


CS4FN

Computer Science for Fun

Issue 7



Special Issue on

Computer Science and the Image

Turning Silicon into Gold

'Googling' Movies

*Florence Nightingale:
Rebel with a Cause*

Special Issue on

Computer Science and the Image

Many people think that Science and Art are opposites. Leonardo da Vinci would never have agreed. He was both scientist and artist, using his science to guide his art and his art as inspiration for his science.


Art meets Science didn't just stop with Leonardo though. Science is a very creative endeavour of course, and this is especially so for Computer Science. Now, in the age of multimedia, computer science is there in the middle of it all. It gives artists new media to work in, and even new artists to do the art. CGI – Computer Generated Imagery is now a central part of filmmaking and computer scientists are still pushing the boundaries of what can be done.

The vast quantities of multimedia on the web and in personal collections on our laptops and mobiles are giving computer scientists new challenges in finding ways to ensure we don't drown in all the material. How do you find photos or specific scenes in a film? Artists, images and image-based algorithms can even help the scientific process. It is an exciting time to be a creative Computer Scientist.

Arabesque Art

You don't need a paintbrush to create winning art. It's possible to create images by writing a program. Raghd Rostom demonstrated how beautiful and intricate patterns can be created in this way. Her programmed art based on an Arabesque style led to her winning BrainAcademy 2007, a competition about computer science and creativity.

Raghd used a drawing system called GeomLab to create her winning images. Geomlab was developed by Mike Spivey, a computer scientist at Oxford University, as a fun way for people to learn to program in a style known as functional programming. It gives an easy way of creating patterns made up from repeated tiles. The program really is a picture.



In this issue:

- Pg 8** **The Colour of Music**
- Pg 10** **'Googling' Movies**
- Pg 16** **Turning Silicon into Gold**
- Pg 19** **Florence Nightingale:
Rebel with a Cause**

Randomly rich art

There are lots of ways computer science techniques can be used to create innovative art. Fractals of the kind illustrated here are one example. They are created using simple rules that, when applied over and over, lead to beautiful and very natural looking images. They are self-similar in the sense that if you take any part of the image, however small, and blow it up you get something very similar to the original. Fractals are an easy way to create very realistic digital trees and ferns because the fractal process mirrors their natural development processes in the real world.

Software designer and visual artist Simon Schofield's work is another example of computer science leading to novel art. He creates incredibly complex digital images made up of vast numbers of tiny drawn images. He has developed a set of special software techniques that put the full image together by following simple sets of rules that have a form of controlled randomness at their core. This leads to a kind of endlessly rich repetition. The results are amazing, capturing the detail and richness of natural textures but with an alien quality. Follow the links from the webzine, Magazine+ to find out more.



Beauty beyond measure

Did you ever think a number could be beautiful? How does 3.1419 grab you? Is 6.238673 hot? What about 1.61803399, does it catch your eye? Strange as it may seem the number 1.61803399 is responsible for some of the most beautiful things in the world around us. It's called the Golden Ratio and it's been used for centuries in art and architecture. So what's so special about this number, what is its link to beauty and why might it lead to a moral dilemma?

The Ratio all around

Take a range of different sized rectangles and ask people to say which looks the best proportioned, that is the width and height look just right, and easy on the eye. They will tend to pick the rectangle where the length and width are in the golden ratio. That is if you divide the length by the width the answer is close to 1.61803399. That is why credit cards are the shape they are. Leonardo Da Vinci and other Renaissance painters used the golden ratio to structure their pictures. For example, the proportions in the Mona Lisa are related to this 'beautiful' number. The Parthenon in Greece also has the Golden Ratio built into its shape. In nature sunflowers, pinecones and other exquisite natural structures exhibit the golden ratio. Many have argued this is because the Golden Ratio arises from using something called a Fibonacci sequence (see the webzine magazine+ for more on this), and a Fibonacci sequence is one of the best ways to pack things tightly together: beautiful and efficient. Of course beauty, or aesthetics, are often in the eye of the beholder and not all scientists are convinced about the 'beautiful' properties of the Golden Ratio.

The perfect face?

American plastic surgeon Dr. Stephen Marquardt has studied human beauty for years. He believes that he has discovered a fundamental rule for facial beauty. Whatever the race, period of history or gender, beautiful faces exhibit the Golden Ratio in their shape. He has developed a 'beauty mask': a grid of lines in the form of a face where the line's geometry is determined using the Golden Ratio. By overlaying this grid on a face and seeing how well an individual's features match up to the 'perfect' Golden Ratio face, you can 'calculate' how beautiful a person is. Dr Marquardt has suggested the mask could be used to help with creating cosmetic makeup, surgery and dentistry for people who have suffered facial disfigurement.

Beautiful idea or not?

There have even been attempts to automate the measurement of how a face matches up to the mask, and so give a numerical 'score' for a person's beauty. Developed for use on a camera phone it's not available in the shops just yet. Would you be interested in using it? Perhaps you think that having people use such an application would be a really bad idea. After all, reducing a person to a single number kind of throws away lots of the other things that make us people. This only looks at geometry, and there is surely more to beauty than that.

How would it be used? Could it cause people to be bullied? Is the science behind it correct? Computer Scientists need to think about these sorts of issues. Ethics, whether something is right or wrong, is an important part of a computer scientist's training. After all they have the power to create new applications, but what applications should they create? Are there applications that would be wrong to work on? With great power comes great responsibility.

Try the beauty mask for yourself by following the links from the Magazine+ part of the cs4fn webzine, but remember there is much, much more to being a 'beautiful' person than a beautiful face.

Aaron and the Art of Art

Aaron is a successful American painter. Aaron's delicate and colourful compositions on canvas sell well in the American art market, and have been exhibited worldwide, in London's Tate Modern Gallery and the San Francisco Museum of Modern Art for example... Oh and by the way, Aaron is a robot!

Yes, Aaron is a robot, controlled by artificial intelligence, and part of a lifelong experiment undertaken by Professor Harold Cohen to create a creative machine. Aaron never paints the same picture twice; it doesn't simply recall pictures from some big database. Instead Aaron has been programmed to work autonomously. That is, once it starts there is no further human intervention, Aaron just draws and paints following the rules for art that it has been taught.

Perfecting the art of painting

Aaron's computer program has grown and developed over the years, and like other famous painters, has passed through several artistic periods. Back in the early 1970's all Aaron could do was draw simple shapes, albeit shapes that looked hand drawn, not the sorts of precise geometric shapes that normal computer graphics produce. No, Aaron was going to be a creative artist. In the late 1970's Aaron learned something about artistic perspective. Objects in the foreground are larger than objects in a picture's background. In the late 80's Aaron could start to draw human figures, knowing how the various shapes of the human body were joined together, and then learning how to change these shapes as a body move in three dimensions. Now Aaron knows how to add colour to its drawings, to get those

clever compositions of shades just spot on and to produce bold, unique pictures, painted with brush on canvas by its robotic arm.

It's what you know that counts

When creating a new painting Aaron draws on two types of knowledge. First Aaron knows about things in the real world: the shapes that make up the human body, or a simple tree. This so called declarative (declared) knowledge is encoded in rules in Aaron's programming. It's a little like human memory: you know something about how the different shapes in the world work. This information is stored somewhere in your brain. The second type of knowledge Aaron uses is called procedural knowledge. Procedural knowledge allows you to move (process) from a start to an end through a chain of connected steps. Aaron for example knows how to proceed through painting

Like other famous painters, Aaron has passed through several artistic periods

areas of the scene to get the colour balance correct, and in particular, getting the tone or brightness of the colour right. That is often more artistically important than the actual colours themselves. Inside Aaron's computer program these two types of knowledge, declarative and procedural, are continuously interacting with each other in complex ways. Perhaps this blending of the two types of knowledge is the root of artistic creativity?

Creating creativity

Though a successful artist, and capable of producing pleasing and creative pictures, Aaron's computer program still has many limitations. Though the pictures look impressive, that's not enough. To really understand creativity we need to examine the process by which they have been made. We have an 'artist' that we can take to pieces and examine in detail. By studying what Aaron can do, given we know exactly what's been programmed into it, allows us to start to examine human creativity. What about it is different from the way humans paint, for example? What would we need to add to Aaron to make its process of painting more similar to human creativity?

Not quite human

Unlike a human artist Aaron cannot go back and correct what it does. Studies of great artist's paintings often show that under the top layer of paint there are many other parts of the picture that have been painted out, or initial sketches that have been redrawn as the artist progresses through the work, perfecting it as they go. Aaron always starts in the foreground of the picture and moves toward painting the background later, whereas human artists can chop and change which part of a picture to work on to get it just right. Perhaps in the future with Harold Cohen's help Aaron will develop new human-like painting skills and produce even better paintings. Until then the art world will need to content itself with Aaron's early period work.

To see Aaron in action follow the links from the cs4fn Magazine+ website.

Behind You!

In 3D computer graphics it's vital to be able to produce realistic looking images. Scenes are often created by arranging 3D shapes in the right pattern. The problem is that sometimes this means that parts of things behind should be hidden by parts of things in the front. Calculating the geometry over which bits of objects are covered isn't easy. There is a simple way to solve this so-called 'hidden surface' problem though: the Painter's algorithm. The computer acts like a simple, but orderly artist who paints the distant parts of a scene first and then splashes on more layers of colour, covering the background by painting those parts that are nearer giving a simple, but often effective, image.



Ed v Zach

Aaron (see page 4) isn't the only computer with artistic pretensions. In a unique experiment at the Institute of Contemporary Arts, London in 2005 two digital artists, Ed Burton and Zach Leiberman created the ultimate cyber head to head. In Ed v Zach two computers were programmed to observe each other's screen and take it in turns to draw the other's output. One computer ran a program called "Ed" by Ed Burton and the other a program called "Zach" written by Zach Leiberman. Placing your hand on one screen for example would make the other computer try to draw this shape. Over a period of days the two computer programs drew each other endlessly, creating a slow and beautiful feedback process and generating wonderful interactive art.

Even computers can draw inspiration from each other!



Sketching Science

Scientists often find themselves doodling on scraps of paper, or sketching out ideas on a whiteboard, but that's not really what science is about ... is it?

In fact, scientists have to be highly creative and all that doodling and sketching both helps them share ideas and plays a vital role in producing flashes of inspiration in the first place.

Bongo playing, Nobel winning physicist Richard Feynman's development of 'Feynman Diagrams' is a classic example – it helped revolutionise the way scientists think of quantum physics: the science of subatomic particles. Feynman Diagrams are just a simple way of using pictures to visualise the quantum level processes that occur.

Queen Mary computer scientist Pat Healey and Tessa Garland, an artist from the Bow Arts Trust, are studying just how big an impact drawing can have on science.

Tessa is spending time exploring the way scientists use drawings in their work. They draw in a whole bunch of ways: on paper, whiteboards and computers. Computer scientists have borrowed observational techniques like this from the social sciences. It's a way they gather data about the way people behave in real situations so that those activities can be supported using computers.

Pat and Tessa are analysing the kind of images and styles scientists use to represent their results. From this, Tessa explores different ways to draw the same things, even creating 3D versions. The idea is to provide the scientists with new techniques for exploring the problem spaces they are working on. Drawing is being used to uncover new facets of scientific problems and to uncover new ways to get to solutions.

The ultimate aim is to communicate the results both to other scientists and to the public in exciting ways, using, for example,



interactive drawing games. Some of the ideas were tried out as part of the "Big Draw" last year run by the Campaign for Drawing, allowing kids to get involved. For example, children used the program *Stop Motion Pro* to explore developmental processes with artist Albert Potrony by creating their own hybrid creatures, and animating them developing from a fertilized egg to a baby.

Not all scientists may realise it but drawing is an important part of science so if you are stuck for an idea, pick up a pencil. That missing flash of creativity might be in your next doodle.

Drawing Dvořák

Pat Healey of Queen Mary has also looked at the role of sketching in multimodal interaction. As they pervade everyday life computers increasingly have to be multimodal – that just means they communicate, for example, with both sound and images. My satnav is multimodal. It both gives me voice commands and has a map showing the route. The combination is more effective than either alone. Pat explored how well people can communicate music by drawing. What kind of drawings do people do to visualise Beethoven, for example, and in what ways do they look different to ones of jazz music, say? The results help computer scientists understand how we naturally link sounds and images and so how best to combine them when designing devices that use multi-modality more naturally.



The Mind Boggling Memory Experiment

After years of training, world memory champions can memorise packs of randomly shuffled cards and recall them in perfect order. In this experiment you will be able to emulate a small-scale feat of memory magic, impressing your friends with your mind's memory power.

Memory magic

Hand a pack of cards to your friend who gives it a good shuffle to ensure there is no set up. Get them to split the pack and divide it roughly into two. They then chose which part of the pack to give you. Pick up those cards and quickly scan them, face-up, for about 5 seconds, saying that you are memorising the order of the pack. Now place these cards out of sight in your pocket.

Name a card?

Your friend then names any card in the pack (not the joker). Suppose they name the Ten of Diamonds. You think carefully, and state that the Ten of Diamonds isn't in your pocket. However, you say you can still prove you can picture the whole deck in your mind. First you pull out a card of the same suit as their freely chosen card. You follow that up with cards that add to the value they gave.

Cards from pocket

You put your hand in your pocket. 'Hmmm, 12th card from the top is the 8 of Diamonds I seem to remember', you say, and you pull out the Eight of Diamonds, just as you said. You've matched their suit with a Diamond. 'Of course', you continue, 'there was a one in four chance of getting a Diamond at random, so I will pull out another card to make the value of your chosen card, the 10 of Diamonds'. Without much effort you remember there was a Two of Hearts 'five cards from the bottom' and pull it out to complete your card memory challenge: Eight of Diamonds and Two of Hearts, $8+2=10$, proof you have amazing skill to memorise cards in an instant, or is it?

If you want to find out how to do the trick and its link to computer science, go to the cs4fn webzine, Magazine+ section.



The new Hampton Court Maze

Many of the cs4fn webzine pages are linked together using a hyperlinked maze. Somewhere out there is the middle of the maze. Can you get to it by navigating through the cs4fn pages? Just like a real maze you can ignore the interesting things you pass that live in the maze and just head for the centre as fast as you can ... or you can take your time and enjoy the things you stumble across along the way.

Our most recent cs4fn maze is modelled on the real Maze at Hampton Court. If you draw your route as you navigate the webpage maze you will end up with a map that will get you to the centre of the real Hampton Court Maze too. The kind of map you need is what a computer scientist calls a 'graph'. You should only include the essentials - draw a dot (a 'node') for each decision point and straight lines ('edges') between them to show how they connect together. By treating things as graphs computer scientists can use the same programs to solve lots of apparently different problems. For example, if you were to write a program to explore a maze and find the shortest route to the centre, then that same program could be used by a satnav program to find the shortest route for drivers to follow between cities. Both the road network and a maze can be represented by a graph: the problem is the same and so is the solution.



The Colour of Music

Music is all about emotion. From a slushy love ballad to the head-banging, no nonsense roller coaster of thrash metal, the songs we listen to can make or change our mood. We have favourite tracks to play when we need cheering up, even maybe some music that keeps melancholic company with us when we are lost in thought. Emotions are what it's all about, and with portable music players built into many of today's mobile phones we can listen to music whenever and wherever we want. So how do you ensure that the music you are listening to is right for the mood you are in? Of course you have play lists, lists of your favourite songs, which you can flick through to find the right song for the right mood, but is there a better way? This is the question Queen Mary computer science student Daniel Fiegenschuh asked, and his answer was a smart and colourful one.

I'm in the mood for some colours

Dan realised that we tend to associate particular colours with particular emotions. Red, for example, might be associated with feeling angry, blue feeling chilled, yellow feeling happy, and so on. We each tend to give colours an emotional 'tag'. That is, we associate them with a mood. Of course we all tag colours differently. Red might be a colour that makes *you* feel happy, and blue makes *you* feel sad. It's something very personal. So Dan decided to write some software to let people tag their favourite songs with a colour. For example,

your chill-out songs might all be green. That's an emotional association that works for you. Dan's software, which works on a mobile phone, creates a picture filled with clouds of the different colours you chose that are important to you, and lets you automatically create a play list of tunes depending on your mood. You move the cursor over say the green cloud, click and you create a play list of a set of the songs you tagged as green. The software comes up with the play list, so you'll maybe have some surprise tracks in there. Unlike a play list you created yourself, the system chooses colour tagged songs near the point you selected.

Phone a friend

Dan tested that his system was easy to use by giving it to his mates to try out for real over several days. He then asked them what they thought of it. This process is called user evaluation and all software should go through this to check it is easy to use, and does what it's supposed to (though in a professional setting it wouldn't just be your mates who tested it of course).

The feedback was that they liked it, it was easy to use and also looked cool, but they didn't like the idea of having to tag all their tracks by hand (too much hassle). They also thought that sometimes a track was in the wrong colour cloud. It wasn't such a red track after all. There were several times that they skipped tracks in the playlist because of that. Armed with these comments Dan took the project to the next level, and in came some artificial intelligence to make the phone smart.

Clever colours

Dan used a simple computer learning technique called negative selection to add some smartness to the phone. When people skipped a track or just played a little of it, the tune was obviously in the wrong colour cloud. So the duff tracks now start to move round on the screen. They are chucked out of say the green cloud as they are negatively





selected (skipped) and have to find another better cloud to join. The distance they are moved depends on how long they were played for; shorter plays get them chucked further away from the original colour cloud centre. This way the tracks that aren't green end up in a different colour cloud. After playing with the system for a bit Dan's friends found that the artificial intelligence had moved any wrongly tagged tracks into the right colour clouds.



Result!

This negative selection idea also meant that any tracks that hadn't been tagged at the beginning could be tagged with a random colour and then would rattle around through different clouds, like a ball in a pinball machine, till they settled in the right place. The artificial intelligence had solved the problem of automatic tagging of songs and also correcting the tagging of songs in one go. Result! His mates agreed the new system worked and after a while it was generating exciting mood-enhancing play lists for them whenever they wanted. Maybe one day soon Dan's smart colour cloud software will float onto your phone. When it does let's hope it's a yellow day: a happy day that is.

Optimisation and all that jazz

Computer Scientists can find inspiration in the strangest of places to help solve some of the toughest mathematical problems. One such mathematical problem is called optimisation; you do it every day when you decide what favourite songs go on your play list, or work out the shortest route to school. Optimisation is about finding a best solution: deciding on the best choice from a range of options. Optimisation is important commercially too. How do you connect all the houses on a new estate together with the least amount of electrical wiring or fewest sewer pipes, for example? What's the best route for your delivery trucks to take round the shops with the minimum mileage ... while trying to avoid using the motorway at rush hour ... and so on. Solving these problems isn't easy, particularly when you have many, many options to choose from.

Music, and in particular jazz improvisation, might seem an odd place to look to help solve this classical optimisation problem, but that's just where a rather cunning method called 'harmony search' found its inspiration...visit the webzine Magazine+ to find out more.



'Googling' Movies

Google-like search engines have been massively successful. Trouble is you have to use text to search...but the world has gone multimedia. The web is now full of images and movies. Youtube has 70 million videos. Flickr has 2 billion images with a million more added every day. Even searching personal photo collections can already be daunting - and photos only went digital a few years ago!

The problems are increasingly becoming everyday ones. Just as 'googling' has become normal activity (how did we manage before?) so much more will be possible if the problems of searching images can be cracked. Professor Andrew Zisserman's Visual Geometry research team at Oxford University have some solutions.


First, let's look at some everyday scenarios. Suppose you do the PR for a company selling a revolutionary new camera design, and have done a product placement deal. The latest Bond film will feature the new camera at least 20 times in the final cut of the film. You have paid a lot of money. Now you want to be able to check that it really does happen. How? Is the only way to sit someone down in front of the film and count? What if they counted 19? Perhaps they just got too engrossed in the film and missed one!

Here's another scenario. You run a successful webzine. You take great pride in using stunning images and need an image to illustrate a story about an iPod. You know you took a perfect image a couple of years back. Fine except you have tens of thousands of photos and you can't remember when you took it. You can't just search for the word 'iPod'. That would only have worked if every time you took a photo you had labelled it with good descriptions. Life is too short.

Find that photo!

Here is a more futuristic scenario (actually not that futuristic anymore, but that's another story). Suppose you are now one of millions of people who wear a video button. It records everything you see and do and saves it forever. All the time. It means you never need to forget an experience. You are now standing in front of the Bridge of Sighs in Oxford with your boyfriend. You claim you've both been here before. He says it must have been someone else. Exactly the kind of argument the video button should solve...but how do you rewind to that particular video memory given you don't remember when it was?

What you need in all these cases is a way to 'google' images...using images. Currently you type in text to search for and someone must have previously labeled the image with that text. That isn't good enough for our tasks. We want to find images when there is no text in sight. Wouldn't it be good to 'google' images by giving an image of an object (like the camera, the iPod or the bridge) and the search engine just returns all the scenes or images that match and in order of how well they match? How could such a system work?



Bridge? What bridge?

There are three different but related problems you could want to solve if searching for images using images in this way. You could, given a picture of a bridge want images of a bridge, any bridge. Alternatively you could be looking for a near duplicate version of that actual image. Perhaps you have the low-resolution version and want the original high-quality version. A third possibility is you could want to find different photos containing that same object. You just want images of the Bridge of Sighs in Oxford.

The last option might sound easy at first - can't you just cut out the object you are interested in and have the computer check if it matches any portion of each image in the collection being searched, pixel by pixel? Well no. In the other images, it won't be at exactly the same angle. The size will be different: it might be a



tiny object in the background. If taken at a different time, the light will have changed. There might even be other things in front of it, partially hiding it.

The Oxford team set out to answer two questions. Is it possible to use a similar approach to the Google text searching methods to search for objects in images? Secondly, can you retrieve all occurrences of an object from an image collection or video?



Find that word

How do search engines search web pages so quickly? If you think about it, it is amazing. You can type in any arbitrary word and Google can search the whole of the web for it instantly! How on earth is it done so quickly?

Search engines send out web crawling programs that go from page to page, taking a copy of the pages and processing them. A first thing they can do is called 'stemming'. This means chopping down words to their roots. If 'walks', 'walking' and 'walked' are all converted to 'walk' they can be treated as though they were all the same. Common words like 'the', 'of' and 'a' are also thrown away as uninteresting. Finally what is stored back at base is an 'inverted file'. That just means that each word found is stored with a list of all the places it appears. The process is basically the same as creating

the index for a book. The final index is an inverted version of the book - each word is followed by a list of page numbers where it appeared. For web pages, the web addresses are stored rather than page numbers. Each search engine company stores a gigantic inverted version of the web on its servers. All this work has been done off-line long before you ever did a search. When you do a search, all that is needed is to look up the word in this vast index and return the web addresses.

Spot the visual word

If we are going to do something similar for images, then we need some image equivalent of words to search for. Pictures of whole objects don't work because they change too much. What does work are small patches of images. You have to choose the patches with care. They need to be interesting points with distinctive texture. Essentially that just means there is a distinctive pattern of edges in the patch. Choosing interesting patches is a bit like the way Google only stores interesting words for text search.

You don't just use the patch of the image as the 'word' though. You turn it into a series of numbers recording things about the texture. In the Oxford system this involves recording 128 numbers for each patch. Storing these properties of the patch rather than the image itself means we don't have to worry about corresponding patches on different images being different sizes or orientations. Magnify or rotate the patch and it will still give the same 128 numbers. We now have the basis of our 'visual words'. There is one more step though, a little similar to chopping off the ends of words. Patches with very similar sets of numbers are grouped together into a single point. How do we decide they are similar? Suppose it was just 3 numbers not 128. The numbers could represent a point in space. Similar sets of numbers are those that cluster together in space. The idea is just the same but in a 128 dimensional space (OK - so that is hard to imagine in reality, but scientists think of multi-dimensional spaces all the

time). Points that cluster together in the 128 dimensional space are treated as a single point. We now have our visual words to search for.

Suppose we are searching for iPods in our photo collection. Before we start the system will have calculated all the visual words on every photo. This might be a thousand or more visual words for each one - a thousand or more interesting patches. Those are stored as an inverted file. Remember that is just an index allowing you to look up any visual word that appears anywhere and gives a list of all the photos it occurs in.

Now we start with any picture containing an iPod and select an area round the iPod so that it is the only part of the image that we search for. The system first works out what all the visual words in that area are. It then looks them up in its big index and returns all the images with matches. The more patches that match, the closer the fit. With text, the search process might stop there. With an image there is more you can do though. Images have structure - even if an object is at a different angle, or a different size, the relative positions of patches on the same object will stay the same.

Your eyes don't jump to the same side of your nose just because you turn your head! The images that contain the same object will therefore not just have matching visual words, but those visual words will also stay in similar relative positions. Put the images side by side and suppose the second is just smaller or shifted. If you draw lines between the corresponding patches then they will not, on the whole, cross. Those that do are just similar looking patches on something different, so can be ruled out. You are left with a set of pictures from your collection that contain a match with the iPod. They can be ranked, for example, by the number of matches left. Go through quickly checking it's not a false match, counting them and you know how many times the iPod appeared.





A video is just a sequence of still images. Finding the camera in the Bond film is therefore really just the same problem, solved in the same way. One improvement is that in films, the stills are grouped into scenes where not that much changes from one to the next. Single 'key frames' can be used to represent whole sequences.

Of course you can just as quickly search for any other object once the index has been made – all the scenes where James Bond's bow tie appears perhaps, or the ones with Blofeld's cat in.

By bringing together techniques from computer vision and information retrieval research, the Oxford team has shown you can search images in the way current search engines search text. It should therefore scale to gigantic collections in the same way that Google can search all the vastness that is the web. Their experiments have also shown you can pull out all the occurrences of a specific object in this way. What next? Well, there are lots of ways that might improve the system further. The visual words take no account of colour, for example. The system can also only search for objects. What about actions: a search for images of 'running' or 'eating' perhaps. Can that be done in a similar way?

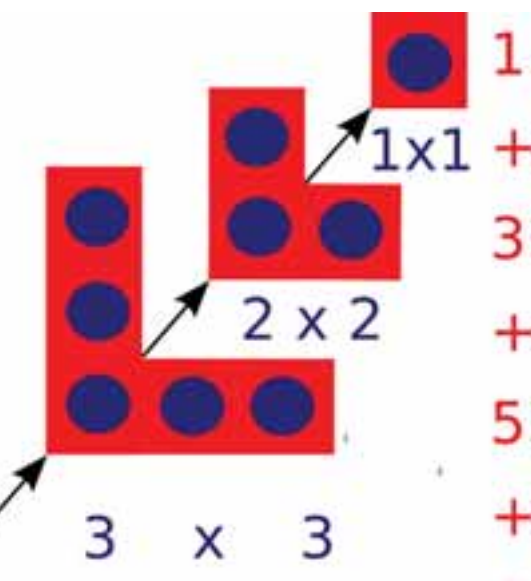
There is lots more research to be done but the system is already impressive – and you can try it out for yourself – both searching films and searching photo collections. You might just be amazed at how fast you can find an object in a film not to mention how accurately it does it...and it is all down to some very clever algorithms and a lot of research.

Oh! And if you really were filming your whole life you could of course do a video search to find all the times you had stood looking at the Bridge of Sighs and see who was with you...though perhaps some things are better left unchecked...just in case.

This article is based on a seminar Andrew Zisserman gave at Queen Mary, University of London, about his work with Ondrej Chum, Michael Isard, James Philbin and Josef Sivic.



Proof without words




Graphic news images often help sway public opinion. Images of famine in Africa led to LiveAid. Images from war zones of civilians can be disturbing enough that war leaders lose or gain political support as a result (depending on who did the bad stuff). Images can have far more power than words to argue a case...to persuade.

to leave no element of doubt that some fact is true, not for emotional reasons but by logic...and they can be just a picture too. To find out more go to the webzine Magazine+, but in the meantime ponder the diagram on the left along with the mathematical 'fact':

Mathematical proofs are also just arguments intended to persuade. They aim

"The square of any number is equal to the sum of consecutive odd numbers."



StoryBank: crossing the global divide

You can sit in front of a browser and the whole world is at your fingertips. You can find anything in Google. The whole idea of the new generation of web tools is that now everyone can create content, everyone can join in – you can blog, network with the world on Facebook and share your photos on Flickr. Everyone can now have a web presence! Everyone's views and concerns can be heard. Except they can't....trouble is all that content is actually only from people like you and me.

Most people in the Third World don't have direct access to electricity never mind computers or the web. Many can't read or write, never mind blog. As a result westerners often jump to conclusions about what needs fixing. Western technologists often barge in with their solutions without actually finding out what the local problems really are first. When solutions aren't based on an understanding of what really matters, things can easily be made worse rather than better. Great that you turned up with hi-tech water pumps to give people drinking water. Trouble is you didn't think about what would happen after you left and there were no affordable spare parts. It's such a shame when you were only trying to help. If only you had understood the lives of those people you were trying to help better. If only the worldwide web did reflect everyone's concerns and worries. If only those people where the social injustice was the worst could create their own web content to tell their stories...

Matt Jones of the University of Swansea has been telling us about a project to change all that. A project intended to really make a difference.

So, picture this. An elderly Indian lady is sitting on the steps of her house. She's throwing scraps to a group of scavenging, chattering chickens. Now, listen carefully, very carefully: There's the bustle of village life, the shouts of children, an occasional motor and, ever so gently, occasionally, a soft word. She's talking to the birds. Is she in urgent need of computing resources? How might technology enhance her life? How would your solutions fit this context? Perhaps they would, perhaps they wouldn't, but if not perhaps just hearing her story will spark an insight into what you could do that would help.

The power of such vignettes to engage, stimulate, and challenge is the motivation behind the StoryBank project (www.cs.swan.ac.uk/storybank). The StoryBank team's aim is to build an online collection of stories: stories told aloud, stories told in pictures, stories that capture the essence of local community life. The stories will give engineers and designers insights into the storytellers' needs as well as relating their prior experiences of technology: what went right, what went wrong, how the technology actually fits into their lives.

The StoryBank isn't focused on the past. The intention is for it to be an engine of change to shape the future. If it works it will help engage people in constructive conversations about the real issues that matter. It will help engineers really understand the needs of cultures that aren't their own.

By using spoken stories illustrated by still pictures, the StoryBank team hope to give a stronger voice to people who can't read or write, never mind use the Internet directly, to tell their stories. The first problem to solve is how to collect the stories. Some simple reliable technology is needed that is cheap, easily transported around, works anywhere, and records both sound and images. Sometimes the time is just right. Camera phones are the answer!

Now the StoryBank team are gathering stories, building the technology to make them easily available and testing the ideas. So next time you snap a photo of the antics going on at a party on your mobile and send it round your friends, remember the same technology may also be helping some of the most disadvantaged people too. Really making a difference.

As Dark as Night

If you look at a TV screen when it's switched off, it looks a sort of dark grey. When you're watching the latest sci-fi space movie blockbuster, the starfield background looks a deep, dark black, so how is this possible? Does your TV suck light into the screen to make the grey screen look black? No, this effect is caused by the way your brain computes brightness in the world around us. The brain doesn't simply measure the amount of light entering the eye; it does a comparison with other sources of light in the image. In fact it's the bright stars that cause the space between to look darker, to look black rather than grey.

There is more to space
than meets the eyes!

Enterprise in space

Think you can come up with an out-of-the-world idea to make money in space? Then why not get involved in the national Space Enterprise Challenge competition. Open to UK students aged 11-19 you need to come up with a cunning concept and a sound business plan. For more details warp over to www.cs4fn.org/competitions/ The closing date is June 20th 2008.

Watch out for the next issue of cs4fn on computer science in space.



Under pressure

What have a jellyfish and some computer science techniques developed to allow computers to 'see,' got to do with helping people recover from dance or football injuries?

Bushra Akhtar, a computer scientist and biologist at Queen Mary, University of London is working on it!

You are made up of bits, little biological bits, called cells.

When your cells work the way they should it's great and you're healthy, but in some medical conditions your cells undergo a change that stops them working the way they should. An example is with cells you use every day walking to school, dancing, playing football or riding a bike. These are the cells in your cartilage, and they have a tough, high-pressure job to do. Cartilage cells live in your joints: the joints of your arms, your legs, everywhere. Every step you take, your cartilage squashes and so do the cartilage cells inside the tissue. The cells detect these 'mechanical signals' and respond by producing all the things that keep the cartilage tissue healthy and strong. If this process goes wrong due to disease or if you stop squashing your cartilage (like weightless astronauts in space) then your cartilage breaks down.

The study of the way cells change shape when they are pushed, prodded or squashed, together with the biological consequences, is called 'mechanobiology'. It's a fascinating subject that involves understanding how all the clever engineering and physics inside a cell work, letting it keep its proper shape and

perform its biological function. Examining the cells' 'mechanotransduction' (how cells 'sense and respond' to mechanical forces like pressure) allows us to understand what goes wrong, and with that knowledge we can start to develop new medical treatments to help those who suffer these debilitating diseases.

Let's have a look

One of the ways researchers are trying to understand how cartilage cells work is by looking at how they respond to being squashed. A clever technique called 'confocal microscopy' is often used to get a picture of what's happening in these tiny cells. In a normal microscope you flood the

slide under the microscope with light and have a look at what bounces back, magnified through the lenses. This conventional way to look at cells has been very useful in medicine and biology. The problem, though, is that you see it all from one particular viewpoint. Cells are actually 3D. It would be useful to be able to see the detail in a slice through them.

Enter confocal microscopy. It uses the laws of optics, and a tiny pinhole, to remove the light that's coming from anything other than the depth you are interested in. In effect it lets you see the information from a slice through your big and bulky cell. Using this method you can see what's happening inside a cell at some particular level. Like a

lift moving up through an office block, you can select the level in the cell to look at the information from. What's more you can put all those slices together to make a 3D image that you can view from any angle.

It's all about labels and jellyfish

Clever as it is confocal microscopy can't see the invisible. If you are interested in what gives a cell its mechanical properties, you have a problem. The cytoskeleton (the 'skeleton of the cell') is what gives the cell its shape and physical strength like your skeleton does for you. Unfortunately it can't be seen using normal microscopy.

We need another clever idea. That's where labelling comes in. You can add chemicals to the cells that stick to the parts you're interested in, like the cytoskeleton. These labels can be created by a bit of genetic manipulation. Jellyfish, those floating blobs that can spoil a trip to the seaside, have given science a really useful tool. Some species of jellyfish 'glow', just like false teeth or dandruff do at a disco. You can take the jellyfish genes that make them glow, (called 'Green Fluorescence Protein' or 'GFP'), and add this genetic instruction into other cells so that parts of them glow in the dark when a light is shone at them. You

can see where the jellyfish marker is attached and with confocal microscopy take slices through this new luminous label, showing up the parts of the cell you're interested in.

Actin apart, what's the bigger picture?


Actin is a part of the cell cytoskeleton that helps give it its shape, and it can be labelled with the GFP jellyfish gene. Being able to see how the shape of the actin in a cell changes under mechanical stresses is a first step to being able to understand how cartilage cells work (or don't work) properly.

Cells going down the tube

You can get the images, but what do they tell you? In comes the computer science, to help measure how the cell is doing and so understand the way a cartilage cell changes. Enter the work of computer scientist and biologist, Bushra Akhtar. She wants to help us find out the answer. She is working towards her PhD developing better ways to produce improved images of cells subjected to mechanical forces, and to create computer software to let us automatically measure these changes. She is exploring the use of a method called 'Digital Image Correlation'. It's an image processing method that looks for parts of an image that are the same; we say that same-looking parts are correlated. One use of the technique is to automatically spot suspicious behaviour, such as a passenger leaving a bag on a busy station platform.

By looking at slices produced by confocal microscopy her techniques can compute how the labelled cell parts change over time. They can extract precise numbers for how the cell parts move and deform when they are subjected to a mechanical load. Bushra sucks cells into a thin glass tube, called a pipette, and calculates how much different parts of the cell structure change as they are pulled in. This gives useful information on how cells respond to pressures and may in future help us understand mechanotransduction better. That will hopefully lead to new medical treatments: all thanks to a jellyfish, a very small hole in a microscope and some novel computer software.

It could just help shape the future of cells!





Turning silicon into gold

The ability to turn things from one material to another is the stuff of legend. Everyone knows of the Midas touch: named after the unfortunate king whose touch turned everything to gold. The Queen of Narnia's most powerful weapon was her wand. Why? Because it turned people to stone. Alchemists too, long looked for ways to turn lead to gold.

Whilst in reality such transmutation isn't possible, in the virtual world there are lots of situations where you might want to do it. Take that Narnia example, for instance. How do you do you make an object turn to stone on the screen? The ideal is that it could be done in a 'photoshop' way: the digital artist indicates a real object in a scene to transmute and the computer instantly changes it to realistically look like stone – or any other material desired for that matter. That way the film makers can make objects out of whatever is most convenient and use the digital wizardry to get the final effect – changing their mind at the touch of a button.

Erik Reinhard of the University of Bristol tells us about a way that Erum Khan, with help from Erik, Roland Fleming and Heinrich Buelthoff came up with to do it.

In computer graphics, we are normally concerned with the generation of images on the basis of 3 dimensional (3D) models that

are created by the computer. We create a scene full of geometric shapes, then specify where the light is coming from, and also specify how each object in the scene reflects the light based on its material's properties.

Once we have a scene made of 3D virtual objects, we can then use all this information to simulate how light would travel through it. This is called 'rendering' the image. It involves simulating photons (individual particles of light). You take lots of imaginary photons and trace their route as they bounce around the scene to determine which would be seen by a viewer of that scene. This technique is called 'ray tracing'. It is as if we just place a virtual camera in our virtual environment, and then compute what image that camera would capture.

The problem with ray tracing is that it's not always practical to painstakingly model a 3D environment, and simulate the very



large number of photons needed. It takes a lot of creative effort by skilled artists to define the 3D environment, and then an enormous amount of computer time to render a good image. Of course if we are making a film, it's not just one image we are after but one for every frame of the film.

For some applications, it's easier to just photograph a real environment, and then edit the photograph afterwards to achieve the desired effect rather than digitally create the whole thing from scratch. For instance, in the film industry, we could build a set that is mostly accurate (i.e. it is as the director would like it to be). When the director would like objects to be made out of a different material, or we would like the lighting to be different from what is practically achievable we only then use special effects, editing the images afterwards. This could be done to cut costs,

but also to achieve Narnia-like material transformations that are part of the story.

Examples include not only changing the material, but also applying arbitrary patterns to objects, or even making an object transparent. For example the Fairy Godmother might be supposed to wave a wand to magically switch through lots of different patterns on the dress the Princess is trying on for the ball. It would be a drag if the actress had to be filmed in 20 different dresses all individually made to be identical apart from the pattern. Or maybe the slippers the Princess is wearing need to be glass. Pretty uncomfortable if the actress had to dance in real glass slippers!

On the technical side, the problem is that to change the material appearance of an object based on a single image, we really need to know both the 3D shape of the object and how it is illuminated (the colour of the light, how strong the light is, and from which directions it comes). With this information we can then use the standard computer graphics techniques.

Unfortunately, to compute the 3D shape of an object from a 2D image, we need to invent information that is not actually there in the image. This is similar to the problem that our visual system appears to be solving though: the image that falls on your retina is 2D, and yet somehow you perceive the world as 3D. We could, therefore, try to mimic the human visual system, and build a computational model of it: create a computer program that does the same thing. In its simplest form, this is exactly what our team at the University of Bristol, working with researchers from the University of Central Florida, have done. We have also developed ways to apply various simple transformations on the

image to approximate the way the object is lit.

Once we have constructed a 3D version of an image and can also illuminate it, we can then render a new object with the same geometry and illumination, but with a different material. The reason that the results look anywhere near plausible is that humans are bad at predicting exactly how light bounces around complicated scenes. For example, we are obviously very good at detecting material properties like transparency, but are actually very bad at then predicting the exact way light passes through a real transparent object. That's one of the reasons we fall for illusions like the stick that seems to be bend when we put it into water. Our brain can't compensate for the bending of the light to make us see what the stick is really like.

In our system for changing the materials of objects in images, we actively exploit these limitations of human vision to arrive at results that are physically wrong, but perceptually plausible. If you took a picture of an actual transparent object and placed it side-by-side with a picture of the fake one you would be able to tell the difference, but if you just saw the fake one you would be unlikely to notice it wasn't real.

At the end of the day techniques like this, even though faking it, could turn computer silicon, at least, into gold. After all special effects are big business. Techniques that make digital effects more realistic and save money for the producers of a film or computer game are likely to make big money. That's why computer graphics is such a thriving research area.

Now where's my wand...




I am Legend

We've all seen the movies. A virus mutates and rapidly spreads. Will there one day be a real Will Smith... a real hero, 'I am Legend' style?... The ultimate last man standing?

Outbreaks of disease spread in unexpected and poorly understood ways, and their growth can be explosive. They can kill large numbers if they do...and diseases don't respect country boundaries of course. A disease starting in the UK could kill anywhere in the world. We can't teleport 'Jumper'-style but we can easily jump between continents in a matter of hours, taking disease with us. The last big flu pandemic that started in Hong Kong in 1968, killed 34,000 in the United States. The 'Spanish flu' of 1918, which this time started in the US, killed around 17 million in India...that was before airliners and jet set holidays for all.

To stop the spread of an epidemic before it gets out of control, health workers must find out who infected patients have been in contact with – they must identify the next wave of potential disease transmission. If these contacts have themselves become unwell, then the disease is still spreading. You still have a problem!

The World Health Organisation (WHO) have the crucial task of managing outbreaks of



people and FIMS helps manage the follow-up, the tests, the samples, the reports. The management of an epidemic generates a huge amount of this kind of difficult-to-co-ordinate administrative paperwork. Painstaking detective work is then needed as health workers collate all of the information they gather each day, and make decisions about priorities for the next day. With the huge amounts of data involved it can be hard for the health workers to see the wood for the trees. This is where computer science in the form of visualisation research can help.

Conrad Hughes and Stuart Anderson at Edinburgh University School of

infectious diseases like Ebola and Marburg Haemorrhagic Fevers or Bird Flu. To do that they could benefit from computerised help: the Field Information Management System (FIMS) is a tool developed by the Alert and Response Operations team at the WHO to do just that. WHO workers use FIMS to track cases. They visit people who have been in contact with the infected

Informatics are working on ways to visualise the spread of epidemics. Florence Nightingale showed that pictures could help persuade decision-makers (see page 19) but their work shows that turning the data into simple pictures can also make it easier for health workers to understand how a disease is spreading in the first place.



Rebel with a cause

Florence Nightingale, the most famous female Victorian after Queen Victoria, is known for her commitment to nursing especially in the Crimean War. She rebelled against convention to become a nurse at a time when nursing was seen as a lowly job, not suitable for 'ladies'.

Transforming the data into simple graphs replaces hours of tedious labour on the part of hard-pressed field teams by a few seconds of computation...as long as an appropriate way to visualise all that data is found.

By representing the people (and places) as nodes (the points) in a graph (see page 7), with the lines between them showing who has been in contact with who, the spread of the disease can be followed. Workers can then easily see how each case connects so that people who may have been infected can be treated in time. It's also simple to spot patients about whom not enough is known. That is crucial so that information can be gathered quickly about the people they might have infected. Building on experience during outbreaks of Ebola in Uganda and Marburg in Angola, there is now interest in the system from a whole range of governments who may one day need to use a system like this to manage an outbreak in their own countries – including Britain.

Spanish Flu disappeared on its own within 18 months. Future epidemics may not materialise in the first place because we know how to control them...due to the hard work of health workers and some cunning ways to visualise data. I may not need to *Be Legend* after all.

She broke convention in another less well-known, but much more significant way too. She was a mathematician – the first woman to be elected a member of the Royal Statistical Society. She also pioneered the use of pictures to present the statistical data that she collected about causes of war deaths and issues of sanitation and health.

Soldiers were dying in vast numbers in the field hospital she worked in... not directly from their original wounds but from the poor conditions. But how do you persuade people of something that (at least then) is so unintuitive? Even she originally got the cause of the deaths wrong, thinking they were due to poor nutrition, rather than the hospital conditions as her statistics later showed. Politicians, the people with power to take action, then, and probably now, were incapable of understanding statistical reports full of numbers. She needed a way to present the information so that the facts would jump out to anyone. Only then could she turn her numbers into life-saving action. Her solution was to use pictures, often presenting her statistics as books of pie charts and circular histograms.

Whilst she didn't invent them, Florence Nightingale certainly was responsible for demonstrating how effective they could be in promoting change, and so subsequently popularising their use. She undoubtedly saved more lives with her statistics than from her solitary rounds at night by lamplight.

Data visualisation is now an important area of computer science. As computers allow us to collect and store ever more data, it becomes harder and harder to make any sense of it all – to pick out the important nuggets of information that matter. Raw numbers are little use if you can't turn them into knowledge, or better still wisdom. The right kind of picture for the right kind of data can do just that as Florence Nightingale showed.

'The Lady of the Lamp': more than just a nurse, but a remarkable statistician and pioneer of a field of computer science...a Lady who made a difference by rebelling with a cause.



Back (page) in the picture

Artists come and artists go. Styles change, fashions change. Sometimes it is new materials becoming available that drive the styles. Artists in centuries past painted or sculpted. Now they can work digitally too – where the opportunities to experiment are enormous. A new generation of ultra talented digital artists are making waves in the world of contemporary art, combining creativity and design with mathematical and programming skills. What will become of today's art? Who will be the artists of the future?

Back to the original

Digitisation is seen as a great way to save art and historic photograph collections. After all you can then make the images available to all on the Internet and stop worrying about all those dirty fingers handling your precious originals...or maybe after digitization you no longer need the original anyway. Hmm! Would you really throw away the Mona Lisa because you had a photo of it? The same goes for historic photos too. A digitized copy, however accurate, is missing lots from the original, both in concrete terms and in the values it contains. Even the technology used to take or make the original is historically important in defining the object.

Lesson for the future:
Look after the original!

Backup

You can look at a Leonardo picture today almost exactly as he painted it (though time and light fades the colours). Today's digital artists create in a world of formats. Their images are JPEGs, their moving sculptures are MPEGs. All these formats are just codes that store the information in images in a particular way. These formats need computer programs to translate the instructions into the image we can see and enjoy. The question is how will these formats go forward? Until the 1950's records were still being made at 78rpm (the turntable moved at 78 revolutions per minute). Now you can't buy a turntable that does this. The format has died. So will the formats of art today still be viewable in a hundred years' time? If not then we will lose images and creativity from our generation, simply because the artists chose to go digital.

Lesson for the future:
To preserve digital art, you need to preserve more than the image!



Back in action

Originally created as an exercise to learn how to program, Sodaconstructor became an international art phenomenon. The Sodaconstructor (www.sodaplay.com) website is a construction kit for interactive creations using masses and springs. By altering physical properties like gravity, friction, and speed, all manner of weird and wonderful models can be made to walk, climb, wriggle, jiggle, or collapse into a writhing heap. It went on to spawn a set of other applications like Sodarace and Moovl and was so successful that it won the 2001 Interactive Arts BAFTA Award, proving that it is indeed a top arty cyber celebrity.

Lesson for the future:
Digital art doesn't have to stay put.

Back to basics

The first episode of cult TV cartoon series South Park took about three months to make as all the bits of the characters were cut out from paper and animated by hand, frame by frame. Today the series uses state of the art computer graphics to duplicate the original, amateurish look, allowing shows to be produced far more quickly and so be able to

respond to current events. The system allows the artists to decide on key frames in the animation such as when Cartman moves from the left of the screen (key frame 1) to the right (key frame 2). Computers can then fill in the sequence of images between these key frames in a process called tweening, short for in-betweening.

Lesson for the future:
Computers can help behind and between the scenes in animation.

cs4fn is written and edited by Paul Curzon and Peter McOwan of Queen Mary, University of London.

For more on the topics in this issue go to the Magazine+ section of the webzine www.cs4fn.org/magazine/

For a full list of our partners see www.cs4fn.org/

cs4fn partners who contributed articles to this issue: Queen Mary, University of London, the University of Bristol and the University of Edinburgh.

cs4fn is funded by EPSRC and Queen Mary, University of London. Additional funding is provided by Microsoft, ARM, Intel, Stop Motion Pro and Chromacolour (www.chromacolour.co.uk).



Microsoft

ARM

Intel



**Queen Mary**
University of London

Passionate about computer science?
www.cs4fn.org