Object Recognition by a 'Smart' Tactile Sensor.

R. Andrew Russell
Intelligent Robotics Research Centre
Department of Electrical and Computer Systems Engineering
Monash University, Clayton, VIC 3800
AUSTRALIA

Abstract

This paper describes a new approach to object recognition based on the time evolution of tactile images produced by an array tactile sensor. Tactile sensation from the fingertip skin conveys a lot of information about the shape of a touched object. It is usually considered that this should be combined with proprioceptive information to give a haptic interpretation of the data as the result of an active search. The method described in this paper determines the position, orientation and identity of objects from a sequence of tactile images gathered over time. External information such as finger joint angles and finger position are not used. It is envisaged that this algorithm could be incorporated into a self-contained 'smart' tactile sensor. Such a sensor would provide information for a robotic manipulation system including object identity. This information would not require significant additional processing in order to be useful. Implementation details are described in this paper together with experimental results.

1 Introduction

The human hand is remarkably good at grasping things, manipulating them, determining many of their physical properties and recognising them. It is considered that these abilities are largely due to haptic sensing [Gibson, 1962; Lederman and Browse, 1988]. Contact sensation provided by the skin is combined with information about finger positions and applied forces. Haptic sensing uses this sensory data to guide a purposeful tactile exploration gathering information about many attributes of the object including surface shape, weight, articulation between parts of the object, material compliance, texture and thermal properties.

In early work on tactile sensing static images were analysed in an attempt to recognise touched objects [Ghani and Rzepczynski, 1987; Briot, 1979]. This approach could only provide useful results for groups of objects that each produced unique patterns of contact. More recently the idea of haptic exploration has been taken as a model for the recognition of objects by touch. In an early example of haptic recognition Hillis [1981] employed a robotic finger to measure the stability of small objects by trying to roll them. This stability information was used to help recognise members of a group of nuts, bolts and other fastening devices. A good cross-section of more recent work in this area is referenced in the paper by Okamura, Turner and Cutkosky [Okamura et al., 1997].

In the project described in this paper objects are recognised using tactile information but without employing haptic exploration. It is assumed that the tactile sensor is located on a robotic finger or in some other location where objects are contacted. However, no control is exerted over the manipulation process and in fact the only information used is that provided by the tactile sensor itself. Objects are recognised based on the time evolution of the tactile sensor output. It is intended that this scheme could be integrated into a self-contained intelligent sensor which as well as providing raw tactile images and statistics derived from those images could also recognise the touched objects. Such a sensor would be relatively easy to integrate into a robotic system, especially as it does not require any external sensory information.

2 The Experimental Equipment

For this project a Bonneville Scientific force sensing array model No. 300-16x16B was used for gathering tactile images. The array contains 256 sensor elements arranged in 16 rows of 16 columns. Overall size of the array is 1.11 inch by 1.11 inch (2.8 cm by 2.8 cm). The sensor measures the thickness of a rubber sheet and each taxel consists of a time-of-flight ultrasonic thickness gauge [Grahn and Astle, 1984]. Being based on measurement of time the sensor is very stable and can resolve deflections in the rubber sheet of 6 microns. The maximum scanning rate for the sensor is 240 scans/sec.

To demonstrate the feasibility of this form of object recognition six objects were selected with shapes

that could be easily distinguished. The objects were a cube, sphere, cone, truncated cone, cylinder and tetrahedron. A photograph of the tactile sensor and test objects is shown in Figure 1.

3 Image Processing

In order to simplify the interpretation of the tactile images a number of restrictions were placed on the objects contacting the tactile sensor. In later stages of the project these restrictions will be relaxed. Currently it is assumed that only one object at a time will be touching the sensor. This removes the necessity for segmenting the tactile image and keeping track of multiple objects. When an object moves while in contact with the sensor only rolling is permitted, sliding is not allowed. It is envisaged that at a later stage slip sensing could be incorporated into the sensor and that this information would permit the restriction on sliding to be removed.

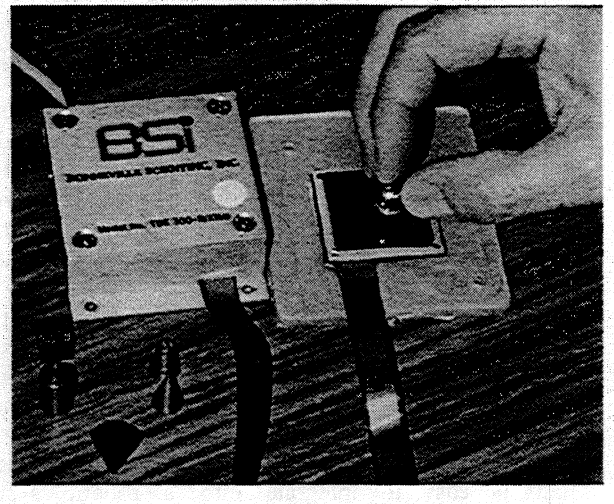

Figure 1: The Bonneville Scientific tactile sensor and six test objects.

Processing of the image follows a 'bottom up' scheme. Images captured by the array tactile sensor are passed through a sequence of processing levels which successively produce more abstract representations of the data. A key feature is the use of the time evolution of the tactile image. Figure 2 shows a schematic representation of the four processing levels and the information flow between them.

Level 0 – Image Capture

At the lowest level a sequence of tactile images is captured by sampling the sensor at regular intervals. During program development the sample rate was set at one per second so that the processing scheme could be monitored in real-time. In a practical application the sampling rate could be 10 to 50 times faster.

Level 1 – Generate Image Statistics

Each image is analysed to generate statistics which are used at higher levels in the analysis. The statistics are based on the moments of area that are commonly used in binary vision and other schemes for analysing tactile images [Winston and Horn, 1984]. The output of the tactile sensor is converted to a binary image using a fixed threshold level where $\tau_{ij} = 0$ if a taxel does not make contact with the touched object and $\tau_{ij} = 1$ if there is contact.

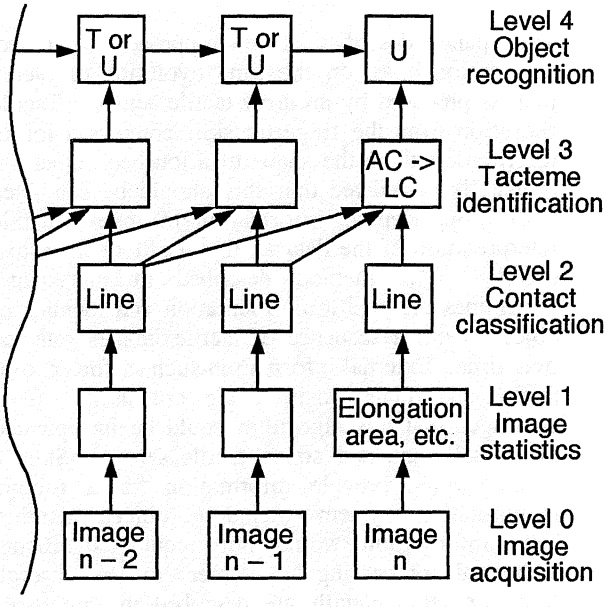

Figure 2: Information flow in the smart sensor.

The general moment of area m_{pq} is calculated as follows:

$$m_{pq} = \sum_{i=0}^{n-1} \sum_{j=0}^{m-1} i^p j^q \tau_{ij} \qquad (1)$$

Area of contact is the moment of area m_{00}. Using other moments of area the coordinates of the centre of area of contact i_0 and j_0 can be found together with the axis of minimum moment of inertia θ_0 (which indicates orientation) and *Elongation* (which indicates the aspect-ratio of the area of contact).

$$i_0 = \frac{m_{10}}{m_{00}} \qquad (2)$$

$$j_0 = \frac{m_{01}}{m_{00}} \qquad (3)$$

$$\theta_0 = \frac{\tan^{-1}}{2}\left[\frac{2(m_{00}.m_{11} - m_{10}.m_{01})}{(m_{00}.m_{20} - m_{10}^2) - (m_{00}.m_{02} - m_{01}^2)}\right] \quad (4)$$

$$a = m_{20} - \frac{m_{10}^2}{m_{00}^2} \quad (5)$$

$$b = m_{11} - \frac{(m_{10}.m_{01})}{m_{00}^2} \quad (6)$$

$$c = m_{02} - \frac{m_{01}^2}{m_{00}^2} \quad (7)$$

$$Elongation = \frac{\sqrt{b^2 + (a-c)^2}}{m_{00}^2} \quad (8)$$

From these statistics the area, centre of area, orientation and elongation are passed to the higher levels of the processing scheme.

Level 2 – Classification of the Area of Contact into a Contact Type

With the aid of the statistics calculated in Level 1 each image is classified according to contact type. The allowable contact types are none, indeterminate, point, line and area. Classification of the point, line and area contacts is performed by a three layer neural network.

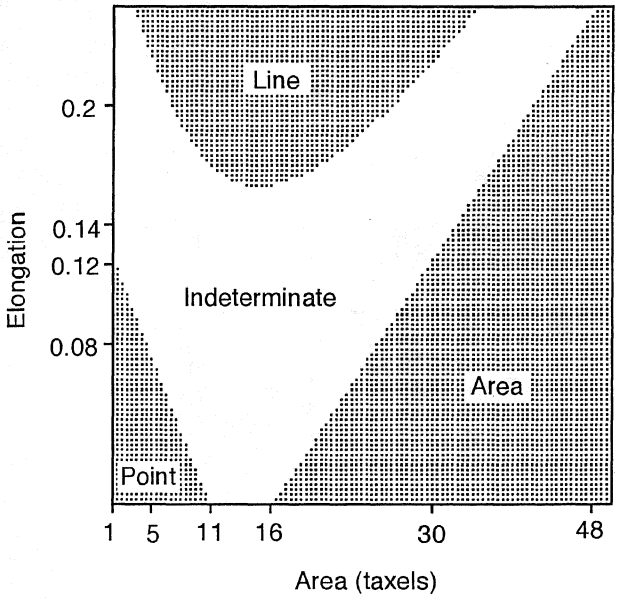

Figure 3: Tactile image classification.

When the area of contact m_{00} is zero there is obviously no contact and all other images that are not point, line or area contacts are classified as indeterminate. Figure 3 shows the way the two dimensional space of area and elongation is partitioned by the neural network to represent the different contact types. For this simple case the weights on the neural network inputs were calculated by hand using experimentally determined information about the values of area and elongation corresponding to each contact type [Trembley and Cutkosky, 1993]. However, having formulated this part of the recognition algorithm as a neural network an increased number of inputs or more complicated decision regions can be accommodated by training the neural network using back-propagation.

Once the contact type is determined by the neural network, that is as far as the processing of individual tactile images is taken. The higher processing levels consider the time evolution of the images.

Level 3 – The generation of 'tactemes'

By analogy with the phoneme I have called significant tactile images or transitions between images 'tactemes'. As an object rolls over a tactile sensor the contact type will change producing a sequence of tactemes If the contact type of an image is classified as indeterminate then it is assumed that this image is part of the current tacteme or put another way an indeterminate image does not indicate the start of a new tacteme.

A region of indeterminate classification separates the valid classifications in Figure 3. This helps to reject the spurious images that tend to occur as contact between sensor and object changes. However, some of these spurious images will be classified as valid points, lines or areas. These images do not persist for very long and are rejected by not accepting a new classification as point, line or area until it has been detected a certain number of times. In current experiments with a sampling rate of one second a new image type has to be repeated three times before it is accepted.

Level 4 – Object Recognition

For any particular object there will be some tactemes that cannot be generated. For instance, a sphere is unable to make a line contact and a cube cannot produce a point contact that moves over the tactile sensor surface (because sliding is not allowed). Thus detection of most tactemes allows us to narrow down the possible identity of the current object touching the sensor. Table 1 lists all significant tactemes that are identified by the object recognition program. The objects that are eliminated by each tacteme are indicated by a '✘'.

Before first contact is made it is assumed that the next object to touch the sensor could be any of the six objects. Successive tactemes are used to eliminate possible interpretations until only one remains. In that

case the remaining interpretation is assumed to identify the object. If the object is removed from the sensor before the object is recognised a partial classification is available listing all interpretations that have not been eliminated. It is possible to eliminate all possible interpretations if the object is allowed to slide over the sensor or if the shape does not correspond to one of the six known shapes. In these cases the object will be classified as 'unknown'.

Tacteme	S	Y	U	O	R	T
point contact (PC)						
line contact (LC)	✘					
area contact (AC)	✘					
PC moves			✘			✘
PC takes an orthogonal turn		✘	✘	✘	✘	✘
LC moves (LCM)	✘		✘			✘
AC moves	✘	✘	✘	✘	✘	✘
AC rotates	✘	✘	✘	✘	✘	✘
LCM parallel (LCMP)	✘		✘	✘	✘	✘
LCM and turns (LCMT)	✘	✘	✘			✘
LCMT about one end	✘	✘	✘		✘	✘
LCMT in an arc	✘	✘	✘	✘		✘
LC -> AC	✘	✘		✘	✘	
AC -> LC	✘	✘		✘	✘	
LC -> AC -> LC (two LCs are parallel or 90°)	✘	✘		✘	✘	✘
LC -> AC -> LC (two LCs are at 60° or 120°)	✘	✘	✘	✘	✘	

S = sphere
Y = cylinder
U = cube
O = cone
R = truncated cone
T = tetrahedron

Table 1: Tacteme-object dependencies.

4 Results

Figure 1 shows the Bonneville force sensing array and six prototype objects. The box on the left of the photograph is the tactile sensor electronics module that performs sensor scanning and pulse amplification.

For three tests during which each of the six objects was moved in contact with the sensor the object was recognised successfully 83% of the time. It is expected that further refinement of the level 2 contact classification process will improve this figure. Recorded data from one of the tests is given in Figure 3. This figure shows a sequence of binary images taken at one second intervals. Beneath each image is the Level 2 area of contact classification for that image. Table 2 give information derived at levels 2, 3 and 4 that was used to identify the object contacting the sensor. Level 2 classified the first contact as an area. After this had been established for three successive images the first tacteme was identified as an area contact. Contact then changed from an area contact to a point contact. During the change a spurious line classification was made for image number 8. This was ignored and so were the point classifications until they had been stable for three successive images. At that time a point tacteme (PC) was identified. In image 17 the change to a line contact was established producing a line contact (LC) tacteme. Finally, the line contact was seen to move parallel to its initial orientation (LCMP). This tacteme provided the unique identification of the object as a cylinder.

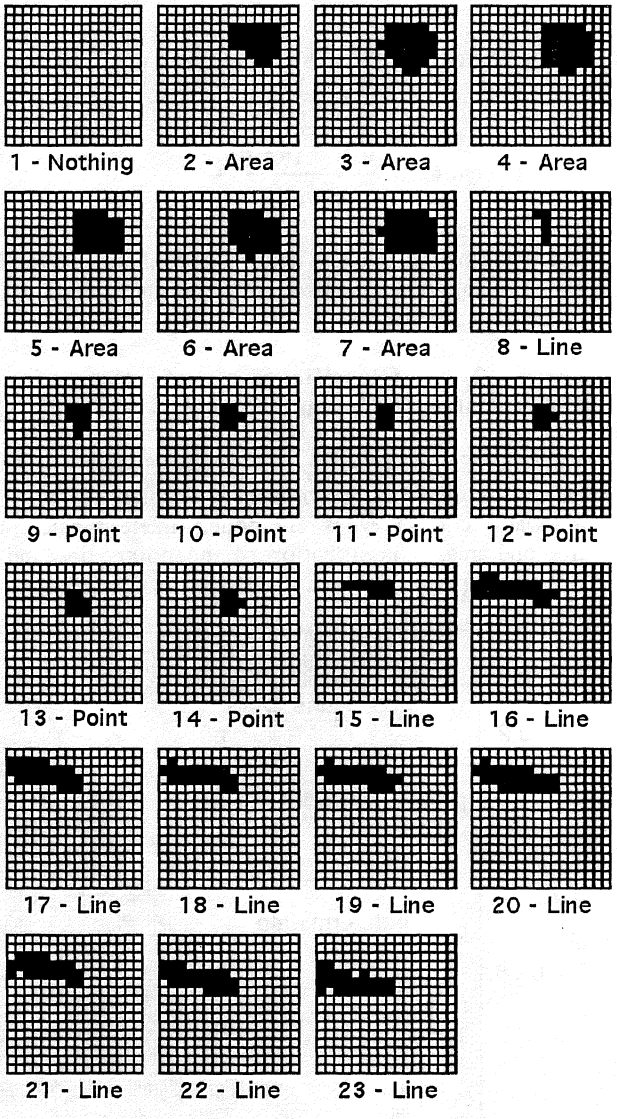

Figure 3: A sequence of tactile images leading to recognition of a cylinder.

The movements of the cylinder that produced the data recorded by the tactile sensor are shown in Figure 4.

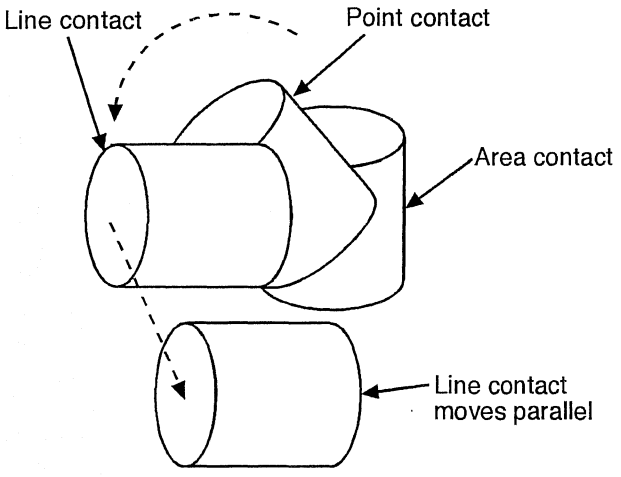

Figure 4: Movement of the cylinder on the tactile sensor.

Image	Classification	Tacteme	Remaining objects
1	nothing	–	S,Y,U,O,R,T
2	area	–	S,Y,U,O,R,T
3	area	–	S,Y,U,O,R,T
4	area	AC	Y,U,O,R,T
5	area	–	Y,U,O,R,T
6	area	–	Y,U,O,R,T
7	area	–	Y,U,O,R,T
8	line	–	Y,U,O,R,T
9	point	–	Y,U,O,R,T
10	point	–	Y,U,O,R,T
11	point	PC	Y,U,O,R,T
12	point	–	Y,U,O,R,T
13	point	–	Y,U,O,R,T
14	point	–	Y,U,O,R,T
15	line	–	Y,U,O,R,T
16	line	–	Y,U,O,R,T
17	line	LC	Y,U,O,R,T
18	line	–	Y,U,O,R,T
19	line	–	Y,U,O,R,T
20	line	–	Y,U,O,R,T
21	line	–	Y,U,O,R,T
22	line	–	Y,U,O,R,T
23	line	LCMP	Y

Table 2: The object recognition process using the images shown in Figure 3 (notation taken from Table 1)

5 Conclusions

An object recognition scheme has been proposed which uses information gathered from the time evolution of tactile images. Proprioceptive information required for haptic recognition, such as finger joint angles, is not used. Because this processing scheme only requires tactile images it opens up the possibility of constructing a self-contained smart tactile sensor which, as well as providing raw tactile images and image statistics, can also identify touched objects. The recognition scheme relies upon identification of object features rather than matching with stored templates. For this reason the dimensions of an object can vary and its shape will still be recognised.

Currently the prototype system has a number of limitations. However, further development should make the tactile image processing scheme more generally applicable. At the moment there is a restriction which only allows one object to touch the sensor at any particular time. By implementing standard image segmentation techniques this will allow multiple objects to be tracked and identified. Another restriction is that objects in contact with the tactile sensor must not slide. It may be possible to relax this restriction if the presence of slip can be detected. Recent work has demonstrated sensors that can measure slip [Trembley and Cutkosky, 1993] and it may be possible to combine these techniques with the current Bonneville sensor. It is also expected that by considering additional information derived from the time evolution of tactile images a wider range of shapes can be recognised.

The identity of touch objects can be hidden by accidentally or deliberately making sure that the tactemes required for recognition do not occur. In these cases the system will still be able to provide a list of currently possible interpretations.

Acknowledgments

This project has been supported by a grant from the Harold Armstrong Fund. I would like to thank Paul Beeley for writing the code which implements these ideas.

References

[Briot, 1979] Briot, M., 'The utilization of an 'artificial skin' sensor for the identification of solid objects', *Proceedings of the 9th International Symposium on Industrial Robots*, Washington, March 13-15, 1979, pp.529-547.

[Ghani and Rzepezynski, 1987] Ghani, N., and Rzepczynski, Z.G., 'A tactile sensing system for robotics', *in Intelligent Autonomous Systems*, L.O. Hertberger and F.C.A. Groen (eds.), North-Holland, Amsterdam, 1987, pp. 241-245.

[Hillis, 1981] Hillis, W.D., *A.I. Memo 629: Active Touch Sensing*, MIT AI Laboratory, April 1981.

[Gibson, 1962] Gibson, J.J., 'Observations on active touch', *Psychological Review*, Vol. 69, No. 6, November, 1962, pp. 477-491.

[Grahn and Astle, 1984] Grahn A. R., and Astle, L., 'Robotic Ultrasonic Force Sensor Arrays', *Proceedings of the Robots 8 Conference*, Detroit, Michigan, 4-7 June 1984, pp. 21.1-21.18.

[Johnson and Picton, 1995] Johnson, J. and Picton, P., *Mechatronics: Designing Intelligent Machines, Vol. 2 Concepts in Artificial Intelligence*, Butterworth-Heinemann, 1995.

[Lederman and Browse, 1988] Lederman, S.J. and Browse, R.A., 'The physiology and psychophysics of touch', in P. Dario (ed.), *Sensors and Sensory Systems for Advanced Robots,*, Springer-Verlag, 1988, pp. 71-91.

[Okamura et al., 1997] Okamura, A.M., Turner, M.L. and Cutkosky, M.R., 'Haptic exploration of objects with rolling and sliding', *Proceedings of the IEEE International Conference on Robotics and Automation'*, Albuquerque, April 1997, pp. 2485-2490.

[Trembley and Cutkosky, 1993] Tremblay, M.R. and Cutkosky, M.R., 'Estimating friction using incipient slip sensing during manipulation task', *Proceedings of the IEEE International Conference on Robotics and Automation*, Atlanta, May 2-6 1993, pp. 429-434.

[Winston and Horn, 1984] Winston, P. H., and Horn, B. K. P., *Lisp*, (2nd Ed.), Addison-Wesley, Reading, Massachusetts, 1984.