Bimanual Ultrasound Mid-Air Haptics for Virtual Reality Manipulation

Lendy Mulot, Thomas Howard, Guillaume Gicquel, Claudio Pacchierotti, Maud Marchal

Bimanual Ultrasound Mid-Air Haptics for Virtual Reality Manipulation

Lendy Mulot†, Thomas Howard†, Guillaume Gicquel, Claudio Pacchierotti and Maud Marchal

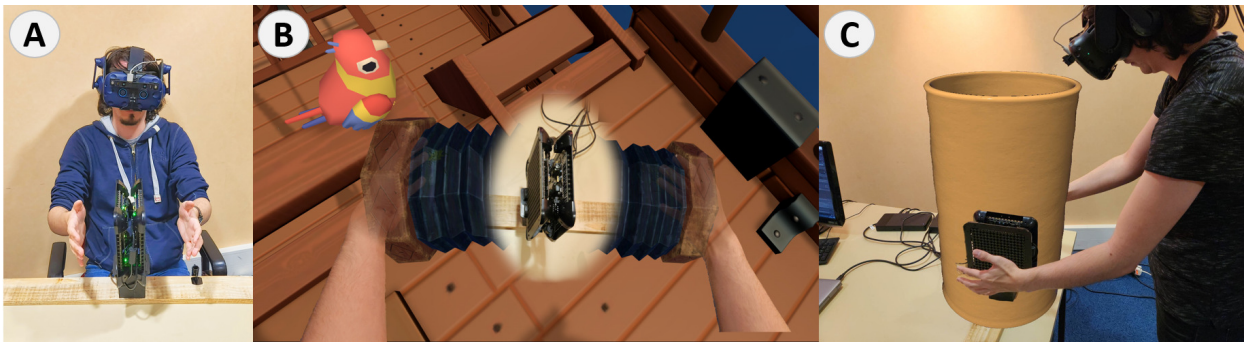

Fig. 1: (A) We propose a novel approach to bimanual ultrasound mid-air haptics in VR using two ultrasound devices mounted back-to-back. This enables bimanual haptic feedback during bimanual virtual object manipulation tasks, such as (B) the virtual concertina playing interaction or (C) the virtual pottery scenarios pictured here.

Abstract— The ability to manipulate and physically feel virtual objects without any real object being present and without equipping the user has been a long-standing goal in virtual reality (VR). Emerging ultrasound mid-air haptics (UMH) technology could potentially address this challenge, as it enables remote tactile stimulation of unequipped users. However, to date, UMH has received limited attention in the field of haptic exploration and manipulation in virtual environments. Existing work has primarily focused on interactions requiring a single hand and thus the delivery of unimanual haptic feedback. Despite being fundamental to a large part of haptic interactions with our environments, bimanual tasks have rarely been studied in the field of UMH interaction in VR. In this paper, we propose the use of non-coplanar mid-air haptic devices for providing simultaneous tactile feedback to both hands during bimanual VR manipulation. We discuss coupling schemes and haptic rendering algorithms for providing bimanual haptic feedback in bimanual interactions with virtual environments. We then present two human participant studies, assessing the benefits of bimanual ultrasound haptic feedback in a two-handed grasping and holding task and in a shape exploration task. Results suggest that the use of multiple non-coplanar UMH devices could be an interesting approach for enriching unencumbered haptic manipulation in virtual environments.

Index Terms—Mid-Air Haptics, Ultrasound, VR, Bimanual

1 Introduction

Ultrasound mid-air haptic (UMH) devices can enrich haptic experiences of virtual environments (VEs) without compromising user comfort, as they provide the opportunity to deliver a variety of freely configurable haptic sensations without physically equipping or constraining users [4, 16, 30]. These devices use arrays of ultrasonic speakers whose phase shift can be individually controlled, allowing steering and focusing of the ultrasound, to deliver tactile stimuli in 3D space. When coupled with appropriate modulation methods, complex vibrotactile stimuli can be rendered on a user's skin at a distance [30].

Many everyday interactions with real environments, such as picking up and manipulating large objects, are performed with two hands. Because of this, bimanual haptic feedback has also been shown to be beneficial in interaction with VEs [34]. However, enabling bimanual haptic manipulation with UMH devices present certain challenges. To optimally deliver ultrasound haptic stimuli, these devices require the user's hands to remain mostly parallel to the array plane. Furthermore, UMH devices can deliver stimuli at relatively large distances, but have a restricted lateral workspace [15]. Finally, a more general problem that needs to be tackled is that of manipulating a single object with both hands simultaneously, as this requires algorithmic and hardware solutions for reflecting the consequences of the actions of one hand onto the other, mediated by the manipulated virtual object. Consequently, applications of ultrasound mid-air haptics to virtual reality (VR) have usually focused on unimanual interactions [1, 3, 7, 9, 15, 17, 18, 23, 32]. Bimanual haptics in general, and bimanual ultrasound mid-air haptics in particular, remain underexplored in the field of VR manipulation.

A few works have looked into bimanual ultrasound haptic interactions with VEs, but they are limited to either alternating the use of hands [10] or simultaneous use of both hands [24] in a coplanar configuration, avoiding the issues mentioned above. Also, to the best of our knowledge, none of these works explicitly evaluates the potential benefit of bimanual haptics over the state-of-the-art unimanual haptics in such interactions.

We propose a novel ultrasound mid-air haptic feedback system and its associated coupling schemes, enabling grasping and simultaneous manipulation of a virtual object with both hands (see Figure 1). In this paper, we address the question of whether, in VR, bimanual haptic feedback using a dedicated UMH device for each hand provides a benefit over state-of-the-art unimanual haptic feedback. We investigate this question through two tasks: a bimanual 3D object grasping and manipulation task, and a bimanual 3D shape exploration task.

Our contributions are:

- A novel system and its associated haptic coupling schemes, which allow multi-device bimanual UMH feedback during bimanual object manipulation in VR (section 3).

- *L. Mulot, T. Howard, G. Gicquel and M. Marchal are with Univ Rennes, INSA, IRISA, France.*
- *C. Pacchierotti is with CNRS, Univ Rennes, Inria, IRISA, France.*
- *M. Marchal is also with IUF, France.*
- *E-mail:* `{firstname.lastname}@irisa.fr`
- †`Both authors contributed equally to this work.`

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

- A user study evaluating the benefits of bimanual UMH feedback in a bimanual virtual object grasping and holding task (section 4).
- A user study investigating the advantages of bimanual UMH feedback for the perception of virtual object shapes, where the feedback was designed based on a state-of-the-art 3D shape rendering algorithm [25], that we customized to bimanual shape rendering (section 5).
- Three use-cases illustrating bimanual UMH for different manipulation tasks in VR environments. (section 7).

2 Related Work

Virtual object manipulation, which consists in the manipulation of an object's properties (e.g., its position) by the user, forms one of the four fundamental classes of interactions with 3D user interfaces [19]. In the present work, we focus on direct manipulation, i.e. manipulation effected by a user's virtual hand coming into contact with a virtual object, and more specifically, on direct manipulation with haptic feedback [28], which we hereinafter refer to as *"haptic manipulation"*. Contrary to direct manipulation of objects in the real world, where actions directly induce visual, haptic, and other sensory feedback, in VR manipulation, perception and action can be decoupled [22]. Manipulation can occur without haptic feedback [35], with full haptic feedback (i.e. any part of the body acting on a virtual object receives haptic feedback) [5], or anything in between. For example, object manipulation with one hand can cause unimanual haptic feedback on the other hand [6].

Bimanual haptic technologies have been developed with different technologies, such as wearable gloves embedded with vibrotactile actuators (e.g. Senso Glove[1]), force-feedback gloves (e.g. SenseGlove[2]), or grounded encountered type devices [31]. While they are able to apply greater forces, none of these has the benefits of mid-air feedback. In ultrasound mid-air haptic manipulation in VR, unimanual tasks and, thus unimanual haptic feedback, have been the focus of attention [1,3,7,9,15,17,18,23,32]. In this case, unimanual haptic feedback has been shown to be beneficial to VR manipulation, improving task performance [9] and effectiveness of VR-based applications [18], improving user immersion [1,32] and enjoyment [9,17,32], and increasing the perceived realism of the virtual environment [17].

Yet, bimanual interaction with objects can provide many advantages over unimanual interaction. Talvas *et al.* [34] provide a detailed discussion of these in the context of interactions with virtual environments. First, a wider range of interactions is possible as two hands allow exploration, grasping and manipulation of larger objects. Second, the possibility of bimanual task execution can enable better task integration for complex tasks [13], and often better harnesses existing skills developed in everyday life [11]. If the tasks performed by both hands are sufficiently integrated, bimanual action can ensure faster and more accurate task execution [34]. Finally, under certain circumstances, bimanual haptic perception outperforms unimanual haptic perception for a number of reasons, including improved estimation of sizes and distances thanks to proprioception [11], specialized stereotypical bimanual haptic exploration strategies [20], and sensory integration of complementary and redundant stimuli from both hands [29]. Generally speaking, many tasks can thus benefit from the possibility of bimanual interaction to better resemble real-world haptic interactions, increase user's speed and accuracy in task execution, and reduce the task's cognitive load [34]. These literature findings motivate our choices of hypotheses for the experiments presented in section 4 and section 5.

Bimanual manipulation in VR with UMH feedback has received limited attention in the literature. A few application prototypes make use of both hands in interactions which either alternate between hands [10] or use both hands simultaneously in a co-planar configuration [24], such that both hands can be stimulated by a single array. However, these co-planar hand configurations are incompatible with many natural bimanual haptic interactions such as grasping, deforming and moving objects, or exploring their 3D shape [20]. This motivates our present study into the applicability of multi-device multi-directional UMH feedback to bimanual VR manipulation.

Our use of multiple non-coplanar UMH devices is somewhat inspired by prior work using similar configurations for augmented reality [26], although these proofs-of-concepts have been mainly used to explore rendering algorithms in unimanual interactions without constraining hand orientation. To the best of our knowledge, no work has yet explored the idea of using a pair of devices back to back, to allow bimanual haptic interaction with opposing hands. Our present work thus focuses on this possibility as a first step in evaluating the usefulness of bimanual UMH feedback in VR.

While it appears sensible that bimanual manipulation in VR would entail bimanual haptic feedback as it does in real environments, the previously discussed decoupling of perception and action in VR poses two questions when considering haptic manipulations with UMH feedback: **(1)** Could unimanual haptic feedback also benefit bimanual manipulation?; **(2)** Does bimanual haptic feedback, as we propose it, provide any advantage over unimanual haptic feedback in bimanual tasks? Therefore, this paper comparatively investigates bimanual VR manipulations without haptic feedback, with unimanual UMH feedback, and with bimanual UMH feedback.

3 Bimanual UHM Manipulation System

We propose a novel system using a pair of UMH devices mounted back to back, around which the virtual environment is positioned so as to enable bimanual interaction with bimanual haptic feedback when the hands are non-coplanar (see Figure 2).

The proposed system is composed of five parts: The virtual environment (VE), hardware and software for simulating and viewing the VE, hardware and software components for tracking the user, hardware for delivering UMH feedback, and a software coupling scheme for linking tracking data to visual and haptic feedback generation (see Figure 3).

3.1 System hardware

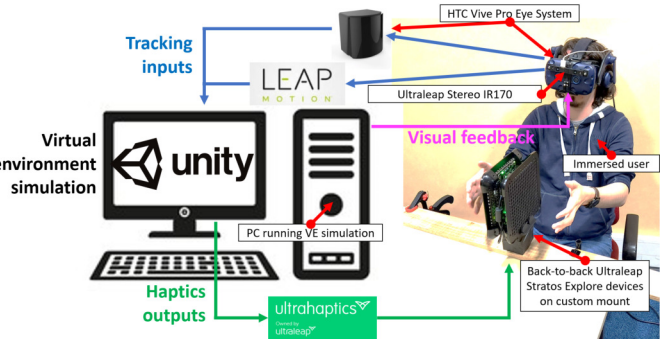

Fig. 2: Our proposed system for bimanual UMH manipulation in VR. Users view the virtual environment through an HTC Vive Pro Eye head-mounted display. Their hands are tracked using an Ultraleap Stereo IR170 camera mounted on the headset, allowing the user to interact with virtual objects haptically rendered using a pair of Ultraleap Stratos Explore devices mounted vertically back-to-back with the array planes parallel to the user's sagittal plane.

Users view the virtual environment through an HTC Vive Pro Eye head-mounted display (HMD) (pink arrows in Figure 2).

Tracking components are represented with blue arrows in Figure 2. The user's viewpoint is tracked using the HTC Vive's Lighthouse system. An Ultraleap Stereo IR170 camera mounted to the front of the HMD tracks the user's hands within the scene. The fixed camera pose with respect to the HMD enables a conversion of the hand tracking data to the virtual environment's world coordinate system defined around the HTC Vive's tracking region.

Ultrasound mid-air haptic feedback is provided by a pair of Ultraleap Stratos Explore arrays mounted vertically back-to-back on a custom 3D-printed mount placed on a table, such that the array planes face outward

[1] https://senso.me/
[2] https://www.senseglove.com/

and lie approximately parallel to the user's sagittal plane during interaction. The mount is designed to leave minimum space between the array faces (64 mm) so as to limit physical obstructions of the interaction region as much as possible. Software methods for dealing with this issue are described in section 3.2. The mount raises the arrays 10 cm above the table, to ensure that the full lateral workspace of the arrays remains usable and unobstructed. The haptic interaction workspace's dimensions are approximately 120x50x50 cm (WxHxD) [15].

Mounted in this way, the UMH devices are effectively independent, each one delivering haptic stimuli to one half of the interaction region, with no overlap between the acoustic fields generated by each device. Extending our proposed method to other arrangements of non-coplanar arrays mounted at angles with possibly overlapping workspaces would not fundamentally change the approach proposed here, but would require additional algorithms for generating haptic stimuli in the overlap between both devices' workspaces. Such an approach could be implemented using a priority system, or with algorithms for joint generation of focal points, as e.g. in the work by Matsubayashi et al. [26].

3.2 Coupling schemes

At the heart of the system lies a software coupling scheme linking the hand tracking data to the state of the virtual environment and to the generation of UMH feedback stimuli. Here, we consider two forms of manipulation interactions: those where the user can physically act upon the virtual environment, and receiving visual and haptic feedback according to the effected changes (orange arrow path in Figure 3), and those where the user can only explore the virtual environment, receiving haptic feedback of the encountered virtual objects (green arrow path in Figure 3. Each of these interaction types involves a slightly different coupling scheme, the specifics of which are respectively described below.

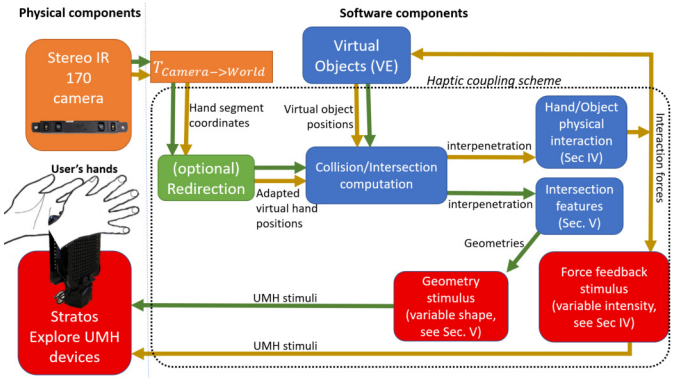

Fig. 3: General overview of the coupling schemes relating bimanual hand tracking to bimanual UMH rendering of interactions with the VE. In this paper we explore two cases: the case where the hands physically act on the VE, and thus reciprocally act on each other (section 4), and the case where the hands do not physically act on the VE and purely serve the purpose of exploring it (section 5). Tracking components are orange, haptics components are red, optional redirection components are green and VE objects and physics components are blue.

Redirection: The first common element in both coupling schemes is the optional hand redirection component.

Set up back to back on the custom-designed mount, the Ultraleap Stratos Explore UMH devices physically occupy a region which is 64 mm wide. In addition, UMH stimuli of a reasonable quality can only be achieved beyond a minimum distance of 50 mm from the front face of the Ultraleap Stratos Explore device. Therefore, a 164 mm wide "dead space" exists at the center of the interaction region. This could be dealt with by strategically placing a wide solid virtual object in the center of the workspace so as to keep users from bringing their hands together and going below the minimum haptic rendering distance or colliding with the arrays (as in our "pneumatic box" use case, see section 7). Another approach which poses fewer constraints on the design of the virtual environment is to use redirection. For our experiments (see section 4 and section 5), we offset the virtual hands by a fixed vector of 70 mm towards the sagittal plane. Virtual hands were thus 140 mm closer to each other than the participant's real hands were. In this way, even when the participants brought both their virtual hands in contact with, the physical hands remained about 40 mm away from the ultrasound arrays, and could thus receive UMH feedback. Having a fixed offset ensured that hand motions perceived visually matched those perceived kinaesthetically, which we expect will not disturb proprioceptive mechanisms involved in bimanual haptic perception.

Definition of hand interaction zones: Once the position of the virtual hand is computed, we can define a set of areas that will be able to interact with the VE and receive haptic feedback. For our experiments, we defined 6 areas approximated as bounded planes, corresponding to the palm and five fingertips, i.e. the hand regions where ultrasound feedback is felt with the highest intensity [30, 33]. Haptic feedback computed with these hand interaction regions can then be shifted back to the position of the non-redirected virtual hand, and then transformed into the reference frame of the haptic interfaces to be rendered.

Hand-object physical interaction: We consider physical interactions with a deformable object, assuming that the object can break if squeezed too tightly between hands. The object undergoes elastic deformation as soon as one or both hands interpenetrate it (see Figure 5). The computed spring force causes a visual deformation at the level of the contacting hand and serves as input for controlling the intensity of the UMH stimulus. We consider two modes for this object. If both hands are in contact, the object is attracted to the center of the hands using a critically damped spring-damper system, allowing the user to move a grasped object. But if at least one hand is not in contact, the object remains static and cannot break, although it can still be deformed, and can still generate feedback for the hand in contact. A formalization of this behaviour, corresponding to the orange path in Figure 3, is presented in Figure 4, and its application is illustrated in Figure 5 for a bubble manipulation (see section 4).

Exploration: computing intersection geometries To render the explored 3D shapes with UMH, we use an adaptation of an algorithm proposed by Martinez et al. [25] who showed that rendering both hand-object intersection contours as well as salient object features such as edges and vertices yielded good shape perception results. This algorithm, corresponding to the green path in Figure 3, is formalized in Figure 6. Applications of this algorithm are illustrated in Figure 7.

4 Object Grasping and Holding with Bimanual Haptics

As discussed in subsection 3.2 a first category of direct haptic manipulation in VR concerns interactions where the user physically acts on virtual objects, and thus the feedback reflects the reciprocal action of those virtual objects on the user. We performed an experiment to investigate whether our proposed system for delivering bimanual UMH feedback was functional in a bimanual object grasping and holding task, and whether bimanual haptic feedback provided a benefit over unimanual haptic feedback in this context. The experiment made use of the system described in section 3, following the coupling scheme presented by the orange arrows in Figure 3.

4.1 Hypotheses

We hypothesized that:

- **H1** **a**) Manipulation accuracy will be better in the presence of haptic feedback, and **b**) will be better when bimanual haptic feedback is provided than when unimanual haptic feedback is provided.
- **H2** **a**) Bimanual manipulation will be more efficient, i.e. faster at equal accuracy, in the presence of haptic feedback, and **b**) will be more efficient with bimanual haptic feedback than with unimanual haptic feedback.
- **H3** **a**) Manipulation task workload will be lower in the presence of haptic feedback, and **b**) be lower in the presence of bimanual haptic feedback than in the presence of unimanual haptic feedback.

Algorithm: Physical feedback

Input: The object *obj* interacted with
The two hands *leftHand*, *rightHand*
The threshold force f_{max} that makes the object break
The minimum haptic intensity $i_{min} \in [0,1]$ rendered when initiating contact
The spring-damper system *sd*

```
If leftHand.CollidesWith(obj) and rightHand.CollidesWith(obj) do
    /* Move the object */
    targetPosition ← (leftHand.position + rightHand.position) / 2
    obj.ApplyForce(sd.ComputeForce(obj.position, targetPosition))
    /* Compute the haptic feedback */
    leftForce ← AppliedForce(leftHand, obj)
    rightForce ← AppliedForce(rightHand, obj)
    meanForce ← (leftForce + rightForce) / 2
    If meanForce > f_max do
        Return /* The object broke, we do not render anything */
    Else do
        intensity ← Lerp(i_min, 1, meanForce / f_max)
        RenderFeedbackOnHand(Circle(leftHand.size, intensity), leftHand)
        RenderFeedbackOnHand(Circle(rightHand.size, intensity), rightHand)
Else do /* At most one hand is in contact */
    closestHand ← ClosestHand(obj, leftHand, rightHand)
    closestForce ← AppliedForce(closestHand, obj)
    If closestForce = 0 do /* No hand is in contact */
        Return
    Else do
        intensity ← Lerp(i_min, 1, closestForce / f_max)
        intensity ← Clamp(intensity, 0, 1)
        RenderFeedbackOnHand(Circle(closestHand.size, intensity), closestHand)
```

Fig. 4: Algorithm for rendering the haptic feedback for a physical interaction with an object. Here, `ApplyForce`, `CollidesWith`, and `AppliedForce` are parts of the simulation of the VE, which respectively apply a physical force to an object, determine if a hand is colliding with an object, and compute the force applied by the hand on said object. `ClosestHand` returns the hand that is the closest to the object. `RenderFeedbackOnHand` takes the computed geometry, expressed in the reference frame of the virtual hand, and converts it into haptic coordinates so that it can be rendered on the corresponding physical hand.

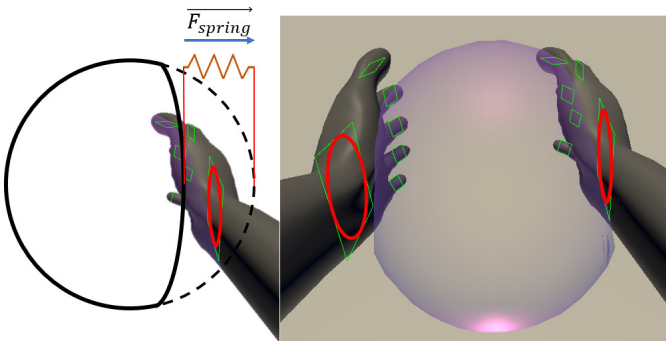

Fig. 5: Tactile force feedback: When a hand interpenetrates a deformable object, a spring force $\overrightarrow{F_{spring}}$ is computed based on the interpenetration between the hand and non-deformed object, the magnitude of which is used to proportionally adjust the intensity of a spatio-temporally modulated UMH circle (in red) projected on the hand interaction region (green rectangle) at the location of interpenetration.

H4 **a**) Participants will prefer performing manipulations in the presence of haptic feedback, and **b**) will prefer when this haptic feedback is bimanual.

H5 **a**) Participants will rate the manipulation experience as more coherent when haptic feedback is provided, and **b**) will find bimanual haptic feedback more coherent than unimanual haptic feedback.

4.2 Materials and methods

Task: To test these hypotheses, participants performed a grasping and holding task in VR (see Figure 8). A video of the task is provided as supplemental material. Participants began with both their hands in a fixed starting position, close to their body and far from the haptic interaction region. A virtual bubble appeared at a short distance in front of them, and they had to carefully bring their hands in contact with it in order to grasp it. Contact between a virtual hand and bubble applied a force to the bubble, which caused it to visually deform. If both hands were in contact with the bubble and applied excessive force, causing a deformation of more than 7% of the bubble radius, the bubble burst. In this case, participants had to return their hands to the starting position, so as to respawn the bubble and continue the task. After pilot testing, we set the haptic intensity when initiating contact with the bubble to 10% of the device capabilities, yielding a sound pressure of around 790 Pa. This is a good compromise between lower values that may not be perceivable, and higher values that would lead to a harder perception of the intensity change. This value is superior to the mean STM pattern detection thresholds for at least 75% of the population [14]. This deformation limit was chosen empirically to ensure that both haptic and visual deformations of the bubble were sufficiently perceivable,

Algorithm: Object exploration

Input: The object *obj* interacted with
The two hands *leftHand*, *rightHand*
The trigger distance *trigDist* for the feature rendering
The minimum required distance *betweenDist* between parts of the feedback

```
Foreach hand ∈ {leftHand, rightHand} do
    rawFeedback ← ∅
    Foreach part ∈ hands.parts do
        /* Dealing with the features */
        Foreach feature ∈ obj.features do
            rawFeedback ← rawFeedback ∪ CutFurther(feature, part, trigDist)

        /* Dealing with the intersection */
        rawFeedback ← rawFeedback ∪ OutlineIntersection(obj, part)

    /* Removing the parts that are too close to each other */
    cleanFeedback ← PruneFeedback(rawFeedback, betweenDist)
    RenderFeedbackOnHand(cleanFeedback, hand)
```

Fig. 6: Algorithm for rendering the haptic feedback for an object exploration interaction. Here, `features` represent the polylines associated to all of the object's features. For our experiments, linear and circular edges were represented as lines and circles of the same sizes, and vertices were represented as small circles with a 1 cm radius. `parts` contains all of the hand interaction parts. `CutFurther(a, b, c)` removes the parts of `a` that are at a distance greater than `c` from `b`. `OutlineIntersection(obj, part)` computes the outline of the intersection between the object `obj` and the hand interaction part `part`. `PruneFeedback(f, d)` optimizes the rendering by removing parts of the feedback `f` that are at a distance lower than `d`. `RenderFeedbackOnHand` takes the computed geometry, expressed in the reference frame of the virtual hand, and converts it into haptic coordinates so that it can be rendered on the corresponding physical hand.

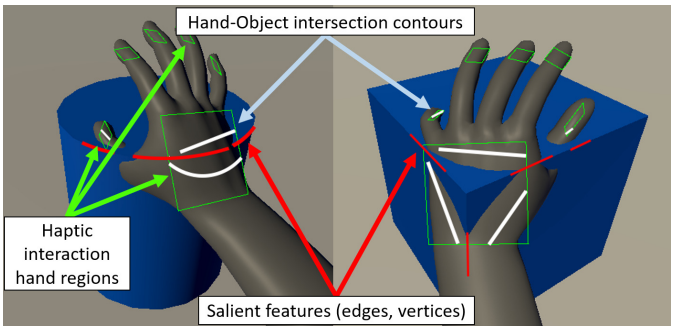

Fig. 7: Visualization of the haptic shape rendering approach. Proximity of the hand interaction regions (green rectangles) to salient features (red) are coupled with intersection contours of the interaction regions with the shape (white) to form sets of line segments rendered with spatio-temporal modulation.

without making the task too easy. The trial ended if participants successfully held the bubble between both hands, without bursting it, for a consecutive duration of 4 s. This duration was also empirically chosen to be sufficiently long to ensure that participants had to make an effort to succeed, while keeping the overall experiment duration reasonably short. After explaining how a trial works, the experimenter asked the participants to grasp the bubble as fast as possible, while trying to avoid bursting it.

Experimental procedure: participants performed the task in three blocks, one for each of the following feedback conditions, whose order was counterbalanced across participants:

- **N** No haptic feedback, in which only visual deformation and motion of the bubble informed about the force applied.
- **U** Unimanual haptic feedback, in which the participants' dominant hand received additional tactile feedback in the form of a spatio-temporally modulated circle [8] rendered on their palm, whose intensity was proportional to the force applied to the bubble (see section 3.2).
- **B** Bimanual haptic feedback, in which both participants' hands received simultaneous tactile feedback identical to that delivered in the **U** condition.

Within each block, participants performed 18 trials which randomly alternated between a large (22 cm radius) and a small (12 cm radius) sphere, each appearing 9 times. These object sizes were chosen to cover a sufficiently wide span of bimanual grasping motions within the haptic interaction region, without compromising hand tracking performance.

Collected data: During each trial, we measured manipulation accuracy by recording the number of bubble bursts as well as the total time during which participants accidentally break contact with the bubble (release time) while trying to hold it. As measures of manipulation efficiency, we recorded the total time of the task and the time users were holding the bubble. The lower these two times, the more we considered the manipulation to be efficient. It should be noted that both times were necessarily greater than 4 s, which was the minimum consecutive holding time required to succeed in the task.

At the end of each block, participants filled out a raw NASA-TLX questionnaire [12], from which a global workload measure was calculated. Participants also responded to a subjective questionnaire assessing perceived visuo-haptic coherence during the manipulation task for two criteria: perceived coherence between the sensory feedback and expectations given the scene appearance (referred to as **c-FB/EXP**), and perceived coherence between the applied force and bubble deformation (referred to as **c-F/D**), both using 7-point Likert scales. They were asked the following two questions: "Rate the coherence between the haptic feedback and what you would expect from the virtual scene", and "Rate the coherence between the deformation and the "force" you felt".

At the end of the experiment, participants were asked to rank the feedback conditions in order of overall preference, as well as based on how easily they were able to establish a strategy for performing the task. We also asked participants to rate the perceived coherence between the motion of their real and virtual hands after each block to ensure that the hand redirection method described in section 3.2 did not have a detrimental impact on the interaction.

Population: The study involved 12 participants (10 M, 1 F, 1 Non-binary) aged 22 to 32 years (m: 24.9y, sd: 3.1y). Ten participants were right-handed, 1 left-handed and 1 ambidextrous. Half of the sample were experienced in VR, while the other was not. Three participants had extensive experience in haptics, while 7 only had limited haptics experience.

4.3 Results

Mean results for each evaluated criterion are summarized in Table 1. Key results for total task time, number of bubble bursts, and subjective workload estimates are also shown in Figure 9.

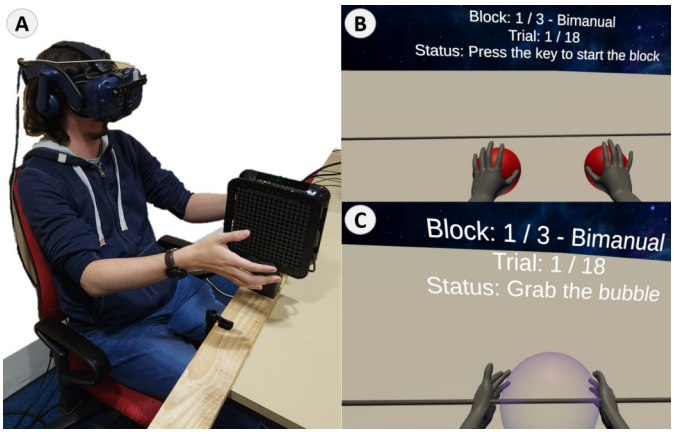

Fig. 8: (A) Participants were seated in front of the system described in section 3, viewing the VE in which they performed the task. (B) At the start of each trial, participants placed their hands at fixed starting positions indicated by a pair of red spheres. (C) A virtual bubble which participants were to grasp and hold then spawned in front of them. Instructions and feedback on progress were provided via a text panel in the VE. A horizontal bar also indicated the preferred axis along which to grasp the bubble to ensure both hands were used in a repeatable manner across trials and participants.

Objective metrics For each participant, we averaged all the objective metrics per type of feedback, leading to 36 (3 blocks, 12 participants) values per metric. Total interaction time and bubble burst distributions are illustrated in Figure 9 (left and center).

For each of the duration metrics (release time, holding time, and total interaction time), either the normal distribution hypothesis was not rejected (Shapiro test, $p > 0.05$ for each feedback), or the Q-Q plot was good enough to perform an ANOVA, given its robustness to moderately non-normal data. Homogeneity of variances was also verified for all three metrics (Levene test, $p > 0.05$). We thus performed an ANOVA, with the duration as dependant variable, the feedback as independent variable, and the participant as random effect variable. Significant effects on holding time ($p = 0.01$) and total interaction time ($p < 0.001$) were found, but none was found on release time ($p > 0.05$). Post-hoc Holm-corrected pairwise t-tests showed that **U** led to significantly lower holding time than **N** ($p = 0.02$), and that **B** and **U** yielded significantly lower interaction time than **N** ($p = 0.006$ for both comparisons).

On the other hand, Shapiro tests and Q-Q plots rejected the normal distribution hypothesis ($p < 0.05$) for the number of bubble bursts. We therefore applied a Friedman test which revealed a significant effect of the feedback condition ($p = 0.003$). Post-hoc pairwise Wilcoxon tests with Holm corrections revealed that both **B** and **U** led to significantly less bubble bursts than **N** ($p = 0.015$ for both comparisons).

Subjective metrics For all subjective metrics, we gathered one value per participant, per feedback condition, leading to 36 values per metric. The workload score distribution is shown in Figure 9 (Right).

After verifying its assumptions (Shapiro test, Levene test, $p > 0.05$, an ANOVA showed a significant effect of the feedback on the score ($p < 0.001$), and a post-hoc analysis with pairwise Holm-corrected t-tests revealed that the workload under condition **B** was significantly lower than that under condition **U** ($p = 0.007$), which in turn was significantly lower than that under condition **N** ($p = 0.001$).

For the **c-FB/EXP** and **c-F/D** perceived coherence scores, we performed a Friedman test, which revealed a significant effect of the feedback ($p < 0.001$ in both cases). In both cases, pairwise Holm-corrected Wilcoxon tests revealed that **B** and **U** were rated as more coherent between visuals and haptics than **N** ($p = 0.007$ for both **c-FB/EXP** comparisons, $p = 0.009$ for both **c-F/D** comparisons).

Friedman tests ($p < 0.001$) followed by pairwise Holm-corrected Wilcoxon post-hoc tests on the feedback condition ranking data showed

Table 1: Mean results and standard deviation for the object grasping and holding experiment. For the first four metrics, standard deviations are computed on the dataset containing all trials, including repetitions. "Bubble bursts" is the number of times the participants applied too much force during grasping, resulting in the bubble bursting (lower is better). "Release time" is the cumulative time during which participants accidentally break contact with the object during grasping (lower is better). "Holding time" is the cumulative time during which participants grasped the object before successfully completing the task (lower is better). "Total time" is the total task execution time, i.e., release time + holding time + time to get hands into position (lower is better). "Workload" is the TLX score (lower is better). "c-FB/EXP" is the coherence rating between sensory feedback and expectations given the scene's appearance (7-point Likert scale, higher is better). "c-F/D" is the coherence rating between the force applied and the bubble deformation (7-point Likert, higher is better). Finally, we consider two rankings of the conditions based on "overall preference" for performing the task and "ease of establishing a (grasping and holding) strategy" (lower is better).

Feedback	Bubble bursts	Release time (s)	Holding time (s)	Total time (s)	Workload (/ 100)	c-FB/EXP (/ 7)	c-F/D (/ 7)	Ranking: overall preference (/ 3)	Ranking: ease of establishing a strategy (/ 3)
B	0.33 ± 0.75	0.51 ± 1.86	5.80 ± 5.21	11.27 ± 11.32	31.25 ± 12.12	6.17 ± 0.83	4.92 ± 1.44	1.25 ± 0.62	1.17 ± 0.58
U	0.44 ± 1.56	0.70 ± 1.75	5.95 ± 4.40	12.48 ± 12.56	39.93 ± 14.35	5.42 ± 1.51	5.42 ± 1.00	1.92 ± 0.51	2.08 ± 0.29
N	1.12 ± 2.95	2.51 ± 9.43	7.67 ± 6.68	21.07 ± 25.42	53.82 ± 16.45	1.25 ± 0.87	1.42 ± 1.44	2.83 ± 0.39	2.75 ± 0.62

that **B** and **U** were significantly preferred to **N** ($p = 0.008$ and $p = 0.01$ respectively). Participants also found it significantly easier (Friedman test, $p < 0.001$) to establish a grasping strategy in the **B** condition than in the **U** condition ($p = 0.02$). Establishing a grasping strategy was also significantly easier in the **U** condition than in the **N** condition ($p = 0.02$).

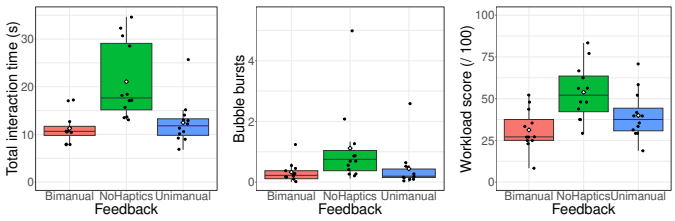

Fig. 9: Distributions of the mean interaction time per trial (Left), mean bubble burst count per trial (Middle), and workload (Right) across participants, for each condition. Black dots represent single participant results, and white diamonds represent the mean value for each condition.

Finally, according to the questionnaires, the hand redirection method employed in the experiment does not appear to have disturbed users (hand coherence rated 6.17 / 7 on average).

4.4 Discussion

Regarding hypotheses **H1** (manipulation accuracy), we observed no difference in release time, but significantly more bubble bursts in **N** than in **B** and **U**. These results tend to support **H1.a**. Post-hoc testing revealed no significant differences were found between **B** and **U**, thus **H1.b** is not supported.

When it comes to manipulation efficiency, both **B** and **U** required significantly less total time than **N**. This was mostly because of the difference in approach time which can be traced back to the significantly larger number of bubble bursts in **N**, although there was also a significant difference in grasp times between **U** and **N** which also contributed. Coupled with the significantly lower bubble burst rates in these conditions, total task times support **H2.a**. However, **H2.b** was not supported, as we found no significant differences between **B** and **U**.

Looking at the NASA-TLX scores, **B** scored significantly lower than **U**, which in turn scored significantly lower than **N**. This supports **H3** entirely, indicating that bimanual haptic feedback yielded the lowest subjective workload for the task at hand.

Participants significantly preferred **B** and **U** over **N**. This tends to support **H4.a**. but does not support **H4.b**. They also found it significantly easier to establish a grasping strategy in **B** than in **U**, while also finding it significantly easier to do so in **U** than in **N**, fully supporting **H4**. In summary, **H4.a** appears fully supported, but **H4.b** is only partially supported by our results.

Conditions **B** and **U** both yield significantly higher perceived coherence than **N**, supporting **H5.a**. However, our results do not support **H5.b**.

Overall, for all tested metrics it appears that the bimanual haptic interaction performed as good as or better than unimanual haptic feedback, and consistently better than no haptic feedback. We can therefore confidently conclude that our proposed novel approach to bimanual UMH is functional in this type of 3D object manipulation task with tactile force feedback. The benefit of bimanual UMH over unimanual UMH is mainly apparent when it comes to subjective ease of use and workload measures. We can therefore conclude that our approach can be particularly suited to improving user experience in VR manipulation, but does not necessarily change much in terms of objective performance metrics when compared to simpler approaches to UMH feedback. However, bimanual UMH appears to conserve all performance benefits that can also be obtained through unimanual UMH, making it at least as good of a candidate for a haptic feedback method as the state of the art.

Overall, these results appear promising, and future work should be conducted to assess to what extent the observed benefits may transfer to a wider range of VR manipulations where users physically act on virtual objects.

5 3D Object Shape Exploration with Bimanual Haptics

A second category of direct haptic manipulations in VR concerns interactions where the user merely explores the virtual environment, receiving feedback but not physically acting on virtual objects. We performed an experiment to investigate whether our proposed bimanual UMH feedback system was functional in bimanual 3D object shape exploration tasks, and whether such haptic exploration with bimanual UMH feedback provided any benefit over what could be achieved with a single UMH device. The experiment once again made use of the system described in section 3, following the coupling scheme presented by the green arrows in Figure 3.

5.1 Hypotheses

We hypothesized that:

H6 Shape identification accuracy will be higher when simultaneous bimanual haptic feedback is provided than when unimanual haptic feedback is provided.

H7 Shape identification speed will be higher when simultaneous bimanual haptic feedback is provided than when unimanual haptic feedback is provided.

H8 Shape identification task workload will be lower than when simultaneous bimanual haptic feedback is provided than when unimanual haptic feedback is provided.

H9 Participants will prefer performing the task when bimanual haptic feedback is provided rather than when unimanual haptic feedback is provided.

H10 Participants will perceive the virtual shapes as equally coherent when exploring them unimanually or bimanually.

H11 In alternating unimanual exploration with both hands (pseudo-bimanual, see **P** below), shape identification accuracy will be

higher due to redundancy in perception between hands [34]. Shape identification speed will be lower than in single-hand unimanual exploration because of the added sequential exploratory motions. For all other metrics, pseudo-bimanual feedback will perform similarly to unimanual feedback, since the haptic exploration task remains essentially unimanual in both cases.

5.2 Materials and Methods

Task: To test these hypotheses, participants performed a virtual 3D shape exploration task. A virtual glovebox was placed in front of the user, as shown in Figure 10-(B, C, D). Two holes on the front side of the glovebox enabled the users to insert their hands into the box to haptically explore its content. An invisible 3D object was generated within the glovebox. The ultrasound haptic interfaces were located at the center of the glovebox, mounted vertically back-to-back (see section 3 and Figure 10-A), and providing haptic feedback of the invisible object to the exploring hand(s). Participants received feedback on the exploring hand(s) according to the rendering scheme discussed in subsection 3.2. The task was to identify the shape from a set of four possible shapes as accurately as possible, and as fast as possible. Participants responded by selecting a 3D shape from a menu of the four possible shapes using the index finger of their dominant hand as a pointer. The order of shapes in the menu was randomized at each trial. To cap the duration of the experiment to a reasonable time, no more haptic feedback was provided after a duration of 30 s, forcing participants to make a choice if they had not yet responded. They then similarly rated their level of confidence on a 7-points Likert scale. After being informed of the trial sequence, participants were instructed to answer as quickly as possible once they had a reasonable idea of what the shape was. A video of the task is provided as supplemental material.

Experimental procedure: Participants performed the task in three blocks, one for each of the following exploration and feedback conditions, which were counterbalances across participants:

- **U** Unimanual, in which the participants' dominant hand was used alone, and received tactile feedback. This condition represents the current state-of-the-art for UMH virtual shape exploration [21, 25, 26].
- **P** Pseudo-bimanual, in which participants' hands *alternately* explore the shape, and can only receive tactile feedback one at a time. This condition was chosen to investigate the effectiveness of involving both hands in a way where UMH feedback could be provided using conventional single-board approaches (e.g. [3, 15]).
- **B** Bimanual haptic feedback, in which both participants' hands simultaneously explore the shape, both receiving tactile feedback at the same time.

It is worth noting that the **P** condition was not present in the first experiment (section 4). Indeed, this experiment required the interactions of both hands simultaneously. While alternating feedback could have been implemented, any alternating rule would have been rather arbitrary and unnatural. For this reasons, it was excluded from the first experiment. Similarly, the **N** condition from the first experiment was not considered for this second experiment, since in the absence of visual cues, the task could not be performed without haptic cues. Future work can focus on comparing bimanual feedback to pseudo-bimanual feedback for other manipulation tasks.

Four possible 3D objects (Sphere, Cylinder, Cone, Cube) were presented, such that they were symmetrical around the seated participant's sagittal plane. In this way, the same quantity of information about the objects was accessible in all conditions. We presented each shape in two sizes: large (20 cm width and height) or small (13 cm width and height). To ensure all relevant features of each shape could be perceived in all conditions, the shape orientations were held constant and corresponded to those of the visual shapes in the response menu. participants were informed that the objects could have one of two different sizes, but that they only needed to indicate which shape they felt, not its size. Within each block, each shape was presented four times in total, twice in the small size and twice in the large size. The order of presentation of shapes was fully randomized.

Collected data: For each trial, we recorded task completion time as the sum of time spent in contact with the shape. This value was capped at the maximum exploration time of 30 s, and lower values indicate better performances.

For each condition, we measured precision and recall for shape identification, with higher values indicating better performances. At the end of each block, participants filled out a raw NASA-TLX questionnaire [12] to assess task workload. We also assessed perceived coherence **(c-H/V)** between the explored haptic shapes and the visual representation that was provided in the response menu. Finally, we inquired about the strategies that participants developed to perform the task, and whether they changed their strategies during a block or between blocks.

At the end of the experiment, participants were asked to rank the feedback conditions in order of overall preference, as well as based on how easily they were able to establish a strategy for performing the task. We also asked participants to rate the perceived coherence between the motion of their real and virtual hands after each block to ensure that the hand redirection method described in section 3.2 did not have a detrimental impact on the interaction.

Population: The study involved 18 participants (16 M, 2 F) aged 18y to 53y (m: 25.8y, sd: 7.3y). 16 participants were right-handed, 1 left-handed and 1 ambidextrous. Four participants had extensive VR experience, while 11 only had limited VR experience. Two participants had extensive experience in haptics, while 10 only had limited haptics experience.

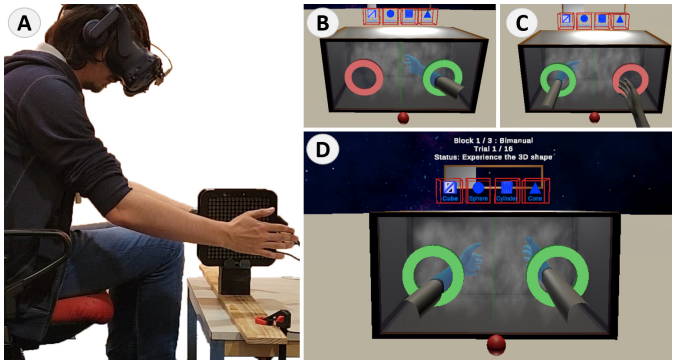

Fig. 10: (A) Participants were seated in front of the setup described in section 3, and viewed a virtual glovebox in which an invisible 3D object could haptically be explored either (B) unimanually, (C) with both hands alternating one at a time (pseudo-bimanual), or (D) with both hands simultaneously (bimanual). Above the glovebox, participants could view instructions for the trial as well as use a menu of shapes to identify.

5.3 Results

Mean results for each evaluated criterion are summarized in Table 2.

Objective metrics One precision and recall value was computed for each condition and each participant, leading to 54 values in total. One participant was removed from the precision data, as their data contained a block where one shape was never chosen, meaning that the precision computation was not possible. Figure 11 shows the confusion matrices for each condition. Task execution times were averaged across all trials within each condition, for each participant, resulting in a total of 54 values (see Figure 12 (Left)).

The ANOVA assumptions were verified (Shapiro test, Levene test, $p > 0.05$), but the ANOVA failed to reveal any effect of feedback on precision and recall ($p > 0.05$). However, Holm-corrected t-tests revealed that both these values were significantly above the 25% chance level ($p < 0.001$).

Q-Q plots and a Levene test ($p > 0.05$) on the task completion times (see Figure 12 (Left)) allowed us to perform an ANOVA which revealed

Table 2: Mean results and standard deviation for the 3D object shape exploration experiment. For the interaction time, the standard deviation is computed on the dataset containing all trials, including repetitions. “Precision” and “Recall” measure the shape identification accuracy (higher is better). “Interaction time” is the time spent close to the experienced shape (lower is better). “Confidence" is the score given by the participant about the confidence they have in each of their answers (7-point Likert scale, higher is better). “Workload” is the TLX score (lower is better). “Coherence" is the coherence rating between the perceived stimulation and the visually rendered shape (higher is better). Finally, we consider two rankings of the conditions based on “overall preference” for performing the task and “ease of establishing an exploration strategy” (lower is better).

Feedback	Precision	Recall	Interaction time (s)	Confidence (/ 7)	Workload (/ 100)	Coherence (/ 7)	Ranking: overall preference (/3)	Ranking: ease of establishing a strategy (/3)
B	0.52 ± 0.18	0.50 ± 0.16	24.32 ± 6.57	4.39 ± 0.81	52.08 ± 15.41	4.28 ± 1.23	1.39 ± 0.70	1.44 ± 0.70
U	0.53 ± 0.16	0.52 ± 0.15	23.32 ± 6.54	4.25 ± 0.94	54.28 ± 14.05	3.78 ± 1.22	2.39 ± 0.70	2.33 ± 0.84
P	0.50 ± 0.22	0.48 ± 0.19	25.11 ± 6.33	4.16 ± 0.81	54.05 ± 16.87	4.06 ± 1.11	2.22 ± 0.73	2.22 ± 0.65

significant effects of feedback ($p = 0.04$). Holm-corrected pairwise t-tests showed that participants performed significantly faster in the **U** condition than in the **P** condition ($p = 0.04$).

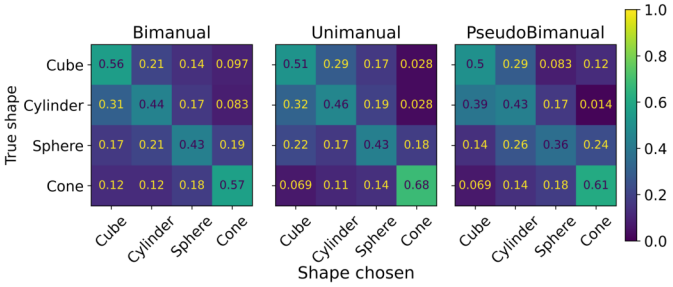

Fig. 11: Confusion matrices for the 3D shape selection. Values correspond to recall for the three conditions.

Subjective metrics We recorded one confidence value per trial, and averaged them per block. For all other subjective metrics, we gathered one value per participant, per feedback condition. In all case, this lead to 54 values per metric (3 blocks, 18 participants).

Confidence and workload values had a good-enough Q-Q plot, and similar variances between feedbacks (Levene test, $p > 0.05$), allowing us to perform an ANOVA, which showed no significant effect of the feedback over any of the two factors ($p > 0.05$).

A Friedman test also failed to identify any significant difference in perceived visuo-haptic shape coherence ($p > 0.05$).

That being said, we found a significant effect of the feedback on the participant’s preferences (Friedman test, $p = 0.006$), showing that **B** was preferred to both **U** and **P** ($p = 0.04$ and $p = 0.02$ respectively).

Participants also found it easier to establish a strategy in the **B** condition, than in the **U** and **P** conditions ($p = 0.04$ for both comparisons).

Finally, according to the questionnaires, the hand redirection method employed in the experiment was not perceived as too disturbing by users (rated 5.22 / 7 on average).

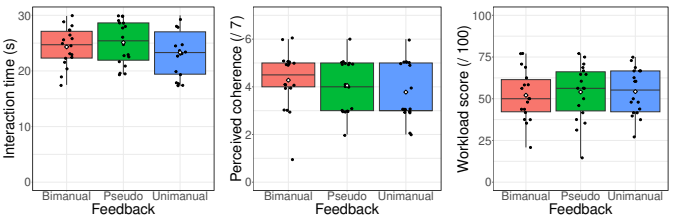

Fig. 12: Distributions of the mean interaction time per trial (Left), perceived coherence between the shapes and haptic feedback (Middle), and workload (right) across participants, for each condition. Black dots represent participant results, and white diamonds represent the mean value of each condition.

5.4 Discussion

All conditions yielded performances significantly above chance levels, meaning all methods allowed effective shape identification. However, no significant differences were observed between conditions when it comes to precision or recall. Therefore, we cannot accept **H6**, as simultaneous bimanual haptic feedback appears to yield identical performance to unimanual and pseudo-bimanual haptics.

The only significant difference in shape identification speed was observed between **U** and **P**. Thus, **H7** is not supported.

The NASA-TLX scores showed no significant differences between conditions, failing to support **H8**.

In terms of participant preferences, **B** was deemed significantly better than both **U** and **P**, both in terms of overall preference for accomplishing the task and in terms of ease of establishing an exploration strategy. **H9** is therefore fully supported.

No significant differences were found between conditions in terms of perceived coherence between the haptic and visual representation of shapes, supporting **H10**.

H11 appears partly supported, as **U** significantly outperformed **P** in terms of shape identification speed and otherwise no significant differences were observed between both. However, **P** did not perform significantly better than **U** with respect to shape identification accuracy.

Similarly to the experiment on object grasping and holding (section 4), fully bimanual UMH, as enabled by our system, appears to mainly be beneficial in terms of subjective preference, and not in terms of objective performance in haptic shape exploration tasks. However, it should be noted that the lack of observed differences may however mask a potential advantage of bimanual haptics in a more realistic use scenario, where objects are not necessarily symmetrically arranged around the user’s sagittal plane. Further investigations in less constrained scenarios may be warranted to assess this. We will also conduct similar studies with different 3D object rendering techniques, which may impact recognition performances. Also, while participants did not report any disturbance in task execution linked to the hand redirection used in the VR scene, additional investigation into possible impacts of redirection on bimanual perception would be required to rule out any detrimental impact on task performance.

Contrary to our expectations, involving both hands with a feedback scheme which could be implemented using existing mobile single-board UMH systems [3, 15] does not improve shape identification performances when compared to conventional unimanual UMH using a fixed array. Overall, it appears that if both hands are to be involved in UMH shape exploration interactions in VR, they should be involved in such a way that they can both receive simultaneous tactile feedback so as to maximise benefits in terms of user experience.

In the future, we will design and evaluate different variations of the shape rendering technique in the hopes of improving identification accuracy, which may then show deeper differences between the three conditions.

6 Global Discussion

Our system enabled simultaneous generation of UMH feedback on both hands. Our results showed that, both for a grasping and exploration task, bimanual feedback led to performances at least as good as with unimanual feedback. Therefore these techniques are at least as good as the state-of-the-art and they show that simultaneous UMH feedback on both hands does not seem to degrade one's integration of the stimuli.

The absence of significant differences between the bimanual and unimanual feedback may be explained by the relatively low difficulty of the tasks, as evidenced by the low amount of bubble pops in the grasping task (section 4). We expect that more difficult tasks would accentuate such differences, so future work will be conducted to assess performances on a wider range of manipulation tasks, and to measure the level of difficulty required to have significant improvements generated by the bimanual feedback.

That said, bimanual UMH feedback offers two significant advantages. First, our results show that UMH feedback significantly improves the user experience, by reducing the workload (section 4), being preferred by participants (section 5), and making the establishment of an interaction strategy easier (section 4, section 5). Secondly, bimanual feedback enables the completion of less constrained tasks that would be impossible to perform with only unimanual feedback. For instance, let us consider our exploration task (section 5), but with a different set of shapes. We could for example consider two-parts shapes where the right side is composed of a cylinder, and the left part is composed of either of a cone, a cylinder, or a semi-sphere, with both parts being joined at their base, lying on the sagittal plane. Given that all these shapes have the same right-part, any unimanual exploration with the right hand would fail to determine with certainty which of these three shapes is being explored. But bimanual feedback would enable the rendering and perception of the left part of the shape, which could then make discrimination possible. Future work will also study the bimanual performances for such tasks.

It is worth noting that our experiments used populations composed of 10 out of 12 and 16 out of 18 males respectively. This is not representative of the global population, meaning that our results are to be interpreted and used carefully. Future work will be focused on expanding our studies to a wider and more diverse population, in order to strengthen our results and improve their generalizability.

Then, it is also worth noting that the presence of our system in the interaction space is in itself a potential limitation. When interacting with small objects in front of the user, they may need to have both hands close to each other, which result in contacts with the system. To avoid this issue, we proposed to use hand redirection (section 3.2). Our results show a strong hand coherence perception, which leads us to believe that redirection was not an issue in our case. That said, the redirection offset would need to be increased when interacting with even smaller objects, which would enlarge the proprioceptive discrepancy, and may then be perceived by users. Future work will be conducted to assess the range of acceptable redirection offset, as well as to design interaction techniques to allow a wider workspace with lower risk of colliding with the system.

Overall, this first bimanual UMH system appears very promising to improve user experience in haptic VR manipulation tasks, while also enabling a whole new range of interactions. Future work will be focused further improving the system, while also studying its performances for other types of interactions and other hand configurations.

7 Use cases

We tested the proposed bimanual ultrasound mid-air haptics approach in the three immersive VR scenarios (see Figure 13) described below. Videos of the use-cases are also provided as supplemental material.

Pneumatic puzzle box: In this first use-case, users face a pneumatic puzzle, box where the aim is to direct a ball through a circuit and into a goal using combinations of air jets. Users can press different pistons and connect sets of nozzles and stoppers to manipulate the air flow in order to achieve this objective.

This use-case exemplifies a VR interaction where the optional hand redirection components discussed in subsection 3.2 are not necessary, as a wide solid virtual object obstructs the region occupied by the pair of UMH devices. This naturally guides users away from the region where collisions between their hand and the devices could occur.

In this interaction, the hands can perform independent haptic manipulations, showing how our system can also be used to render different haptic sensations due to different interactions between each hand and the virtual environment. For example, a spatio-temporally modulated circle of varying intensity can be rendered on the right hand as it presses down on a piston (similar to the tactile force feedback scheme discussed in section 4, while the left hand explores an outgoing air jet, receiving a spatio-temporally modulated circle whose radius and intensity randomly oscillate around fixed values, simulating turbulent air flow.

Virtual pottery: In this second use-case, users can rotate a virtual block of clay and shape it as they desire by pressing their hands into it and deforming it. Tactile feedback is provided based on contact and interpenetration between the hand interaction regions (see section 3.2) and the virtual clay. The edges of the hand interaction regions which come into contact with the clay are added to a pool of path points for a global spatio-temporally modulated polyline drawn across all hand parts in contact with the clay. In addition, the intensity of each line segment is increased from a baseline value, proportionally to the local interpenetration between hand and clay.

This tactile force feedback scheme is similar to that described in section 3.2 and used in the experiment presented in section 4, although the manipulated virtual object in this case is fixed in the virtual environment and can only incur plastic deformation instead of being both mobile in the environment and incurring elastic deformation.

While the realism of the haptic feedback may not be perfect, partly due to the relative softness of ultrasound-based vibrotactile feedback compared to the physical properties of actual clay, it has been shown that using UMH feedback for clay modelling can actually provide very rich interactions, while helping the participants better control their hand movements [2].

Virtual concertina: In this last use-case, we provide an example of how our approach to UMH feedback can be applied to guiding a user's motions. Users manipulate a virtual concertina and must move their hands in rhythmical back-and-forth motions to trigger bars of music. Spatio-temporally modulated circles which grow and shrink in accordance with the ideal opening and closing motions to be made by the hands are projected onto the users' palms, providing a tactile guide for playing the melody with the correct rhythm.

8 Conclusion

In this paper, we described a novel approach for providing bimanual ultrasound mid-air haptic (UMH) feedback during bimanual manipulation in VR. By using a pair of UMH devices back-to-back, we enabled simultaneous bimanual manipulation of single virtual objects in opposing hand configurations. We investigated the benefits of bimanual UMH feedback in two representative human participant experiments studying a bimanual object grasping and holding task, and a bimanual 3D shape exploration and identification task.

Results show that our approach yields performances at least on par with state-of-the-art unimanual UMH feedback, and significantly improves user experience (lower workload, higher preference, greater ease to establish an interaction strategy) during such bimanual manipulation when compared to unimanual UMH feedback. These results confirm the viability of our approach, complementing the current state-of-the-art for bimanual ultrasound haptics in VR, which thus far has been limited to simple demonstrations of interactions with coplanar hands. In addition, many tasks would be impossible with unimanual feedback, but our system would enable their completion, although future work is required to evaluate performances in such unconstrained environments.

While we provide an interesting first study of bimanual feedback with UMH, the results are currently limited to a quite male-skewed population. Future work will expand our studies to a wider and more diverse population to make our results more generalizable.

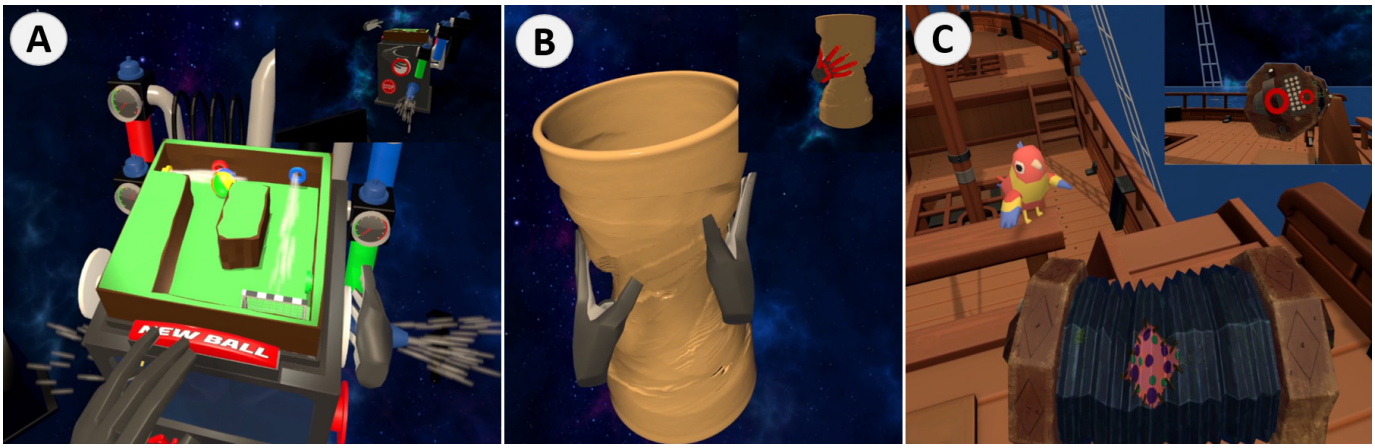

Fig. 13: Three VR scenarios exemplifying variations on applications and implementations of our approach. The figures show the VR view in first-person perspective, while the miniatures in the top right corners show a third person view with the tactile feedback patterns highlighted in red. (A) Pneumatic puzzle box use-case: users receive diverse tactile feedback while exploring various air-jets, pressing on pistons and manipulating nozzles. (B) Virtual pottery: tactile force feedback is provided across the hand as users mould a block of virtual clay into shape. (C) Virtual concertina: tactile rhythm cues are provided simultaneously to both hands to guide a user's motions when playing a tune on the virtual instrument.

A general limitation of our work was that we looked at a subset of VR manipulation tasks involving static virtual objects. In future work, it would be interesting to investigate potential benefits to user performance and experience in a wider range of manipulation tasks, in particular tasks involving object motions.

In the future, we also plan to work on key limitations of this system, in particular its workspace, and intend to provide the ability to enable both coplanar and opposing hand configuration, e.g. by actuating the pair of UMH devices [3, 15]. We will optimize the coupling scheme by experimenting with different tracking systems. We also plan on studying the perceptual implications of the number of interaction regions on the hand on users' experiences. This in turn will allow us to investigate potential benefits of bimanual UMH to user experience and performance when executing a wider range of bimanual manipulation tasks in VR. In parallel, we plan to extend the scope of haptic feedback beyond purely physics-based feedback as was investigated in this paper, to encompass e.g. guidance cues [27], which may yield greater benefit in terms of user performances. This work will hopefully open novel perspectives for applications of ultrasound mid-air haptics to VR interaction, and broaden the scope of possibilities for enriching users' experiences in VR with haptics.

Acknowledgments

This project has received funding under the ANR project "MIMESIS" and the European Union ERC project "ADVHANDTURE" (101088708).

References

[1] P. Balint et al. Medical virtual reality palpation training using ultrasound based haptics and image processing. *Proc. Jt. Work. New Technol. Comput. Assist. Surg*, 2018. 1, 2

[2] H. Barreiro et al. Natural tactile interaction with virtual clay. In *Proc. IEEE World Haptics Conf.*, pp. 403–408, 7 2021. 9

[3] D. Brice et al. A proof of concept integrated multi-systems approach for large scale tactile feedback in vr. In *Proc. AVR Conf.*, pp. 120–137. Springer, 2019. 1, 2, 7, 8, 10

[4] D. Brice et al. Using ultrasonic haptics within an immersive spider exposure environment to provide a multi-sensorial experience. *Frontiers in Virtual Reality*, 2:707731, 2021. 1

[5] G. Cirio et al. The virtual crepe factory: 6dof haptic interaction with fluids. In *ACM SIGGRAPH Emerging Tech.*, pp. 1–1. 2011. 2

[6] X. De Tinguy et al. Enhancing the stiffness perception of tangible objects in mixed reality using wearable haptics. In *Proc. IEEE VR Conf.*, pp. 81–90. IEEE, 2018. 2

[7] T. Duan et al. Flyinghand: extending the range of haptic feedback on virtual hand using drone-based object recognition. In *SIGGRAPH Asia 2018 Technical Briefs*, 2018. 1, 2

[8] W. Frier et al. Using spatiotemporal modulation to draw tactile patterns in mid-air. In *Proc. Eurohaptics Conf.*, pp. 270–281. Springer, 2018. 5

[9] M. Frutos-Pascual et al. Evaluation of ultrasound haptics as a supplementary feedback cue for grasping in virtual environments. In *Proc. Int. Conference Multimodal Interaction*, pp. 310–318, 2019. 1, 2

[10] O. Georgiou et al. Touchless haptic feedback for vr rhythm games. In *Proc. IEEE VR Conf.*, pp. 553–554, 2018. 1, 2

[11] Y. Guiard. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *J. Mot. Behav.*, 19(4):486–517, 1987. 2

[12] S. G. Hart. Nasa-task load index (nasa-tlx); 20 years later. In *Proc. Hum. Fact. Erg. Soc. Ann. Meet.*, vol. 50, pp. 904–908. Sage publications, 2006. 5, 7

[13] K. Hinckley et al. Attention and visual feedback: the bimanual frame of reference. In *Proc. Symp. Interactive 3D Graph.*, pp. 121–ff, 1997. 2

[14] T. Howard et al. Investigating the recognition of local shapes using mid-air ultrasound haptics. In *Proc. IEEE World Haptics Conf.*, pp. 503–508, 2019. 4

[15] T. Howard et al. Pumah: Pan-tilt ultrasound mid-air haptics for larger interaction workspace in virtual reality. *IEEE Trans. Haptics*, 13(1):38–44, 2019. 1, 2, 3, 7, 8, 10

[16] T. Howard et al. Ultrasound mid-air tactile feedback for immersive virtual reality interaction. In *Ultrasound Mid-Air Haptics for Touchless Interfaces*, pp. 147–183. Springer, 2022. 1

[17] I. Hwang et al. Airpiano: Enhancing music playing experience in virtual reality with mid-air haptic feedback. In *Proc. IEEE World Haptics Conf.*, pp. 213–218, 2017. 1, 2

[18] G. Karafotias et al. Mid-air tactile stimulation for pain distraction. *IEEE Trans. Haptics*, 11(2):185–191, 2017. 1, 2

[19] J. J. LaViola Jr et al. *3D user interfaces: theory and practice*. Addison-Wesley, 2017. 2

[20] S. J. Lederman et al. Extracting object properties through haptic exploration. *Acta psychologica*, 84(1):29–40, 1993. 2

[21] B. Long et al. Rendering volumetric haptic shapes in mid-air using ultrasound. *ACM Trans. Graph.*, 33(6):1–10, 2014. 7

[22] J. M. Loomis et al. Immersive virtual environment technology as a basic research tool in psychology. *Behav. Res. Meth. Instrum. Comput.*, 31(4):557–564, 1999. 2

[23] M. Marchal et al. Can stiffness sensations be rendered in virtual reality using mid-air ultrasound haptic technologies? In *Proc. Eurohaptics Conf.*, pp. 297–306. Springer, 2020. 1, 2

[24] J. Martinez et al. Touchless haptic feedback for supernatural vr experiences. In *Proc. IEEE VR Conf.*, pp. 629–630. IEEE, 2018. 1, 2

[25] J. Martinez et al. Mid-air haptic algorithms for rendering 3d shapes. In *2019 IEEE International Symposium on Haptic, Audio and Visual*

Environments and Games (HAVE), pp. 1–6. IEEE, 2019. 2, 3, 7

[26] A. Matsubayashi et al. Direct finger manipulation of 3d object image with ultrasound haptic feedback. In *Proc. CHI Conf. Hum. Factors Comput. Syst.*, pp. 1–11, 2019. 2, 3, 7

[27] S. H. McAmis et al. Simultaneous perception of forces and motions using bimanual interactions. *IEEE Trans. Haptics*, 5(3):220–230, 2012. 10

[28] M. A. Otaduy et al. Haptic technologies for direct touch in virtual reality. In *ACM SIGGRAPH Courses*, pp. 1–123. 2016. 2

[29] M. A. Plaisier et al. Two hands perceive better than one. In *Proc. Eurohaptics Conf.*, pp. 127–132. Springer, 2012. 2

[30] I. Rakkolainen et al. A survey of mid-air ultrasound haptics and its applications. *IEEE Trans. Haptics*, 14(1):2–19, 2021. 1, 3

[31] M. Sato. Spidar and virtual reality. In *Proc. 5th Biannual World Automation Congress*, vol. 13, pp. 17–23, 2002. 2

[32] Y. Singhal et al. Mid-air thermo-tactile feedback using ultrasound haptic display. In *Proc. ACM Symp. Virt. Real. Softw. & Tech.*, pp. 1–11, 2021. 1, 2

[33] C. Sun et al. Tactile sensitivity in ultrasonic haptics: Do different parts of hand and different rendering methods have an impact on perceptual threshold? *Virtual Reality & Intelligent Hardware*, 1(3):265–275, 2019. 3

[34] A. Talvas et al. A survey on bimanual haptic interaction. *IEEE Trans. Haptics*, 2014. 1, 2, 7

[35] G. Wilson et al. Object manipulation in virtual reality under increasing levels of translational gain. In *Proc. CHI Conf. Hum. Fact. Comput. Syst.*, pp. 1–13, 2018. 2