Philosophy

Computer graphics is an important component of modern scientific analysis, storytelling, and communication of ideas. As the role of computers as mediator between people and the world grows, computer generated visual stimulus will become increasingly critical in our daily lives. My long term research agenda is to ensure that this computed imagery allows for the same depth of understanding as does direct perception of the natural world.

Computer graphics make use of a variety of underlying models to produce imagery. It is precisely these models which lend computed images their power to aid analysis and understanding. In contrast to the static collection of pixels collected by a camera, computed images can be manipulated and transformed, allowing an engineer to easily see whether a design will work when the shape of a part is changed, and a movie director to modify the motion of a dance to best convey a story.

Modern rendering systems have progressed to a level of richness such that it is often impossible to distinguish complex computed images from real images collected on a camera. Unfortunately, methods for creating the complex graphical models used for rendering have not kept pace with our ability to render imagery from these models. As the level of realism increases, so does the need to specify greater and greater detail in the underlying models. Movie studios achieve the level of realism that they do, primarily by hiring a larger and larger number of skilled artist, and building tools to increase the efficiency of these artists. This endeavor is expensive and time consuming.

An alternate approach to creating graphical models is to acquire models directly from the real world, leveraging the phenomenal complexity of nature. While manual creation of graphical models is appropriate in some circumstances, it is impossible in others. Scientific visualizations often represent some aspect of the real world in a form that can be manipulated and analyzed. In these cases, the models must be acquired directly, rather than subjected to artistic interpretation. For example, an anthropologist studying the skull structure of chimpanzees almost certainly requires a graphical skull model based on real data.

Unfortunately, many complex graphical models are not yet possible to derive from the natural world. Examples include: the precise shape of wrinkles while an old woman smiles, the involuntary reactive motion of a medical patient in response to neurological brain probes, and the muscular deformation of a runners legs while sprinting down the track.

My primary research direction is to develop tools for easily creating complex and usable graphical models. This focus lies at the intersection of computer graphics and computer vision, and I have found myself an active participant in both communities. Methods for acquiring raw data from the environment are necessary. In addition, systems that transform this raw data into useful graphical models are required. These models enable scientist to analyze and experiment during the course of their visualizations, and artist to explore in the pursuit of their craft.

Contributions

I have developed a number of specific techniques allowing more accurate graphical models to be recovered from the world. This work includes methods both for acquiring raw data, and transforming the raw data into usable models.

Models of motion play an important role in any animated visualization. Commercial motion capture systems are in common use among both researchers analyzing human bio-mechanics, and movie studios producing their next film. My dissertation work focused on addressing the scaling difficulties inherent in traditional motion capture systems. Higher resolution and larger working volumes have demanded an ever increasing number of cameras to be deployed. Rather than distribute resolution uniformly throughout the working volume, I introduced a framework that allows for multiple scales of recovery. Pan-tilt cameras are used to follow the action, providing high resolution only in a dynamically variable subset of the entire working volume. For a particular resolution requirement, this results in an order of magnitude greater efficiency in terms of physical and computational resources.

Models of shape are at the core of most graphic visualizations. The digital Michelangelo project explicitly set out to explore complexity. While participating on this team, we obtained shape measurements of ten statues, each accurate to within 0.25 mm. The largest statue, Michelangelo's David, stands five meters high and resulted in a model with over a billion polygons. The polygonal shape models are being used by art historians to study Michelangelo's work and to plan the restoration of these priceless sculptures. In addition, the models produced are now among the most commonly used data sets by computer graphics researchers.

Among the lessons learned during the course of the digital Michelangelo project were the facts that complex mechanical equipment does not stay properly calibrated for long periods, and no amount of scanning will produce a complete hole free model. Following this project I developed a new design for laser scanners that does not include calibrated moving components. This design provides an improved method for acquiring raw data that is both cheaper and more reliable. In addition I developed a volumetric method for filling the inevitable holes that remain when acquiring geometric models. This hole filling method is automated, efficient, and guaranteed to produce a physically plausible solution.

Motion and shape are deeply related. Unfortunately, technologies for acquiring raw measurements in these domains have traditionally been separated. A key challenge in bringing useful graphical visualizations to many scientific investigations is developing a method for accurately recovering the shape of moving objects. To address this need, I introduced the space-time stereo framework, unifying existing triangulation based methods for shape recovery. This framework has allowed the shape of moving objects to be recovered at higher quality than was previously possible. An alternate technology for recovering the shape of moving objects is shuttered light pulse imaging. Unfortunately, several deficiencies currently limit its applicability to graphical modeling. I have addressed several of these, introducing methods for simplifying calibration, increasing precision, and enhancing robustness to ambient lighting.

Scene lighting is often modeled as an environment map. One technology for creating environment maps, stitching several photos into a single composite, has moved from the research lab into the consumer marketplace. Many digital cameras now include panoramic mosaicing features. I developed one of the first algorithms for globally aligning large collections of images. Additionally, I introduced a procedure for creating a consistent composite image, even when some objects are moving. The methods from this work were commercialized as part of the Sony PictureGear product.

User interfaces are often a necessary tool through which artists and scientist interact with complex graphical models. I developed a pen based interface system for wall sized displays. The system allows multiple users to work and interact simultaneously on the display. In another project, I developed an interface for sketch based animation that quickly turns an animator's sketched stick figures into rough 3D animation. This interface provides full artistic control, but is substantially easier to use than conventional inverse kinematics modeling tools.

Future directions

My driving belief is that technology as a whole should be used in the service of humanity. Applied properly it will enhance health, education, safety, justice, and economic equality. Many of these emerging social technologies will rely heavily on visual communication and computed imagery. This imagery must be sufficiently complex and rich that it allows for deep understanding, analysis, and communication.

The work described above is an initial step towards acquiring complex models of the real world. In the short term I am intent on improving the quality and speed with which shape and motion can be recovered. However, since the value of complex graphical visualizations stands with its applications, rather than with the technology itself, I have established collaborations with a variety of researchers outside of computer graphics and vision. For example, I am working with robotics researchers to enable real-time shape sensing on mobile platforms. Their interest lies both with automated exploration and map building, and with imitation learning of humanlike motion for humanoid robots. Another collaboration, with machine learning researchers, seeks to automatically determine the kinematic structure of objects directly from range data. I am collaborating with biomechanics researchers who seek to understand the motion of human joints under conditions in which traditional marker based motion capture is unsuitable. Finally I have been approached regarding an application of this technology to recovering the deformation of human tissue during medical procedures.

Each of the above applications is a specific instance of recovering raw data from the world and transforming it into a graphical model which is useful for accomplishing a particular task. I anticipate that graphical models will grow to incorporate an ever wider class of natural phenomena. My research program will continue designing methods to acquire, manipulate, merge, transform, parameterize, and decompose these models, ensuring that they are usable by the artist, scientist and social causes that require them.

In the long term I would like to extend the impact that academic researchers have, by establishing a center for social technology, bringing together researchers from a broad range of fields. This center would explicitly identify future applications of technology in the service of humanity. Further it would encourage the research directions and collaborations necessary to enable those applications. Remote surgery on a patient halfway around the world, digital archival and preservation of priceless art, virtual communities for the isolated, and education in third world villages are all examples of realizable goals.