popularising the science and researchers of PASCAL

Peter J Bentley

p.bentley@cs.ucl.ac.uk

http://www.cs.ucl.ac.uk/staff/p.bentley/

Communication is one of the most fundamental skills required by contemporary scientists.

Without appropriate and effective dissemination, even the very best research can be overlooked.

Yet, the frenetic world of 24-hour media and fast-paced journalism often precludes science in favour of the meaningless sound-bite.

Such are the barriers of communication, that it is extremely difficult for scientists to communicate their work beyond their specialist conferences and journals.

For those that break the barriers, the rewards are clear: popularised science that catches the imagination of others can create new opportunities.

An enthusiastic portrayal of successes can lead to growth in that field, as new researchers are attracted to it and new funders become keen to exploit the advances.

There are several routes to broader dissemination.

One is simply to expand traditional scientific outputs, inviting delegates from broader scientific fields to workshops, or producing web pages with “friendlier” interfaces.

However, science writers and journalists rarely attend scientific events, and with only a few hours to write their articles, rarely have the time to read a broad range of scientific publications.

Today, the media look for easy-to-digest “news stories” – examples of successes that have been proven through peer review, and which also have a broad impact and interest.

This project exploits the services of an experienced science communicator in order to find and publicise such stories, to improve the visibility of PASCAL research and researchers.

UCL


THE DESIGN CYCLE FORTNIGHTLY WWW.NEON.CO.UK
New Electronics
11 OCTOBER 2005
Evolutionary activity
Can digital biology help to create electronic devices that truly deserve to be called intelligent?
Special Report:
Expert Panel


Evolutionary activity
Can digital biology help to create truly intelligent electronic devices?
By David Boothroyd.


DIGITAL BIOLOGY
COVER STORY


DIGITAL BIOLOGY
COVER STORY

UCL

New Scientist | Space | Technology | Environment | New Scientist Jobs | Subscribe to New S

NewScientist SEARCH Tips

NEWS | EXPLORE BY SUBJECT | SPECIAL REPORTS | LAST WORD | SUBSCRIBE | BLOGS | VID

BREAKING NEWS

The World's No.1 Science & Technology News Se

LATEST HEADLINES

- 'Dark field' X-rays reveal bodies in new detail
- Has Messenger revealed lava flows on Mercury?
- Tremors keep crust-dwelling microbes alive
- Origami spaceplane aims for space station descent
- Invention: 3D tissue printer
- Coffee can double risk of miscarriage
- First subglacial eruption found in Antarctica
- Cosmic strings observed in background radiation
- ALL LATEST NEWS

PRINT EDITION

Subscribe

Fast cars could be tuned by evolution

16:03 18 June 2004
Exclusive from New Scientist Print Edition. Subscribe and get 4 free issues.
Will Knight

Can evolutionary theory help rival Formula 1 teams break Michael Schumacher's seemingly unassailable hold on top-flight motor racing? Computer scientists at University College London think it might.

Peter Bentley and Krzysztof Wloch have used genetic algorithms software that mimics evolution's drive for fitness to breed the best tuning configurations for racing cars.

Using the technique, ... achieved by an expe... conference on evolut...

Genetic algorithms m...

NewScientist SEARCH Tips

NEWS | EXPLORE BY SUBJECT | SPECIAL REPORTS | LAST WORD | SUBSCRIBE | BLOGS | VID

BREAKING NEWS

The World's No.1 Science & Technology News S

LATEST HEADLINES

- 'Dark field' X-rays reveal bodies in new detail
- Has Messenger revealed lava flows on Mercury?
- Tremors keep crust-dwelling microbes alive
- Origami spaceplane aims for space station descent
- Invention: 3D tissue printer
- Coffee can double risk of miscarriage
- First subglacial eruption found in Antarctica
- Cosmic strings observed in background radiation
- ALL LATEST NEWS

PRINT EDITION

Subscribe

Insect swarming inspires jazz software

08:00 06 May 2002
Exclusive from New Scientist Print Edition. Subscribe and get 4 free issues.
Duncan Graham-Rowe

Jazz musicians who enjoy freeform improvisation may soon be using computer software to accompany themselves. A team at University College London has written a program that mimics insect swarming to "fly around" the sequence of notes the musician is playing and improvise a related tune of its own.

The Swarm Music program is the creation of computer scientists Tim Blackwell and Peter Bentley, who study how natural processes can be modelled in software. The pair believe that improvised music is self-organising in the way swarms of insects and flocks of birds are.

Their software works by treating music as a type of 3D space, ...

NewScientist SEARCH Tips

NEWS | EXPLORE BY SUBJECT | SPECIAL REPORTS | LAST WORD | SUBSCRIBE

BREAKING NEWS

The World's No.1 Science & Technology News Service

LATEST HEADLINES

- 'Dark field' X-rays reveal bodies in new detail
- Has Messenger revealed lava flows on Mercury?
- Tremors keep crust-dwelling microbes alive
- Origami spaceplane aims for space station descent
- Invention: 3D tissue printer
- Coffee can double risk of miscarriage
- First subglacial eruption found in Antarctica
- Cosmic strings observed in background radiation
- ALL LATEST NEWS

PRINT EDITION

Subscribe

Robot spy can survive battlefield damage

PRINT SEND RSS FEEDS SYNDICATE

19:00 20 August 2003
Exclusive from New Scientist Print Edition. Subscribe and get 4 free issues.
Duncan Graham-Rowe

A military reconnaissance robot being developed at a British lab can keep moving even if it gets damaged on the battlefield. When any of the snake-like robot's "muscle" segments are damaged, clever software "evolves" a different way for it to wriggle across any terrain.

The serpentine spy is a research project funded by aerospace company BAE Systems to make a low-cost military robot that can be dropped out of helicopters to carry out reconnaissance missions. Because it is not wheeled, the low-profile, ground-hugging snakebot should make a versatile battlefield spy. The team behind it has also developed a shape-changing antenna that broadcasts high-quality video and audio.

Enlarge image
Anatomy of a snakebot

UCL


EVOLUTIONARY DESIGN BY COMPUTERS
EDITED BY PETER J. BENTLEY
"DARWIN WOULD LOVE THIS BOOK"
RICHARD DAWKINS
EDITED BY
J. BENTLEY & DAVID W. CORNE
The PhD Application Handbook
Peter J. Bentley

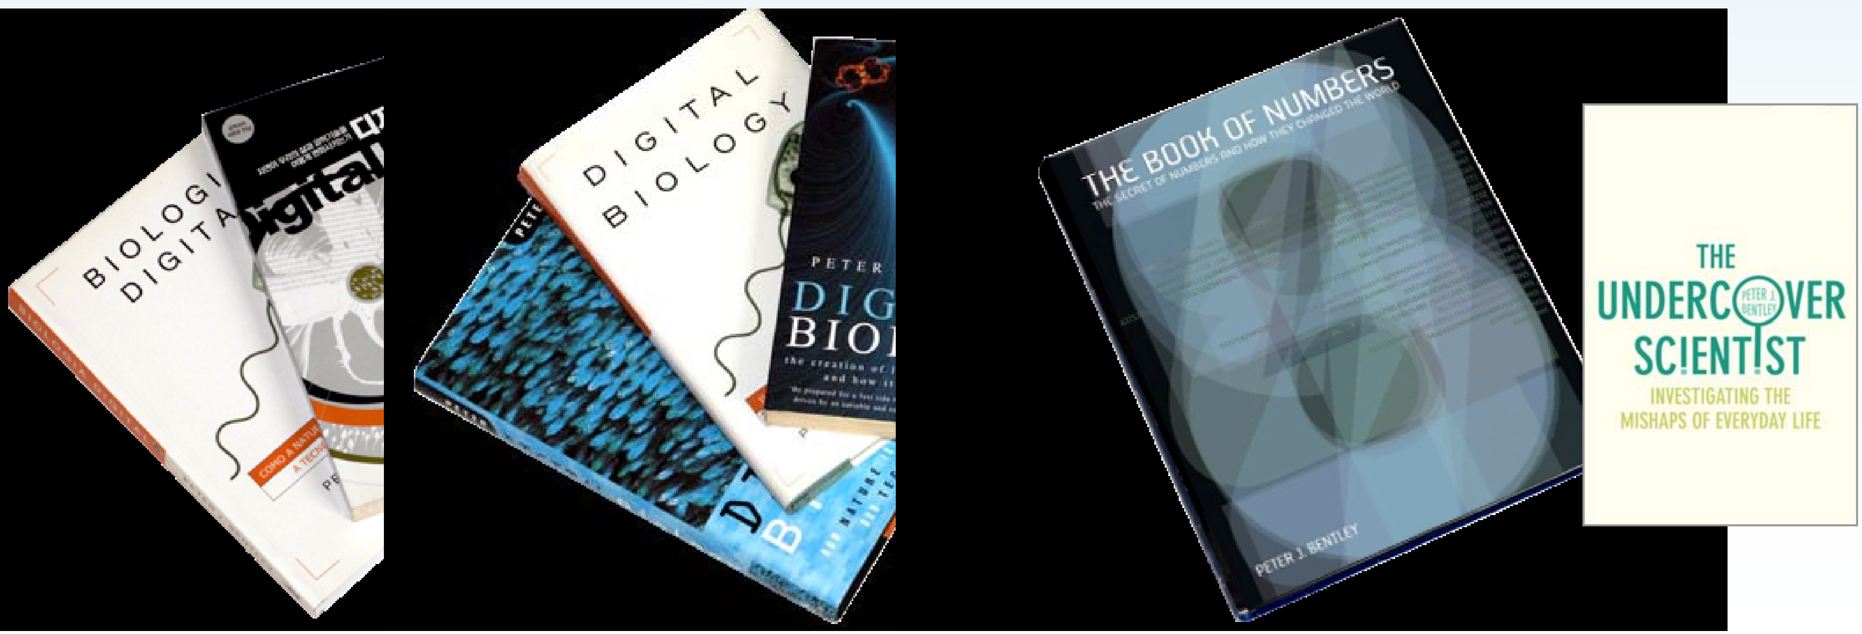
DIGITAL BIOLOGY
PETER
THE BOOK OF NUMBERS
THE SECRET OF NUMBERS AND HOW THEY CHANGED THE WORLD
PETER J. BENTLEY
THE UNDERCOVER SCIENTIST
PETER J. BENTLEY
INVESTIGATING THE MISHAPS OF EVERYDAY LIFE

UCL


Technology

The garden where perfect software grows

It's guaranteed bug-free and carries on working even if chunks of its code are lost. Nobody could write such a program, says software engineer **Peter Bentley**, but that doesn't matter because he has a way to make uncrashable code write itself

"In place of programmers there will be digital horticulturalists who plant, prune and grow software from seeds that they have cultured"

IN THE BEGINNING...

GROWING MACHINES...

"Conventional software may be provably correct but that doesn't mean it won't fall to pieces when faced with the unpredictable circumstances of the real world"

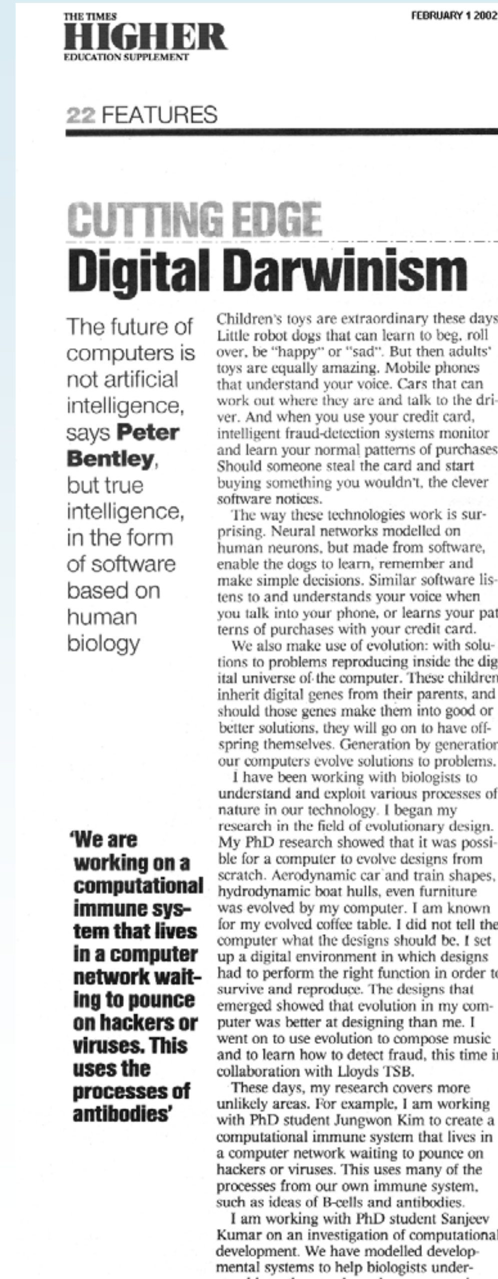
THE TIMES HIGHER EDUCATION SUPPLEMENT FEBRUARY 1 2002

CUTTING EDGE

Digital Darwinism

The future of computers is not artificial intelligence, says **Peter Bentley**, but true intelligence, in the form of software based on human biology

Children's toys are extraordinary these days. Little robot dogs that can learn to beg, roll over, be "happy" or "sad". But then adults' toys are equally amazing. Mobile phones that understand your voice. Cars that can work out where they are and talk to the driver. And when you use your credit card, intelligent fraud-detection systems monitor and learn your normal patterns of purchases. Should someone steal the card and start buying something you wouldn't, the clever software notices.

The way these technologies work is surprising. Neural networks modelled on human neurons, but made from software, enable the dogs to learn, remember and make simple decisions. Similar software listens to and understands your voice when you talk into your phone, or learns your patterns of purchases with your credit card.

We also make use of evolution: with solutions to problems reproducing inside the digital universe of the computer. These children inherit digital genes from their parents, and should those genes make them into good or better solutions, they will go on to have offspring themselves. Generation by generation, our computers evolve solutions to problems.

I have been working with biologists to understand and exploit various processes of nature in our technology. I began my research in the field of evolutionary design. My PhD research showed that it was possible for a computer to evolve designs from scratch. Aerodynamic car and train shapes, hydrodynamic boat hulls, even furniture was evolved by my computer. I am known for my evolved coffee table. I did not tell the computer what the designs should be. I set up a digital environment in which designs had to perform the right function in order to survive and reproduce. The designs that emerged showed that evolution in my computer was better at designing than me. I went on to use evolution to compose music and to learn how to detect fraud, this time in collaboration with Lloyds TSB.

'We are working on a computational immune system that lives in a computer network waiting to pounce on hackers or viruses. This uses the processes of antibodies'

These days, my research covers more unlikely areas. For example, I am working with PhD student Jungwon Kim to create a computational immune system that lives in a computer network waiting to pounce on hackers or viruses. This uses many of the processes from our own immune system, such as ideas of B-cells and antibodies.

I am working with PhD student Sanjeev Kumar on an investigation of computational development. We have modelled developmental systems to help biologists under-

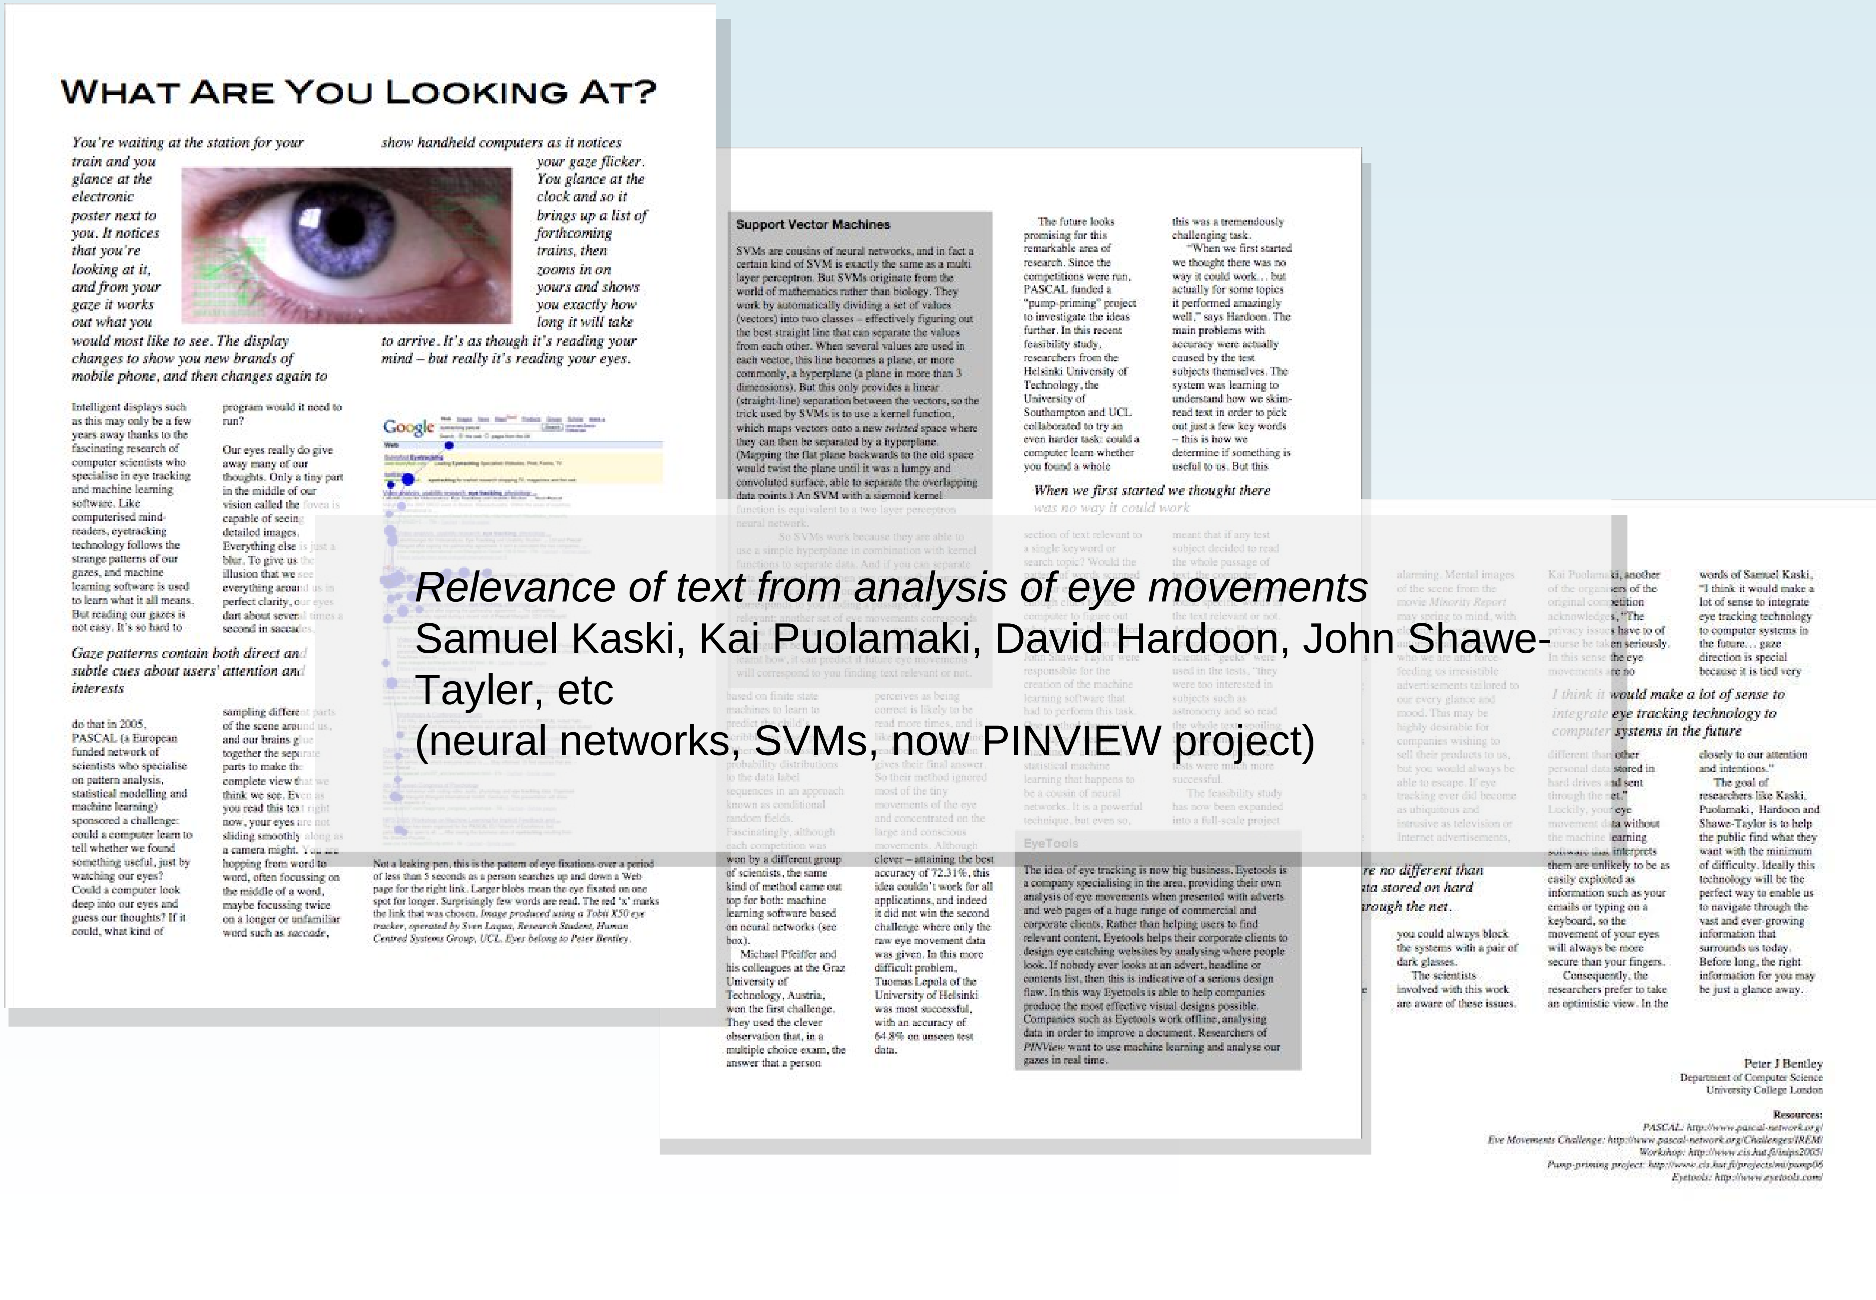

Relevance of text from analysis of eye movements

Samuel Kaski, Kai Puolamaki, David Hardoon, John Shawe-Tayler, etc

(neural networks, SVMs, now PINVIEW project)

Hex-o-Spell: analysis of EEG signals combined with novel Hex interface

Benjamin Blankertz, John Williamson, Roderick Murray-Smith, etc.

CSP (common spatial patterns) & LDA (linear discriminant analysis)

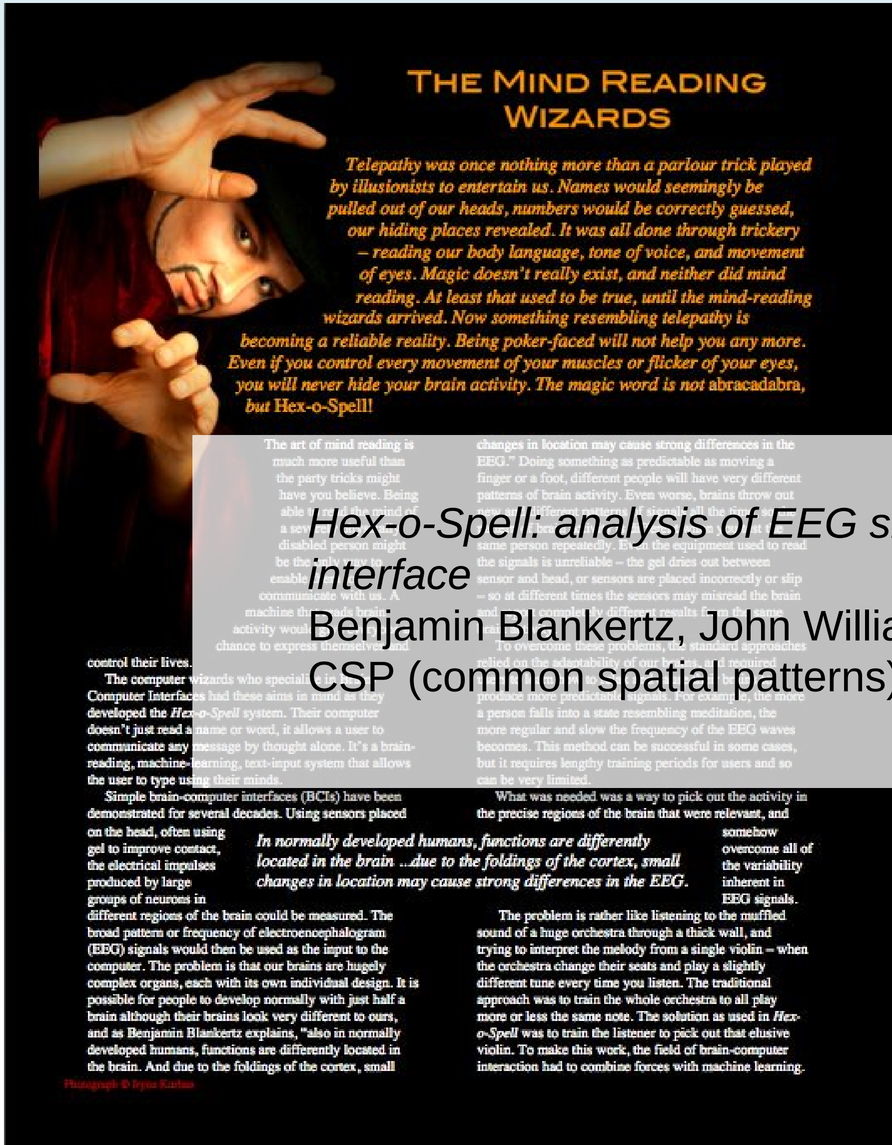

The Mind Reading Wizards

Telepathy was once nothing more than a parlour trick played by illusionists to entertain us. Names would seemingly be pulled out of our heads, numbers would be correctly guessed, our hiding places revealed. It was all done through trickery – reading our body language, tone of voice, and movement of eyes. Magic doesn't really exist, and neither did mind reading. At least that used to be true, until the mind-reading wizards arrived. Now something resembling telepathy is becoming a reliable reality. Being poker-faced will not help you any more. Even if you control every movement of your muscles or flicker of your eyes, you will never hide your brain activity. The magic word is not abracadabra, *but* Hex-o-Spell!

The art of mind reading is much more useful than the party tricks might have you believe. Being able to read the mind of a severely [illegible] disabled person might be the only way to enable [illegible] communicate with us. A machine that reads brain activity would [illegible] chance to express themselves and control their lives.

The computer wizards who specialise in Brain Computer Interfaces had these aims in mind as they developed the *Hex-o-Spell* system. Their computer doesn't just read a name or word, it allows a user to communicate any message by thought alone. It's a brain-reading, machine-learning, text-input system that allows the user to type using their minds.

Simple brain-computer interfaces (BCIs) have been demonstrated for several decades. Using sensors placed on the head, often using gel to improve contact, the electrical impulses produced by large groups of neurons in different regions of the brain could be measured. The broad pattern or frequency of electroencephalogram (EEG) signals would then be used as the input to the computer. The problem is that our brains are hugely complex organs, each with its own individual design. It is possible for people to develop normally with just half a brain although their brains look very different to ours, and as Benjamin Blankertz explains, "also in normally developed humans, functions are differently located in the brain. And due to the foldings of the cortex, small changes in location may cause strong differences in the EEG." Doing something as predictable as moving a finger or a foot, different people will have very different patterns of brain activity. Even worse, brains throw out new and different patterns of signals all the time [illegible] same person repeatedly. Even the equipment used to read the signals is unreliable – the gel dries out between sensor and head, or sensors are placed incorrectly or slip – so at different times the sensors may misread the brain and [illegible] completely different results from the same [illegible]

To overcome these problems, the standard approaches relied on the adaptability of our brains, and required [illegible] produce more predictable signals. For example, the more a person falls into a state resembling meditation, the more regular and slow the frequency of the EEG waves becomes. This method can be successful in some cases, but it requires lengthy training periods for users and so can be very limited.

What was needed was a way to pick out the activity in the precise regions of the brain that were relevant, and somehow overcome all of the variability inherent in EEG signals.

In normally developed humans, functions are differently located in the brain ...due to the foldings of the cortex, small changes in location may cause strong differences in the EEG.

The problem is rather like listening to the muffled sound of a huge orchestra through a thick wall, and trying to interpret the melody from a single violin – when the orchestra change their seats and play a slightly different tune every time you listen. The traditional approach was to train the whole orchestra to all play more or less the same note. The solution as used in *Hex-o-Spell* was to train the listener to pick out that elusive violin. To make this work, the field of brain-computer interaction had to combine forces with machine learning.

Photograph © Irynа Kurhan

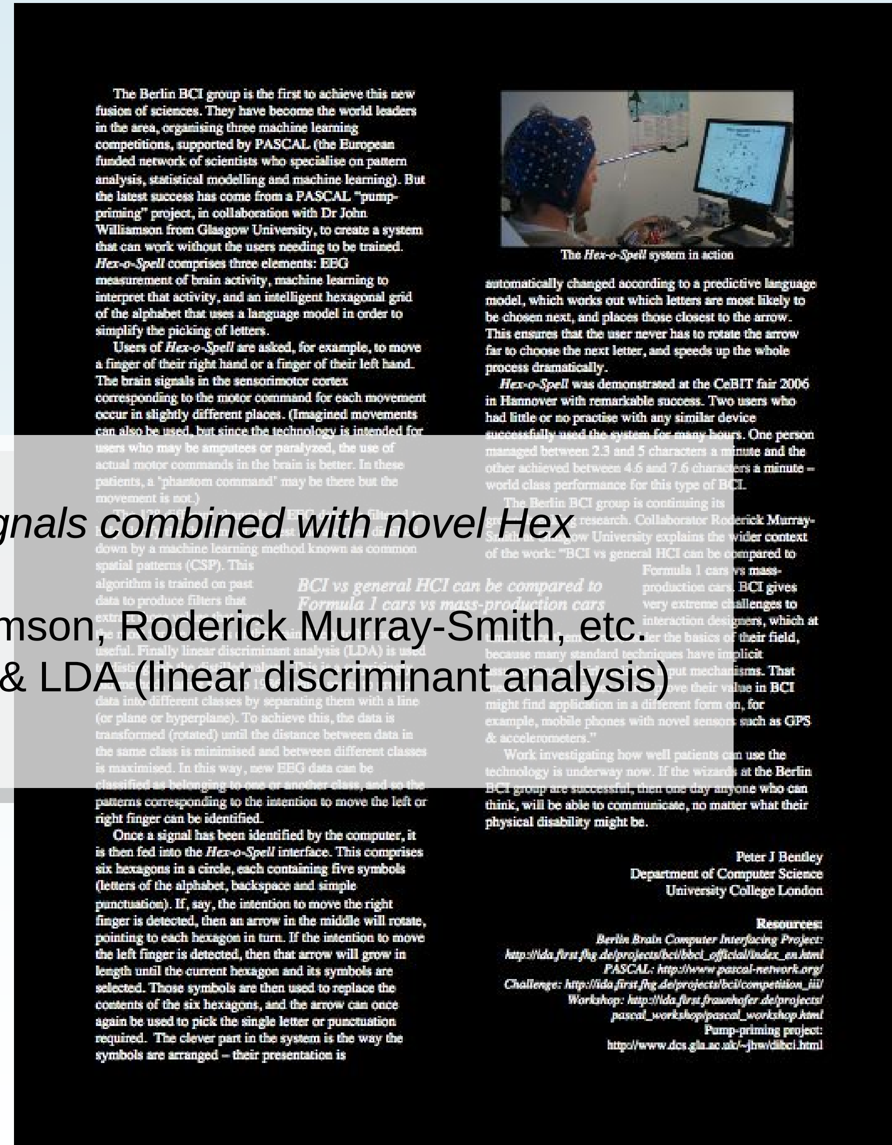

The Berlin BCI group is the first to achieve this new fusion of sciences. They have become the world leaders in the area, organising three machine learning competitions, supported by PASCAL (the European funded network of scientists who specialise on pattern analysis, statistical modelling and machine learning). But the latest success has come from a PASCAL "pump-priming" project, in collaboration with Dr John Williamson from Glasgow University, to create a system that can work without the users needing to be trained. *Hex-o-Spell* comprises three elements: EEG measurement of brain activity, machine learning to interpret that activity, and an intelligent hexagonal grid of the alphabet that uses a language model in order to simplify the picking of letters.

Users of *Hex-o-Spell* are asked, for example, to move a finger of their right hand or a finger of their left hand. The brain signals in the sensorimotor cortex corresponding to the motor command for each movement occur in slightly different places. (Imagined movements can also be used, but since the technology is intended for users who may be amputees or paralyzed, the use of actual motor commands in the brain is better. In these patients, a 'phantom command' may be there but the movement is not.)

The 128 different channels of EEG data are filtered to [illegible] down by a machine learning method known as common spatial patterns (CSP). This algorithm is trained on past data to produce filters that extract [illegible] useful. Finally linear discriminant analysis (LDA) is used to distinguish the distilled values. [illegible] data into different classes by separating them with a line (or plane or hyperplane). To achieve this, the data is transformed (rotated) until the distance between data in the same class is minimised and between different classes is maximised. In this way, new EEG data can be classified as belonging to one or another class, and so the patterns corresponding to the intention to move the left or right finger can be identified.

Once a signal has been identified by the computer, it is then fed into the *Hex-o-Spell* interface. This comprises six hexagons in a circle, each containing five symbols (letters of the alphabet, backspace and simple punctuation). If, say, the intention to move the right finger is detected, then an arrow in the middle will rotate, pointing to each hexagon in turn. If the intention to move the left finger is detected, then that arrow will grow in length until the current hexagon and its symbols are selected. Those symbols are then used to replace the contents of the six hexagons, and the arrow can once again be used to pick the single letter or punctuation required. The clever part in the system is the way the symbols are arranged – their presentation is automatically changed according to a predictive language model, which works out which letters are most likely to be chosen next, and places those closest to the arrow. This ensures that the user never has to rotate the arrow far to choose the next letter, and speeds up the whole process dramatically.

The *Hex-o-Spell* system in action

Hex-o-Spell was demonstrated at the CeBIT fair 2006 in Hannover with remarkable success. Two users who had little or no practise with any similar device successfully used the system for many hours. One person managed between 2.3 and 5 characters a minute and the other achieved between 4.6 and 7.6 characters a minute – world class performance for this type of BCI.

The Berlin BCI group is continuing its groundbreaking research. Collaborator Roderick Murray-Smith at Glasgow University explains the wider context of the work: "BCI vs general HCI can be compared to Formula 1 cars vs mass-production cars. BCI gives very extreme challenges to interaction designers, which at [illegible] the basics of their field, because many standard techniques have implicit [illegible] input mechanisms. That [illegible] prove their value in BCI might find application in a different form on, for example, mobile phones with novel sensors such as GPS & accelerometers."

BCI vs general HCI can be compared to Formula 1 cars vs mass-production cars

Work investigating how well patients can use the technology is underway now. If the wizards at the Berlin BCI group are successful, then one day anyone who can think, will be able to communicate, no matter what their physical disability might be.

Peter J Bentley
Department of Computer Science
University College London

Resources:

Berlin Brain Computer Interfacing Project: http://ida.first.fhg.de/projects/bci/bbci_official/index_en.html
PASCAL: http://www.pascal-network.org/
Challenge: http://ida.first.fhg.de/projects/bci/competition_iii/
Workshop: http://ida.first.fraunhofer.de/projects/pascal_workshop/pascal_workshop.html
Pump-priming project: http://www.dcs.gla.ac.uk/~jhw/dibci.html

Promedas: PRObabilistic MEdical Diagnostic Advisory System

Wim Wiegerinck, Martijn Leisink, Bert Kappen, Jan Neijt

Bayesian inference rules

THE DESKTOP DOCTOR

Hundreds of years of medical experience. An infinite patience and the ability to take every symptom into account. Precise and logical, up-to-date, and never short on ideas. All just casually sitting on your doctor's desk. It may not have much of a bedside manner, but then its job is not to meet patients. Its role is to brainstorm with human doctors and help them look at all possibilities when diagnosing illnesses. It is Promedas, the desktop doctor.

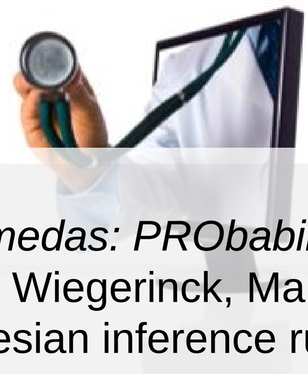

Promedas is a PRObabilistic MEdical Diagnostic Advisory System – a computer program that figures out how likely certain medical diagnoses are from a given set of symptoms and test results. It addresses the hugely complex task faced by doctors every day: how to ensure that even unusual or obscure causes of symptoms are not overlooked. Whether the doctor is a specialist in one area or a general medical practitioner, keeping up-to-date and remembering the myriad relationships between illness and symptoms is often tremendously challenging. Promedas does not make a diagnosis, but it suggests alternatives to the doctor and based on the information it has been given, shows exactly how likely it is for the patient to be suffering from each condition.

Since the beginnings of Artificial Intelligence, there has been a dream of creating a computerised expert, filled with our knowledge, and able to infer new conclusions of its own. These computer programs became known as expert systems, often using elaborate trees constructed of rules: "if symptom A and not symptom B then ask about C, if symptom C then diagnosis is D". But unlocking the knowledge held within scientific journals and coded in the neurons of specialists was not so easy. In medicine, for example, medical conditions and symptoms are not either true or false. Not all people experience the same symptoms, and the appearance of some symptoms is hugely more significant than others. There can also be a large number of symptoms, leading to rules with excessive numbers of variables. So the old "decision tree" methods were often cumbersome and made bad decisions, sometimes ruling out possibilities for no other reason than data being presented in an unexpected order.

The failure of these rationalist "good old fashioned AI" (GOFAI) methods led to many scientists rethinking the ideas. Clearly an expert system needed to represent knowledge in some form, and clearly that knowledge needed to be used with data to infer some form of decision. But how best to achieve these goals?

One common solution was to use fuzzy logic, where the binary true or false rules were turned into linguistic variables such as "partially true" or "mostly false". But even fuzzy logic still suffered from problems: it might enable the expert system to define degrees of truth, but in fields such as medicine, a partially true diagnosis is not ideal. Instead, it would be much more useful if the probability of a diagnosis being true was provided. For example, while a runny nose might result in a "partially true" diagnosis that you have a virus in a fuzzy logic system, it would be more helpful if the system could infer that you have a certain probability of suffering from several different illnesses, some much more probable than others. The solution, as exemplified by the medical decision support system Promedas, was to use Bayesian inference rules.

Thomas Bayes was a mathematician born in London in 1702. Amongst his works, he wrote about probability. Instead of being concerned with, say, the probability of drawing a black ball from a bag of a

Photograph © Yazik Chauvin 2007

certain number of black and white balls, Bayes was interested in the inverse probability of the event. In other words, if you had drawn more black balls compared to white balls from the bag, what was the probability that the bag contained more black balls and so you would draw another black one next? Given a hypothesis (for example that the next ball will be black) and some information about which balls have been picked previously, Bayes figured out the maths to infer the probability that the hypothesis was true or not. Several decades later, Frenchman Laplace developed these ideas further, creating a more general version of the Bayes theorem for use in astronomy and physics.

Amazingly, two centuries later, systems such as Promedas now use Bayesian inference to achieve decision support. It's ideal for an application such as medicine where we have plenty of evidence that certain symptoms tend to be observed for specific illnesses.

Promedas is the result of many years of development by researchers such as Wim Wiegerinck, led by Martijn Leisink and Bert Kappen of Radboud University Nijmegen in the Netherlands. It now contains Bayesian inference rules that [illegible] advantage. In the words of Martijn Leisink: "The clear advantage of the probabilistic representation is the [illegible] each other. If one finding leads to two different diagnoses these are difficult to merge in a rule based system. At least one additional rule is necessary. In the probabilistic setting, making use of the basic probabilistic rules, it is immediately clear how to combine evidence and variables."

This is the future for medicine with all the sub specialists. They need a program that looks further away than their sub specialisation.

The success of Promedas relies on its careful structuring of the dependencies (which findings imply which diseases). It organises its information as a tree in a three-layered noisy "OR" model. The layers of the tree correspond to risk factors such as occupation or drug use, possible diseases and the tests and symptoms. Each node in the network (either a risk factor, disease or test result) is linked according to specific probabilities of cause and effects, with some risk factors likely to cause some diseases, and some diseases likely to cause some symptoms and results of tests. By representing each node as a "noisy OR" Promedas simplifies and speeds up the inference process, making the assumption that each cause can behave independently but multiple causes combine to make outcomes more probable. With a given model and data about the risk factors and test results, the probabilities of different diseases can then be inferred.

As with all such systems, often the main bottleneck is simply the input of data. Many patient records are incomplete or are not in the right kind of format to enable easy input, and typing in the results of tens or hundreds of different observations and tests can be laborious. As more patient records become stored electronically, this will become less of an issue but in the near future the use of Promedas is likely to be restricted to those mysterious cases where the doctor needs some new ideas. It's an important role, for specialists who have chosen to focus on one specific area of medicine may become less knowledgeable about other areas. The vast amount of expert knowledge in Promedas means that all physicians of all specialisations will have access to the same up-to-date specialist information. There can be no doubt that a list of possible diagnoses of [illegible] for a patient makes an excellent decision support system, for it may suggest rarer alternatives that could be confirmed by additional tests.

The first large-scale trial will begin in early 2008 at the University of Utrecht, Netherlands. While results are still not perfect, initial experiences by doctors are very positive. In the words of Dr Jan Neijt of the University of Utrecht (the physician who has so far provided all of the medical knowledge for Promedas), "…used with reason it is always helpful… This is the future for medicine with all the sub specialists. They need a program that looks further away than their sub specialisation."

Perhaps one day systems such as Promedas will become as ubiquitous as the stethoscope, with their databases updated as new medical findings are published. We can never remove uncertainties from medicine, but with decision support systems we can ensure that all decisions made by our doctors are as well-informed as possible.

Peter J Bentley
Department of Computer Science
University College London

Resources:

Promedas (including online demonstration): *http://www.promedas.nl/*

Promedas publications: *http://www.snn.ru.nl/nijmegen/publicatie.php3?projekt=Promedas*

PASCAL: *http://www.pascal-network.org/*

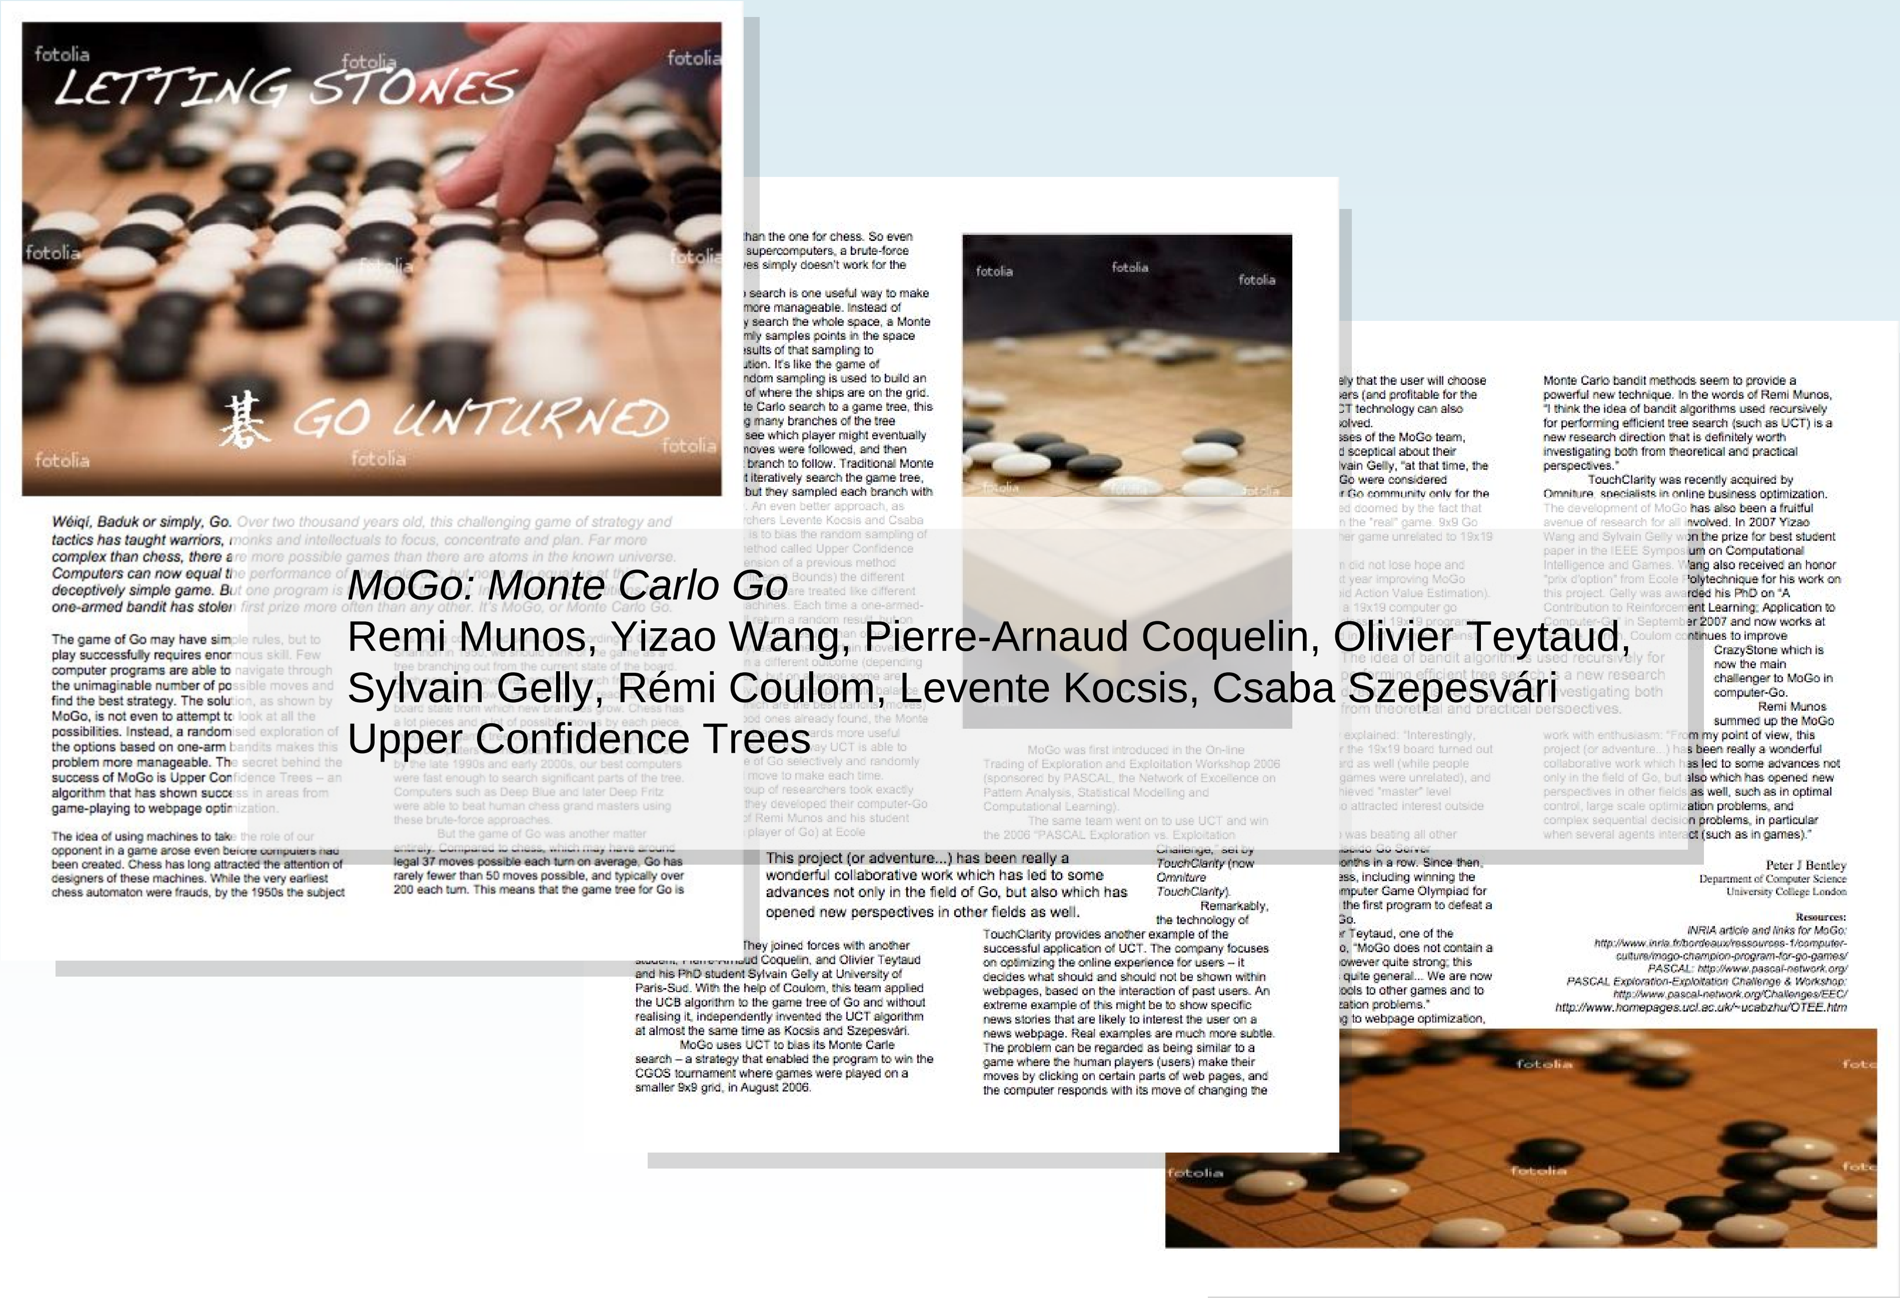

MoGo: Monte Carlo Go

Remi Munos, Yizao Wang, Pierre-Arnaud Coquelin, Olivier Teytaud, Sylvain Gelly, Rémi Coulom, Levente Kocsis, Csaba Szepesvári

Upper Confidence Trees

Things you can do to help:

- always try to make yourself available and be friendly to journalists; they have tight deadlines and if they can't talk to you they'll give up or use someone you might prefer they didn't.

- provide general information on websites: biographies, summaries of research, pointers to important new papers and results.

- learn how to explain your work using everyday language (it often makes the concepts crystal clear in your own mind if you can talk without any terminology)

- think of a major new result or discovery which has a broad impact as a news event and write a press release with the help of your university.

Think about how to communicate with a more general audience - it’ll benefit you.

It’s nice to talk to peers, but there’s a whole world out there.

Questions?