Robotic Perception of Material

Eric Krotkov
Robotics Institute, Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, Pennsylvania 15213
epk@cs.cmu.edu

Abstract

In this paper, we develop a conceptual framework in which acts of manipulation are undertaken for the sake of perceiving material. Within this framework, we disambiguate different materials by actively contacting and probing them, and by sensing the resulting forces, displacements, and sounds. We report experimental results from four separate implementations of this framework using a variety of sensory modalities, including force, vision, and audition. For each implementation, we identify sensor-derived measures that are diagnostic of material properties, and use those measures to categorize objects by their material class. Based on the experimental results, we conclude that the issue of shape-in variance is of critical importance for future work.

1 Introduction

Most past and present research in robotic perception concerns shape and position. Relatively little research addresses material properties, such as compliance and density. For robots to go beyond their current limits and reach their potential, they must understand not only where things are and what they are shaped like, but also what they are made of.

By definition, a material property is independent of the size and shape of a particular sample. Although there are visual cues to material properties (e.g., surface luminance is a cue to the coefficient of friction), reliable determination of the material composition of an unknown object generally requires contact with it. Humans who wish to determine material properties show stereotypical patterns of manual exploration; they press, poke, tap, heft, squeeze, shake, rub, and strike, according to the type of information desired [Lederman and Klatzky, 1987].

To formalize this kind of information gathering, and to understand it well enough to program a robot to perform it, we must answer some fundamental questions. What physical principles relate the acts of manipulation to percepts such as forces, torques, displacements, and vibrations? What material properties can be inferred from a given act of manipulation such as tapping? What acts of manipulation best reveal a given material property? Which material properties can be determined only by contact sensing? Which combinations of sensing modalities are most informative about a given material property? How can material and shape properties be deconflated?

Let us elaborate further this last question by way of an example. Consider a metal rod and a wooden rod of exactly the same length. When you strike the metal rod with your knuckle, it rings; when you strike the wooden rod, it produces a much shorter "thud" sound. This difference in the sound despite the same excitation is due to the difference in the way that the materials vibrate, which in turn is due to stress/strain properties. The rods sound different because they have fundamentally different material properties, so sound waves travel through them quite differently.

Now consider two metal rods that are identical except that one is twice as long as the other. Given the same excitation, the shorter rod will "ring*' at a higher frequency than the longer rod. The rods sound different because the waves travel different distances inside them.

How can these differences in the way things sound, one due to material and one due to shape, be resolved? Or, in other words, What perceptual information is diagnostic of material, but invariant over object shape?

We will address this question in Section 4.4. For the moment, we content ourselves with the observation that this fundamental question (which is just one of the several posed earlier) has barely been explored in the perception literature. One reason is that, compared to the perception of shape or position, perception of material properties is a field in its early infancy, unsure of its vocabulary, models, and strategies.

To supply answers to the fundamental questions, we have been advancing toward a theory that specifies constructively the joint perception and manipulation operations that are necessary and sufficient to identify what things are made of. As a first step toward such a theory, we have begun to develop and demonstrate the robotic capability to distinguish a wide variety of materials.

In our approach, materials are disambiguated by actively contacting and probing them and by sensing the resulting forces, displacements, and sounds. One can visualize this capability by imagining a game of non-verbal 'Twenty Questions," in which one player is the robot and the other player is any object placed in the robot workspace. The robot probes (presses, pokes, taps, etc) the object, in effect asking questions about the object stiffness, strength, density, and other material properties. At the end of the game the robot announces its decision about the material composition of the object.

In this paper, we report on our first steps in achieving this capability. In the next section, we describe potential applications of the capability to perceive material, and in the following section we review related research. In Section 4, we present results in identifying materials and their properties from several robotic probing tests implemented with the diverse sensing modalities of force, vision, and audition. We conclude the paper with a critical discussion of progress to date, and a statement of future work.

2 Potential Applications

Knowledge of material properties and classes can improve performance of many tasks involving physical interaction. In many real-world scenarios, such knowledge is not given in advance; instead, it must be determined at a worksite or in the field, without jigs or fixtures.

Grasping Grasping is a generic capability required for a wide variety of manipulation tasks. Typically, to plan grasping actions, current robots analyze geometric properties of the hand and of the object. However, there are classes of objects, among which are slippery objects like a wet bar of soap or deformable objects like a ripe peach, for which the geometric information is not sufficient. Robots that perform initial probing actions to learn the object's material properties can exploit that non-geometric information to successfully plan and execute grasping tasks for a larger class of objects under a larger class of conditions.

Non-Destructive Evaluation and Inspection Consider a manufacturing process that employs a finite element CAD model for producing a part. Once manufactured, the question is whether the manufactured part meets specification tolerances. To determine the answer to that question, a robotic inspection strategy can be automatically generated from the CAD model. First, the robot would derive from the CAD model the mass, damping, and stiffness matrices of the object. These matrices specify the relationships between each point within the part. Next, by applying the finite element method and using the matrices, predict how the part will respond to impinging forces. Then, command the robot to probe the part (for instance by tapping it) to excite a characteristic response (for instance a vibration). Finally, compare the predicted and observed responses, which will agree if the part is well-formed.

Reasoning about Functionality Perception of material types and properties will contribute significantly to the emerging area of research on reasoning about object functionality. Most of this research has concentrated on geometric reasoning, for example, recognizing a pair of scissors by analyzing visual image sequences showing the articulations of its handles and blades [Stark, 1994]. The ability to classify objects by their material properties will permit deeper reasoning, for example, recognizing that a hard-heeled shoe could substitute for a hammer, even though their shapes differ dramatically.

Handling Hazardous Waste One of the world's urgent needs is to clean up hazardous waste. A key challenge is to sort and process mixed waste, including more than 11,000 cubic yards of low level solid waste and 1,000 cubic yards of transuranic waste. Stored in drums and boxes, much of the waste is wrapped in multiple layers of plastic film, which makes non-contact analysis difficult or impossible. Robots capable of probing the waste items and classifying them based on their material properties could serve as the basis of a sorting workcell. This advance will permit human operators to avoid these dangerous substances.

Recycling An emerging national priority is to recycle non-hazardous waste. In most current recycling systems, human operators sort cans, bottles, and newspaper. Recycling robots will not be able to rely on shape, because cans of different composition (aluminum, tin, steel) may have identical shapes, as may bottles of different composition (glass, plastic). Nor will they be able to rely on appearance, because labels may cover the object surface. To automate recycling and free human operators from working with garbage, the robots must be capable of probing objects and discriminating between diverse substances.

Excavating The excavation industry must handle many different materials, ranging from soil to pipes to boulders. An excavator must respond to the mechanical properties of both surface and sub-surface materials. When it encounters buried obstacles such as rocks or timbers or utility pipes, it must decide whether to lift them out, slide them away, or work around them. A robot excavator capable of probing an obstacle and inferring its material type has a solid basis for selecting the proper actions. This could relieve human operators from long duty in harsh, dirty environments.

Traversing Natural Terrain Consider an outdoor mobile robot that must traverse rugged natural terrain. Without prepared surfaces such as floors and roads, the robot risks traveling over unstable or hazardous regions. With active probing techniques, the robot can evaluate the trafficability of the terrain, and follow routes that avoid treacherous regions.

3 Related Research

The artificial intelligence, robotics, civil engineering, mechanical engineering, and materials literature documents two families of techniques to estimate mechanical and mass properties, one employing non-contact sensing, the other employing contact sensing.

3.1 Non-Contact Sensing

Non-contact sensing can provide significant information about the material composition of objects. Indeed, this is an area of intensive research in physics-based vision. The work of [Wolff and Boult, 1991] exemplifies this field. They developed methods to analyze the magnitudes of the polarization components of reflected light, which permitted them to segment material surfaces according to varying levels of relative electrical conductivity, and in particular to distinguish dielectrics from metals. [Caillas, 1990] developed analytical techniques operating on thermal images to produce estimates of material grain size sufficient to distinguish between dust, sand, and rock. The remote sensing literature cites many other techniques, such as a back-scattering and impulse radar.

Non-contact methods suffer from two fundamental deficiencies: they are superficial and indeterminate. They are superficial to the extent that they are sensitive only to the surface of the object. An adhesive label or a coat of paint can mask the underlying material composition and properties. They are indeterminate in the degree that they cannot directly measure properties such as density and friction. The properties of interest must be observed using methods providing richer information than non-contact sensing can provide.

3.2 Contact Sensing

Contact sensing techniques draw from diverse sensing modalities, including force, tactile, vision, and temperature.

Force sensing has a long history in the recovery of material properties. Load cells are used in civil engineering laboratory devices such as shear boxes and tri-axial test cells. The shortcoming is that the sample must be collected and placed inside the instruments. The cone penetrometer is appropriate for field tests, but it estimates only a few mechanical properties [Bekker, 1969], and does not permit identification of mass properties.

In recent research with force sensing, [Pentland and Williams, 1989] formulated estimators for stiffness, strength, and mass (they used both contact and non-contact sensors). [Bicchi *et al.*, 1989] assessed the deformation of rubber blocks and estimated coefficients of static friction with an instrumented leg-ankle-foot system. In a separate effort, [Bicchi *et al.,* 1991] evaluated friction characteristics with an articulated robotic hand. By using a compliant wrist mounted on a Puma 260 manipulator, [Sinha *et al.,* 1990] implemented procedures to recover mechanical properties such as penetrability, hardness, and surface roughness. By using a robotic backhoe for excavation, [Bernold, 1993] estimated soil properties from cutting forces and moments.

Employing tactile sensors and kinesthetic feedback, [Stansfield, 1991] developed exploratory procedures to extract naptic primitives such as compliance and elasticity. Using vision sensors, [Gandolfo *etal,* 1991] developed methods for determining the stability and rigidity of a struck object by analyzing optical flow fields. Applying temperature sensors, [Campos *et al.,* 1991] describe an exploratory procedure that returns the thermal property of an unknown object, thereby enabling discrimination of materials such as aluminum and wood.

3.3 Unresolved Issues

These efforts in contact and non-contact sensing each advance the state-of-the-art. Still, the capability to sense and manipulate varying materials and substances remains to be either formulated or demonstrated. We believe there are two shortcomings of current work responsible for this state of affairs.

First, the surveyed works described actions such as looking, pressing, shearing, and heating. These actions are dainty and timid and do not reveal enough about the state of nature. We need to develop more aggressive probing actions such as striking, tapping, and poking.

Second, the current generation of exploration procedures is piecemeal and partial. In most cases they estimate no more than one or two of the many material properties. We must extend the robotic repertoire by developing techniques that take into account all of the key variables, instead of only a few. Further, in most cases, they execute according to a preplanned script. We need to develop procedures that use sensor feedback on-line to guide further exploration.

4 Experiments

Before mounting an extensive effort to address the outstanding issues, we conducted preliminary experimental investigations to better understand the general nature of the material perception problem, and to learn about the constraints imposed by using different sensors and different materials. We were particularly motivated to gain practical insight into sensor characteristics, and the transduction of material properties into sensor signals. This led us to conduct four investigations, using force, vision, and acoustic sensors.

4.1 "Whack and Watch"

The objective of the "Whack and Watch" effort is to estimate the mass and friction properties of an unknown object. In this preliminary work, we strike a variety of objects with a home-made wooden pendulum, and observe their trajectories with a camera. Then we estimate the object's mass and the coefficient of sliding friction.

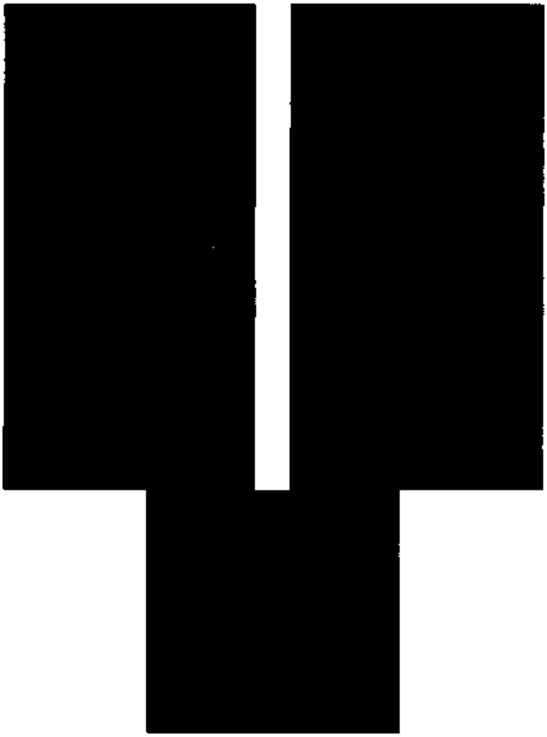

Figure 1: Image sequence of pendulum striking object

We use the pendulum because its dynamics are simple enough to permit advance calculation of the instantaneous force, eliminating the need for force sensing. We acquire a sequence of images (Figure 1) of the pendulum making contact and of the object sliding. For each image, we employ basic blob-finding techniques to determine the position of the

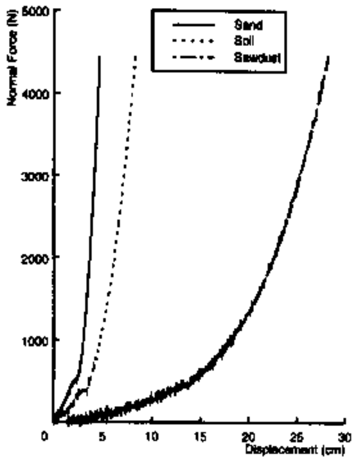

Figure 2: Force-displacement characteristics for sand, soil, and sawdust

struck object. From the sequence of positions we derive the velocity and acceleration of the object.

We determine the mass of the object using the equality between the impulse of the force and the momentum of the object after impact

$$\int f\, dt = mv$$

We determine the coefficient of sliding friction u_s using the equation of motion under constant acceleration as

$$\mu_s = \frac{-v}{gt} .$$

(Alternatively, we could use the definition of work as the change in kinetic energy, and determine the coefficient of sliding friction a $\mu_s = \frac{v^2}{2gl}$ \r e / is the distance the body slides.) The computed masses are within a factor of two of the true masses, and the coefficients of friction are within 25 percent of the values listed in the CRC Handbook of Chemistry and Physics. For the purposes of these preliminary trials, we performed neither detailed modeling nor calibration, so the factor of two accuracy can easily be improved.

4.2 "Step and Feel"

The objective of the "Step and Feel" effort is to identify the compliance of natural terrain. In this work, a robot leg steps on different terrains and measures normal forces exerted on the foot. Then we interpret the forces in terms of the compliance of the terrain. [Krotkov, 1990] presents the methods and results in detail.

We use one leg of a walking robot as the testbed because it is instrumented with a six-axis force-torque sensor and is readily available. We suspend the leg above the ground, and attach cables to prevent it from moving in the plane parallel to the ground.

We measured the applied normal force / and the vertical displacement *z* while stepping on different materials: sand, sawdust, and soil. We drive the foot into the terrain at low velocity (typically 1 mm/s) in order to minimize dynamic effects such as strain-rate hardening. We sample the normal force and vertical position at a fixed rate (typically 60 Hz) and continue stepping until reaching a specified load (typically 4450 N).

The observed vertical force-displacement response (Figure 2) differs dramatically for the three materials tested, demonstrating that the samples can be discriminated without great difficulty.

43 "Hit and Listen" with Fixed-Shape Objects

The objective of the "Hit and Listen" effort is to classify objects from the sounds they make when struck. In this preliminary work, we hit a variety of objects by dropping a blind person's cane from fixed heights, and listen to the sounds of impact with a microphone. Then we classify the different objects by analyzing patterns in the frequency domain. [Durst and Krotkov, 1993] describe the approach in detail.

The apparatus includes a cane, selected because of its simplicity and low cost, and a cylindrical guide through which it is dropped. A condenser microphone element transduces the acoustic energy, and low-pass filters cut off the signal at 20 kHz. The objects to be classified include a wooden block, a concrete brick, a clay brick, a zinc ingot, and a ceramic tile.

We digitize the microphone signal and extract spike features from its power spectrum (Figure 3). The feature extraction proceeds in two steps. First, it eliminates all points that are not local maxima in their immediate (three-point) neighborhood. Second, it eliminates all of those surviving points that are not local maxima in a larger neighborhood, whose size is derived from a resolvable frequency tolerance (approximately 60 Hz), which in turn is derived from the signal-to-noise ratio for the apparatus.

Based on these features, we classify the test object as one of the five objects with a hybrid minimum-distance and decision-tree classifier. Like a decision tree, the hybrid classifier refines its set of possible class assignments based on analysis of different features at different times. This analysis itself is a minimum-distance classification scheme: at each point, the hybrid classifier iteratively finds the nearest-neighbor prototypes in defined intervals around the test vector. The test sample is assigned to the closest matching prototype's class iff there are prototypes from a single class in the nearest occupied interval. If more than one classes' prototypes are in the nearest occupied interval, the hybrid classifier moves to the test vector's next feature. This is analogous to moving down a decision tree. The classifier achieves 97 percent accuracy on 580 training samples, and 94 percent accuracy on 240 test samples.

4.4 "Hit and Listen" with Variable-Shape Objects

The objective of the "Hit and Listen" with variable-shape objects effort is to formulate acoustic measures that are diagnostic of material properties, and hence invariant with respect to shape and size, and further, to use those measures to categorize objects by their material class.

[Wildes and Richards, 1988] have advanced a theoretical approach to recovering the material type of an object from the sound generated when it is struck. They restrict their attention to anelastic solids, and study the modulus of compliance as the key to understanding the vibration of the struck object.

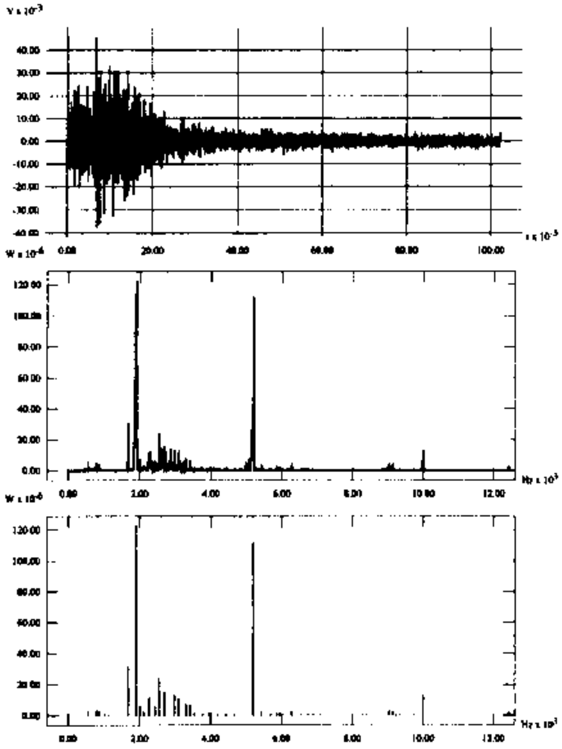

Figure 3: Acoustic signal (top), power spectrum (middle), and spikes

Following classical analysis they relate the modulus of compliance to the angle of internal friction. This is an intrinsic (shape-invariant) property of a given material.

They propose two methods for determining the angle of internal friction of an unknown sample: one that impulsively excites the sample and then measures the acoustic decay rate; another that periodically (say, sinusoidally) excites the sample and then identifies the bandwidth of the acoustic signals. They did not experimentally verify either method, although they did cite evidence from earlier empirical studies [Gemant and Jackson, 1937].

We have explored in some detail the decay rate approach. In this approach, the angle of internal friction 0 is determined by the time t_e it takes the amplitude of vibration to decay to 1/e of its original value after the material sample is struck. According to Wildes and Richards,

$$\tan\phi = \frac{1}{\pi f\, t_e}, \qquad (1)$$

where f is the observed frequency associated with the amplitude. Assuming exponential decay, it follows that

$$t_e = -\frac{1}{\log\theta}, \qquad (2)$$

where θ is a retention parameter representing the proportion of the amplitude present at time t_i that is still present at t_{i+1}. For an exponential process, θ is constant for all i.

To assess the validity of the decay rate approach to identifying the angle of internal friction, we produced thin rods of wood, brass, aluminum, glass, and plastic. For each material, we produced two rods, one of length L = 15 cm and one of length $2L$ = 30 cm. We suspend each rod, strike it impulsively, and digitally record the resultant sound. We analyze the digital signal *x[n]* as follows:

1. Compute the spectrogram of the signal, which describes the distribution of the signal energy in the time-frequency plane. The spectrogram *S[l, k]* is the squared modulus of *X[l, k),* where

   $$X[l,k] = \sum_{n=-\infty}^{+\infty} x[n]g[n-l]e^{-j\frac{2\pi k}{n}l}$$

   is the discrete-time Fourier transform of a windowed version x[n]g[n - /] of the original signal.

   As an example, Figure 4 shows the original signal recorded after striking an aluminum rod, and the spectrogram of that signal. The energy is concentrated in three main bands, at approximately 2500,4000, and 6000 Hz.
2. Determine where the signal begins. Due to the impulsive contact, the early part of the signal contains energy at all frequencies. This transient effect (or click) does not convey meaningful modal information, so we desire to exclude this segment of the signal from analysis.

   For this, we examine temporal correlations between the spectrogram magnitudes. During the click, the spectrogram magnitudes are highly correlated from instant to instant. Immediately after the click, as energy begins to concentrate in relatively narrow bands, the spectrogram magnitudes are poorly correlated. Well after the click, the spectrogram magnitudes are again highly correlated. Based on these observations, we determine where the signal begins by identifying when the correlation coefficients rise from the dip due to the end of the click.
3. Find bands of concentrated energy. First, for each time step, we threshold out those spectrogram magnitudes that contribute little to the total power at that time. Next, we analyze the connectivity between surviving magnitudes. Finally, for each connected component, we identify the minimum bounding rectangle, and call that a band. Figure 4 shows three bands overlaid on the spectrogram.
4. For each band, fit a line to the within-band log power. The slope of each line determines the log *6* term in (2). Now it is possible to determine tan nellset for each frequency band by substituting (2) into (1).

   Figure 4 illustrates the total power associated with the three bands (band 1 is at 6000 Hz, band 2 is at 4000 Hz, and band 3 is 2500 Hz), and also the background (defined as all spectrogram magnitudes that did not pass the test in Step 3). In addition, the figure shows the lines fit to the three power curves. Bands 1 and 2 are reasonably linear (r values above 95 percent, that is, the fit accounts for more than 90 percent of the variance), and band 3 is not (r value of 70 percent).

The results of these preliminary experiments indicate that the differences in slope between materials are significant, and that the differences in slope between lengths of the same

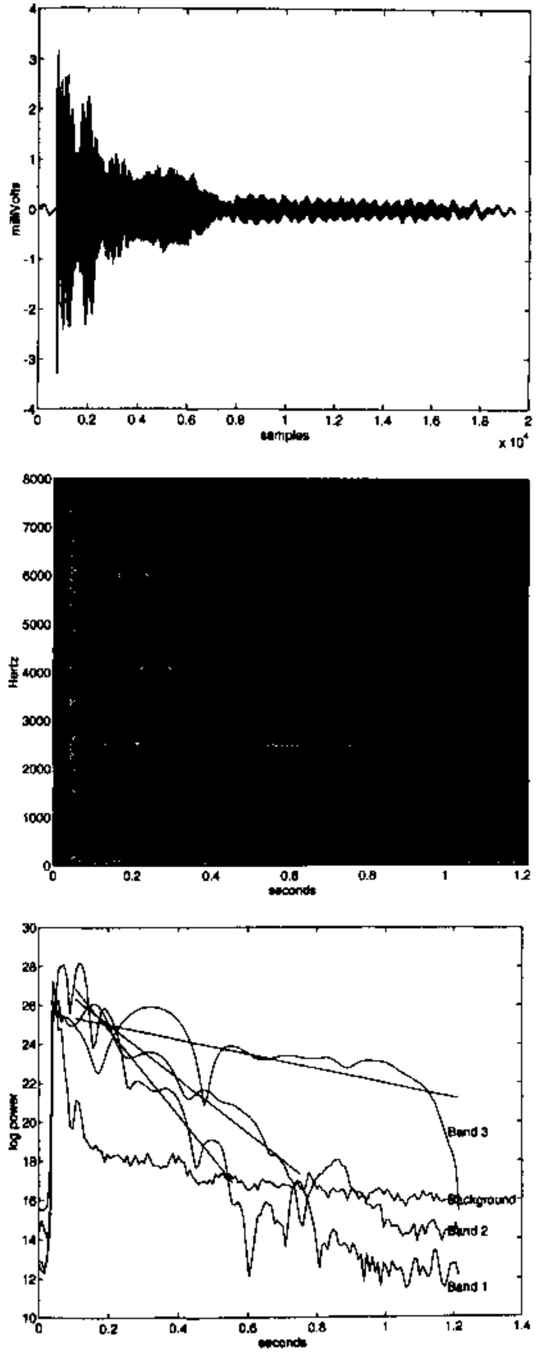

Figure 4: Analysis of sounds generated by striking a thin aluminum rod: microphone output (top), spectrogram with bands identified (middle), linear fits to log power within bands (bottom).

material are far less significant. It remains to perform the analysis described above on a statistically significant number of samples; only then will we be able to assess the practical effectiveness of the approach.

5 Discussion

In this paper, we have formulated a research agenda for the new field of perceiving materials. We explored the use of several sensory modalities, including force, vision, and audition. For each, we identified sensor-derived measures that are diagnostic of material properties, and used those measures to categorize objects by their material class.

We view the principal contributions of the work to be the following:

- Articulation of fundamental questions to be answered by future work on perception of material
- Development of a conceptual framework in which acts of manipulation are undertaken for the sake of perceiving material. Within this framework, we disambiguate different materials by actively contacting and probing them, and by sensing the resulting forces, displacements, and sounds.
- New robotic probing techniques to recover material properties.

Although we have made progress, much work remains to advance the new field of perception of material from its early infancy to childhood. We will concentrate our efforts in the near future on the auditory sensory modality, primarily because of the promising results from our preliminary investigations (Sections 4.3 and 4.4), and secondarily for the relative simplicity of the experimental effort.

We take the primary challenge to be the formulation of shape-invariant acoustic measures of material properties. (As an aside, we observe a parallel between this challenge and one of the key challenges in visual perception. In audition, numerous factors are conflated in sound formation; specifically, the decay rate of an impact sound depends both on material composition and on geometry. In geometric vision, numerous factors are conflated in perspective image formation; for example, the size of the projection of an object depends both on its size and on its distance from the viewer. In both material-seeking audition and shape-seeking vision, the key problem is disentangling or deconflating the different factors.) We take the secondary challenge to be object categorization by material class.

In order to achieve shape-in variance, we relied heavily upon the concepts of the modulus of compliance and angle of internal friction in the work reported in Section 4.4. However, there are numerous alternative approaches that merit exploration.

We will explore work in computer vision on view-invariant object recognition. The underpinnings of that work involve group theory and projective transformations. We will search for formulations of shape-invariance that parallel their formulations of view-invariance.

We will explore work in optics that studies point scattering within a medium and refraction at the interface between media. This may provide useful analogs for the acoustic phenomena of wave propagation within the solid material and impedance matching as the wave exits the solid and enters the surrounding air as it travels to the microphone.

We will also explore higher-dimensional representations of the acquired acoustic data. For example, we will search for structure in spectrograms, and such three-dimensional plots as material (angle of internal friction) versus structure (sample length) versus highest-frequency vibration.

Still more alternative approaches need to be pursued before the new field of perception of material emerges from its early infancy, and realizes its potential for revolutionizing robotic interaction with the real world.

Acknowledgments

Robert Durst, Wes Huang, Roberta Klatzky, Richard Stern, and Nina Zumel have made significant contributions to the work reported here.

References

[Bekker, 1969] M. Bekker. *Introduction to Terrain-Vehicle Systems.* University of Michigan Press, Ann Arbor, Michigan, 1969.

[Bernold, 1993] L Bernold. Motion and Path Control for Robotic Excavation. *J. Aerospace Engineering,* 6(1), January 1993.

[Bicchi *etaL,* 1989] A. Bicchi, L. Bosio, P. Dario, M. Guiggiani, E. Manfredi, and P. Pinotti. Leg-ankle-foot system for investigating sensor-based legged locomotion. In *Proc. 1EEE/RSJ Intl. Workshop Intelligent Robots and Systems,* pages 634-638, Tsukuba, Japan, September 1989.

[Bicchi et al, 1991] A. Bicchi, J. Salisbury, and D. Brock. Experimental evaluation of friction characteristics with an articulated robotic hand. In R. Chatila and G. Hirzinger, editors, *Experimental Robotics II,* pages 153-167. Springer-Verlag, Toulouse, France, 1991. Lecture Notes in Control and Information Sciences 190.

[Caillas, 1990] C. Caillas. Thermal Imaging for Autonomous Vehicle in Outdoor Scenes. In *Proc. IEEE Intl. Workshop on Intelligent Robots and Systems,* pages 651-658, Tsuchiura, Japan, July 1990.

[Campos *et al,* 1991] M. Campos, R. Bajcsy, and V. Kumar. Exploratory procedures for material properties: The temperature perception. In *Proc. Intl. Conf. Advanced Robotics,* pages 205-210, Pisa, Italy, June 1991.

[Durst and Krotkov, 1993] R. Durst and E. Krotkov. Object Classification from Acoustic Analysis of Impact. Technical Report CMU-RI-TR-93-16, Robotics Institute, Carnegie Mellon University, June 1993.

[Gandolfo et al, 1991] F. Gandolfo, G. Sandini, and M. Tistarelli. Towards vision guided manipulation. In *Proc. Intl. Conf. Advanced Robotics,* pages 661-667, Pisa, Italy, June 1991.

[Gemant and Jackson, 1937] A. Gemant and W. Jackson. The measurement of internal friction in some solid materials. *Philosophical Magazine,* 157:960-983,1937.

[Krotkov, 1990] E. Krotkov. Active Perception for Legged Locomotion: Every Step is an Experiment. In *Proc. IEEE Intl. Symposium on Intelligent Control,* pages 227-232, Philadelphia, Pennsylvania, September 1990.

[Lederman and Klatzky, 1987] S. Lederman and R. Klatzky. Hand movements: A window into haptic object recognition. *Cognitive Psychology,* 19:342-368,1987.

[Pentland and Williams, 1989] A. Pentland and J. Williams. Perception of non-rigid motion: Inference of shape, material, and force. In *Proc. Intl. Joint Conf. Artificial Intelligence,* pages 1565-1570, August 1989.

[Sinha *et al,* 1990] P. Sinha, R. Bajcsy, and R. Paul. How does a robot know where to step? Measuring the hardness and roughness of surfaces. In *Proc. IEEE/RSJ Intl. Workshop Intelligent Robots and Systems,* pages 121-126, Tsuchiura, Japan, July 1990.

[Stansfield, 1991] S. Stansfield. A haptic system for a multifingered hand. In *Proc. IEEE Intl. Conf. Robotics and Automation,* pages 658-664, Sacramento, California, April 1991.

[Stark, 1994] Louise Stark. Recognizing object function through reasoning about 3-d shape and dynamic physical properties. In *Proc. IEEE Computer Vision and Pattern Recognition,* pages 546-553, Seattle, June 1994.

[Wildes and Richards, 1988] R. Wildes and W Richards. Recovering material properties from sound. In W. Richards, editor, *Natural Computation,* pages 357-363. MIT Press, Cambridge, Massachusetts, 1988.

[Wolff and Boult, 1991] L. Wolff and T. Boult. Constraining object features using a polarization reflectance model. *IEEE Trans. Pattern Analysis and Machine Intelligence,* 13(7):635-657,1991.