

Research Statement: Srinivasa G. Narasimhan

Underlying all of computer vision research is a model of how light interacts with a scene and then reaches a camera to form images. Light propagates through a scene in complex ways — inter-reflections between scene points, diffusion beneath the surface of translucent materials like skin and marble, and scattering through media like the atmosphere and murky water. Despite this complexity, the vision community has defined the brightness of a pixel in the image as solely due to the light reflected from a single point in the world. This characterization is correct only for convex opaque objects in vacuum! Modeling this wide variety of optical phenomena is crucial for effective scene understanding in all real-world environments. In my research, I am solving this challenging problem by carefully collecting data in the laboratory and in the field, constructing models (both analytic and data-driven) and developing algorithms that effectively use these new models of light transport to solve problems in computer vision, graphics, medical imaging and robotics.




Figure 1: Results: Removing haze & murky water; rendering liquids, 2D drop display, reconstructing wreaths & bones.

Vision in bad weather and murky water. Studies have shown that we do not judge distances accurately in bad weather, which often leads to accidents. In these conditions, autonomous systems that warn the drivers of impending collisions and lane departures could save lives. I was the first to show how to exploit light scattering models to enhance visibility in foggy and hazy images. A surprising by-product was that I was also able to estimate relative scene distances from fog or haze. In addition, I have developed real-time simulators of weather effects (distributed with NVIDIA's graphics cards), that can be easily included in video games and movies. I am currently developing methods to handle more dynamic conditions like rain, snow, smoke and dust. I am also applying my ideas to measure and improve visibility in murky waters.

Immersive volumetric displays. Current display technologies all require a 2D display surface. Even prototype holographic devices are small and require complicated optics. I plan to create novel 3D displays by using ordinary water drops and mist as display media. Water drops (and mist) can be generated and controlled fast and with my scattering models it should be possible to use projectors to provide light that will be scattered through the medium to create the appropriate 3D image for the viewer. I have tested the feasibility of this idea by building a 2D system that accurately controls both an array of water drops and light rays that hit them.

Acquiring 3D models of natural objects. I am interested in 3D modeling of humans and natural objects like plants and trees. Today's laser scanners and stereo algorithms have demonstrated impressive results on stationary (or slowly moving) objects with simple shapes and diffuse materials. By analyzing light transport and exploiting the duality between illumination and imaging, I demonstrated the acquisition of complex feathery shapes and translucent objects, entirely automatically for the first time. I also developed a method that uses micro-mirror arrays for kHz rates motion capture that can lead to a better understanding of how athletes and animals move.

Endoscopic image analysis. Automatic analysis of endoscopic images could significantly reduce human error during patient examinations. A detailed understanding of light transport through bones, tissues and blood is necessary here as the distinction between a nascent tumor and healthy tissue can be quite subtle. My goal is to combine critical medical applications with a formal understanding of light transport. I have demonstrated reconstruction of bones from endoscopic images, which serves as a positioning tool during minimally invasive orthopedic surgery. I am now analyzing the translucency of bronchial tissues for early detection of lung cancer.

In summary, defining an image purely as “a geometric mapping from 3D to 2D” is far too simplistic and discards information that is essential in most real world applications. Light transport must not be viewed as noise, but rather as an additional form of information about the environment. This insight has led to new research directions that will impact not only vision but also such domains as graphics, robotics, oceanography and medicine. I have had considerable success in each of these areas already (receiving three Best Paper awards) and along with my four graduate students, I am excited about the challenges and possibilities that lie ahead.