

Lecture 5: Advanced user interface techniques

This lecture summaries recent and current HCI research into advanced interaction technologies, using a variety of projects to review the principles introduced elsewhere in the course.

Virtual reality (VR)

The term *virtual reality* originally applied only to *full immersion VR*, in which simulated world is projected onto all walls of a room (*CAVE* – a recursive acronym for CAVE Automatic Virtual Environment), or via a *head-mounted display* (HMD) which uses motion-tracking to change the view as you turn your head. Interaction was always a challenge – *data gloves* could supposedly be used to pick up and interact with objects in the virtual scene. ‘Natural’ navigation in the real world is achieved by walking, but CAVEs were never large enough to walk far, and HMDs with motion tracking were normally tethered by cables. In practice, the illusion was always fairly limited, unlike the Matrix-style science fiction ideal that motivated it. Marketing creep has meant that any interactive 3D environment might get called VR, even if presented on a standard monitor, and controlled by a mouse. As games players know very well, control of view and camera angle, unless constrained by a script, can make arbitrary action in 3D scenes complex.

Augmented reality

Augmented reality (AR) systems overlay digital information onto the real world, either using partially-transparent head mounted displays as with Google project Glass, or by taking a video feed of an actual scene, and compositing it with computer generated elements. A key technical problem is registration – relatively recently, this had to be done by integrating GPS, compass orientation, accelerometer for gravity orientation, and often gyroscopes, into the HMD. Now that all these peripherals are available on high end mobile phones, Mobile AR is becoming a major marketing buzzword, possibly with the same loss of actual functionality that occurred when VR shifted from research ambition to marketing buzzword.

Tangible user interfaces

Tangible user interfaces (TUIs) use physical objects to control the computer, most often a collection of objects arranged on a tabletop to act as ‘physical icons’. An immediate problem is that physical objects don’t change their visible state very easily. You can include motors and displays in each object (expensive), or project overlaid AR information onto them, or just use them as multiple specialized mice/pucks that control elements of the display on a separate screen. In this case, it is necessary to track their positions, perhaps by using a large tablet device. If they are just being used as tokens to select a particular function or piece of data, an embedded RFID chip can be used to sense when they are placed within a certain distance of a reader.

Machine vision

Machine vision is a key technology for both AR and TUIs, as a way of recognizing real world objects such as buildings (in the case of outdoor AR) or objects on a desk (used for TUIs). Many current AR prototypes recognize distinctive objects from a large number of low-level visual features. Key problems are to maintain a sufficiently large database of object features, track them fast enough to give user feedback that responds to camera, gesture or object motion in real time, and do both of these in varying lighting conditions. An alternative is ***fiducial markers*** – simple visual markers such as barcodes that can be used to more reliably identify and track objects from camera input. They are more robust to changes in camera angle and lighting than object recognition algorithms.

Paper interfaces

Inspired by the research conducted by Abigail Sellen and Richard Harper, whose book ‘The Myth of the Paperless Office’ analyses the ways in which the properties of paper are preferable to computers for many kinds of activity. The book remains a useful resource for designers of mobile devices substituting for paper (phones and tablets), but has also inspired research in which paper is integrated with digital systems, for example with fiducial markers on the page that can be traced by cameras (the ***Anoto digital pen*** can

perhaps be considered an extreme example of gesture recognition implemented with fiducial markers).

Eye tracking and gaze control

Originally developed for psychological research into visual attention processes, *eye trackers* are now used fairly routinely in HCI research to study what position on the screen users are looking at. A high resolution close-up camera is used to capture video of one of the user's eyes, and the precise position where they are looking is deduced from the position of the pupil, often combined with reflections from a pair of small infrared (LED) spotlights. One company sells a device with the camera and spotlights integrated into the surround of a monitor, to be unobtrusive. However, almost all systems like this require the user to sit fairly still, and to undergo a calibration procedure in which they look at points on the screen in sequence. Performance can be poor when there is strong ambient lighting, when the user wears spectacles, has watery eyes or shiny skin. Often practice is required to get good results.

In practice, natural eye movement of *fixations* and *saccades* can confuse the eye-tracker inference algorithms, it is hard work to keep your eyes fixated on control locations for substantial periods of time, and the natural temptation to glance elsewhere (check work in progress, look at the time, look down at your hands etc) or to blink excessively must be constantly fought.

Surface and tabletop interaction

Surface/tabletop interaction uses large display areas, usually projected, on a flat surface such as a wall or table. User interaction takes place by touching, gesturing, or pointing at the display. Many of these systems use camera input, with more accurate recognition of the users hands possible by using infrared, rather than visible light. A low-powered infrared spotlight is often used to illuminate the scene, rather than relying on body heat. A popular technique at present is frustrated total internal reflection (*FTIR*), where infrared light is shone inside a flat transparent medium such as a glass panel, and anything touching the surface causes infrared to be scattered. This technique can be used

to recognize fingertip touches, or gestures involving more skin contact, such as multiple fingers or even a flat palm. A motivating scenario for many of these systems has been the gesture-controlled projection interface in the movie *Minority Report*.

Embodied interaction

A user sitting at a desk, in front of a screen, with a keyboard and mouse on the surface, was the default assumption in most classical user interface designs. Most of the new technologies described above are used in other positions, making it necessary to take account of how users stand or move around. Machine vision, tracking of infrared markers, use of accelerometer data, ‘smart’ fabrics and clothing, and many other techniques can be used to analyse and track body positions.