STAG: Smart Tools and Applications in Graphics (2017)
A. Giachetti, P. Pingi and F. Stanco (Editors)

Single-Handed vs. Two Handed Manipulation in Virtual Reality: A Novel Metaphor and Experimental Comparisons

F. M. Caputo, M. Emporio, A. Giachetti

Department of Computer Science, University of Verona, Italy

Abstract

In this paper we present a novel solution for single-handed deviceless object manipulation (e.g. picking/translating, rotating and scaling) in immersive visualization environments. The new method is based on degree of freedom (DOF) separation and on the idea of activating unambiguous gesture recognition when the hand is close to the object, giving visual feedback about gesture realization and available transitions. Furthermore, it introduces a novel metaphor, the "knob", to map hand rotation onto object rotation around selected axes.

The solution was tested with users on a classical visualization task related to finding a point of interest in a 3D object and compared with the well known "Handlebar" metaphor.

The metaphor shows a reasonable usability, even if not comparable with the bi- manual solution, particularly suited for the tested task. However, given the relevant improvements with task repetitions and the technical issues that can be solved improving the performances, the method seems to be a viable solution for deviceless single-hand manipulation.

CCS Concepts

•***Human-centered computing*** → ***Gestural input;***

1. Introduction

The recent technological advancements of devices suited for Virtual Reality (VR) experiences such as Head-Mounted Displays (HMD) and different kinds of input devices ranging from wireless controllers to depth sensors lead to the development of a substantial number of VR applications, some already sold as products, for different uses such as gaming, simulation, scientific visualization, etc.

So far, the use of mid-air manipulation techniques, to interact with Virtual Environments (VE) hasn't been a very popular design choice due to the many challenges still standing and preventing both the interaction and the overall user experience to be perceived as "natural".

The manipulation of virtual objects appears particularly challenging with the use of hand trackers (e.g. Leap Motion, Intel RealSense, etc.). These devices and related APIs provide real time hand and finger position mapping and some gesture recognition capabilities, that, even if not always stable and accurate due to occlusions and noise, can be directly used to build gesture-based interface allowing "direct" object manipulation similarly to the finger-based manipulation on touch-screens.

However, even basic actions such as grabbing, translating, rotating and scaling become non-trivial when trying to combine learnabilty, high-accuracy, robustness and easiness of use in a single technique based on these devices.

Simple and widely accepted interaction techniques are still missing and, therefore, intuitive and easy to use applications are difficult to find.

In this paper we propose a novel method to perform object manipulation in VE, separating control over 9 DOFs in different subtasks (translation, rotation, scaling) with each of them featuring a specific design to enhance general usability and precision. To cope with the most challenging subtask, e.g. rotation, we then propose a novel metaphor (knob metaphor) providing intuitive control with a single hand.

We compared the performance of the proposed method with the well known bi-manual solution based on the the Handlebar metaphor [SGH*12, BGG*07], that is probably the most popular deviceless manipulation metaphor and has been recently applied in several VR applications. The main limit of the Handlebar method is that it requires two hands and large motions so that is not suitable for generic applications.

The main contributions of our work are:

- A novel solution for deviceless and single-hand manipulation without the need of semaphoric gestures or other input to select a specific mode. Its implementation satisfies the desired re-

DOI: 10.2312/stag.20171225

quirement of rotation separation (and also axis decomposition) suggested in [BMA*14].
- Novel feedback ideas, e.g. release anchors to simplify object drop in the expected position after translation
- A novel metaphor (the knob metaphor) to perform arbitrary 3D rotations with simple and natural actions

In the following sections will describe the related work, the proposed solution and the design and implementation of the experimental test with the relative results.

2. Related Work

There are several works in the literature that cover mid-air manipulation in various form. Some papers feature analysis and comparison tests between different manipulation techniques, providing performance results and often reports of users' preferences. Unfortunately, due to the variety of different setups and different devices used for tracking and control, it is still difficult to derive design guidelines and further research is necessary to understand which are the best options for specific Virtual Environments.

Mendes et al. [MFA*14] compared different manipulation methods using physical devices and a 2D touch based approach and mid-air bi-manual methods using two techniques. The first is 6DOF Hand, that lets the dominant hand (DH) take control of 6 DOFs (i.e. 3 DOFs for translation and 3 DOFs for rotation) while the other hand enables control of the 3 DOFs about scaling, with the distance from the DH. The second is 3DOF Hand, in which rotation control is moved over from the DH to the non-DH. Vuibert et al. [VSC15] also made a comparison of different methods including physically constrained devices, a wand-like device, tangible object replica and finger-based interaction. An extensive work about direct mid-air manipulation is presented in [Cha14]. Here are analyzed different aspects of the presented manipulation techniques: the techniques themselves, using different kind of inputs such as a 6 degree of freedom flystick, a finger-tracking system, grasp/release heuristics, finite state machine for object manipulation, handling of complex hands/fingers interactions with objects.

Caputo et al. [CG15] features comparison of mid-air manipulation techniques with a hand tracking device and comparing successful techniques already present in literature along with novel ones. The work includes performance and user preferences data, collected during the test.

These comparison works provide enough data to make deviceless interaction look promising. In fact despite the still standing issues related to accuracy, noise, occlusion, hand fatigue, etc. bound to the use of hand tracking sensors, performance are still somewhat good and more importantly the results showed a tangible user preference towards deviceless setups and techniques. Such approach to mid-air manipulation becomes even more important when considering immersive virtual environments (VE) as the use of head-mounted displays (HMD) could arise more issues when dealing with physical objects as a mean of input or in other kind of setups like autosteroscopic monitors or light field monitors where the interaction device as a whole might only rely on hand tracking [ASSJ15].

This branch of works and techniques revolves around the idea of "natural" interaction and "natural" interfaces, that would ideally provide the best experience in VE for users, by offering a way of interaction greatly resembling the naturalness of interaction in the real world [BMR12]. The success of techniques aiming for this kind of interaction is the determined by their ability to address not only the aforementioned issues about tracking devices, but also some intrinsic problems related to design choices about the implemented techniques such as detecting gestures start and end, object selection, gesture recognition and decomposition. More issues are to be addressed when using low cost trackers like the Leapmotion [BWR14], and also the difficulty of controlling hands in mid-air movements requires to deal with relevant limits like the selection accuracy and the impossibility of directly tracking wide rotations [NLB14].

Song et al. presented the handlebar metaphor [SGH*12], this bi-manual technique provides 9 DOFs control by anchoring the selected object to the middle point between user hands and applying translation directly by moving the anchoring point, rotation by revolving hands around it and scaling with distance between hands without any DOF separation mean. A virtual Handle with a Grabbing Metaphor has been proposed by Kim and Park [KP14], based on a bounding sphere around the selected object. and the creation of a local reference frame for the following transformation. The crank handle metaphor for rotation has been presented in [BMA*14], separating rotation from translation, and using selected primary axes and a circular hand gesture to perform rotation. The rotation mode is, however, started with a non-intuitive gesture. In the same paper another metaphor (Grasping object) derived from [HCC07] is also proposed. After a pointing-based object selection the object is anchored to the movement of the hand both for translation and rotation. Finger-point rotation, introduced in [CG15] is a uni-manual technique that uses constrained 2D rotations projected on a selected plane using gesture recognition (i.e. a single extended finger) for transformation start and stop. In [SS*11], rotation around the center of mass and zooming is made with freehand gesture after a selection gesture and other gestures are used for panning and zooming although gesture transitions and accurate control are not covered.

In regard to uni-manual and bi-manual techniques and DOF separation, Bossavit et al [BMA*14] derived some useful guidelines highlighting the benefits of separating translation and rotation, of the decomposition of rotation in primary axes and of the use of metaphors and the use of sing-handed methods.

Further investigation in "natural" interaction with focus on manipulation is to be expected in future works. Many issues are still open and the provided solutions up till now are not always sufficiently effective to consider a technique usable in real applications or when the situation requires different hardware setups in both input device type or cost. In this work we focus on some specific issues by presenting some novel technique as a possible solution and present data about comparison results.

3. The Proposed Solution

Our goal was to create as simple as possible solutions to allow an intuitive and easy manipulation of objects, given hand and finger information captured by a cheap sensor (e.g. Leapmotion) placed

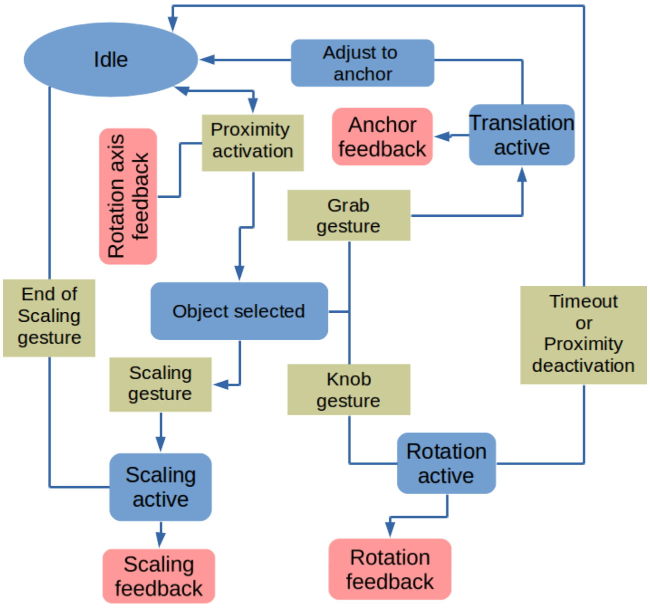

Figure 1: *Block diagram of the proposed interaction method. Blue rounded corner rectangles represent system states, while green rectangles indicate hand gestures triggering transitions.*

on the Head Mounted display when using an immersive VR application.

The interaction system acquires the data stream provided by the tracker and processes hand/finger trajectories to determine internal state changes, activate manipulation modes and feedback visualization.

Following the hints coming from the literature, we aimed at:

- performing the manipulation control with a single hand
- separating translation from rotation and constrain rotation axes
- being robust against limit of tracking and gesture segmentation accuracy provided by the chosen device/API, making difficult to control object release in the desired position
- finding intuitive metaphors for rotation and scaling
- avoiding as much as possible the necessity of non-intuitive gestures or use of different modalities to switch interaction modes

The scheme of the designed interaction solution is represented in Figure 1. The system starts in idle state, when hands are tracked and their position is displayed in the scene but no interaction with the objects is activated. However, if the hand falls in the region of interaction with an object, specific gestures can trigger the switch to different modes. Furthermore, if the hand is in the right position to start a rotation gesture on the object, a slightly different state is enabled, giving also visual feedback about the enabled rotation axis. In this state a knob rotation gestures starts the rotation as described in the following subsections, locking a rotation direction. Translation and scaling can be activated as well, when the hand is in the interaction region, with unambiguous gestures.

In this way basic manipulation modes are completely decoupled, but can be activated easily with a single hand gesture. In the following we described the proposed solutions to manage the specific manipulation modes.

3.1. Translation: Release Anchor and Visual Feedback

If we separate translation from rotation and scaling, the natural solution to grab objects and translate them is to check for hand closure and then put the system in translation mode, anchoring the displacement of the virtual object to the displacement of the hand (palm). This method is intuitive and also clearly outperformed other methods in many experiments, e.g. in docking tasks [BMA*14, CG15]. The only problem with this method, not fully addressed in the literature, is how to deal with object release in uni-manual, deviceless interaction.

The grab detection is currently realized with a simple thresholding $t = 0.5$ on the "grab strength" value, accessible through the Leapmotion API, when the hand is superimposed to the object. This choice provides reasonable accuracy, but it is not free from missed and false detections, so that further investigations will surely be performed to improve the usability of the system.

The biggest issue with this grab/release mechanism is that, if we recognize a simple finger gesture to exit from translation mode (e.g. opposite thresholding of the strength value), we could have a large displacement from the palm position before the release gesture and the tracked. A drop position detector based on trajectory speed analysis can reduce this problem giving a batter localization, however, it could create other problems: the visual feedback is not synchronized with the internal object position status and this is particularly bad when the user changes idea and, after stopping the translation, he does not release the grabbing and starts again translating the object.

For this reason we not only designed a translation stop/restart detector based on hand speed, but also a specific feedback to show, in this case the detected release position (visual anchor) and the current object with standard rendering.

The release anchor detection and update is realized as follows: once the translation mode has been initiated by the user, the system will constantly try to detect and guess the intended drop point. If the velocity is above a small threshold v_t, the release anchor is constantly moved constrained to the detected palm position at time $t - dt$ where dt is a speed dependent value avoiding unwanted displacements due to the release gesture. When the speed is under the threshold v_t the detected drop point is fixed in the position of the palm at the threshold crossing time. In this case a semi-transparent and monochromatic replica of the object is displayed (visual anchor) as showed in Figure 2, so that the user can see where the real one would snap to in space. This way the user can decide if they want to confirm the drop position by releasing there the object, by simply interrupting the grab gesture, or keep moving the object. If the object is not released, the anchor is kept for a fixed time, then the tracking is restored. If the velocity remains slow, the visual feedback of the release anchor position is maintained for some time also to allow a more accurate positioning in a feedback loop.

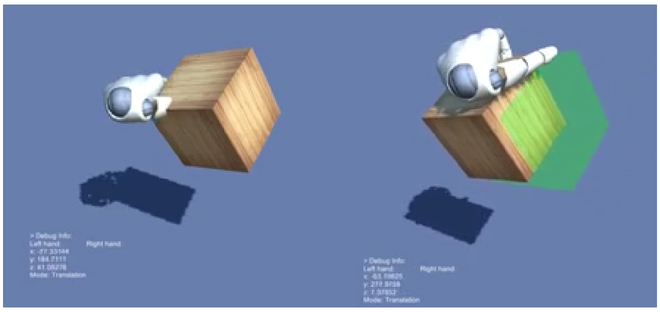

Figure 2: *The translation action can be started with a simple grab gesture (left side). The release anchor location is showed through a visual feedback (right side).*

3.2. Rotation: the Knob Metaphor

To perform rotation we introduced a novel metaphor that we called "knob" metaphor. The most typical uni-manual rotation gesture, in fact, is the action we do when turning a volume control. A similar single-handed gesture emerged also as the most intuitive rotation gesture in the elicitation study by Aigner et al. [AWB*12], with a clear preference over bi-manual gestures or non-iconic ones.

The idea is to map this gesture on the rotation of the object (selected by proximity when the gesture is started), automatically selecting a constrained rotation axis on the basis of the palm orientation. To make the system simpler to understand we force rotation to be around the z-axis or x-axis. As in the real word interaction, users can then perform multiple rotations of the knob in sequence to obtain large rotations. To avoid unwanted forward and back rotations, the system is locked in a specific rotation direction after the gesture starts, and exits from the lock only after a fixed time step passes without further rotation.

This method allows the realization of arbitrary rotations with sequences of x and z rotations. Despite the fact that in specific cases in which the user would like to rotate exactly about the y-axis, it can take a fairly long sequence of x/z rotations we believe the fluency provided by this rotation mechanism is enough to compensate for the extra number of steps needed to obtain the desired pose.

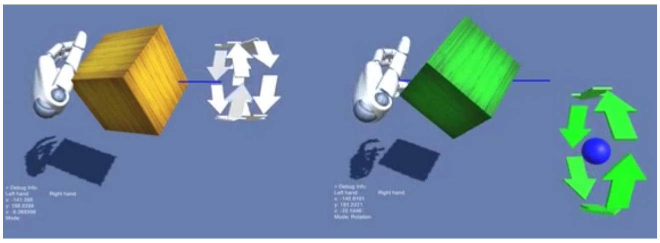

Figure 3: *The visual feedback for rotation shows the possible directions before the knob gesture is applied (left side). A visual feedback is provided showing the current locked direction of rotation (right side).*

3.3. Rotation Feedback

A further help for the understanding of the interaction mechanism without training is given by the visual feedback of potential rotation (Figure 3). This is designed as follows: when the system is in idle state and the user's hand enters the object's selection range with the gesture associated with the knob rotation method, a double set of arrows rotating around the selected axis is displayed. This notifies the user that both directions for rotation are available. Upon rotating the hand, one of the two directions is confirmed and the system will only keep the corresponding set of arrows and discard the other one to make clear which direction the user has locked onto.

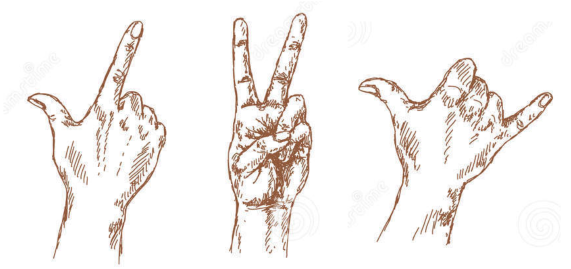

Figure 4: *Equivalent hand poses starting scaling mode when hand is over the object. Temporal coherence avoids unwanted manipulation due to sensor errors.*

3.4. Scaling: Iconic Gesture

The solution found to perform object scaling without using the second hand consists of using a reasonably intuitive and non ambiguous hand pose that the system will recognize in order to switch to the scaling mode and then the actual scaling is mapped on horizontal hand translations while the recognized hand pose is kept. As we want to drive the selection of rotation/translation/scaling with just one hand over the object, the gesture should not be similar to object grabbing. We opted for a gesture consisting of putting two fingers straight and extended. This choice is due to the sufficiently low stress deriving from the pose [BRR*04], the possibility of realizing it in different ways (see Figure 4, allowing an easy user adaptation. To avoid ambiguities and problems due to finger tracking errors, to switch to the scaling mode, we require that the fingers' position is maintained for at leas 1 second before switching mode. After this time, horizontal motion is mapped into scale: moving the hand to the right the object is enlarges, moving to the left it is downscaled. Visual feedback of the scaling activation is provided by slightly increasing the brightness of the selected object color. To exit this mode it is sufficient to change hand pose.

4. Implementation and Results

We implemented versions of the interaction solutions proposed in Virtual Reality environments realized with Unity 5.3 (`www.unity3d.com`), using the Leapmotion (`www.leapmotion.com`) controller and Orion beta API along with the OVR plug-in (`developer.oculus.com`) in order to enable the use of a HMD for stereoscopic display through an Oculus Rift DK2. Since the

the task presented to our test users is designed to be performed on a workbench-like virtual scene, we opted for a fixed position of the Leapmotion controller on the desk, so that its detection volume could include all the relevant space required by the user's hands for the task completion.

4.1. Software Application and Task

The whole application used for the test was built with Unity 5.3 (`www.unity3d.com`) framework. Our application features the full implementation of our Knob metaphor along with a Handlebar implementation, following the description in the original papers [BGG*07, SGH*12] regarding the single-object manipulation. Besides the techniques, the application also features a task used to evaluate and compare users' performances. We designed an information searching task, with the idea of stressing on both accuracy and interaction fluency. The idea was to have users performing the same task multiple times in series and the task can be summarized in two main parts:

- **I**nformation retrieval: A colored cube is presented at the center of the scene between 4 numbered tiles that are used later for the "information check" sub-task (Figure 5a). A different texture is loaded on the cube from a pool of 4 with each run. All the 4 textures in the pool look exactly the same besides a small detail, and are designed to color each of the cube faces cube with the same tint. They only variant between them is a small digit, ranging from 1 to 4, and appearing on a different face for each texture. Once the test begins, the user is asked to use all the necessary transformations actions available for the current mode (i.e. knob metaphor or handlebar) to manipulate the cube in order search for the small number. The initial position of the cube is set so that at least a rotation and a scaling action (to magnify the number) are required to identify the digit value.
- **I**nformation Check: Once the value is discovered the user has to position the cube on one of the 4 tiles, also numbered from 1 to 4. If the number on the tile matches the number on the cube texture, the task is successfully completed (Figure 5b). The correct tile will also change its color from white to green once cube is placed, but only after it has been released, in order to prevent the user from getting unwanted hints on the solution by simply hovering the cube across the tiles. Furthermore, to avoid accidental activation of the tiles during the execution of the first part of the task due to an excessive scaling of the cube, the collision detection is done relatively to the center of the cube rather than its bounding volume.

4.2. Experiment Overview

To perform the experiment, we selected a group of 24 subjects, with no experience of 3D interfaces and VR in general. Each user performed the described task in 2 series of 4 runs each. Half the subjects started with the Knob metaphor as their first technique while the second half started with the Handlebar. For both techniques, regardless of the order, all users were given a 5 minutes time to familiarize with the manipulation of a neutral object (i.e. a white cube) with no specific task to perform and a small break in between the two series. To further remove biases, each user, for each technique, performed the 4 runs going through all the 4 different textures. The order in which they were assigned to the runs was based on a latin square design.

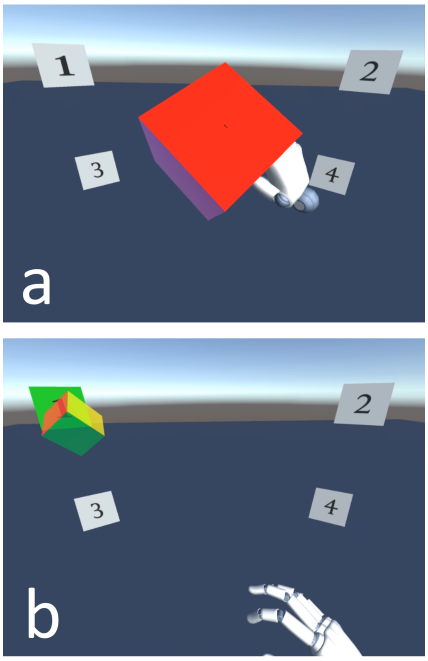

Figure 5: *Task sequence: a) Information retrieval, a small number "1" is present on the red face. b) Information check, the cube is correctly placed on the tile labeled with the corresponding number.*

During the experiments the application developed for the task automatically collected data about execution performance such as the execution, the time spent in idle, and for the Knob metaphor, the whole sequence of transitions between actions with associated time stamps. This recording was not possible with enough reliability for the Handlebar technique due to overall uncertainty of the action in progress derived by not having separated DOFs.

At the end of the two series each user provided further data with questionnaires surveying their test experience. In particular we collected data for both techniques about: learnability, as in how easy/hard was to understand the interaction mechanism, usability in terms of how practical they were in order to complete the task and gestural comfort to estimate the perceived fatigue. The questions were asked so that, by using a Likert scale ranging from 1 to 5 the lower value was always considered the negative one for the judged parameters.

4.3. Results

We examined the data collected by the task application. For each run of each series/technique we averaged the times of all users, thus having 4 averages for each techniques. These results are shown in Figure 6. Users took on average 17.6 s to complete the task with the Handlebar method and 44.6 with the single-handed one. It is a large difference, but the trend suggested by the curve indicates that, while

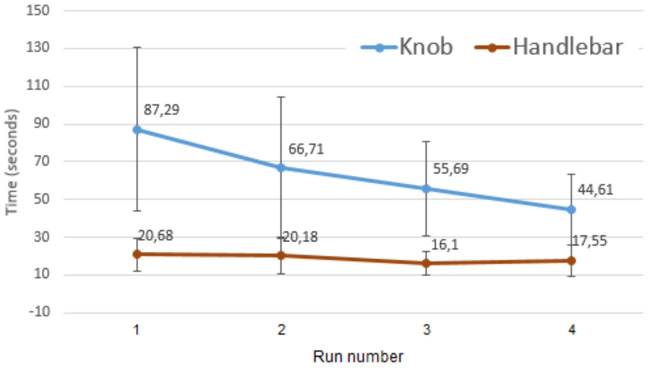

Figure 6: *Task completion times (Knob vs Handlebar)*

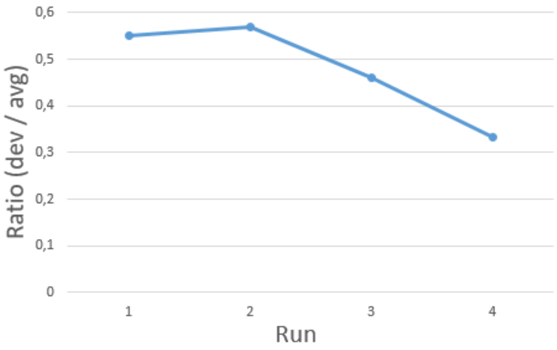

Figure 7: *Times averages deviations excluding idle time.*

in successive trials the handlebar performances are approximately constant, there is a steep decrease in the task completion time with the single-handed method. This suggests that, while the handlebar is more intuitive and immediately learned, the single-handed manipulation proposed require a small training phase, but then could be probably similarly efficient. This argument is further supported by the decrease of deviation in subsequent runs. In fact, by considering completion times excluding the "idle" time segments of each run, the averages of times decrease uniformly across the runs of both technique of a factor of 0.6, but more interestingly the variability of data lowers with subsequent runs (Figure 7), suggesting that not only the performances of the Knob method improve with more runs but it improves overall across all the subjects. This effect might be due to the difference in learning speed between users but, nonetheless, data suggest a certain convergence towards improved performances regardless of that.

An interesting observation can be made on the task completion recordings is that user did not follow a fixed sequence of actions on the different manipulation DOFs, as it could be expected. The action sequence could in principle have spitted in a single rotation and a single scaling for the information retrieval part, followed by a final translation for the information check part.

Users performing the task with the Knob technique did not act in this way and performed on average sequences of 15.9 different actions (excluding idle) and in different orders to complete the task, with an average deviation of 5.6.

This means that they did not take full advantage of the DOF separation and that the improvements of the performances with repeated trials could have been obtained with an optimization of the sequences.

Regarding questionnaires data, we report the averages of answers to questions previously described, in Figure 8. Despite the outcome not being very encouraging for the Knob metaphor we can still notice, based on the Likert scale ranging from 1 to 5 that our method hasn't received overall negative scores. This leads us to believe that the method is still viable and could probably be better evaluated when applied to tasks taking advantage of the free hand to perform other actions simultaneously, compared against bi-manual techniques like Handlebar, requiring such actions to be performed in sequence.

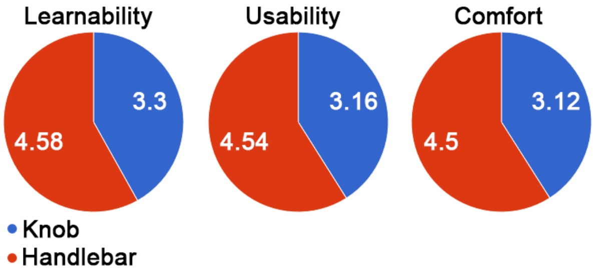

Figure 8: *Questionnaires answers regarding the evaluation of techniques.*

5. Discussion

We presented a novel technique trying to provide a viable unimanual technique to perform mid-air manipulation in immersive VE (i.e. the Knob metaphor). To ensure the method provides enough accuracy for generic tasks we designed the interaction with minimal interferences between the different actions through DOFs separation. We evaluated the performance of our technique by comparing against the Handlebar metaphor its performance in a information searching task and its viability through evaluation questionnaires. Our findings lead to believe that the proposed technique is sufficiently usable, even if, without training, it is not as effective as the bi-manual manipulation with the Handlebar metaphor in classic manipulation tasks. However, experimental results also shows that the difference in performance clearly scales down as the users become more accustomed to the Knob technique. Considering these results it's reasonable to believe that not only the Knob technique is at least a viable technique in general, but also that it can be a method of choice for all the applications combining manipulation with other actions potentially performed with the non-dominant hand.

Further research is certainly needed to optimize our single-handed manipulation method and evaluating its benefits. For this reason we plan to continue our activity following two main directions. The first is related to the improvement of the proposed techniques allowing more efficient manipulation selection and mode transition. The use of specific widgets will be investigated for this purpose and compared with the current gesture-based selection. The second one is to develop more complex testbeds mimicking real-world applications where the necessity of performing multiple tasks can more clearly show the advantages of uni-manual methods.

References

[ASSJ15] ADHIKARLA V. K., SODNIK J., SZOLGAY P., JAKUS G.: Exploring direct 3d interaction for full horizontal parallax light field displays using leap motion controller. *Sensors 15*, 4 (2015), 8642–8663. 2

[AWB*12] AIGNER R., WIGDOR D., BENKO H., HALLER M., LINDBAUER D., ION A., ZHAO S., KOH J.: Understanding mid-air hand gestures: A study of human preferences in usage of gesture types for hci. *Microsoft Research TechReport MSR-TR-2012-111* (2012). 4

[BGG*07] BETTIO F., GIACHETTI A., GOBBETTI E., MARTON F., PINTORE G.: A practical vision based approach to unencumbered direct spatial manipulation in virtual worlds. In *Eurographics Italian Chapter Conference* (2007), pp. 145–150. 1, 5

[BMA*14] BOSSAVIT B., MARZO A., ARDAIZ O., CERIO L. D. D., PINA A.: Design choices and their implications for 3d mid-air manipulation techniques. *Presence 23*, 4 (Nov 2014), 377–392. `doi:10.1162/PRES_a_00207`. 2, 3

[BMR12] BOWMAN D. A., MCMAHAN R. P., RAGAN E. D.: Questioning naturalism in 3d user interfaces. *Communications of the ACM 55*, 9 (2012), 78–88. 2

[BRR*04] BOUSSEMART Y., RIOUX F., RUDZICZ F., WOZNIEWSKI M., COOPERSTOCK J. R.: A framework for 3d visualisation and manipulation in an immersive space using an untethered bimanual gestural interface. In *Proceedings of the ACM symposium on Virtual reality software and technology* (2004), ACM, pp. 162–165. 4

[BWR14] BACHMANN D., WEICHERT F., RINKENAUER G.: Evaluation of the leap motion controller as a new contact-free pointing device. *Sensors 15*, 1 (2014), 214–233. 2

[CG15] CAPUTO F. M., GIACHETTI A.: Evaluation of basic object manipulation modes for low-cost immersive virtual reality. In *CHITALY* (2015). 2, 3

[Cha14] CHAPOULIE E.: *Gestures and direct manipulation for immersive virtual reality*. PhD thesis, Université Nice Sophia Antipolis, 2014. 2

[HCC07] HANCOCK M., CARPENDALE S., COCKBURN A.: Shallow-depth 3d interaction: design and evaluation of one-, two-and three-touch techniques. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (2007), ACM, pp. 1147–1156. 2

[KP14] KIM T., PARK J.: 3d object manipulation using virtual handles with a grabbing metaphor. *IEEE Computer Graphics and Applications*, 3 (2014), 30–38. 2

[MFA*14] MENDES D., FONSECA F., ARAUJO B., FERREIRA A., JORGE J.: Mid-air interactions above stereoscopic interactive tables. In *3D User Interfaces (3DUI), 2014 IEEE Symposium on* (2014), IEEE, pp. 3–10. 2

[NLB14] NABIYOUNI M., LAHA B., BOWMAN D. A.: Poster: Designing effective travel techniques with bare-hand interaction. In *3D User Interfaces (3DUI), 2014 IEEE Symposium on* (2014), IEEE, pp. 139–140. 2

[SGH*12] SONG P., GOH W. B., HUTAMA W., FU C.-W., LIU X.: A handle bar metaphor for virtual object manipulation with mid-air interaction. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems* (New York, NY, USA, 2012), CHI '12, ACM, pp. 1297–1306. URL: `http://doi.acm.org/10.1145/2208516.2208585`, `doi:10.1145/2208516.2208585`. 1, 2, 5

[SS*11] SORO A., SCATENI R., ET AL.: Natural exploration of 3d models. In *Proceedings of the 9th ACM SIGCHI Italian Chapter International Conference on Computer-Human Interaction: Facing Complexity* (2011), ACM, pp. 118–121. 2

[VSC15] VUIBERT V., STUERZLINGER W., COOPERSTOCK J. R.: Evaluation of docking task performance using mid-air interaction techniques. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction* (2015), ACM, pp. 44–52. 2