

Reality-based Modeling with ACME: *

A Progress Report

Dinesh K. Pai, Jochen Lang, John E. Lloyd, and Joshua L. Richmond
Department of Computer Science
University of British Columbia
Vancouver, Canada
{pai|jlang|lloyd|jlrichmo}@cs.ubc.ca

Abstract: We describe the current state of ACME, the UBC Active Measurement facility, a telerobotic facility designed for building computational models of everyday physical objects. We show how ACME is being used to acquire models of contact texture (including friction and roughness), contact sounds, and contact deformation.

1. Introduction

The UBC Active Measurement Facility (ACME) is a robotic measurement facility designed for building reality-based models (i.e., models of real objects constructed from measurements). These reality-based models are intended for applications in virtual environments, computer animation, computer games, and e-commerce; the focus is on easily creating rich multi-modal models of objects suitable for human interaction, rather than on highly accurate models of specific attributes.

At ISER 99, we described the design of the system and initial results [9]. In this paper, we describe the current state of the ACME facility (§2), and how ACME has been used to build new types of models, including contact texture models (§3), contact sound models (§4), and contact deformation models (§5).

2. ACME Facility

ACME consists of a variety of sensors and actuators (Fig. 1), all of which can be controlled using small Java programs called **Experiments**. More details are available in [9]. Briefly, the subsystems include: a 3 DOF Test Station used for precise planar positioning of the test object; a Field Measurement System (FMS), consisting of a Triclops trinocular stereo vision system, a high resolution RGB camera, and a microphone, all mounted on a 5 DOF positioning gantry; and a Contact Measurement System (CMS) consisting of a Puma 260 robot equipped with a force/torque sensor, mounted on a linear stage. The entire facility can be teleprogrammed in Java from any location on the Internet.

The current ACME system has been enhanced in a number of ways. A special **Sensor** interface (Sect. 2.1) has been added to facilitate the real-time

*Supported in part by grants from NSERC and IRIS NCE.



Figure 1. ACME facility being used to map surface friction and roughness of a bottle.




Figure 2. ACME simulation detail showing the bounding boxes. Boxes involved in a collision are highlighted.

collection of data from the system's various sensors, and the ACME simulator has been redesigned to facilitate the easy addition of classes for simulating sensor input. The entire ACME facility is now modeled using an enhanced Java3D scene graph, containing all the geometric, positional, and appearance information used in simulation, run-time checking, and graphic display. One of the features of this scene graph is a collision detector, which performs rapid on-line collision detection using oriented bounding boxes [4] placed around principal system components (Fig. 2).

2.1. Sensor interface

An ACME experiment collects data from the system's sensors (e.g., cameras, actuator position sensors, force/torque sensors, and microphones) through implementations of a `Sensor` interface. This provides an abstract interface to the underlying sensing framework which has several important features. Most importantly, it enables asynchronous real-time data acquisition from a Java program. The user's Java `Experiment` controls data acquisition by specifying trigger conditions (which can be sophisticated methods written in Java) to decide when to start and stop data acquisition based on the sensed data. The interface provides a number of utility methods that permit the experiment to adjust the frequency at which data is returned from the sensor, enable/disable the time stamping of data, and adjust the buffering within the `SensorInputStream`. Finally, the framework allows the `Experiment` to connect either to a remote `SensorServer` object running on the main ACME server, or (when an experiment is run in simulation) to a `SensorServer` object running in the simulator on the ACME client.

3. Contact Texture Modeling

Friction and roughness are important components of haptic texture and are required for contact simulation. Real objects usually have multiple materials and surface finishes that vary over the surface of the object. We can measure these in ACME using the CMS and force sensor, by analyzing the force profile produced by stroking the surface with a round-tipped probe (Fig. 1). The task is non-trivial for several reasons. First, friction estimates must account for variation in the surface normal due to curvature and uncertain geometry. Second, we need to be able to extract from the noisy force profile a simple model of surface roughness. Third, exploring the entire surface can be very time consuming if done in an exhaustive fashion, and so a more efficient approach is needed.

3.1. Friction Estimation

The problems in estimating friction can be handled using a differential measurement approach. The probe strokes the surface along a short path using a compliant motion that applies pressure to the surface. This will result in a normal force component \mathbf{f}_n and a tangential component of magnitude $\|\mu\mathbf{f}_n\|$ (where μ is the coefficient of friction) in a direction opposite to that of the motion. If \mathbf{f}_n were known, then μ could be computed directly. Unfortunately, \mathbf{f}_n is not accurately known because the surface normal is itself not known accurately and may also vary along the path. To compensate for this, we stroke the surface again, along the same path but in the opposite direction. At any point along the path, we then have a force value \mathbf{f} which was measured during the forward motion, and another value \mathbf{f}^* which was measured during the reverse motion.

Now referring to Fig. 3, \mathbf{f} and \mathbf{f}^* each have components parallel to the surface normal \mathbf{n} , and friction components \mathbf{f}_f and \mathbf{f}_f^* which lie opposite to the directions of motion and are perpendicular to \mathbf{n} . Now even if \mathbf{n} is unknown, and the magnitudes of \mathbf{f} and \mathbf{f}^* differ, we can still estimate μ by calculating the angle between \mathbf{f} and \mathbf{f}^* :

$$\mu = \tan(\theta/2). \quad (1)$$

This calculation is independent of the speed of travel, the force of the probe against the surface (in either direction), the orientation of the probe, and of course the surface normal itself. By averaging the values μ obtained at various points along the path, a reasonable estimate of the friction coefficient may be obtained.

3.2. Roughness characterization

The force profiles obtained by stroking the surface can also be used to identify the surface roughness (Fig. 4). We wish to characterize this roughness using




Figure 3. Forces associated with motion in two different directions along an object surface.




Figure 4. Force profiles for two different types of material: rubber gasket sheeting (top), and masonite (bottom).




Figure 5. Zero-meaned AR(2) simulations based on models extracted from the profiles in Fig. 4. Top: rubber gasket sheeting, with $a_1 = 0.60$, $a_2 = 0.25$, and $\sigma = 0.11$. Bottom: masonite, with $a_1 = 1.61$, $a_2 = -0.73$, and $\sigma = 0.16$.

a simple model that is suitable for easy simulation, particularly in haptic displays. Intuitively, we can see that the roughness is largely indicated by the “noisiness” of the profile, as characterized by the noise amplitude (variance), and spatial density. In addition, the profile will often contain a non-random periodic component, particularly in man-made materials, due to manufacturing processes, or functional or aesthetic requirements (as evidenced by the masonite in Fig. 4).

These qualities can be captured by treating the force profile $f(x)$ as a discrete random series described by a second-order autoregressive model, denoted AR(2):

$$f(kL) \equiv f_{(k)} \approx a_1 f_{(k-1)} + a_2 f_{(k-2)} + \sigma \epsilon_{(k)}, \quad (2)$$

where k is the sample number, L is the length of the spatial discretization, a_i are the model coefficients, and $\sigma \epsilon_{(k)}$ is a sample from a white noise sequence with standard deviation σ .

The model of Eq. 2 is obviously very efficient to simulate, and is completely defined by the parameters a_1 , a_2 , and σ . These can be readily determined, for instance by autoregressive parameter estimation via the covariance method (e.g., the `arconv` function in Matlab). The parameters and associated simulated signals for the profiles in Fig. 4 are shown in Fig. 5. Higher order AR(p) models or ARMA models could also be used if necessary.

3.3. Efficient surface exploration

To avoid having to sample contact texture exhaustively over the entire surface, we employ a hierarchical exploration approach, similar to that used in the acoustical modeling described in Sect. 4 of this paper. The surface geometry is described using a Loop subdivision mesh [8], and the robot initially samples

friction and/or roughness at locations corresponding to the vertices of coarsest resolution mesh. If the contact texture properties at adjacent vertices differ significantly, we then sample between them, at a vertex of the next higher resolution mesh.

This method was used to map the surface friction for the bottle shown in Fig. 1, to which sandpaper patches had been applied to provide regions of high friction. The resulting map is shown in Fig. 6; the sample locations are shown by small line segments. The sampling decisions were made entirely by the planner: notice that the hierarchical exploration correctly clusters the samples near the boundaries between areas of constant friction.




Figure 6. Measured friction map of bottle. Dark regions represent low friction ($\mu \approx 0.13$ for the glass), and light regions represent high friction ($\mu \approx 0.5$ for the sandpaper patches shown).




Figure 7. Contact sound measurement

4. Contact Sound Modeling

Contact sounds provide important cues for the shape and material properties of objects [6], and can be efficiently simulated [13]. By “contact sounds”, we mean the sounds produced by the interaction of two objects’ surfaces (e.g., scraping, knocking, rolling, etc.). Such sounds are useful in many applications, including interactive virtual environments, physical simulations, tele-operative controls and robotic perception of material [7].

A sound model may be viewed as the acoustic analog to a texture map in graphics. The parameters of contact sound models vary over an object’s surface as they are affected by shape, material, composition and surface texture. For most everyday objects, their complex geometry and material composition render analytical solutions to the sound model difficult, if not impossible. An empirical solution is more tractable and practical, particularly when construct-

ing sound models for many arbitrary objects.

Many techniques for acquiring the measurements needed to create empirical sound models have been tried, especially in the realm of musical instrument modeling [1, 3, 14]. Many of these techniques, though appropriate for their application, are inaccurate in either location or impact profile. Also, acquiring densely sampled measurements by these manual techniques is often tedious and subject to human error. ACME was recently equipped to perform the first automatic measurements and estimations of contact sound models [10].

We measure the sound response of an object in ACME using the “sound effector” device [10]. This push-type solenoid mounted on the end-effector of the CMS robot is used to apply light near-impulsive impacts at specified locations on the surface, and the resulting sound is recorded using a microphone (see Fig. 7).

As in our contact texture modeling experiments (§3), we employ an adaptive algorithm to select sample locations automatically. The vertices of the object’s surface are selected as an initial set of sample points. As sound models are constructed at adjoining vertices, the sounds are compared, and new sample locations added if the difference between the two sounds models is perceptually relevant. Models are compared by a sound metric derived from perceptual studies of contact sounds [6].

An additional advantage of using ACME to perform sound model acquisition is repeatability; multiple measurements can be recorded at the same surface location. Spectral averaging is performed across samples at the same location to produce a prototypical sound. Spectral averaging yields a 10 to 20% reduction in parameter estimation error [11].

The parameters of the sound model are estimated from the prototypical sound at each location. Complete sound models have been created for a variety of objects using ACME. Figure 8 demonstrates the fidelity of the model at one location on the surface of the glass bottle of Fig. 1. Note that the amplitude, decay and frequency of the most dominant frequency modes have been correctly identified in the model.

5. Deformation Modeling

Modeling the deformation behavior of real objects is a major challenge. There has been recent progress in real-time simulation of linear elastic objects [5, 2]. To construct such models requires measuring the deformation of an object in response to applied forces.

In particular, a linear elastic map Ξ relates known (prescribed) tractions and displacements x at the boundary to the unknown complementary set of tractions and displacements b .

$$\Xi x = b$$

The matrix Ξ represents the discrete Green’s functions of the boundary value problem associated with the object’s elastostatic behavior. These Green’s functions can be used to quickly solve boundary value problems encountered in haptic simulation, e.g. produced by touching the object and moving the point




Figure 8. A spectral comparison of a recorded acoustic impulse (a) and an acoustic impulse synthesized from a prototypical model (b). The prototypical model was created by spectral averaging over three recordings of the bottle at the same surface location. The initial amplitude, decay rate and frequency of the forty most-dominant modes were estimated to create the model. These modes are illustrated by the horizontal peaks in the spectrogram.

of contact over the surface [5]. The Green’s functions for an object with known geometry and known homogeneous, isotropic linear-elastic material properties can be calculated in various ways. One approach is to apply the boundary element method and obtain the map directly (for details and further references, see [5]); another approach is to discretize the object’s volume and apply a finite element method followed by a condensation step [2].

In practice, for most everyday objects, the geometry and material properties are not known. We would like to directly estimate the Green’s functions based on active deformations performed using ACME.

The Green’s functions correspond to a particular reference boundary configuration. In the reference configuration, the object’s fixed support surface has a displacement boundary condition, and the free surface has a traction boundary condition. The contact measurement system (CMS) in ACME applies a sequence of point-like forces to the free surface. The CMS’s sensors provide a local measurement of the displacement of the contact point and the applied force. The global object deformation is sensed using the Triclops stereo vision system¹ mounted on the field measurement system (FMS).

We have developed a novel surface deformation measurement technique utilizing the Triclops trinocular stereo-vision system. The measurement of surface deformation is based on three-dimensional scene flow. Optical flow (see Fig. 9) is calculated simultaneously through a sequence of calibrated stereo images. The stereo images are also matched utilizing the Triclops Stereo Vision library. Combining the optical flow in the sequence of stereo triplets with the depth information from stereo provides three-dimensional scene flow (see Fig. 10). A geometric model of the undeformed surface of the object under test

¹Point Grey Research, Vancouver, Canada, <http://www.ptgrey.com/>.




Figure 9. Segmented image flow during deformation. (Zoom shows every 4th flow vector.)

is either assumed known or acquired beforehand. This geometric model and the tracking capabilities of the system allow the three-dimensional flow to be segmented into surface flow and non-surface flow.

Fig. 9 shows the segmented image flow for deforming a test object; here a stuffed toy tiger. The initial geometry is represented as a triangular mesh (Fig. 10). The mesh is estimated from Triclops data with the aid of a volumetric surface reconstruction software from the National Research Council of Canada [12]. The resulting displacement at mesh nodes is shown in Fig. 10. The applied force and local displacement is simultaneously measured using the CMS.

In finding an object's Green's functions, an estimation step based on the measured local and global deformations is next. We are currently exploring ways to robustly perform this estimation. Questions which ACME will enable us to explore include estimation robustness, the adequacy of linear elastic maps for non-homogeneous bodies and a comparison of simulation and measurement.

6. Conclusions

The ACME facility has matured and is demonstrating the promise of reality-based modeling. Several improvements make the system more robust and easier to program. The abstract sensor framework allows the acquisition of large amounts of data in a timely fashion. The improved simulation protects the facility from many operator errors. ACME can now be used to acquire models of a wide variety of everyday objects from metallic vases to stuffed toys. The object's response in multiple modalities can be excited and measured, enabling the modeling of contact texture, sound, and deformation.




Figure 10. Mesh node displacements. (Darker nodes indicate larger displacements.)

References

- [1] Perry R. Cook and Dan Trueman. A database of measured musical instrument body radiation impulse responses, and computer applications for exploring and utilizing the measured filter functions. International Symposium on Musical Acoustics, Acoustical Society of America, Woodbury, NY, 1998.
- [2] S. Cotin, H. Delingette, J-M Clement, V. Tassetti, J. Marescaux, and N. Ayache. Geometric and physical representations for a simulator of hepatic surgery. In *Proceedings of Medicine Meets Virtual Reality IV*, pages 139–151. IOS Press, January 1996.
- [3] Robert S. Durst and Eric P. Krotkov. Object classification from analysis of impact acoustics. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1:90–95, 1995.
- [4] Stefan Gottschalk, Ming Lin, and Dinesh Manocha. Obb-tree: A hierarchical structure for rapid interference detection. In *Computer Graphics (ACM SIGGRAPH 96 Conference Proceedings)*, pages 171–180, August 1996.
- [5] Doug L. James and Dinesh K. Pai. ARTDEFO, accurate real time deformable objects. In *Computer Graphics (ACM SIGGRAPH 99 Conference Proceedings)*, pages 65–72, August 1999.
- [6] R. L. Klatzky, D. K. Pai, and E. P. Krotkov. Hearing material: Perception of material from contact sounds. *PRESENCE: Teleoperators and Virtual Environments*, 9:4, pages 399–410, The MIT Press, August 2000.
- [7] Eric Krotkov. Robotic perception of material. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 88–94, 1995.

- [8] Michael Lounsbery, Tony D. DeRose, and Joe Warren. Multiresolution analysis for surfaces of arbitrary topological type. *ACM Transactions on Graphics*, 16(1), pages 34–73, January 1997.
- [9] Dinesh K. Pai, Jochen Lang, John E. Lloyd, and Robert J. Woodham. Acme, a telerobotic active measurement facility. In *Proceedings of the Sixth Intl. Symp. on Experimental Robotics*, 1999.
- [10] J. L. Richmond and D. K. Pai. Active measurement and modeling of contact sounds. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, pages 2146–2152, San Francisco, April 2000.
- [11] Joshua L. Richmond *Automatic measurement and modeling of contact sounds*. MSc. thesis, University of British Columbia, August, 2000.
- [12] G. Roth and E. Wibowoo. An efficient volumetric method for building closed triangular meshes from 3-D image and point data. In *Proceedings Graphics Interface*, pages 173–180, 1997.
- [13] Kees van den Doel and Dinesh K. Pai. The sounds of physical shapes. *Presence*, 7(4), pages 382–395, 1998.
- [14] Kees van den Doel *Sound synthesis for virtual reality and computer games*. PhD thesis, University of British Columbia, May, 1999.