ORIGINAL ARTICLE

Human-as-a-robot performance in mixed reality teleultrasound

David Black[1] · **Septimiu Salcudean**[1]

Abstract

Purpose In "human teleoperation" (HT), mixed reality (MR) and haptics are used to tightly couple an expert leader to a human follower [1]. To determine the feasibility of HT for teleultrasound, we quantify the ability of humans to track a position and/or force trajectory via MR cues. The human response time, precision, frequency response, and step response were characterized, and several rendering methods were compared.

Methods Volunteers (n=11) performed a series of tasks as the follower in our HT system. The tasks involved tracking pre-recorded series of motions and forces while pose and force were recorded. The volunteers then performed frequency response tests and filled out a questionnaire.

Results Following force and pose simultaneously was more difficult but did not lead to significant performance degradation versus following one at a time. On average, subjects tracked positions, orientations, and forces with RMS tracking errors of 6.2 ± 1.9 mm, $5.9 \pm 1.9°$, 1.0 ± 0.3 N, steady-state errors of 2.8 ± 2.1 mm, 0.26 ± 0.2 N, and lags of 345.5 ± 87.6 ms, respectively. Performance decreased with input frequency, depending on the input amplitude.

Conclusion Teleoperating a person through MR is a novel concept with many possible applications. However, it is unknown what performance is achievable and which approaches work best. This paper thus characterizes human tracking ability in MR HT for teleultrasound, which is important for designing future tightly coupled guidance and training systems using MR.

Keywords Teleoperation · Human–computer interaction · Ultrasound · Robotics · AR/VR

Introduction

Teleguidance is becoming more important in a wide range of fields, from remote maintenance and monitoring [2] to telemedicine [3]. The latter has grown in importance particularly in the last years due to the COVID-19 pandemic. One important and heavily studied form of telemedicine is teleultrasound (TUS). In the last 2 years alone, this has found applications in rural or under-resourced environments, for COVID-19 safety, and for training [4, 5]. To improve teleultrasound, a novel system dubbed "human teleoperation" (HT) was introduced based on mixed reality (MR) [1]. In brief, a novice person carries out an US procedure by following the motion and force of a virtual US probe projected onto the patient via a mixed reality headset. The virtual probe is controlled in real time by a remote expert who sees the US images and a video stream of the patient. This improves performance over video conference-based teleguidance by creating a tighter coupling between the leader and follower, yet is more flexible and accessible than robotic TUS.

MR has long been used to augment a person's sensory capabilities by overlaying medical images or volumetric models in situ for ultrasound-guided needle biopsies [6], robot-assisted and laparoscopic surgery [7], and more. Some implementations also include static guides illustrating where to align an instrument [8]. However, to our knowledge no other system has used dynamic hand-over-hand guidance to enable an expert to effectively teleoperate a novice person. This has traditionally been the task of robots, whose performance characteristics are precisely known, but the performance characteristics of the human follower in HT have never been studied.

Similar research has been carried out for the development of haptic devices. Of particular interest are the visual-motor response time (RT), frequency response of the hand, and accuracy in positioning and force exertion. Tan compiled

✉ David Black
  dgblack@ece.ubc.ca

  Septimiu Salcudean
  tims@ece.ubc.ca

1 Department of Electrical and Computer Engineering,
  University of British Columbia, Vancouver, BC, Canada

several measures including an approximate JND for force sensing of 7%, a JND for wrist angle of 2°, and a resolution for applied forces of 0.36 N [9]. The spatial resolution of hand movements is approximately 0.5–2.5 mm [10]. It was found that the somatosensory system can perceive vibrotactile stimuli up to 1 kHz, though the ability to apply forces is far more limited, with a force control bandwidth of approximately 20–30 Hz and closer to 7 Hz in practice [11]. Various studies have measured a visual-motor RT of 190–300 ms [12].

While these values constitute approximate limits in expected human performance, they do not describe the ability of a human to track an input signal, whether visual, haptic, or auditory. This ability to track has great implications for AR/VR interfaces and teleguidance systems, enabling technologies such as our HT system. This paper therefore characterizes human tracking performance, using mixed reality input to follow positions, orientations, and forces, both individually and simultaneously. We determine frequency responses, tracking bandwidth, root-mean-square (RMS) error, steady state error, reaction time, and rise time, among other important factors. Step response tests with the system are described in [13].

These results are important not only to determine the feasibility and expected effectiveness of our specific TUS system, but also more broadly to inform future mixed or virtual reality human–computer interfaces. The expected performance measures give designers a baseline in building new applications, including how fast a person will react or follow virtual cues, how precisely a person will move or apply forces, and what type of rendering the human brain reacts to best. The results also allow such a system to be modeled mathematically, which will be essential in designing teleoperation controllers.

Methods

The HT prototype used in this paper was previously introduced [1, 14]. HT for ultrasound (US) involves an expert sonographer at a medical center (the "expert"), and a novice "follower" in a remote location who carries out the US examination. The follower wears a mixed reality headset (Microsoft HoloLens 2) which projects a three-dimensional virtual US probe into the leader's visual field. This virtual probe is controlled in real time using a haptic device (3DSystems Touch X) by the expert, who sees the US images and video feed from the HoloLens in real time, and receives force feedback via the haptic device. The follower matches the position and orientation (pose) and force of the virtual probe as it moves according to the expert's input, thus achieving teleoperation.

The system logs all data during the tests and can play back prerecorded or generated motion and force sequences on the follower side. An overview is shown in Fig. 1. The follower was given a 3D printed object shaped exactly like a Clarius C3HD3 probe, but with a 6-axis force/torque sensor at the tip (ATI Nano25), and an electromagnetic pose sensor (driveBAY, Northern Digital) embedded in the hand grip, as far from the metal force sensor as possible. The pose and force of the follower were thus measured throughout the tests.

The sensors were connected to a PC running a .NET program which read in the data and sent it to the HoloLens via WiFi, over a WebRTC connection, which adds up to 5 ms delay to the measurements [13]. This delay was measured in real time, and the reported latencies in the results were adjusted accordingly. The electromagnetic force sensor readings were transformed to the HoloLens frame by placing two ArUco markers [15] in known positions on the magnetic transmitter, which defines the sensor frame. The HoloLens detected the pose of both markers (Mitchell Doughty GitHub[1] and found the optimal transform between the frames. The electromagnetic tracking is accurate to 1.4 mm and 0.5°, and the force sensor has 0.02–0.06 N resolution.

To test rendering schemes for the human–computer interaction, 4 renderings of the virtual probe were tested. For pose teleoperation, an US probe shape was tested, as was the same shape with the middle removed (Fig. 1), to test whether occlusion of the real probe by the hologram affects control accuracy. For forces, two schemes were tested: (1) changing the color of the virtual probe continuously from blue (too little force) to green (good force) to red (too much force) and (2) using an error-bar (EB) which grew away from the patient and turned blue (too little force), shrunk down and turned green (good force), and grew toward the patient and turned slowly red (too much force) (Fig. 1). The calculation involved subtracting the measured force from the desired one.

Volunteers ($n = 11$, aged 20–64 years, mean 32. 36% female) were recruited to perform a series of tasks. By design, the volunteers came from a variety of backgrounds including engineering graduate and undergraduate students, a medical student, a physicist, a surgeon, an engineering professor, a lawyer, and an architecture student. They were first given a demonstration of how to use the system and allowed to try it for 1 min. The tasks included following 2 sequences of prerecorded motions and 2 sequences of prerecorded forces (single parameter tracking), and 4 sequences in which motions (tangent to surface) and forces (normal to rigid surface) were tracked simultaneously (dual parameter tracking). Each sequence was 1-3 min long, and each subject carried out the exact same sequences. During each single parameter test, a different one of the four rendering schemes was used. During the dual parameter tests, all four combinations of the rendering schemes were used. The order of

---

[1] https://github.com/doughtmw/ArUcoDetectionHoloLens-Unity.

Fig. 1 System overview: **a** electromagnetic tracker (on device and separate, with thumb tack for scale); **b** dummy US probe; **c** force sensor; **d** electromagnetic sensing transmitter with ArUco markers for registration; **e** follower with HoloLens 2; **f** follower side user interface; **g** force rendering schemes (color and error-bar) with full and partial probe

rendering schemes was randomized to avoid learning effects or bias due to the prerecorded motions being slightly different from each other.

Next, the subjects performed a frequency response test ("Frequency response" section) for both pose and force. An equally relevant step response test is described in [13]. The frequency response consisted of a sinusoidal input signal which increased in frequency every 5 oscillations until the follower was unable to track it.

All desired and measured force, position, and orientation data were saved with timestamps at a sampling rate of 36 Hz on the HoloLens to avoid clock drift between devices and were analyzed using MATLAB. Desired and measured values were aligned using the timestamps and resampled so each sample lined up in time. Resampling was performed by linear interpolation for position and force data, and spherical interpolation for quaternions (MATLAB slerp.m). The desired and measured signals were then subtracted from each other element-wise to determine errors. Tracking lags were calculated by finding the time delay that maximized the normalized cross-correlation between the measured and input signal.

In addition to the numerical data, the volunteers filled out a questionnaire. This aimed to determine how the users perceived the tasks in terms of physical and cognitive difficulty, and which rendering schemes they preferred. This could then be compared to the numerical results and gives useful insight into the challenge of HT from the follower's perspective. Statistical significance was measured using the two-sample Kolmogorov–Smirnov test.

Results

Tracking tests

Every volunteer followed a series of prerecorded motions and forces for the first eight tests. Example position and orientation tracking results are shown in Fig. 2a, b for an average user. Example force tracking is shown in Fig. 2c, d. Numerical tracking results are outlined in Table 1.

There is no statistically significant difference between the full and partial rendering for pose tracking accuracy or lag, though the mean lag is slightly better with full probe, while the mean tracking accuracy is better with partial probe. This indicates that the full probe does cause some accuracy problems by occluding the real probe, which several users commented on. This is very dependent on the ambient light; in a better-lit environment, the occlusion was much less. Additionally, the user interface includes an opacity setting in which the user can adjust the alpha level of the probe color. From the questionnaire, testers consistently had no preference between partial and full.

On the other hand, for force tracking, the EB is clearly superior, which can be seen in Fig. 2. On the questionnaire, all users indicated a strong preference for the EB, with all users commenting that it was much easier to see when the force error was zero by seeing the EB disappear as opposed to trying to see if the probe color was truly green, or had a slight blue or red tinge. Additionally, the tracking lag is significantly better with EB than with color force.

The next four tests involved tracking force and pose simultaneously. The overall results are found in Table 1. Despite the large differences in the single parameter tracking, none
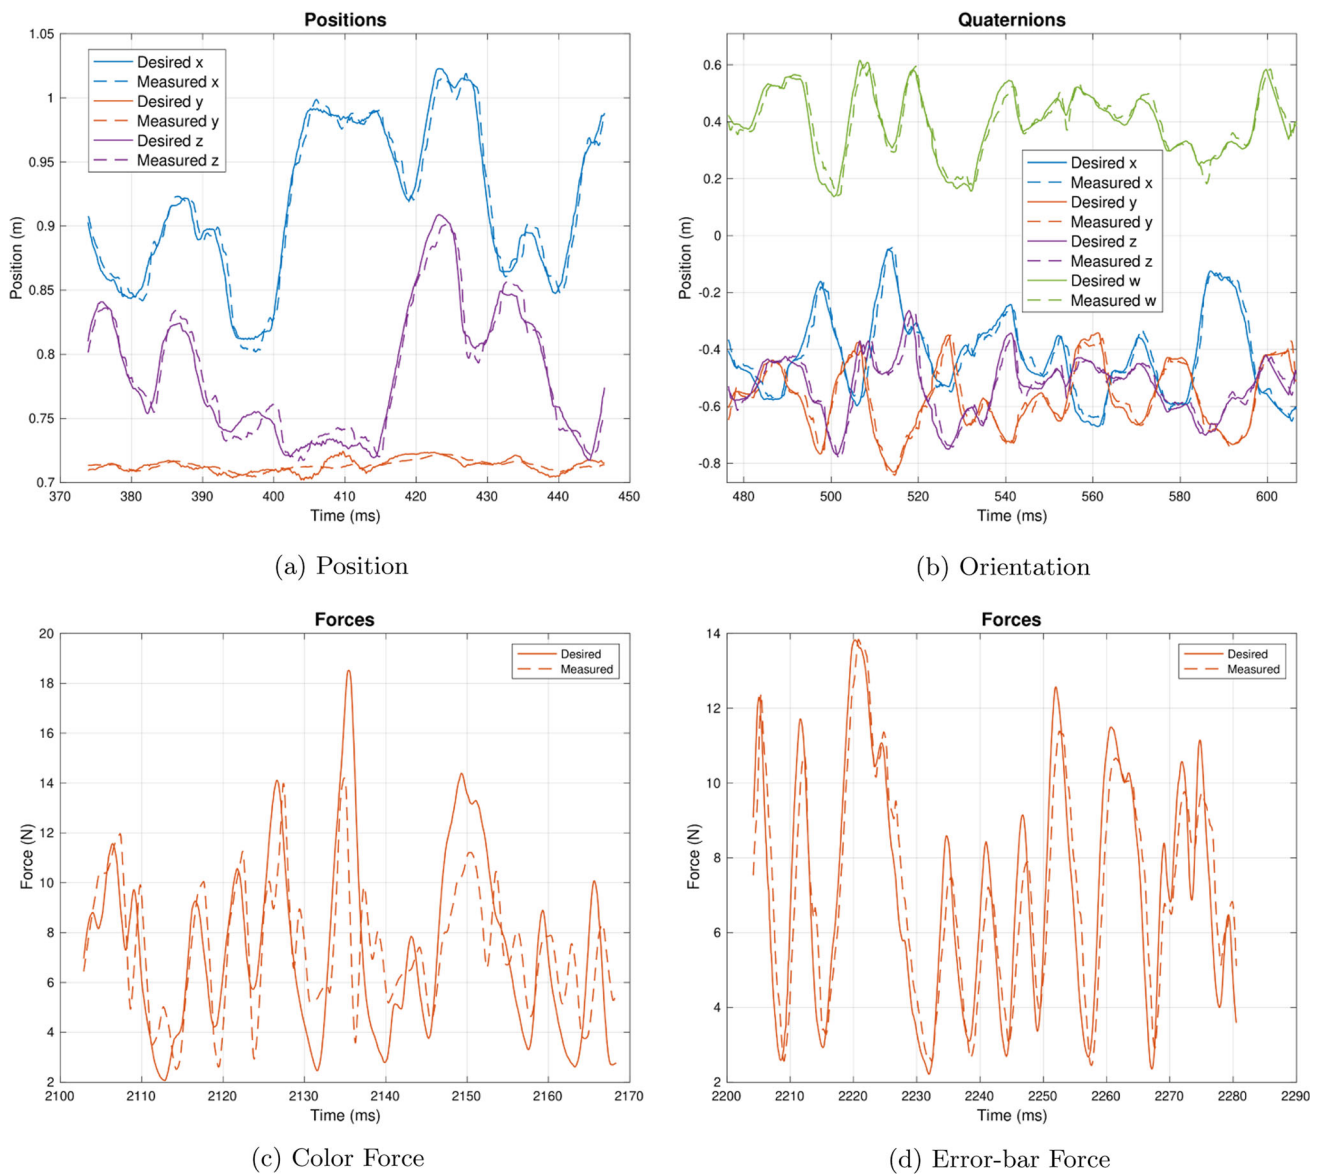
(a) Position

(b) Orientation

(c) Color Force

(d) Error-bar Force

Fig. 2 Example motion tracking test **a, b** and force tracking test **c, d** from one user whose results were approximately average. Solid lines are the input signal, while dotted lines are the measured motion or force.

The force tracking is plotted for two different rendering schemes, showing clear performance improvement using the error-bar

of the differences between the four rendering methods is significant in dual parameter tracking, except that Full+EB is significantly faster than Partial+EB ($p = 0.014$) and is the fastest on average. It is possible that with the Partial+EB rendering, the follower is presented with too much information, which is good for accuracy but impacts the tracking lag.

The performance differences between single and dual parameter tracking are in Table 2, showing some performance detriment from tracking both parameters at once, though only rotation error is statistically significant. Looking at the individual rendering schemes, EB alone was significantly faster than Full+Color ($p = 0.045$), Partial+Color ($p = 0.050$),

and Partial+EB ($p < 0.001$), but not Full+EB. Similarly, dual parameter renderings involving color force were significantly slower and less accurate than all single parameter tracking modes ($p = 0.01$ to 0.02). Conversely, Full+EB performs about as well as tracking just force. Hence, with Full+EB tracking, accuracy and speed are not significantly worse than single parameter tracking.

All participants stated in the questionnaire that dual parameter tracking constituted a greater cognitive load than single parameter tracking (average score $(4.36 \pm 0.81)/5$ where $5 =$ much harder and $1 =$ much easier). Additionally, force tracking was more mentally demanding (($3.36 \pm$

Table 1 Single and dual parameter tracking results (mean ± std. deviation)

Single param	Full	Partial	Color	Error-bar
Lag (ms)	346.23 ± 118.15	358.91 ± 57.53	469 ± 107.10	255 ± 118.88
p-value	0.76		< 0.001	
RMS error	8.1 ± 1.7 mm 7.7 ± 2.5°	6.24 ± 1.93 mm 5.93 ± 1.85°	1.69 ± 0.43 N	0.99 ± 0.29 N
p-value	0.38		0.001	
Dual param	**Full+Color**	**Full+EB**	**Partial+Color**	**Partial+EB**
Lag (ms)	408.72 ± 175.36	345.5 ± 87.60	435.28 ± 219.90	478.13 ± 126.46
RMS error	8.7 ± 1.6 mm 7.00 ± 2.31° 1.46 ± 0.33N	8.5 ± 1.4 mm 7.27 ± 2.25° 1.25 ± 0.33N	8.9 ± 1.9 mm 8.11 ± 2.65° 1.40 ± 0.23 N	8.6 ± 2.5 mm 7.48 ± 1.97° 1.26 ± 0.20 N

0.50)/5) than pose tracking $(2.18 \pm 0.60)/5$). The system scored $(2.36 \pm 0.81)/5$ for physical demand (where 5 = very demanding), and no participant became dizzy.

Frequency response

The frequency responses are plotted in Fig. 3, with a fitted average response and 95% confidence interval. The fitted curves are second degree polynomials, which were found to give the best fit. We see that in tracking position, some users managed to get to 2Hz, although with large phase lag. For phase lag less than 180°, users could follow up to 1 Hz. In this range, the phase delay trend was uniform among users and the gain was > −3 dB. Tracking forces was substantially more difficult despite not actually having to move. This is likely due to the less directly intuitive visual force control. While phase delay remained fairly small, it quickly became impossible to follow signals faster than about 0.35 Hz. There was a strong trend among all participants where phase angle initially decreased slightly with frequency. This is possibly because users were initially careful to be as accurate as possible, but quickly switched their focus to following quickly enough. It is unlikely that the users learned and improved during the low frequency, since these tests were carried out last. Force tracking gain decreased rapidly with frequency as users failed to match the desired force quickly enough before it changed again. Relatively good force performance could be achieved at less than about 0.25 Hz.

We next look at the dependency of the frequency response on the input signal amplitude. Stocco found that the human hand force bandwidth was 7 Hz for small motions but dropped past 5 Hz for larger motions [16]. From preliminary tests with a subset of the volunteers, we obtained the responses in Fig. 4. We see that the users were able to continue following positions for longer as the motion amplitude decreased, but that the trend in gain and phase was about the same, irrespective of amplitude. This implies that the users were physically able to make fast enough motions (as shown in [16]), but that their ability to follow was cognitively or visually limited as the virtual probe began to move too fast. With force tracking, we see the opposite trend because no motion is required. Instead, as the desired force differences became very small and approached the human hand's JND, it became hard to match them precisely, so the smaller forces were harder to follow.

Closely related tests of step response were carried out with the same volunteers and measurement system and are presented in [13]. Due to space constraints, they could not be included here, but human response time was found to be 628.3 ± 102.3 ms and 171.5 ± 85.9 ms, with steady state RMS error of 2.8 ± 2.1 mm and 0.26 ± 0.16 N for position and force, respectively.

Discussion

In this paper, we have introduced a measurement system based on the HT concept [1] and used it to perform a number of measurements of human performance in mixed reality pose and force teleoperation. Additionally, the human response to different rendering schemes was compared.

It was found that the error-bar force rendering was superior to color-based force, and there was little difference between full and partial probe rendering, but that the opacity of the probe is a matter of personal preference and should thus continue to be adjustable. With full probe + error-bar rendering, tracking both pose and force simultaneously did not lead to a statistically significant decrease in performance, though it was more mentally demanding. This shows that high performing HT, in which poses and forces are tracked simultaneously, accurately, and quickly, is feasible. For US procedures, motions are generally slow and do not coincide temporally with large force variations, so the requirements for simultaneous pose and force tracking are much less strict

Table 2 Overall single versus dual parameter tracking

	Single	Dual	*p*-value
Position error (mm)	7.9 ± 2.1	8.6 ± 1.9	0.16
Rotation error (°)	6.08 ± 1.84	7.44 ± 2.18	0.02
Force error (N)	1.32 ± 0.28	1.34 ± 0.50	0.81
Lag (ms)	362.29 ± 125.73	431.81 ± 183.93	0.062

(a) Force (b) Position
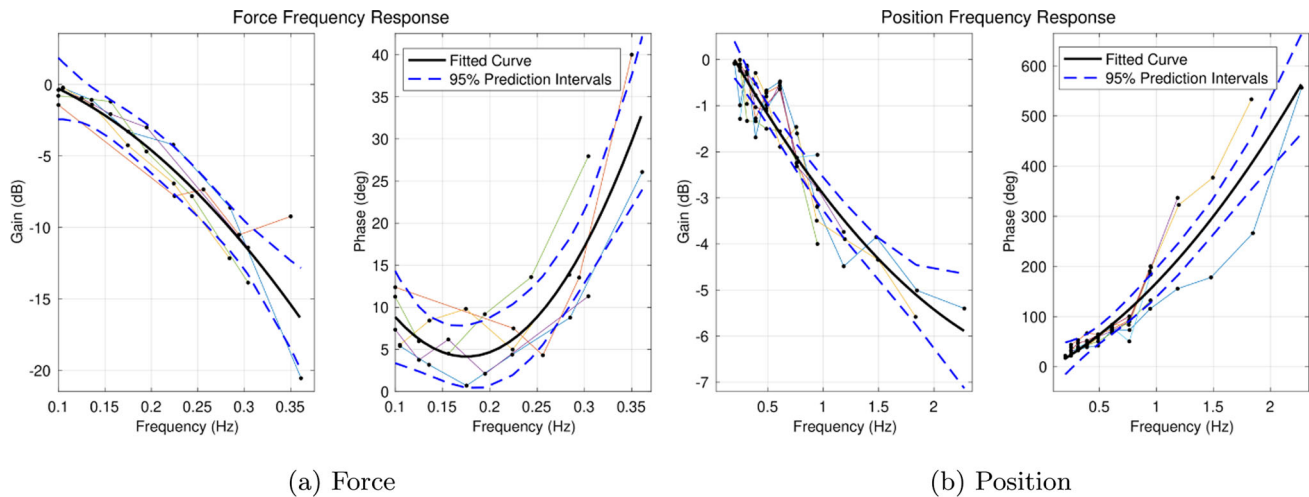
Fig. 3 Frequency responses. The black line is the fitted curve, and the dotted line is 95% CI. Phase refers to phase lag, where the follower lags the expert

than those considered in this paper. However, for other applications of HT, requirements may be more stringent. While this study shows the potential for good performance in HT, it does not assess the feasibility in real-world scenarios including sonographers and patients. Future work will include a clinical human study to build off the preliminary work in [1].

The visual force renderings tested here are not the only possibilities. For example, a second virtual probe could be rendered, offset from the primary one in a direction and by an offset proportional to the force error. Alternatively, an arrow could be used instead of the error-bar. Both of these methods include directional information which was lacking in the renderings tested here. However, forces in US are generally normal to the surface of the patient. Thus, single-axis force rendering may be sufficient or even preferable as excess information may slow down the follower, as seen in the Partial+EB rendering in "Tracking Tests" Section. Having a second virtual probe could also lead to confusion. This was the motivation for the two tested force rendering schemes. They are relatively simple and allow the user to keep their eyes on the US probe.

A third alternative which will be explored in future work is to offset the virtual probe itself from the desired position by an amount proportional to the force error and an estimated impedance of the patient tissue. In this way, by matching the rendered pose, the follower would automatically match the desired force as well, while effectively only tracking a single parameter at a time. However, for relatively stiff tissue, this would rely on the follower tracking very small differences in position. As shown in [13], the position resolution is approximately 3 mm. Finally, it is possible that using a color map other than red/green/blue would lead to better tracking, depending on the preferential sensitivity of the human eye to certain changes in color [17]. Conversely, it has been shown that the human auditory system has a bandwidth of 20 kHz compared to 50 Hz for visual perception [11]. Thus, force feed-forward could potentially also be achieved through auditory signals.

Several frequency response measures are outlined in "Introduction" section which can be compared to the values found in this study. The bandwidth of the human hand was found to be approximately 5–7 Hz [11, 16]. While this measures how fast a human hand can move or control forces, our study explored the ability of people to follow a desired signal, which is likely limited by the visual and cognitive systems rather than the hand. Indeed, for decent performance, the volunteers tracked up to 1 Hz input signals. This is approximately in the same order of magnitude as the maximum hand capabilities.

Several users commented that they were becoming fatigued by the end of the tests, after prolonged intense focus and visual stimulation. It is therefore likely that a well-rested
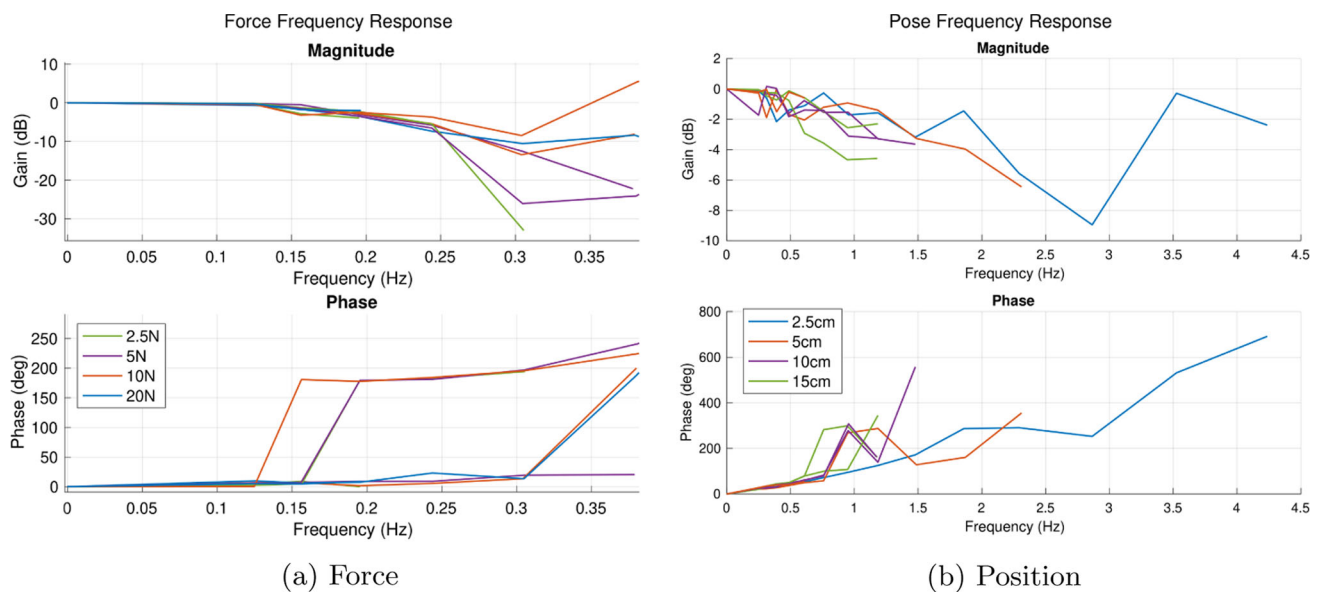
Fig. 4 Frequency responses at different input signal amplitudes. For position, smaller signal amplitudes are generally easier to follow, while the trend is less clear for force

group of volunteers performing the frequency and step responses again would perform better, although the difference would likely be relatively small. This is also one reason why we did not test further rendering schemes. Additionally, it has been shown that this type of behavior is a learned skill [18], so with training we might also expect improved results.

The described tests were performed on a TUS system. Correct pose and force are important in US to obtain good quality images and are very difficult for a novice user to achieve without tightly-coupled expert guidance. While pose ensures the correct anatomy is visualized, force is equally important. The sonographer must avoid deformation of some anatomies or push down hard under the ribs to view others. The force control thus requires accurate feedback of the forces or replication of the patient on the expert side, which will be the topic of future work and is discussed in [1]. This paper focuses on the human MR tracking ability in TUS, but the only US-specific factor in the presented tests is the shape of the virtual probe. Additionally, most motion and force sequences were faster and more difficult to follow than typical US exams. Thus, the results are expected to generalize well to other applications of HT. Indeed, there are many other applications where such teleoperation could be useful, including remote maintenance, inspection, and teaching.

Conclusion

In this paper, we have measured human performance in tracking a desired pose and force sequence through an MR interface. While tracking was found to be strongly dependent

on the rendering scheme, humans can track pose and force simultaneously with a lag of 345.5 ± 87.6 ms, and RMS tracking errors of 8.5 ± 1.4 mm, $7.27 \pm 2.25°$, and 1.25 ± 0.33 N. Steady state errors are significantly better, at 2.8 ± 2.1 mm and 0.26 ± 0.16 N. Tracking of signals with good performance is possible up to a bandwidth of about 1 Hz for position and 0.25 Hz for forces, both of which depend on the magnitude of the input signal. Rendering the full ultrasound probe with error-bar-based force feedback was most effective, and tracking both force and pose simultaneously was cognitively more difficult but did not lead to statistically significant degradation in performance. Ultimately, these values and results can serve as a guide for the design of future MR (or AR/VR) interfaces and demonstrate that HT for teleultrasound and likely other applications can provide good performance.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose beyond the funding sources listed above.

Ethical approval This study was performed in line with the principles of the Declaration of Helsinki and approval of the University of British Columbia Behavioural Research Ethics Board (Approval No: H22-01195; June 20, 2022).

Informed consent Informed consent was obtained from all individual participants in the study.

References

1. Black D, Yazdi YO, Hosseinabadi AHH, Salcudean S (2022) Human teleoperation–a haptically enabled mixed reality system for teleultrasound. Hum Comput Interact 24:1521

2. Yamamoto T, Otsuki M, Kuzuoka H, Suzuki Y (2018) Tele-guidance system to support anticipation during communication. Multimod Technol Interact 2(3):55

3. Gajarawala SN, Pelkowski JN (2021) Telehealth benefits and barriers. J Nurse Practitioners 17(2):218–221

4. Uschnig C, Recker F, Blaivas M, Dong Y, Dietrich CF (2022) Tele-ultrasound in the era of Covid-19: A practical guide. Ultrasound Med Biol 193:1253

5. Salerno A, Tupchong K, Verceles AC, McCurdy MT (2020) Point-of-care teleultrasound: a systematic review. Telemed e-Health 26(11):1314–1321

6. Rosenthal M, State A, Lee J, Hirota G, Ackerman J, Keller K, Pisano ED, Jiroutek M, Muller K, Fuchs H (2002) Augmented reality guidance for needle biopsies: an initial randomized, controlled trial in phantoms. Med Image Anal 6(3):313–320

7. Navab N, Traub J, Sielhorst T, Feuerstein M, Bichlmeier C (2007) Action-and workflow-driven augmented reality for computer-aided medical procedures. IEEE Comput Graph Appl 27(5):10–14

8. Chen L, Day TW, Tang W, John NW (2017) Recent developments and future challenges in medical mixed reality. In: 2017 IEEE international symposium on mixed and augmented reality (ISMAR), pp 123–135, IEEE

9. Tan HZ, Srinivasan MA, Eberman B, Cheng B (1994) Human factors for the design of force-reflecting haptic interfaces. Dyn Syst Control 55(1):353–359

10. Kaber DB, Zhang T (2011) Human factors in virtual reality system design for mobility and haptic task performance. Rev Human Factors Ergonom 7(1):323–366

11. Brooks TL (1990) Telerobotic response requirements. In: 1990 IEEE International conference on systems, man, and cybernetics conference proceedings, pp 113–120, IEEE

12. Thorpe S, Fize D, Marlot C (1996) Speed of processing in the human visual system. Nature 381(6582):520–522

13. Black D, Salcudean S (2022) A mixed reality system for human teleoperation in tele-ultrasound: communication performance. TechRxiv

14. Black D, Salcudean S (2022) A mixed reality system for human teleoperation in tele-ultrasound. In: Hamlyn symposium for medical robotics, pp 91–92

15. Garrido-Jurado S, Muñoz-Salinas R, Madrid-Cuevas FJ, Marín-Jiménez MJ (2014) Automatic generation and detection of highly reliable fiducial markers under occlusion. Pattern Recognit 47(6):42554

16. Stocco L, Salcudean SE (1996) A coarse-fine approach to force-reflecting hand controller design. In: Proceedings of IEEE international conference on robotics and automation, vol 1, pp 404–410. IEEE

17. Jameson D, Hurvich LM (1964) Theory of brightness and color contrast in human vision. Vision Res 4(1–2):135–154

18. Roenker DL, Cissell GM, Ball KK, Wadley VG, Edwards JD (2003) Speed-of-processing and driving simulator training result in improved driving performance. Hum Factors 45(2):218–233