

Human-like Robot Interaction: A Framework for Generating Natural Robot Dynamics Using Bilateral Control and Transformer Neural Networks

ANONYMOUS AUTHOR(S)




Fig. 1. A schematic representation representing the work described in this paper. A 4-channel teleoperation setup allows the collection of force data during a displacement task. Data are then fed to a transformer, which learns to replicate human-generated resistive forces. Finally, a modified Turing test shows that the forces generated by the transformer are indistinguishable for human generated forces.

In the advancing field of human-robot interaction, the ability of robots to authentically replicate human behavior is essential for their integration into industrial and everyday settings. This paper introduces a framework to enable robotic mimicry of human dynamics. We utilize force-based bilateral control to enable bidirectional, robot-mediated human-to-human interactions. Interaction data are collected and fed to transformer models capable of learning and replicating human dynamic responses. The approach is validated through a modified Turing test, where participants are unable to distinguish between human-operated and the movements generated by our framework. These results confirm the model's ability to mimic human-generated interaction forces convincingly. We discuss the implications of our results, suggesting significant potential for the application of our framework in scenarios requiring robots mediated human-to-human interaction to deliver forces in a natural, human-like manner.

CCS Concepts: • **Do Not Use This Code** → **Generate the Correct Terms for Your Paper**; *Generate the Correct Terms for Your Paper*; Generate the Correct Terms for Your Paper; Generate the Correct Terms for Your Paper.

Additional Key Words and Phrases: Do, Not, Us, This, Code, Put, the, Correct, Terms, for, Your, Paper

ACM Reference Format:

Anonymous Author(s). 2018. Human-like Robot Interaction: A Framework for Generating Natural Robot Dynamics Using Bilateral Control and Transformer Neural Networks. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 15 pages. <https://doi.org/XXXXXXX.XXXXXXX>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

Manuscript submitted to ACM

1 Introduction

The ability of robots to mimic human behaviours is pivotal for advancing both industrial applications and everyday human-machine interactions. Human-like movements and behaviors in robots can enhance interaction by making the robots' actions more predictable and relatable for humans [27], [10].

Current robotic systems have made significant advances in achieving human-like level of performance for numerous tasks, with imitation learning (IL) being one of the most successful ways to achieve such feat.

Interestingly, the great majority of IL literature focuses on obtaining human-like performances, such as biologically plausible motion trajectories [21] or controlling forces in a pincher grasp well enough to pick up objects [31]. However, to this date, no IL-based system or framework focuses on generating robots that not only *perform* like humans but *feel* human-like as well.

Yet, generating seamless human-like robot interactions has been a complex issue often revisited by researchers in the field of haptics, which strived to achieve human likeness of the haptic feedback on generic impedance devices[3] , [25] as well as more task-oriented ones, such as handshake devices [33] or even kissing devices [9] In this work we propose an IL-based framework that allows a robotic arm (Franka) to respond to human-generated forces in a displacement task in a way that is indistinguishable from a human. The IL in this work focuses specifically on reaction forces, which our framework can learn directly and decoupled from position information thanks to the use of data acquisition protocol leveraging a 4-channel bilateral control [13].

This robot's ability to deliver haptic(force) feedback that feels human-generated has potential repercussions in many areas, as for example social haptics, where affective robots have been used to soothe and emotionally support users [46]. It can also be useful in human-robot collaboration, where biologically plausible haptic feedback could work alongside human-like trajectories to facilitate object handover asks [?]. Most relevant in the case of our work, human-like haptic feedback could be used in medical training [11] and remote medical treatment [37].

Our work is particularly relevant to the field of robotic telerehabilitation and AI-powered healthcare. In robotic rehabilitation, robots can assist and even replace physiotherapists delivering treatments to patients [37]. While this approach has often yielded outstanding results (e.g. [30]), recent works highlighted the limitation of a fully robotic telerehabilitation system, showing that the presence of a therapist in the treatment can increase motivation and engagement in patients [32].

Moreover, telerehabilitation protocols using robotic systems have been shown to lack adaptability to the patients' variable performances [2]. Artificial Intelligence has been often proposed to obviate this issue, and it has been applied to the identification of a patient's state and the optimization of treatment plans [8]. Our framework leverages AI (and more specifically, transformers [44] to develop the very control of the robot and allows the system to resist displacement attempts in a human-like manner. Such resistive effort is often used by therapists to train the patient's force and motor control in training exercises where the patient is asked to extend or curl their limbs while the therapist attempts to keep them in place (Manual Resistance Techniques, MRTs, [28].

By presenting a framework able to successfully replicate human resistive forces to a displacement task, this work is the first crucial step toward a new way of developing control for telerehabilitation devices. The contributions of this work spur forward a vision where the robots' feedback is able to seamlessly mimic the feedback from a therapist, therefore mitigating the limitations concerning the adaptability of fully robotics systems and even supporting the design of AI therapists able to engage patients with seamless human-like interactions. In summary, the contributions of this work are as below:

- We present a framework to replicate human reaction forces in a displacement task. This framework is the first of its kind to leverage both a 4-channel bilateral control to record force feedback decoupled from the position and a transformer model to generate a control capable of imitating human-generated force feedback.
- We show that this framework can successfully generate an IL-based control providing human-like reaction and resistive force feedback in a displacement task.
- We discuss the results of our test as highly relevant for the field of robotic telerehabilitation, defending that this work represents a first step towards a novel way to design telerehabilitation controls.

Section I introduces the work to the reader, its contribution, and the context of impact. Section II discusses the technological foundation of our model providing an overview of the literature related to its components, namely bilateral control and imitation learning. It also provides a short review of telerehabilitation systems, providing the foundation to understand the selection of the task in our experiment. Section III includes the first contribution of this work, describing our framework, including the bilateral 4-channel controller and the transformer-based IL. Section IV describes the details of the experimental setup and evaluates reports on the results of our study. It represents the second contribution of this work. Section V explores the implications of these findings, while Section VI concludes by highlighting the limitations of our current setup and potential research directions of interest for the HCI community

2 Related Works

2.1 Telerehabilitation systems

Our framework aims at simulating human behavior and force responses in robot-mediated, human-to-human interaction. This type of interaction is common in the field of human-robot and human-computer interaction, appearing in contexts such as social haptic (e.g., [7]), Virtual Reality (e.g., [3]) and healthcare (e.g., [23]).

On the latter, robot-mediated human-to-human interaction finds multiple applications. For example, surgical robots can be teleoperated by medical professionals to carry out complex procedures [42]. Crucially, these systems can also provide haptic feedback to their users, conveying the feeling of the forces acting on them from the patient's side [36]. Alongside teleoperation systems, telerobotic rehabilitation has recently gained momentum, due to the recent pandemic [6]. Telerobotic rehabilitation systems allow therapists to manipulate, train, and even perceive patients' reactions via robotic proxy [23].

Many systems include the delivery of force feedback to the therapist, allowing them to adapt treatment dynamically within sessions based on patient performance (e.g., [4, 17]). Some systems bypass direct feedback to the therapist and instead leverage adaptive technologies, such as AI [29] or soft, adaptive actuation [38] to vary the difficulty of tasks in real time. Crucially, these latter implementations, while adaptable, largely exclude the therapist from the treatment loop. This has the obvious advantage of providing patients with treatment anytime it is needed, without expending the time and resources of a healthcare professional. On the other hand, the role of the therapist goes beyond simply granting adaptability. In [10] authors highlight the role of therapists in maintaining patients' engagement and motivation throughout the treatment. The work presented in this paper constitutes a first step towards an adaptable system that, despite removing the therapist from the treatment loop, allows users to experience treatments as if a human therapist delivered them, therefore maintaining the patients' motivation and engagement.

2.2 Bilateral controllers

In our research, we developed a framework to first record, and then learn, the resistive forces in a specific robot-mediated displacement task. The exchange and the recording of force data between 2 users during this task is enabled through bilateral control of the robots the users interact with.

The application of bilateral controllers in teleoperation has been well-documented, leading to multiple control strategies and applications in various domains [16, 18, 20]. Bilateral controls are particularly effective when the system needs to convey users with haptic feedback, for example in robotic surgery [35] and remote maintenance in hazardous environments (including space exploration [43]).

As in the case of this study, bilateral controllers have also been used to enable human-to-human interaction via robotic proxy. For example, [3] describes a system where two users interact in a virtual environment, each receiving force feedback based on the other’s input. In [4] multiple examples are included where bilateral controls are used in telerehabilitation settings, with therapists and users providing each other haptic feedback as a consequence of their movements.

In this work, we adopt a 4-channel bilateral controller. 4-channel bilateral control systems utilize four distinct feedback channels: position and force control in both master-to-slave and slave-to-master directions. The master system transmits position/velocity commands to the slave, which replicates the operator’s movements. Simultaneously, the slave measures and sends interaction forces (from the environment) back to the master. These forces are reflected to the operator via the master device, providing haptic feedback [43]. The key advantage is the decoupling of position and force feedback, allowing independent control and tuning of both, which enhances transparency (fidelity of feedback) and stability, especially in the presence of time delays or dynamic environments [26], as well as improving the operational symmetry and interaction fidelity between the operator and the robotic system.

Importantly, in our framework, the decoupling of force feedback from position information serves a further purpose, as it allows our transformer to directly learn force information from the interaction data recorded during our experimental task.

2.3 Human-like robot behaviour via imitation learning

The concept of Imitation Learning (IL) in robotics has been extensively explored to enable robots to perform tasks by mimicking human actions. Seminal works such as [1] laid foundational principles for this domain, highlighting the potential of learning from the demonstration as a method to teach robots complex behaviours without explicit programming.

Recent advancements have shown an increasing interest in the integration of force dynamics within imitation learning frameworks. For instance, [47] demonstrated the significance of force feedback in enhancing the accuracy and reliability of robotic tasks involving delicate human interactions, such as in medical surgeries or caregiving. The significance of incorporating force data is further emphasized by numerous studies employing diverse techniques to achieve this objective. In [24] for example, reaction forces from robot-mediated interaction are learned via kinaesthetic teaching through a haptic device, while in [45] a reinforcement learning algorithm aimed at precisely controlling forces complements the imitation learning of robot’s trajectories.

In our case, our framework feeds the force data directly collected from the bilateral control system to a transformer network able to imitate human behaviour in a displacement task. Initially popularized by [44] in the field of natural language processing, transformer models have been adapted for robotics as in the case of [14], which successfully

employed these models to predict and replicate complex human-robot interactions. Transformers have also demonstrated remarkable capabilities in predicting users' intent from trajectory data [21] and in anticipating human movements and forces, as shown in [15] where a GPT-2 Transformer \cite{radford2019language} was employed to predict up to 800 ms into the future.

3 A framework to replicate reaction forces in a robot-mediated displacement task

This section highlights the first contribution of our paper, that is: The creation of a framework leveraging on 4-channel bilateral control and transformer-based imitation learning to generate forces in human-robot interaction tasks that feel human-generated. It first describes the components separately and later dives into the application of the framework to learn the reaction forces exerted by users in a displacement task.

3.1 Framework components

3.1.1 Robots setup and bilateral control. Our setup includes two identical Franka Emika Panda robots (hereafter referred to as Franka) configured to operate side by side within a shared workspace. Each comes equipped with eight links and seven joints giving the robot seven degrees of freedom (7-DOF) therefore allowing the dynamics of the robots to closely mimic the dexterity and range of motion of a human arm. Each of these seven joints comes equipped with a highly sensitive torque sensor, allowing the robot to accurately measure and respond to external Cartesian forces. Ethernet cables and an offline router connect the robots to a dedicated control machine, ensuring stable, high-speed communication above the operational threshold of 1 kHz. The underlying control architecture of the system is built upon the C++ Franka Libraries. The Robot Operating System (ROS) [39] is used to facilitate the synchronization and data exchange between robots and two ROS nodes are used to call the control functions and exchange Cartesian Force Measurements. A moving average of five datapoints is used to reduce the noise affecting the external Cartesian forces applied to the robots. The underlying torque control algorithm is based on a simplified version of a 4-channel bilateral controller, with only the force sent bilaterally. The rest of the control structure is based on a purely damped PD controller. This structure enables a smooth bilateral exchange of force allowing for a torque-based control:

$$\tau_{k,1} = \mathbf{K}_p \left(\theta_{(k-1),1} - \theta_{k,1} \right) - \mathbf{K}_d \cdot \dot{\theta}_{k,1} + \mathbf{J}_1(\theta)^T \mathbf{F}_2 + \mathbf{g} \quad (1)$$

$$\tau_{k,2} = \mathbf{K}_p \left(\theta_{(k-1),2} - \theta_{k,2} \right) - \mathbf{K}_d \cdot \dot{\theta}_{k,2} + \mathbf{J}_2(\theta)^T \mathbf{F}_1 + \mathbf{g} \quad (2)$$

- $\tau_{k,i}$: is the joint control torque applied on each robot at time step k on the robot i
- $\theta_{k,i}$ and $\theta_{(k-1),i}$ are the current and previous joint positions of robot i , respectively
- $\mathbf{K}_p$ and $\mathbf{K}_d$ are the proportional and derivative gains, respectively
- $\mathbf{J}_i(\theta)$ is the end-effector Jacobian matrix for robot i
- $\mathbf{F}_1$ and $\mathbf{F}_2$ are the Cartesian force vector component applied at the end-effector of the first and second robot respectively
- $\mathbf{g}$ corresponds to the gravity compensation torques at the joint level

This control structure can then be used to specifically gather force data arising from an exchange between two humans during a displacement task.

3.1.2 Transformer architecture: In our work, the collected force data is used to train a Transformer model. Transformers [44] are deep learning models that have achieved significant success in handling sequential data such as time series.

Table 1. Proportional and Derivative Gains

Joint	K_p	K_d
1	80.0	8.0
2	60.0	6.0
3	80.0	8.0
4	100.0	8.0
5	100.0	8.0
6	100.0	8.0
7	50.0	6.0




Fig. 2. Block diagram of the implemented bilateral PD + bilateral force controller




Fig. 3. Transformer architecture. The multivariate input series is split into multiple univariate series, which are then fed into the Transformer backbone. Within the Transformer, these univariate series are divided into batches and normalized before entering the encoder, where multi-headed attention scores are computed.

Their popularity is partly due to the introduction of self-attention that captures dependencies between the elements in the sequence. This mechanism mitigates the vanishing gradient issue that affects traditional sequence models such as Recurrent Neural Networks (RNNs) [5].

The specific transformer used in our system is adapted from [34]. A model designed for time-series analysis that has consistently outperformed most state-of-the-art (SoTA) Transformers. This architecture, shown in figure 3, introduced two mechanisms into Transformers: 1) Patching, and 2) Channel-independence. This improves the network performance in numerical sequence analysis. Patching consists of aggregating multiple time steps into sub-series patches to capture local semantic information that might be lost if elements were analyzed independently. Channel-independence manages multivariate time series as multichannel signals, with each input token containing information from a single channel, while the embeddings and weights are shared across all channels.




Fig. 4. Point cloud representing the totality of the input data fed to the Transformer. Each point represents a force recorded in time across the x, y, and z axes

4 Applying the framework on a displacement task

We tested the framework described above by applying it to a robot-mediated human-to-human displacement task. This task involved 2 users, each interacting with a Franka robot: one *responder*, tasked with maintaining the initial position of their robot's end-effector, and one *disturber*, whose task was to perturb the responder's position. Our bilateral control allowed users to teleoperate with each other's robot, effectively providing them with haptic feedback. Thus, the responder could feel the disturber's pushes and pulls and react to them by counteracting the perceived forces to keep their end-effector in place. The disturber, on the other hand, attempted to move their robot to displace their partner through teleoperation and felt the reaction forces exerted by the responder. \textcolor{red}, noteworthy, in the context of telerehabilitation, the role of the *disturber* is akin to that of the patient, which attempts to move their limb while being restrained by the therapist (in the context of this task referred as *responder*) in an MRT task. However, the use of the terms patients and physiotherapist would be imprecise in this instance, as healthy, untrained subjects took part in the data collection and in the following study (see Limitations and future research directions below)

4.1 Data collection

The first step in our framework consisted of collecting task-related data to feed the transformer. As such, 2 male users (ages 39 and 23) performed the displacement task in a single session of the duration of 11 minutes. Data were collected at 1 kHz and included the forces exerted in the Cartesian space as shown in figure 4 and a rotation vector (roll, pitch, and yaw) at the end-effectors' level for both robots.

Notably, the collection of continuous interaction data for a sustained period of time allowed the transformer to exploit its self-attention mechanism to capture long-range dependencies, consequentially improving the accuracy of its prediction [44].

4.2 Transformer training

Data collected were first downsampled by averaging every 10 data points. The resulting dataset was then split into training (70%), validation (20%), and test data (10%). The transformer aimed at learning the responder's strategies, enabling the system to produce reaction forces in a human-like manner when disturbed. The model consists of a single

Table 2. Training parameters

Transformer parameters	Chosen values
Number of Encoder layers	1
Attention Heads	4
Model dimension	256
Dropout rate	0.05
Kernel size	6
Training epochs	200

encoder layer with four attention heads, a complete list of training parameters is provided in Table 2. Training took place across 200 epochs and the network took 30 minutes on average to train. Prediction time was on average 0.9 ms. Both training and testing were performed on a Nvidia Geforce RTX 4090 machine. Predicted force errors across the 3 axes averaged 1.81N, with the biggest prediction errors found on the y-axis (2.16N), followed by the x-axis (1.78N) and then by the z-axis, where we observed the least error (1.50N). Errors in rotation averaged (0.29N).

Figure 5 shows a comparison between test and predicted data on the x-axis within a 40-second interval.

5 Testing the human-likeness of the control algorithm

In this section we report the second contribution of our work and demonstrate that our framework can generate resistive forces that feel human-generated, seamlessly mimicking the feedback delivered by the responder in the displacement task

5.1 Conditions:

Our framework provides a way to generate human-like robot behaviour in robot-mediated human-to-human interactions. To test the extent to which the force feedback generated by our framework was considered human-like, we ran a modified Turing test where participants played the part of the "disturber" in the displacement task as shown in Figure 1. The role of the "responder" was taken on by either:

- (1) The experimenter (Human-Based control). The robot was operated by an examiner, in a classic robot-mediated human-to-human interaction setup using the developed bilateral force controller (1), (2).
- (2) The Transformer-Based Control. The robot reaction forces were generated autonomously using the transformer model. In this case, the torques are computed as

$$\tau_k = \mathbf{K}_p (\theta_{(k-1)} - \theta_k) - \mathbf{K}_d \dot{\theta}_k + \mathbf{J}_1(\theta)^T \mathbf{F}_T + \mathbf{g}, \quad (3)$$

where $\mathbf{F}_T$ is the Cartesian force vector outputted by the Transformer.

- (3) A proportional-derivative (PD) Control. The reaction forces to the disturber's attempt at displacement followed the formula below:

$$\tau_k = \mathbf{K}_p (\theta_d - \theta_k) - \mathbf{K}_d \dot{\theta}_k + \mathbf{g}. \quad (4)$$




Fig. 5. A direct comparison of human-generated and transformer-predicted forces. For simplicity and to avoid redundancy, the picture reports the forces in the x axes as an example.

where $\theta_d = \theta_{(0)}$ is the desired reference position, equal to the initial robot joint position in our case.

Due to the nature of the chosen task, the third controller was selected to act as a mass-spring damper under pure damping. Its architecture allows it to implicitly generate a pseudo-force profile that increases smoothly and proportionally to the positional error. To maintain a standardized baseline and ensure that controller behavior was not all gains were maintained to be same for all controllers. The addition of this third controller to the test served as a lower anchor. The comparison between Human-based and PD controllers ensured that participants could indeed discriminate between a human and a robotic responder. Each participant taking part in the experiment interacted with all three controllers according to the procedure described below. Participants were not informed of the number of controllers in the test.

5.2 Methods

Participants were invited to take part to the experiment on a volunteer basis and compensation was offered for their participation. 15 participants (13 males) were recruited among staff and students of the Computer Science department at -anonymous-. Of these, one participant had direct experience with working with Franka robots. Each participant was welcomed by the experimenter, introduced to the setup, and instructed about the necessary safety protocols. The experimental setup replicated the conditions used during data collection, with the Franka robots mounted on two




Fig. 6. Left: count of responses for each controller for the question "was the robot on the other side operated by a human?", pooled in data. Right: Confidence ratings for each controller and response to the Turing test, pooled in data.

separate tables and positioned side-by-side. A screen was mounted between the robotic arms, restricting the field of view (FOV) of each participant as shown in Fig. 7. Furthermore, a pair of headphones playing music was used during the trials to prevent participants from noticing any sounds or movements from across the screen. The test began with a short familiarization, where participants were invited to interact with the experimenter (human-based control). This provided them with a clear reference of what mediated human-to-human interaction felt like in our setup. The familiarization was followed by 15 trials (five for each controller) of 10 seconds each, presented in a fully randomised order. Each trial was initiated verbally and stopped via a visual signal (light indicator) by the experimenter. After each trial, users were asked 2 simple questions: "was the robot on the other side operated by a human?" and "To what degree of confidence do you think so?". This latter question was answered on a Likert scale from 1 (no degree) to 7 (maximum degree). At the end of the 15 questions, users were asked about their thought process in discriminating between robotic and human-generated responses.

5.3 Results

Participants' responses were pooled and analyzed using a Chi-Square Test for Homogeneity, comparing the frequency of instances where participants thought they were interacting with a human across all three conditions (Human-Based control, Transformer-based control, and fixed reference PD control). Analysis of pooled data showed that the fixed reference PD controller was identified as human with a significantly lower frequency (0.36) compared to both the human-based (0.64, p -value < 0.01) and the transformer-based controls (0.60, p -value < 0.01). No statistically significant difference was found when directly comparing the human-based and the transformer-based controller ($p = 0.74$). A more in-depth analysis was then conducted to account for participants' individual differences. Cochran's Q Test confirmed the results of the Chi-Square Test for Homogeneity ($p < 0.05$), and subsequent pairwise comparison using McNemar's Test confirmed the differences between the three conditions, with only the fixed reference PD controller being identified as human significantly less often ($p < 0.05$). Confidence data were analysed using a Cumulative Logistic Mixed Model to

account for the ordinal nature of the answer and the participants' differences, which were included in the model as random effects. Results showed no statistically significant effect of the controller on the confidence ratings. Moreover, no statistically significant effect on users' belief concerning the identity of the responder (human or robot) was found. Finally, users were asked, at the end of the test, what they based their answer on when deciding whether the responder was human or robot. 13/15 participants mentioned how human responses were characterized by a short delay in the reaction force. 11/15 participants also mentioned that the force profiles generated by what they believed were human responders were more variable, less predictable, and less precise.

5.4 Discussion

The results of our test show that the forces exerted by our system were identified as "generated by a human" significantly more often than those delivered by a classic PD controller. Results also showed no statistically significant difference in the perceived "human-likeness" between human-based control and our transformer-based control. While the lack of statistically significant difference is very promising, more focused tests at higher statistical power are needed to conclude without doubt that our system actions are indistinguishable from those of a human [19]. Despite that, it is clear that, at the very least, our system feels "more human" than a classic, not imitation-based controller. Based on users' feedback, the transformer-based control ability to replicate delays and variability typical of human responses was the key factor in driving its "human-likeness". This is in line with the evidence on "naturalness" in other technological systems. For example, in the realm of artificial voices (Text-To-Speech), the role of the variability of the signal in making the system perceived more human-like has been known for decades [22]. Interestingly, results also showed a number of trials where the human responder was perceived as robotic and somewhere the fixed reference PD controller was perceived as human. Average confidence in these trials was not lower than in any other trial ($p = 0.6$), suggesting that users were likely to identify human-like aspects in the fixed reference PD controller behaviour and robot-like aspects in human behaviour to motivate their responses with an acceptable degree of certainty. These trials, while the minority, constitute an important source of information that might help us understand what aspects of haptic feedback are interpreted as characteristics of human beings. Further data gathering and a focused investigation will likely shed light on the matter.

6 General Discussion

In this paper, we introduced a novel framework to replicate the forces exerted by a human-controlled robot in a displacement task. Our test shows that forces generated by our framework are considered human-generated at a frequency comparable to that of forces truly generated by a human.

To the best of our knowledge, this work is the first case of IL combining the direct recording of force data through a 4-channel bilateral control with the learning of reaction force dynamics via Transformer. In fact, These results reflect the richness and fidelity of data that can be obtained by recording interactions mediated by devices equipped with bilateral controllers as well as the flexibility and effectiveness of transformers in handling time series in this specific situation.

To our knowledge, this is also the first study leveraging IL to deliver robotic force feedback that feels human-generated. In fact, IL mostly concerns itself with achieving human-like performance [21] rather than human-like feeling. It is worth mentioning that human-like artificial force feedback has been achieved in the past [e.g.[40]]. However, in those instances, the problem was greatly simplified to force delivered by a 1DoF device and computed based on an analytical model, rather than learned. The relevance of our work for the HCI community is strictly related to the displacement task over




Fig. 7. User testing setup

which we evaluated our system. In fact, while generalizing the performance of the framework to any force feedback scenario and any robot-mediated interaction might be farfetched, the successful evidence in our displacement task allows us to speculate on the specific application of our research in the field of healthcare and robotic telerehabilitation. The displacement task adopted in the paper is strongly inspired by rehabilitation practices, and in particular by MRTs, where physiotherapists ask patients to exert forces that they then counteract to achieve isometric contraction of the users' muscles. Our framework enables robots to deliver resistive reaction forces to users that can serve the same purpose. Moreover, the resistive forces delivered by our framework feel human-generated. This specific aspect of our results is of high importance as it allows us to paint the vision of a future where AI-powered virtual physiotherapists, indistinguishable from human ones, can deliver telerehabilitation protocols to patients providing flexibility and control in the delivery of the treatment, while at the same time maintaining the characteristics of adaptability to patient's performance and patient's engagement which are the biggest advantages of human-supervised telerehabilitation.

7 Limitations and future challenges

Despite its potential and the promising results, our work is not without limitations.

While having a potential impact on the field of telerehabilitation, our system has not been tested with therapists or patients. While there is no reason to believe that the system will perform differently when replicating a professional application of MRTs, further specialized testing on a larger, gender balanced sample is needed to consolidate the

evidence reported here. Specifically on gender balance, it is worth reporting that our sample was primarily constituted by male subjects. While this is unlikely to have affected the results (as results in gender differences in kinesthetic and force perception report contradictory results [27], [12]), future studies will strive to achieve gender balance in testing. On a more technical note, it is also worth noting that our framework only works with Franka or similar devices capable of accurate torque control and force reading, while most telerehabilitation systems leverage on haptic or custom devices. However, it is worth noting that the Franka robot possesses 7DoF, and could be seen as a generalized device whose DoF could be constrained to simulate the behaviour of more traditional 3DoF and 6DoF haptic devices.

Another advantage of our hardware was its sensing capability, which allowed us to collect force data. Nonetheless, this framework could still be used on other robotic arms, even in the absence of dedicated force sensors or load cells, leveraging the results from [41], which introduces a method to compute the interaction force of sensorless robots. Future work will aim at consolidating the framework by applying it to a system composed of different devices.

Alongside expanding the framework integrating different devices, we will also aim at exploring more complex tasks integrating joint position, velocity, and Cartesian force profiles into the transformer's input. This would potentially allow the robot to exhibit more complex behaviours, for example, generating variable motion in Cartesian space while also exerting human-like force profiles. Furthermore, other sensory feedback (e.g., tactile) could be integrated to increase the human-likeness of the forces delivered by our framework. Finally, further avenues of research could include a further improvement of the transformer architecture. More complex transformers have been used in the past for the prediction of complex human motion [Kedia2024interact]. In the future, we plan to apply such models to force data to improve the capability of our system.

References

- [1] Brenna Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57 (05 2009), 469–483. <https://doi.org/10.1016/j.robot.2008.10.024>
- [2] Adriana Daniela Banyai and Cornel Brişan. 2024. Robotics in physical rehabilitation: Systematic Review. In *Healthcare*, Vol. 12. MDPI, 1720.
- [3] Cagatay Basdogan, Hao Ho, Mandayam Srinivasan, and Mel Slater. 2001. An Experimental Study on the Role of Touch in Shared Virtual Environments. *ACM Transactions on Computer-Human Interaction* 7 (07 2001). <https://doi.org/10.1145/365058.365082>
- [4] Georgeta Bauer and Ya-Jun Pan. 2020. Review of Control Methods for Upper Limb Telerehabilitation With Robotic Exoskeletons. *IEEE Access* 8 (2020), 203382–203397. <https://doi.org/10.1109/ACCESS.2020.3036596>
- [5] Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. *Advances in neural information processing systems* 28 (2015).
- [6] Janet Bettger and Linda Resnik. 2020. Telerehabilitation in the Age of COVID-19: An Opportunity for Learning Health System Research. *Physical therapy* 100 (08 2020). <https://doi.org/10.1093/ptj/pzaa151>
- [7] Scott Brave and Andrew Dahley. 1997. inTouch: a medium for haptic interpersonal communication (*CHI EA '97*). Association for Computing Machinery, New York, NY, USA, 363–364. <https://doi.org/10.1145/1120212.1120435>
- [8] Andrea Calderone, Desiree Latella, Mirjam Bonanno, Angelo Quartarone, Sepehr Mojdehdehbaheer, Antonio Celesti, and Rocco Salvatore Calabrò. 2024. Towards Transforming Neurorehabilitation: The Impact of Artificial Intelligence on Diagnosis and Treatment of Neurological Disorders. *Biomedicine* 12, 10 (2024), 2415.
- [9] Adrian David Cheok, Emma Yann Zhang, Adrian David Cheok, and Emma Yann Zhang. 2019. Kissenger: Transmitting kiss through the internet. *Human-Robot Intimate Relationships* (2019), 77–97.
- [10] Colleen O'Brien Cherry, Neale R Chumbler, Kimberly Richards, Amber Huff, David Wu, Laura M Tilghman, and Andrew Butler. 2017. Expanding stroke telerehabilitation services to rural veterans: a qualitative study on patient experiences using the robotic stroke therapy delivery and monitoring system program. *Disability and Rehabilitation: Assistive Technology* 12, 1 (2017), 21–27.
- [11] W Randolph Chitwood Jr, L Wiley Nifong, William HH Chapman, Jason E Felger, B Marcus Bailey, Tara Ballint, Kim G Mendleson, Victor B Kim, James A Young, and Robert A Albrecht. 2001. Robotic surgical training in an academic institution. *Annals of surgery* 234, 4 (2001), 475–486.
- [12] Min-Chi Chiu and Mao-Jiun Wang. 2007. The effect of gait speed and gender on perceived exertion, muscle activity, joint motion of lower extremity, ground reaction force and heart rate during normal walking. *Gait & posture* 25, 3 (2007), 385–392.
- [13] Seyed Ali Mohamad Dehghan, Hamid Reza Koofgar, and Mohsen Ekramian. 2016. Nonlinear bilateral control of 4-channel teleoperation systems using adaptive force estimator. In *2016 20th International Conference on System Theory, Control and Computing (ICSTCC)*. IEEE, 102–106.

- [14] Weiyang Deng, Barbara Sargent, Nina S. Bradley, Lauren Klein, Marcelo Rosales, José Carlos Pulido, Maja J. Matarić, and Beth A. Smith. 2021. Using Socially Assistive Robot Feedback to Reinforce Infant Leg Movement Acceleration. In *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*. 749–756. <https://doi.org/10.1109/RO-MAN50785.2021.9515537>
- [15] Alessia Fusco, Valerio Modugno, Dimitrios Kanoulas, Alessandro Rizzo, and Marco Cognetti. 2024. Transformer-Based Prediction of Human Motions and Contact Forces for Physical Human-Robot Interaction. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. 3161–3167. <https://doi.org/10.1109/ICRA57147.2024.10611211>
- [16] Guido Gioioso, Mostafa Mohammadi, Antonio Franchi, and Domenico Prattichizzo. 2015. A force-based bilateral teleoperation framework for aerial robots in contact with the environment. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 318–324.
- [17] PA Diluka Harischandra and AM Harsha S Abeykoon. 2019. Upper-limb tele-rehabilitation system with force sensorless dynamic gravity compensation. *International Journal of Social Robotics* 11 (2019), 621–630.
- [18] Peter F. Hokayem and Mark W. Spong. 2006. Bilateral teleoperation: An historical survey. *Automatica* 42, 12 (2006), 2035–2057. <https://doi.org/10.1016/j.automatica.2006.06.027>
- [19] Hui Jiang and Fuqing Gao. 2008. Large deviations and moderate deviations for the chi-square test in type II error. *Wuhan University Journal of Natural Sciences* 13, 2 (2008), 129–132.
- [20] Karam Dad Kallu, Jie Wang, Saad Jamshed Abbasi, and Min Cheol Lee. 2018. Estimated reaction force-based bilateral control between 3DOF master and hydraulic slave manipulators for dismantlement. *Electronics* 7, 10 (2018), 256.
- [21] Kushal Kedia, Atiksh Bhardwaj, Prithwish Dan, and Sanjiban Choudhury. 2024. Interact: Transformer models for human intent prediction conditioned on robot actions. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 621–628.
- [22] L. G. Kersta, Peter D. Bricker, and Edward E. David. 1960. Human or Machine?—a Study of Voice Naturalness. *Journal of the Acoustical Society of America* 32 (1960), 1502–1502. <https://api.semanticscholar.org/CorpusID:120470087>
- [23] Min Koh, Sheng-Che Yen, Lester Leung, Sarah Gans, Keri Sullivan, Yasaman Adibnia, Misha Pavel, and Christopher Hasson. 2021. Exploiting telerobotics for sensorimotor rehabilitation: a locomotor embodiment. *Journal of NeuroEngineering and Rehabilitation* 18 (04 2021). <https://doi.org/10.1186/s12984-021-00856-w>
- [24] Petar Kormushev, Sylvain Calinon, and Darwin G Caldwell. 2011. Imitation learning of positional and force skills demonstrated via kinesthetic teaching and haptic input. *Advanced Robotics* 25, 5 (2011), 581–603.
- [25] Peter Kudry and Michael Cohen. 2023. Enhanced Wearable Force-Feedback Mechanism for Free-Range Haptic Experience Extended by Pass-Through Mixed Reality. *Electronics* 12, 17 (2023). <https://doi.org/10.3390/electronics12173659>
- [26] D.A. Lawrence. 1993. Stability and transparency in bilateral teleoperation. *IEEE Transactions on Robotics and Automation* 9, 5 (1993), 624–637. <https://doi.org/10.1109/70.258054>
- [27] Yan-Xia Li, Lin Li, Zhi-Qiang Cai, Xu-Xue Zhou, Xiu-Jun Hao, and Li Li. 2022. Exploring sex differences and force level effects on grip force perception in healthy adults. *Motor control* 26, 2 (2022), 241–257.
- [28] Craig Liebenson and Curtis Thor Rigney. 2023. Manual resistance techniques in rehabilitation. *Chaitow's Muscle Energy Techniques E-Book: Chaitow's Muscle Energy Techniques E-Book* (2023), 273.
- [29] Jun Hong Lim, Kaibo He, Zeji Yi, Chen Hou, Chen Zhang, Yanan Sui, and Luming Li. 2023. Adaptive Learning based Upper-Limb Rehabilitation Training System with Collaborative Robot. In *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 1–5.
- [30] Albert C Lo, Peter D Guarino, Lorie G Richards, Jodie K Haselkorn, George F Wittenberg, Daniel G Federman, Robert J Ringer, Todd H Wagner, Hermano I Krebs, Bruce T Volpe, et al. 2010. Robot-assisted therapy for long-term upper-limb impairment after stroke. *New England Journal of Medicine* 362, 19 (2010), 1772–1783.
- [31] Xiaofeng Mao, Yucheng Xu, Ruoshi Wen, Mohammadreza Kasaei, Wanming Yu, Efi Psomopoulou, Nathan F Lepora, and Zhibin Li. 2023. Learning fine pinch-grasp skills using tactile sensing from real demonstration data. *arXiv preprint arXiv:2307.04619* (2023).
- [32] Maxi Miciak, Maria Mayan, Cary Brown, Anthony S Joyce, and Douglas P Gross. 2018. The necessary conditions of engagement for the therapeutic relationship in physiotherapy: an interpretive description study. *Archives of physiotherapy* 8 (2018), 1–12.
- [33] Hideyuki Nakanishi, Kazuaki Tanaka, and Yuya Wada. 2014. Remote handshaking: Touch enhances video-mediated social telepresence. *Conference on Human Factors in Computing Systems - Proceedings* (04 2014). <https://doi.org/10.1145/2556288.2557169>
- [34] Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *International Conference on Learning Representations*.
- [35] Carolina Passenberg, Angelika Peer, and Martin Buss. 2010. A survey of environment-, operator-, and task-adapted controllers for teleoperation systems. *Mechatronics* 20, 7 (2010), 787–801. <https://doi.org/10.1016/j.mechatronics.2010.04.005> Special Issue on Design and Control Methodologies in Telerobotics.
- [36] Rajni V. Patel, S. Farokh Atashzar, and Mahdi Tavakoli. 2022. Haptic Feedback and Force-Based Teleoperation in Surgical Robotics. *Proc. IEEE* 110, 7 (2022), 1012–1027. <https://doi.org/10.1109/JPROC.2022.3180052>
- [37] Alessandro Peretti, Francesco Amenta, Seyed Khosrow Tayebati, Giulio Nittari, Syed Sarosh Mahdi, et al. 2017. Telerehabilitation: review of the state-of-the-art and areas of application. *JMIR rehabilitation and assistive technologies* 4, 2 (2017), e7511.
- [38] Tommaso Proietti, Kristin Nuckols, Jesse Grupper, Diogo Scherz de Lucena, Bianca Inirio, Kelley Porazinski, Diana Wagner, Tazzy Cole, Christina Glover, Sarah Mendelowitz, et al. 2024. Combining soft robotics and telerehabilitation for improving motor function after stroke. *Wearable*

- Technologies* 5 (2024), e1.
- [39] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, Andrew Y Ng, et al. 2009. ROS: an open-source Robot Operating System. In *ICRA workshop on open source software*, Vol. 3. Kobe, Japan, 5.
- [40] Kyle B Reed and Michael A Peshkin. 2008. Physical collaboration of human-human and human-robot teams. *IEEE transactions on haptics* 1, 2 (2008), 108–120.
- [41] Loris Roveda and Dario Piga. 2020. Interaction Force Computation Exploiting Environment Stiffness Estimation for Sensorless Robot Applications. In *2020 IEEE International Workshop on Metrology for Industry 4.0 & IoT*. 360–363. <https://doi.org/10.1109/MetroInd4.0IoT48571.2020.9138189>
- [42] Mahya Shahbazi, Seyed Farokh Atashzar, and Rajni V. Patel. 2018. A Systematic Review of Multilateral Teleoperation Systems. *IEEE Transactions on Haptics* 11, 3 (2018), 338–356. <https://doi.org/10.1109/TOH.2018.2818134>
- [43] Bruno Siciliano, Lorenzo Sciavicco, Luigi Villani, and Giuseppe Oriolo. 2008. *Robotics: Modelling, Planning and Control* (1st ed.). Springer Publishing Company, Incorporated.
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [45] Yan Wang, Cristian C Beltran-Hernandez, Weiwei Wan, and Kensuke Harada. 2021. Robotic imitation of human assembly skills using hybrid trajectory and force learning. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 11278–11284.
- [46] Steve Yohanan and Karon E MacLean. 2008. The haptic creature project: Social human-robot interaction through affective touch. In *Proceedings of the AISB 2008 Symposium on the Reign of Catz & Dogs: The Second AISB Symposium on the Role of Virtual Creatures in a Computerised Society*, Vol. 1. Citeseer, 7–11.
- [47] Jianjun Zhang, Weidong Liu, Li'e Gao, and Haifeng Zheng. 2017. The force feedback instrument based on stepper motor drive and its application in underwater teleoperation. In *OCEANS 2017 - Aberdeen*. 1–7. <https://doi.org/10.1109/OCEANSE.2017.8084887>