BMVA News

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

Volume 13 Number 3
March 2003

Editor: Professor Roy Davies
Department of Physics
Royal Holloway, University of London
Egham, Surrey, TW20 0EX
tel: +44(0)1784 443497
fax: +44(0)1784 472794
email: e.r.davies@rhul.ac.uk

http://www.bmva.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 1 June 2003.

Contents

[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.

Editorial: *Books and the Organisation of Knowledge*

Over the past few years we have had a number of book reviews in these pages. This time we are fortunate to have two more, and interestingly, the short biography on our 2002 Distinguished Fellow, Professor Josef Kittler, recognises the value of his much earlier book on Statistical Pattern Recognition. Such mentions and reviews are useful in bringing information on readily accessible knowledge to all of us. Of course the growth of knowledge and data is something of a curse, so much so that we have to find a delicate balance between getting as much information as we can, and making sense of it. Really, it is only *organised* knowledge that matters: a 'splurge' of data is not much use to anyone. In fact, good books are indispensible both to initial learning and to keeping up with the expansion of knowledge. (And of course books are able to integrate and organise knowledge in a way that journals and conference proceedings cannot.)

So how can we cope with the rapidly growing numbers of books on computer and machine vision, and with those covering closely related pattern recognition and image processing areas? And are they really in competition, continually displacing each other in an ever more fearsome market place? And when will it stop? Well, clearly it won't stop,

but I see a more steady state situation in which there will be a core of 'general purpose' basic books and increasing numbers of more specialised volumes. I expect it to go that way because machine vision is now in an expansionary, more mature phase where so much is known that many potential applications are able to be implemented cost-effectively, with the result that many more people will become involved in the subject. At the same time, the maturing nature of the subject will continue adding more sophisticated theory, much of it being specialised and thus not needed by everyone. So the general purpose basic books that I mentioned earlier must be profound enough to lead people as far as the more advanced and specialised volumes they may need for individual applications.

By way of example, volumes such as Faugeras and Luong are not easy to read without fair background—preferably in subjects such as projective geometry, which is a solid mathematical topic. What is more, it has deep implications and considerable subtleties: coupled with which it cannot be considered on its own but has to be examined in conjunction with error propagation, robust statistics, problems of aliasing, and so on. All this may not be too difficult for those in (for example) Faugeras' research group, where there is a wealth of 'floating' knowledge, but it is difficult for a good many starting students to make up the gap between the written word and the practice of the subject.

These considerations suggest that there is a third stratum of targeted mathematics books that is needed in our discipline. Such books would for example, take people in easy stages through projective geometry, and lead them into Faugeras' world with less headache than either Faugeras and Luong or even Hartley and Zisserman (CUP, 2000)—the problem in both of these cases is that they are aiming too quickly at computer vision rather than the mathematical concepts I have in mind. Could it be that these and other praiseworthy authors have in fact too narrowly addressed the subject as a whole and failed to present its mathematical foundations to an adequate extent for proper learning?

Unfortunately, as an Editor, I have so far failed to get people writing in with opinions on such topics, but I do hope that this time some readers will help me to clear the air and confirm or deny my prognosis on this vital subject.

Professor Roy Davies
Editor, BMVA News
email: e.r.davies@rhul.ac.uk

UKCRC Grand Challenges Discussion Phase

The UK Computing Research Committee (UKCRC) recently polled the UK computing community in order to try to define the "Grand Challenges" which computing research must address in the next couple of decades. The idea of a Grand Challenge is a research target which it is widely agreed we should meet, which we're not near now, but which we can imagine reaching in 15 years. If we can provide a clear statement to other scientific fields and to the broader population about what our research is *for*, why it's valuable, and of course why it's hard, we stand a much better chance of persuading the general population why it should be supported.

Proposals from the community were discussed at a workshop in December, and narrowed down to a few draft proposals which we hope represent valuable challenges to the research community. In order to find if they do, we need to hear from active researchers in the field, so a discussion forum has opened at:
http://umbriel.dcs.gla.ac.uk/NeSC/general/esi
/events/Grand_Challenges/Proposals/index.html .

Your input is really important to make this process represent more views than those of the hundred or so researchers whose proposals went into the initial process.

The vision people present were mostly involved in preparing the challenge in the area of human-oriented computing, which is called "Memories for life" (or "dealing with data" for short). The challenge is to research techniques for taming the enormous quantities of multi-format information with which people are increasingly overloaded. Of course, some of this is done—Google and Citeseer index my own papers better than my memory can, but we need to manage images and audio at least as well as we can currently deal with text.

Please look at the proposals on the above webpage, and criticise, support or otherwise comment. Your input may have a lot of influence on how computer science is perceived in the UK. You can post directly to mailto:GCProposal-3@nesc.ac.uk once you've looked at the challenge.

Dr Andrew Fitzgibbon
Oxford University
email: awf@robots.ox.ac.uk

Professor Josef Kittler: BMVA Distinguished Fellow 2002

Starting from a provincial town in Czechoslovakia and making it as one of the top professors in the Western world is an epic of coincidence, adventure and achievement! When the Russian tanks entered Prague, Josef Kittler was an undergraduate student of Radio Electronics at the Czech Technical University in Prague, and found himself on a brief visit abroad.

That summer Josef came to the United Kingdom to spend his holidays doing voluntary work in a hospital in Northern Ireland. Faced with the sudden overcrowding in his native country he decided to continue his education in the UK. However, before being able to resume his studies, he had a major task to accomplish: he had to learn the language of his new country. One year after the Prague Spring had ended, Josef Kittler was accepted as a final year student in the Engineering Department of Cambridge University.

He did well enough to be accepted as a PhD student in the Control Engineering Group headed by Professor John Coales FREng, to work on the topic of Statistical Pattern Recognition under the supervision of Professor Peter Young.

From then on, his academic life has been a series of successes and triumphs: his PhD in 1974; his book in 1982 published by Prentice Hall being considered for years the "Bible" of Statistical Pattern Recognition; his SERC fellowship at Cambridge; his Royal Society European Research Fellowship at Ecole Nationale Superieure des Telecommunications, Paris; his IBM and subsequent Rutherford Research Fellowships at Oxford; his job with Rutherford Appleton Lab; and finally his Readership at Surrey University in 1986 with the task to set up research in the area of Pattern Recognition and Machine Intelligence. He achieved all this, and set up one of the most successful groups in the country, managing to raise some £9,000,000 in grants over the years, produce more than 500 publications, more than 150 of which were in journals, edit several books, chair many committees, sit on the editorial boards of several journals, offer consultancy to many companies and finally set up his own spin-off company, OmniPerception Ltd, in 2001 to exploit his research output in Biometrics.

His professional activities and successes were rewarded by many scientific honours he has received: ScD by Cambridge in 1991, "Honorary Medal" from the Electrotechnical Faculty of the Czech Technical University of Prague in 1995, honorary Doctorate by Lappeenranta University of Technology in 1999, elected president of the International Association of Pattern Recognition in 1994–1996, IAPR Fellow in 1998, Fellow of the IEE in 1999, Fellow of the Royal Academy of Engineering in 2000, IEE Achievements Medal for outstanding contributions to visual information engineering in 2002, Fellow of the Institute of Mathematics and its Applications in 2003, and Distinguished Fellow of the BMVA in 2002.

Reverse Engineering the Human Vision System

I seem to remember a year or so ago that there were quotes in the press saying that the dawn of the new millennium would be remembered as the century of the brain, the time that the human race finally understood the most complicated system known to man. Certainly it isn't going to be for a lack of effort if this doesn't happen. There is not just one but several entire scientific disciplines (including psychology and neuroscience) dedicated to understanding brain structure and function, and many whose goals require an understanding of brain function, not to mention the medical disciplines such as psychiatry.

Despite some quite compelling early results (which have convinced us that the brain is indeed a computational engine), we are still a very long way from any real understanding of the computational processes involved in our conscious perception of the world. The difference now is that there are new imaging tools available (such as fMRI, PET and EEG) which for the first time give us the ability to monitor some aspects of brain function non-invasively. Such techniques offer the possibility of testing hypotheses for computational models of modular brain function as never before. First, however, we need some models.

The presentations at the meeting covered a broad range and illustrated both the diversity and difficulties of this area. Roberta Piroddi (University of Surrey) and Sumitha Balasuriya (University of Glasgow) gave talks on the related area of sampling data using an irregularly spaced sensor (such as the retina). The topic then changed to that of computing saliency in image data with a talk by Zhaoping Li (University College), with a computational model which looked similar to that suggested several years ago by Stephen Grossberg (to account for a wide variety of visual illusions) for segmentation of structures in early visual processing.

Aaron Sloman (University of Birmingham) has long been an advocate for re-introducing the grand challenge, of understanding the human brain, back into the machine vision community. Certainly, it's hard to argue with his observation that machine vision has focussed increasingly on narrower and narrower problems and his opinion that a broader set of goals may bring better results. His talk covered a description of the kind of architecture which he thinks might be needed to produce the level of perceptual sophistication we experience. As such, you could say that it makes an attempt to identify the kind of computational models which we may expect to find in the brain. The issue of whether these modules would ever be sufficiently localised in order to identify them is a separate question which can only be solved by experiment.

Masud Husain (Imperial College) and Bob Fisher (University of Edinburgh) talked on the topic of visual attention. Bob Fisher's talk presented results from a system designed to generate the visual search characteristics seen in humans, raising the intersecting issue of how the data from such a system could be used efficiently for scene interpretation. This was followed by work from the University of Southampton, investigating the recognition of human gait.

Lewis Griffin (Kings College) suggested that the current failure to generate machine vision systems with capabilities similar to human performance may be in part due to a poverty of features, in particular the over-reliance on edges, and described a new way of identifying feature classes. The debate that ensued was followed by a related talk on the use of wavelets from Imperial College.

My own interests in this area were represented by Chris Atherton, a former PhD student of mine who has recently taken a post at the University of Central Lancashire. She presented results from a collaborative study between ourselves and Charles Leek in the University of Bangor. The project attempted to identify areas of the brain responsible for the construction of a rotation invariant representation of shape, using fMRI. The motivation for this work came both from biologically plausible computational models for shape recognition and psychophysics results on the recognition of line drawings.

The issue of computational modelling of brain function, paired with the central role that vision appears to have in both brain structure and perception, suggests that the subject of machine vision could be of fundamental importance to understanding key issues in neural processing. However, since early attempts to directly focus on these ideas, as long as twenty years ago, machine vision as a subject seems to have carefully avoided the area of comparative biological modelling. Entire conferences in the area of Machine Vision can come and go without the slightest whiff of the topic.

This meeting at the Royal Statistical Society, chaired by Maria Petrou, is the first that I can remember which makes an attempt to alter this trend and just looking at the attendance list is enough to convince me that there is a potential critical mass of interest within the BMVA. I look forward to more meetings of this kind in the future and the hope that this topic may one day be given podium space at the BMVC.

Neil Thacker
University of Manchester
email: neil.thacker@man.ac.uk

BMVC 2003

The British Machine Vision Conference (BMVC) is the main UK conference on machine vision and related areas. Organised by the British Machine Vision Association, the 14th BMVC will be held in September 2003 at the University of East Anglia, Norwich, UK. Papers will be refereed on their originality, presentation, empirical results, and quality of evaluation.

Conference topics include (but are not limited to):

- image features and coding
- multi-view and stereo vision
- grouping and segmentation
- texture, shading and colour
- object recognition
- real-time and active vision
- shape and surface geometry
- medical and industrial applications.

BMVC 2003 will be a single-track meeting with oral and poster presentations. The proceedings will be available to delegates at the conference in hard copy and on CD and a selection of the best papers will be published separately in a special issue of the journal Image and Vision Computing.

In addition to the contributed papers BMVC 2003 will include presentations by invited speakers and a pre-conference tutorial programme on Monday 2 September 2002.

BMVC 2003 will incorporate a designated Industry Day on Wednesday 10 September. Presentations on Industry Day will focus on the practical applications of machine vision, and we are particularly keen to see presentations of well evaluated applications which clearly show thorough understanding of the underlying principles.

Delegates will be able to view poster presentations and see demonstrations by both industrial exhibitors and researchers demonstrating their academic work.

Important Dates:

Deadline for paper submission: 11 April 2003.
Notification of acceptance: 16 June 2003.
Deadline for camera-ready electronic copy: 11 July 2003.

For further details please see the conference website at: www.sys.uea.ac.uk/bmvc2003 .

Dr Richard Harvey
University of East Anglia
email: r.w.harvey@uea.ac.uk

'Algorithms to Art' exhibition at BMVC 2003

Computer vision algorithms are being used in an increasing number of applications for creating art. BMVC 2003 will provide an opportunity to exhibit art created with the help of computer vision algorithms.

Coffee, posters, commercial exhibits and the 'Algorithms to Art' show will run concurrently throughout the conference in a room 33 m by 10 m about 50 m from the lecture room.

'Algorithms to Art' entry requirements:

1. Attendees of BMVC 2003 are entitled to exhibit.

2. Artwork must be created either entirely automatically, or semi-automatically, or from a photograph, or through manual tools based on vision algorithms.

3. Each exhibit shall be accompanied by a single page account of how the artwork was produced. Minimum font size of 16 point. The account to include the Title, author(s) and key algorithms.

4. Submission of a related paper is not required, but encouraged.

To enter, please notify the Conference Organiser by 1 May 2003.

Dr Richard Harvey
University of East Anglia
email: r.w.harvey@uea.ac.uk

Ninth EPSRC Summer School on Computer Vision

University of Surrey, Guildford, 23–27 June 2003.

The Summer School begins at 2 pm on 23 June 2003 and ends at lunchtime on 27 June 2003. The registration fee includes attendance at the School, a copy of the course notes, accommodation on the nights of 23, 24, 25, 26 June 2003, and all meals.

All applicants must complete an application form, which must be signed by their research supervisors. Application forms are available from Ann Levers at the University of Surrey (tel: 01483 689801, email: a.levers@eim.surrey.ac.uk).

Application forms should be returned to:

Professor M Petrou, EPSRC Summer School on Computer Vision, Centre for Vision, Speech and Signal Processing Research, School of Electronics and Physical Sciences, University of Surrey, Guildford, Surrey, GU2 7XH.

Professor Maria Petrou
University of Surrey
email: m.petrou@surrey.ac.uk

Book Review—Faugeras and Luong

Olivier Faugeras and Quang-Tuan Luong, *The Geometry of Multiple Images: The Laws That Govern the Formation of Multiple Images of a Scene and Some of Their Applications*, MIT Press, March 2001, ISBN: 0-262-06220-8, Hardback £43, 644 pp.

Projective Geometry, the geometry that best models image formation, provides a unified framework for tackling many geometric problems relevant to vision. This book formalises and analyses the geometric laws that relate multiple views of a scene. Many theorems and their proofs are provided, along with some algorithms and examples to aid the reader in their implementation.

Olivier Faugeras has been a leading figure in the field of Computer Vision for many years; he is Research Director and head of the Robiovis group at INRIA and an Adjunct Professor at MIT. His previous book, the hefty *Three-Dimensional Computer Vision* (MIT Press, 1993), provided a rigorous exposition of the broader area of 3D vision.

The book begins by providing a (relatively) gentle introduction to the principle areas of projective geometry before, in effect, summarising the key areas it addresses. Chapter 2 describes the tools of Projective Geometry and the important concept that Euclidean and Affine geometry are merely special cases of Projective geometry. Chapter 3 then discusses the complementary algebraic viewpoint and introduces Grassman–Cayley algebra.

After these initial chapters the following chapters deal with the geometry of increasing numbers of views. Chapter 4 looks at the case of one view and introduces the geometric representation of the pinhole camera, the perspective projection matrix. Chapter 5 examines the two-view case, introducing epipolar geometry and the Fundamental Matrix, with the estimation of the Fundamental Matrix being dealt with separately in Chapter 6. The next chapter discusses the different levels of reconstruction; Projective, Affine and Euclidean. The amount of information needed to upgrade the reconstruction to the next level is shown along with what properties are gained by that upgrade. Chapters 8 and 9 go on to deal with the theory and estimation of the 3-view equivalent of the Fundamental Matrix—the Trifocal Tensor. Chapter 10 then looks at the N-view case, before the final chapter deals with the important area of Self-calibration.

Overall, this book provides a thorough discussion of the application of Projective Geometry to the

Structure from Motion problem. Unfortunately, the style of the book is very much like a traditional mathematical text, full of lemmas, propositions and proofs. Whilst this certainly formalises the topic, it also renders the book much less readable and, in my opinion, the style will put off many readers looking for an introduction to the field or those hoping to use it as an aid to implementing some of these theories.

It would be very hard to review this book without drawing direct comparisons with *Multiple View Geometry* by Hartley and Zisserman (Cambridge University Press, 2000). Although both books tackle exactly the same topic, *Multiple View Geometry* is much more accessible; the explanations are clearer and many algorithms are provided in a form that makes them easy to implement. *The Geometry of Multiple Images* would provide an ideal companion to Hartley and Zisserman's book as it does approach the same problem from a slightly different standpoint and in a more formal manner.

Oli Cooper
University of Bristol
email: `oli.cooper@bristol.ac.uk`

Book Review—Nixon and Aguado

Mark Nixon and Alberto Aguado, *Feature Extraction & Image Processing*, Newnes, 2002, ISBN: 0-7506-5078-8

As you start reading this review, you're probably thinking to yourself "not another book on computer vision". Well, that's precisely the question the authors pose in their preface, and they answer it by pointing out that new vision techniques are continually being developed, partly fuelled by the rapid improvement in computer hardware in recent years. They also point out that, as the field is now so wide, it makes sense for a new book to concentrate on a specific topic. As there is a relative dearth of books on feature extraction, this is the area that they emphasise.

The book starts *ab initio*, so it would be useful for an undergraduate course in the area. Rather than illustrate algorithms with arbitrary pseudo-code or pages of impenetrable C, C++ or Java code, the authors choose to present algorithms in the much more compact Mathcad or Matlab forms. This makes it fairly easy to see the parallels between mathematics and code. The writing style is approachable and readable, though some of the section titles are a bit trite (such as "Hello Mathcad, hello images").

The book starts by presenting an overview of the eye and the camera, a rationale for using Mathcad and Matlab, and a section on books, journals and web resources.

The technical material starts with a consideration of images, sampling and Fourier analysis. This choice seems a little strange to this reviewer, as it may well put off students with backgrounds in, for example, computer science rather than engineering or physics. The chapter, and indeed the book as a whole, does not shy away from including mathematics, though the level should not be beyond the capabilities of a good science undergraduate.

Following that, Chapter 3 introduces conventional image processing: point operators, histograms, convolution, and related techniques. The following chapter concentrates on low-level feature extraction, emphasising edge detection. As well as describing *de facto* standards such as Sobel and Canny, it touches on less familiar ones such and Spacek and Petrou. The chapter also introduces the measurement of image curvature and motion. All these are supported by appropriate mathematics and, in many cases, code.

Chapter 5 considers feature extraction by shape matching and, not surprisingly, includes a lot of material on various forms of the Hough transform. The following chapter considers active contours, and includes a complete implementation of a snake. (The latter seems to have inspired the book's striking cover, which features a green snake slithering over a leaf.) The code examples here are a little less clear than in the earlier chapters, perhaps because more sophisticated algorithms do not lend themselves as well to clear exposition in matrix-oriented programming languages.

Chapter 7 moves on to consider the descriptions of objects, firstly by their boundaries and then by region-oriented descriptors. Finally, Chapter 8 gives a bird's-eye view of texture segmentation and classification.

Did I like the book? Yes, I did: I'd have no hesitation in recommending it as a text for a lecture course on image analysis. Indeed, some of my research students have used the book to help them get started with some classes of feature extraction, and they were quite complimentary about it. My two lasting impressions from this book are, however, the same

as when I read any other image analysis or computer vision text: that a tremendous amount of our discipline is *ad hoc*, with few underlying principles; and that we lack a common software framework in which we can express and share algorithms.

Dr Adrian F Clark
University of Essex
email: alien@essex.ac.uk

Book Review Corner

Book for Review

People seem to be missing an opportunity to review an interesting book—

J.-L. Starck and F. Murtagh *Astronomical Image and Data Analysis*, Springer, 2002, ISBN 3 540 42885 2, hardback , xi + 289 pp.

Book currently under review

Hartley and Zisserman *Multiple View Geometry in Computer Vision* (Martin Lewin).

This is scheduled for publication in the next issue of BMVA News.

Professor Roy Davies
Editor, BMVA News
email: e.r.davies@rhul.ac.uk

Information on Meetings

For up-to-date information on forthcoming BMVA meetings, events and news, subscribe to the BMVA mailing list at:

http://www.jiscmail.ac.uk/lists/bmva.html/.

This is a moderated list: you will only receive official BMVA postings.

Dr Richard Bowden
University of Surrey
email: r.bowden@surrey.ac.uk

April Deadlines

BMVC 2003

The deadline for submitting papers to BMVC 2003 is **11 April 2003**: for further details see p. 5.

BMVA Distinguished Prize 2003

The BMVA Executive Committee seeks nominations for the Distinguished Fellow 2003 award. This prestigious award is given to one person only each year in recognition to his/her services to the British Machine Vision community. The nominees must be distinguished researchers, based in the UK, who have contributed significantly to the field of research and the reputation of the British Machine Vision Community both nationally and internationally. Nominations, with a few lines of rationale, should be sent to Dr Tim Cootes (t.cootes@man.ac.uk), by **30 April 2003**.

The Sullivan Doctoral Thesis Prize 2003

The BMVA Executive Committee seeks nominations for the Sullivan Doctoral thesis prize. The 2003 prize will be awarded to the best nominated thesis which was submitted and examined during the calendar year 2002. Thesis examiners and supervisors may act as nominators, but the committee would like to receive an accompanying report and endorsement of the nomination from the thesis external examiner.

Please send any nominations to the Secretary of the Association, Dr Dave Marshall (Dave.Marshall@cs.cf.ac.uk) by **30 April 2003**. The theses nominated should be made available through a web page. In addition, two hard copies (not necessarily bound) should be sent to Dr David Marshall, Department of Computer Science, Cardiff University, Queen's Buildings, Newport Road, Cardiff, CF24 3XF.

Solution to Christmas Crossword

Overleaf you will find the solution to Richard Harvey's Christmas crossword. I'm sure many of us have enjoyed tackling it and will join with me in thanking him!—Ed.

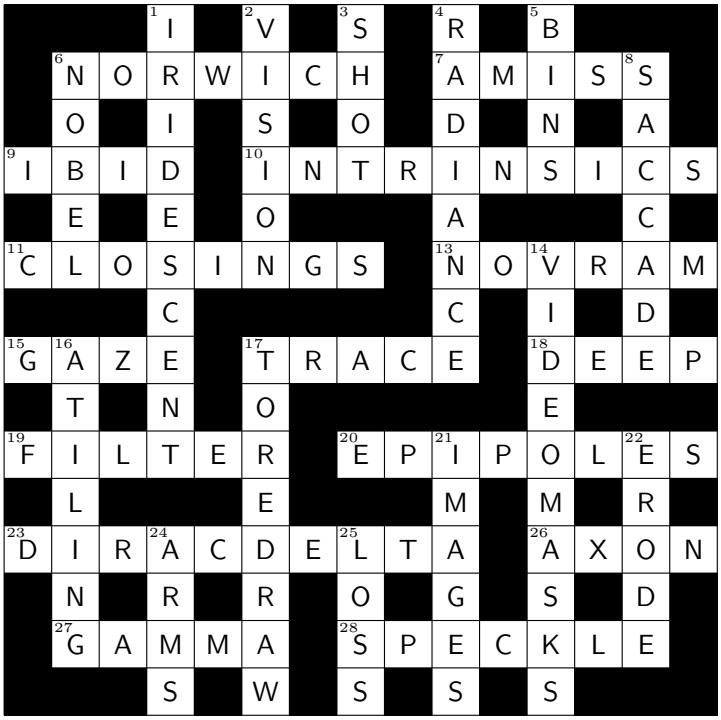

ACROSS

6 Confused chairwoman, missing Active Appearance Model, tells us the location of BMVC 2003. (7)

7 Sam is puzzled when this morphological transform doesn't hit the foreground. (1,4)

9 In the same place, instinctive impulse about head of bigot. (4)

10 Camera parameters are integers around Royal Institution's points with one charge. (10)

11 Approaches blockings. (8)

13 Head of Novisad fronts press explanation of why graphics card doesn't work? (2,3)

15 Examination in Malaga zealotry. (4)

17 Draw around the sum of the eigenvalues? (5)

18 Urinated around Koenderink's structure. (4)

19 Fir cigarette-end? (6)

20 $2.7183\ldots+3.1416\ldots+\frac{1}{z^2-1}\Big|_{z=\pm 1}$ = stereo intersections. (7)

23 Paul's impulse? (5,5)

26 Nervous part. Sounds like the Play's started! (4)

27 This intensity normalisation is implemented with genetic algorithm followed by one-thousand degree annealing. (5)

28 Remove arsenic from crazy Cocker Spaniels' noise. (7)

DOWN

1 Rainbow-like colour incites red mix-up. (10)

2 Number six is back atop our subject. (6)

3 Spirit photograph? (4)

4 Measure of light intricate dance air. (8)

5 Throw-away histogram. (4)

6 This prize sounds like a dead-ringer! (5)

8 Marquis takes on head accountant for eye movements. (7)

14 Eccentrically makes voids for mattes. (5,5)

16 Penrose has it on his floor? (1,6)

17 Re-render frantic rowed rat. (2,6)

21 Digital pictures of unreasonable ageism. (6)

22 Reduce the sound of Yorkshire retired jockey's former activity. (5)

24 Sounds like Cockney wounds robotic manipulators. (4)

25 Unprofitable function of estimation theory. (4)