



Tracking and Modeling the Human Hand

Jamie Shotton
Microsoft Research Cambridge, UK

Abstract

The recent availability of consumer-grade virtual and augmented reality devices is fueling a renewed interest in vision systems for tracking the detailed articulated motion of the human hand. The dream is that we should be able to interact with virtual objects in much the same way as we interact with real objects: reach out in VR and see your avatar hand; grasp or poke at virtual objects and they react as if they were real. Unfortunately, our hands pose huge problems for vision systems: they rotate freely in 3D, they often self-occlude, and the fingers look very similar locally. Our hand movements can also be incredibly precise, and ultimately sub-millimeter precision will be required. This talk will highlight some of history and algorithms behind tracking and modelling the human hand. We'll cover recent breakthroughs in the vision, learning, and graphics communities that are finally bringing real-time tracking of the human hand close to consumer readiness. We'll cover both pose and shape estimation, and look at a range of techniques including generative model-fitting approaches and discriminative learning approaches. We'll also touch on some of the incredible possibilities that this research will enable in next-generation user interface design.

Keywords

Tracking, 3D human hand modeling, pose and shape estimation, generative model fitting, discriminative learning, VR, user interface design