Multi-modal Modeling with the UBC Active Measurement Facility

Dinesh K. Pai
Department of Computer Science
University of British Columbia
Vancouver, Canada
pai@cs.ubc.ca

Abstract

Multi-modal models of physical objects are important in many application areas, but constructing such models has been difficult. ACME, the UBC Active Measurement facility, is a robotic facility we have developed for constructing such models. We describe the construction of multi-modal models, including models of shape, contact sounds, and surface roughness.

1 Introduction

Constructing realistic models of the objects in the physical world is important in many application areas, including computer graphics, computer games, animation, e-commerce, telepresence and training simulators. Increasingly, these applications require not simply models of a specific aspect of the object but integrated, multi-modal models.

For instance, consider the simulation of an object being dropped on a table, bouncing, and rolling to a stop. Such a simulation would obviously require a geometric model of the boundary and the inertial properties of the object. What may be less obvious, until one thinks about it, is that in reality such a contact interaction is inevitably accompanied by contact sounds, first of the object's impact with the table, and then the sounds of continuous rolling contact. In fact, such sounds are so important that they are currently added as a post-production step by "Foley artists" in the movie industry (and even in videos produced at computer graphics conferences, intended to depict a visual simulation!). To support such a simulation, we will need a model for sound synthesis registered with the geometric model, so that when an impact occurs at a particular location we know what kind of sound would be produced. Now suppose that instead of dropping the object on the table, we want to interactively grasp a virtual object and slide it on the table. In this case we will also need to have a model of surface friction and roughness, registered with the geometric and sound models, to simulate the contact forces to be "displayed" using a force-feedback haptic device like the PHANToM^{TM}.

More generally, we believe multi-modal models are important for all human-information interfaces, as they can exploit a broad range of sensory modalities available to a human user, to convey information in a natural way. Sutherland's vision of the "The Ultimate Display" [15] is an early example of this view.

To make accurate models of existing objects, we require that the models are "reality-based" (i.e., based on measurements of the real objects, rather than constructed a priori).

Traditionally, however, reality-based modeling has focused on specific properties of objects, for instance, shape modeling using computer vision techniques. Using such special purpose measurement techniques it is extremely difficult to build comprehensive multi-modal simulations such the one depicted above. These systems are thus "measurement-oriented" in the sense that they are oriented towards the needs of the measurement device.

In this paper we describe the opposite approach, which may be called "object-oriented" in the sense that it is oriented towards building multi-modal models of an object. This approach has been the motivation for the design of ACME, the UBC Active Measurement Facility [8, 9]. In Section 2 we give a brief overview of ACME, our Active Measurement Facility, its measurement capabilities and software architecture. In Section 3 we show how ACME is used to construct multi-modal models. Section 4 contains concluding remarks.

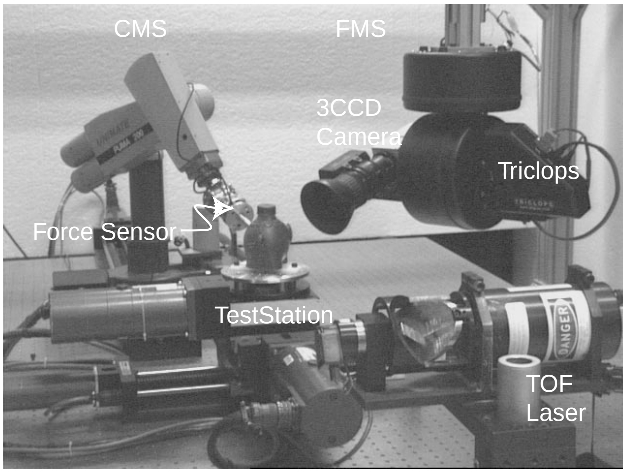

Figure 1: ACME Facility Overview

2 The Active Measurement Facility

The UBC Active Measurement Facility (ACME) [8] is an integrated robotic measurement facility for acquiring multi-modal reality-based models. ACME is designed to acquire a large number of registered measurements of a given object including shape, sound, and contact texture. See Figure 1.

At the center of the facility is the TestStation, a three-dof robotic subsystem that can provide precise motions of a test object in the horizontal plane. Range measurements for building a geometric model can be acquired with a laser range finder (Acuity AccuRange 3000 LIR) or with a color Triclops trinocular vision system from Point Grey Research [10]. The Triclops, a 3-CCD color video camera (Sony DXC 950), microphones, and other sensors are mounted at the end of a 5 degrees of freedom robotic subsystem called the Field Measurement System (FMS). The FMS provides an active and flexible method to place sensors around the test object and sample the light and sound field around the object. The Contact Measurement System (CMS) is used to make measurements of surface texture and also to excite contact sounds from the test object. It consists of a 6 DOF PUMA 260 robot arm (placed on a long motion stage to increase its range of motion), equipped with a 6 axes force/torque sensor (ATI Mini 40) and a variety of specialized end-effectors described below.

Since ACME is a robotic system with fifteen degrees of freedom and a large number of sensors, a great deal of attention was paid to the software design and system architecture, to make this complex system easy to use. First, almost everything in ACME, ranging from camera settings to sound sample acquisition, is under software control. This has the advantage of making experiments repeatable even after a period of time – the precise measurement conditions could be recreated merely by re-running a program. Second, ACME was designed to be a shared telerobotic facility, accessible from anywhere on the Internet. Third, we designed for distributed measurement and control, to provide scalability and real-time performance as new sensors and actuators are added to ACME.

Specifically, high-level Java objects encapsulate ACME `Devices` (`Sensors` and `Actuators`) and `Motion Plans`. These objects provide a layer of abstraction that hides the distributed implementation of ACME devices. Asynchronous interfaces are provided between the high-level Java objects and real-time sensing and control. For instance, `Sensors` use an event-based data acquisition model. Users specify events for triggering data collection activities, which are performed outside the JVM for real-time response.

An extensive teleprogramming system was developed to make it possible to control ACME from the Internet, despite variable latency and bandwidth. The key to performing real-time data acquisition and closed loop robot control from a remote location is a piece of mobile software called the ACME `Experiment`. The ACME server performs real-time tasks such as control and data acquisition. The ACME client, at a remote location, can create an `Experiment` class in Java. `Experiment`s can thus make on-line decisions based on acquired data, such as what measurements to perform next and when to stop. Compiled `Experiment`s can be be dynamically loaded into a running server or a simulation. Collision detection and sensor simulation are available for `Experiment` verification. A verified `Experiment` can then be loaded, transparently and without any modification or recompilation, into a running server's Java virtual machine and executed on the real ACME hardware to acquire real data.

3 Multi-modal modeling

To construct a multi-modal model of a test object, it is fixed on the TestStation. The fixturing required to hold the test object must not interfere with the measurement process. This remains somewhat of an art and must be carefully planned. Obviously the fixturing should not occlude important parts of the test object and allow the surface to be reached by the end-effector of the CMS. A more subtle issue is that clamping the test object changes the boundary conditions for the vibration the object and therefore can change

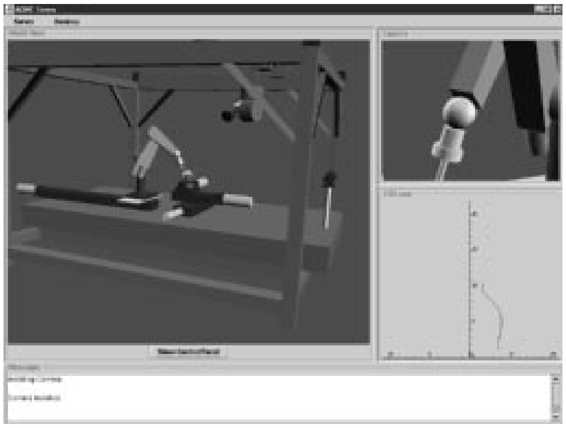

Figure 2: ACME Simulation. The panels on the right show simulated sensor data from the camera and laser scanner

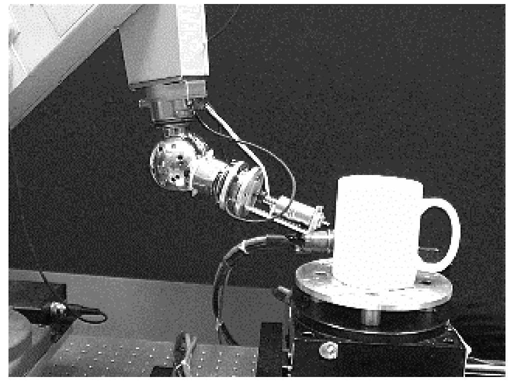

Figure 3: Sound Response Modeling

the sound.

Our approach is to register all measurements to a parametrized model of the surface of the object. Therefore, the first task is to acquire a geometric model using range measurements provided by ACME. Since most work on modeling to date has focused on this area, we have currently chosen to use existing solutions from the literature for combining range data into a single mesh (e.g., [1]; see also [11]) and to convert the data into a Loop subdivision surface (e.g., [4]). Range data can be acquired either using the Acuity Laser or the Triclops system mentioned above. ACME provides the ability to actively scan the surface of the object, primarily using the TestStation to move the test object in front of the sensors.

One of the most novel aspects of modeling is the acquisition of contact sound models [12]. A sound model consists of a model structure and model parameters. Several sound model structures have been proposed which are suitable for real time sound synthesis. We use a physically-based modal synthesis model, proposed by [16]. Real objects can have complex shapes and material composition, hidden internal structure, and subtle boundary conditions, and therefore the parameters of this sound model change over the surface. Our approach is to associate sound synthesis parameters at each vertex of a subdivision surface.

In our work [12, 13], the test object is struck with a special robot end effector called the *sound effector* (see Figure 3). A special effector is used rather than simply hitting the test object with the entire robot arm because we would like light, near-impulsive impacts on the test object and to avoid damaging the robot arm. The sound effector consists of a push type solenoid with a return spring, and microphone. This design delivers a low-inertia impact to objects 5 mm in front of the solenoid, at any orientation. The ACME robot moves the sound effector to different locations near the surface of the test object, locating the surface using force controlled guarded moves. At each location, the solenoid is fired and the impact sound recorded digitally. A large number of registered measurements can thus be systematically acquired. These measurements are then used to estimate realistic parameters of a sound model. Recently, Richmond [13] has developed an adaptive sampling strategy, which samples at coarse vertices of a subdivision mesh and then refines the sampling based on the estimated perceptual distance between sound models at adjacent mesh vertices.

We can also use ACME for contact texture measurement. Contact textures are important for haptic displays (e.g., [7, 3]). Perceived texture is a multidimensional quantity [6, 5]; surface roughness is one of the most important dimensions. Characterizations of surface roughness in engineering metrology do not appear to be directly useful for predicting the perception of roughness. Perceptual studies [6, 5] indicate that the spacing between geometric "elements" strongly determines the perception of roughness during exploration with a bare finger; surprisingly, surface friction is less important. Friction is also more difficult to characterize, since it is not a property of the surface alone, and depends on the contact pair.

We have developed a system to measure the distribution of geometric elements with an ordinary industrial robot (Puma 260) in ACME, equipped with a 6-axis force/torque sensor. A variety of rigid aluminum probes were designed for contact measurements. Each probe has a steel ball at the tip to provide a spherical contact surface with low friction. The tip diameters range from 2.37mm to 6.33mm. We developed calibration procedures to locate the center of the ball relative to the force sensor. We have developed a non-linear filtering method using wavelet thresholding which reduces the measurement noise while maintaining sharp discontinuities in the forces. From contact force data, grooves of width 0.5mm, etched in a steel test plate, were estimated with mean errors much smaller than 0.1mm.

The advantage of using such a setup with a commercial robot arm is versatility and low cost. The robot, mounted on an additional long motion stage and used in conjunction with the 3 DOF TestStation, can make contact over different parts of an intact test object. This is in contrast with devices like profilometers and stylus measuring instruments, which require a small surface sample to be removed from the object and placed in the instrument.

Finally, we should mention that procedures similar to our contact sound modeling can be used for measuring deformation due to contact. Ph.D. student J. Lang is currently developing ACME Experiments to estimate Green's functions for elastostatic deformation using the CMS and the Triclops stereo vision. Surface reflectance can also be estimated with ACME (for instance, using the technique of [14]; see also [2]).

4 Conclusions

We believe multi-modal models are important in many applications and provide better human-information interfaces. However, acquiring such models of existing objects, carefully registered with respect to each other, has been extremely difficult. ACME, the UBC Active Measurement facility makes it easier to acquire such models by providing one-stop shopping for measurements required for reality-based modeling. We described how ACME is used to acquire models which include the object shape, contact sound response, and surface roughness. At the conference we will show examples of simulations using multi-modal models acquired using ACME. In the near future, we plan to extend these multimodal models to include surface deformation and better reflectance models.

Acknowledgments

Several people have contributed significantly to the development of ACME at the University of British Columbia and to the multimodal modeling effort described here. They are, in alphabetical order: R. Barman, C. Chiu, J. Fong, A. Fournier, L. Ke, J. Lang, J. E. Lloyd, S. Kingdon, and A. K. Mackworth, J. L. Richmond and R. J. Woodham. Financial support was provided in part by grants from NSERC and IRIS NCE.

References

[1] B. Curless and M. Levoy, "A Volumetric Method for Building Complex Models from Range Images," ACM SIGGRAPH 96, pp. 303–312.

[2] K.J. Dana, B. van Ginneken, S.K. Nayar, J.J. Koenderink "Reflectance and Texture of Real World Surfaces," CVPR '97.

[3] D. F. Green, "Haptic Simulation of Naturally Occuring Textures and Soil Properties," M.S. Thesis, M. I. T., 1998.

[4] H. Hoppe, T. DeRose, T. Duchamp, M. Halstead, H. Jin, J. McDonald, J. Schweitzer, and W. Stuetzle, "Piecewise Smooth Surface Reconstruction." In Computer Graphics (SIGGRAPH Proceedings), pp. 295–302, 1994.

[5] Roberta L. Klatzky and Susan J. Lederman. Tactile perception with a rigid link from surface to skin. *Perception and Psychophysics*, 1999.

[6] J. Loomis and S. J. Lederman. Tactual perception. in K. Boff, L. Kaufman and J. Thomas (Eds.) "Handbook of Human Perception and Performance," Wiley, NY, 1986. pp. 1–41.

[7] A.M. Okamura, M.A. Costa, M.L. Turner, C. Richard, and M.R. Cutkosky, "Haptic Exploration of Surfaces," International Symposium on Experimental Robotics, Sydney, Australia, 1999.

[8] D. K. Pai, J. Lang, J. E. Lloyd, and R. J. Woodham. "ACME, A Telerobotic Active Measurement Facility." International Symposium on Experimental Robotics, Sydney, Australia, 1999.

[9] D. K. Pai, "Robotics in Reality-based Modeling," International Symposium on Robotics Research, Snowbird, UT, October 1999.

[10] Triclops On-line Manual, *http://www.ptgrey.com/*, Point Grey Research, Vancouver, Canada.

[11] M. Reed, P. Allen, and I. Stamos, "3-D Modeling from Range Imagery: An Incremental Method with a Planning Component," In Proceedings of the International Conference on Advances in 3-D

Digital Imaging and Modeling, Ontario, Canada, 1997.

[12] J. L. Richmond and D. K. Pai. Active measurement of contact sounds. In proceedings of the 2000 IEEE International Conference on Robotics and Automation, San Francisco, April 2000.

[13] J. L. Richmond Automatic measurement and modeling of contact sounds. M.Sc. Thesis, Dept. of Computer Science, UBC. Expected August 2000.

[14] Y. Sato, M. D.Wheeler, and K. Ikeuchi, "Object shape and reflectance modeling from observation" ACM SIGGRAPH 1997, Page 379-387.

[15] Ivan Sutherland, "The Ultimate Display," Proc. IFIP, 1965.

[16] K. van den Doel and D. K. Pai, "The Sounds of Physical Shapes," *Presence*, The MIT Press, 1998, pp. 382-395, (Vol. 7, No. 4).