Earthquake Safety Training through Virtual Drills

Changyang Li, Wei Liang, Chris Quigley, Yibiao Zhao and Lap-Fai Yu, *Member, IEEE*

Fig. 1: The user undergoes a training in a virtual environment to learn survival skills applicable during an earthquake. Left: An office scene used for training. Right: The office scene during a simulated earthquake. The user learns to detect potential danger and to protect himself through an immersive training experience.

Abstract—Recent popularity of consumer-grade virtual reality devices, such as the Oculus Rift and the HTC Vive, has enabled household users to experience highly immersive virtual environments. We take advantage of the commercial availability of these devices to provide an immersive and novel virtual reality training approach, designed to teach individuals how to survive earthquakes, in common indoor environments. Our approach makes use of virtual environments realistically populated with furniture objects for training. During a training, a virtual earthquake is simulated. The user navigates in, and manipulates with, the virtual environments to avoid getting hurt, while learning the observation and self-protection skills to survive an earthquake. We demonstrated our approach for common scene types such as offices, living rooms and dining rooms. To test the effectiveness of our approach, we conducted an evaluation by asking users to train in several rooms of a given scene type and then test in a new room of the same type. Evaluation results show that our virtual reality training approach is effective, with the participants who are trained by our approach performing better, on average, than those trained by alternative approaches in terms of the capabilities to avoid physical damage and to detect potentially dangerous objects.

Index Terms—Virtual reality, modeling and simulation, virtual worlds training simulations

1 Introduction

Earthquake safety is a major issue in many parts of the world. According to the report of the Seismological Society of America, Nevada and California experience over 5,000 earthquakes annually. Of these, over 100 are rated between 6 and 6.75 on the Richter scale. An additional 20 occur which rank between 7 and 7.7 [15]. In areas where earthquakes are this frequent, it is important for an individual to know how to protect himself or herself in the case of an emergency.

Traditionally, earthquake safety has been taught through simulated drills, frequently mandated at schools located in regions with a high risk of earthquakes. However, a recent study conducted by Ramirez et al. [26] found that this method of drilling commonly suffers from the problem of being non-standardized, and that, as a whole, the drills conducted at many schools have not been effective in improving preparedness of students for emergency situations such as earthquakes. One key suggestion for improvements is developing a more realistic simulated exercise drill.

Our work explores using virtual reality to provide this realistic simulated experience to an individual. Figure 1 illustrates this idea. The user navigates in a virtual environment mimicking a common indoor scene such as an office, which is populated with common objects whose masses and physical properties have been realistically assigned. An earthquake simulation is then applied to the environment while the user tries to protect himself to prevent his avatar from being hurt in the virtual environment. Through this immersive experience in several different rooms, the user is trained to gain observation and self-protection skills to survive an earthquake.

The fact that the user is not physically harmed during the simulation allows us to include features in our earthquake scenarios that might be considered dangerous or impractical in a real-world simulated drill (e.g. the breaking of windows, the shaking of the furniture and walls, and the falling of various objects.). Additionally, our training approach is applied based on a consumer-grade VR headset (the HTC Vive) and hence lends itself well to standardized distribution.

In this paper we present an earthquake scenario to users, but in doing so, we show that more generally a virtual simulation of a disaster scenario can be used to train individuals to respond properly in the case

- *Changyang Li is with the Beijing Institute of Technology and the University of Massachusetts Boston. E-mail: 1120131747@bit.edu.cn.*
- *Wei Liang is with the Beijing Institute of Technology. E-mail: liangwei@bit.edu.cn.*
- *Chris Quigley is with the University of Massachusetts Boston. E-mail: Christopher.Quigl001@umb.edu.*
- *Yibiao Zhao is with the Massachusetts Institute of Technology. E-mail: ybz@mit.edu.*
- *Lap-Fai Yu is with the University of Massachusetts Boston. E-mail: craigyu@cs.umb.edu.*

Manuscript received 19 Sept. 2016; accepted 10 Jan. 2017.
Date of publication 26 Jan. 2017; date of current version 18 Mar. 2017.
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TVCG.2017.2656958

of a real emergency. This may provide insights to future researchers and developers who wish to create virtual training scenarios for other types of emergency situation. The major contributions of our work include the following:

- Demonstrating that virtual environments based on a consumer-grade VR headset can be used for earthquake safety training.
- Providing the technical details about how such virtual environments can be modeled and how user interaction can be designed, to enable realistic earthquake simulation in indoor scenes for training purposes.
- Evaluating the effectiveness of our approach and comparing it with other training methods.

2 Related Work

We provide a succinct overview of the traditional earthquake safety training approaches and review previous efforts in using virtual environments for different training purposes.

2.1 Traditional Earthquake Safety Training

We focus our discussion on safety training for common indoor spaces, which our approach focuses on. Studies found that, during an earthquake, the greatest potential danger present to someone in a room is getting hit by falling or flying objects (e.g., light fixtures, mirrors, hanging decorations) [17, 41], or heavy furniture that could fall (e.g., high shelves, bookcases, cabinets). A sudden and intense earthquake shaking of several feet per second can easily cause unsecured object to topple, fall or become airborne. In fact, studies [13, 17] found that it is more likely for someone to get injured by the falling objects than to get killed in a collapsed building, providing that the building was constructed following seismic code regulations. Therefore, the skills to quickly assess the potential falling risks of different objects and identify a safe spot are keys to avoid major injuries during an earthquake, which our approach focuses on training the user with.

One common technique to reduce chances of injuries during an earthquake is to apply the "drop, cover and hold on" strategy [13,17,41] to protect oneself: *drop* means quickly moving to a spot safe from falling objects and then dropping to the floor; *cover* means protecting the head and neck, the critical and vulnerable body regions, with arms and hands; it is also advisable to take shelter under a sturdy desk or table if there is one nearby; *hold* means holding onto the shelter until the shaking stops. In our training approach, through an immersive experience, the user will learn to protect himself or herself with a similar technique. To mimic possible injuries in a real-world scenario, our approach computes the injuries caused by falling objects hitting the user's body in the virtual environment, with the user's head and neck modeled to be more vulnerable to emphasize the importance of protecting these body regions by applying the *cover* step.

Traditional methods of earthquake safety training include conducting earthquake drills [22, 26, 34], reading earthquake safety manuals [6, 13] (e.g., the ShakeOut Drill Manual) and watching training videos. The goal of training is to reinforce preparedness and safe behavior, such that when an earthquake occurs, people can respond quickly without hesitating or trying to remember what they are supposed to do [22]. Our approach aims to achieve the same goal by exposing the user to a simulated earthquake in a virtual environment. The engaging, immersive experience helps the user to remember the earthquake safety techniques, which they can apply in a new earthquake scenario, as we show in our evaluation experiments.

2.2 Virtual Environments for Safety Training

The increasingly widespread use of virtual reality devices demonstrates its great potential in various fields, such as for medical [1, 27, 36] and safety training purposes. We discuss some of the recent work. For instance, simulated virtual environments have been used for teaching pedestrian and road safety. Schwebel et al. [29] and McComas et al. [20] used virtual environments to train children to cross roads safely. Child participants were asked to go through the training in virtual environments and then their road crossing behavior was tracked in the real world. Results showed that using virtual reality for such training is highly effective. On the other hand, Backlund et al. [3] developed a serious game similar to a driving simulator to teach safe driving skills. One advantage of using game-based simulated environments for training is that they generally appeal to the participants (especially children), making them more engaged in the training process as compared to traditional training methods such as reading training manuals or watching training videos. We also devise our approach in a serious game setting in order to make the training process more engaging to users.

Virtual reality has also been used for studying human evacuation behavior and for evacuation training in emergency conditions. Some of these training applications are targeted for professional practitioners. For example, virtual reality has been used for firefighter training and simulation [2, 7, 38]. Other training applications target the general public, to teach how to escape from an emergency condition. For example, virtual reality has been used for teaching people to evacuate during a fire accident [21,39]. A major advantage of using virtual reality and simulation for training is that it enables practice under hazardous conditions. In our current approach, we focus on training people how to protect themselves during an earthquake.

An important consideration in using virtual environments for training is whether the knowledge learned in virtual environments can be transferred to tackle similar real-world scenarios. Such knowledge transfer is demonstrated to be possible in previous work. For instance, in a study on pedestrian safety [20] which utilized a virtual reality training regimen for training school children to cross streets, it was found that training in the virtual environments led to significant improvement in real-world street-crossing behavior. Another study on using virtual reality for teaching fire evacuation skills [23] also found the knowledge transfer effective: at a follow-up test, all the training participants successfully completed each of the taught safety steps in a real world simulation. Recently, Chittaro and Buttussi conducted an interesting study [10] to compare knowledge retention of teaching aviation safety through an immersive virtual environment versus a traditional training method (using safety cards). Their results show that training through an immersive environment leads to more superior knowledge retention. These findings motivate us to proceed under a similar assumption that knowledge transfer from virtual environments to real-world environments is feasible. We evaluate the performance of the users who have received training in a follow-up simulation test.

2.3 Earthquake Simulations in Virtual Environments

Compared to other virtual reality training applications, using virtual environments to perform earthquake simulation and training is less frequently attempted. Tarnanas and Manos [37] used virtual reality to teach pre-school children and children with Down Syndrome to cope with emergencies, where a virtual earthquake was used as a showcase. Sinha et al. [35] described an approach for generating an earthquake disaster scenario in a 3D environment. Since the focus of their approach is to provide a realistic visualization of an earthquake rather than an interactive training experience, in their approach, the camera path of the user is scripted and fixed, and there is no interaction between the user and the objects in the environment. Very recently, a company called PulseVR released a demo video showing how virtual reality can be used to hint to people about the safety precautions to take before and during an earthquake [25], in a step-by-step manner. Compared to the previous work, our approach focuses on providing an highly interactive training experience in the guise of a serious game. The user needs to figure out the paths to take and poses to make in order to minimize injury, which will be tracked by our setup to evaluate the user's success. By enabling rich user interactions with the virtual environment, we believe our approach will give the user a more engaging learning experience.

3 Overview

The goal of our work is to provide an earthquake safety training approach by consumer-grade virtual reality technology. The user learns effective observation, navigation and self-protection skills through a realistic earthquake simulation in an immersive virtual environment. In particular, our approach makes use of the HTC Vive virtual reality device, which allows the user to navigate in a virtual environment and manipulate virtual objects through two hand motion controllers, while it closely tracks the user's head and hand positions. Figure 2 shows an office scene which we use to illustrate our approach. A virtual earthquake is simulated in the scene, and the user's goal in the simulation is to protect himself from injuries (e.g., due to falling objects) by navigating and posing himself appropriately. A human model is used to represent the user in the simulation, with different colliders added for collision detections with virtual objects based on which the level of injury is computed.

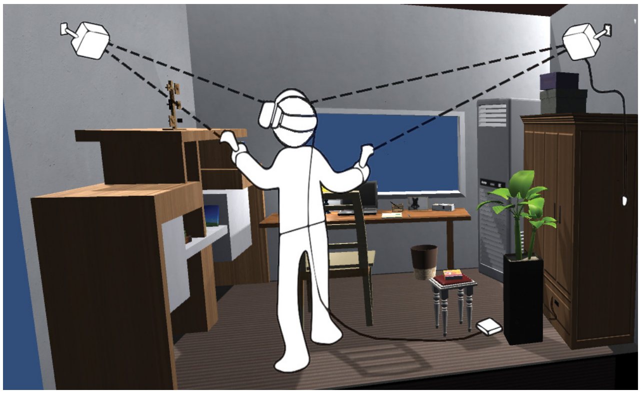

Fig. 2: An office scene used as an illustrative example of our approach. The player experiences a virtual earthquake through the HTC Vive.

4 Technical Approach

Our approach consists of three major components: virtual environment modeling, human model and physics simulation. We provide technical details of each component in the following sections.

4.1 Virtual Environment Modeling

We construct the virtual environments in Unity 5. The rooms and objects are represented as 3D meshes. We create three types of scenes: dining rooms, living rooms and offices, based on the assumption that self-protection strategies may vary with the scene type, considering the fact that each type of scenes is associated with some typical objects and layouts. For example, it might be a good strategy to hide under a dining table in a dining room during an earthquake. However, as tables are uncommon in a bedroom, strategies to protect oneself in a bedroom could be quite different. We show in our supplementary material the numbers of different types of objects in different scenes used in our experiments. In the scenes we used, living rooms tend to have more props, while offices tend to have less breakable objects. Table 1 shows the amount of physical damage the participants experienced in different types of scenes in our experiments. As shown, the participants could be more vulnerable to physical damage in certain types of scenes (e.g., dining rooms). Therefore we analyze user performance separately in different types of scenes.

For each room, we place furniture and objects commonly available in a room of the corresponding scene type according to scene statistics from the SUN Database [40], like in the work of the Clutterpalette [44]. For example, an office scene is populated with desks, computers and books. A living room usually has a television, a couch and a lamp.

To present the objects in the virtual environments realistically, we scale the objects to realistic dimensions manually. Alternatively, an automatic scaling technique [28] could be applied. The objects are also assigned with materials and physical properties, e.g., masses, such that Newtonian physics can be applied for realistic simulation using Unity's built-in physics engine.

4.1.1 Objects

Figure 4 shows examples of different types of objects used in our scenes. The objects can be classified into three categories: *Structures*, *Furniture* and *Props*, following conventions in previous scene modeling [43, 44] and scene understanding work such as the NYU Kinect Dataset [33]. We provide more details for each category:

- **Structures.** These refer to the objects used to construct the room, including floor, walls, columns and ceiling. Similar to previous work [19, 43], we organize the structure objects hierarchically, with the floor being the root, and the walls and the ceiling being its children. When an earthquake is simulated, the walls and ceiling shake together with the floor. For simplicity, we do not consider the collapse of structures due to a very strong earthquake. Therefore, in our rooms the structures are always attached to each other.
  As we use the HTC Vive for our experiments, we create rooms with a rectangular floor of 3m $\times$ 4m following the space specifications of the HTC Vive's play area. The height of a room is set as 3m, similar to that of common apartments.
- **Furniture.** These refer to the movable objects that generally lie on top of the floor, such as couches, chairs, tables, cabinets and bookcases. These objects may move if acted upon by a strong enough force, but most of them are relatively stable in an earthquake due to their heavy weights. A big and sturdy piece of furniture can sometimes serve as a good shelter to protect people from getting hit by falling clutter objects, which is the reason why people are suggested to take shelter under a table following the "drop, cover and hold on" self-protection strategy [13, 17, 41]. The user can apply a similar strategy to protect himself during a simulation.
- **Props.** These refer to the small, movable objects that are generally placed on top of a furniture object. Examples include cups and plates on a table, mobile phones and laptops on a desk, and books on a bookshelf. As these objects are generally small and light, they can easily fall when pushed by a force. Depending on the shape and material of the objects, falling props can sometimes cause considerable physical damage. For example, getting hit in the head by a falling, sharp objects such as a pair of scissors or a knife is definitely dangerous. The user will learn to avoid and protect his head by his arms from dangerous falling objects in the training process.
  Some of the props are hanging on a wall or from a ceiling instead of lying on a piece of furniture, similar to some of the props in the NYU Kinect dataset. Examples include paintings and televisions attached to a wall and chandeliers hanging from a ceiling. For simplicity, for these kinds of *hanging* props, our approach assumes that the connector holding the prop will be broken if it experiences a force larger than a certain threshold (two times the estimated weight of the prop) and the prop will fall down due to gravity. As noted in earthquake safety literature [13, 41], getting hit by these types of falling objects is a common cause of injury during an earthquake, and the user will learn to avoid them.

(a) Input Scene

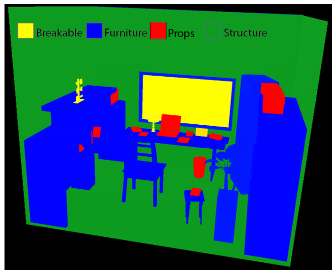

(b) Object Type

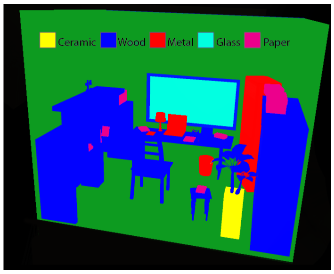

(c) Material Type

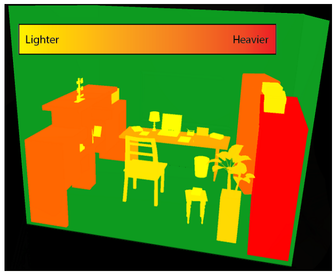

(d) Mass

Fig. 3: (a) Input scene and color maps of the scene by (b) object type, (c) material type and (d) mass.

Figure 3 visualizes the object types, material types and masses of different objects in the illustrative scene by color maps.

4.1.2 Material

To enable realistic physics simulation which is discussed in Section 4.3, each object is assigned a material. For example, a bottle is assigned with the material "glass" and a chair is assigned with the material "wood". We use eight materials in our scenes: metal, glass, ceramics, wood, plastic, fiber, leather and paper. We obtain the material for each object type from OpenSurfaces [4], which stores the common materials for objects in common real-world indoor scenes. To facilitate the material assignment process, we automatically assigned each object with a common material ranked within top three in OpenSurfaces. If an object should be composed of multiple materials (e.g., a potted tree is composed of "ceramics" and "wood"), the object is divided into several smaller objects attached together, each of which is assigned with an appropriate material.

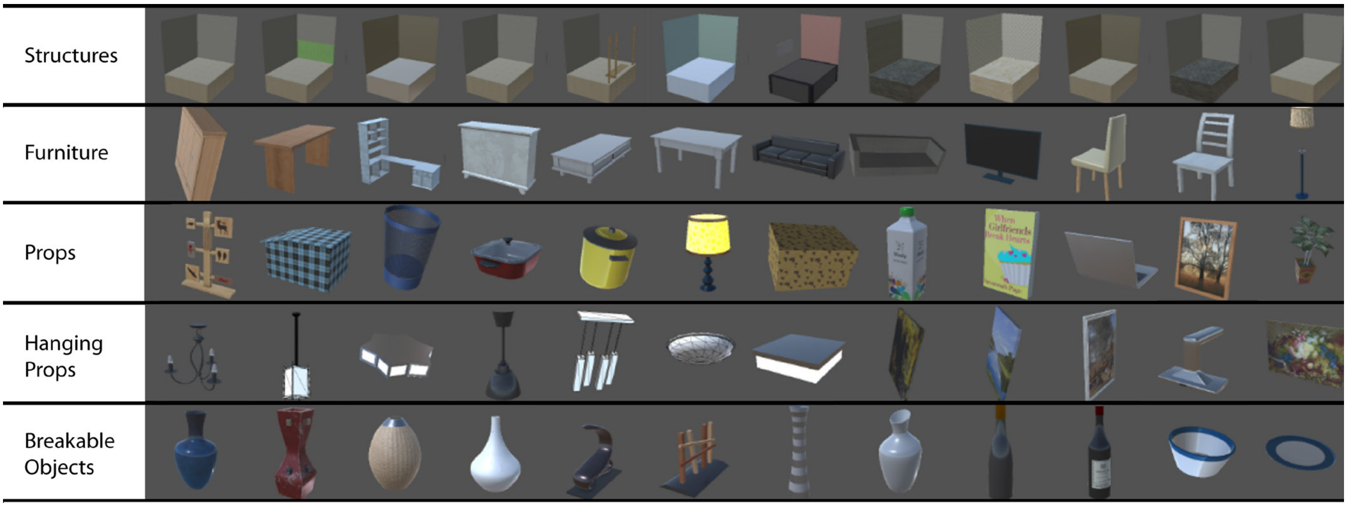

Fig. 4: Example objects used in our scenes.

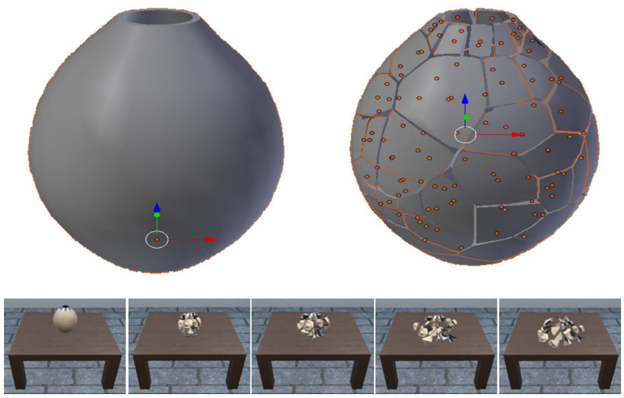

Fig. 6: Top: the fragments of a breakable, ceramic pot are precomputed using the cell fracturing method. Bottom: this example shows how the pot breaks during a simulation.

4.1.3 Mass

Our approach computes an approximate mass for the object for realistic physics simulation based on the object's material. First, the volume of the object is computed by summing the signed volumes of the constituent tetrahedrons of the object's mesh [45]. Then our approach multiplies the volume by the material's density (looked up from [24]) to compute the object's mass. Note that the masses of the structures (i.e., the floor, walls and ceiling) are set to be very large such that they are only moved by the shake of the earthquake.

4.1.4 Breakable Objects

We set the objects made with certain types of materials, such as glass and ceramics, to be *breakable* to enhance realism in calculating the physical damage these objects might cause. For example, a glass bottle falling off a shelf and hitting the floor can break into pieces.

To use breakable objects in our simulation, we precompute how these breakable objects may break into fragments using the cell fracturing method provided by Blender. During a simulation, when a breakable object collides with another object and the impulse exceeds a certain threshold, it will break into the precomputed fragments, which will then follow Newtonian physics to fly and fall in the scene. Figure 6 illustrates how a ceramic vase is modeled and how it is broken during a simulation. As in the real world, getting hit by these sharp fragments (e.g., a piece of glass) can cause serious injury. We discuss how to compute physical damage in Section 4.3.

4.2 Human Model

We describe the virtual human model which represents the user in the virtual environment. During a simulation, the user controls his viewpoint through the HTC Vive headset, and his hands through the hand motion controllers. His head and hands locations are used to control the pose of the virtual character, whose body parts are attached with colliders for detecting collision with objects in the virtual environment, in order for our approach to compute the physical damage that has been incurred on the body (e.g., due to a falling object hitting the body).

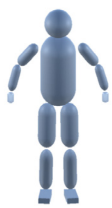

Fig. 5: The human model used to represent the user.

4.2.1 Representation

Figure 5 depicts the virtual human model that we use. It consists of twelve body parts: head, torso, upper arms, fore arms, thighs, calves and feet. Each body part may collide with objects in the environment. To enable collision detection, each body part is associated with a collider that approximates its shape. We use a sphere collider for the head, box colliders for the feet, and capsule colliders for all other body parts. Note that the hands correspond to the HTC Vive hand motion controllers. The user will frequently use his hands to manipulate objects in the scene (e.g., grasping an object, pushing an object), and our approach does not consider such “collision” of his hands with objects in the scene as a physical damage.

4.2.2 Durability

Each body part is given a durability value $\mu \in [0,1]$ corresponding to how durable it is to physical damage. The more durable a body part is, the less physical damage the body part receives when hit by an object. For example, we set the arms to have a higher durability value than the head, such that the physical damage caused by a falling lamp hitting the arms is less than that caused by the lamp hitting the head. Hence it would be preferable for the user to protect his head with his arms, akin to the general advice in earthquake safety literature [13, 17, 41]. The overall physical damage caused to the body throughout the earthquake simulation will be used to account for how well the user survives an earthquake. We set the durability values of the head, the torso, the upperarms, the forearms, the thighs, the calves and the feet as 0.1, 0.6, 0.8, 0.7, 0.9, 0.8 and 0.7 respectively. Hence the head is the least durable and most critical to protect. We give more details about computing the physical damage based on these durability values in Section 4.3.

4.2.3 Locomotion and Tracking

During a simulation, the user can rotate his head to see the virtual environment from a different viewpoint via the HTC Vive headset. The user can also move his hands in the virtual environment via the hand motion controllers. At every frame, the positions and orientations of the headset and the two hand motion controllers are tracked by the HTC Vive lighthouse base stations, so that the location of the user can be estimated.

To enable the user to control the virtual human through his body motion, the virtual human's pose should resemble that of the user at every frame. The challenge is that we only have the tracked positions of

(a) Standing (b) Leaning Forward (c) Squatting

Fig. 7: Inferring the user's pose with the IK algorithm taking the tracked headset and hand motion controller positions as constraints. In each example, the user's pose is shown on the left, and the inferred pose is shown on the right. The inferred poses mimic the user's poses reasonably, allowing the user to control the virtual human model in the earthquake simulation.

the head and the two hands rather than those of the whole body at every frame. To this end, we apply the Inverse Kinematics algorithm [30] to infer the pose of the user using the tracked positions and orientations of the head and the two hands as constraints, which correspond to the head joint and the two hand joints of the virtual human. The Inverse Kinematics algorithm is applied to infer the positions of all the other joints of the virtual human, which can be solved by the Jacobian inverse technique in real time [5]. By doing this, the pose of the virtual human is updated in real time to mimic that of the user. Figure 7 shows some examples. For instance, when the user squats down and covers her head by her hands, the virtual human poses itself similarly. The estimated pose provides a fairly accurate approximation for controlling the virtual human.

4.2.4 Object Manipulation

The user can manipulate the objects in the scene during a simulation. For example, he can push a chair aside; he can also grab a chair and put it elsewhere. For simplicity, our approach assumes that the user pushes with a force of 400N and lifts with a force of 100N with a single hand. These parameters can be alternatively set based on the strength of the person that the virtual human represents. To apply a force to a virtual object, the user can simply push his hands against the target object. If the force is large enough to overcome the static friction, the object will start to slide. Similarly, to lift an object, the user can grab the top or the sides of the object with his hands while holding the triggers of the hand motion controllers to signal the intention to lift. The object will be lifted if the lifting force overcomes the object's weight.

If the user's force is not large enough to overcome the static friction (when pushing) or the object's weight (when lifting), the collisions between the object and the user's hands will simply be ignored by our system. On one hand, the user is not supposed to be able to move heavy objects. On the other hand, the user will not be blocked by virtual objects due to the absence of haptic feedback in reality.

Following this manipulation model, the user can push or grab a light object (e.g., a chair) but not a heavy object (e.g., a cupboard) in the scene, similarly as in the real world. Whenever an object is manipulated by the user, the HTC Vive controller held by the hand manipulating the object vibrates slightly to notify the user of the manipulation.

4.3 Physics Simulation

We simulate an earthquake by shaking the floor according to historical earthquake data from the real world. The shake propagates from the floor to all the other objects in the scene according to Newtonian mechanics computed by the Unity's physics engine. If an object hits the virtual human representing the user, our approach will compute the physical damage caused by the hit. The physical damage will be used as a metric to evaluate how well the user has survived the earthquake.

4.3.1 Earthquake Simulation

Our approach assumes that the source of the earthquake is far away from the room. Therefore the whole floor shakes as a single entity along the same direction at any time. To simulate a realistic earthquake, we apply the historical earthquake data provided by the PEER Ground Motion database [9]. In our experiments, we use the data of the 1952 Kern County earthquake, which occurred at Los Angeles with a magnitude of 7.3. Figure 9 shows the velocity of the ground during the earthquake. The three plots show the shaking speeds along the x, y (vertical) and z axes respectively. The data is given every 0.005 second, with a total duration of about 70 seconds. The data is used to set the velocity of the floor over time in our simulation. The motion of the objects in the scene are then computed and updated by the physics engine. Note that other real-world earthquake data can be downloaded from the database and can be used to generate a corresponding earthquake simulation.

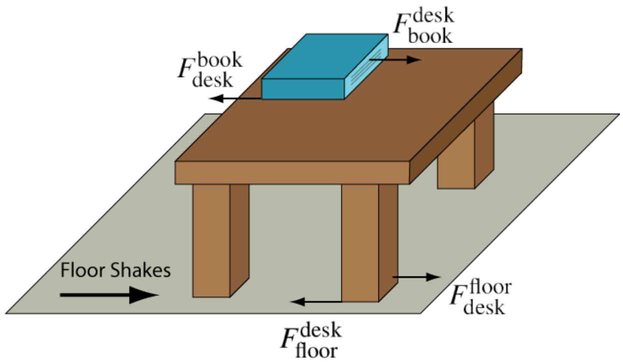

Fig. 8: A shake propagates from the floor to the objects in the room.

Figure 8 illustrates how a shake propagates from the floor to the objects in the room. Consider an object standing on another object (e.g., a desk standing on the floor). The objects' static and kinetic frictional coefficients are set according to their assigned materials [24] (Section 4.1). The physics simulation engine makes use of these coefficients when computing the movement of objects. The static frictional coefficient determines how much force is needed to overcome the static friction force such that the object on top starts sliding. The kinetic frictional coefficient determines how much kinetic friction force the object experiences while sliding. In this example, as the floor moves to the right, the desk exerts a kinetic friction force $F^{\text{desk}}_{\text{floor}}$ to the floor and receives a reaction force $F^{\text{floor}}_{\text{desk}}$ exerted by the floor pushing it to the right. Similarly, suppose the desk has overcome the static friction force; as the desk slides to the right, the book exerts a kinetic friction force $F^{\text{book}}_{\text{desk}}$ to the desk and receives a reaction force $F^{\text{desk}}_{\text{book}}$ exerted by the desk pushing it to the right. The above results in a chain reaction of motion of objects, which is computed by the physics engine. As the book slides over the edge of the desk, it falls down by gravity.

4.3.2 Computing Physical Damage

We describe how our approach computes the physical damage caused to the user when he is hit by a virtual object. Suppose the user's shoulder is hit by a flying vase from the front in the virtual environment. Realistically, the user should be pushed backwards by the vase. However, because the current virtual reality setup does not provide such haptic feedback, the user will not feel any pushing force at his shoulder and will not be pushed backwards in reality. In other words, the user will just stay at the same location after being hit. Our approach simply assumes that all kinetic energy of the object is passed to the user during the hit. The physical damage D is computed by $D = (1-\mu)s|\mathbf{J}|$. $|\mathbf{J}|$ is the magnitude of the impulse of the collision computed by the physics engine. μ is the durability value of the body part being hit. Note that the less durable the body part being hit is, the more physical damage the hit will cause. $s = 9s_f + 1$ relates the physical damage with how sharp the hitting object is, where $s_f \in [0,1]$ is the sharpness value of

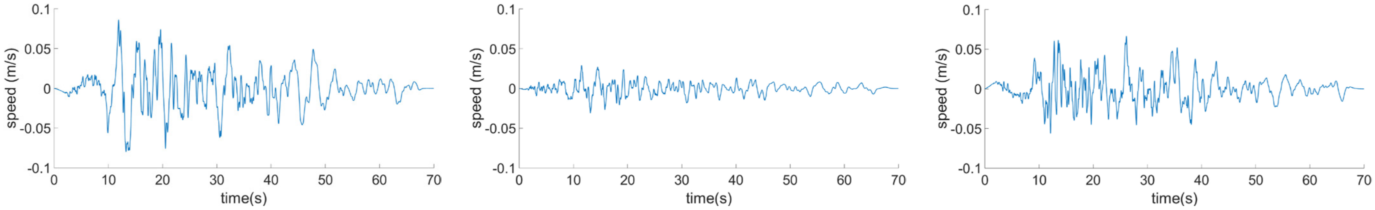

(a) x-axis (b) y-axis (c) z-axis

Fig. 9: Shaking speeds of the ground over time along the x, y, and z axes based on the Kern County earthquake data.

the object's face f that hits the human body and is computed by the approach of Chen and Cheng [8] based on face normals. For example, a sharp fragment hitting the human body will cause more damage. The physical damage of a hit is accumulated to calculate the overall physical damage throughout the earthquake simulation.

5 Experiments

5.1 Implementation

We implemented our approach using C# and Unity 5. We ran our experiments on a PC equipped with 16GB of RAM, an Nvidia Titan X graphics card with 12GB of memory, and a 2.60GHz Intel i7-5820K processor. The program ran steadily at about 60 frames per second. The user experienced the simulation via the HTC Vive in an empty space of 3m $\times$ 4m, the largest play area it allows. We include a picture of our setup in the supplementary material. Please refer to our supplementary video for a demo of the training process.

5.2 Scene Data

To test our approach, we created 12 scenes, which include 4 scenes for each of the 3 scene types: *Dining Room*, *Living Room* and *Office*. Please refer to Figure 10 for the screenshots of the scenes and the supplementary material for the scene statistics.

Fig. 10: The scenes used in our experiments.

5.3 Evaluation

We evaluate the effectiveness of our virtual reality training approach. Specifically, we want to test how well the users perform in a simulated earthquake in a new scene after four different training conditions:

1) *VR*: receiving training through out virtual reality approach;
2) *Video*: receiving training through watching an earthquake safety training video;
3) *Manual*: receiving training through reading a manual about earthquake safety in an indoor environment;
4) *None*: receiving no training.

5.3.1 Participants

We recruited 96 participants, whose ages ranged from 20 to 30. They were undergraduate and graduate students from different majors. The participants were randomly divided into 4 groups of 24 people, with each group corresponding to a training condition described above. For each group, 8 participants were randomly assigned to each of the 3 scene types: *Living Room*, *Dining Room* and *Office*.

5.3.2 Training

We describe the training procedure under each of the training conditions. For the *VR* group, each participant was asked to train with scenes 1, 2 and 3 (refer to Figure 10) of the scene type he was assigned to. In each training, the participant went through an earthquake simulation, where his objective was to avoid getting hurt in the scene as best as he could. He was told that his head was the most vulnerable and hence he should try his best to protect it. Note that no specific strategy on how to perform well, such as hiding under a desk or holding something for self-protection, was taught. The participants in this group were supposed to come up with self-protection strategies in the training process. After training in a scene, they were informed of their performance (in terms of physical damages received) so that they could evaluate their strategies according to the results. For the *Video* group, each participant was asked to watch an earthquake safety training video provided by the Southern California Earthquake Center. The video showed the safety steps to take during an earthquake in an indoor environment, including demonstration of the "drop, cover and hold on" technique. The participant could watch the video for as many times as he wanted to remember the details of the instructions. For the *Manual* group, each participant was asked to read an earthquake safety training manual from the Earthquake Country Alliance [12]. It provided details and pictorial illustrations about the safety steps to take during an earthquake. For example, it described the "drop, cover and hold on" technique and also mentioned about protecting the head and neck from falling objects. The participant was asked to read the manual carefully such that he understood the steps to protect himself during an earthquake. He could read the manual for as long as he wanted. For the *None* group, the participants did not undergo any training.

5.3.3 Tests

Each participant was asked to do two tests where he would try protecting himself from injury throughout an earthquake simulation in a virtual environment. The first test was done right after the training. The second test was done a week after the training. Each participant was asked to do the tests in scene 4 (Figure 10) of the scene type he was assigned to. Note that these scenes were not used for training.

To eliminate the potential bias towards the participants who underwent the Virtual Reality training, due to the fact that these people had immersive experiences before the tests while others did not, we included a familiarization process for all participants. Before a test, no matter which training condition the participant went through, he was asked to familiarize himself with the control of the HTC Vive device in a warm-up session. He was told how to use the HTC Vive headset and controllers, and how he could navigate in the scene and manipulate objects similarly as in the upcoming test. He was asked to play with the device in a living room scene (not used in training or testing). Note that no earthquake simulation was done in this scene.

(a) Results Immediately after the Training

Scene Type	Training Condition	P1	P2	P3	P4	P5	P6	P7	P8	Mean	Median	Standard Deviation
Living Room	VR	43	19	55	104	123	67	38	24	**59**	43	**34.86**
Living Room	Video	192	51	319	32	18	153	29	32	103	**32**	101.59
Living Room	Manual	348	238	6	217	246	295	46	101	187	217	114.35
Living Room	None	389	161	146	365	360	222	427	55	266	222	128.24
Dining Room	VR	35	49	0	9	0	32	138	50	**39**	**32**	**41.88**
Dining Room	Video	141	0	1	146	167	60	119	51	86	60	62.05
Dining Room	Manual	15	18	377	72	43	0	228	81	104	43	123.14
Dining Room	None	370	80	28	284	2	351	34	255	176	80	140.80
Office	VR	21	69	4	0	1	3	19	102	**27**	**4**	35.35
Office	Video	33	94	1	26	3	78	46	65	43	33	**31.76**
Office	Manual	225	34	41	67	85	294	16	0	95	41	99.55
Office	None	61	73	121	382	150	23	166	71	131	73	105.01

(b) Results One Week after the Training

Scene Type	Training Condition	P1	P2	P3	P4	P5	P6	P7	P8	Mean	Median	Standard Deviation
Living Room	VR	111	0	60	14	46	34	118	28	**51**	**34**	**40.28**
Living Room	Video	41	8	305	122	332	150	176	48	148	122	112.37
Living Room	Manual	13	415	289	52	348	297	25	125	196	125	149.34
Living Room	None	449	133	247	325	47	141	322	261	241	247	120.74
Dining Room	VR	48	20	43	159	0	26	63	0	**45**	**26**	**47.92**
Dining Room	Video	125	208	53	6	0	142	10	298	105	53	101.36
Dining Room	Manual	33	86	102	202	77	319	81	42	118	81	89.99
Dining Room	None	281	16	115	161	146	347	320	55	180	146	115.08
Office	VR	25	1	19	0	41	20	78	74	**32**	**20**	**28.08**
Office	Video	1	44	115	0	7	143	67	103	60	44	52.39
Office	Manual	212	75	33	233	25	110	156	37	110	75	76.94
Office	None	127	99	333	157	169	13	64	58	128	99	91.93

Table 1: Physical damage results of the tests conducted immediately and one week after the training. For each scene type and training condition, the results of the 8 participants, and the mean, median and standard deviation of the results are shown. A smaller physical damage value refers to a better performance. The smallest mean, median and standard deviation of each scene type are in bold

The participant was asked to play with the device for as long as he wanted, until he confirmed that he was familiar and comfortable with the control. Therefore, different groups of participants should have the same level of familiarity with the control before they took the tests, and differences in performance among different groups should mainly be attributed to how effectively the participants learned under their safety training conditions. This familiarization process typically took about 5 to 10 minutes.

In each test, the participant started at a predefined position in the open space of the scene. After about 10 seconds, an earthquake simulation would start, which lasted for about 70 seconds. The participant was asked to protect himself similarly as he would in a real earthquake.

5.3.4 Metrics

We collected the following metrics to evaluate and analyze the performance of the participants in the tests:

1) *Physical Damage*: We recorded the physical damage the participants received during the earthquake simulation as described in Section 4.3. The less physical damage the participant received, the better he was at surviving the simulated earthquake.
2) *Visual Attention*: We also tracked the participant's visual attention in the scene to analyze if he was aware of the potential dangers during the simulated earthquake. To achieve this, our approach tracked the headset position and orientation at every time frame. To determine the objects the participant was looking at, a ray was cast from the headset position along the direction where the headset was facing. We assume that an object was being noticed it was within a circular cone centered about this ray, with the cone's apex aligned with the headset position and the apex angle set as 60 degree to mimic the near peripheral vision of a human [14] (see Figure 11).

We want to check how well the participant could notice the dangerous objects around him at each moment during the earthquake. More specifically, we consider an object as dangerous if it would fall within 0.5 meter from the participant in the next 2 seconds according to the simulation, had the participant stayed at his spot. Following the above definitions, our approach counted the percentage k of dangerous objects noticed by the participant at all time frames of the simulation.

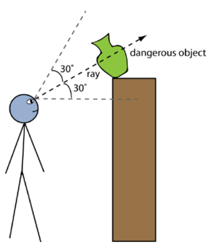

Fig. 11: A dangerous object noticed.

6 Results and Discussion

6.1 Physical Damage

Figure 12 shows the physical damage received throughout the test by participants trained under different conditions. For each test scene, each participant's result, and the mean and median of the results under each training condition, are shown. Table 1 shows the numeric results.

In general, in terms of the means and medians of the results, the participants who went through the virtual reality training received the least amount of physical damage throughout the test, followed by those who were trained with a video, trained with a safety manual and untrained. The results suggest that the virtual reality training approach is more effective than the other approaches in terms of the physical damage metric. For the *Living Room* scene, the participants who were untrained performed particularly badly compared to those who went through any other form of training. For the *Dining Room* scene. the differences in physical damage of participants trained with different conditions are not substantial, though training with the virtual reality approach achieves the least amount of physical damage in general. For the *Office* scene, training with a safety manual does not seem to be effective, as those participants trained with the manual achieved similar performance as those who were untrained.

It is interesting to look at the standard deviations of the results in Table 1. In general, the participants trained with the VR approach performed consistently better as reflected by smaller standard deviations, while those trained under other conditions had more fluctuating performance as reflected by larger standard deviations. This may suggest that the good performance of the participants trained under other conditions attribute more to the participants' prior individual skills in surviving a simulated earthquake which are not necessarily learned from the training. For example, as the results of the *Dining Room* show, there are a few untrained participants who could even considerably outperform the other participants who underwent a training.

We believe that the immersive experience provided by the virtual reality training approach may account for its higher effectiveness compared to the traditional training approaches based on a safety manual or a video, which may suffer from a gap between theory and practice. For example, a participant may learn from a safety manual that he should avoid falling objects in an earthquake, but without a real practice he may not be able to estimate which objects will probably fall during an earthquake and how he should position himself in a room to avoid them. Comparatively, a virtual reality training provides a more direct learning experience, allowing the participant to learn from practicing.

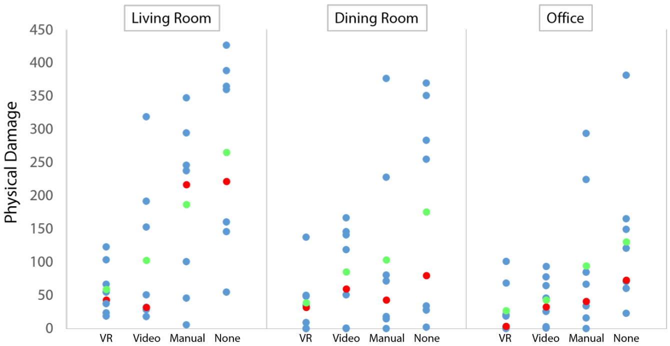

Fig. 12: Physical damage results immediately after the training under different conditions. Each blue dot refers to the result of a participant. Each green dot and each red dot respectively refer to the mean and median of the results under each training condition. The participants trained by the virtual reality approach generally received a lower physical damage.

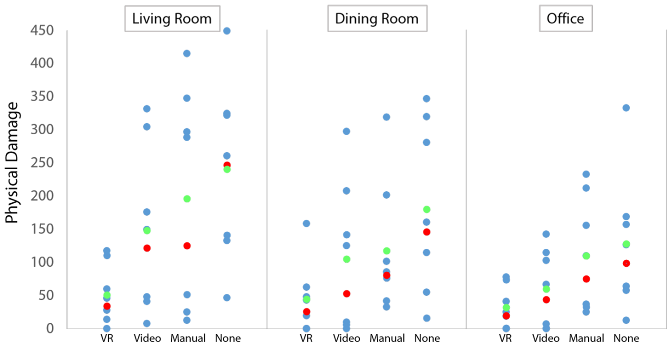

Fig. 13: Physical damage results one week after the training.

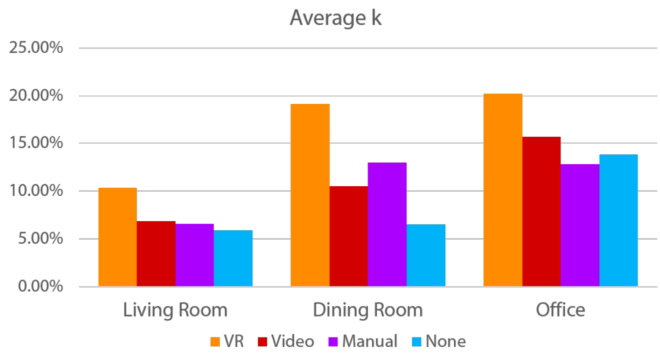

Fig. 14: Visual attention results immediately after the training. The plots show the average percentage *k* of dangerous objects noticed by the participants in all time frames of the simulation. Participants trained by the virtual reality approach are more attentive to dangerous objects that can potentially hit them during the simulation.

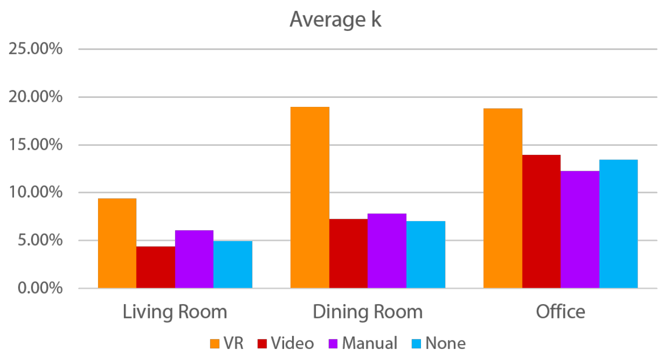

Fig. 15: Visual attention results one week after the training.

6.2 Visual Attention

Figure 14 shows the average percentage *k* of dangerous objects noticed by the participants in all time frames of the simulation. For each test scene, the results of the participants trained under different conditions are shown.

As the results show, the participants trained with the virtual reality approach achieved the best performance in noticing the dangerous objects, which would hit on them shortly had they stayed in the same location. By taking the VR training, the participants learned to be more attentive to the potential danger around them during an earthquake, which might help them to move more effectively to avoid injury.

It may seem counter-intuitive that the highest value of *k* achieved is only about 20%. We provide further clarification. First, recall that we define *k* as the number of dangerous objects noticeable in all time frames during the simulation. Suppose a participant notices a dangerous object at a certain time frame. At the next time frame, he may turn his head to look at the other side of the room trying to figure out how to escape, and the dangerous object he noticed previously could be out of sight even though he knows about its existence. As our approach does not consider the participant's memory of what has already been noticed, it would simply count the dangerous object as a miss at the second time frame. This results in a relatively small *k*. We decide not to consider the participant's memory in defining *k* to keep the definition simple and intuitive. Second, according to our definition, an object is only counted as being noticed if it falls within the peripheral vision of the participant, which just spans 60 degrees (out of 360 degrees). Therefore, at any particular time frame, the region of the scene that can be noticed by the participant is quite limited, and it is expected that not many dangerous objects can be noticed simultaneously.

6.3 Re-testing

To investigate how well the participants retain the knowledge they learned under different training conditions, we conducted the test again one week after the training session.

Figure 13 shows the new physical damage results. Please also refer to Table 1 for the numeric results. The participants who underwent the virtual reality training achieved similar performance as they did one week ago. The performances of the participants trained by a video or a safety manual dropped in general, while they still showed some improvement over those who were untrained.

Figure 15 shows the visual attention results. The participants who underwent the virtual reality training maintained almost the same level of visual attention to dangerous objects. The participants who underwent a video or safety manual training became less attentive to dangerous objects, and their levels of visual attention almost fell to the level of those who were untrained.

According to the results, the immersive experience of the virtual reality training approach seems to help the participants to form a more long-lasting memory about how to detect and avoid the potential danger in an earthquake, and hence we believe the virtual reality training is more effective than the other approaches compared.

6.4 User Feedback

We spoke with the participants after the evaluation experiments. With regard to the interaction experience, a majority reported that it was intuitive to navigate in the scene and to manipulate objects. Some participants felt that the absence of haptic feedback when manipulating objects rendered the interaction unrealistic. For example, a participant complained that the feeling of pushing a heavy chair was unrealistic because in reality he would have needed to push really hard to move the chair, and he should have felt a large reaction force from the chair. We believe advancement of haptics technology such as haptic gloves for virtual reality will enhance the interaction experience.

Several participants reported that they were dazzled by the shaking of the objects and felt a bit dizzy during the simulation, but the immersive experience was still tolerable as it only lasted for about one minute. A few participants commented that the earthquake simulation appeared scary at times, for example, when a light suddenly fell down or when a window broke into pieces. We believe that such psychological effects may affect a participant's performance. When a participant feels scared, he may not be able to stay calm and make rational navigation choices. This problem also occurs in a real-world earthquake.

With regard to the realism of the earthquake simulation, most of the participants reported that the simulation felt realistic to them, because the objects fell down in the scene as they would expect during a real earthquake. We attribute the realism to the realistic physics simulation we employed in our approach. A few participants complained that the breaking of some objects was not realistic. For example, the breaking of a window should have depended on where another object struck the window (our approach assumed that the window was struck at its center). While physics-based simulation of fracturing could be expensive and difficult to employ in real time, a simple workaround may enhance realism. For example, our approach may generate multiple fracturing for a breakable object in a pre-processing step, and apply the closest one based on the point of strike during the simulation. Finally, most participants who underwent the virtual reality training found the training process interesting and engaging; the immersive training let them experience how an earthquake might feel like and made them aware of the potential dangers (e.g., falling objects) to watch out during an earthquake. Most of them commented that the training experience still lingered well when they did the test one week after training.

7 SUMMARY

We introduced a virtual reality-based approach for earthquake safety training, which exposes a user to simulated earthquakes in realistically-modeled scenes. The evaluation experiments show that the proposed virtual reality training approach can more effectively train a user to avoid physical damage and to be aware of potential danger in a simulated earthquake, compared to traditional training approaches by a safety manual or a video.

As the participants of our evaluation experiments did not go through a real earthquake, we cannot firmly conclude about their performance in a real earthquake under different training conditions. However, we believe that our evaluation experiments and results are still meaningful and indicative, because our experiments were conducted in realistically-modeled scenes and the physics simulation was also realistic. The simulated earthquake was generated using real-world earthquake data. The physical damage evaluation metric was based on the measure of how often a participant got hit and the durability of the body parts being hit, which should correspond to the amount of injury the participant would experience in a real earthquake.

We believe there are several major benefits of using a virtual reality training approach over traditional approaches. First, the realistic, immersive training experience provided by virtual reality allows the participant to learn by practicing directly, hence avoiding the gap between theory and practice in traditional training approaches. We believe that our virtual reality training approach can complement traditional training approaches. Second, a virtual reality training is very engaging to the participant as he no longer sees the outside world during the training session and has to deal with the virtual danger presented seriously. This is in contrast to traditional training approaches of having a participant read a safety manual or watch a video, where the participant can easily get distracted and may not remember the details of the training material. Third, different from safety manuals and videos which lack interactivity, a virtual reality training approach which features user interaction is often more appealing to the participant. By conducting the virtual reality training in the form of a serious game, the participant may feel more enthusiastic and motivated about doing the training.

7.1 Limitations

We discuss the limitations of our virtual reality training approach and some ideas for future extension.

For achieving interactivity, our approach is based on the simplified physics simulations provided by the Unity game engine rather than on highly realistic physics simulations. The latter will likely be expensive to compute, and may introduce lags and motion sickness in the interactive experience. The physics simulations provided by Unity can achieve high efficiency and frame rates; the simulations of collisions between rigid bodies are realistic, yet the simulations of soft bodies and particles are more rough and may result in inaccuracy. For example, when a soft object (e.g., a pillow) hits the user, our approach based on Unity's physics simulations could not take the elastic deformation into account when calculating the physical damage caused.

Our approach which is based on the HTC Vive does not track the full human body. Only the head and the two hands of the user are tracked, and the positions of the remaining joints are estimated by the IK algorithm. As can be seen from Figure 7, the estimated pose may not be precise. The advantage is that our setup is simple. The downside is that the imprecise pose may result in inaccurate calculation of physical damage. A consumer-grade motion capture suit such as the PrioVR suit may provide a solution for more accurate pose estimation for a household user of our system.

Our approach does not provide the user with realistic haptic feedback. For example, when an object (e.g., a cup) hits the user, the user does not feel the collision. On the other hand, when the user pushes an object (e.g., a table), he cannot feel the reaction force exerted by the object neither. This limitation may result in an unrealistic user experience. For example, the user may keep trying to push a table even though it is too heavy and infeasible to push, because he cannot judge how heavy the table is without feeling the reaction force. Moreover, because in reality the space for conducting the experiment is empty, the user may unrealistically walk to a physical location which is occupied by an object in the virtual environment (our approach does not count this as a collision with physical damage). Some novel navigation metaphors [11] could be incorporated into our application to ensure user safety. A haptic suit that provides haptic feedback to the body may also be used to enhance realism by notifying the user of a collision with a virtual object.

In the evaluation experiments, we estimated the visual attention of the user by shooting rays based on the position and orientation of the user's headset. This assumes that the user was looking straight ahead at any time, while in reality his eyes could shift to look at things at his side. A virtual reality headset with eye-tracking capabilities can more accurately measure the user's visual attention.

Our approach does not consider building collapse. Our earthquake simulation assumes that the structures of a room such as the ceiling and the wall do not break. This assumption usually holds for buildings that are constructed strictly following construction safety regulations [6, 22], and in that case falling objects pose one of the greatest threats in an indoor environment during an earthquake, which our approach addresses. Our virtual training is conducted in a single room rather than a whole apartment. In other words, the user is not allowed to leave the room through a door or a window. We believe that this assumption is reasonable as earthquake safety practices recommend not to leave a room during an earthquake because such attempts are usually risky [12, 22] in reality.

7.2 Future Work

One possible extension is to incorporate realistic sound simulations to hint the user about what is happening during the simulated earthquake, in line with what recent research has found about the important role sound could play in an immersive virtual experience [31, 32]. For example, if a glass bottle falls off from a shelf behind the user, hits the ground and breaks into pieces, the user could hear that and sense the potential danger without seeing it. Moreover, to compensate for the absence of haptic feedback, it could be helpful to include visual hints to alert the user about his status during the simulation. For example, the user may see a flicker if hit by an object in the virtual scene.

We use synthetic scenes for training and testing. As 3D room scanning technology becomes more mature and popular, in future work it would be interesting to perform training and testing in 3D-reconstructed scenes captured from the real world [16, 18, 42]. For example, a user can do the training in a 3D-reconstructed model of his or her apartment. Such kind of training, similar in spirit to a conventional earthquake drill conducted at home, would allow the user to get well-prepared in case an earthquake occurs while he or she is at home. Our training approach can be similarly employed in a 3D-reconstructed scene.

We demonstrate that the proposed virtual reality approach can be employed for earthquake safety training. In future it would also be interesting to explore the use of a similar approach for safety training of other disasters. With the growing popularity of consumer-grade virtual reality devices, using a virtual reality approach for disaster safety training becomes both a cost-effective and scalable choice. As virtual reality technology matures, we believe that virtual drills can be widely conducted like traditional drills to minimize causalities in disasters.

ACKNOWLEDGMENTS

The authors would like to thank Ana Aravena for her help on narration. We thank Mengyao Jia for serving as the model of the illustrations in the main paper and video. We also thank Haikun Huang and Yumeng Wang for helping with our experiments. This research is supported by the National Science Foundation under award number 1565978. This research is supported by the University of Massachusetts Boston StartUp Grant P20150000029280 and by the Joseph P. Healey Research Grant Program provided by the Office of the Vice Provost for Research and Strategic Initiatives & Dean of Graduate Studies of the University of Massachusetts Boston. We also acknowledge NVIDIA Corporation for graphics card donation.

REFERENCES

[1] P. B. Andreatta, E. Maslowski, S. Petty, W. Shim, M. Marsh, T. Hall, S. Stern, and J. Frankel. Virtual reality triage training provides a viable solution for disaster-preparedness. *Academic Emergency Medicine*, 17(8):870–876, 2010. doi: 10.1111/j.1553-2712.2010.00728.x

[2] P. Backlund, H. Engstrom, C. Hammar, M. Johannesson, and M. Lebram. Sidh-a game based firefighter training simulation. In *Information Visualization, 2007. IV'07. 11th International Conference*, pp. 899–907. IEEE, 2007.

[3] P. Backlund, H. Engstrom, M. Johannesson, and M. Lebram. Games and traffic safety-an experimental study in a game-based simulation environment. In *Information Visualization, 2007. IV'07. 11th International Conference*, pp. 908–916. IEEE, 2007.

[4] S. Bell, P. Upchurch, N. Snavely, and K. Bala. Opensurfaces: A richly annotated catalog of surface appearance. *ACM Trans. Graph.*, 32(4):111:1–111:17, July 2013. doi: 10.1145/2461912.2462002

[5] S. R. Buss. Introduction to inverse kinematics with jacobian transpose, pseudoinverse and damped least squares methods. *IEEE Journal of Robotics and Automation*, 17(1-19):16, 2004.

[6] California Seismic Safety Commission. *Homeowner's Guide to Earthquake Safety*. The Commission, 1997.

[7] M. Cha, S. Han, J. Lee, and B. Choi. A virtual reality based fire training simulator integrated with fire dynamics data. *Fire Safety Journal*, 50:12–24, 2012.

[8] C.-Y. Chen and K.-Y. Cheng. A sharpness-dependent filter for recovering sharp features in repaired 3d mesh models. *IEEE Transactions on visualization and computer graphics*, 14(1):200–212, 2008.

[9] B. Chiou, R. Darragh, N. Gregor, and W. Silva. Nga project strong-motion database. *Earthquake Spectra*, 24(1):23–44, 2008.

[10] L. Chittaro and F. Buttussi. Assessing knowledge retention of an immersive serious game vs. a traditional education method in aviation safety. *IEEE Transactions on Visualization and Computer Graphics*, 21(4):529–538, 2015.

[11] G. Cirio, P. Vangorp, E. Chapoulie, M. Marchal, A. Lécuyer, and G. Drettakis. Walking in a cube: Novel metaphors for safely navigating large virtual environments in restricted real workspaces. *IEEE Transactions on visualization and computer graphics*, 18(4), 2012.

[12] Earthquake Country Alliance. Seven steps to earthquake safety. `http://earthquakecountry.org/survive/`. [Online; accessed 2-September-2016].

[13] L. B. Grant. Living with earthquakes in california: A survivor's guide. *Eos, Transactions American Geophysical Union*, 82(49):611–611, 2001.

[14] T. Grosvenor and T. P. Grosvenor. *Primary care optometry*. Elsevier Health Sciences, 2007.

[15] B. Gutenberg and C. F. Richter. Frequency of earthquakes in california. *Bulletin of the Seismological Society of America*, 34(4):186–188, 1944.

[16] B.-S. Hua, Q.-H. Pham, D. T. Nguyen, M.-K. Tran, L.-F. Yu, and S.-K. Yeung. Scenenn: a scene meshes dataset with annotations. *International Conference on 3D Vision (3DV)*, 2016.

[17] L. Jones and M. Benthien. Putting down roots in earthquake country. *Southern California Earthquake Center and US Geological Survey*, 2011.

[18] W. Liang, Y. Zhao, Y. Zhu, and S.-C. Zhu. Evaluating human cognition of containing relations with physical simulation. In *Proceedings of the 37th annual conference of the cognitive science society*, 2015.

[19] W. Liang, Y. Zhao, Y. Zhu, and S.-C. Zhu. What is where: Inferring containment relations from videos. In *25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

[20] J. McComas, M. MacKay, and J. Pivik. Effectiveness of virtual reality for teaching pedestrian safety. *CyberPsychology & Behavior*, 5(3):185–190, 2002.

[21] A. C. A. Mól, C. A. F. Jorge, and P. M. Couto. Using a game engine for vr simulations in evacuation planning. *IEEE Computer Graphics and Applications*, 28(3):6–12, 2008.

[22] Occupational Safety and Health Administration, US Department of Labor. Earthquake preparedness and response. `https://www.osha.gov/dts/earthquakes/preparedness.html`. [Online; accessed 27-August-2016].

[23] L. S. Padgett, D. Strickland, and C. D. Coles. Case study: using a virtual reality computer game to teach fire safety skills to children diagnosed with fetal alcohol syndrome. *Journal of pediatric psychology*, 31(1):65–70, 2006.

[24] E. R. Parker. *Materials data book*. McGraw-Hill. New York. New York. USA, 1967.

[25] PulseVR. Pulse vr. `http://www.pulsevr.co/`, 2016. Accessed 08-27-2016.

[26] M. Ramirez, K. Kubicek, C. Peek-Asa, and M. Wong. Accountability and assessment of emergency drill performance at schools. *Family and Community Health*, 32(2), 2009.

[27] M. Reznek, P. Harter, and T. Krummel. Virtual reality and simulation: Training the future emergency physician. *Academic Emergency Medicine*, 9(1):78–87, 2002. doi: 10.1197/aemj.9.1.78

[28] M. Savva, A. X. Chang, G. Bernstein, C. D. Manning, and P. Hanrahan. On being the right scale: Sizing large collections of 3D models. In *SIGGRAPH Asia 2014 Workshop on Indoor Scene Understanding: Where Graphics meets Vision*, 2014.

[29] D. C. Schwebel, J. Gaines, and J. Severson. Validation of virtual reality as a tool to understand and prevent child pedestrian injury. *Accident Analysis & Prevention*, 40(4):1394–1400, 2008.

[30] L. Sciavicco and B. Siciliano. *Modeling and control of robot manipulators*, vol. 8. McGraw-Hill New York, 1996.

[31] S. Sikström, A.-M. Jürgensen, M. Haghighi, D. Månsson, D. Smedilik, and T. Habekost. Self-rated attentiveness interacts with transcranial direct current stimulation and noise stimulation in reaction time in a go/no-go task. *Neural plasticity*, 2016, 2016.

[32] E. Sikstrm, N. Nilsson, A. Gtzen, and S. Serafin. *Is this bridge safe? Evaluation of Audiovisual Cues for a Walk on a Small Bridge Over a Canyon*, pp. 285–286. IEEE, United States, 2016.

[33] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012.

[34] D. M. Simpson. Earthquake drills and simulations in community-based training and preparedness programmes. *Disasters*, 26(1):55–69, 2002.

[35] R. Sinha, A. Sapre, A. Patil, A. Singhvi, M. Sathe, and V. Rathi. Earthquake disaster simulation in immersive 3d environment. In *15th World Conference on Earthquake Engineering*, 2012.

[36] S. Stansfield, D. Shawver, A. Sobel, M. Prasad, and L. Tapia. Design and implementation of a virtual reality system and its application to training medical first responders. *Presence: Teleoperators and Virtual Environments*, 9(6):524–556, 2000.

[37] I. Tarnanas and G. C. Manos. Using virtual reality to teach special populations how to cope in crisis: the case of a virtual earthquake. *Studies in Health Technology and Informatics*, 81:495–501, 2000.

[38] D. L. Tate, L. Sibert, and T. King. Virtual environments for shipboard firefighting training. In *Virtual Reality Annual International Symposium, 1997., IEEE 1997*, pp. 61–68. IEEE, 1997.

[39] M. Xi and S. P. Smith. Simulating cooperative fire evacuation training in a virtual environment using gaming technology. In *2014 IEEE Virtual Reality (VR)*, pp. 139–140. IEEE, 2014.

[40] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *Computer vision and pattern recognition (CVPR), 2010 IEEE conference on*, pp. 3485–3492. IEEE, 2010.

[41] P. Yanev and A. C. Thompson. *Peace of Mind in Earthquake Country: How to Save Your Home, Business, and Life*. Chronicle Books, 2009.

[42] L.-F. Yu, S.-K. Yeung, Y.-W. Tai, and S. Lin. Shading-based shape refinement of rgb-d images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.

[43] L. F. Yu, S. K. Yeung, C. K. Tang, D. Terzopoulos, T. F. Chan, and S. J. Osher. Make it home: automatic optimization of furniture arrangement. *ACM Transactions on Graphics (TOG)-Proceedings of ACM SIGGRAPH 2011, v. 30, no. 4, July 2011, article no. 86*, 2011.

[44] L.-F. Yu, S.-K. Yeung, and D. Terzopoulos. The clutterpalette: An interactive tool for detailing indoor scenes. *IEEE Transactions on Visualization and Computer Graphics*, 22(2):1138–1148, 2016.

[45] C. Zhang and T. Chen. Efficient feature extraction for 2d/3d objects in mesh representation. In *Image Processing, 2001. Proceedings. 2001 International Conference on*, vol. 3, pp. 935–938. IEEE, 2001.