

A close-up photograph of a woman with dark, curly hair wearing a white and black VR headset. She is looking upwards and smiling, suggesting an immersive and enjoyable virtual experience. The background is a soft, out-of-focus green, possibly representing a virtual landscape.

CAPTIVE: A New Direction for Virtual Reality Controllers

Professor Christopher G. Healey
Zeyuan Chen

CAPTIVE: A NEW DIRECTION FOR VIRTUAL REALITY CONTROLLERS

Virtual and augmented reality systems are transforming how we view and manipulate 3D objects. One of the biggest challenges to developers is to design fast, accurate, yet affordable controllers. Now, **Professor Christopher Healey and Zeyuan Chen** at North Carolina State University have announced the creation of CAPTIVE – a controller system they believe will do just that.



One of the most important elements in virtual and augmented reality systems is their controller. As an architect or a doctor, you may wish to accurately inspect complex 3D objects from every angle. If you play video games, you want your in-game character to perform fluid, realistic movements through virtual worlds. Images that lag noticeably and don't accurately reflect how you manipulate the controller can be frustrating and disorientating, making the systems difficult to use.

When deciding on a controller for a virtual/augmented reality headset, you may be faced with a trade-off. A 3D mouse like 3Dconnexion's SpaceNavigator may be your first choice – with its pressure sensitive handle, it is a fairly simple, cheap option, but is difficult to control. Like flying a kite, any slight of your hand could send your image careering off in the wrong direction, and you

could spend hours learning to control the device properly. On the other end, free-moving devices like the HTC Vive are much easier to learn to control, but their complex arrangements of sensors and base systems make them much more expensive.

Recent developments have brought about desirable alternatives – systems such as MagicMouse and ARToolKit use a single camera to track 2D marks on a physical controller. These systems are simple, cheap and easy to control, but their tracking algorithms have a long way to go. Transitions between frames of moving augmented reality images can be clunky and inaccurate, making the experience disorientating to users. Changes in lighting and objects in the background can also make it hard for the system to detect where the controller even is, which could render the system unusable.


Professor Christopher Healey and Zeyuan Chen at North Carolina State's Department of Computer Science have created a system they believe will be cheap, easy to control, and functional in many lighting environments, all while operating at 63 frames per second (<https://youtu.be/gRN5bYtYe3M>). 'Our system consists of a 3D-printed wireframe cube, a regular RGB camera and our program,' explains Chen.

Tracking Movements

When you move a computer mouse, two primary motions, or 'degrees of freedom' affect the position of the cursor on the screen: forwards/backwards and left/right. Professor Healey and Chen's controller has a much more varied range of motion. As well as those two initial degrees of freedom, their software detects when the user is moving the cube up or down, or rotating it about all three axes (motions called yaw, pitch and roll). This brings the number of degrees of freedom up to six – reflecting any possible way you could move a rigid cube through space.

The first step that CAPTIVE's software takes is to map out a rough region containing the cube within each frame the camera captures. Using previous frames, and the system's knowledge of the shape of the cube, it can change the shape and location of the region based on previous trends in the cube's motion. It uses

‘The primary advantage of CAPTIVE is that it is efficient. There are a number of tools on the market that can be used to manipulate 3-D virtual objects, but CAPTIVE allows users to perform these tasks much more quickly.’



these trends to predict where the region should be in the next frame, ensuring smooth transitions between frames. At 63 frames per second, ‘there’s basically no latency, no detectable lag time between what the user is doing and what they see on screen,’ says Chen.

The next task is to detect the exact pose of the cube within the region. CAPTIVE does this using the eight uniquely-coloured balls that comprise the eight corners of the cube. Within the rough region in each frame, the software locates the positions of each corner of the cube. It then compares these positions with a virtual cube stored in the software.

Like the physical cube, this virtual image has a different colour assigned to each of its corners, meaning that you can always tell what angle you are viewing it from, based on the orientation of its corners. The size of the virtual cube is defined by a 3D coordinate system, with its origin located at the centre, and each corner positioned at a unique coordinate.

By aligning the images of the real cube and the virtual cube, CAPTIVE calculates the exact pose of the real cube in physical space and establishes a new coordinate system in each frame. Each new coordinate system acts as a canvas on which the frames of virtual

reality objects are projected. Because new coordinates are established in each frame, the user will see 3D objects with movements that reflect how they move the cube in real life. Like the cube, these objects will also possess six degrees of freedom, meaning users can manipulate them in any way they wish.

The researchers also designed innovated methods to use CAPTIVE as a joystick-like controller by mapping the calculated 3D information to velocity of virtual objects. Because CAPTIVE’s cube is free-moving, it doesn’t have a ‘default’ position like most controllers. This means that, in the ‘joystick’ mode, when the user moves the cube, the 3D object will carry on moving when the cube stops, based on the trends in motion that the tracking software picks up. To stop the object from moving, bringing it back to a ‘neutral’ position, the research team has implemented a ‘pedal’, which the user can press any time they want.

‘The primary advantage of CAPTIVE is that it is efficient,’ Chen explains. ‘There are a number of tools on the market that can be used to manipulate 3-D virtual objects, but CAPTIVE allows users to perform these tasks much more quickly.’



Imperfect Lighting and Noisy Environments

Professor Healey and Chen's tracking system is simple but effective, yet the combined task of CAPTIVE's camera and its software isn't so easy. In practical use, different lighting environments can significantly change the colours of the eight corners as they appear to the camera, and other objects in the background can clutter the image, making it harder for the software to pick out the cube's position. These issues have caused problems for previous controllers, but the researchers believe they have found solutions in their software.

To combat the challenges posed by unpredictable lighting, the team has made use of Hue, Saturation and Value (HSV) colour space. HSV space consists of a colour wheel in which colours fade into each other, with darker shades around the edge fading to white towards the centre. In different lighting environments, the colours of the cube's corners will appear to vary, corresponding to different positions in HSV space. So rather than picking up single colours, the software is programmed to detect a corner if its colour appears in a particular small region of HSV space. Thanks to this approach, users don't need to worry about the CAPTIVE becoming unresponsive when lighting conditions change.

Virtual reality systems can also become confused when their cameras pick up objects other than the controller, but the software can deal with this too. If it is initially confused by background clutter, it can recognise when pixels in the frames fall inside circle-like regions, as the balls comprising the cube's corners, appear to the camera. If any pixel doesn't meet these criteria, it is blocked out, making only the cube's corners visible to the software. So instead of needing to set up the system in a sterilised lab environment, users can use CAPTIVE in more practical settings.

Practical Uses

CAPTIVE could offer enormous benefits to fields including medicine, architecture, archaeology and geology, to name just a few. By simply integrating CAPTIVE into other stereo systems, users can project their models and 3D objects onto the cube as augmented reality images. They can literally hold projections in their hand, and easily and

accurately inspect them from any angle. Furthermore, the system could be set up in locations ranging from operating theatres to remote wildernesses, making it much more practical than other devices. With further improvements, users could even use tools like a stylus to directly interact with and manipulate the images.

Video game creators have also begun to invest heavily in the technology, but various controller-related difficulties have prevented any titles from becoming hugely popular so far. Motions including rotation have been difficult to translate from the user's hand to the virtual world, often making movements slow and disorientating. CAPTIVE's six degrees of freedom could solve this issue, enabling players to move and rotate objects fluidly and realistically. The use of the technology as a joystick controller was the subject of an experiment the researchers carried out themselves.

Experimenting with Real Users

Chen, Professor Healey and their colleagues tested CAPTIVE's usability with the help of 15 of North Carolina State's computer science graduate students. Projecting a virtual reality image of a house inside the cube, they asked the participants to rotate the house to a target position as quickly and accurately as possible. They essentially used the cube as an input device. But unlike similar previous tests, the participants' view was affected by their head movements, which added an extra layer of difficulty to the task. The researchers were confident that CAPTIVE's ease of control, and rapid frame rate, would offset this difficulty.

On average, the students completed the task in just 2.79 seconds and missed out on a perfect rotation by an angle of 6.2 degrees. The results were a drastic improvement on experiments with similar controllers, with participants completing rotations almost twice as fast. They almost matched the performance of a much more complex and expensive system, known as a surround screen virtual environment, though they were not quite as accurate.

Despite the experiment's success, not everything went perfectly. Strangely, three students had difficulty in mentally connecting the 3D image they saw with the cube they were holding. The researchers decided that more research would be needed to understand why the technology could be more cognitively difficult to grasp for certain people.

'We are improving the accuracy and robustness of this tool, especially in outdoor environments, using deep learning methods,' Chen claims. 'Additionally, we are combining this tool with mobile devices (e.g. iPhones) and professional virtual and augmented reality devices (e.g. the Microsoft HoloLens).' The team would also like to develop a cube with a more translucent wire frame. Augmented reality images can be disorientating to users when the cube's frame visibly passes in around their image. A less intrusive cube design is desirable but would require more work to implement. In the future, they would like CAPTIVE to project two separate frames for user's eyes to improve their depth perception, which would make 3D images more realistic. The technology still has a way to go, but for any virtual reality enthusiast, CAPTIVE is a truly exciting development.



Meet the researchers

Professor Christopher G. Healey
Department of Computer Science
North Carolina State University
Raleigh, NC
USA

Zeyuan Chen
Department of Computer Science
North Carolina State University
Raleigh, NC
USA

Professor Christopher G. Healey is a Professor in both the Department of Computer Science and the Institute for Advanced Analytics at North Carolina State University. He completed his PhD in computer graphics at the University of British Columbia in Vancouver and has since received \$5.5 million in funding from both the government and industrial partners, for his world-leading research in computer science. Professor Healey's accomplishments have earned him the National Science Foundation's CAREER Early Faculty Development Award, and the Outstanding Instructor Award from North Carolina State University.

CONTACT

E: healey@ncsu.edu
W: www.csc.ncsu.edu/faculty/healey

Zeyuan Chen is currently completing his PhD in the Department of Computer Science at North Carolina State University, where his areas of research include visualisation, computer vision, and human-computer interaction. Prior to starting his PhD research, Chen studied Electronics and Information Engineering and Information Engineering at Huazhong University of Science and Technology in China, and then completed his Master's degree in Electrical Engineering at Cornell University. He has a keen interest in designing computer graphics, and a selection of his creations can be found on his website.

CONTACT

E: zchen23@ncsu.edu
W: www.zeyuanchen.com

FUNDING

US National Science Foundation

REFERENCES

Z Chen, CG Healey and RS Amant, Performance Characteristics of a Camera-Based Tangible Input Device for Manipulation of 3D Information, in Proceedings of the 43rd Graphics Interface Conference, Canadian Human-Computer Communications Society, 2017, pp. 74-81.

**NC STATE
UNIVERSITY**