Institute *of* Physics

Newsletter

of

The Computational Physics Group

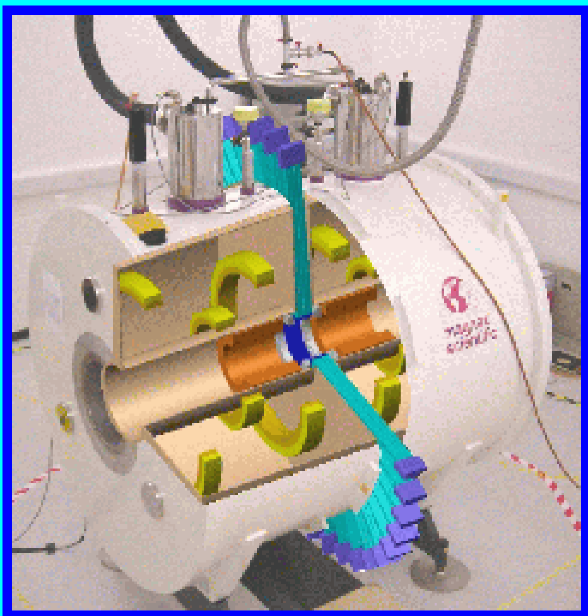

Spring 2005

MEMBERS OF THE COMMITTEE

Peter Borcherds (chairman)	p.h.borcherds@birmingham.ac.uk
Allan Boardman (treasurer, secretary)	a.d.boardman@salford.ac.uk
Roger Barrett	R.Barrett@surrey.ac.uk
Alan DuSautoy	alan.dusautoy@npl.co.uk
Andrew Horsfield (newsletter)	a.horsfield@ucl.ac.uk
Dave Hutchings	d.hutchings@elec.gla.ac.uk
Geraint Lewis	dg.lewis@physics.org
Jim McNiff	jemc@tinyonline.co.uk
Ian Morrison (web page)	i.morrison@salford.ac.uk
Massimo Noro	Massimo.Noro@Unilever.com
Matt Probert	mijp1@york.ac.uk

Our web page can be found here: http://www.iop.org/IOP/Groups/CP/

Comments about the newsletter should be sent to Andrew Horsfield

Front cover picture courtesy of Richard Ansorge, Cambridge University.

Contents

Image Modelling

Maria Petrou, University of Surrey

For most people an image is an icon, a picture. When people see an image they see flowers, trees, the sky, the sea... When computers see an image they see an array of numbers. If only computers could manage to decipher the language of images, an array of those numbers would tell them more than a thousand words! For a start, these numbers convey their message not only by their relative position, but also by their relative values. So, if we want to decipher that message we have to understand and use both these ways of conveying information. And there is an awful lot of information conveyed by an image. To appreciate that, all you have to do is to stand in front of a picture in the National gallery and let a guide talk to you about it. He/she may be talking for a very long time on the same image! How do we get all that information out of a single two dimensional piece of canvas?

This is one of the great unsolved problems of artificial vision: we cannot produce a vision system that sees pictures in the same way humans do. The problem is unsolved because it is difficult, and it is difficult because it is an inverse problem: a large part of the history of evolution of the world itself may be hidden behind an image, collapsed and captured on a 2D finite piece of canvas: all the processes that contrived to put the objects in the world in the relative position with each other so that when painted (or photographed) they produced the particular combination of colours and flat 2D patches next to each other on the image plane, and all the processes in the world which placed you, and your camera, in the particular position with respect to the lighting source, at the particular distance from the objects, with the particular camera, with its particular resolution and noise levels, and made you press the button at the particular time, to capture the scene for ever! It is rather too ambitious to hope that we shall be able to recover all these items of information from a 2D image! But, it is characteristic of the human nature not to give up easily and to think of the unthinkable! Engineers know how to deal with problems when the direct solution is difficult: they go backwards! They start from a model and they check whether it fits the data or not.

The simplest thing one tries to model in image processing is of course the image noise. It is a common practice to assume that the noise in an image is Gaussian, appealing to the central limit theorem. In practice, noise is often Gaussian, but not always. In synthetic aperture radar images for example, the noise is shot noise consisting of pixels that simply have very wrong values due to reflections and occlusions. Having a model of the noise field then opens up the possibilities of reducing noise levels, and detecting features buried in high levels of noise. This is the basis of processes like low pass filtering and optimal filter design for edge detection. Canny's filters have become classical in image processing exactly because he started by modelling what he was looking for: a step edge immersed in additive white Gaussian noise.

However, the real challenges start when one wishes to go beyond modelling the noise field. A large part of image processing relies on linear operations: it relies on the definition of bases of elementary images useful for expressing an image as a linear combination of them. Of course one may use singular value decomposition and identify the optimal such basis for an image. (Optimality is often a loosely used term and one should be careful what it means in each context. In this context it means that if we truncate the expansion, we create the least square error approximation of the image.) Such intrinsic image bases are useful in some applications, but very restrictive: we have a different basis for each image. Not very practical. At the other extreme we have the universal bases in terms of which any image may be expressed as a weighted sum, with the coefficients of the expansion fully defining the image. An example of such a basis is the 2D discrete Fourier transform. Alternative universal bases may be constructed from the vector outer products of the discrete versions of the Walsh or the Haar functions.

Universal bases are good, as everybody knows them. However, when these expansions are truncated, the approximations created are not optimal in any sense. When you truncate a Fourier expansion, you omit the higher harmonics. However, nobody guarantees that these harmonics have the least energy, and even more, there is no way to know the error you commit. (In the case of the singular value decomposition you always know the error you commit by truncation because it is equal to the sum of the omitted eigenvalues.) There is a need, therefore, to do something in between: can we construct a basis that is optimal for a particular class of images only?

First we have to define what we mean by class. Intuitively, we understand what a class of images is. For example, as far as applications are concerned, if a hospital wants to store in a compressed form digitised chest X-rays of its patients, all these chest X-rays constitute a class of images. If we could have at our disposal the full collection of images which the hospital archive will ever encounter, we would be able to perform principal component analysis over all these members of the ensemble and identify a basis which would be optimal in the **mean** least square error sense for approximating the images. Unfortunately, however, we do not have the full ensemble in our disposal, so we have to bypass this problem.

Here comes the first leap of imagination: we make the assumption that these images, which belong to the same **semantic** class, also belong to the same **statistical** class, by sharing, for example, the same statistical properties. Not only that, but we invoke the theory of ergodicity: we assume that the ensemble statistics of the images are the same as the spatial statistics of any image in the collection, in terms of the mean and the covariance matrix. By no stretch of the imagination is this true: each image has to be pretty much a pure random noise field to make such an assumption correct. The surprising thing is that this approach, known as Karhunen-Loeve expansion, does lead to reasonably good results!

So, we started by simply modelling the noise in an image and moved into modelling the covariance matrix of an ensemble of images by invoking ergodicity. In all cases useful results were produced, but we are still very far from recovering the lost information when we moved from the 3D world into the 2D recording medium. Before we discuss this final frontier, we may get a first taste of what an inverse problem is by thinking about image restoration. Let us say that we have a bad image. Let us say that it looks blurred. Can we undo the damage it suffered? The answer is "maybe" if we know exactly what happened.

Of course, we shall never know exactly what happened, but if the picture of a static scene was, for example, captured from a car moving with constant velocity, we can easily model the blurring effect. We may even model the blurring effect if the image was captured from a uniformly accelerating car. Blurring with constant velocity, or even constant acceleration, is a linear process. All we have to do is deconvolve the image, having modelled the degradation process by its point spread function.

And then the fun starts! What if the point spread function of the degradation process has killed some frequencies in the image? Obviously, that information has been lost for ever, and those frequencies can never be recovered.

Yes, but those frequencies **will** be present in the degraded image! How? But due to noise of course! Noise is white! Noise goes everywhere! Noise makes such efforts of deconvolution fail clearly and loudly! Even if we are resigned from ever recovering those frequencies, and we deliberately remove them from the reconstructed image, the result is still bad: once the model of the point spread function starts killing frequencies, it becomes unreliable. Best results are obtained when all frequencies beyond the first nulled frequency by the degradation point spread function are assumed unrecoverable. Statistical modelling comes to rescue again: In Wiener filtering, you assume that somehow you have modelled the power spectral density of the image you wish to recover, and the power spectral density of the noise you have, and arrive at the filter you should use to perform the restoration. Not much chance for all this power spectra modelling! In practice, all you do is to add a constant in the denominator of your filter (deciding it by trial and error) so that when the degradation point spread function at some frequency becomes too small, you are dominated by that constant, and things do not blow up! And what if you cannot model the degradation phenomenon because you hardly understand it? What do you do, for example, if the degradation has been caused by atmospheric turbulence, or some other combination of complex phenomena? Then you simply do not bother to try to understand what happens! You stop acting like a physicist and start acting like an engineer! You model the result, as opposed to understanding and modelling the process that caused it!

For example, in an astronomical image, the image of a distant star is nothing more than the point spread function of the degradation process served ready in a plate: Measure the profile of the star (which should be a Mathematical point if imaging conditions were perfect, but it is a blob instead) and you have the point spread function of your imaging set up, with the atmospheric turbulence, the instrumental defects and the interstellar distortions all bundled up together in a single convolution kernel! And what if God did not oblige us, and the degradation process is not linear?

We clearly cannot model it by a point spread function which supposingly convolved the perfect image! Then we have to make another leap of faith! We decide that we want to find the most probable configuration of values which, within the limits of noise, is consistent with the data AND the degradation model (which is assumed to have been worked out somehow, by hook or by crook).

For example, if a pixel has value 25 and the noise has standard deviation 2, and we know that the degradation process was such that the square root of the true value was recorded instead of the true value directly, then we can easily infer that for the true value to have a square root which agrees with the recorded value within one standard deviation of the noise, it must be in the range 23^2 to 27^2. Now, let us say that we have an image with 500×500 pixels and for each pixel we have 5 possible true values. This makes $5^{250,000}$ possible solutions, consistent with the information we have. If you feel like trying them all in order to find the best, forget it! The age of the Universe was only 5^{40} seconds at the last count!

Here is then when we start hand-waving: The world around us is smooth! Peer pressure makes odd ones out to comply with their neighbours! Conflicts and cultural boundaries are small in extent in comparison with the vast areas of harmonious co-existence and mutual understanding! Why not assume that the same happens in images? Why not **impose** on our solution space a structure which favours those configurations where pixels next to each other agree in their values? Why don't we impose harmony to chaos? And thus the model driven approaches are born! Effectively, we decide that the solution we shall seek will be chosen according to some prior probability that favours images with pixels next to each other being similar. Then some appropriate optimisation method is chosen which finds the best compromise between remaining faithful to the data and satisfying this prior model. And it works! It works in image restoration, in stereo problems, where two images are combined to recover the elusive third dimension, and it works in many other problems in computer vision where we are trying to solve some sort of inverse problem.

Does it give always the right answer? Certainly not! You are making a compromise between your prejudice about the world (the model) and the facts (the data). It is up to you how much you believe your data and how much you believe your prejudices as expressed by the prior model. All comes down to a parameter that weighs their relative importance and all comes down to experience, trial and error! Good luck with it! And if you want to know the ultimate challenge for modelling, you need not go far! The current buzz words are "cognitive vision". In other words, recognition of what the image shows. How can you decide that a rectangle is a window and not the top of a table, or a piece of paper lying on the floor? Very simply: by the context, by its relative position with respect to other objects in the image, the semantic meaning assigned to the other objects, and so on. Model then these semantic relations and you will hit gold!

Visual Sign and Gesture Recognition

Richard Bowden, University of Surrey

Introduction

Over the last 3 years the Centre for Vision Speech and Signal Processing (University of Surrey) has been working with the Robotics Research Group (University of Oxford) on Sign language Recognition as part of the EU FW5 Project CogViSys. The objective of this work is to efficiently and accurately recognise signed words, from Sign Language, using a minimal number of training examples. Furthermore, our aim was to use natural image sequences, without the signer having to wear data gloves or coloured gloves, and to be able to recognise hundreds of signs. The proposed approach should be language independent providing a generic solution to Sign recognition regardless of country. The motivation for this work is to provide a real time interface so that signers can easily and quickly communicate with non-signers.

Why is it hard?

1. Each country has its own sign language with different vocabularies and grammar. Any system that is to be of use must be independent of the specific language model used.

2. Individuals are different shapes and sizes and will vary the way in which a specific sign is performed, for example someone who is new to signing

will sign slower and with a larger sign space (the volume in which the sign is performed) with minimal co-articulation between signs. A fluent signer will be far faster with heavy coarticulation and typically in a far smaller sign space. This is similar to speech where a fluent speaker, with their own dialect, will blur words together and use slang and abbreviations to communicate faster. In addition to these fundamental variations our sensor modality is video. Cameras have different lenses, responses and an individual may be arbitrarily placed relative to the camera which further complicates matters.

3. Traditional approaches such as those used in speech recognition (such as the HMM) require large amounts of labelled data in order to generalise about the points raised in 2 above, such as the feature space, sign variation and co-articulation artefacts. No such databses for sign exist (unlike in speech). Considering the storage requirements of video and the task of labelling this data the acquisition of labelled data becomes a limiting factor in the size of lexicon that can be addressed. Obviously this limitation also has serious implications to the issues raised in point 1 above, one would have to generate training data for each sign language to be learnt.

How do we do it?

We break the problem down into 2 areas:

1. Generic tracking of the human, regardless of size, camera type and placement.
2. A novel 2 stage classification architecture which reduces training requirements by generating a high level feature description based upon sign linguistics.

An overview of the system is given in figure 1.

The novelty of our approach is that we structure the classification model around a linguistic definition of signed words, rather than a HMM. This enables signs to be learnt reliably from just a handful of training examples. The classification process is divided into two stages. The 1st generates a description of hand shape and movement at the level of 'the hand has shape 5 (an

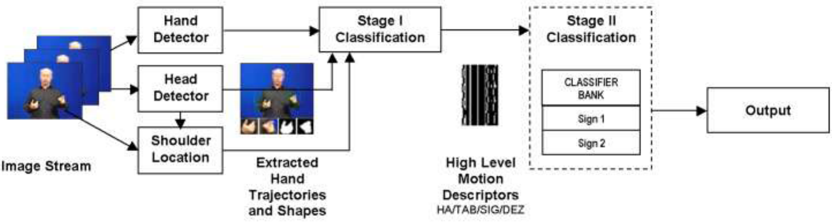

Figure 1: An overview of the system

open hand) and is over the left shoulder moving right'. This level of feature is based directly upon those used within sign linguistics to document signs. Its broad description aids in generalisation and therefore significantly reduces the requirements of further stages of classification. In the second stage, we apply Independent Component Analysis (ICA) to separate the channels of information from uncorrelated noise. Final (stage II) classification uses a bank of Markov models to recognise the temporal transitions of individual words/signs.

Tracking - Detecting Faces and Hands using Boosted Classifiers

Boosting is a general method that can be used for improving the accuracy of a given learning algorithm [3]. More specifically, it is based on the principle that a highly accurate or "strong" classifier can be produced through the linear combination of many inaccurate or "weak" classifiers.

Classifier efficiency is increased by organising the weak classifiers into a 'cascade', where the number of weak classifiers in each layer of the cascade increases with depth. The purpose of this is that initial layers (which are required to test many possible hypotheses) are simple to compute. They should reject large numbers of hypotheses such that later layers (which are more complex and therefore computationally expensive) need only be applied to a small subset of the original hypotheses. In this fashion an exhaustive search over all positions and scales is possible as in excess of 90% of possible hypotheses can be rejected at each stage of the cascade.

Model	N' Layers	Cascade Layers	+ve Training Eg	-ve Training Eg
Head	10	2,5,5,20,50,50 150,150,150,150	2500	6000
Hands	8	2,5,5,20,50,50 150,150	2400	6000

Table 1: Details of classifiers for the head and hands

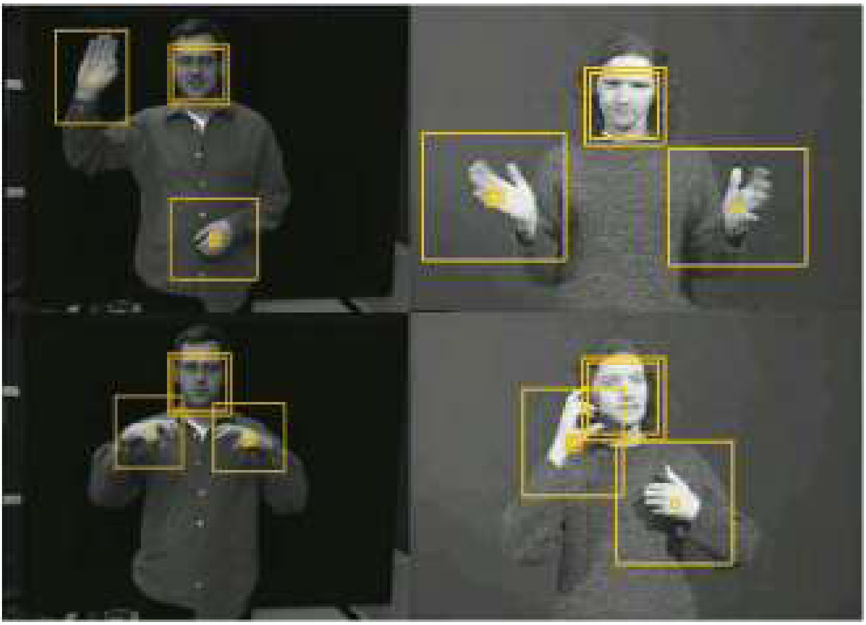

Figure 2: Sample images taken from two sequences

In order to perform detection, we train two different strong classifiers for the head and hands respectively. The details of which are given in the table 1. The entire image is searched over all position and scale. To perform detection on a section of an image the weak classifiers are transformed such that they are applied to that image section. If the image content does not contain the object of interest (head or hands), it will be rejected by a particular layer on the cascade. Otherwise it will filter down to the final layer to be accepted by it. The strongest detected head and hand hypotheses are then passed to stage I classification. Figure 2 shows some sample images taken from 2 sequences with the detected location of head and hands outlined.

Signers generally face the viewer as directly as possible to ease understanding and remove ambiguities and occlusions that occur at more oblique angles. The system uses the boosted detectors coupled with a contour model of the head and shoulders. This provides a bodycentred co-ordinate system in which to describe the position and motion of the hands. The 2D contour is a coarse approximation to the shape of the shoulders and head and consists of 18 connected points. The contour is a mathematical mean shape taken from a number of sample images of signers.

The contour is fitted to the image by estimating the similarity transform which minimises the contour's distance to local image features. Estimates for key body locations, are placed relative to the location of the head contour. This means that as the contour is transformed to fit the location of the user within the video stream, so the approximate locations of the key body components are also transformed.

Classification

Our approach is based upon a novel two stage classification:

Classification stage I

Raw image sequences are segmented in order to extract the shapes and trajectories of the hands in the monocular image sequence. The initial classification stage converts these into a viseme representation (the visual equivalent of a phoneme) taken from sign linguistics:

1. HA Position of the hands relative to each other

2. TAB Position of hands relative to key body locations

3. SIG Relative movement of the hands

4. DEZ The shape of the hand(s)

This HA/TAB/SIG/DEZ notation provides a high-level feature descriptor that broadly specifies events in terms such as hands move apart, hands touch or right hand on left shoulder. This description of scene content naturally generalises temporal events, hence reduces training requirements.

The 5 HA, 13 TAB, 10 SIG and 12 DEZ states we currently use are listed in the table below and are computed as follows: HA: the relative position of the hands to each other is derived directly from deterministic rules on the relative x and y co-ordinates of the centroids of the hands and their approximate area in pixels.

TAB: the position of the hands is categorised in terms of their proximity to key body locations using the Mahalanobis distance computed from the approximate variance of these body parts gained from contour location.

SIG: the movement of the hands is determined using the approximate size of the hand as a threshold to discard ambient movement and noise. The motion is then coarsely quantised into the 10 categories listed in the table.

DEZ: British Sign Language has 57 unique hand-shapes (excluding finger-spelling) which may be further organised into 22 main groups. A visual exemplar approach is used to classify the hand shape into twelve (of the 22) groups. This is described in more detail in our papers.

Figure 3 shows the features generated by the system over time. The horizontal binary vector shows HA, SIG, TAB and DEZ in that order delineated by grey bands. The consistency in features produced can clearly be seen between examples of the same word. It is also possible to decode the vectors back into a textual description of the sign in the same way one would with a dictionary. The feature vector naturally generalises the motion without loss in descriptive ability. Figure 3 shows the word 'different' being performed by two different people along with the binary feature vector produced. The similarity is clear, and the signed words are correctly classified. Linguistic evidence points to the fact that sign recognition is primarily performed upon the dominant hand (which conveys the majority of information) we therefore currently discard the non dominant hand and concatenate HA, TAB, SIG and DEZ features

HA	TAB	SIG	DEZ
1. Right hand high	1. The neutral space	1. Hand Makes no movement	1. 5
2. Left hand high	2. Face	2. Hand moves up	2. A
3. Hands are side by side	3. Left side of face	3. Hand moves down	3. B
4. Hands are in contact	4. right side of face	4. Hand moves left	4. C
5. Hands are crossed	5. chain	5. Hand moves right	5. F
	6. R Shoulder	6. Hands move apart	6. G
	7. L Shoulder	7. Hands move together	7. H
	8. Chest	8. Hands move in unison	8. I
	9. Stomach		9. P
	10. Right Hip		10. V
	11. Left Hip		11. W
	12. Right Elbow		12. Y
	13. Left Elbow		

Table 2:

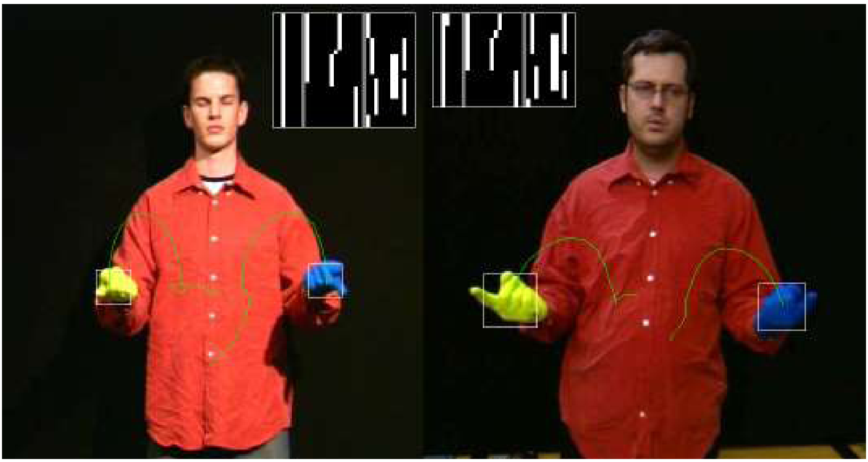

Figure 3: Features generated by the system over time

together to produce a 40 dimensional binary vector which describes the shape and motion in a single frame of video.

Classification stage II

Each sign is modelled as a 1st order Markov chain in which each state in the chain represents a particular set of feature vectors (denoted symbols below) from the stage I classification. The Markov chain encodes temporal transitions of the signer's hands. During classification, the chain which produces the highest probability of describing the observation sequence is deemed to be the recognised word. In the training stage, these Markov chains may be learnt from a single training example.

Robust Symbol Selection: An appropriate mapping from stage I feature vectors to symbols (representing the states in the Markov chains) must be selected. If signs were produced by signers without any variability, or if the stage I classification was perfect, then (aside from any computational concerns) one could simply use a one-to-one mapping; that is, each unique feature vector that occurs in the course of a sign is assigned a corresponding state in the chain. However, the HA/TAB/SIG/DEZ representation we employ is binary and signers do not exhibit perfect repeatability. Minor variations over sign instances appear as perturbations in the feature vector degrading classification performance.

For example the BSL sign for 'Television', 'Computer' or 'Picture' all involve an iconic drawing of a square with both hands in front of the signer. The hands move apart (for the top of the square) and then down (for the side) etc. Ideally, a HMM could be learnt to represent the appropriate sequence of HA/TAB/SIG/DEZ representations for these motions. However the exact position/size of the square and velocity of the hands vary between individual signers as does the context in which they are using the sign. This results in subtle variations in any feature vector however successfully it attempts to generalise the motion.

To achieve an optimal feature-to-symbol mapping we apply Independent Component Analysis (ICA). Termed feature selection, this takes advantage of the separation of correlated features and noise in an ICA transformed space and removes those dimensions that correspond to noise. More details are given in our papers.

Results

The system is a system capable of running in real-time, and generating extremely high recognition rates for large lexicons with as little as a single training instance per sign. We have demonstrated classification rates as high as 92% for a lexicon of 164 signs with extremely low training requirements outperforming previous approaches where thousands of training examples are required.

How do I find out more? We have published a number of papers in this area, the most recent that describes the system being [2] winning the BMVC05 Industrial Paper Prize. For details of the booting see [1] and [6]. For discussions about the feature selection process see [4]. For body tracking and estimating elbows see [3]. An extended journal publication on the entire system is currently under preparation. A demonstration of the system was performed at [8][5][1]. For older work on hand modelling see [9] and [10]. Failing all that contact r.bowden@eim.surrey.ac.uk. Im happy to give invited talks on this work or others areas such as Visual Surveillance, HCI, Artificial Life, Cognitive Vision and tracking. See http://www/ee/surrey.ac.uk/Personal/R.Bowden for more details.

Publications and further information

[1] Bowden R, Progress in Sign and Gesture Recognition. Invited Speaker, AMDO2004, Third International Workshop of Articulated Motion and Deformable Models, Palma de Mallorca, Spain.
[2] Bowden R, Kadir T, Ong E, Windridge D, Zisserman A, Brady M. Minimal Training, Large Lexicon, Unconstrained Sign Language Recognition. In Proc. BMVC04. Kingston UK. Sept 2004. Vol 2, pp939-948.
[3] Micilotta A, Bowden R, View-based Location and Tracking of Body Parts for Visual Interaction. In Proc. BMVC04, Kingston UK. Sept 2004. Vol 2, pp849-858.
[4] Windridge D, Bowden R, A General Strategy for Hidden Markov Chain Parameterisation in Composite Feature Space. To appear in Proc. SSPR04 Syntactical and Structural Pattern Recognition 2004.
[5] Bowden R, Windridge D, Kadir T, Zisserman A, Brady M. A Linguistic Feature Vector for the Visual Interpretation of Sign Language, In Tomas Pajdla, Jiri Matas (Eds), Proc. 8th European Conference on Computer Vision,

ECCV04. LNCS3022, Springer-Verlag (2004), Volume 1, pp391- 401.
[6] Ong E, Bowden R, Detection and Segmentation of hand shapes using Boosted Classifiers, In Proc. 6th Int Conf on Automatic Face and Gesture Recognition, FGR'04, IEEE Comp Soc TC PAMI, Korea 2004, pp889-894.
[7] Windridge D, Bowden R, Induced Decision Fusion In Automatic Sign Language Interpretation: Using ICA to Isolate the Underlying Components of Sign. In 5th International Workshop on Multiple Classifier Systems, MCS04, Cagliari, Italy, 2004. pp
[8] Bowden R, Zisserman A, Kadir T, Brady M. Vision based Interpretation of Natural Sign Languages. Exhibition at ICVS03, The 3rd International Conference on Computer Vision Systems, Graz. Austria, April 2003.
[9] Bowden R, Sarhadi M, A non-linear Model of Shape and Motion for Tracking Finger Spelt American Sign Language, Image and Vision Computing, vol 20/9-10, pp597-607, Aug 2002, Elsevier Science Ltd
[10] Bowden R, Sarhadi M, Building Temporal Models for Gesture Recognition, In Proc. BMVC'00, M Mirmehdi & Barry Thomas Ed,

Combining PET and MRI

Richard Ansorge, Cambridge University

Positron Emission Tomography (PET) is has been around for at least 20 years but is only now beginning to come into widespread clinical use for oncology. PET is one of a number of "body scanning" techniques; others include CT, MRI and Ultrasound. One feature all these modalities have in common is the extensive computation required to convert the raw signals from the detectors into an image suitable for clinical use. Final images are typically three (or four) dimensional arrays of volume elements (voxels) and resolutions of $256 \times 256 \times 256$ are typical; thus the storage and display of images also presents interesting computational challenges. The steady growth of cheap computing power has been a key enabling feature for all these modalities; Beowulf PC clusters can now be found tucked away in odd corners of a surprising number of hospitals.

Most of the imaging modalities reveal the anatomy of the subject; X-rays are superb for bones while MRI is excellent for soft tissue discrimination. MRI can also provide some limited functional information. PET however is unique in imaging function rather than structure. The most common tracer used is an ^{18}F labelled glucose analogue, florodeoxyglucose or FDG, this accumulates in sites of active metabolism and hence is extensively used to find secondary cancers. Unfortunately the spatial resolution of PET is much worse that CT or MRI. Recently combined PET-CT scanners have come into use which offer a way round the limited resolution of PET. However because CT used X-rays only a single CT scan can be done at the start of an examination but each of each PET scan will take many minutes during which the patient may move.

Combining PET and MRI is an attractive alternative but the high magnetic fields required for MRI cause major problems for the detectors required in PET. A collaboration between the University of Cambridge Cavendish Laboratory and Wolfson Brain Imaging Centre has now developed a unique spilt-coil MRI magnet which in principle overcomes this difficulty. The group headed by Dr Adrian Carpenter and Dr Richard Ansorge has recently obtained a new grant from EPSRC which will allow them to build the world's first combined PET-MR system with a high sensitivity multi-ring PET detector. A unique feature of the instrument is that it will be capable of simultaneously imaging with both PET and MRI. Some further details are available on http://www.pet-mr.com

Upcoming Computational Physics Group Events

A Gentle Introduction to Biological Modelling

The Computational Physics Group of the Institute of Physics is holding a one-day workshop on the modelling of biological systems on Thursday, 15 September 2005 at the Institute of Physics, 76 Portland Place, London. This meeting is targeted at physicists who have little or no knowledge about biological modelling but who wish to know more. There will be four talks, each one hour long, covering the following topics

1. The biology of the cell
2. The biology of proteins and DNA
3. Modelling proteins and membranes
4. Bioinformatics

Web page: http://conferences.iop.org/GIM/

The 4^{th} Annual Computational Physics Thesis Prize

The Committee of the Institute of Physics Computational Group has endowed two annual prizes. £500 will be awarded to the author of the PhD thesis that contributes most strongly to the advancement of computational physics. Two runners-up will receive £250. There will be free group membership for 2006 for *all* entrants. The Committee will select the recipients and its remit will be very broad, in order to capture a broad spectrum of modelling activity.

- The deadline for applications is December 31^{st}, 2005. The competition is open to all students whose PhD examination has taken place in 2005.
- The submission format is a 4 page (A4) abstract together with a citation (max. 500 words) from the PhD supervisor.
- The submission address is:
  Dr M Probert
  Department of Physics

University of York
York, YO10 5DD

Applicants must have carried out their thesis work at a University in the United Kingdom or the Republic of Ireland.

The following were the winners of this year's competition:

- 1st place:
- 2nd place:
- 3rd place:

Other Upcoming Meetings

The 11th International Conference on the Applications of Density Functional Theory in Chemistry and Physics

The 11th International Conference on the Applications of Density Functional Theory in Chemistry and Physics (DFT2005) which be held September 11-15, 2005, in Geneva, Switzerland. This edition succeeds previous such biennial events organized in Brussels (2003), Madrid (2001), Rome (1999), Vienna (1997), ... The DFT2005 Conference will be devoted to both fundamental and applied aspects of Density Functional Theory (DFT) in Chemistry and Physics. The main topics covered by the Conference will be recent advances in:

- Fundamental aspects of DFT
- New exchange-correlation and kinetic energy functionals
- Conceptual DFT
- Weak interactions
- Magnetic properties
- Photochemical and photophysical properties
- Reacrivity

- Materials
- Biosystems
- Dynamics

All the informations concerning Conference program, Committees, Registration, Accommodation, etc., may be found on the web site.

Web page: http://DFT2005.unige.ch