

SPIE's
International
Technical
Group
Newsletter

Electronic Imaging

Automatic crowd scene generation

Animation has always been a labor-intensive task. In its current stage, computer-assistance frees the individual from frame by frame animation of characters, by allowing animators to specify key frames and leaving the software to interpolate between them. This has prompted developers to look for further advances in automatic animation in the form of enhanced character autonomy. A good test case is the ability to generate automatic crowd scenes in which characters stroll and mingle without collisions.

One approach for crowd-scene generation is to give the characters some minimal intelligence, together with the means to sense the virtual world that surrounds them and actuators through which they can interact with objects they encounter (see Figure 1).

The only information we have to provide is then the world model and the starting parameters for each character forming part of the crowd. Each character is granted up-to-date knowledge of its surroundings, including the current position of other characters. A longer-term ambition is that characters be equipped with their own perceptual systems with which to update their personal knowledge base.

By having awareness of the environment, much longer and more complex scenes can be automated. Having been given initial goals, characters can take their own decisions about how to react to obstacles based on a set of rules designed to generate realistic behavior. This calls for regular sensing and modification of plans (a list of actions culminating in a desired goal state) to take account of changing and unanticipated situations.

In our project at Kingston University, also involving WhiteSpace Studio and Cinesite, we are building an AI (Artificial Intelligence) engine to provide the necessary goal and planning mechanisms for such scenes. A key fea-




Figure 1. Shot from one of the automatic animations in which we see a group of characters walking around on a New York City street.

ture of our approach is that all modelling is done using the Unified Modelling Language (UML), a standard controlled by the Object Management Group and soon to become an ISO standard. This applies not only to the modelling we have done ourselves of the virtual world and the characters within it, but also by the characters as they construct their own plans and models of the environment.

More complex scene generation is now pos-

sible allowing the animator to concentrate on other creative aspects of the work.

Juan Amiguet-Vercher

Research assistant
WhiteSpace Studios, Grove House
Kingston University, Penrhyn Road
Surrey KT1 2EE, United Kingdom
E-mail: K995263@kingston.ac.uk
<http://www.kingston.ac.uk/whitespace>

Newsletter now available on-line

Technical Group members are being offered the option to receive the Electronic Imaging Newsletter in an electronic format. An e-mail notice is being sent to all group members advising you of the web site location for this issue and asking you to choose between the electronic or printed version for future issues. If you have not yet received this e-mail message, then SPIE does not have your correct e-mail address in our database. To receive future issues of this newsletter in the electronic format please send your e-mail address to spie-membership@spie.org with the word EI in the subject line of the message and the words "Electronic version" in the body of the message.

If you prefer to continue to receive the newsletter in the printed format, but want to send your correct e-mail address for our database, include the words "Print version preferred" in the body of your message.

Efficient fitting of Non-Uniform Rational B-Spline surfaces using non-organized 3D data

The problem of approximating a surface from a cloud of non-ordered 3D points appears in many areas including computer vision, computer-aided design and object recognition. With advancements in technology allowing fast digitization of a large number of points on an object, there is a clear need for new methods for surface fitting that can handle large amounts of data in acceptable time and memory space. Today, the preferred surface modelling entity is NURBS (Non-Uniform Rational BSpline),¹ because it is the primary entity in modern CAD systems and it is widely supported by the modern standards for graphics (OpenGL) and geometric data exchange (IGES). Furthermore, trimmed NURBS are extensively used in industry.

Least-squares-fitting of NURBS surfaces to a non-organized point set is generally regarded to be difficult.² The process is considered to be computationally expensive, both in terms of computing time and memory storage, posing severe practical limitations on the number of data points and the number of surface control points that can be effectively handled. Also, the fitting problem is known to suffer from ill-conditioning, which essentially occurs when the positions of some control points are not defined sufficiently well by the data. As Figure 1 illustrates, this is particularly acute in situations involving knot insertion, when additional control points are needed to provide modelling flexibility in some regions (the central portion in this case), but this leads to instability in other regions (the outer portion) due to insufficient data density.

These are some of the reasons why the majority of the available packages demand triangulation of the input data as the necessary first step. The triangular mesh is then re-sampled to produce a uniform grid of points with the required properties. However, triangulation of non-ordered 3D data is in itself a non-trivial task.

Here we report the development of a new method to generate, or update, a NURBS model using a set of non-organized 3D points. The main

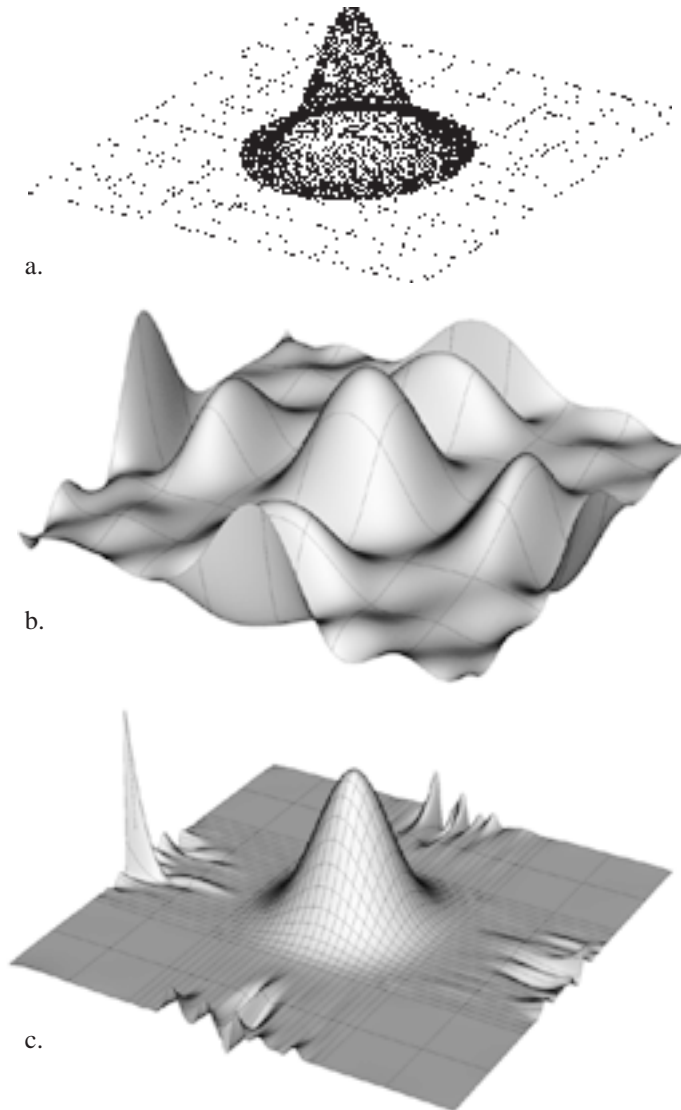


Figure 1. Surface fitting and knot insertion: (a) initial surface and unevenly distributed data cloud; (b) low flexibility of the fitted surface incapable of representing the shape; (c) knot insertion improves flexibility, allowing the surface to follow the shape, but causes instability in the sparsely measured outer regions.

features of the approach are the following:

- Solves the problem of ill-conditioning through regularization, by adopting additional fitting criteria
- Exploits banded-matrix-based algorithms to gain in computational efficiency
- Eliminates the need for pre-processing/triangulation and directly generates NURBS

Iterative NURBS fitting

The first step in the fitting process is to adopt a suitable base surface, from which parameterization of the data points can be derived, together with the knot vectors and the weights. This allows the setting up of a linear optimization problem with control points as the unknowns. We propose using a pre-defined CAD model for this purpose, since there are numerous applications in industrial manufacture when such a model is readily available.³ Alternatively, the base surface may be generated automatically. If data is single-valued (for example a single-view range image) then a rectangle can be employed, while if the surface is of a closed, tubular geometry, then a generalized cylinder⁴ can be used. The flexibility of the base surface can be further increased, if necessary, by knot insertion during the fitting procedure.

In developing the solution for the ill-conditioning problem, it was noted that when the system becomes unstable, the control points associated with the areas with insufficient data move in an uncontrollable fashion, away from the surface. Our principal concept is based on the fact that control points should approximate the surface, and this is achieved by introducing an additional “a-criterion” which minimizes the sum of the squared distances between the control points and their corresponding surface points, given by the Greville points.¹ This solution was found to produce very good results, particularly when large deformations of the base surface are needed. However it does not strictly guarantee that the fitting problem is well-posed and for this reason an additional “b-criterion” was introduced that minimizes the combined movement of all control points. In the proposed regularization, the quality of the fitted surface may be controlled by a trade-off between the weights a and b.

In attempting to improve the computational

continued on p. 6

Defect detection in plastic lids by active infrared thermography

Here we present a machine vision setup for automatic detection of glue occlusions inside plastic lids by active infrared thermography.¹ The sample (a lid) being tested is made of two parts (see Figure 1). The top part is supposedly glued onto the bottom part (see Figure 1a) which is screwed onto the cream pot. As shown in Figure 1b, the glue spots are sandwiched between the top and the bottom part of the lid. If there is not enough glue, the two parts can fall apart while unscrewing or screwing them from the pot.

Previously the control procedure consisted of weighing each cap. According to this measurement, the presence or the lack of glue could be inferred. However, the glue weight is in some cases lower than the weight variation of the plastic lid and as much as 22% of lids with glue on them could be rejected using this method.

In order to reduce this high rate of misclassified products, a prototype machine vision system was developed based on active infrared thermography. This system involves rapidly and uniformly heating the surface of the lid using a flash lamp (150W), then visualizing and finally analyzing the cap with an IR imager, hooked up to the PC via a frame grabber. The schematic of such a setup is presented in Figure 2.

Different attempts were made to find the most suitable time for heating, and then the delay time before imaging the lid. A 5-8 second heating time is optimal, and a further 8s is required before imaging in order to obtain the best thermal contrast.²


Figures 3a-c are images of three different caps, where dark areas signify the presence of glue. Indeed, whenever a glue spot is present, it creates a heat sink that can be seen from the surface layer. These images are useful and could be sufficient for an operator to control the process. However, in order to fully and automatically control the process (without any operator involvement), an autonomous machine vision system with the following input parameters was developed

As seen in figures 3a-c, the thermal images are noisy, the luminance is not homogenous over the whole image, and glue spots—which are as small as few square millimeters—can be found anywhere. Various processing methods were applied to the raw thermal images. Firstly, to remove the nonhomogeneous background, a reference image of a cap without any glue spots is subtracted from each analyzed image.

On the resulting images, a segmentation using Wen's threshold and a filtering process (to remove the noise) are successively implemented. Images 4a-c present the effect of the filter after the segmentation and clearly show the efficiency of the correction. Once the images are processed, a simple count of the pixel pertaining to the glue is realized. As a prototype test, this process was applied on a test set, and all the lids without any glue were detected. Moreover, no misclassification occurs during the process classification.

The main advantage of this method is the low processing time required to analyze and classify the images.

We are also currently working on a set of images that represent the thermal radiation of the cap surface as a function of time, in order to infer the glue spot



a.



b.

Figure 1. a) Side view of the lid. b) Top view of the lid.




Figure 2. Schematic of the experimental setup.

continued on p. 7

New pixel structure for full-color OLED

By having different colors on the faces of a single pyramid-shaped pixel, organiclight-emitting-diode (OLED) arrays may be easier and cheaper to manufacture: especially for large displays. According to researchers at the University of California, Los Angeles, the new structure has the advantage of allowing different colors to be deposited without the use of external shadow masks, thus reducing fabrication costs. Also, the full-color pixels can be smaller and more numerous than conventional red, green and blue triads, which means that higher resolution is possible than with conventional devices. Work is still at an early stage, but initial experiments have been successful.

The basic idea is very simple.¹ First, thin film transistor (TFT) addressing electronics are deposited on a transparent substrate. Next, plastic molding techniques are used to create pyramid structures over the TFT electrical contact pads, with openings for the electrical connection to be made. After that, the next step is to deposit the red, green, and blue OLEDs on the different faces of the pyramids.

To do this, the indium tin oxide (ITO) transparent electrodes must first be deposited on all three faces. For this to work properly in the final device, it is important that there be no electrical contact between the layer on one face and that on the next. UCLA scientists Yang Yang and Shun-Chi Chang have designed a pyramid structure that contains walls or ridges along its edges in order to ensure this isolation. After the ITO has been laid down, two more layers are necessary. First is the hole-transport layer, which is again deposited on all three faces, and next is the electron transport layer or ETL.

This ETL layer—through which the light is actually emitted—is crucial because the material used determines what color is

emitted by that face of the pyramid. It can be deposited selectively without masking simply by having the substrate held at the right angle within the deposition chamber. This discrimination is achieved because the ETL is only de-

posited within the line-of-sight of the vapor source. Thus, each set of faces can be deposited with a different material without affecting the others.

So far, the UCLA team have produced two demonstrations of the technology. The first showed that the selectivity of the material deposition could be made to work in practice. Though they were not able to use pyramid structures (because the necessary substrates were not available) they were able to demonstrate the self-masking principle using ridge-shaped pixels (see Figure 1).

The second demonstration involved showing how the light from the red, green and blue faces could be combined to produce a desired output color. A large pixel was created by putting polymer LEDs inside a an optical prism (see Figure 2). The white emission was obtained by choosing the PLEDs to get the best possible coverage of the CIE chromaticity diagram (see Figure 3).

For further information please see:

<http://www.seas.ucla.edu/ms/faculty1/yang-yang.html>

Sunny Bains

References

1. Yang Yang and Shun-Chi Chang, *Pyramid-shaped pixels for full-color organic emissive displays*, *Applied Physics Letters* **77** (7), pp. 936-938, 14 August 2000.

SUNNY BAINS is a scientist and journalist based in London, UK.

<http://www.summybains.com>




Figure 1. a) A thin film of aluminum (III) tris(8hydroxyquinolate)—known as Alq_3 —was selectively deposited on one side of the ridge structures shown. b) The substrate was cut with a razor blade and the photoluminescence from the Alq_3 layer was captured using



Figure 2. A mock up of the pyramid-shaped pixel structure was created by gluing polymer LEDs onto the inside of this prism (a scaled-up version of the design shape).

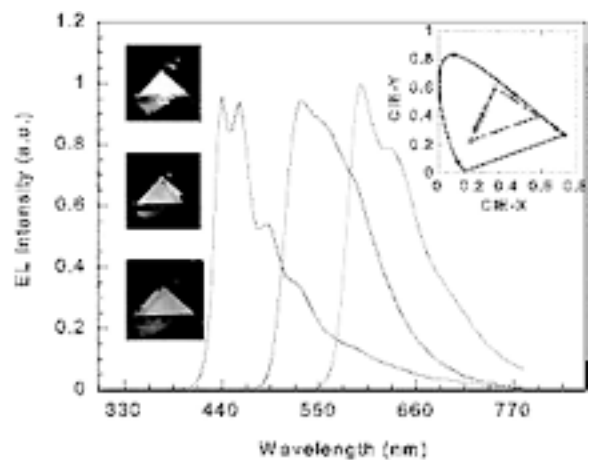



Figure 3. The spectra of the three LEDs used in the mock-up shown in Figure 2, and what they cover on the chromaticity diagram.

Neural network learns optimal feature set: sees what it wants to

Researchers in London have developed a new kind of neural network for vision and other machine understanding tasks. The technique has the advantage of first separating the problem into perception and recognition, which means the neural network can be trained with less outside intervention. Also, the perception network developed is both biologically valid and much more efficient than its predecessors. The Product of Experts approach is lead by Geoffrey Hinton at the Gatsby Computational Neuroscience Unit, University College London. If the work is successful, the network models produced may not only help us to build machines that can see, but to understand our own vision systems better.

Neural networks are systems where the processing and storage of information are performed together. In engineering terms, conventional networks can simply be thought of as complex filters that take incoming signals A, B, and C (which could, for instance, be images of faces) and give out the right answers X, Y, and Z (which could be names) respectively. They have major advantages in real-world applications like face recognition because A, B and C cannot be known exactly in advance: faces change with age, lighting, etc. and no person will present the exact same image to a camera twice. Unlike more conventional algorithmic systems, neural networks reconfigure themselves based on incoming data: they learn that various different images map to X and find their similarities, while at the same time determining how they are different from those images that map to Y. This learning is what makes them so powerful.

Structurally, neural networks are both simple and complicated. They consist of a number of neurons or processing elements that sum and then perform some function (like the sigmoid function) on incoming data. In the most common type of backpropagation network, the first layer of neurons take their information from the outside world: in an image processing application, this means that each neuron



1. A POE network with a hundred hidden units was trained on 16×16 images of the number two. Shown are the interconnection weights to the hidden units. Sets of weights such as those in column 1, row 7, look like edge detectors, where others may look like Gabor filters etc.. Hinton suggests that the best way to think of these are as ways to deform incoming images and make them more like some idealized digit that the network believes in.

looks at the signal coming in from a single pixel. The next layer, the hidden layer, can have any number of neurons, and each of these are connected to all of the pixels in the first. These neurons are then connected to the output. During training, after the image signals have propagated through the network, being processed by the neurons and attenuated/amplified by the interconnection weights, the “answer” is compared with the label already assigned to the data. The network learns by changing the weights or strengths of the various neural connections to minimize the difference between these two, and it is how well this process works that makes one neural network architecture or configuration better than another.

Overworked supervisor

The problem with this basic approach is that supervised learning (teaching) is required: the neural network is shown lots of examples of objects that have been labelled in advance, and

so eventually begins to associate the input (face) with the label (name). It is known as response learning because it directly links the inputs with the outputs. From a practical point of view, it is inefficient because, not only is lots of training data necessary in order to fully represent the “fuzziness” of the problem at hand, but all that training data has to be labelled somehow: presumably by a human.

Hinton and his colleagues¹ have chosen to concentrate on another approach: perceptual learning. Instead of concentrating on recognizing, images it’s job is simply to learn to be good at perceiving a given type of data without (at this point) assigning any meaning to it. An example of this kind of learning in humans is the fact that people brought up to speak different languages are better at distinguishing between different sets of sounds, even outside the context of a meaningful word or sentence. All the perceptual neural net does is to decompose the data into a particular combination of features: the better adapted the feature set, which is stored in the hid-

den unit interconnection weights, the more efficient and accurate the network will be for that problem.

By using this perceptual learning as a “front end” to a pattern recognition system, the second part—the response learning—becomes a more tractable problem. If the perceptual network has done its job well, a class of objects should now be represented by a relatively small number of feature combinations (compared to the number of images that went into defining the features). Essentially, because the fuzziness of the original images has been encapsulated in the hidden units, much less labelled training data should be necessary for the second stage.

Finding a cause

One way of determining whether a network is good at perceiving incoming data is to look and see what kind of data it would generate. So-

continued from p. 8

Efficient fitting of Non-Uniform Rational B-Spline

continued from p. 2

performance of the fitting algorithms, it was noticed that the sparse banded structure of main matrices can be exploited in order to achieve substantial savings. Table 1 provides an indication of the realized performance in terms of computational speed and memory requirements (Pentium III, 400MHz).

Examples

Figure 2 shows how rectangular base surface was employed for the reconstruction of the Igea artefact, while Figure 3 demonstrates the use of a cylindrical base. Finally, Figure 4 demonstrates the use of a pre-defined trimmed CAD model for the reconstruction of a car windscreen. The quality of the results and the realized processing speed are highly encouraging, leading to the conclusion that NURBS surface fitting can be readily employed as a practical tool in a wide range of applications.

Mihailo Ristic

Imperial College of Science,
Technology & Medicine
Department of Mechanical Engineering
Exhibition Road
London SW7 2BX
United Kingdom
Phone: +44 20 7594-7048
Fax: +44 20 7584-7239
E-mail: m.ristic@ic.ac.uk

References

1. G.E. Farin, *Curves and surfaces for computer-aided geometric design: a practical guide*, 4th ed., Academic Press, San Diego, London, 1997.
2. P. Dierckx, *Curve and surface fitting with splines*, Oxford, Clarendon, 1993.
3. M. Ristic, D. Brujic, and I. Ainsworth, *Precision Reconstruction of Manufactured Free-Form Components*, *Proc. SPIE 3966*, January 2000.
4. M. Kim, E. Park, and H. Lee, *Modelling and Animation of Generalized Cylinders with variable radius offset space curves*, *J. Visualization and Computer Animation 5*, pp. 189-207, 1994.

Table: Computational performance in terms of time (seconds) and memory (Kb).

Number of data points	Number of control points			
	10	100	1000	10000
100	0.03s 42.6Kb	0.09s 146 Kb	—	—
1,000	0.08s 42.6Kb	0.14s 146 Kb	0.78s 1054Kb	—
10,000	0.62s 42.6Kb	0.69s 146 Kb	1.39s 1054Kb	7.2s 9126Kb
100,000	6.17s 42.6 Kb	6.20s 146 Kb	6.97s 1054Kb	13s 9126Kb
1,000,000	61.8s 42.6 Kb	61.5s 146 Kb	63.0s 1054Kb	69s 9126Kb




Figure 2. Reconstruction of Igea artefact: (a) data and planar base surface; (b) result of knot insertion and fitting; (c) trimmed surface. (Data courtesy of University of Thessalonica, available from <http://www.cyberware.com>).




Figure 4. Fitting of a trimmed NURBS CAD model surface.




Figure 3. (a) Measured points and base surface, (b) smooth B-spline surface after fitting, (c) triangulation of data points gives inferior smoothness.

Defect detection in plastic lids by active infrared thermography

continued from p. 3

size by an inverse method. This kind of method has been used in pulsed infrared thermography.³ However, in our case, the glue spots are embedded at the same depth but may have different sizes. The current study could also be applied to control procedures in the plastic industry, where the flaws to be detected have a thermal response different from the sound material (such as air bubbles trapped in plastic, cracks, etc.).

F. Meriaudeau, A.-C. Legrand, and P. Gorria

Le2i, 12, rue de la Fonderie
71200 Le Creusot
Phone: +33 3 85 73 10 90
Fax: +33 3 85 73 10 97
E-mail: {f.meriaudeau,
ac.legrand}@iutlecreusot.u-bourgogne.fr
[http:// www-le2i.u-bourgogne.fr](http://www-le2i.u-bourgogne.fr)

References

1. Herbert Kaplan, **Practical applications of infrared thermal sensing and imaging equipment (SPIE Press TT34)**, 2nd edition, 1999.
2. A.-C. Legrand, P. Gorria, and F. Meriaudeau, *Infrared vision system for glue detection in plastic lids*, **Proc. Inframation**, Orlando, pp. 205-210, September 2000.
3. S. Vallerand and X. Maldague, *Defect characterization in pulsed infrared thermography: a statistical method compared with Kohonen and Perceptron neural networks*, **NDT&E International 33**, pp. 307-315, 2000.




Figure 3. Thermal images of: a) and b) lids with glue occlusions; c) lid without glue occlusions.




Figure 4. Thermal images after the image processing: a) and b) lids with glue occlusions; c) lid without glue occlusions.

Neural network learns optimal feature set: sees what it wants to

continued from p. 5

called generative models, where the neural network is essentially run backwards and the hidden units (features) are stimulated to produce their own "input" have the advantage that they can show what a network "believes in".

Unfortunately the more powerful nonlinear generative models have traditionally been difficult to work with. They fall into two classes, both of which have their disadvantages. The first, known as causal models, can be compared to computer graphics. Though it is easy to generate pictures from, say, a 3-D model, it is not easy to reconstruct the 3-D model from this data. In fact, this is why machine vision is so difficult in the first place. There are all sorts of ambiguities: what goes with what (image segmentation), how big are objects (as opposed to how far away they are), where is occlusion taking place. With such models, it is easy to generate images, yet difficult to extract meaning from those images afterwards.

Hinton's approach is called the Product of Experts (POE), and had always been thought to have the opposite problem. Here, in order to generate a dream image from scratch, every hidden unit on must agree on every pixel. To achieve this, each hidden unit is either active or not based on its own statistics and, if on, it then "votes" on whether each pixel should be on or off based on the relevant interconnect weights. If the vote is not unanimous, the dice are effectively rolled again to select which hidden units should be active until a set is found that can agree on what image should be produced.

Being practical

This is less impossible to achieve than it sounds because not all of the hidden units care about every pixel in the image: they specialize on detecting certain things and ignoring others. Therefore, some features will be complimentary to others and will be able to "co-exist". On the other hand, it is still very difficult to find those winning combinations in the first place, which is why such models were thought to be impractical.

What Hinton realized was that there is no need to start from scratch. It is actually relatively easy to figure out how to do it if you have the answer in front of you (this, in fact, is why such systems are good at inferring meaning from dream data even though they are not good at generating it). Hinton was able to exploit this in the learning algorithm for a Product of Experts system he calls a Restricted Boltzmann Machine (RBM). Data comes in



2. A demonstration of how important well-adapted feature detectors are to pattern recognition. The center row shows a set of 2s that were previously unseen by two networks, one trained on 2s, the other trained on 3s. The top row demonstrates how the network that "believes in 2s" can regenerate the numbers reasonably well (and so will be able to recognize them easily). The network trained on 3s, on the other hand, keeps trying to "fix" the 2s to look like the number it's expecting.

through the input, which excites various units (features) in the hidden layer. These are then used to generate an image, which is compared with the original data to produce an error signal, which in turn is used to update the interconnections weights between the input pixels and hidden units.

This way, with each new piece of data, the feature set is refined to more perfectly generate images like that in the training set. Effectively, the neural network is accepting the excited hidden units, caused by the data, as the "answer" to the question "which features should I use to generate this image," and then trying to optimize the features that have been turned on to create a better image next time. Over time, the system efficiently creates a set of optimized features that can accurately regenerate the training data, and that allows easy inference between generated image and feature combination.

Biological validity

In experiments where the technique was applied to both handwriting recognition (see figure 1) and face recognition, similar detectors to those found in our own early vision systems—such as those that look for lines or edges at different orientations—emerged naturally as hidden units. The fact that many of the feature detectors look alike despite being specialized (which is also true of human feature detectors), is particular to the POE approach. In causal models, the features or hidden units are independent and compete to effect the image: very similar features are therefore less likely to co-exist. In Hinton's networks, because features

work cooperatively, it is natural to evolve many that will effect the image in a similar way.

Other ways in which the Hinton's approach is biologically credible include its speed, and the fact that it doesn't require synapses (neural interconnections) to work backwards (as is necessary for backpropagation). In the POE system, the data, hidden unit, reconstruction, and weight modification operate in a loop. Hinton has patented the technique and is continuing its development.

This biological basis is important in the context of The Gatsby Computational Neuroscience Unit, as it was set up two years ago—with funding to last a decade—to study neural computation theories of perception and action with an emphasis on learning. The unit is funded by the Gatsby Charitable Foundation (named after the book by F. Scott Fitzgerald) established by Lord Sainsbury (who is both Science Minister in the British Government and a supermarket magnate).

For further information please see:
<http://www.gatsby.ucl.ac.uk/Hinton/>

Sunny Bains


References

1. Geoffrey E. Hinton, *Training Products of Experts by Minimizing Contrastive Divergence*, Technical Report GCNU TR2000-004, Gatsby Computational Neuroscience Unit, University College London, 2000.

SUNNY BAINS is a scientist and journalist based in London, UK.
<http://www.sunnybains.com>

HANDBOOK OF

Medical Imaging



The three-volume *Handbook of Medical Imaging* is a comprehensive treatment of the theory and current practice of medical imaging, one hundred years after the discovery of the x-ray.

SPECIAL SET PRICE • Volumes 1-3

Medical Imaging Handbook 3-Volume Set
\$255 (SPIE Mem.) / \$300 (Tech. Group Mem.) / \$315 (Nonmem.)

Volume 1. Physics and Psychophysics

Editors: Jacob Beutel, Harold L. Kundel, Richard L. Van Metter

Examines x-ray imaging physics and reviews linear systems theory and its application to signal and noise propagation. The chapters in the first half of the book address the physics of important imaging modalities now in use: ultrasound, CT, MRI, and the recently emerging flat panel x-ray detectors and their application to mammography. The second half describes the relationship between image quality metrics and visual perception of the diagnostic information carried by medical images.

SPIE PRESS Vol. **PM79** • 949 pages (*hardcover*) • February 2000
\$90 (SPIE Mem.) / \$105 (Tech. Group Mem.) / \$110 (Nonmem.)

Volume 2. Medical Image Processing and Analysis

Editors: Milan Sonka and J. Michael Fitzpatrick

Addresses the methods in use or in development for enhancing the visual perception of digital medical images obtained by a wide variety of imaging modalities and for image analysis as an aid to detection and diagnosis.

SPIE PRESS Vol. **PM80** • 1,108 pages (*hardcover*) • July 2000
\$105 (SPIE Mem.) / \$120 (Tech. Group Mem.) / \$130 (Nonmem.)

Volume 3. Display and PACS

Editors: Yongmin Kim and Steven C. Horii

Describes concurrent engineering developments that affect or are expected to influence future development of digital diagnostic imaging. Also covers current developments in Picture Archiving and Communications System (PACS) technology, with particular emphasis on integration of emerging imaging technologies into the hospital environment.

SPIE PRESS Vol. **PM81** • 750 pages (*hardcover*) • October 2000
\$90 (SPIE Mem.) / \$105 (Tech. Group Mem.) / \$110 (Nonmem.)

See SPIE Bookstore for full tables of contents and online ordering:

www.spie.org/bookstore/

Phone (1)360/676-3290 • Fax (1)360/647-1445

E-mail bookorders@spie.org • Web: www.spie.org/bookstore/

CALENDAR

2001



Photonics West

20-26 January

San Jose Convention Center
San Jose, California, USA

Including international symposia on:

- LASE '01—High-Power Lasers and Applications
- OPTOELECTRONICS '01—Integrated Devices and Applications
- BIOS '01—International Biomedical Optics Symposium
- SPIE/IS&T's EI '01—Electronic Imaging: Science and Technology.

Education Program and Short Courses

Technical Exhibit: 20-21, 23-25 January



Medical Imaging

18-23 February

San Diego, California, USA

Technical Exhibit: 20-22 February



Optical Science and Technology

SPIE's Annual Meeting

29 July - 3 August

San Diego Convention Center
San Diego, California, USA

Technical Exhibit: 31 July-2 August

For More Information Contact

SPIE • PO Box 10, Bellingham, WA 98227-0010

Phone (1) 360/676-3290 • Fax (1) 360 647-1445 • E-mail spie@spie.org • Web www.spie.org

Configuration assistant

continued from p. 12

shown to be very efficient in terms of the average time needed to reach a score.

The developed assistant derives great benefit from its ability to provide near-optimal solutions in a short period of time. Activated in association with a learning set of example circuits, it quickly delivers near-optimal solutions that are selected for production inspection. An additional benefit of the proposed configuration assistant is the availability of a score value, which gives a hint of its confidence in the configuration.

Olivier Hüsser and Heinz Hügli

Institut de Microtechnique
Université de Neuchâtel
Rue A.-L. Breguet 2
2000 Neuchâtel, Switzerland
<http://www-imt.unine.ch/>

Reference

1. O.Hüsser and H. Hügli, *A configuration assistant for versatile vision-based inspection systems*, **Proc. SPIE 3966**, pp. 259-269, February 2000.

Special
On Site Prices!

Continuing Education
the way YOU want it.
Convenient! • Comprehensive! • Searchable!
Choose the format that works for you!

Can't take the time to attend live short courses? SPIE has over 90 courses on video and CD-ROM, and new for 2001, online courses taught by world-renowned instructors, with new releases scheduled throughout the year! Short courses on video or CD-ROM and now online, include optics, optoelectronics, and imaging. Learn at your own pace. Discover new approaches and problem-solving skills; enhance your in-house training and improve your overall performance.

Easily accessed from your desktop, SPIE's online courses and CD-ROMs run on any computer with a CD drive, using a web browser to access information. Videos are available in NTSC or PAL formats, please specify when ordering. Each course comes with a set of course notes or textbook.

Make the most of your valuable time, and gain the competitive advantage you need in the ever-changing world of optics!



Order today!

Call 1-360-676-3290 for customer service; send email to distance-ed@spie.org

or order online at <http://www.spie.org/info/education>

Shipping and handling fees apply. SPIE members, take 10% off the list price.



Join the SPIE/IS&T Technical Group

...and receive this newsletter

This newsletter is produced twice yearly and is available only as a benefit of membership in the SPIE/IS&T Electronic Imaging Technical Group.

IS&T—The Society for Imaging Science and Technology has joined with SPIE to form a technical group structure that provides a worldwide communication network and is advantageous to the memberships.

Join the Electronic Imaging Technical Group for US\$30. Technical Group members receive these benefits:

- *Electronic Imaging Newsletter*
- SPIE's monthly publication, *OE Magazine*
- annual list of Electronic Imaging Technical Group members
- discounts on registration fees for IS&T and SPIE meetings and on books and other selected publications related to electronic imaging.

Persons who are already members of IS&T or SPIE are invited to join the Electronic Imaging Technical Group for the reduced member fee of US\$15.

Please Print Prof. Dr. Mr. Miss Mrs. Ms.

First (Given) Name _____ Middle Initial _____

Last (Family) Name _____

Position _____

Business Affiliation _____

Dept./Bldg./Mail Stop/etc. _____

Street Address or P.O. Box _____

City _____ State or Province _____

Zip or Postal Code _____ Country _____

Phone _____ Fax _____

E-mail _____

Technical Group Membership fee is \$30/year, or \$15/year for full SPIE and IS&T Members.

Amount enclosed for Technical Group membership \$ _____

I also want to subscribe to IS&T/SPIE's *Journal of Electronic Imaging (JEI)* \$ _____
(see prices below)

Total \$ _____

Check enclosed. Payment in U.S. dollars (by draft on a U.S. bank, or international money order) is required. Do not send currency. Transfers from banks must include a copy of the transfer order.

Charge to my: VISA MasterCard American Express Diners Club Discover

Account # _____ Expiration date _____

Signature _____
(required for credit card orders)

2447

JEI 2000 subscription rates (4 issues):	U.S.	Non-U.S.
Individual SPIE or IS&T member	\$ 50	\$ 50
Individual nonmember and institutions	\$215	\$235

Your subscription begins with the first issue of the year. Subscriptions are entered on a calendar-year basis. Orders received after 1 September 2000 will begin January 2001 unless a 2000 subscription is specified.

Send this form (or photocopy) to:
SPIE • P.O. Box 10
Bellingham, WA 98227-0010 USA
Phone: (1) 360/676-3290
Fax: (1) 360/647-1445
E-mail: membership@spie.org

Please send me:

- Information about full SPIE membership
- Information about full IS&T membership
- Information about other SPIE technical groups
- FREE technical publications catalog

Electronic Imaging

The *Electronic Imaging* newsletter is published by SPIE—The International Society for Optical Engineering and IS&T—The Society for Imaging Science and Technology. The newsletter is the official publication of the International Technical Group on Electronic Imaging.

<i>Technical Group Chair</i>	Arthur Weeks
<i>Technical Editor</i>	Sunny Bains
<i>Managing Editor/Graphics</i>	Linda DeLano
<i>Advertising Sales</i>	Roy Overstreet

Articles in this newsletter do not necessarily constitute endorsement or the opinions of the editors, SPIE, or IS&T. Advertising and copy are subject to acceptance by the editors.

SPIE is an international technical society dedicated to advancing engineering, scientific, and commercial applications of optical, photonic, imaging, electronic, and optoelectronic technologies.

IS&T is an international nonprofit society whose goal is to keep members aware of the latest scientific and technological developments in the fields of imaging through conferences, journals and other publications.

SPIE—The International Society for Optical Engineering, P.O. Box 10, Bellingham, WA 98227-0010 USA. Phone: (1) 360/676-3290. Fax: (1) 360/647-1445. E-mail: spie@spie.org.

IS&T—The Society for Imaging Science and Technology, 7003 Kilworth Lane, Springfield, VA 22151 USA. Phone: 703/642-9090. Fax: 703/642-9094.

© 2000 SPIE. All rights reserved.

Want it published?

Submit work, information, or announcements for publication in future issues of this newsletter by e-mail to Sunny Bains. For full submission information, length and format required, deadline, etc., please go to:

<http://www.sunnybains.com/newslet.html>.

All materials are subject to approval and may be edited. Please include the name and e-mail address of someone who can answer technical questions. Calendar listings should be sent at least eight months prior to the event.

EIONLINE

New Electronic Imaging Web Discussion Forum launched

You are invited to participate in SPIE's new online discussion forum on Electronic Imaging. The **INFO-EI** mailing list is being "retired" as we move the discussions to the more full-featured web forums. We hope you will participate. To post a message, log in to create a user account. For options see "subscribe to this forum."

You'll find our forums well-designed and easy to use, with many helpful features such as automated email notifications, easy-to-follow "threads," and searchability. There is a full FAQ for more details on how to use the forums.

Main link to the new Electronic Imaging forum: <http://spie.org/app/forums/tech/>
 Related questions or suggestions can be sent to forums@spie.org.



Configuration assistant for vision-based inspection

Nowadays, vision-based inspection systems are present in many stages of the industrial manufacturing process. Versatile vision systems are used to accommodate the broad range of inspection requirements but limitations appear due to the time-consuming system setup performed at each production change. This work aims at providing a configuration assistant that helps to quicken the lengthy manual system setup.

A versatile visual inspection system possesses a certain number of *measurement methods* (e.g. dimensions, positions, similarities, ...) and some associated *treatments* (e.g. decision, storage, alerts, ...). Some of them can be selected for customizing the system for a specific inspection task. This selection, called the *visible configuration*, is usually rather straightforward. Selecting the *hidden configuration*, which consists of a set of parameters that tune the selected methods and treatments for the custom application, is more difficult however. Choosing a good set of parameters is often a lengthy, tedious process during which the operator empirically tries to tune various dependant variables. The goal of the assistant presented here is to provide hidden configurations automatically.

A first aspect of the assistant concerns the performance rating of a specific hidden configuration. In essence, it uses a performance function and applies to a set of images S obtained from good and bad objects to be inspected. The performance function F takes positive values only when the inspection system performs without mistakes on S and has a magnitude that increases with the confidence in the decision.

A second aspect concerns the hidden configuration selection. It proceeds by maximizing F with




Figure 1. Marking images after preprocessing with different configurations.




Figure 2. Averaged score over 100 trials.

a search over permitted ranges for the parameters. The search requires smart and robust heuristics because the configuration space is typically high dimensional, and also because the performance function is unpredictable. Specifically, a simulated annealing algorithm and a genetic algorithm are considered for this task. The search is iterated until a stopping criterion (time limit, minimal F-value) is satisfied. Upon termination, it delivers the optimal configuration characterized by a set of parameters, as well as the associated F-value that provides information about its quality.

One application is to aid in the inspection process of markings found on top of molded integrated circuits. This process faces frequent change in the

circuits to be inspected and also large variations in the marking type and quality. Thus, system setup is essential and frequent. The application-visible configuration system uses various methods that generate two measurements and end up with a single accept/reject decision. The choice of a well-suited performance function has been studied and is described in Reference 1. The hidden configuration consists of as many as eight parameters: two thresholds for the decision, two thresholds for binarizing the image, three morphological parameters and one positional parameter. Figure 1 shows an example of an integrated circuit marking after preprocessing. The left image reflects the effect of a bad hidden configuration whereas a good one is reflected in the right image.

In this example, there are approximately 1013 possible hidden configurations! The huge size of the configuration space clearly rules out an exhaustive search. A search with intelligent methods, however, was shown to reliably find a near-optimal configuration

within a short time. Indeed, it requires only a small fraction of the time that would be necessary to try all configurations exhaustively (typically less than 1% in this application). Figure 2 shows the averaged score (normalized F -value) of configurations obtained with two search heuristics, namely a genetic algorithm (GA) and simulated annealing (SA), after a given number of trials (end criterion). For comparison, a curve representing a pure random search method is also given. Although all methods will eventually find the best configuration as a consequence of monotonicity of the score function, genetic and simulated annealing methods have been

continued on p. 10