9807-28
Variable Viewpoint Reality

Progress Report
January 1 — June 30, 1999

Paul Viola and Eric Grimson

Project Overview

In the foreseeable future, sporting events will be recorded in super high fidelity from hundreds or even thousands of cameras. Currently the nature of television broadcasting demands that only a single viewpoint be shown, at any particular time. This viewpoint is necessarily a compromise and is typically designed to displease the fewest number of viewers.

In this project we are creating a new viewing paradigm which will take advantage of recent and emerging methods in computer vision, virtual reality and computer graphics technology, together with the computational capabilities likely to be available on next generation machines and networks. This new paradigm will allow each viewer the ability to view the field from any arbitrary viewpoint — from the point of view of the ball headed toward the soccer goal; or from that of the goalie defending the goal; as the quarterback dropping back to pass; or as a hitter waiting for a pitch. In this way, the viewer can observe exactly those portions of the game which most interest him, and from the viewpoint that most interests him (e.g. some fans may want to have the best view of Michael Jordan as he sails toward the basket; others may want to see the world from his point of view).

Progress to Date

We have made rapid progress on a number of problems related to the goals of the Variable Viewpoint Reality project:

• We have developed a number of basic algorithms for 3D reconstruction. One approach is designed to work in real time on many cameras. Another is a bit slower, but is designed to yield higher quality results. A third attempts to find the arm, leg and body positions of a human being from one or multiple camera views.

Each of these algorithms is in its very earliest stages. We are exploring these algorithms now: constructing implementations, testing assumptions, etc.

• We have designed and setup a multiple camera systems for acquiring data in real-time. This system was designed to be flexible and to work indoors. Right now we have 12 cameras working in synchrony. We would like to setup more.

• We have acquired a great deal of multi-camera data. This is allowing us to test our algorithms and to develop new ideas.

• In collaboration with students working on another project we have been observing outdoor activities. This system provides coarse tracking information of multiple people and cars. The system can also recognize simple activities.

Timeline for Future Work

August 1999:

Demonstrate real-time 3D reconstruction and visualization from 12 cameras distributed around an indoor 3 meter cube. This space will be large enough for a single person to move and take actions. (We believe that this system is the first step along the way toward a more general purpose, larger scale system.)

December 1999:

We will extend the August system to include additional functionality:

• Tracking of people using articulated body models.

• Improved texture mapping of body models.

• First results on action interpretation.