BMVA News

The Newsletter of the British Machine Vision Association and Society for Pattern Recognition

Volume 9 Number 2
November 1998

Editor: Dr Paul Rosin
Department of Information Systems & Computing
Brunel University
Middlesex UB8 3PH
Tel: (01895) 203392
Fax: (01895) 251686
Email: Paul.Rosin@brunel.ac.uk

BMVA News[1] is published every three months. Contributions on any activity related to machine vision or pattern recognition are eagerly sought. These could include reports on technical activities such as conferences, workshops or other meetings. Items of timely or topical interest are also particularly welcome; these might include details of funding initiatives, programmatic reports from ongoing projects and standards activities. Items for the next edition should reach the editor by 18th January 1999.

Contents

Editorial

We are constantly being told that we are in the Information Age. Everyone seems obsessed by generating, manipulating, cataloguing, and storing gargantuan amounts of data. So it's not surprising that a favourite pastime is making up lists of useful and, more often, useless facts. Here's my contribution (with help from various colleagues) to the information glut:

A List of Superlative Computer Vision Systems

most unpopular – car number plate recognition
most unpleasant – robotic animal carcass processing
most boring – bin picking
most sophisticated – sorry, couldn't understand it
most secret – XXXXX
most expensive – CMU's 51 camera VR system
most frightening – all AGV systems
most politically (in)correct – naked flesh detection
most unlikely – various contenders here:
    robot sheep shearer
    robot sheep dog
    fish weighing
cutest – Lena (OK, not a vision system, but ...)
most useless – again several possibilities:
    the 1966 World Cup Final analyser
    robot sheep dog (again)
most suspicious – intruder detection
most big–brotherish - fingerprint identification
most distant – Mars rover

Naturally these are always other people's systems. I've still get quite a few categories that I couldn't fill, for instance: most untested, slowest, least reliable, etc. I'm sure many of you out there must have stumbled across a few such systems to your cost. So if you can, throw discretion to the winds, and send in your favourites.

[1] The British Machine Vision Association and Society for Pattern Recognition is a Company limited by guarantee, No. 2543446, registered in England and Wales. Registered Office: Granta Lodge, 71 Graham Road, Malvern, WR14 2JS. The Association is a non-profit-making body and is registered as charity No. 1002307. See http://www.bmva.ac.uk/ for more information about the BMVA.

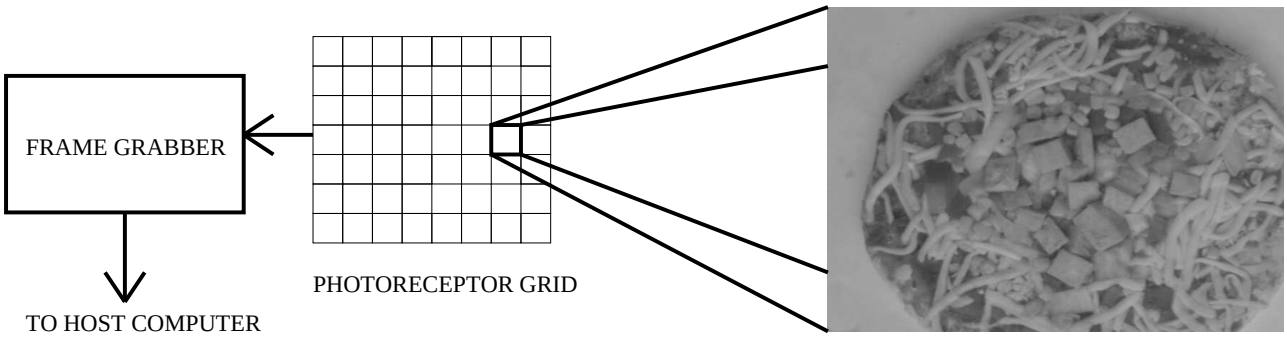

Figure 1: Figure 1: architecture of an imaging system based on a smart pizzel sensor, showing magnified pizzel. The camera and optics needed to form an image on the sensor are not shown. Our prototype MOS grid is 20×10 mm, comparable with the size of a CCD sensor[1].

Smart Pizzels

Introduction

We introduce a smart optical sensor based on specialised photoreceptors called *pizzels*. The sensor is optimised for detecting pizzas in all possible states, except already digested and Alabama. Pizzels are packed in a quasi-regular[1] grid, similarly to a CCD array[1]. This new sensor pushes computer vision[1] into an applicative sector spinning billions of dollars[1] and pizza bases worldwide each year. For reasons of space (and also of time, dinner is approaching fast) we cannot include our impressive mole of results testifying the high-quality performance of our prototype. We show, however, our impressive mole.

The system

The pizzel-based smart sensor is an imaging device implementing matched filtering, that is, pizzels respond to impinging light distributions matching their own internal structure. In our prototype we used MOS technology (Mozzarella, Oregano and Salami)[1,1], thus achieving maximum response for matched MOS pizzas, and good responses for other toppings selections. A grid of pizzels of several types, packed together at an appropriate resolution, will recognise most pizzas within a brief, uncompensated response period (BURP). Figure 1 sketches the architecture of a pizzel-based imaging system. ANOVA analysis[1] of data collected during field research[1] indicate that MOS pizzels are indeed matched to the three most frequently asked toppings. Experimental analysis shows that a MOS sensor suffices for most practical purposes, except guessing lottery numbers.

Experimental results

Synthetic tests. These were conducted with 20,896 synthetic images generated on a Tarragon Graphics running OS Ole MIO. Each image contained 1 to 5 pizzas of 10 possible types, selected at random from the 1998 menu of Pizza Hut. We added increasing amounts of salt-and-pepper noise and parmisan cheese. The system failed to recognise pizzas only when the parmisan was added in the presence of fish toppings, which is perfectly consistent with the perception of any pizza *gourmet* worth his capers.

Real data. Near misses and false alarms rates were tested systematically with real images of quiches, tacos, focaccia, assorted pizzas (types as before), the Spicy Girls, Sergeant Pepper, the Tower of Pizza, and the author's face. Typical results are shown in Figure 2. Table 2 summarises our impressive mole of results achieved with the MOS prototype. Table 1 is missing. Figure 3 shows our impressive mole.

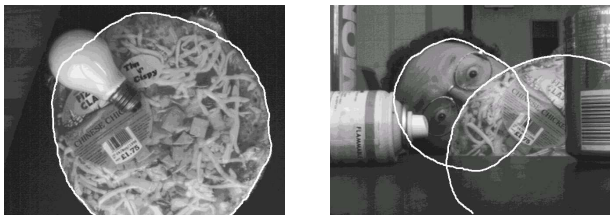

(b) (c)

Figure 1: Results of automatic pizza detection. Left: successful recognition in noisy image (occlusion, packaging). Right: successful recognition in the presence of strong occlusion (right, partial circle). Notice the false alarm (centre circle), indicating that the system insinuates the author looks like a pizza.

Jolly good.

Table 1: Quantitative summary of our impressive mole of experimental results.

Figure 2: Our impressive mole

Conclusions and future work

We have introduced a new, smart sensor which opens the (oven) door to the application of computer vision in an important industrial sector. For the future, we plan to build a MOSFET pizzel sensor (Mozzarella, Oregano, Salami, Fennel, Ecstasy pills and Tomato) to maximise performance over, arguably, an extended set of toppings of commercial and psychedelic consequence. We also plan a trip to the pub.

Acknowledgements (the closest I can get)

I am indebted to Pasquale's Pizza Lounge, as I still have to pay for my last pizza. This work was partially supported by my bicycle, on which I rested most drafts.

Notes

1. Actually, very *quasi*.
1. Only tastier.
1. E. Trucco and A. Verri, **Introductory Techniques to 3-D Computer Vision**, Prentice-Hall, 1998.
1. Pink Floyd, *Money*, in **The Dark Side of the Moon**, MFSL UDCD 517, 1973.
1. Delia Smith, **Complete Cookery Course**, BBC Books, London, 1992.
1. S. S. Cheng, **MOS Technology**, World Scientific, 1987.
1. L. L. Lapin, **Probability and Statistics for Modern Engineering**, Wadsworth Pub Co., 1992.
1. The data were abstracted from orders processed by an internationally renowned pizza restaurant, Pasquale's Pizza Lounge, Edinburgh, UK, during the summer season. The only reliable figures were the invoice numbers, so they constituted the basis of our analysis.

Emanuele Trucco
Department of Eclectic and Scented Engineering
Heriot-Watt University
Edinburgh

New Professorial Appointment

Hot off the electronic press comes the news that starting this December Edwin Hancock will be taking a Chair in Computer Vision at the Department of Computer Science at the University of York. Congratulations Edwin!

The Ninth British Machine Vision Conference

BMVC98 was hosted by the Image Speech and Intelligent Systems (ISIS) research group from Electronics and Computer Science at the University of Southampton between September 14–17th. The conference is national in name, but international in nature; at peak time there were 270 delegates. It was very enjoyable indeed.

The plenary talks were given by Shree Nayar from Columbia University USA, on Sensors for Computational Vision and Wolfgang Frstner from the University of Bonn Germany, on Sensing Scenes for Seeing Things. Further, there was an exciting tutorial programme on Computer Vision and Virtual Reality, given by John Illingworth, from the University of Surrey, and Adrian Clark, from the University of

Essex. New to BMVC this year, there were demonstrations where delegates presented working implementations of new vision techniques. Again, one day was designated as Industry Day with an excellent industrial exhibition where leading companies exhibited new vision equipment. The UK Industrial Vision Association (UKIVA) General Meeting was also held on Industry Day. Other new items included poster spotlight sessions where each poster presenter gave a 1 minute/1 slide presentation and electronic print Proceedings which had the quality of offset litho at the cost of photocopy, via the University of Southampton print centre's new machine.

The prizes included:

- the Demonstration Prize (sponsored by UKIVA) which was awarded to Manfred Prantl ("Active Object Recognition in Parametric Eigenspace") from the Technical University Graz, Austria;
- the Science Prize which was awarded to Etienne Grossmann and Jose Santos Victor ("The Precision of 3D Reconstruction from Uncalibrated Views") from Instituto Superior Tecnico Lisbon, Portugal;
- the Industry Prize (sponsored by CRS) which was awarded to David Nicholls and David Murray ("Applying Visual Processing to GPS Mapping of Trackside Structures") from the University of Oxford, UK; and
- the Poster Prize which was awarded to Xinquan Shen and Mike Spann ("3D Shape Modelling through a Constrained Estimation of a Bicubic B-spline Surface") from the University of Birmingham, UK.

I chaired the conference; Paul Lewis was Programme Chair; John Carter was Media Chair; Joanne May and Amanda Goodacre ran the registration very smoothly, and Mike Grant, Bob Roddis and Dave Hurley helped delegates with a/v, machines, posters and with general information, and didn't even complain about the yellow jackets they had to wear!

I enjoyed the conference a great deal, even if I couldn't get to all the papers. I was very impressed by the quality of the presentations I saw, both poster and oral and the demonstrations. I was very grateful to members of the programme committee not only for their efforts in reviewing the papers, but also for their chairing of the conference sessions, especially to Edwin Hancock, John Illingworth and Paul Lewis who chaired the new poster spotlight sessions.

The conference dinner was good fun (even if I lost my voice, perhaps it was Tim Ellis' jokes about my waistcoat!), and the wine was especially drinkable. Apart from sorting out the special edition of Image and Vision Computing, I still haven't found out why Luo, Cross and Hancock's paper was printed twice. Like I say, we enjoyed having you here. Time moves on now to BMVC99 in Nottingham. It's your round Tony!

Mark Nixon
ISIS Research Group
Electronics and Computer Science
University of Southampton
email: msn@ecs.soton.ac.uk

Editor's note:

Shree Nayar is currently on sabbatical and so could not write a full review of his impressions of the conference. He did however send the following:

I will say this: I was very impressed with the quality of the BMVC conference. For a domestic conference, the standard of the papers and presentations was very high. In my opinion, a good number of the papers at BMVC would have had no problem being accepted to the leading international conferences in the field. Certainly, a lot of the credit must go to the organisers and the program committee.

Shree Nayar
Department of Computer Science
Columbia University

Workshop on Handwriting Analysis and Recognition

The Third IEE European Workshop on Handwriting Analysis and Recognition was held in Brussels on 14-15th July, 1998, previous events in the series having been held in 1994 (also in Brussels) and in 1996 (in London).

The Workshop aimed to achieve two principal goals:

1. To provide a forum for European researchers in the handwriting analysis field to present and review their work.
2. To create an informal yet structured setting in which ideas could be shared, progress discussed, new interactions facilitated, and future research directions identified.

The workshop was intended to complement - at the European level - the major international events in this field, namely IWFHR (International Workshop on Frontiers in Handwriting Recognition), ICDAR (International Conference on Document Analysis and Recognition) and ICPR (International Conference on Pattern Recognition). In 1998 IWFHR was in Korea and ICPR was in Australia, so the IEE Handwriting workshop was the only opportunity for European researchers to present their work locally.

The mix of attendees was very encouraging and the numbers (33) were ideal for an interactive workshop. While this may not seem large compared with the ICDAR or IWFHR conferences, those present included not only representatives of all the major U.K. handwriting research groups, but also wide European representation including contributions from France, Belgium, Holland, Greece and Germany. We even had participants from Russia and the USA ! The mix of Universities, Research Institutions and, encouragingly, several industrial organisations, gave an ideal cross-section of activity in the field. Almost all participants (with the exception of some of the industrial participants) presented at least one paper on their work, ensuring that the essential 'workshop' character of lively interaction and discussion was maintained throughout the meeting, not only during the formal workshop sessions, but also late into the evening in the bars and brasseries of Brussels.

The initial focus of the workshop was defined by two invited Keynote Presentations from Lambert Schomaker (NICI, University of Nijmegen) and Thomas Bayer (Siemens Electrocom). Dr. Schomaker presented an overview of Pen computing, which reviewed the experience of both research and commercial applications in this field. The performance of handwriting recognition software is only one aspect of pen computers, and the market appeal of such products depends equally on their overall conceptual integrity and practical usefulness, as evidenced by experience gained with products such as Pen Windows, the Apple Newton and, most recently, the Palm Pilot palmtop computer. Nevertheless, continuing miniaturisation and the increasing introduction of wearable computers may yet provide a major outlet for pen and gesture-based computing systems in the future. Subsequent sessions on the first day followed up this overview with detailed reports on current research into writer characteristics and biometric applications of handwriting, and on on-line and off-line character and cursive handwriting recognition.

On the second day, still replete from our Workshop dinner the night before (see below), Dr. Thomas Bayer brought our minds back to the workshop topic with an expert and comprehensive review of the state of the art in off-line handwriting applications. His industrial perspective and unparalleled experience of research, development and commercial applications at Daimler-Benz (now taken over by Siemens) provided valuable insights and gave a stimulating and helpful focus to the discussions, which were now directed at the systems level and concerned with document and text processing tools and current and potential future applications. Of particular interest were presentations from Parascript (Russia and USA) and A2iA (France) describing commercial applications of offline handwriting in bank cheque processing, but our own award for the most informative and entertaining presentation at the workshop goes to Professor Dave Elliman for his inspiring paper on 'Reading Auntie Pauline's Christmas Letter'.

The venue chosen for the meeting, the Metropole Hotel Brussels, was an impressively high class hotel with spectacular decor in the traditional Brussels architectural style, and the meals from the hotel's five-star kitchen would have satisfied the most demanding gourmet. The unanimous view of participants was that they had received real value for money both in the content and style of the Workshop.

When the first IEE workshop on this topic was held in 1994, it represented a significant new initiative for the IEE in organising overseas events. Brussels was chosen as the venue partly because of its geographical location which allowed easy access from all parts of Europe, but also because of its symbolic significance within the European Community. Judging by the success of this third workshop, a very successful and unique formula for effective interaction between researchers and industry has now been established, and there is every reason to believe that a Fourth successful Workshop will be arranged in around two years time.

Mike Fairhurst
University of Kent
email: `M.C.Fairhurst@ukc.ac.uk`
Andy Downton
University of Essex
email: `acd@essex.ac.uk`
(Workshop Organisers)

Semantic Networks for Understanding Scenes

Gerhard Sagerer and Heinrich Niemann

ISBN 0-306-45704-0, Plenum Press

Knowledge representation and utilisation is the key factor in any discipline. Here too, in this book, Knowledge is the key word. The nine chapters of the book revolve around issues related to the role and use of knowledge in understanding scenes. The Introduction chapter outlines the world of computer vision analysis using a series of postulations, and abstractions of information and knowledge from image processing to image understanding. Chapter 2 is titled Segmentation and discusses the types of results that can be achieved at a stage which is data-driven and requires no task-specific knowledge. While this philosophy is widely accepted and practised, the lower levels of image analysis can benefit from some knowledge of the goal at hand. For example, through feedback, higher levels of understanding could direct lower levels of processing for more robust output in order to further focus on the hypotheses that are under verification.

The next chapter, Knowledge Representation, examines diverse issues such as different views of knowledge and their representations, and the basic approaches to the design and specification of knowledge representation languages. Of course one category of such languages is Semantic Networks which encapsulate the idea of knowledge storage in terms of a graph. Chapter 4 defines such a knowledge representation language. The language, along with related control algorithms form the knowledge representation system ERNEST (Erlangen semantic NEtwork System and Tools).

Important prerequisites of goal directed search among competing alternatives are the issues of quantitative judgement and control of processing. These two topics are covered in chapters 5 and 6. Acquisition of knowledge is the subject of the next chapter and uses examples to construct, generalise, and specialise models. These are all within the ERNEST system and Chapter 8 follows with a brief description of its User Interface. The final chapter considers several applications such as knee joint diagnosis, speech understanding, and description of industrial objects.

Examples throughout the book illustrate the theories and definitions presented. The book is also supplemented by a thorough set of references. However, a spell-check of the text, in addition to some typesetting corrections, is well advised if there is to be a second edition! In fact a thorough proof-reading is necessary to remove numerous mistakes. Overall, this book is useful for most vision system designers even if they do not plan to design their system using Semantic Networks. The concepts regarding knowledge representation are bound to have a positive effect in the design of most vision systems.

Majid Mirmehdi
Department of Computer Science
Bristol University
email: M.Mirmehdi@compsci.bristol.ac.uk