolozzi, D. Pucci, A. Wykowska, and G. Metta, “iCub: The not-yet-finished story of building a robot child,” *Science Robotics*, vol. 2, no. 13, p. eaaq1026, 2017.
- [35] T. Høiland-Jørgensen, B. Ahlgren, P. Hurgig, and A. Brunstrom, “Measuring latency variation in the internet,” in *Proc. Int. Conf. emerging Networking EXperim. and Techn.*, 2016, pp. 473–480.
- [36] X. Zhao, S. Chumkamon, S. Duan, J. Rojas, and J. Pan, “Collaborative human-robot motion generation using LSTM-RNN,” in *IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*, 2018, pp. 1–9.
- [37] J. Bütepage, H. Kjellström, and D. Kragic, “Anticipating many futures: Online human motion prediction and generation for human-robot interaction,” in *IEEE ICRA*, 2018.

Supplementary Material

Prescient teleoperation of humanoid robots

Accepted at IEEE HUMANOIDS 2023

Luigi Penco¹, Jean-Baptiste Mouret² and Serena Ivaldi²

¹Luigi Penco is with the Florida Institute for Human and Machine Cognition. lpenco@ihmc.org

²Jean-Baptiste Mouret and Serena Ivaldi are with Inria Nancy — Grand Est, CNRS and Université of Lorraine.

{[jean-baptiste.mouret](mailto:jean-baptiste.mouret@inria.fr), [serena.ivaldi](mailto:serena.ivaldi@inria.fr)}@inria.fr

Overall, our *prescient* whole-body teleoperation system relies on the following components (Fig. 2):

- a whole-body controller based on quadratic optimization;
- a dataset of whole-body trajectories retargeted from human motions;
- a set of ProMPs that can predict future trajectories given observations;
- a computation of the round-trip delay to select the appropriate commands from the prediction, so that the robot anticipates both the operator-to-robot and the robot-to-operator delays;
- a blending to keep the trajectory smooth, at the start of a trajectory or in case of changes of delays;
- a video streaming system that uses a jitter buffer to cope with the stochastic part of the backward delay.

APPENDIX I

HUMANOID ROBOT AND WHOLE-BODY CONTROLLER

A. The *iCub* humanoid robot

iCub [34] is a research-grade open-source humanoid robot designed by the Italian Institute of Technology (IIT) to experiment with embodied artificial intelligence. It measures 104 cm in height and weighs 22 kg, which roughly corresponds to the body dimensions of a five-year-old child. *iCub* has 53 actuated degrees of freedom: 7 in each arm, 9 in each hand (3 for the thumb, 2 for the index, 2 for the middle finger, 1 for the coupled ring and little finger, 1 for the adduction/abduction), 6 in the head (3 for the neck and 3 for the cameras), 3 in the torso/waist, 6 in each leg. In this work, we do not use the fingers of the hands and we do not move the eyes, therefore we control 32 degrees of freedom. The head has stereo cameras in a swivel mounting in the corresponding location of the human eye sockets. *iCub* is also equipped with six 6-axial force/torque (F/T) sensors in the upper arms, legs and feet, an IMU in the head, and a distributed tactile skin.

B. Whole-body controller

The whole-body motion of the robot is defined by the following trajectories: center of mass ground projection, waist height, hands positions, arms postures, neck posture, torso posture, which are either given by the delayed retargeted human motion, or generated by the delay compensation

algorithm during the execution of the main task. Each sample of each of these trajectories represents a control reference $\hat{\mathbf{y}}$. Given $\hat{\mathbf{y}}$, the robot commands \mathbf{q} are generated by solving the redundant inverse kinematics, which can be formulated as a constrained quadratic programming problem with equality and inequality constraints [2], [38]:

$$\begin{aligned} \arg \min_{\dot{\mathbf{q}}} \quad & \sum_i w_i f_i + \sum_j w_j g_j + \mathbf{C}\dot{\mathbf{q}} \\ & f_i = \|\mathbf{J}_i \dot{\mathbf{q}} - \dot{\mathbf{x}}_i\|^2 \\ & g_j = \|\dot{\mathbf{q}}_j - \dot{\mathbf{q}}_j^r\|^2 \\ \text{subject to} \quad & \mathbf{J}\dot{\mathbf{q}} = \dot{\mathbf{x}} \\ & \mathbf{A}\dot{\mathbf{q}} \leq \mathbf{b} \end{aligned} \quad (\text{S1})$$

The cost function consists of terms f_i with relative weight w_i concerning the pose of a specific body link i , where \mathbf{J}_i is the Jacobian matrix for body link i and $\dot{\mathbf{x}}_i = \hat{\mathbf{y}}_i$ are the reference velocities for body link i . Additionally terms g_j with relative weight w_j concern the posture of certain joints j , where $\dot{\mathbf{q}}_j^r = \hat{\mathbf{y}}_j$ are the reference joint velocities for joints j . $\mathbf{C}\dot{\mathbf{q}}$ is a regularization term used to get a unique solution and avoid singularities, where \mathbf{C} is a linear cost vector.

In our implementation, we considered in the terms f_i the hand positions with $w_i = 1$ and the waist height with $w_i = 0.65$. Instead, the terms g_j include the head posture with $w_j = 1$ and the torso posture with $w_j = 0.72$, the elbow and wrist postures with $w_j = 0.11$. We computed optimal priorities with a multi-objective stochastic optimization that was run in simulation [2]. More details about the whole-body controller can be found in [2].

The equality constraints correspond to the highest priority task, which should be solved exactly. In our implementation, these include the center of mass x position and the feet poses. The inequality constraints contain the robot joint velocity bounds and zero moment point bounds, which is constrained to stay inside the support polygon.

Our controller is based on the OpenSOT framework [40] and the qpOASES quadratic programming solver [42].

This controller is run at 100 Hz, which is also the frequency of the motor commands.

C. Motion retargeting

The captured human information cannot be directly used as a reference for the humanoid, due to differences in kinematics (e.g. joint limits, body dimensions) and dynamics (e.g. mass distribution) between the human and the robot.

TABLE S1

PREDICTION ERROR AFTER OBSERVING DIFFERENT PORTIONS OF THE COMMANDED TRAJECTORIES (DATASET MULTIPLE TASKS). WE EVALUATED THE DIFFERENCE BETWEEN THE ACTUAL TRAJECTORY (COMMANDS RETARGETED FROM THE OPERATOR) AND THE PREDICTED TRAJECTORY FOR THE 20 TESTING MOTIONS FROM THE BOTTLE REACHING SCENARIO (FIG. S1) AND THE 21 TESTING MOTIONS FROM THE BOX HANDLING SCENARIO (FIG. S2). TO UNDERSTAND THE INFLUENCE OF THE CONDITIONING OF THE PROMPS, WE COMPUTED THE MEAN ERROR BY FOLLOWING THE MEAN OF THE PROMP SELECTED BY HAND ('NO OBS'), AFTER THE INITIAL RECOGNITION ('RECOGNITION') THAT TAKES ABOUT 1S, AFTER A FOURTH OF THE MOTION (1/4 MOTION) AND AFTER HALF OF THE MOTION (1/2 MOTION). THANKS TO THE CONDITIONING, WHEN MORE DATA IS USED, THE PREDICTION IS MORE ACCURATE, WHICH MEANS THAT THE PREDICTION IS ADJUSTED TO SUIT THE PARTICULAR MOTIONS OF THE OPERATOR (THAT IS, THE ROBOT DOES NOT SIMPLY FOLLOW THE MEAN TRAJECTORY ONCE IT HAS RECOGNIZED IT). EXAMPLES OF PREDICTED TRAJECTORIES ARE DISPLAYED IN FIG. S1 AND FIG. S2.

Trajectory	Box handling RMS error [rad]				Bottle reaching RMS error [rad]			
	no obs	recognition	1/4 motion	1/2 motion	no obs	recognition	1/4 motion	1/2 motion
head yaw	0.112±0.049	0.080±0.028	0.045±0.012	0.012±0.009	0.083±0.038	0.040±0.015	0.023±0.011	0.010±0.008
torso pitch	0.155±0.064	0.120±0.046	0.054±0.030	0.018±0.008	0.103±0.056	0.061±0.022	0.044±0.015	0.019±0.008
torso roll	0.119±0.049	0.103±0.038	0.055±0.028	0.017±0.008	0.082±0.042	0.054±0.016	0.032±0.010	0.016±0.008
torso yaw	0.168±0.053	0.136±0.049	0.079±0.032	0.049±0.019	0.088±0.046	0.065±0.039	0.049±0.023	0.033±0.018
r. should. yaw	0.164±0.059	0.109±0.050	0.053±0.019	0.040±0.011	0.172±0.059	0.114±0.056	0.055±0.024	0.027±0.011
r. elbow	0.169±0.072	0.146±0.038	0.097±0.032	0.033±0.010	0.220±0.083	0.097±0.032	0.065±0.017	0.034±0.011
r. wrist pros.	0.113±0.049	0.092±0.022	0.043±0.012	0.026±0.008	0.124±0.052	0.071±0.022	0.048±0.016	0.021±0.008
	RMS error [cm]				RMS error [cm]			
	no obs	recognition	1/4 motion	1/2 motion	no obs	recognition	1/4 motion	1/2 motion
r. hand x	3.76±0.91	2.10±0.48	1.57±0.49	0.78±0.12	3.54±0.88	1.62±0.65	0.61±0.33	0.48±0.14
r. hand y	3.55±0.93	1.91±0.40	0.97±0.32	0.52±0.10	3.71±0.89	2.02±0.48	0.89±0.19	0.35±0.11
r. hand z	2.98±0.91	2.34±0.83	1.22±0.27	0.66±0.19	2.71±0.80	1.99±0.79	0.77±0.29	0.50±0.13
com x	2.40±0.76	0.98±0.26	0.57±0.21	0.23±0.14	0.78±0.53	0.39±0.22	0.29±0.12	0.12±0.10
com y	1.12±0.55	0.58±0.26	0.29±0.18	0.20±0.13	0.80±0.44	0.52±0.29	0.29±0.12	0.11±0.10
waist z	3.53±1.04	2.74±0.93	1.68±0.57	0.65±0.22	0.95±0.36	0.65±0.22	0.39±0.19	0.28±0.14

Hence, motion retargeting is employed to map the human information into feasible reference trajectories for the robot. For transferring the translational movements of the end-effectors we used a fixed scaling factor (0.4). For transferring postures, the joint angles of the human joints are manually identified and mapped to the corresponding joints of the robot [1]. The instantaneous reference value of the robot is then computed as:

$$\Delta q_{iR} = q_{0R} + (q_{iH} - q_{0H}) \quad (S2)$$

where q is the vector of current joint positions, Δq is the vector of joint variations with respect to the initial posture, the indices 0 and i refer to measurements at an initial time and at time i , and the subindices H and R indicate measurements on human and robot, respectively. The same applies to the Cartesian positions of the end-effectors.

For the center of mass, the normalized offset-based reconstruction is used [1]. We consider the ground projection of the human center of mass \mathbf{p}_{com}^g . Its position with respect to the left foot is projected onto the line connecting the two feet. The result is then normalized to obtain an offset $o \in [0, 1]$

$$o = \frac{(\mathbf{p}_{com}^g - \mathbf{p}_{lFoot}^g) \cdot (\mathbf{p}_{rFoot}^g - \mathbf{p}_{lFoot}^g)}{\|\mathbf{p}_{rFoot}^g - \mathbf{p}_{lFoot}^g\|^2}$$

where \mathbf{p}_{lFoot}^g and \mathbf{p}_{rFoot}^g are the ground projections of the left and right foot respectively. When the human is in a symmetric pose, the offset o has a value around 0.5 and when the human stands on a single foot, it is either 0 (left foot) or 1 (right foot). The robot center of mass ground projection is then reconstructed on the line connecting its feet by means of this offset value. To also retarget changes of the center of mass that are not on the line connecting the feet, we

can apply the same concept while considering the maximum backward and forward center of mass displacement in the orthogonal direction of the line connecting the feet as done in [1].

APPENDIX II DATASETS

To train our method, we teleoperated the robot in an ideal network without any delay and recorded the corresponding robot motion. Every demonstration contains several Cartesian and postural trajectories that determine the whole-body motion (Fig. S2): the center of mass ground projection, the waist height, the hand positions, the arms posture (shoulder rotation, elbow flexion, forearm rotation), the neck posture (flexion and rotation) and the torso posture (flexion, rotation and abduction). We record the retargeted trajectories, that is the reference trajectories for the whole-body controller of the robot.

We used three different datasets (Table I), each one divided into a training set and a test set. The first dataset (Dataset Multiple Tasks) is designed to test how well the robot recognizes tasks and deals with the intrinsic variability of the operator's movements. This is the dataset used to perform the experiments on the real robot. The second one (Dataset Obstacles) is designed to evaluate the approach with unexpected obstacles. The third dataset (Dataset Goals) is designed to evaluate the approach with novel goal positions.

The datasets are available online at <https://doi.org/10.5281/zenodo.5913573>.

A. Dataset Multiple Tasks

This dataset covers two scenarios: reaching a bottle with the right hand (Fig. S1) and handling a box (Fig. S2).

The bottle task consists of demonstrations of 2 distinct whole-body reaching primitives: one primitive is for reaching the bottle on the table, the other one is for reaching the bottle on the top of the box. For each primitive, we recorded 6 repetitions of the task for training, for a total of 12 training whole-body demonstrations with an average duration of 6.1s. Every demonstration provided by the human operator is different, since it is not possible to exactly reproduce the same whole-body movement twice; to further increase the variance of the demonstrated movements, in 3 repetitions out of the 6, an obstacle was placed in between the robot and the bottle. To assess the quality of the predictions, 10 different testing repetitions were recorded for each of the two primitives; 5 with the obstacle in between the robot and the bottle, and 5 without any obstacle, for a total of 20 motions. In this dataset, the obstacles are at the same positions in both the training set and the testing set (see the dataset ‘‘Obstacles’’ below).

The second scenario consists of demonstrations of 7 distinct whole-body box handling primitives: 3 for picking up the box — from a low position, from a mid-height and from the table; 4 for placing the box at a specific location — on the floor, on the table, inside a container, or in a person’s hands. For each primitive, we recorded 6 different repetitions for training, for a total of 42 training whole-body demonstrations, with an average duration of 7.2 s for the pick-up motions and of 5.8 s for the box-placing motions. For the test set, 3 new different repetitions of the 7 motions were recorded, for a total of 21 testing motions.

B. Dataset Obstacles

The training set is composed of 6 demonstrations of bottle reaching motions with 3 different locations of an obstacle (Fig. S3): 2 repetitions without obstacles, 2 with an obstacle in between the robot and the bottle, and 2 with a different obstacle. The average duration of the demonstrations is 6.9s. The test set consists of motions for the same task but with obstacles at different locations (Fig. S3b): 3 repetitions for each of the 3 distinct scenarios with different obstacles.

C. Dataset Goals

The training set is composed of 7 demonstrations of bottle reaching motions (Fig. S4), with the goal located in 7 different positions. The average duration of the demonstrations is 6.1s. The test set consists of motions reaching the same bottle but at 10 different locations (Fig. S4b).

APPENDIX III

DELAYED TELEOPERATION

A. Hardware and communication setup

The human motion is captured by the Xsens MVN system [44], which considers a human model comprising 66 degrees of freedom (corresponding to 22 spherical joints). The user teleoperating the robot is equipped with the wearable motion

capture suit MVN Link, consisting of a Lycra suit with 17 inertial measurement units (IMUs) and a wireless data link transmitting at a frequency of 240Hz. Our compensation method receives the delayed data from the motion capture system at 100Hz and transmits to the robot controller at 50Hz.

The user teleoperating the robot is also equipped with a VR Oculus headset. Through the headset, the operator can visualize the delayed images from both an external camera at the robot side, as a third-person view of the teleoperated robot, and the robot cameras, for a first-person immersive experience. The communication protocol employed by the network is UDP with a bandwidth of 3Mbps. The forward delay is artificially generated, using the standard way to delay packets in Linux with the ‘‘netem’’ scheduling policy, which is based on the ‘‘iproute2’’ set of tools [45].

The images from the cameras at the robot side are delayed by using the open-source application Kinovea [46], which allows the user to set a constant delay for the streaming of the video. The resulting delayed streaming is projected onto the VR headset through the application Virtual Desktop [47].

B. Delay generation

The round-trip delay $\tau(t)$ at time-step t is divided into a forward $\tau_f(t)$ (operator to robot) and a backward delay $\tau_b(t)$ (robot to operator):

$$\tau(t) = \tau_f(t) + \tau_b(t)$$

Each one-way delay is composed of two parts, one deterministic and one stochastic [25]. The deterministic component corresponds to the transmission and propagation delays. It does not change when all the transmitted packets have the same format and size and use the same physical link [25]. The stochastic part, often called the ‘‘jitter’’, is mainly associated with the queueing delay [49] and varies from packet to packet, even when the packets have the same size and format.

If we denote by $\tau_{f,D}$ the deterministic part of τ_f and by $\tau_{f,S}$ the stochastic part:

$$\tau_f(t) = \tau_{f,D} + \tau_{f,S} \quad (\text{S3})$$

$$\tau_b(t) = \tau_{b,D} + \tau_{b,S} \quad (\text{S4})$$

In our experiments, we generate a forward delay that follows a normal distribution:

$$\tau_f(t) = \tau_{f,D} + \mathcal{N}(0, \sigma_{\tau_f})$$

For both simulations and real experiments, we set the deterministic part of the forward delay $\tau_{f,D}$ between 100ms to 1s (depending on the experiment, see the captions of each figure) and the jitter σ_{τ_f} to $\frac{2}{15}\tau_{f,D}$, which is in line with what the jitter usually represents [50].

$$\begin{aligned} \tau_{f,D} &\in [100, 1000] \text{ (depending on the experiment)} \\ \sigma_{\tau_f} &= \frac{2}{15}\tau_{f,D} \end{aligned} \quad (\text{S5})$$

(S6)

For the stochastic part of the backward delay, we assume that the robot uses a video streaming software that implements a jitter buffer (sections “Delay estimation by the robot” and “Jitter buffer”), which is the case of all the modern video streaming systems. If we set the jitter buffer length to d , then this buffer adds an additional deterministic delay of d : all the packets that arrive before d seconds are re-ordered and packets that are not arrived are dropped (dropping a few frames has little consequence for a state-of-the-art video codec). As a consequence, from the operator’s perspective, the backward delay is constant. This is why, for both simulations and real experiments, we generate a constant backward delay:

$$\tau_b(t) = \tau_{b,D}(t) \quad (S7)$$

For simplicity, we set $\tau_b(t) = \tau_{f,D}$ in all our experiments, but nothing in our system requires these two delays to be equal; in particular, it would be equivalent to set $\tau_b(t) = \tau_{b,D}(t) + K$ with K any constant delay caused by the jitter buffer.

C. Delay estimation by the robot

We assume that the clocks of the robot and of the computer of the operator are synchronized. In our real experiments, we synchronize the clock using the NTP system [26], which is the standard Unix protocol for time synchronization. The two clocks typically differ from less than 1ms on a local network [52]. Alternatively, GPS receivers can provide a highly accurate and absolute clock reference with an error of a few nano-seconds [52].

We add a time-stamp to each of the packets sent by the operator, which makes it possible for the robot to compute the forward delay $\tau_f(t)$ (this includes both the deterministic and stochastic part) when a packet is received at time t :

$$\tau_f(t) = \text{clock}_{\text{robot}}(t) - \text{timestamp}_{\text{operator}}(t) \quad (S8)$$

Please note that this does not assume that the delay follows a normal distribution. If necessary, the robot can re-order the packets according to the time-stamps to condition the trajectory predictions.

While the robot needs an estimate of the backward delay, it cannot know in advance the stochastic part before sending it. Our approach is to rely on the jitter buffer (section “Jitter buffer”) implemented in the video streaming system to make the backward delay deterministic: if we set the jitter buffer length to d s on the operator receiving side, then we know that the delay caused by the jitter will be exactly equal to d .

In our experiment, we therefore assume that the backward delay is known and constant (100 ms, 250ms, etc. depending on the experiments). To keep the implementation simple and easy to reproduce, we assumed that the deterministic backward delay was always equal to the deterministic forward delay (a reasonable assumption given that the same link is used for both directions) and that the stochastic part is negligible (because we chose to not add any jitter on the backward delay in the robot experiment, see the Jitter buffer section below):

$$\tau_b(t) = \overline{\tau_f(t)} = \tau_{f,D}(t) \quad (S9)$$

$$\tau_{b,S}(t) = 0 \quad (S10)$$

In a deployed setup, the robot would benefit from a better estimate of the average backward delay (the deterministic part) and the length of the jitter buffer. To do so, most video streaming systems use the RTCP protocol [28] to get out-of-band statistics and control information for a video streaming session that follows the RTP protocol [54]. This data would need to be sent periodically from the operator’s computer to the robot so that the robot knows both d and $\tau_{b,D}$ (which are not expected to change at high speed). Alternatively, a custom system can be designed by using time-stamps on the image packets to gather statistics about the delay.

D. Jitter buffer

The jitter buffer is the component of a video streaming system [27], [55] that re-orders packets if they are delayed by less than the length of the buffer and drops packets that are too late. Much work has been dedicated to adapt its length automatically [27]: if it is too small, then the video is jittery, but if it is too large, delays are added, which is detrimental to the user (in particular during video calls). In our system, we assume that the length is fixed and known to the robot, for simplicity.

We did not implement a jitter buffer because we wanted to avoid modifying the video streaming system: reordering or dropping packets would require a considerable expertise in the internals of both the encodings (e.g., MP4) and the video streaming software. Instead, we consider that video streaming with delay and jitter is a problem that is well solved by all the current video streaming systems, as experienced by the million of users who watch videos online on smartphones with non-ideal connections.

To summarize, from the point of view of our system, the jitter buffer results in an additional but deterministic delay. However, we assume that the robot knows the value of this additional delay as well as the deterministic part of the delay.

APPENDIX IV

PROBABILISTIC MOTION PRIMITIVES (PROMPS) FOR PRESCIENT TELEOPERATION

A. Definition of Probabilist Motion Primitives

A ProMP [31] is a probabilistic model for representing a trajectory distribution. The movement primitive representation models the time-varying mean and variance of the trajectories and is based on basis functions. A single trajectory is represented by a weight vector $\mathbf{w} \in \mathbb{R}^m$. The probability of observing a trajectory \mathbf{y} given the underlying weight vector is given as a linear basis function model

$$\xi_t = \Phi_t \mathbf{w} + \epsilon_\xi, \quad (S11)$$

$$p(\mathbf{y}|\mathbf{w}) = \prod_t \mathcal{N}(\xi_t | \Phi_t \mathbf{w}, \Sigma_\xi), \quad (S12)$$

where Σ_ξ is the observation noise variance, $\epsilon_\xi \sim \mathcal{N}(0, \Sigma_\xi)$ is the trajectory noise. The matrix $\Phi_t \in \mathbb{R}^m$ corresponds to

the m normalized radial basis functions evaluated at time t , with

$$\phi_c(t) = \frac{\exp\left(-\frac{1}{2}\left(t - \frac{c-1}{m-1}\right)^2\right)}{\sum_{c=1}^m \exp\left(-\frac{1}{2}\left(t - \frac{c-1}{m-1}\right)^2\right)}, \quad (\text{S13})$$

where the variable $c \in \{1, 2, \dots, m\}$ specifies the center of each basis function. The distribution $p(\mathbf{w}; \boldsymbol{\theta})$ over the weight vector \mathbf{w} is Gaussian, with parameters $\boldsymbol{\theta} = \{\boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w\}$ specifying the mean and the variance of \mathbf{w} . The trajectory distribution $p(\mathbf{y}; \boldsymbol{\theta})$ is obtained by marginalizing out the weight vector \mathbf{w} , i.e.

$$p(\mathbf{y}, \boldsymbol{\theta}) = \int p(\mathbf{y}|\mathbf{w})p(\mathbf{w}; \boldsymbol{\theta})d\mathbf{w}. \quad (\text{S14})$$

B. Learning ProMPs from demonstrations

The demonstrations are trajectories retargeted from the human. These are recorded in an "offline phase", while the user teleoperates the robot within a local network (approximately without delays) to perform a variety of tasks. Since the duration of each recorded trajectory may be different, a phase variable $v \in [0, 1]$ is introduced to decouple the movement from the time signal, obtaining a common representation in terms of primitives that is duration independent [32]. For each task, the modulated trajectories $\boldsymbol{\xi}_i(v)$ are then used to learn a ProMP. The parameters $\boldsymbol{\theta} = \{\boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w\}$ of the ProMP are estimated by using a simple maximum likelihood estimation algorithm. For each demonstration i , we compute the weights with linear ridge regression, i.e.

$$\mathbf{w}_i = (\boldsymbol{\Phi}_v^\top \boldsymbol{\Phi}_v + \lambda)^{-1} \boldsymbol{\Phi}_v^\top \boldsymbol{\xi}_i(v), \quad (\text{S15})$$

where the ridge factor λ is generally set to a very small value, typically $\lambda = 10^{-12}$ as in our case, as larger values degrade the estimation of the trajectory distribution. Assuming Normal distributions $p(\mathbf{w}) \sim \mathcal{N}(\boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w)$, the mean $\boldsymbol{\mu}_w$ and covariance $\boldsymbol{\Sigma}_w$ can be computed from the samples \mathbf{w}_i :

$$\boldsymbol{\mu}_w = \frac{1}{D} \sum_{i=1}^D \mathbf{w}_i, \quad \boldsymbol{\Sigma}_w = \frac{1}{D} \sum_{i=1}^D (\mathbf{w}_i - \boldsymbol{\mu}_w)(\mathbf{w}_i - \boldsymbol{\mu}_w)^\top, \quad (\text{S16})$$

where D is the number of demonstrations. Since a whole-body trajectory is represented by N trajectories (x , y , z position of the center of mass, of the hands, etc.), we learn a ProMP for each of the N trajectories. These ProMPs all together encode the same task k .

C. Recognizing the category of motion

Each learned k -th set of N ProMPs encodes different whole-body trajectories to accomplish a given task like picking up a box or squatting. To recognize to which set k the current teleoperated motion belongs to, we can minimize the distance between the first n_{obs} delayed observations and the mean of the N ProMPs of a group k , as done in [32]:

$$\hat{k} = \arg \min_{k \in [1:K]} \left[\sum_{n=1}^N \sum_{t \in T_{obs}} |\mathbf{y}_n(t - \tau_f(t)) - \Phi_{n, t - \tau_f(t)} \boldsymbol{\mu}_{n, \mathbf{w}_k}| \right], \quad (\text{S17})$$

where K is the number of tasks in the dataset and $T_{obs} = \{t_1, \dots, t_{n_{obs}}\}$ is the set of timesteps associated to the n_{obs} early observations. While computing \hat{k} , the ProMPs are modulated to have a duration equal to the mean duration of the demonstrations. The recognition (S17) starts whenever a motion is detected, i.e. the derivative of the observed end-effector trajectories exceeds a given threshold. The distance in (S17) is continuously computed after having identified the current motion. In this way, we can verify that the observations do not diverge from the demonstrations (exceed by a given threshold the demonstrated variance), in which case a gradual switch to delayed teleoperation is performed.

D. Time-modulation of the ProMPs

During motion recognition, we assumed that the duration of the observed trajectories is equal to the mean duration of the demonstrated trajectories, which might not be true. To match as closely as possible the exact speed at which the movement is being executed by the human operator, we have to estimate the actual trajectory duration. More specifically, we want to find the time modulation α , that maps the actual duration of a given (observed) trajectory to the mean duration of the associated demonstrated trajectories.

During the learning step, for each k -th set of ProMPs we record the set of α parameters associated to the demonstrations: $S_{\alpha k} = \{\alpha_1, \dots, \alpha_n\}$. Then, from this set, we can estimate which α better fits the current movement speed. We considered the best $\hat{\alpha}$ to be the one that minimizes the difference between the observed trajectory and the predicted trajectory for the first n_{obs} datapoints:

$$\hat{\alpha} = \arg \min_{\alpha \in S_{\alpha k}} \left\{ \sum_{t \in T_{obs}} |\mathbf{y}(t - \tau_f(t)) - \Phi_{\alpha(t - \tau_f(t))} \boldsymbol{\mu}_{\mathbf{w}_k}| \right\}. \quad (\text{S18})$$

E. Updating the posterior distribution of the ProMPs

Once the \hat{k} -th most likely set of ProMPs and their duration has been identified, we continuously update their posterior distribution to take into account the observations that arrive at the robot side. Each ProMP has to be conditioned to reach a certain observed state \mathbf{y}_t^* . The conditioning for a given observation $\mathbf{x}_t^* = \{\mathbf{y}_t^*, \boldsymbol{\Sigma}_y^*\}$ (with $\boldsymbol{\Sigma}_y^*$ being the accuracy of the desired observation) is performed by applying Bayes theorem

$$p(\mathbf{w}_{\hat{k}} | \mathbf{x}_t^*) \propto \mathcal{N}(\mathbf{y}_t^* | \Phi_{\hat{\alpha}t} \mathbf{w}_{\hat{k}}, \boldsymbol{\Sigma}_y^*) p(\mathbf{w}_{\hat{k}}). \quad (\text{S19})$$

The conditional distribution of $p(\mathbf{w}_{\hat{k}} | \mathbf{x}_t^*)$ is Gaussian with mean and variance

$$\hat{\boldsymbol{\mu}}_{\mathbf{w}_{\hat{k}}} = \boldsymbol{\mu}_{\mathbf{w}_{\hat{k}}} + \mathbf{L}(\mathbf{y}_t^* - \Phi_{\hat{\alpha}t}^\top \boldsymbol{\mu}_{\mathbf{w}_{\hat{k}}}), \quad (\text{S20})$$

$$\hat{\boldsymbol{\Sigma}}_{\mathbf{w}_{\hat{k}}} = \boldsymbol{\Sigma}_{\mathbf{w}_{\hat{k}}} - \mathbf{L} \Phi_{\hat{\alpha}t}^\top \boldsymbol{\Sigma}_{\mathbf{w}_{\hat{k}}}, \quad (\text{S21})$$

where

$$\mathbf{L} = \boldsymbol{\Sigma}_{\mathbf{w}_{\hat{k}}} \Phi_{\hat{\alpha}t} (\boldsymbol{\Sigma}_y^* + \Phi_{\hat{\alpha}t}^\top \boldsymbol{\Sigma}_{\mathbf{w}_{\hat{k}}} \Phi_{\hat{\alpha}t})^{-1}. \quad (\text{S22})$$

Given the delay in the transmitted data $\tau_f(t)$, we can compute the timestep t^* at which the ProMP has to be

conditioned to a certain observation \mathbf{x}_t^* :

$$t^* = t - \tau_f(t) - t_0. \quad (\text{S23})$$

where t_0 is the starting time of the current motion.

F. Motion anticipation

The references for the robot controller are generated at each time based on the updated ProMPs' mean trajectories $\hat{\boldsymbol{\mu}}_{\mathbf{w}_k}$. For a given ProMP, the sample $\hat{\boldsymbol{\mu}}_{\mathbf{w}_k}(t^*)$ corresponding to the last conditioned observation, is a reconstruction of the past retargeted human input

$$\hat{\boldsymbol{\mu}}_{\mathbf{w}_k}(t^*) = \hat{\mathbf{y}}(t - \tau_f(t)). \quad (\text{S24})$$

The sample $\hat{\boldsymbol{\mu}}_{\mathbf{w}_k}(t^* + \tau_f(t))$ is an estimate of the current retargeted human input

$$\hat{\boldsymbol{\mu}}_{\mathbf{w}_k}(t^* + \tau_f(t)) = \hat{\mathbf{y}}(t), \quad (\text{S25})$$

and can be used to synchronize the human movement with that of the robot, compensating only the forward delay (see Fig. 2). In our case, we want to synchronize the motion of the human operator with what is seen from the robot side, thus compensating for both the forward and backward delays. To do so, we select the sample $\hat{\boldsymbol{\mu}}_{\mathbf{w}_k}(t^* + \tau_f(t) + \hat{\tau}_b(t))$ as a control reference, which corresponds to a future prediction of the retargeted human movements:

$$\hat{\boldsymbol{\mu}}_{\mathbf{w}_k}(t^* + \tau_f(t) + \hat{\tau}_b(t)) = \hat{\mathbf{y}}(t + \hat{\tau}_b(t)). \quad (\text{S26})$$

The remaining samples $[\hat{\boldsymbol{\mu}}_{\mathbf{w}_k}(t^* + \tau_f(t) + \hat{\tau}_b(t) + 1), \hat{\boldsymbol{\mu}}_{\mathbf{w}_k}(t^* + \tau_f(t) + \hat{\tau}_b(t) + 2), \dots]$ are also given to the controller. They are used as control references if a new reference cannot be computed in the next control step due to packet losses or jitter.

After generating a first prediction, the transition from delayed to predicted references can be discontinuous (Fig. ??). To smoothen the transition, a policy blending arbitrates the delayed received references $\mathbf{y}(t - \tau_f(t))$ and the predicted ones $\hat{\mathbf{y}}(t + \hat{\tau}_b(t)|t - \tau_f(t))$, determining the adjusted reference (Fig. ??):

$$\hat{\mathbf{y}}'(t + \hat{\tau}_b(t)|t - \tau_f) = (1 - \boldsymbol{\beta})\mathbf{y}_d + \boldsymbol{\beta}\mathbf{y}_p, \quad (\text{S27})$$

where $\mathbf{y}_d = \mathbf{y}(t - \tau_f(t))$, $\mathbf{y}_p = \hat{\mathbf{y}}(t + \hat{\tau}_b(t)|t - \tau_f(t))$, $\boldsymbol{\beta} = \{\beta_0, \dots, \beta_n, \dots, \beta_N\}^\top$ with $\beta_n \in]0, 1[$

$$\beta_n = \frac{1}{1 + e^{-12(\frac{i}{\Delta y_n} - \frac{1}{2})}}, \quad (\text{S28})$$

$i = \{0, 1, \dots, \Delta y_n\}$ and Δy_n is the initial difference between a delayed reference and the corresponding prediction expressed in *mm* (for Cartesian trajectories) or *deg* $\times 10^{-1}$ (for postural trajectories).

G. Teleoperation under unexpected circumstances

If something unexpected happens, or if the operator suddenly changes their mind about what to do and the ongoing motion cannot be completed or is significantly altered, the prescient teleoperation is transitioned back to the delayed teleoperation. The transition from predicted to delayed references is triggered whenever the distance between the current

observation and learned mean exceeds a given threshold Δ_σ , which in the experiments was fixed equal to the learned variance plus 5cm and considered for each of the x, y, z trajectories of the hands. Since the transition can be discontinuous, a policy blending arbitrates the last predicted sample $\hat{\mathbf{y}}(t_{last} + \hat{\tau}_b(t_{last})|t_{last} - \tau_f(t_{last}))$ and the delayed received references $\mathbf{y}(t - \tau_f(t))$:

$$\hat{\mathbf{y}}'(t + \hat{\tau}_b(t)|t - \tau_f) = (1 - \boldsymbol{\beta})\mathbf{Y}_p + \boldsymbol{\beta}\mathbf{y}_d, \quad (\text{S29})$$

where $\mathbf{y}_d = \mathbf{y}(t - \tau_f(t))$, $\mathbf{Y}_p = \hat{\mathbf{y}}(t_{last} + \hat{\tau}_b(t_{last})|t_{last} - \tau_f(t_{last}))$, $\boldsymbol{\beta} = \{\beta_0, \dots, \beta_n, \dots, \beta_N\}^\top$ with β_n defined as in (S28).

H. Comparison between ProMPs and LSTM

We implemented a LSTM (Long Short Term Memory network) [59] using the Pytorch library [60]. This LSTM predicts the next k time-steps given the previous k time-steps. To keep the network small, the trajectories are sub-sampled from 100Hz to 10Hz, so that 10 time-steps correspond to 1 second of motion. For instance, to predict the next 2 seconds of motions, the LSTM has 20 inputs and 20 outputs. The network has 10 hidden nodes (preliminary experiments showed that increasing this number did not have any qualitative effect on the predictions). It is trained with the Adam optimizer using mini-batches of size 16, for 100 epochs (our experiments show that this is enough to reach the minimum loss), and the Mean Squared loss function. To account for the stochasticity of the initialization and the stochastic gradient descent, 10 LSTMs are trained with different seeds.

Like with ProMPs, one LSTM is trained for each dimension of the prediction (e.g., a LSTM for the x-coordinate of the left hand, another one for the y-coordinate, etc.); however, contrary to ProMPs, the whole ‘‘Multi tasks’’ training set is used to train each LSTM. A different set of LSTMs is trained for 1-second prediction (Fig. S6) and 2-second prediction (Fig. S7). The LSTMs that were trained for 2 seconds could have been used to predict 1-second ahead, but at the risk of being lower-performing than a network specialized in 1-second predictions.

For this comparison, the ProMPs are conditioned at each time-step with all the points since the beginning of the trajectory, and the ProMP is queried to predict the value at $t + 2$ seconds (or $t + 1$ s). As an additional baseline, the ‘‘delayed’’ trajectory is the operator’s trajectory shifted in time by either one or two seconds, which corresponds to what the robot would do if there were no compensation.

For a particular trajectory T and a particular dimension D (e.g., x-position of the hand), the prediction error is the sum of the differences between the prediction and the recorded trajectory:

$$e_{T,D} = \frac{1}{N} \sum_{i=1}^{i=N} \left| x_i^{(pred)} - x_i^{(gt)} \right| \quad (\text{S30})$$

where N is the number of points in the trajectory, $x_i^{(pred)}$ the prediction and $x_i^{(gt)}$ the point of the trajectory performed by the operator at time-step i (unknown to the robot at this time-step).

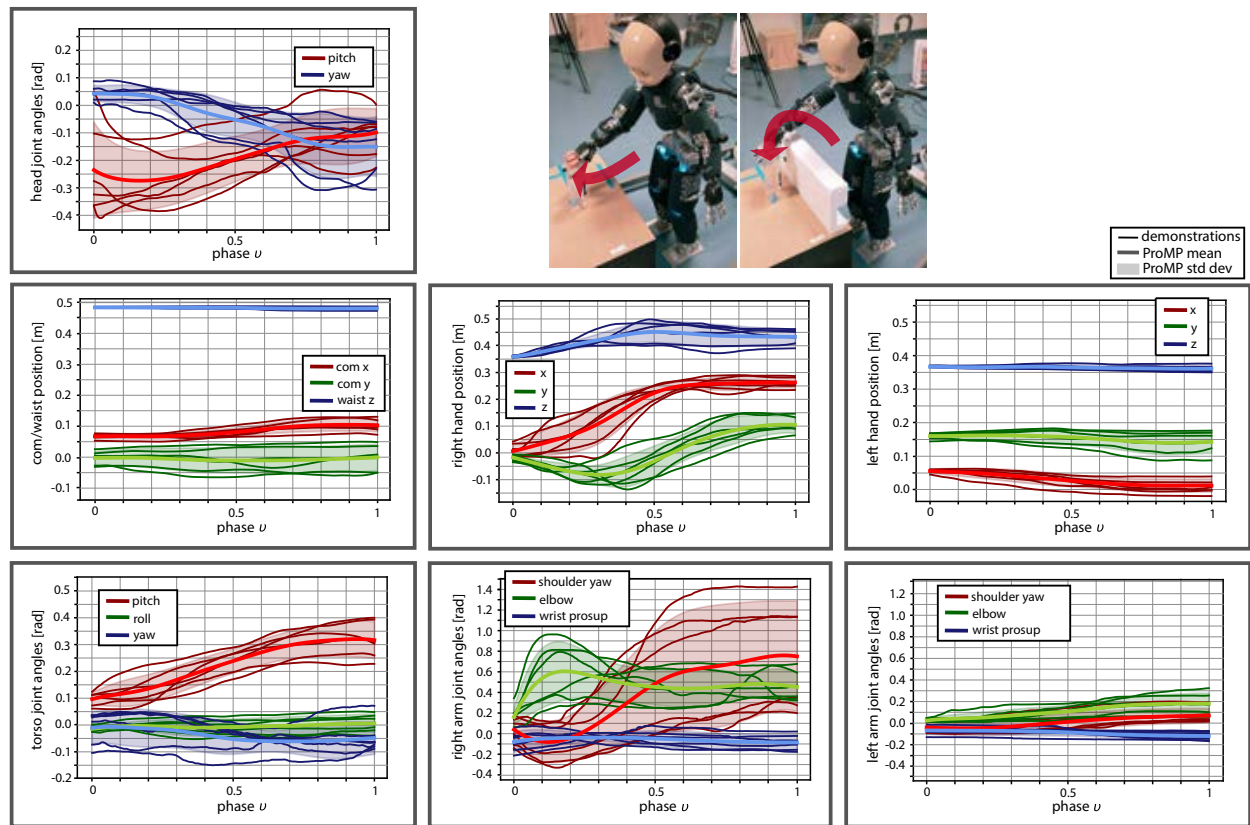
APPENDIX V SOURCE CODE

The code of the ProMP library is available at <https://doi.org/10.5281/zenodo.7438257> and on <https://github.com/hucebot/promp>. The comparison between ProMP and LSTM is available as a python notebook at <https://doi.org/10.5281/zenodo.7441367> and on https://github.com/hucebot/lstm_vs_promp (this code relies on the ProMP library).

REFERENCES

- [38] A. Escande, N. Mansard, and P.-B. Wieber, "Hierarchical quadratic programming: Fast online humanoid-robot motion generation," *The International Journal of Robotics Research*, vol. 33, no. 7, pp. 1006–1028, 2014.
- [39] L. Penco, E. M. Hoffman, V. Modugno, W. Gomes, J. Mouret, and S. Ivaldi, "Learning robust task priorities and gains for control of redundant robots," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2626–2633, 2020.
- [40] A. Rocchi, E. M. Hoffman, D. G. Caldwell, and N. G. Tsagarakis, "Opensot: a whole-body control library for the compliant humanoid robot coman," in *IEEE ICRA*, 2015, pp. 6248–6253.
- [41] L. Natale, C. Bartolozzi, D. Pucci, A. Wykowska, and G. Metta, "iCub: The not-yet-finished story of building a robot child," *Science Robotics*, vol. 2, no. 13, p. eaaq1026, 2017.
- [42] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl, "qpOASES: A parametric active-set algorithm for quadratic programming," *Mathematical Programming Computation*, vol. 6, no. 4, pp. 327–363, 2014.
- [43] L. Penco *et al.*, "Robust real-time whole-body motion retargeting from human to humanoid," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2018, pp. 425–432.
- [44] D. Roetenberg, H. Luinge, and P. Slycke, "Xsens mvn: Full 6dof human motion tracking using miniature inertial sensors," *Xsens Motion Technol. BV Tech. Rep.*, vol. 3, 01 2009.
- [45] T. L. Foundation, "Iproute2." [Online]. Available: <https://wiki.linuxfoundation.org/networking/iproute2>
- [46] "Kinovea. A microscope for your videos," <https://www.kinovea.org/>.
- [47] "Virtual Desktop. Your PC in VR," <https://www.vrdesktop.net/>.
- [48] O. Gurewitz, I. Cidon, and M. Sidi, "One-way delay estimation using network-wide measurements," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2710–2724, 2006.
- [49] M. Garetto and D. Towsley, "Modeling, simulation and measurements of queuing delay under long-tail internet traffic," in *Proceedings of the 2003 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS '03. New York, NY, USA: Association for Computing Machinery, 2003, pp. 47–57. [Online]. Available: <https://doi.org/10.1145/781027.781034>
- [50] A. Alharbi, A. Bahnasse, and M. Talea, "A comparison of vpn performance evaluation on different environments over vpn multipoint network," *International Journal of Computer Science and Network Security*, vol. 17, pp. 123–128, 05 2017.
- [51] D. Mills, J. Martin, J. Burbank, and W. Kasch, "Network time protocol version 4: Protocol and algorithms specification," Internet Requests for Comments, RFC 5905, June 2010.
- [52] D. Veitch, S. Babu, and A. Pásztor, "Robust synchronization of software clocks across the internet," in *Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '04. New York, NY, USA: Association for Computing Machinery, 2004, p. 219–232. [Online]. Available: <https://doi.org/10.1145/1028788.1028817>
- [53] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: A transport protocol for real-time applications," Internet Requests for Comments, RFC Editor, STD 64, 2003.
- [54] J. Ott, S. Wenger, N. Sato, C. Burmeister, and J. Rey, "Extended RTP profile for real-time transport control protocol (RTCP)-based feedback (RTP/AVPF)," *RFC*, vol. 4585, pp. 1–51, 2006.
- [55] M. Claypool and J. Tanner, "The effects of jitter on the perceptual quality of video," in *Proceedings of the seventh ACM international conference on Multimedia (Part 2)*, 1999, pp. 115–118.
- [56] B. Oklander and M. Sidi, "Jitter buffer analysis," in *Proc. Int. Conf. on Comp. Comm. Networks, ICCCN*, 09 2008, pp. 1 – 6.
- [57] A. Paraschos, C. Daniel, J. Peters, and G. Neumann, "Using probabilistic movement primitives in robotics," *Autonomous Robots*, vol. 42, p. 529–551, 07 2018.
- [58] O. Dermy, A. Paraschos, M. Ewerton, J. Peters, F. Charpillat, and S. Ivaldi, "Prediction of intention during interaction with iCub with probabilistic movement primitives," *Frontiers in Robotics and AI*, vol. 4, pp. 45–45, 2017.
- [59] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [60] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, *et al.*, "Pytorch: An imperative style, high-performance deep learning library," *Advances in neural information processing systems*, vol. 32, 2019.

a. Reaching a bottle (training set), dataset Multiple Tasks



b. Reaching a bottle (test set), dataset Multiple Tasks

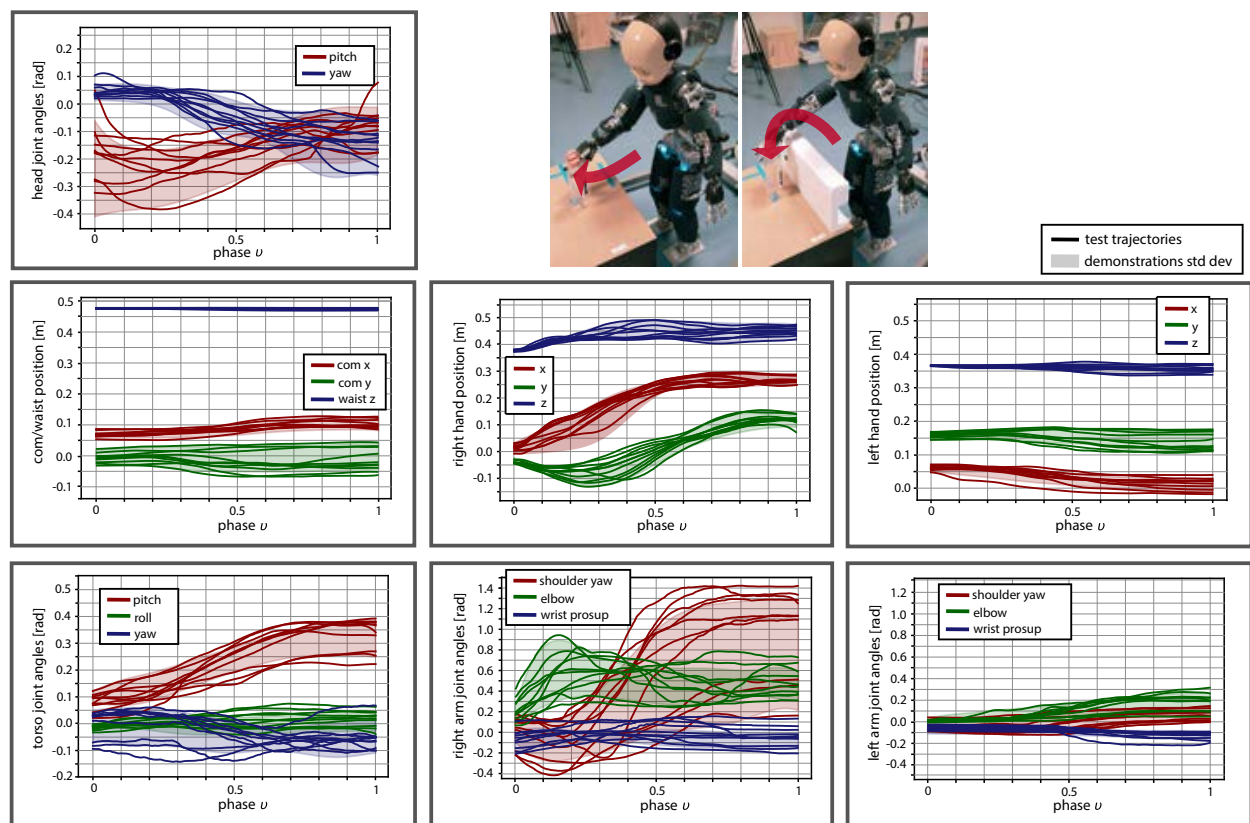
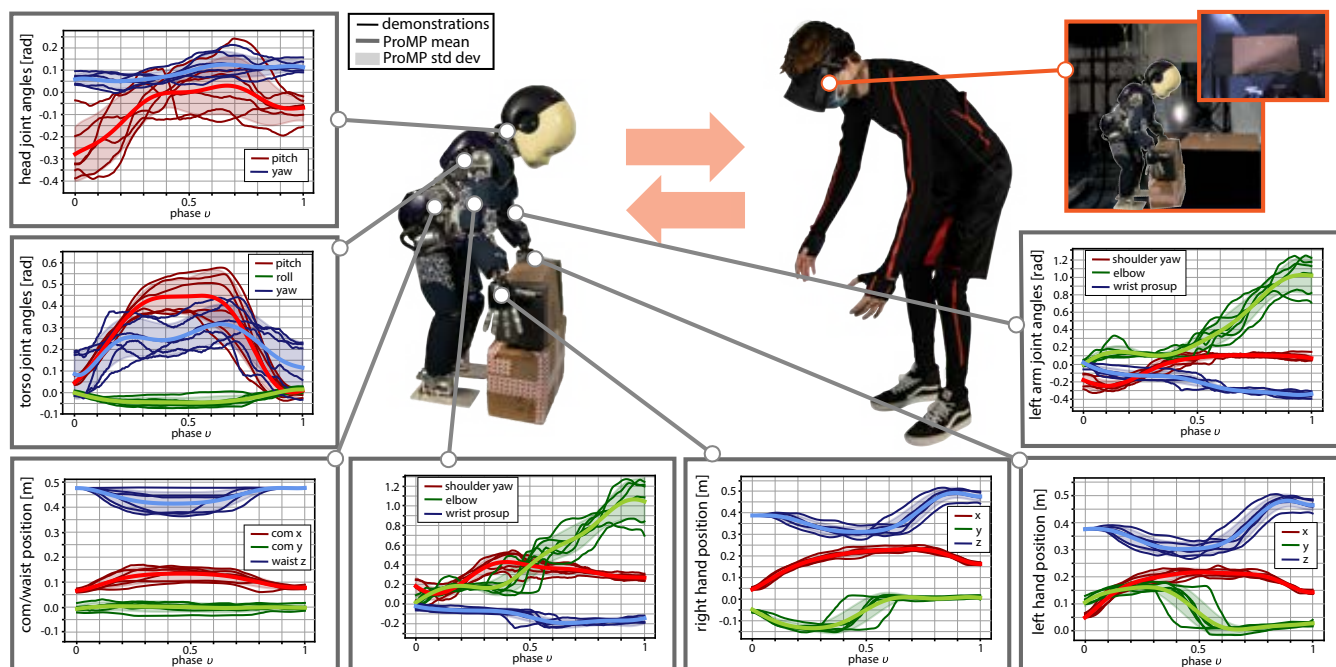


Fig. S1. **Dataset “Multiple Tasks”, scenario reaching a bottle on the table. a. Training trajectories and learned ProMPs.** The whole-body motion of the teleoperated robot is obtained by following the reference trajectories retargeted from the human. We learned a ProMP for each of these trajectories, given 6 demonstrations in a local network without any delay (3 with an obstacle in between the robot and the bottle, and 3 without). **b. Test trajectories.** The test trajectories are different and additional repetitions of the training motions. For the task of reaching a bottle 10 different repetitions were recorded (5 with an obstacle in between the robot and the bottle, and 5 without).

a. Picking up a box (training set), dataset Multiple Tasks



b. Picking up a box (test set), dataset Multiple Tasks

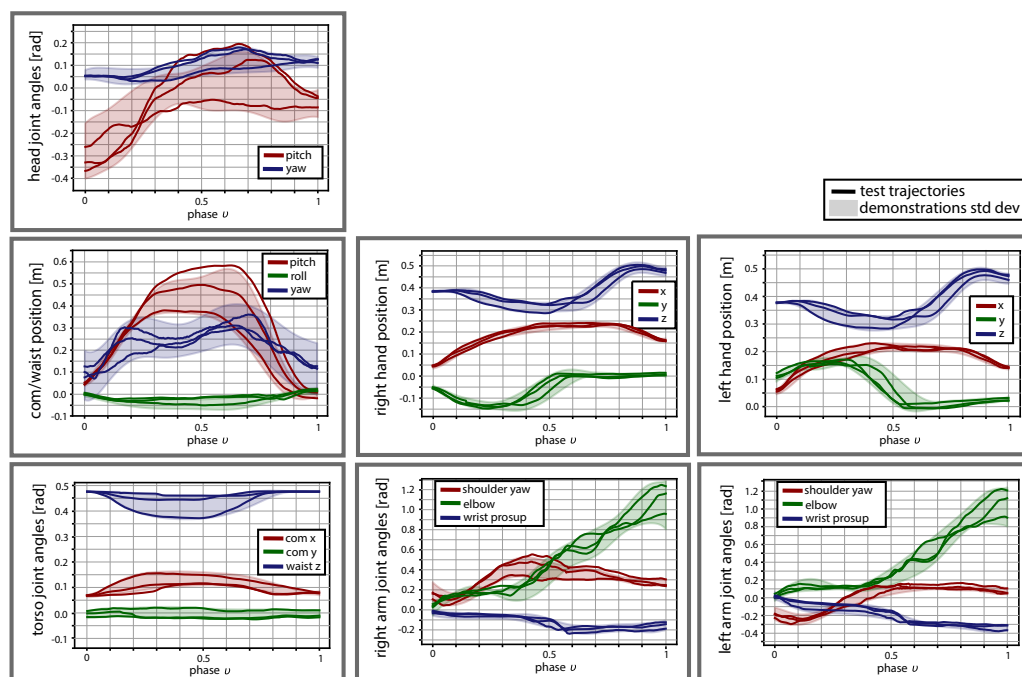


Fig. S2. Dataset “Multiple Tasks”, scenario picking up a box at a low position. **a. Training trajectories and learned ProMPs.** For each of the 6 demonstrations, the motion of the operator is first “retargeted” to the robot using the whole-body controller (ignoring delays). The trajectories of each body/joint of the robot is then recorded. From this set of demonstrations (thin lines), a ProMP is fitted for each trajectory; this ProMP is represented here as a thick line (the mean) and a light zone (the standard deviation). The computed mean is a smooth trajectory that averages all the demonstrations and the standard deviation captures the variability of the demonstrations. **b. Test trajectories.** The test trajectories are different and additional repetitions of the training motions. For the tasks of picking up a box 3 different repetitions were recorded.

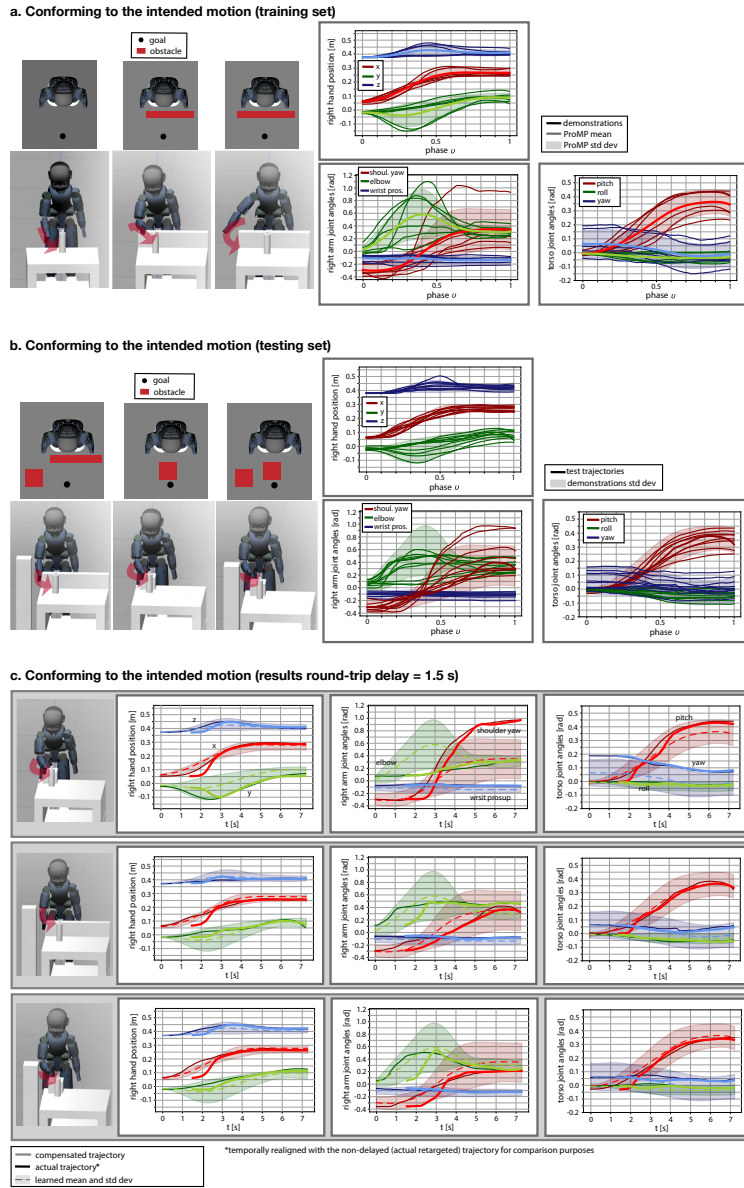
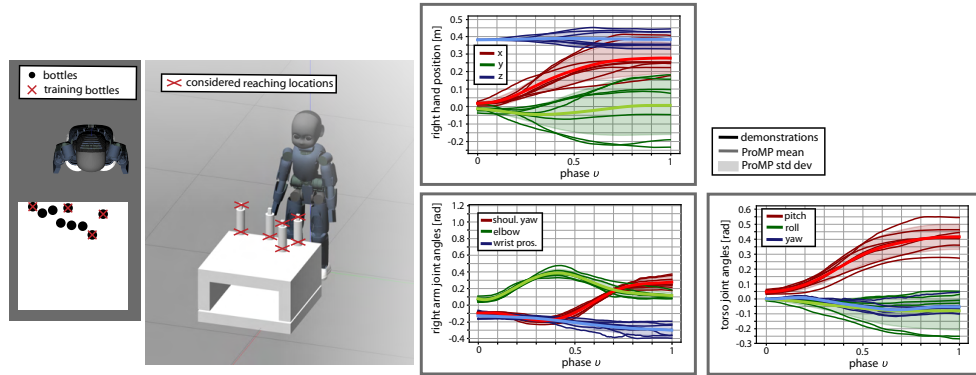
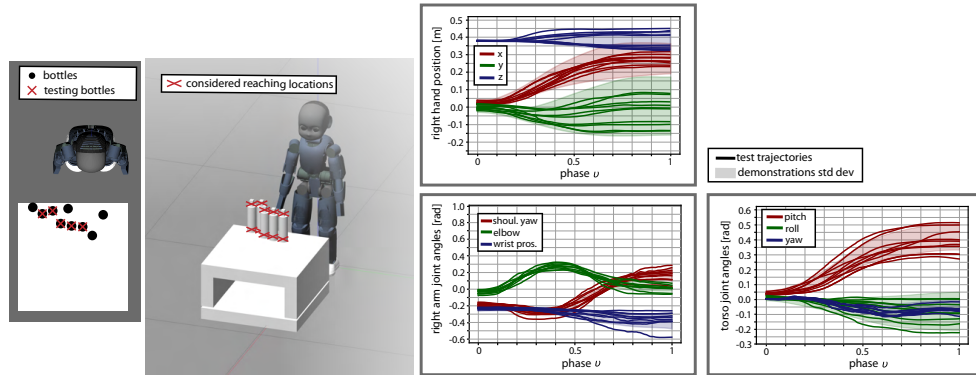


Fig. S3. **Dataset “Obstacles”:** **Conforming the teleoperation to the intended motion.** **a. Training trajectories and learned ProMPs.** The most relevant learned ProMPs and the associated demonstrations are reported. The 6 demonstrations (2 without obstacles, 2 with an obstacle in between the robot and the bottle and other 2 with a different obstacle) have been recorded while teleoperating the robot in simulation, in a local network without any delay. **b. Test trajectories.** The 9 test trajectories are different from those used for training and consist of 3 repetitions of the bottle reaching motion for each of the 3 distinct simulated scenarios with different obstacles, illustrated on the left. **c. Results: comparison between the compensated trajectory and the ideal (non-delayed) trajectory with a mean round-trip delay of 1.5s.** On the top row, there is an unexpected obstacle (a small box) to avoid and the operator approaches the object by the right; on the second row, the obstacle in front of the robot is different; on the bottom row, there is the same obstacle from the top row with in addition a large obstacle on the right of the robot, which forces the operator to move the hand in between the two obstacles. These situations were not in the training set. After the initial recognition period, our approach makes the robot follow the specific way the human is performing the task despite the delay, and even if the robot is asked to perform the task in a way that has not been demonstrated before (but included in the distribution of the demonstrations). This is not the same as following the mean of previously demonstrated motions (here, the dashed line) or letting the robot replicate previously demonstrated motions. The non-delayed trajectories are some of the test trajectories from panel b, where the robot has to reach the bottle on the table in the presence of different obstacles that were not considered during the training.

a. Conforming to new goals (training set)



b. Conforming to new goals (testing set)



c. Conforming to new goals (results round-trip delay = 1.5 s)

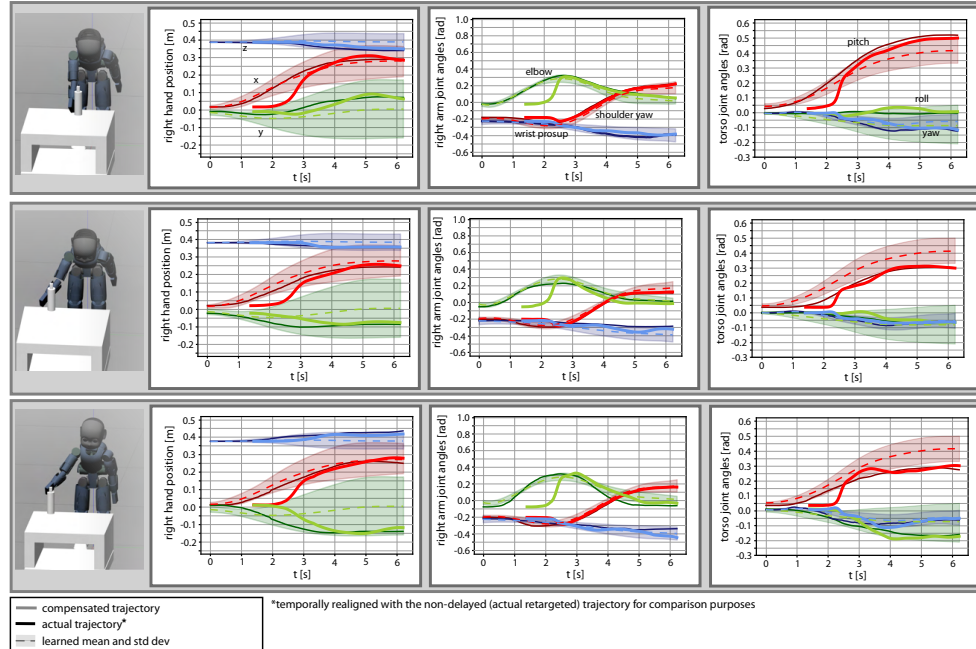


Fig. S4. **Dataset “Goals”:** Conforming the teleoperation to new goals. **a. Training trajectories and learned ProMPs.** The most relevant learned ProMPs and the associated demonstrations are reported. The 7 demonstrations have been recorded while teleoperating the robot in simulation, in a local network without any delay. The different bottle locations are illustrated on the left. **b. Test trajectories.** The 10 test trajectories are different from those used for training. The different bottle locations are illustrated on the left. **c. Results: comparison between the compensated trajectory and the ideal (non-delayed) trajectory with a mean round-trip delay of 1.5s.** The bottles are located on the table in different positions that were not in the training set (but included in the distribution of the demonstrations). The non-delayed trajectories are some of the test trajectories from panel b, where the robot has to reach the bottle on the table in the presence of different obstacles that were not considered during the training.

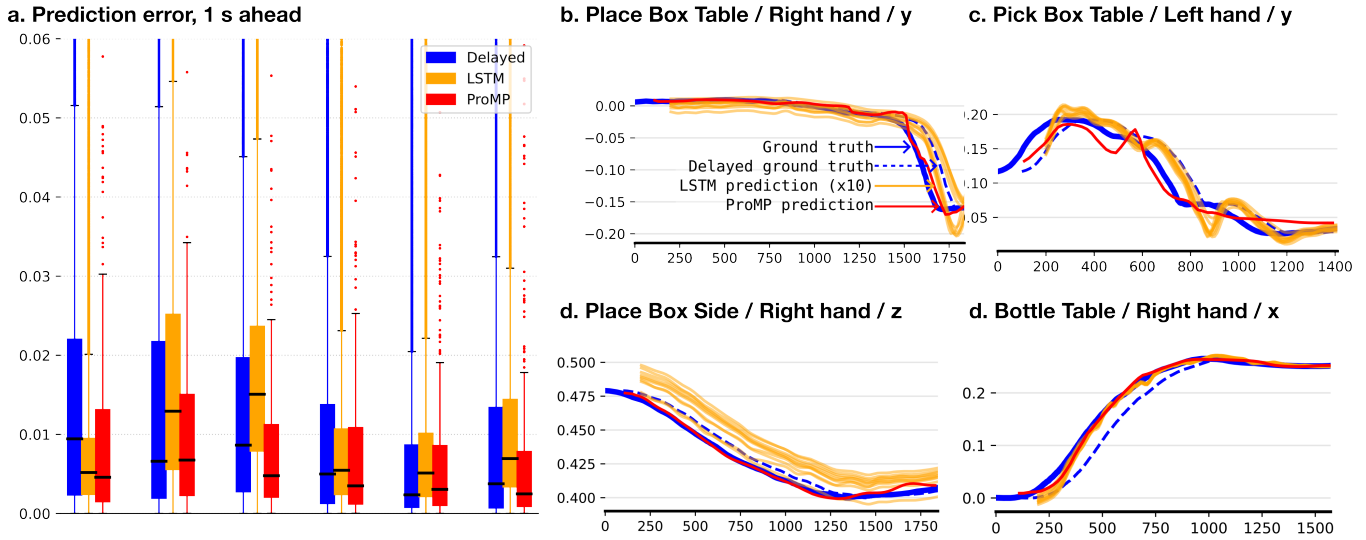


Fig. S6. **Prediction for LSTM and ProMP, 1 second ahead.** **a. Prediction error, 1-s horizon.** The plot shows the average difference between the point predicted 1 second ahead (100 time-steps) and the ground truth (Methods), for each degree coordinate of each hand, for each trajectory of the dataset “Multiple tasks” (Methods). The box extends from the lower to upper quartile values of the data, with a line at the median. For the LSTM, the 10 replicates of the learning process (with different seeds) are considered as independent data (i.e., the variance of the prediction error comes both from the different trajectories and the different seeds). The “Delayed” trajectory corresponds to the original trajectory delayed by 1 second, that is, to what the robot would have done without any prediction system. All comparisons are statistically significant ($p < 10^{-6}$, Mann-Whitney U-test). **b-d. Examples of predicted trajectories.**

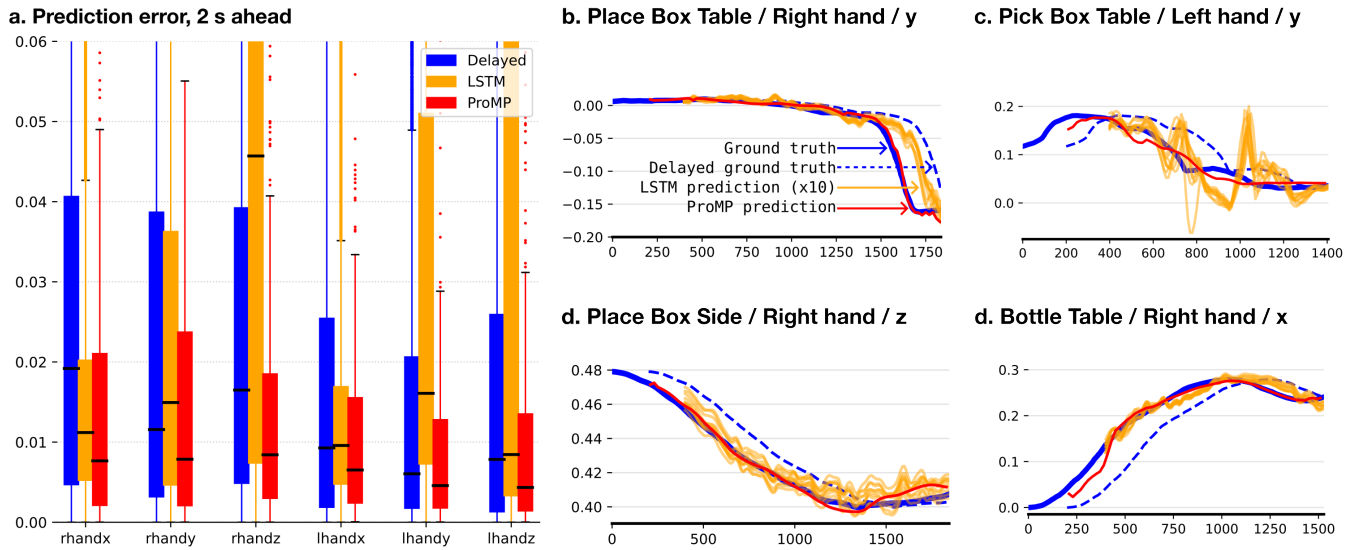


Fig. S7. **Prediction for LSTM and ProMP, 2 seconds ahead.** **a. Prediction error, 2-s horizon.** The plot shows the average difference between the point predicted 2 seconds ahead (200 time-steps) and the ground truth (Methods), for each degree coordinate of each hand, for each trajectory of the dataset “Multiple tasks” (Methods). The box extends from the lower to upper quartile values of the data, with a line at the median. For the LSTM, the 10 replicates of the learning process (with different seeds) are considered as independent data (i.e., the variance of the prediction error comes both from the different trajectories and the different seeds). The “Delayed” trajectory corresponds to the original trajectory delayed by 2 seconds, that is, to what the robot would have done without any prediction system. All comparisons are statistically significant ($p < 10^{-6}$, Mann-Whitney U-test). All comparisons are statistically significant ($p < 10^{-6}$, Mann-Whitney U-test). **b-d. Examples of predicted trajectories.**

TABLE S2

DIFFERENCE (ROOT MEAN SQUARE ERROR) WITH THE NON-DELAYED TRAJECTORIES, FOR BOTH THE COMPENSATED AND THE NON-COMPENSATED (DELAYED) TRAJECTORIES (AVERAGE DELAY: 1.5 s) THE ERROR IS COMPUTED FOR THE 20 TESTING MOTIONS FROM THE BOTTLE REACHING SCENARIO OF THE DATASET MULTIPLE TASKS (FIG. S1), AND FOR THE 21 TESTING MOTIONS FROM THE BOX HANDLING SCENARIO OF THE DATASET MULTIPLE TASKS (FIG. S2). THE TIME-VARYING FORWARD FOLLOWS A NORMAL DISTRIBUTION WITH 750MS AS MEAN AND 100MS AS STANDARD DEVIATION. THE BACKWARD DELAY IS SET EQUAL TO 750MS.

	Box handling RMS error [rad]		Bottle reaching RMS error [rad]	
	compensation	no compensation	compensation	no compensation
head yaw	0.024±0.011	0.035±0.012	0.013±0.007	0.021±0.011
torso pitch	0.045±0.020	0.136±0.064	0.027±0.012	0.041±0.019
torso roll	0.022±0.011	0.089±0.038	0.015±0.008	0.020±0.010
torso yaw	0.069±0.028	0.129±0.055	0.019±0.009	0.032±0.011
r. shoulder yaw	0.071±0.025	0.145±0.051	0.065±0.018	0.221±0.092
r. elbow	0.062±0.020	0.171±0.067	0.096±0.030	0.194±0.071
r. wrist prosup.	0.025±0.007	0.077±0.033	0.054±0.012	0.091±0.041
	RMS error [cm]		RMS error [cm]	
	compensation	no compensation	compensation	no compensation
r. hand x	1.02±0.31	2.95±1.12	1.29±0.33	4.97±1.46
r. hand y	0.90±0.26	3.36±1.21	1.21±0.31	4.33±1.17
r. hand z	0.96±0.30	2.96±0.75	1.11±0.25	4.15±1.13
com x	0.90±0.13	1.24±0.36	0.33±0.07	1.01±0.20
com y	0.79±0.11	1.06±0.39	0.24±0.06	1.01±0.32
waist z	0.88±0.14	2.02±1.22	0.22±0.07	0.61±0.09