



HAL
open science

Toward Universal Tangible Objects: Optimizing Haptic Pinching Sensations in 3D Interaction

Xavier de Tinguy, Claudio Pacchierotti, Maud Marchal, Anatole Lécuyer

► **To cite this version:**

Xavier de Tinguy, Claudio Pacchierotti, Maud Marchal, Anatole Lécuyer. Toward Universal Tangible Objects: Optimizing Haptic Pinching Sensations in 3D Interaction. VR 2019 - 26th IEEE Conference on Virtual Reality and 3D User Interfaces, Mar 2019, Osaka, Japan. pp.321-330, 10.1109/VR.2019.8798205 . hal-02021319

HAL Id: hal-02021319

<https://inria.hal.science/hal-02021319v1>

Submitted on 18 Feb 2019

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

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

Toward Universal Tangible Objects: Optimizing Haptic Pinching Sensations in 3D Interaction*

Xavier de Tinguy[†]

Univ Rennes, INSA, IRISA, Inria,
CNRS, Rennes, France

Claudio Pacchierotti*

CNRS, Univ Rennes, Inria, IRISA
Rennes, France

Maud Marchal*

Univ Rennes, INSA, IRISA, Inria,
CNRS, Rennes, France

Anatole Lécuyer*

Univ Rennes, Inria, CNRS, IRISA
Rennes, France

ABSTRACT

Tangible objects are a simple yet effective way for providing haptic sensations in Virtual Reality. For achieving a compelling illusion, there should be a good correspondence between what users see in the virtual environment and what they touch in the real world. The haptic features of the tangible object should indeed match those of the corresponding virtual one in terms of, e.g., size, local shape, mass, texture. A straightforward solution is to create perfect tangible replicas of *all* the virtual objects in the scene. However, this is often neither feasible nor desirable. This paper presents an innovative approach enabling the use of few tangible objects to render many virtual ones. The proposed algorithm analyzes the available tangible and virtual objects to find the best grasps in terms of matching haptic sensations. It starts by identifying several suitable pinching poses on the considered tangible and virtual objects. Then, for each pose, it evaluates a series of haptically-salient characteristics. Next, it identifies the two most similar pinching poses according to these metrics, one on the tangible and one on the virtual object. Finally, it highlights the chosen pinching pose, which provides the best matching sensation between what users see and touch. The effectiveness of our approach is evaluated through a user study. Results show that the algorithm is able to well combine several haptically-salient object features to find convincing pinches between the given tangible and virtual objects.

Index Terms: Human-centered computing—Human computer interaction—Interaction devices—Haptic devices

1 INTRODUCTION

Haptic sensations can greatly improve the immersiveness and illusion of presence during Virtual Reality (VR) and Augmented Reality (AR) experiences. There are many different ways to provide users with such sensations, including using actuated devices known as kinesthetic or cutaneous feedback devices [22, 35, 44], exploiting perceptual phenomena with cross-modal effects [4, 29], or using passive props also known as tangible objects [5, 9, 25, 37].

In this respect, tangible objects, also known as passive haptic or physical props, have been proven to be an effective - yet simple - solution for conveying global and distributed shape sensations in VR/AR. However, for the illusion to work, the haptic features of the tangibles should match those of the corresponding virtual objects, i.e., there should be a good correspondence between the, e.g., size, local shape, texture, mass of the tangible object with respect to its virtual counterpart. One way to match what users see in the virtual environment to what they touch in the real world is creating tangible versions of all the virtual objects in the scene. However, this solution may be neither feasible, e.g., if users can modify the virtual objects, nor desirable, e.g., if the virtual environment comprises several

objects. It is therefore important to study solutions enabling the use of few tangible objects to render many virtual ones.

Toward this objective, we present an algorithm which analyzes different tangible and virtual objects to find the grasping strategy best matching the resultant haptic pinching sensation. Starting from the meshes of the considered objects, the algorithm guides users towards the grasping pose which best matches what they see in the virtual scene with what they feel when touching the tangible object. By selecting different grasping positions according to the virtual object to render, it is possible to use few tangible objects to render multiple virtual ones (see this idea in Fig. 1). The main contributions of our work can be summarized as follows:

- a novel algorithmic approach for using few tangibles to render multiple virtual objects, maximizing haptic fidelity;
- a preliminary user study assessing the effectiveness of the proposed approach for a set of representative objects.




Figure 1: Illustration of our approach through a carousel of virtual objects that can be grasped using a single “universal” tangible object. The user is able to turn the virtual carousel and manipulate the three virtual objects using the suggested pinch poses (in green). These poses are proposed by our algorithm to best match the corresponding haptic pinching sensations on the tangible object.

2 RELATED WORK

Mesh matching. Matching two meshes over their geometry [46] or topology [51] is a common task in mesh processing, and it is usually performed by extracting representative features that can be compared. Shapira et al. [46] address the problems of partitioning and skeleton extraction of a set of 3D meshes. They use a volume-based shape-function mostly invariant to pose changes of the same object and that retains a similar value in analogue parts of different objects. It uses a measure of the diameter of the object’s volume in the neighborhood of each mesh point. Using this approach, it is possible to partition and create skeletons among multiple meshes. Kortgen et al. [27] present a method to evaluate 3D shape similarity. It captures the distribution over relative positions of other shape points to match different global shapes, even if they possess local differences, topological asperity, or a different scale. Objects can also

*This research has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 801413 for project “H-Reality.”

[†]e-mail: {xavier.de-tinguy, claudio.pacchierotti, maud.marchal, anatole.lecuyer}@irisa.fr.

be segmented by evaluating their similarity to basic shapes. A common approach to locally fit primitives has been presented by Zhou et al. [56], who introduce a quantitative measure of “cylindricity” for a shape and develop an optimization algorithm for generalized cylinder decomposition. Similarly, Yan et al. [54] segment meshes by fitting quadric surfaces, while Vidal et al. [52] present a method for segmenting 2-manifold meshes by decomposing them into local shape primitives such as planes, spheres, or cylinders. A similar approach has been also presented in [32,45]. Descriptors can also be obtained by evaluating the surface properties. In particular, Cazals et al. [8] propose a method to compute mesh differential properties by fitting a polynomial to a neighborhood of points for extracting, e.g., curvature directions. This method is implemented in the Computational Geometry Algorithms Library, CGAL¹. Similarly, Szilvasi et al. [49] propose a method for the construction of quadratic polynomial B-spline patches fitting analytical surfaces and triangle meshes in a neighborhood of points. These simpler patches can then be used for re-triangulation, and mesh simplification.

Grasping and manipulation. Finding feasible grasping poses on known and unknown objects is a common problem in robotics [42], where researchers want to use robotic hands to autonomously pick and place objects in unstructured environments. This issue is quite similar to ours, as we start our algorithm by finding feasible pinching poses on the available virtual and tangible objects. GraspIt! [36] is a popular software able to generate feasible grasps on an object. It starts from known grasps on the object shape’s primitives, which are then combined and tested in simulation for feasibility. More recently, researchers have focused their attention on finding feasible grasps from point cloud data [16, 19, 50]. For this purpose, ten Pas et al. [50] find grasps for a 2-fingered manipulator, without any prior knowledge of the environment, using only a point cloud obtained from a depth camera. Bonilla et al. [7] address the issue of planning grasps for soft hands. As soft hands can easily adapt to known basic shapes, the algorithm decomposes the target object into Minimal Volume Bounding Boxes proven to be graspable by the hand in use. Hand poses are then generated using geometric information extracted from such bounding boxes. Björkman et al. [6] present a probabilistic approach for learning object models based on visual and tactile perception through physical interaction with an object. This model can then be used to retrieve feasible grasp candidates. Generating suitable grasping poses is also a common problem for animation. For example, Li et al. [31] present a data-driven approach to grasp synthesis, starting from a database of captured human grasps. To identify grasp candidates on a new object, they match hand and object shape by identifying features having similar relative placements and surface normals. Similarly, Lui et al. [33] describe a physics-based technique taking an initial grasping pose and partial object trajectory as input, and generating believable hand grasplings as output. To address unexpected dynamic situations during manipulation, the algorithm generates a range of possible hand configurations, taking into account changes in joint configurations, contact points, and grasping force.

Tangibles within Virtual and Mixed Reality. Interacting with tangible objects has been proven to improve the immersiveness of VR and AR systems [5,21,24]. Billinghurst et al. [5] present an AR tangible book. Users can turn its pages, look at the pictures, and read the text as in a normal book. However, if they use an AR display, they can also see 3-dimensional virtual models popping out of the pages. Users can then change the models by physically turning the book pages. Cheok et al. [9] develop a mixed-reality game carried out in a real environment. Players walk around a large room-size area and pick up real objects, as if they were playing a traditional non-virtual game. However, these real objects are augmented with superimposed virtual objects and figures, e.g., a real box is opened and inside it is found a virtual treasure. More recently, Harley et al. [21] present a system for diegetic tangible objects in VR. They

develop four tangible objects prototypes, including a cube, a stuffed animal, a treasure chest, and a wooden boat, providing passive and active haptics. For example, the stuffed animal is a furry, hollow raccoon toy. It has a skeletal cage to give structural integrity to its body and a heartbeat that can be either calm or accelerated. Finally, Dalsgaard and Halskov [13] combine tangible tabletop interaction with 3D projection so that any tangible on a table can be augmented with visual content corresponding to its physical shape, position, and orientation.

As already mentioned, many work in the literature have argued that a mismatch between virtual and physical objects may severely affect the user’s experience and illusion of presence [23, 28, 47]. Thus, it is important to well match what users see in the display and what they feel when interacting with the tangible object. In this respect, Hettiarachchi and Wigdor [24] introduce a wearable system able to scan the surroundings of the user and find a matching between the available virtual and tangible objects. However, it only takes into account the objects’ size, primitive shape, and orientation. HaptoBend [34] is a self-reconfigurable device able to change its shape to match the target virtual object. A similar result has been also achieved using small robots which can assemble into different shapes [55]. Sait et al. [20,43] map a single tangible prop to multiple virtual objects distributed throughout a virtual environment. They combine a redirection approach, that uses rotational adjustments to align the user during virtual locomotion, and a resetting approach, that introduces a discrete rotational update when the user virtually approaches a target for interaction. Also Azmandian et al. [3] use a single tangible prop to provide passive haptics for multiple virtual objects. They achieve this result by dynamically and gradually aligning tangible and virtual objects as the user interacts with the environment. Finally, Eckstein et al. [14] create multiple virtual version of the same real room, with a varying degree of mismatch. They analyze the influence of these environments on the perceived sense of presence and motivation of users.

Our work lies along this line of research and it aims at improving haptic pinching sensations when manipulating a tangible object. With respect to ten Pas et al. [50] and Hettiarachchi and Wigdor [24], our approach focuses on the *local* haptic sensations felt by the user at the grasping location. Although we could retrieve the objects’ meshes in a similar fashion to [24, 50], the proposed approach to generate target grasping poses is substantially different. In this paper, we aim at matching a series of haptically-salient features at the users fingertips, without considering any global shape characteristic.

3 DESCRIPTION OF THE ALGORITHM

The objective of our algorithm is to provide users with coherent sensations, matching what they see on the display to what they touch in the real environment. Given a set of virtual and tangible objects’ meshes, the algorithm identifies two grasping/pinching poses, one on the virtual object and one on the tangible object, maximizing their similarity in terms of haptic perception. The computation is done during the pre-processing phase, and the generated grasping matches are used to provide guidance to the user in VR.

Starting from the meshes of all considered objects, our algorithm behaves in the following way. First, it identifies feasible 2-fingers pinching poses on the considered tangible and virtual objects. Then, for each pose and for each object, the algorithm evaluates a series of haptically-salient characteristics: (i) distance between the two fingertips, (ii) distance of the pinch from the center of mass, (iii) relative orientation between the two fingertips, (iv, v, vi, vii) local surface orientation and curvature. These features have been identified by the authors as a representative - yet non exhaustive - set of important grasping information. Considering a broader range of feature is of paramount importance and will be addressed in the near future (see Sec. 5). Next, the algorithm identifies the two most similar pinching poses, one on the tangible and one on the virtual object, evaluated in terms of the above metrics. Finally, the chosen pinching pose is highlighted in the virtual scene for the user to grasp. It guarantees

¹<https://www.cgal.org/>

the best match between what the user sees on the display and what he feels when grasping the tangible object. Of course, the perceptual result highly depends on the intrinsic characteristics of the available tangible and virtual objects. The algorithm is only able to find the best solution given the available objects in the scene, which is not guaranteed to lead to a good/believable interaction. Obviously, as expected, if a tangible replica of every virtual object is available, the algorithm will make the user grasp the virtual and tangible objects in the same place. A more detailed discussion about this point can be found in Sec. 5, together with an analysis of the current limitations and assumptions of the algorithm (e.g., we consider only 2-fingers pinches, we assume a standard size for the fingertip, we consider the meshes of all objects available).

The remaining part of the Section describes in detail the above-mentioned process, which is summarized by Algorithm 1.

Algorithm 1:

Input: objects' meshes
foreach *object* **do**
 Step #1: find all feasible pinching poses
 foreach *pinching pose* **do**
 Step #2: compute haptic characteristics
 Retrieve the two criteria 7-tuples $[\zeta_1 \ \zeta_2]$
Find all combinations of poses $[V, T]$ such that one pose is on a virtual object and one on a tangible object.
foreach *set of poses* $[V, T]$ **do**
 Step #3: choose the matching pinch
 Evaluate the error between current set of poses $\Delta\zeta$
 Evaluate mean error $\overline{\Delta\zeta} = \sum_{i=1}^{14} \alpha_i \frac{\Delta\zeta}{14}$
Choose the pose set $[V, T]$ minimizing $\overline{\Delta\zeta}$

3.1 Step #1: finding the set of feasible grasping poses

“A grasp is every static hand posture with which an object can be held securely with one hand [15].” In this work, as a starting point, we only consider thumb-index fingers precision grasps, as defined by Cutkosky [12]. We believe that considering a broader range of grasps is important, and it will be addressed in the near future (see Sec. 5).

As only two fingers are involved, we refer to this type of grasp as a “pinch,” and we then refer to the pose of the two fingertips during the grasp as a “pinching pose.” The contact area between each fingertip and the object’s surface is referred to as a “contact patch.” For simplicity, we consider circular patches, centered in a point $p \in \mathbb{R}^3$ on the mesh and including all points on the mesh closer than 8 mm. Once the considered tangible and virtual objects are identified and their mesh is retrieved, we use CGAL to compute the set of all feasible thumb-index fingers precision grasps. Each feasible grasp is composed of two fingertip poses, which in turn define two contact patches. Variables related to pinches on the tangible and virtual objects are indicated with T and V subscripts, respectively. Similarly, the two contact patches on the same object are indicated with 1 and 2 subscripts.

3.2 Step #2: computation of the haptic characteristics

To measure the haptic similarity of two pinching poses, for each contact patch, we evaluate seven representative haptically-salient characteristics, $\zeta = [w \ d \ \phi \ \psi \ \rho \ \kappa_M \ \kappa_m] \in \mathbb{R}^7$. They take into account both global ($w \ d \ \phi$) and local ($\psi \ \rho \ \kappa_M \ \kappa_m$) haptic information. As mentioned at the beginning of this section, we plan to consider more criteria in the near future. This study presents a proof-of-concept of our approach and its implementation.

Table 1 shows an example of each criterion. The blue triangles indicate the two pinching poses p_1, p_2 , the blue line the vector $\overrightarrow{p_1 p_2}$ connecting them, the red line the object’s first principal component axis, and the red dot its center of mass.

3.2.1 Criterion #1: pinch size

Hershberger and Misceo [23] found out that, between vision and haptic information, neither modality inherently dominates the perceived size of an object. Discordant haptic information biases visual size estimates by as much as discordant visual information biases haptic size estimates. Since neither modality captures the other one completely, a discordant stimulus will be perceived as discordant. Flanagan and Mandomir [17] also found that the width of the grasp affect the object’s perceived weight. Therefore, it is very important to well match the size of the tangible object with respect to its virtual counterpart.

To address this important point, we consider the perceived size of the object as our first criterion. For each pinching pose, we measure the distance w between the two contact patches.

3.2.2 Criterion #2: distance from the center of mass

Although an object’s weight can be already estimated when it rests on a stationary hand [30], its perceived weight and size when welded is determined by its resistance to rotational forces of the user’s limbs. Amazeen and Turvey [1] provides a physical measure of this property by means of the inertia tensor.

To take into account this property, we included in our criteria a measure of the distance between the pinching pose and the object’s center of mass. Specifically, we evaluate the distance d between the object’s center of mass and the projection of the center of vector $\overrightarrow{p_1 p_2}$ onto the principal component axis.

3.2.3 Criterion #3: pinch tilt

Proprioception provides humans with information about the relative position of their limbs as well as the effort being employed in movement. Therefore, although it is known that absolute position retargeting is possible [3], it is paramount to well match the relative position of the fingers during the grasp. Proprioceptive accuracy has been measured to be around 20 mm [11, 48], increasing with the length of the movement.

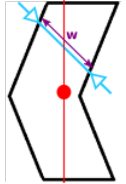
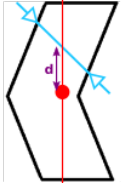
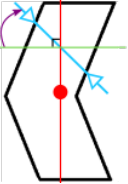

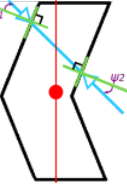
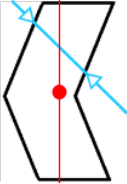
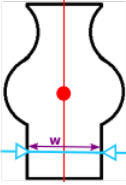
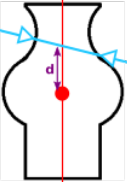
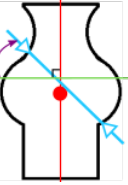
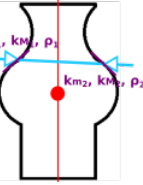

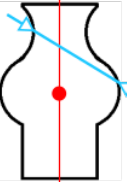
For this reason, we considered the relative orientation between the two pinching locations as our third criterion. Specifically, we measure the angle ϕ between vector $\overrightarrow{p_1 p_2}$ and its projection onto the plane orthogonal to the principal component axis.

3.2.4 Criteria #4,5,6,7: local surface orientation

Several studies in the literature suggest that the orientation of the contact surface is an important source of information for curvature perception [10, 38], a dominant source of information both in static [40] and dynamic [53] conditions. Gordon and Morison [18] found the 75% discrimination threshold of 2-cm-long curved objects to be approximately 2 m^{-1} . Pont et al. [40] found the 84% discrimination threshold of 20-cm-long curved objects to be approximately 0.5 m^{-1} , suggesting that the threshold may not depend directly on the curvature, but rather on the total variation of the surface orientation. In other works, Pont et al. [26, 39] tested static haptic discrimination of curved strips. Results show that the effective stimulus for curvature discrimination is the difference in local surface attitude.

An infinite number of curvatures exists at a given point on a 3-dimensional surface. Principal curvatures are defined as the maximum and minimum of the normal curvature, and they provide information about the maximum and minimum bending of a surface at one point. Similarly, we can also define the normal vector as the vector perpendicular to the normal curvature. As the fourth, fifth, sixth, and seventh criteria, we consider the maximum principal curvature κ_M ; the minimum principal curvature κ_m ; the angle ρ between the maximum principal curvature direction and the principal component projected onto the patch; and the angle ψ between the pinching vector $\overrightarrow{p_1 p_2}$ and the normal vector at the surface.

Table 1: Examples of how criteria are evaluated. The rows show a representative example of how each criterion is used to find a pinch match between two different shapes, i.e., one tangible and one virtual objects. Each column shows one criterion and how the matching pinch would be chosen if considering only that criterion. For each object, the blue triangles indicate the two pinching positions p_1, p_2 , the blue line the vector $\vec{p_1 p_2}$ connecting them, the red line the object’s first principal component, the red dot its center of mass, and the green line a perpendicular to the principal component.

	#1: Pinch size	#2: Distance from the center of mass	#3: Pinch tilt	#4,5,6: Local surface properties	#7: Local surface orientation	Combination of all the criteria
Tangible Object						
Virtual Object						

3.3 Step #3: choosing the best matching pinch

At the end of the previous step, each feasible pinching pose identified in Sec. 3.1 is assigned two 7-tuples $[\zeta_1 \zeta_2]$, one per contact patch. Each 7-tuple contains the seven criteria metrics discussed in Sec. 3.2. For example, if we found 50 feasible pinching poses on a given object, we will end up with 100 7-tuples (two per pose) and 700 metrics (fourteen per pose). However, it is important to notice that the first three criteria $[w \ d \ \phi]$ provide pose-specific information while the last three $[\psi \ \kappa_M \ \kappa_m]$ provide patch-specific information. Therefore, as we generate one 7-tuple per patch (and two per pinching pose), the first three components of $[\zeta_1 \zeta_2]$ will be the same, i.e., $[w_1 \ d_1 \ \phi_1] = [w_2 \ d_2 \ \phi_2]$.

To choose the best matching pinching pose between two objects, i.e., one virtual and one tangible, we perform an exhaustive search between all the feasible poses on the two objects. For each couple of poses, one per object, we calculate the error vector between the haptic criteria, i.e., $\Delta\zeta = [\zeta_{V,1} \ \zeta_{V,2}] - [\zeta_{T,1} \ \zeta_{T,2}] \in \mathbb{R}^{14}$. Referring to our example above, 50 feasible pinching poses on each object will result in $50 \times 50 = 2500$ $\Delta\zeta$ comparisons. To be able to consider all criteria at once and retrieve a single number describing the haptic similarity of two pinching poses in two different objects, we calculate a weighted sum of the error metric $\Delta\zeta$. Then, we compute $\overline{\Delta\zeta} = \sum_{i=1}^{14} \alpha_i \Delta\zeta_i^2$ for each comparison, where α_i are weighting factors. These weights can (and should) be adjusted according to their perceptual importance (see Sec. 5). In our implementation, for the sake of simplicity, we scale each element of $\Delta\zeta$ to have them varying in the same range. Completing our example above, 2500 $\Delta\zeta \in \mathbb{R}^{14}$ comparisons will lead to the same number of $\overline{\Delta\zeta} \in \mathbb{R}$. We are finally ready to choose our best matching pinch as the one showing the smallest composite error $\overline{\Delta\zeta}$.

Of course, we are not limited to choosing the best matching pose only between *two* objects. The proposed algorithm is quite flexible and can take into account a larger number of objects. Imagine an experimental setup composed of n tangible objects and m virtual ones, with $m \gg n$. Our algorithm can analyze all the scene at once and then find the best pinch for each virtual object using the available set of tangible ones.

4 USER STUDY

We conducted a preliminary user study to evaluate the user perception when he is manipulating a tangible object (TO) and visualizing

virtual objects (VO) that do not always match the manipulated tangible object. The participants were able to manipulate three basic tangible shapes: a cone, a cube and a cylinder (see Table 2). The virtual objects were designed as a combination of these three different shapes. The participants were asked to rate the similarity between the manipulated object and the seen object. The suggested grasping positions of our approach were compared to grasping positions suggested by a global registration of the tangible and the virtual objects. The main hypotheses of our user experiment are:

- H1: when the tangible object and the virtual object are the same, the user perceives the objects as similar.
- H2: when the tangible object and the virtual object are different, the pinching positions suggested by our approach improves the perceived similarity between the two objects.
- H3: whatever the user is just touching the tangible object or lifting it, the perceived similarity remains identical.

4.1 Experiment Description

4.1.1 Participants

Twelve participants (10 males, 2 females, $M = 26.67$, $SD = 5.31$) took part to the experiment, all of whom were right-handed.

4.1.2 Experimental Apparatus

The setup is shown in Figure 2. Users wore an HTC Vive headset which displays the virtual scene with the virtual objects superimposed with the tangible one. A Bonita Vicon system combined with the Blade software was used to track their thumb and index fingers, at the back of which we positioned 3D-printed attachments holding a constellation of markers. We made sure that the pad of each finger remained free from the attachments, and to ensure that these did not move we secured them with wig tape. Using this tracking system, participants were able to control index and thumb avatars of their fingers in the virtual environment. This minimalist hand representation [2] has been chosen to avoid occluding the virtual object from the user’s point of view (see inset of Figure 2). The virtual scene is composed of an instruction panel and a pedestal onto which one VO is standing.

We chose not to track the tangible object for not affecting its perception. To know its position, we built a 3D-printed support, fixed to the table, which kept the tangible object at a known location.

To make sure that the TO was correctly replaced at the beginning of each trial, the support and the TO base had complementary shapes and magnets. This support also served as a reference point in the real environment. The virtual scene was then centered on this point using a Vive Tracker. Since Blade had also its own reference frame, we calibrated it using a third constellation. Therefore, the position of the TO was always known, standing on the calibrated support, while the VO was generated onto the TO at the beginning of the interaction. The virtual fingers were calibrated through the tracked constellations by touching a tangible prop at a specific known location. This allowed us to know where the contact between the fingers and the TO was with respect to the reference. In this way, we ensured that, whenever users were touching the TO, they were also touching the VO accordingly. This also resulted in the VO always moving together with the TO.

Visual feedback was used to guide users toward the target grasping pinches. Two virtual green cones indicated where the user was supposed to grasp the virtual (and tangible) objects. This pose was generated following the algorithm presented in Sec. 3. Visual feedback was also used to let participants know when they arrived close enough to the targeted pinch: whenever the users fingers were closer than 5 mm to their target pinch location and their orientation error w.r.t. the pinch main axis was lower than 14° , the grasping was considered successful and the VO turned red [41]. Holding this position for 1 s validated the pinch.



Figure 2: Setup for the experiment. Insets show the virtual environment during the task.

4.1.3 Procedure

During the experiment, participants were asked to perform two tasks: a static task, where they simply grasped the objects; and a dynamic task, where they grasped and lifted the objects.

First, participants were asked to grasp the objects as indicated on the screen, matching their two virtual fingers with the pose suggested by the algorithm via the green virtual cones (see first inset of Figure 2). When the pinch is validated, the object either turns green to confirm the end of the task (static task), or it turns blue to instruct the user to lift it (dynamic task). In this latter case, participants were asked to lift the tangible object up 10 cm and tilt it on each side. After that, participants were asked to put the object down and the task was considered completed. Thanks to the magnets on the support, it was easy to put back the tangible object on its original location. At the end of each trial, participants were asked to answer the following question: “How much does the tangible object correspond to the virtual one?” (see second inset of Figure 2) using a 7-item Likert scale.

The experiment started with an explanation of the procedure and the signature of the consent form. Then, a calibration of the virtual fingers was performed for each participant. Participants went first through all the trials of the static task, and then through the ones of the dynamic task. Before each task, 3 random practice trials were

performed to ensure that the user had understood the interaction, the task, and that the calibration was done properly. For each task, participants went through all the pinch pose registrations of the TO with the 6 VOs, either with or without our algorithm, in a random order. Since the algorithm allows a translation along the principal component axis, some of the computed registrations brought the VOs to be partially below or above the TO. The height of the pedestal was modified accordingly. The experimenter also faked the change of the tangible object between each trial.

4.1.4 Design

We chose a unique TO composed of three basic shapes: a cone, a cube, and a cylinder (see the representation in Table 2). The three shapes have the same height and the same width in the middle. The VOs correspond to a combination of these three different shapes. As a result, we ended up with 6 VOs (see Table 2).

In this experiment, we only considered three possible pinch poses on the TO, located at the middle of each shape (see Table 2, for their localization). Then, we used our algorithm to compute the best match between these poses and one potential pinch pose for each VO. We decided to take a set of 17 potential valid pinch poses per VO among the initial set of infinite pinch poses. The matching positions provided by our algorithm are shown in Table 2.

To better evaluate the influence of Criterion #2 (distance to the center of mass) on the suggested pinches, we divided the experiment in two identical parts: a first one where the participant had just to pinch the object (static task), not manipulating the tangible object and thus feeling any inertia, and a second one where the participant had to lift the object (dynamic condition). This order was not counter-balanced since the dynamic task could bring some additional cues about the properties of the tangible object.

In order to assess the performances of our approach, we chose to compare the registration provided by our algorithm with a reference corresponding to a simple global superimposition of the VO with the TO. For each task (static/dynamic), the participants had to evaluate the registration between the three pinch poses on the TO (cone, cube, cylinder) and three pinch poses on the VO (middle of the object for the *Ref* algorithm, generated following the description in Table 2 for our algorithm). Thus, for 6 different VOs, we obtained a block of 18 pinch poses x 2 algorithms = 36 trials. The different poses were presented in a randomized order within a block. The participant had to do each block three times, for a total of 108 trials for a task, and therefore 216 trials in total for the whole experiment, which lasted 1.5 hours.

4.2 Conditions

Three conditions are considered in our experimental design:


- **C1** is the registration algorithm between the TO and the VO: the algorithm of reference *Ref* corresponds to a superimposition of both objects, without using any information on the object properties; our algorithm is using the cost function described in Section 3 for finding the matching pose on the VO, given a pose on the TO.
- **C2** is the shape manipulated on the TO: cone, cube or cylinder (see Table 2).
- **C3** corresponds to the virtual object. The 6 VOs are represented in Table 2.

4.3 Collected Data

For each trial, we collected the participant’s answers on the correspondence between the VO and the TO using a 7-item Likert scale. Participants also completed a subjective questionnaire at the end of the experiment, following a 7-item Likert scale:

- Q1: It felt like **grabbing** a real object.
- Q2: It felt like **lifting** a real object.

Table 2: Tangible object (left) and the 6 virtual objects (right). All pinch poses are in the frontal plan. The three pinch poses of the TO are: 0 in blue at the middle of the cylinder, 1 in green at the middle of the cube and 2 in magenta at the middle of the cone. The associated pinch poses computed by our algorithm on each virtual object are displayed with the same color code. The "F" means that the registration was performed by flipping the VO upside down.



- Q3: It was **hard to grab** the object at the right location.
- Q4: It felt like I was **seeing my own fingertips**.
- Q5: The **shape** allows me to tell if both objects correspond to each other.
- Q6: The **curvature** allows me to tell if both objects correspond to each other.
- Q7: The **width** allows me to tell if both objects correspond to each other.
- Q8: The **perceived inertia** allows me to tell if both objects correspond to each other.
- Q9: I felt **tired** at the end.
- Q10: It felt as if the displayed objects were not all positioned at the **same height**.

We also asked them the following question at the end of the questionnaire “How many different tangible objects did you interact with in the experiment?”.

4.4 Results

4.4.1 Static Task

The answers of the participants concerning their perceived correspondence between the TO and the VOs are summarized in Table 3.

To study the perceived correspondence between the TO and the VO when they are identical, we first analyzed the ratings for the answers concerning VO #1 only (first column of the table). We combined answers from both algorithms since the pinch poses are identical for VO #1. We found a significant effect of the shape using a linear mixed model on the ratings ($F(2, 202) = 10.02, p < 0.001$). We performed a post-hoc analysis on the condition C2 using a Tukey test. Bonferroni correction was used for all post-hoc analysis. We

found that the correspondence between the cones of the TO and the VO was less perceived than for the cube ($Z = 4.4, p < 0.001$) and for the cylinder ($Z = 2.93, p = 0.009$). We did not find any significant effect between the cube and the cylinder for VO #1.

To study the perceived correspondence in the general case, we used a linear mixed model on the collected data to model the probability of perceived correspondence with respect to the three independent variables C1, C2 and C3 defined in the experimental design. The participants are considered as a random effect in the model. We performed an analysis of variance of the model and we found a significant effect for all conditions C1 ($F = (1, 1274) = 466.48, p < 0.001$), C2 ($F = (2, 1274) = 30.39, p < 0.001$), and C3 ($F = (5, 1274) = 51.62, p < 0.001$) as well as an interaction between C1 and C2 ($F = (2, 1274) = 17.83, p < 0.001$).

We performed a post-hoc analysis on C1 using a pairwise comparison and we found that the participants rated higher the correspondence between TO and VO for the poses proposed by our algorithm compared to the reference algorithm ($Z = 21.6, p < 0.001$), whatever the object ($p < 0.001$ for all with $Z = -8.06$ for the cone, $Z = -16.47$ for the cube and $Z = -12.88$ for the cylinder). We performed also a post-hoc analysis on C2 knowing C1 using a pairwise comparison. For the reference algorithm, we found that the correspondence between the cylinder of the TO and the corresponding VO was rated higher than the cone ($Z = 2.87, p = 0.01$) and the cube ($Z = 3.08, p = 0.005$). For our algorithm, the participants rated lower the correspondence between the cone of the TO and the corresponding pose on the VO compared to the cube ($Z = -8.21, p < 0.001$) and the cylinder ($Z = -7.70, p < 0.001$). Finally, we performed also a post-hoc analysis on C3 and we found that VO #1 was rated higher than the other VOs ($p < 0.001$ for each VO) and

than VO #4 was rated lower than the other VOs ($p < 0.001$ for each VO, except with VO #5 with $p = 0.007$). We found also that VO #2 was rated higher than VO #5 ($p = 0.04$).

Table 3: Ratings (Means and SD) for the static task using a 7-item Likert scale on the correspondence between the TO (first column, three shapes: cone, cube and cylinder) and the six different VOs, for the two compared algorithms (Reference and our approach). Ratings are in bold when TO shape and VO shape were identical.

	VO #1	VO #2	VO #3	VO #4	VO #5	VO #6
Cone <i>Ref</i>	6.06 (1.09)	2.19 (1.28)	3.47 (1.48)	2.11 (1.33)	3.08 (1.73)	5.97 (1.00)
Cone <i>Ours</i>	6.02 (1.03)	5.86 (1.13)	3.22 (1.61)	3.17 (1.52)	6.03 (1.16)	6.06 (1.22)
Cube <i>Ref</i>	6.58 (0.55)	2.44 (1.78)	6.53 (0.69)	2.47 (1.90)	2.28 (1.63)	2.39 (1.91)
Cube <i>Ours</i>	6.50 (0.56)	6.28 (0.91)	6.39 (0.93)	6.28 (0.81)	6.19 (0.86)	6.33 (0.76)
Cyl. <i>Ref</i>	6.31 (0.89)	6.39 (0.69)	3.50 (1.28)	3.81 (1.55)	2.78 (1.82)	2.78 (1.85)
Cyl. <i>Ours</i>	6.44 (0.84)	6.47 (0.84)	6.44 (0.69)	5.89 (1.06)	6.56 (0.56)	5.69 (1.37)

4.4.2 Dynamic Task

The answers of the participants concerning their perceived correspondence between the TO and the VOs are summarized in Table 4.

To study the perceived correspondence between the TO and the VO when they are identical, we first analyzed the ratings for the answers concerning VO #1 only. We combined answers from both algorithms since the pinch poses are identical for VO #1. We found a significant effect of the shape using a linear mixed model on the ratings ($F(2, 202) = 5.18$, $p = 0.006$). We performed a post-hoc analysis on the condition C2 using a Tukey test and we found that the correspondence between the cones of the TO and the VO was less perceived than for the cylinder ($Z = 3.15$, $p = 0.004$). We did not find any significant effect between the cube and the cylinder as well as between the cube and the cone for VO #1.

As for the static task, we used a linear mixed model on the collected data to model the probability of perceived correspondence with respect to the three independent variables C1, C2 and C3 defined in the experimental design for the general case. We performed an analysis of variance of the model and we found also a significant effect for all conditions C1 ($F = (1, 1274) = 422.07$, $p < 0.001$), C2 ($F = (2, 1274) = 48.76$, $p < 0.001$), and C3 ($F = (5, 1274) = 61.09$, $p < 0.001$) as well as an interaction between C1 and C2 ($F = (2, 1274) = 29.68$, $p < 0.001$). We performed a post-hoc analysis on C1 using a pairwise comparison and we found that the participants rated higher the correspondence between TO and VO for the poses proposed by our algorithm compared to the reference algorithm ($Z = 20.54$, $p < 0.001$), whatever the object ($p < 0.001$ for all with $Z = 5.70$ for the cone, $Z = 16.03$ for the cube and $Z = 13.85$ for the cylinder). We performed also a post-hoc analysis on C2 knowing C1 using a pairwise comparison. For the reference algorithm, we found that the correspondence between the cylinder of the TO and the corresponding VO was rated higher than the cone ($Z = 2.43$, $p = 0.04$). For our algorithm, the participants rated lower the correspondence between the cone of the TO and the corresponding pose on the VO compared to the cube ($Z = -10.62$, $p < 0.001$) and the cylinder ($Z = -10.59$, $p < 0.001$). Finally, we performed also a post-hoc analysis on C3 and we found that VO #1 was rated higher than the other VOs ($p < 0.001$ for each VO) and than VO #4 was rated lower than the other VOs ($p < 0.001$ for each VO).

Table 4: Ratings (Means and SD) for the dynamic task using a 7-item Likert scale on the correspondence between the TO (first column, three shapes: cone, cube and cylinder) and the six different VOs, for the two compared algorithms (Reference and our approach). Ratings are in bold when TO shape and VO shape were identical.

	VO #1	VO #2	VO #3	VO #4	VO #5	VO #6
Cone <i>Ref</i>	5.97 (1.30)	2.22 (1.48)	3.11 (1.41)	2.28 (1.37)	2.89 (1.51)	6.19 (0.95)
Cone <i>Ours</i>	5.92 (1.18)	5.03 (1.32)	3.06 (1.47)	2.42 (1.20)	5.42 (1.23)	5.92 (1.30)
Cube <i>Ref</i>	6.28 (1.03)	2.42 (1.50)	6.61 (0.60)	2.50 (1.50)	2.44 (1.48)	2.67 (1.79)
Cube <i>Ours</i>	6.25 (0.91)	6.33 (0.83)	6.44 (0.77)	6.08 (0.81)	6.14 (0.99)	5.97 (1.00)
Cyl. <i>Ref</i>	6.25 (0.84)	6.17 (0.97)	3.42 (1.38)	3.28 (1.49)	2.89 (1.56)	2.83 (1.40)
Cyl. <i>Ours</i>	6.58 (0.55)	6.47 (0.81)	6.28 (1.00)	5.69 (1.31)	6.39 (0.80)	5.78 (1.42)

4.4.3 Questionnaire

Regarding Q1 and Q2, most of the participants felt like grabbing ($M = 5.67$, $SD = 0.75$) or lifting ($M = 5.58$; $SD = 0.86$) the real object during the experiment. Regarding the difficulty to grab the object (Q3), the feeling is split between the participants ($M = 4.00$, $SD = 1.35$) with rates from 2 (2 participants) to 6 (2 participants). Participants mostly felt that they were seeing their own fingertips (Q4, $M = 5.17$, $SD = 1.77$). Figure 3 shows the answers to our subjective questionnaire for the questions regarding the properties of the object (Q5 to Q8). Among the different properties proposed for rating to the participants, they mostly felt that the shape allowed them to tell if both objects correspond to each other (Q5, $M = 5.92$, $SD = 1.03$) as well as the curvature (Q6, $M = 5.42$, $SD = 1.44$). The answers were more mitigated for width (Q7, $M = 3.83$, $SD = 1.77$) and inertia (Q8, $M = 4.75$, $SD = 1.88$). Finally, most of the participants reported overall medium levels of fatigue (Q9, $M = 5.33$; $SD = 1.37$). Most of the them also reported the difference of height between the displayed objects (Q10, $M = 5.33$, $SD = 1.80$). The answer to the question regarding the number of different tangible objects they were interacting with was between 2 to 20 ($M = 7.5$, $SD = 5.34$).




Figure 3: Bar-graph of the answers for questions 5 to 8 concerning the properties of the object that helped the participants to find a correspondence between the TO and the VO. Q5: shape, Q6: curvature, Q7: width, Q8: perceived inertia.

5 DISCUSSION

In this paper, we propose to use few tangible objects to enable 3D interactions with multiple virtual ones. To do so, we devised an algorithm which performs an exhaustive search on all the feasible pinching poses on the available objects, to find the best match in terms of haptic sensations. The algorithm is designed to provide the best available match between what the user sees in the virtual environment and what he feels when grasping the tangible object.

To evaluate the effectiveness of our approach, we carried out a user study. Twelve participants were asked to grasp objects of different shapes in virtual reality, all rendered by the same tangible, and evaluate how much the former corresponds to the latter. We made three hypotheses, described at the beginning of Sec. 4. First, we wanted to understand if participants are actually able to recognize when they are touching the same object in the real and virtual environments (H1). Results show that, indeed, they can. Values in bold in Tables 3 and 4 show users' ratings when provided with the same tangible and virtual shape. As expected, in both conditions (our algorithm and simple superimposition) ratings are very high. The second question we raised is whether, when tangible and virtual objects are different, using our approach improves the perceived grasping sensations with respect to using a standard superimposition (H2). Ratings show that our algorithm was able to well combine the considered haptic features and find convincing pinches between the given tangible and virtual objects. We registered improvements in ratings of 48.8% and 45.1% with respect to a standard superimposition matching technique in the case of static and dynamic grasping tasks, respectively. From Tables 3 and 4 we can see that ratings when using our algorithm are significantly closer to the ground-truth than when using simple superimposition. We also hypothesized (H3) that the above results hold when grasping an object with no subsequent motion (static task) as well as when lifting it up (dynamic task). Results show indeed little difference between the two tasks. At the end of the experiment, participants were also asked to guess how many tangible objects were used during the trials. Answers spanned from 2 to 20, with an average of 7.5 objects, and no one guessed that only one tangible object was used!

Although we have achieved promising results, our work still shows some limitations. First, our algorithm does not consider that humans have fingertips of different size and stiffness, which can significantly affect how they perceive a surface. Moreover, we only considered a reduced number of criteria for our metric. Of course, for future work, it would be important to include other perceptually-relevant properties, such as the object's texture. We also consider pinch grasping only, but extending our approach to other types of grasping should be possible. For example, extending our method to precision grasps with more than 2 fingers would require extending the description of our poses to more contact patches and the computation of the criteria to account for the increased number of contact couples. Also the interaction techniques used to grasp the objects can be improved, to ensure that the user always grasps the objects at the right location. Finally, another limitation of our approach is the tracking. Although we used a state-of-the-art tracking system, it is not easy to consistently and robustly track the fingertips at all time, and any discrepancy between the motion of the user's fingers and the virtual avatar will degrade the experience. However, when the same shapes were registered between the TO and VO, the experiment still gave the best scoring, highlighting the fact that the tracking issues did not hinder the perception of the users.

6 CONCLUSIONS

We presented an innovative haptic approach which enables the use of a reduced number of tangible objects for interacting with multiple different virtual ones. The need for such an approach stems from two conflicting requirements. To achieve a compelling illusion, there should be a good correspondence between what users see in the virtual environment and what they touch in the real world. In other words, the haptic features of the tangible and virtual objects should match. However, this calls for the creation of tangible replicas of all the virtual objects in the scene – which is often not possible. Therefore, it is important to devise approaches maximizing tangible/virtual objects matching, even when these are different. The proposed algorithm addresses this problem. It analyzes the provided tangible and virtual objects to find the best pinch in terms of matching haptic sensations. It starts by identifying several suitable pinching poses on all objects. Then, for each pose, it evaluates a

series of haptically-salient characteristics. Next, it identifies the two most similar pinching poses, one on a tangible and one on a virtual object. Finally, it highlights the chosen pose, which provides the best matching sensation between what users see and touch.

We tested our approach in a user study. Twelve participants were asked to grasp different virtual objects, all rendered by the same tangible one. For every virtual object, our algorithm found the best pinching match on the tangible one, and guided the participant toward that grasp. Results show that our algorithm was able to well combine several haptically-salient object features to find corresponding pinches between the given tangible and virtual objects. At the end of the experiment, participants were also asked to guess how many tangible objects were used during the experiment. No one guessed that we used only one, proof of a convincing experience.

This work then opens several interesting prospects: we could scan the real environment to exploit the real physical surroundings of the user. Our technique would then enable to provide precise pinching sensations. However, of course, this raises the challenge of real-time tracking algorithms. Moreover, using our algorithm, we could analyze the haptic properties of a given virtual scene, and we could try to automatically generate and 3D print, one (or more), *universal* tangible objects able to provide the best possible match in terms of haptic sensations.

Finally, although the results of our user study are promising and statistically significant, we believe that more experiments could be conducted, considering more objects, different tasks, different contexts.

REFERENCES

- [1] E. L. Amazeen and M. Turvey. Weight perception and the haptic size–weight illusion are functions of the inertia tensor. *Journal of Experimental Psychology: Human perception and performance*, 22(1):213, 1996.
- [2] F. Argelaguet, L. Hoyet, M. Trico, and A. Lecuyer. The role of interaction in virtual embodiment: Effects of the virtual hand representation. In *Proc. of IEEE International Conference on Virtual Reality*, pp. 3–10, 2016.
- [3] M. Azmandian, M. Hancock, H. Benko, E. Ofek, and A. D. Wilson. Haptic retargeting: Dynamic repurposing of passive haptics for enhanced virtual reality experiences. In *Proc. of ACM Conference on Human Factors in Computing Systems (CHI)*, pp. 1968–1979, 2016.
- [4] Y. Ban, T. Kajinami, T. Narumi, T. Tanikawa, and M. Hirose. Modifying an identified curved surface shape using pseudo-haptic effect. In *Proc. of IEEE Haptics Symposium*, pp. 211–216, 2012.
- [5] M. Billingham, H. Kato, and I. Poupyrev. The magicbook-moving seamlessly between reality and virtuality. *IEEE Computer Graphics and Applications*, 21(3):6–8, 2001.
- [6] M. Bjorkman, Y. Bekiroglu, V. Hogman, and D. Kragic. Enhancing visual perception of shape through tactile glances. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3180–3186, 2013.
- [7] M. Bonilla, D. Resasco, M. Gabbicini, and A. Bicchi. Grasp planning with soft hands using bounding box object decomposition. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 518–523, 2015.
- [8] F. Cazals and M. Pouget. Estimating differential quantities using polynomial fitting of osculating jets. *Computer Aided Geometric Design*, 22(2):121–146, 2005.
- [9] A. D. Cheok, X. Yang, Z. Z. Ying, M. Billingham, and H. Kato. Touchspace: Mixed reality game space based on ubiquitous, tangible, and social computing. *Personal and ubiquitous computing*, 6(5-6):430–442, 2002.
- [10] F. Chinello, M. Malvezzi, D. Prattichizzo, and C. Pacchierotti. A modular wearable finger interface for cutaneous and kinesthetic interaction: control and evaluation. *IEEE Transactions on Industrial Electronics*, 2019.
- [11] A. Crowe, W. Keessen, W. Kuus, R. Van Vliet, and A. Zegeling. Proprioceptive accuracy in two dimensions. *Perceptual and motor skills*, 64(3):831–846, 1987.

- [12] M. R. Cutkosky. On grasp choice, grasp models, and the design of hands for manufacturing tasks. *IEEE Transactions on robotics and automation*, 5(3):269–279, 1989.
- [13] P. Dalsgaard and K. Halskov. Tangible 3d tabletops. *Interactions*, 21(5):42–47, 2014.
- [14] B. Eckstein, E. Krapp, and B. Lugrin. Towards serious games and applications in smart substitutional reality. In *Proc. of 10th International Conference on Virtual Worlds and Games for Serious Applications (VS-Games)*, pp. 1–8, 2018. doi: 10.1109/VS-Games.2018.8493444
- [15] T. Feix, J. Romero, H. Schmiedmayer, A. M. Dollar, and D. Kragic. The grasp taxonomy of human grasp types. *IEEE Transactions on Human-Machine Systems*, 46(1):66–77, 2016.
- [16] D. Fischinger, M. Vincze, and Y. Jiang. Learning grasps for unknown objects in cluttered scenes. In *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, pp. 609–616, 2013.
- [17] J. R. Flanagan and C. A. Bandomir. Coming to grips with weight perception: effects of grasp configuration on perceived heaviness. *Perception & Psychophysics*, 62(6):1204–1219, 2000.
- [18] I. E. Gordon and V. Morison. The haptic perception of curvature. *Perception & psychophysics*, 31(5):446–450, 1982.
- [19] M. Gualtieri, A. ten Pas, K. Saenko, and R. Platt. High precision grasp pose detection in dense clutter. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 598–605, 2016.
- [20] D. T. Han, M. Suhail, and E. D. Ragan. Evaluating remapped physical reach for hand interactions with passive haptics in virtual reality. *IEEE Transactions on Visualization & Computer Graphics*, (1):1–1, 2018.
- [21] D. Harley, A. P. Tarun, D. Germinario, and A. Mazalek. Tangible vr: Diegetic tangible objects for virtual reality narratives. In *Proc. of Conference on Designing Interactive Systems*, pp. 1253–1263, 2017.
- [22] V. Hayward, O. Astley, M. Cruz-Hernandez, D. Grant, and G. Robles-De-La-Torre. Haptic interfaces and devices. *Sensor Review*, 24(1):16–29, 2004.
- [23] W. A. Hershberger and G. F. Misceo. Touch dominates haptic estimates of discordant visual-haptic size. *Perception & Psychophysics*, 58(7):1124–1132, 1996.
- [24] A. Hettiarachchi and D. Wigdor. Annexing reality: Enabling opportunistic use of everyday objects as tangible proxies in augmented reality. In *Proc. of the ACM Conference on Human Factors in Computing Systems (CHI)*, pp. 1957–1967, 2016.
- [25] B. Insko. *Passive haptics significantly enhances virtual environments*. PhD thesis, University of North Carolina at Chapel Hill, 2001.
- [26] A. M. Kappers and J. J. Koenderink. Haptic unilateral and bilateral discrimination of curved surfaces. *Perception*, 25(6):739–749, 1996.
- [27] M. Körtgen, G.-J. Park, M. Novotni, and R. Klein. 3d shape matching with 3d shape contexts. In *Proc. of the 7th central European seminar on computer graphics*, vol. 3, pp. 5–17. Budmerice, 2003.
- [28] S. Lacey and K. Sathian. Visuo-haptic multisensory object recognition, categorization, and representation. *Frontiers in Psychology*, 5:730, 2014.
- [29] A. Lécuyer. Simulating haptic feedback using vision: A survey of research and applications of pseudo-haptic feedback. *Presence: Teleoperators and Virtual Environments*, 18(1):39–53, 2009.
- [30] S. J. Lederman and R. L. Klatzky. Haptic perception: A tutorial. *Attention, Perception, & Psychophysics*, 71(7):1439–1459, 2009.
- [31] Y. Li, J. L. Fu, and N. S. Pollard. Data-driven grasp synthesis using shape matching and task-based pruning. *IEEE Transactions on Visualization and Computer Graphics*, 13(4):732–747, 2007.
- [32] Y. Li, X. Wu, Y. Chrysathou, A. Sharf, D. Cohen-Or, and N. J. Mitra. Globfit: Consistently fitting primitives by discovering global relations. *ACM Transactions on Graphics*, 30(4):52:1–52:12, 2011.
- [33] C. K. Liu. Dextrous manipulation from a grasping pose. *ACM Transactions on Graphics*, 28(3):59, 2009.
- [34] J. C. McClelland, R. J. Teather, and A. Girouard. Haptobend: Shape-changing passive haptic feedback in virtual reality. In *Proc. 5th Symposium on Spatial User Interaction (SUI)*, pp. 82–90, 2017.
- [35] L. Meli, C. Pacchierotti, G. Salvietti, F. Chinello, M. Maisto, A. De Luca, and D. Prattichizzo. Combining wearable finger haptics and augmented reality: User evaluation using an external camera and the microsoft hololens. *IEEE Robotics and Automation Letters*, 3(4):4297–4304, 2018.
- [36] A. T. Miller, S. Knoop, H. I. Christensen, and P. K. Allen. Automatic grasp planning using shape primitives. In *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2, pp. 1824–1829, 2003.
- [37] T. Ohshima, K. Satoh, H. Yamamoto, and H. Tamura. Ar2 hockey system: A case study of collaborative mixed reality system. *Trans. VRSJ*, 3(2):55–60, 1998.
- [38] C. Pacchierotti, S. Sinclair, M. Solazzi, A. Frisoli, V. Hayward, and D. Prattichizzo. Wearable haptic systems for the fingertip and the hand: taxonomy, review, and perspectives. *IEEE Transactions on Haptics*, 10(4):580–600, 2017.
- [39] S. C. Pont, A. M. Kappers, and J. J. Koenderink. Haptic curvature discrimination at several regions of the hand. *Perception & Psychophysics*, 59(8):1225–1240, 1997.
- [40] S. C. Pont, A. M. Kappers, and J. J. Koenderink. Similar mechanisms underlie curvature comparison by static and dynamic touch. *Perception & Psychophysics*, 61(5):874–894, 1999.
- [41] M. Prachyabrued and C. W. Borst. Visual feedback for virtual grasping. In *Proc. of IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 19–26, 2014.
- [42] D. Prattichizzo and J. C. Trinkle. *Springer Handbook of Robotics*, chap. Grasping, pp. 955–988, 2016.
- [43] M. S. M. Y. Sait, S. P. Sargunam, D. T. Han, and E. D. Ragan. Physical hand interaction for controlling multiple virtual objects in virtual reality. In *Proc. of the 3rd International Workshop on Interactive and Spatial Computing*, pp. 64–74, 2018. doi: 10.1145/3191801.3191814
- [44] S. Scheggi, L. Meli, C. Pacchierotti, and D. Prattichizzo. Touch the virtual reality: using the Leap Motion controller for hand tracking and wearable tactile devices for immersive haptic rendering. In *Proc. ACM SIGGRAPH Conference Posters*, p. 31, 2015.
- [45] R. Schnabel, P. Degener, and R. Klein. Completion and reconstruction with primitive shapes. *Computer Graphics Forum*, 28(2):503–512, 2009.
- [46] L. Shapira, A. Shamir, and D. Cohen-Or. Consistent mesh partitioning and skeletonisation using the shape diameter function. *The Visual Computer*, 24(4):249, 2008.
- [47] A. L. Simeone, E. Velloso, and H.-W. Gellersen. Substitutional reality: Using the physical environment to design virtual reality experiences. In *Proc. of ACM Conference on Human Factors in Computing Systems (CHI)*, pp. 3307–3316, 2015.
- [48] G. E. Stelmach and A. Sirica. Aging and proprioception. *Age*, 9(4):99–103, 1986.
- [49] M. Szilvási-Nagy and S. Béla. B-spline patches fitting on surfaces and triangular meshes. *KoG*, 15(15):17–24, 2011.
- [50] A. ten Pas, M. Gualtieri, K. Saenko, and R. Platt. Grasp pose detection in point clouds. *The International Journal of Robotics Research*, 36(13-14):1455–1473, 2017.
- [51] J. Tierny, J.-P. Vandeborre, and M. Daoudi. Topology driven 3d mesh hierarchical segmentation. In *Proc. of IEEE International Conference on Shape Modeling and Applications*, pp. 215–220, 2007.
- [52] V. Vidal, C. Wolf, and F. Dupont. *Mechanical Mesh Segmentation and Global 3D Shape Extraction*. PhD thesis, Université Lyon 1-Claude Bernard; INSA Lyon, 2014.
- [53] M. W. A. Wijntjes, A. Sato, V. Hayward, and A. M. L. Kappers. Local surface orientation dominates haptic curvature discrimination. *IEEE Transactions on Haptics*, 2(2):94–102, 2009.
- [54] D.-M. Yan, W. Wang, Y. Liu, and Z. Yang. Variational mesh segmentation via quadric surface fitting. *Computer-Aided Design*, 44(11):1072–1082, 2012.
- [55] Y. Zhao, L. H. Kim, Y. Wang, M. Le Goc, and S. Follmer. Robotic assembly of haptic proxy objects for tangible interaction and virtual reality. In *Proc. of the ACM International Conference on Interactive Surfaces and Spaces (ISS)*, pp. 82–91, 2017. doi: 10.1145/3132272.3134143
- [56] Y. Zhou, K. Yin, H. Huang, H. Zhang, M. Gong, and D. Cohen-Or. Generalized cylinder decomposition. *ACM Transactions on Graphics*, 34(6):171–1, 2015.