FUTURE RESEARCH DIRECTIONS

My background in relation to my future research

In the broadest terms, my research interests are in artificial intelligence and computer graphics. These interests did not just arise in graduate school or even as an undergraduate. I began thinking about and experimenting with artificial intelligence in fourth grade and I went to graduate school at MIT specifically to pursue a Ph.D. in artificial intelligence, as I had planned since high school. I also did my Ph.D minor in psychology out of a general interest in how the mind works. I have a life-long interest in art, animation, and computer graphics stemming from the fact that my parents are artists and my father was a professor at the Art Institute of Chicago and with the New Bauhaus at the Illinois Institute of Technology. I wrote my first 3D graphics renderer in eighth grade after attending a SIGGRAPH workshop and I later studied computer graphics under Al Barr at Caltech. After leaving MIT, I worked with a start-up company producing a desk-top video editing system, where I cloned all the image-processing capabilities of Adobe Premiere (the leading video editing system of the time) and Faux Matisse (a sophisticated computer painting program).

My Ph.D thesis research has been in machine vision because this area gives me a chance to do research in both artificial intelligence and computer graphics. There are some additional fields of personal interest to me, namely art, paleontology, astronomy, history, and movie making, that I will actively try to incorporate into my future research.

I have always been able to pursue research in a highly independent manner, both at Wisconsin and in earlier research at MIT, Argonne National Laboratory, Star Media Systems, and during a summer internship at Caltech. I also have significant experience in creating and writing grant proposals, having worked closely with my advisor on two NSF grant proposals.

A few specific research plans

Continued research into automatic scene reconstruction. In addition to further work on screw-transform manifolds and calibration without minimization, I am interested in the problem of automatically recovering complete, detailed models of individual objects, including surface normals for fine-scale features that cannot be modeled otherwise. I am particularly interested in developing simple algorithms for this task.

"Blob" matching, specifically for measuring brain growth. The challenge is to create an algorithm that can find a logical correspondence (e.g., similar to what a person might produce) between two somewhat different deformable, organic shapes. I have already designed and implemented a method for finding correspondences between two 3D models of a maturing brain, and the results are very promising. This is a rich area of research with many important problems to be solved, both in medical imaging and more general machine vision areas; it should keep me busy for years.

Simple, computationally-efficient algorithms for identifying general classes of objects. If a smooth, white blanket or sheet is dropped haphazardly onto the floor, it will land in a pile containing arbitrary folds and curves. If photographed, the crumpled white blanket appears as a grayscale image showing smooth variations from black to white. Each time the blanket is picked up and dropped, a completely new set of folds and curves are created, a set that has never been seen before, and yet a person has no trouble identifying the object for what it is. I am interested in finding algorithms for differentiating between broad classes of objects (e.g., soft object with folds), particularly classes like the crumpled blanket that show infinite variation on a simple theme. I am especially interested in creating algorithms that are basic enough to be embedded into mobile devices with low power consumption and limited space constraints.

Automatic color adjustment for inserting synthetic objects into real images. Many movies have special effects in which synthetic objects are placed into real scenes or real objects that have been photographed under one lighting context are inserted into scenes with different lighting contexts. Ideally, the lighting on the inserted object should match the final scene lighting, but this is impossible if the final lighting is unknown or the object's geometry and reflectance properties are unknown. Furthermore, absolute physical correctness might be unnecessary in some circumstances; sometimes all that is needed is to match the human viewer's expectation of what the correct coloring should be. I plan to acquire statistics about how an object's color changes in relation to its context and about what kinds of recoloring human viewers find realistic, and then create an algorithm that can automatically recolor an object to match these statistics. Such research might be useful in both computer graphics and cognitive science.