

Everyday Objects for Providing Variable Haptic Feedback in Virtual Reality

Yujie Tao
Stanford University
Stanford, CA
yjtao@stanford.edu

1. Abstract

Virtual reality (VR) offers immersive experiences but often lacks physical feedback. In this work, we propose a novel technique that enables everyday objects to provide variable haptic feedback in VR. By reconstituting virtual objects with parts of physical props, we demonstrated that participants tend to grasp the re-composed objects by their newly manipulable parts. This method expands the range of physical feedback achievable with everyday passive objects. Furthermore, by implementing a part segmentation algorithm using a geometric approach, our project achieves automatic segmentation of both virtual and physical objects. Our user study revealed that these re-composed VR objects significantly enhance the haptic experience for users. We concluded the paper with a discussion of limitations and future directions.

2. Introduction

Through Virtual Reality (VR), individuals immerse themselves in environments beyond their physical surroundings and embody virtual representations that are different from their own body appearances. These immersive environments offered by VR unleash new opportunities for learning, entertainment, and social interactions.

With the increasing popularity and availability of commercial VR headsets, such as the Meta Quest and HTC VIVE, a significant challenge remains unsolved: delivering realistic haptic feedback within virtual environments. Haptic technology aims to recreate the sense of touch by applying forces, vibrations, or motions to the user, thereby enhancing the overall immersion and realism of virtual experiences. Despite its potential, the widespread adoption of haptic feedback devices has been limited. These devices are often perceived as bulky and expensive, making them less accessible to general users. Consequently, the challenge lies in developing more user-friendly and cost-effective haptic solutions to bring this technology into everyday use and maximize its potential in various applications.

In this work, we aim to explore the possibility of using

everyday objects—items already present in the user’s environment—to provide haptic feedback in virtual reality. By leveraging common household objects, we hope to make haptic feedback more accessible and cost-effective for a wider range of users. However, a significant challenge in using daily objects for haptic feedback in VR is that physical objects are constrained by their static physical shapes, while virtual objects are dynamic and can change over time. For instance, you might see a virtual mug in front of you in VR but only have a physical scissor as prop. This discrepancy limits the capability of the physical scissor to provide realistic feedback for a VR mug.

Therefore, we focus on scenarios where there is a discrepancy between the shapes of physical props and their virtual counterparts. How can we achieve not only a one-to-one mapping of passive objects to virtual objects for accurate physical feedback but also enable these passive objects to provide variable feedback for multiple virtual objects?

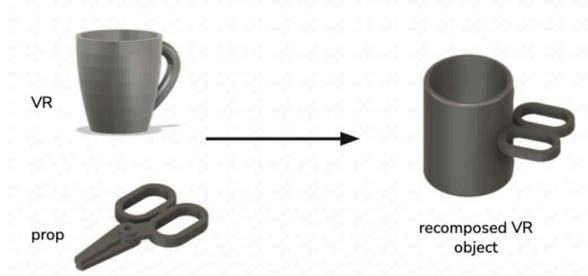


Figure 1. A demonstration of recombining affordance of virtual and physical objects.

In this project, we propose to integrate the affordances of virtual and physical objects to enable the same physical objects to provide haptic feedback for various virtual objects in VR, shown in Figure 1. Here, “affordance” refers to the manipulable part of an object (e.g., the handle of a mug). By leveraging the heuristic that people tend to grasp objects following predictable patterns on certain manipulable parts, we hypothesize that objects with a re-composed

manipulable part will still be grasped on this new part. By composing the affordance part of a physical object with the virtual object while preserving the semantic meaning of the virtual object, we aim to allow physical objects to provide haptic feedback for multiple virtual objects. This approach could enhance the versatility and practicality of using everyday objects for haptic feedback in virtual environments

To demonstrate the general concept, we visualize an example of affordance recomposition in Figure 1.

3. Related Work

To enhance physical feedback in virtual reality, one prominent direction is to leverage physical objects in everyday environment to provide touch feedback for virtual objects, so-called “passive haptics” [8,9]. By utilizing everyday objects, users can interact with virtual items and receive rich tactile feedback—such as texture, softness, and temperature—from the real objects, making it feel as though the sensations are coming from the virtual items [6].

Early work in passive haptics, such as the research conducted by Hinckley et al. [8] and Hoffman et al. [9], focused on the registration of physical objects to virtual objects. Those foundational work explored how real-world items could be precisely aligned and integrated with virtual counterparts to create a cohesive and immersive VR experience. The concept relied on the idea that if a user touches or manipulates a physical object, the virtual environment would reflect this interaction accurately, thus enhancing the sense of immersion and realism.

To further develop and generalize the technique of passive haptics, Cheng et al. [3] proposed a set of haptic proxies. These proxies used a general passive prop, which could be adapted to represent various virtual objects in a VR environment. This innovation aimed to provide a more versatile and scalable approach to integrating passive haptics into VR, allowing for a broader range of interactions and experiences.

The mapping between the virtual and physical worlds can range from single objects to entire rooms. For instance, Sra et al. [13] and Shapira et al. [12] proposed systems that scan and reconstruct users’ surroundings. These systems allow daily objects to serve as haptic props, seamlessly integrating the user’s immediate environment into the VR experience. This capability transforms everyday items into interactive elements within the virtual world. RealityCheck, developed by Wilson et al. [14], offers another innovative approach by combining 3D reconstruction of the real world with the virtual environment. This method not only maps physical space into virtual reality but also supports free VR locomotion. Users can move around their physical environment while experiencing a seamlessly integrated virtual world, enhancing the sense of presence and spatial awareness. Additionally, VRoamer, created by Cheng et al. [4],

enables users to walk in unseen physical spaces. The system dynamically generates a virtual scene on-the-fly, allowing users to explore new and uncharted virtual territories while navigating their real-world environment safely.

Sensory VR, introduced by Harley et al. [7], extended the exploration of passive haptics by incorporating the senses of smell and taste. This approach bridges the gap between the everyday environment and virtual reality, providing a multi-sensory experience that includes not just touch but also olfactory and gustatory feedback.

While passive haptics provide touch feedback in a low-cost and accessible manner, they suffer from being static, which do not always satisfy the need for dynamic virtual scenes. VR scenes often change dynamically, creating experience beyond real-world interactions [1]. In this work, we focus on extending the capability of passive objects to provide not only one-to-one mapping to a virtual object but also variable haptic feedback.

4. Approach

The primary problem statement of this project is to determine whether objects with recomposed parts will be grasped on the substituted manipulable part when used with a functional intent. If validated as true, the same physical object (e.g., a scissor) could then provide haptic feedback for both its virtual counterpart (e.g., a virtual scissor) and the corresponding recomposed virtual objects (e.g., a mug with a scissor handle). Such a technique preserves the functional meaning of the target objects while only swapping the graspable parts to enable everyday objects to provide variable feedback.

To explore this technique, we began with a grasping pattern analysis to determine if there is a predictable behavior pattern for object grasping with a functional intent. Next, we conducted a preliminary user perception test to see if users would interact with re-composed objects on the recomposed manipulable parts. Following this, we implemented a segmentation algorithm to achieve automatic segmentation of multi-part objects. Finally, we utilized recomposed objects in VR to understand the effect of recomposing affordance on the haptic experience.

4.1. Grasping Pattern Analysis

An underlying assumption behind the proposed technique is that people grasp objects in a predictable pattern. We first analyzed the grasping behavior provided in the ContactPose dataset [2] to validate this assumption. ContactPose contains 2306 unique grasps of 25 household objects grasped with two functional intents by 50 participants. We chose this dataset because it provides hand-object contact paired with hand pose and object pose. Such datasets help us understand if there are common grasping patterns and manipulable parts of each provided object.

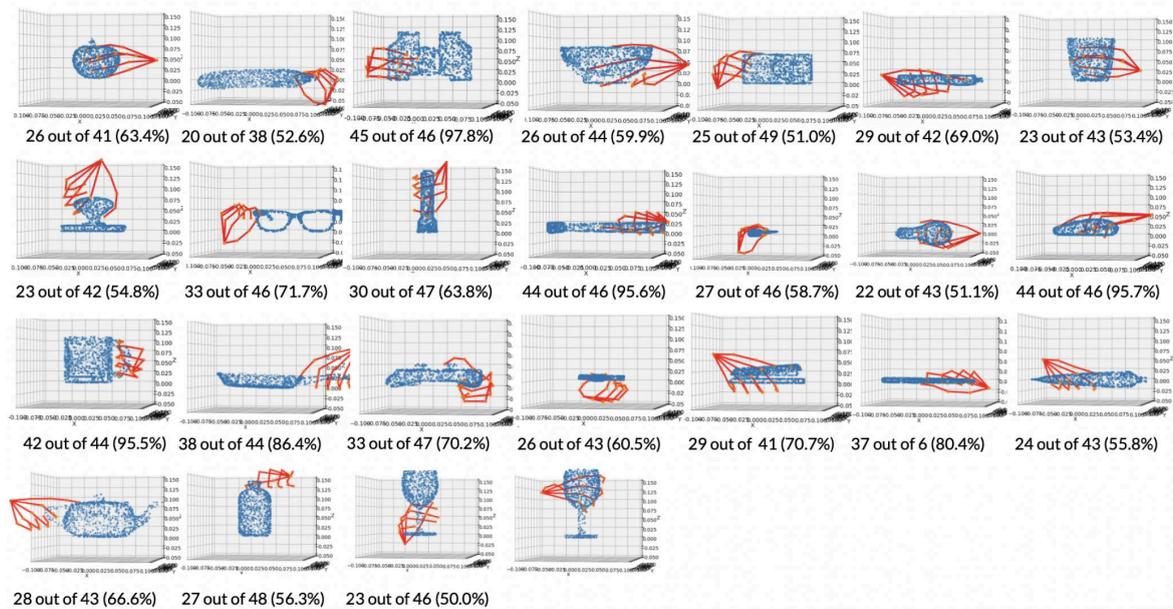


Figure 2. Analysis of common grasping pattern of objects in ContactPose [2].

We applied the clustering algorithm *Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN)* on the hand pose data to derive each object’s most common grasping pattern, as shown in Figure 2. Note that the total number of grasps for each object varies because not all 50 individuals interacted with all 25 objects during the data collection process.

For 24 out of 25 objects, there exists a grasping pose agreed upon by more than half of the individuals who grasped the object in the data collection. For the object without a majority consensus (i.e. wine glass), half of the participants hold it on the stem while the other half hold it on the bowl. The objects with the highest majority votes for grasping is binocular (97.8%), mug (95.5%), mouse (95.7%) and pan (86.4%).

It’s also shown that objects with a multi-part structure (e.g., a teapot or a pan with a handle) elicit a stronger consensus on grasping behavior compared to objects with a single structure (e.g., an apple).

The grasping analysis results showed a predictable grasping behavior pattern given a functional intent. It provides grounding to the proposal of recomposing parts of physical to virtual objects to provide variable feedback. This analysis helps us select an initial set of objects to start with for experimenting on the part recombination. The initial set of the 3D models selected are binoculars, door handles, mugs, knives, and scissors, as shown in Figure 3.

To get an understanding of how people interact with ob-

jects with recomposed affordance, we manually segmented the graspable part of each object following the dataset analysis results and recomposed it onto the mug object. This results in 3D models as shown below in Figure 3.



Figure 3. Selected 3D models for further exploration.

4.2. User Interaction with Re-composed Objects

With the grasping analysis completed, the next question is how users will interact with recomposed objects and whether they will grasp these objects by their recomposed parts under a functional intent. To investigate this, a preliminary user evaluation is needed.

To address this, we conducted a pilot study with one participant to gain initial insights into how people interact with objects featuring recomposed affordances. The participant was invited to the lab and presented with nine objects—five original (base) objects and four recomposed objects—one by one in a randomized order. They were instructed to grasp

each object as they would intend to use it. To mitigate the novelty effect, the entire process was repeated twice.

During the study, the participant’s grasping behavior was recorded using a webcam, with a hand pose estimation algorithm [11] running in real-time. This setup allowed us to capture detailed data on how the participant interacted with both the original and recomposed objects, providing valuable insights into whether the recomposed affordances influenced their grasping patterns. The findings from this pilot study will inform us about user perception of re-composed objects and validate if people will interact with the recomposed objects on the new manipulable parts.

We found that in all four recomposed cases, the participant grasped the recomposed objects following the affordance indicated by the new manipulable part in both trials. While participants grasped the new manipulable parts consistently, the exact grasping behavior might vary by trial. An example is the mug with a scissor handle case, shown in Figure 4. This initial pilot confirmed that objects with recomposed affordance will be grasped on the substituted manipulable part when used with a functional intent.



Figure 4. Grasping behavior from initial pilot

After the task, the participant also shared their experience interacting with the object. Participant reported that all base objects (no recombination) were intuitive grab. The recomposed objects were novel to interact with and some of more intuitive than others to grab. The participant was confident about how to grab the knife-mug and knob-mug recombination but unsure about the scissor mug: “I want to grab it like a scissor handle, but it felt weird. Should I forget that it is a scissor, or should I grab it like with my finger in between?” While recomposed objects might cause confusion at first sight, participants generally find the experience to be fun and interesting.

4.3. Segmentation Algorithm

With validation of the common grasping pattern and user perception of re-composed objects, we are now ready to implement the segmentation pipeline. We implemented the segmentation algorithm based on the method proposed in *Hierarchical Mesh Decomposition using Fuzzy Clustering and Cut* [10].

4.3.1 Initialization and Data Reading

The segmentation process begins with reading a 3D model from a PLY file, extracting vertices and face data. Each face of the mesh is represented by three vertices, and the normal vector of each face is computed. This step sets up the initial data structures required for further processing.

4.3.2 Neighbor Computation

Next, the algorithm computes information about neighboring faces for each face in the mesh. This involves creating edges from face vertices and identifying pairs of faces that share an edge. For each neighboring pair of faces, the algorithm calculates the angular distance and the geodesic distance. These distances are combined into a weighted metric that reflects the overall similarity between faces. This step establishes the relationships between adjacent faces, which is crucial for segmentation.

4.3.3 Shortest Path Calculation

After computing neighbor information, the algorithm calculates the shortest paths between all pairs of faces using Dijkstra’s algorithm. This is done in parallel to speed up the process. The shortest paths are stored in a distance matrix, which will be used later to determine face segments based on proximity and connectivity. We ensure that the algorithm can efficiently find and use the shortest connections between any two faces in the mesh.

4.3.4 Segmentation

The segmentation process performs hierarchical segmentation of the mesh. Initially, representative faces are selected based on their distances to other faces. The algorithm iteratively calculates the probability of each face belonging to a segment and refines these probabilities. Faces are then assigned to segments based on the highest probability.

To resolve ambiguous regions where faces do not clearly belong to a single segment, the algorithm uses a max-flow method based on the Ford-Fulkerson algorithm. It constructs a flow network where each face is a node, and edges represent connections between neighboring faces. The max-flow algorithm is used to ensure clear boundaries between segments by reassigning faces in fuzzy regions.

Then, we refine the segmentation by eliminating ambiguities and ensuring that each face is clearly associated with a single segment. Finally, the segmented mesh is written to a PLY file with each segment colored differently for visualization. The final output from segmentation algorithm is shown in Figure 5.

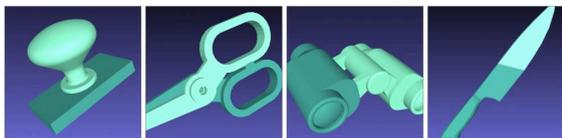


Figure 5. Segmentation results from the pipeline

4.4. Understand impact on VR Haptic Experience

With the implementation of the part segmentation pipeline, we are now prepared to investigate whether the re-composed virtual objects can actually enhance the haptic experience in VR. To do so, we developed a VR application using Unity 3D. In this experiment, participants were invited to the lab and equipped with a VR headset (Meta Quest 2) to interact with virtual objects, as shown in Figure 6. Four physical objects (scissors, mug, knife, and binoculars) were chosen to provide haptic feedback for a virtual mug in the VR environment, composing four sessions of interactions.

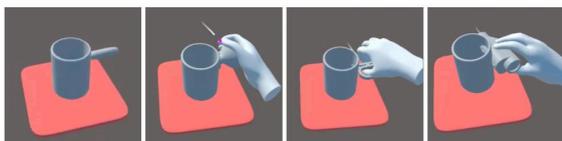


Figure 6. Implementation of the VR scene.

During each session, participants used one of the physical objects as a haptic prop. In the session, participants interacted with two versions of the virtual objects: the original and the re-composed virtual objects. They were instructed to grab the virtual objects as if they were going to use them. In the real-world, they always touch with the same physical objects as the haptic prop. This setup allowed us to evaluate the differences in haptic experience between the original and re-composed virtual objects.

We conducted the study with a single participant. As illustrated in Figure 7, the participant generally reported a stronger sense of haptic realism when interacting with re-composed VR objects compared to the original VR objects. This trend was consistent across all four objects the participant interacted with. These preliminary findings suggest that re-composing affordances between physical and virtual objects can enhance the effectiveness of passive haptic props.

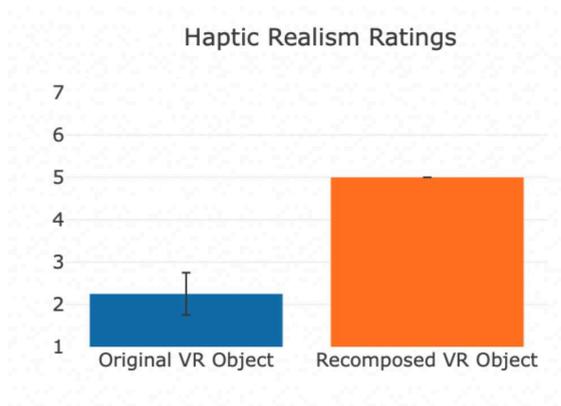


Figure 7. Rated haptic realism from VR study.

5. Future Work

This project serves as a preliminary exploration of enabling everyday objects to provide variable haptic feedback in VR. We see multiple future directions to push forward this project.

5.1. Affordance Segmentation Using Deep-learning Approaches

In this work, we achieved part segmentation using a geometric approach. This method focuses on the spatial and structural aspects of the objects, allowing us to delineate various components based purely on their shapes and spatial relationships. However, this geometric solution has its limitations. It does not consider the semantic meaning of different parts, which is crucial for understanding and interpreting the functionality and interaction potential of these parts in a meaningful context.

Recently, there has been a growing interest in creating datasets of hand-object interactions. These datasets are invaluable as they capture a wide range of real-world scenarios where individuals interact with various objects, providing rich data on how different objects are grasped and manipulated. By analyzing these interactions on a large scale, we can uncover patterns and regularities in grasping behaviors, which can significantly enhance our ability to predict and model how objects are typically handled.

One promising future direction of this project is the development and utilization of models like 3D AffordanceNet [5]. This model offers a comprehensive framework for understanding and predicting the affordances of objects in three-dimensional space. By leveraging the data from hand-object interaction datasets, 3D AffordanceNet can be trained to recognize and predict various grasping patterns, providing a robust foundation for advancing our un-

derstanding of object manipulation. This approach moves beyond mere geometric segmentation, incorporating semantic insights that are essential for developing more intuitive and human-like models of object interaction.

5.2. User Evaluation on a Larger Scale

Moreover, the user evaluation in this project is limited by a very small sample size, which provides only preliminary insights into the effectiveness and usability of the system. While these initial findings are useful, they are not sufficient to draw definitive conclusions about the broader applicability and reliability of the approach. Future evaluations should involve a significantly larger number of participants to enhance the statistical validity and generalizability of the findings. A more diverse participant pool will also help in understanding how the system performs across different user demographics and prior VR experience. Additionally, the test environment should better reflect the type of VR experiences users are likely to engage in, such as specific games that involve tasks like grasping objects.

6. Conclusion

In this work, we propose a new technique to enable everyday objects to provide variable haptic feedback in virtual reality. By recomposing virtual objects with parts of physical props, we demonstrated that participants will grasp the recomposed objects by their newly manipulable parts. This approach enhances the range of physical feedback that can be achieved using everyday passive objects. Additionally, by implementing a part segmentation algorithm using a geometric approach, our project achieves automatic segmentation of virtual and physical objects. Finally, our user study showed that the re-composed VR objects enhance the haptic experience for users.

References

- [1] Parastoo Abtahi, Sidney Q Hough, James A Landay, and Sean Follmer. Beyond being real: A sensorimotor control perspective on interactions in virtual reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–17, 2022. [2](#)
- [2] Samarth Brahmabhatt, Chengcheng Tang, Christopher D Twigg, Charles C Kemp, and James Hays. Contactpose: A dataset of grasps with object contact and hand pose. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, pages 361–378. Springer, 2020. [2](#), [3](#)
- [3] Lung-Pan Cheng, Eyal Ofek, Christian Holz, Hrvoje Benko, and Andrew D Wilson. Sparse haptic proxy: Touch feedback in virtual environments using a general passive prop. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 3718–3728, 2017. [2](#)
- [4] Lung-Pan Cheng, Eyal Ofek, Christian Holz, and Andrew D Wilson. Vroamer: generating on-the-fly vr experiences while walking inside large, unknown real-world building environments. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 359–366. IEEE, 2019. [2](#)
- [5] Shengheng Deng, Xun Xu, Chaozheng Wu, Ke Chen, and Kui Jia. 3d affordancenet: A benchmark for visual object affordance understanding. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1778–1787, 2021. [5](#)
- [6] Cathy Mengying Fang, Ryo Suzuki, and Daniel Leithinger. Vr haptics at home: Repurposing everyday objects and environment for casual and on-demand vr haptic experiences. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–7, 2023. [2](#)
- [7] Daniel Harley, Alexander Verni, Mackenzie Willis, Ashley Ng, Lucas Bozzo, and Ali Mazalek. Sensory vr: Smelling, touching, and eating virtual reality. In *Proceedings of the twelfth international conference on tangible, embedded, and embodied interaction*, pages 386–397, 2018. [2](#)
- [8] Ken Hinckley, Randy Pausch, John C Goble, and Neal F Kassell. Passive real-world interface props for neurosurgical visualization. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 452–458, 1994. [2](#)
- [9] Hunter G Hoffman. Physically touching virtual objects using tactile augmentation enhances the realism of virtual environments. In *Proceedings. IEEE 1998 Virtual Reality Annual International Symposium (Cat. No. 98CB36180)*, pages 59–63. IEEE, 1998. [2](#)
- [10] Sagi Katz and Ayellet Tal. Hierarchical mesh decomposition using fuzzy clustering and cuts. *ACM transactions on graphics (TOG)*, 22(3):954–961, 2003. [4](#)
- [11] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*, 2019. [4](#)
- [12] Lior Shapira and Daniel Freedman. Reality skins: Creating immersive and tactile virtual environments. In *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 115–124. IEEE, 2016. [2](#)
- [13] Misha Sra, Sergio Garrido-Jurado, Chris Schmandt, and Pattie Maes. Procedurally generated virtual reality from 3d reconstructed physical space. In *Proceedings of the 22nd ACM conference on virtual reality software and technology*, pages 191–200, 2016. [2](#)
- [14] Andrew D Wilson, Christian Holz, Eyal Ofek, and Jeremy Hartmann. Blending virtual environments with situated physical reality, May 11 2021. US Patent 11,004,269. [2](#)