

## TECHNICAL NOTE

# Bank note recognition for the vision impaired

A. Hinwood<sup>1</sup>, P. Preston<sup>1</sup>, G. J. Suaning<sup>3,1</sup> and N. H. Lovell<sup>1,2</sup>

<sup>1</sup>*Graduate School of Biomedical Engineering, University of New South Wales, Sydney, Australia*

<sup>2</sup>*National Information and Communications Technology Australia (NICTA) Australian Technology Park, Eveleigh Australia*

<sup>3</sup>*School of Engineering, University of Newcastle, Newcastle, Australia*

### Abstract

Blind Australians find great difficulty in recognising bank notes. Each note has the same feel, with no Braille markings, irregular edges or other tangible features. In Australia, there is only one device available that can assist blind people recognise their notes. Internationally, there are devices available; however they are expensive, complex and have not been developed to cater for Australian currency. This paper discusses a new device, the MoneyTalker that takes advantage of the largely different colours and patterns on each Australian bank note and recognises the notes electronically, using the reflection and transmission properties of light. Different coloured lights are transmitted through the inserted note and the corresponding sensors detect distinct ranges of values depending on the colour of the note. Various classification algorithms were studied and the final algorithm was chosen based on accuracy and speed of recognition. The MoneyTalker has shown an accuracy of more than 99%. A blind subject has tested the device and believes that it is usable, compact and affordable. Based on the devices that are available currently in Australia, the MoneyTalker is an effective alternative in terms of accuracy and usability.

**Key words** vision impairment, technical aid, disabled, classification

### Introduction

The recognition of bank notes presents a huge difficulty for someone who is blind or vision impaired. In the case of Australian bank notes, each note has the same feel, with no Braille markings, irregular edges, distinguishable raised print or other tangible features. The width of each note is the same and there is only a 7 mm difference in length between consecutive note values<sup>1</sup>.

Studies conducted over the last 10 years have found that almost 500,000 Australians are now vision impaired, including 50,000 who are legally blind. This number is expected to double in the next 20 years<sup>2</sup>.

In Australia, there is presently only one device available that can assist blind people to recognise their notes, the CashTest<sup>1</sup>. This device determines the value by a mechanical means relying on the different lengths of each note. The downfall of the CashTest is that it does not allow for the shrinking of notes with time nor the creases or rips

that are common in our notes. A user of the device has reported that the CashTest is inaccurate and difficult to use, although cheap and extremely portable<sup>3</sup>.

Internationally, there are electronic devices available, including the BlindNoteReader and NoteTeller 2 (Brytech Inc., Canada). These are electronic devices developed in Canada that work by raster scanning the face of the bank note with a charge-coupled device (CCD) to obtain an image of the note<sup>4</sup>. The image is compared with stored images in the device memory and if the comparison correlates to a predetermined degree, the value is output. Bank notes can be inserted and read in any orientation. Although these devices are accurate, reliable and easy to operate, they are expensive, complex and have not been developed to cater for Australian currency at this stage.

The MoneyTalker is a new device that recognises notes and announces their value using electronic speech, either through a speaker or through headphones. The aim of this study is to determine the best classification methods for note recognition, and the best hardware configuration to obtain the highest accuracy and efficiency possible. It relies on the largely different colours used on each Australian note, and can recognise all four orientations of an inserted note. This means that it can take a ripped or torn note, provided that the end of the note that is inserted does not contain the rip or tear. It can also recognise faded notes, due to the variety of note conditions that it was trained on. It has been designed to be easy to use, accurate and cheap to manufacture. Another important feature of its design is that it is easily adaptable to recognise new notes, utilising machine learning techniques.

---

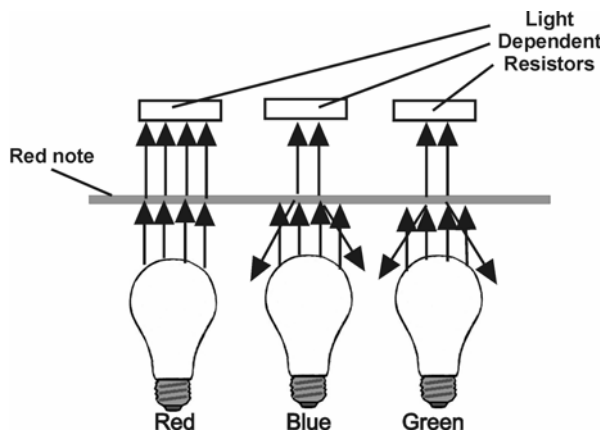
*Corresponding author: N. H. Lovell, Graduate School of Biomedical Engineering, University of New South Wales, Sydney, NSW 2052 Australia*

*Tel: +61 (2) 9385 3922, Fax: +61 (2) 9663 2108*

*Email: N.Lovell@unsw.edu.au*

*Received: 17 June 2005; Accepted: 20 April 2006*

*Copyright © 2006 ACPSEM/EA*



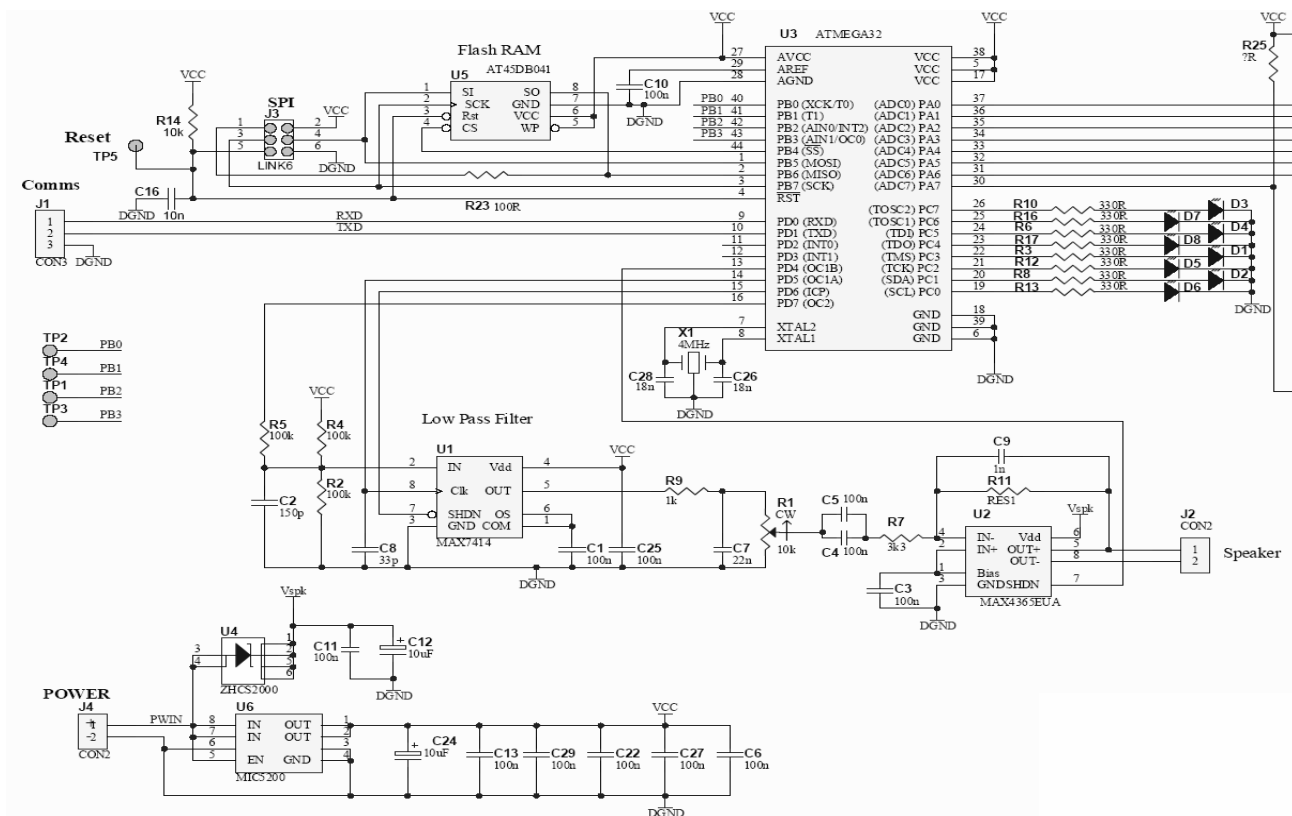
**Figure 1.** The transmission of red, green and blue light through a red note. In the final implementation a total of eight LEDs and eight light-dependent resistors were used.

**Technical description**

The main principle behind the MoneyTalker's recognition technique is the transmission/reflection property of light. Coloured lights are transmitted through the note and as each note is a different colour, the amount of light that is transmitted through the note is different and dependent on the colour of the note. Sensors positioned beneath the coloured lights return values that are dependent on a note's colouring and light transmission characteristics (Figure 1).

Once the sensor readings have been performed, they are compared to values that are stored in the memory for pre-classified examples. If the values in memory correlate to the values read, the appropriate note value is announced using a simple pulse-code digital-to-audio system. The note value is determined using the best possible classification algorithm for the data set.

The system is implemented using a small 8-bit microcontroller that features digital I/O, an 8-channel analogue-to-digital converter (ADC) and a pulse-width modulation (PWM) unit. The Atmel AVR ATmega32 (Atmel Corporation, San Jose, USA) is an 8-bit microcontroller than can run at up to 16 MHz, although in the prototype, its speed was limited to 11 MHz to reduce power consumption. The digital I/O drives high intensity LEDs, the ADC reads the relevant light levels (as transmitted through the note) from a number of light-dependent resistors. The PWM unit is used to create the pulse-code audio output through a filter and amplifier system. Additional to the microcontroller was an Atmel AT45DB041B DataFlash (Atmel Corporation, San Jose, USA) memory component that provided the storage for each spoken sound. This DataFlash memory also stored the data used for note comparison. By utilising the features available in modern microcontrollers, the principal electronics are implemented in two small integrated circuits, maintaining small size, low cost and low power consumption. A schematic of the overall circuit is given in Figure 2.



**Figure 2.** Schematic of the prototype circuit used for controlling the LEDs, sampling the light output and generating the audio output. The system is based around an Atmel AVR ATmega32 microcontroller and an Atmel AT45DB041B DataFlash memory integrated circuit.

The electronics for the prototype were housed in a small, light-tight box to reduce the effects of external light illumination on the recognition process (Figure 3).



**Figure 3.** The prototype Money Talker in use. The system is showed connected via a serial line to a PC for the purposes of testing.

## Methods

### Classification methods

In an initial prototype with only three sensors and three coloured lights, the classification of notes was performed by storing the ranges of values for each sensor, for each note, for each direction. It was found that three sensors were insufficient to reliably differentiate notes in different orientations. The unit (as described herein) uses eight sensors and LEDs.

Given the considerable number of combinations, a simple range comparison of each sensor value for each note in each orientation was considered inappropriate. Hence, classification algorithms were considered. Using these approaches, the actual colour of the LEDs used and their location were found not to be critical. In the prototype, three red, three green, one blue and one white LED were used.

There were many methods considered in the classification of an inserted note. There is a variety of data mining algorithms that exist including rule-based algorithms, classification trees, neural networks, instance-based algorithms and function-based algorithms. As there were such a large number of different techniques, methods were tested using Weka, a program written in Java for solving data mining problems<sup>5</sup>. Weka offered an easy environment for evaluation and testing of a large range of machine learning algorithms without having to hand code each individual methodology.

There are 6 different note categories (\$5, \$5 federation, \$10, \$20, \$50, \$100) in Australia. An analysis sample set consisted of 20 samples of each note category (excluding the \$100 which only had 10 samples). The notes selected were in various conditions of newness. The dataset was

formed by inserting the notes into the device 3 times in each orientation, giving a dataset of 1320 samples. The samples were inserted 3 times as small changes in insertion position can produce slightly different results. Samples were also taken at low, medium and high ambient light levels. The dataset was acquired from the device via an RS232 connection, allowing the device to communicate with a computer (Figure 3). A simple program was written in Java and run on the desktop computer, opening the serial port and then allowing the readings to be automatically written to a file in an appropriate format for Weka to use.

Weka ran each classification algorithms five times, and the average number of correctly and incorrectly classified samples was recorded. The tests were run with the dataset randomly split, using 66% of the dataset as a training set and 33% of the dataset as a test set.

### Testing methods

The dataset created by Weka was transformed into a training set for the device. The training set was stored in the DataFlash memory component because 1320 samples could not be stored in the small flash memory of the Atmel microcontroller.

Another set of currency with 40 representatives from each note category was used to test the system. Each note sample was inserted into the device in all four orientations and the result was recorded in terms of accuracy and timing.

## Results

### Classification results

This stage was designed to select the best algorithm to finally implement in the device. Essentially a sample of notes was presented to Weka, and the accuracy of the various algorithms was recorded.

All of the algorithms that achieved over 90% accuracy are shown in Table 1.

The most accurate algorithms for this dataset are IB1, Multi-Layer Perceptron, K\*, DECORATE, Random Committee, LMT and IBk.

**Table 1.** Accuracy of various classification algorithms using the Weka system for the training and testing dataset of 1320 notes.

Algorithm type	Algorithm	Correct (%)	Standard deviation (%)
LAZY	IBk	99.68	0.11
	K*	99.92	0.34
	IB1	99.92	0.11
TREE	LMT	99.47	0.23
FUNCTION	Multi-layer perceptron	99.8	0.2
	DECORATE	99.72	0.23
META	Random committee	99.68	0.27

**Table 2.** Results of Money Talker using an embedded IB1 algorithm based on a separate dataset of 960 notes.

Note type	Number correct	Number incorrect
\$5 federation	160	0
\$5	160	0
\$10	160	0
\$20	160	0
\$50	157	3
\$100	157	3
Total	954	6

**Table 3.** The breakdown of notes in issue in Australia at the end of June 2004<sup>11</sup>.

Notes on issue	Number	
	Millions	%
\$5	107	13.8
\$10	79	10.2
\$20	127	16.4
\$50	319	41.2
\$100	142	18.4
Total notes	774	100.0

IB1 is an instance-based classifier that finds the single training instance closest to the given test instance using the Euclidean distance metric for the N-dimensional space of the data<sup>5,6</sup>. If multiple instances have the same (smallest) distance to the test instance, the first one found is used. The time taken to classify an instance increases linearly with the number of training instances. Consequently it is necessary to restrict the number of training instances that are kept in the classifier.

IBk is an instance-based classification algorithm that is an implementation of the  $k$  nearest neighbours classifier<sup>5,6</sup>. This algorithm is similar to IB1 however, instead of only finding the single closest test instance, it finds the  $k$  closest test instances and the prediction becomes the most common class in the neighbours list. Again, the time taken to classify a test instance with a nearest neighbour classifier increases linearly with the number of training instances.

K\* is an instance-based learner that uses an distance measure based on entropy<sup>7</sup>. This algorithm defines the distance between instances as the complexity of transforming one instance to another. The mathematics required to perform this algorithm are complicated and presents a large difficulty for a low-cost microcontroller to perform.

A Multi-Layer Perceptron is a neural network algorithm that uses back propagation to classify instances<sup>5,8</sup>. The network can be monitored and modified

during training time. This algorithm requires a large amount of free memory to build the classifier and is not considered feasible for the MoneyTalker hardware.

Logistic Model Trees (LMTs) are classification trees with logistical regression functions at the leaves<sup>9</sup>. Building this type of classifier takes a large amount of time, approximately 16 minutes on a computer with a 1.2GHz CPU. Given that the microcontroller used is approximately 100 times slower, this algorithm is also not feasible in the current device environment.

DECORATE (Diverse Ensemble Creation by Oppositional Relabeling of Artificial Training Examples) is an ensemble meta-learner that directly constructs diverse committees by employing specially-constructed artificial training examples<sup>10</sup>. As this algorithm is a meta-learner, building the classifier is more difficult as it relies on at least two different algorithms, requiring more free memory than a single classifier, and taking more time to classify an instance.

The decision was made to implement IB1 in the device, since it offered the best balance of size, computation speed and accuracy.

### Testing results

The IB1 algorithm was embedded on the microcontroller. An additional set of notes was used to test the system in its final configuration. In this test phase, 40 notes of each denomination were inserted in all four orientations, leading to 160 tests for each denomination. Table 2 documents the system performance in terms of note denomination.

When considering the data set as a whole, the total accuracy of this device is 99.4%. The results were then further analysed in terms of the numbers of note types in circulation. Table 3 illustrates the number of notes in issue in Australia as of June 2004. As there is a variation in the number of notes produced for each denomination, the likelihood is that similar proportions will be used with this device. Using this data, the accuracy, weighted by the percentage breakdown of total notes in issue is 98.9%.

The time taken for each note to be recognised was approximately 3.01s. This time was the time taken from note insertion to the value announced. This time could be further reduced using a faster crystal or a smaller training set, although reducing the training set would decrease the accuracy.

### Discussion

This paper discusses the value and accuracy of the MoneyTalker device, a new device that assists the blind to recognise Australian currency notes. The only similar device available for Australians is a device that is mechanical in nature and becomes unreliable when notes are old and have started to shrink. This is not a problem for the MoneyTalker as it is not based on the size of the notes.

A device of this type must be accurate, easy to use, inexpensive, quick and portable. Sample testing was

performed to determine the feasibility of this device in terms of its accuracy and speed. The non-measurable requirements were tested by a person who was blind.

The results showed that the device had an accuracy of 99.4%, with the accuracy falling slightly on the \$50 and \$100 notes. One reason for the drop in accuracy for the \$100 notes is the smaller training set used. When the accuracy is weighted according to the number of notes on issue, the figure drops to 98.9%.

A shortcoming of the testing that was performed on this device was the limited ability to obtain training and test data. This is due to the large amount of money required to perform these tests adequately. The minimum amount of samples required to obtain feasible statistical results is approximately 40. The maximum available amount of notes available for development of this device was 60 representatives of each of the six note categories (except for the \$100 note which only had 50). These were split into 20 notes for training and 40 notes for testing.

The cost of construction of the device would be approximately \$84 (this assumes around 50 are made at the same time). The cost would decrease to around \$40 if 1000 are made. This compares favourably to the electronic devices overseas, being sold for approximately \$325 each. This device could be sold for one third of the price, delivering a huge saving to blind Australians.

Since this project started, a number of new microcontrollers that would suit this application have become commercially available. High on the list are the new ARM-based single chip microcontrollers (Philips' LPC2xxx family and Texas Instrument's TMS470), Atmel's AT91SAM7 and Analog Devices' AduC7000 series. While there are variations in functionality, speed and peripherals available they all offer similar specifications. The combination of 32-bit arithmetic and 60 MHz operation would lead to very fast computation.

These new microcontrollers also offer much larger memory – in the case of the Philips family, 512 kB of flash memory and 64 kB of RAM, eliminating the need for an extra memory chip. Using one of these microcontrollers would eliminate the need for almost all other chips in the device as many have inbuilt timer and PWM functions – leading to a final main circuit board of just a CPU chip and a voltage regulator. A disadvantage would be a somewhat higher power consumption but this would be countered by a much faster processing speed and the loss of the extra memory chip as needed for the current solution. A second disadvantage could be the slightly higher cost (\$10 compared to \$5).

In the current prototype, battery life has not been optimised. An approximate analysis of power consumption would indicate that with the Atmel microcontroller consuming 30 mA during computation and 50  $\mu$ A when idling and with the LEDs being driven at 5 mA, then a 2500 mAh battery pack would allow 40,000 notes to be read. A known issue with the current design is that when the

batteries are near the end of their useful life, the accuracy decreases because the light sources become dimmer. This could be easily rectified in subsequent designs with additional battery monitoring technology to audibly report battery status. More extensive field trialling of the MoneyTalker is the subject of on-going study. This will be necessary to systematically evaluate such issues as how fading colours in older notes impacts accuracy.

The MoneyTalker has not been thoroughly tested or field trialled. However, a number of vision impaired people were consulted in the design phases as well as stakeholder groups. With the current prototype one blind person has tested and examined the device and has given his approval. He feels that it is easy to use, is small and compact and would be a great device to own<sup>3</sup>. He believes that the main factors that would limit a person buying a device like this would be the price and the accuracy.

The MoneyTalker has proven to have high accuracy (>98%), is inexpensive and is easy for a blind person to use. The speed of recognition is adequate, although the speed could be improved with the use of superior hardware or faster clock speeds at the expense of cost and battery life respectively.

## References

1. Reserve Bank of Australia. *How the RBA assists people with a vision impairment to differentiate notes.* [http://www.rba.gov.au/CurrencyNotes/vision\\_impaired.html](http://www.rba.gov.au/CurrencyNotes/vision_impaired.html) accessed 7/6/2005.
2. Congdon, N., O'Colmain, B., Klaver, C.C., Klein, R., Munoz, B., Friedman, D.S., Kempen, J., Taylor, H.R. and Mitchell, P., *Causes and prevalence of visual impairment among adults in the United States.* Archives of Ophthalmology. 122(4):477-85, 2004.
3. ABC Radio National, *Money Talks* in M. Holstrom, ed., *The Buzz*, 2004.
4. Bryenton, E.L., Brule, D.A. and Bryenton, A.L., *Portable hand-held banknote reader*, US patent 5692068, 1997.
5. Witten, I.H. and Frank, E., *Data mining: practical machine learning tools and techniques with Java implementations*, Morgan Kaufmann, San Francisco, USA, 2000.
6. Aha, D.W., Kibler, D. and Albert, M.K., *Instance-based learning algorithms*, Machine Learning (Historical Archive), 6, pp. 37-66, 1997.
7. Cleary, J.G. and Trigg, L.E. K\*, *An instance-based learner using an entropic distance measure*, 12th International Conference on Machine learning, Morgan Kaufmann, Tahoe City, California, US, pp 108-114, 1995.
8. Russell, S.J. and Norvig, P., *Artificial intelligence: a modern approach*, Prentice Hall, Upper Saddle River, USA, 2003.
9. Landwehr, N., Hall, M. and Frank, E., *Logistic model trees*, Proc Fourteenth European Conference on Machine Learning, LNCS 2837, Cavtat-Dubrovnik, Croatia, pp 241-252, 2003.
10. Melville, P. and Mooney, R.J., *Constructing diverse classifier ensembles using artificial training examples*, IJCAI, Acapulco, Mexico, pp 505-510, 2003.
11. Royal Australian Mint. *Mint issue*, vol 59, 2004.