A system for generating colours and images using 'one-time' cryptography.

David Upton, MA

artist
www.codedwalls.com
e-mail: cogible@gmail.com

Cubist self-portrait: David Upton

Abstract

A generative art system which makes colours and shapes out of other images, using mathematics to combine them, allows the user to find new colours and shapes in the previously unnoticed backgrounds, and to transcend existing views.

I have tried not to make this a mathematical talk: Instead, please just look at the colours I will show you in the next few minutes, and enjoy them.

1. How does an artist achieve?

I was fortunate enough to see a major Turner exhibition at the Tate Gallery in July this year.

Turner achieved extraordinary effects with light and colours, effects which have left painters wondering ever since how he did it. Even in his own lifetime, people used to ask what his 'secret' was, to which he grumpily replied: “The only secret I have got is damned hard work!” (Townsend, 2019, p7). In Turner's day, all materials were hand made: paints, papers, brushes, pencils. The variety was considerable, sometimes even between different batches of the same product from the same artisan. A large part of any conscientious artist's life must have been spent finding out what tools and materials best suited his or her vision.

Turner had certain advantages. Most important, he was a genius. He combined an exceptional vision – what he wanted to see on his canvas – with exceptional ability to make paint do what he wanted.

Trying to understand how Turner used light, shape and colour made me think about my own practice as a computational artist. What was my equivalent to Turner's 'damned hard work?

For the digital artist nowadays, life seems very different. Digital colours are precisely specified and pure. Monitors and projectors differ, of course, but you can buy systems to standardise and calibrate them. Only when you get to making a print of an image – if you get that far – do you encounter a degree if variations in paper types and inks: and even then your options are far fewer and the results more standardised.

At the theoretical level, systems use precisely specified digital values, expressed in pixels. Provided your systems are properly set up, one digital image is precisely the same as any other. This may be the case, but it is more complicated than it seems. There are different 'colour spaces', ways of defining colours, such as RGB and HSV. There are different ways of storing the data in a file (such as .jpg and .png files.)

My own interest in digital light and colour comes from the fact that it is, ultimately, mathematics. It can be manipulated at a mathematical level. A painting in shades of green can be converted to shades of red at a stroke. The machine is colour-blind: it sees only numeric values.

So this paper is largely about my own generative investigation of this mathematics/ colour interface.

The other issue, of course, is that colour is not a singular thing. Colours on their own are strange, unusual, orphaned things: normally they occur next to each other, and what really affects us is their interplay. This was analysed by colour theorists like Chevreuil, and was important in the work of the Impressionists. It has been played with by painters and experimenters like Albers, as well as by Turner.

Colour systems based on on a combination of 3 x 8-bit values only offer 16,777,216 colour options. These are so close to each other that a human eye would not be able to detect the difference between one colour and the next. Take as an example a RGB colour, expressed as three numbers or 'channels'. Black is 0,0,0. It is possible to define another colour, which you might call 'slightly red black', as 1,0,0, shown here to the right of the 'pure' black. However, it is not easy to tell them apart: in effect we have more colours than our eyes need.

Combining colours greatly increases the mathematical possibilities.Once you start putting blocks of colour beside each other, the number of possible combinations between just two colours becomes absurdly large – back from clinical digits to the normal degrees of complexity of the human world, in fact.

Similar complexity comes from the universe of possible shapes. John F Simon Junior's 1996 work 'Every Icon' (Simon, 1996) attempts to draw every possible icon within a space 512 x 512 sqaures large. Each of the 262,144 squares (think of it as a pixel) can be either black or white. Simon calculates it will take 'several hundred trillion years' to formulate and display every possible combination.

Faced with this, the poor human artist has only two options. One is to see and reproduce something, however abstractly; the other is to generate new forms using various techniques such as randomness, feedback, and external contraints, and see what happens. You end up with a pile of images that you could not have imagined by yourself. Some are dull, some are ugly, but a surprisingly large number give you pleasure. As Vera Molnar said:

“there is a thing which can replace intuitions - it's randomness... that will show you billions of possibilities, which you, with your limited imaginations, couldn't have thought of. So it enriches the senses. Therefore randomness has a lot of importance to me, but not in the way of dadaism. It's not to say anything can be art. On the contrary, it helps me to better find what I like. Because when you work with intuition, you do ten, twelve, fourteen tests. At the twentieth, you're tired, and stop. With computers, you can first open the entire specturm, and say this is the part that interests me, and not the rest. So you place the focus, and develop all possibilities within. Afterward, you'll find the interesting part is over here. So you get closer. Its a paradox, but the people who argued at the beginning that using computers dehumanises art, the opposite is true. Because it's thanks to all this technology that we can get very close to what we have imagined, that we might not have found otherwise...” (Molnar, 2019)

This talk is aboutd such a process.

2. Unusual Raw Materials

In 1955 French painter Jean Dubuffet moved from Paris to Vence, for the sake of his wife's health.

Vence is the sort of place people think would be ideal for an artist: picturesque, great light, lots of beautiful views. The trouble with this sort of place, though, is that it is difficult to paint or photograph it without producing 'chocolate box' images, or repeating what every other artist has done there before.

Dubuffet wanted to incorporate the place in his paintings, but not in the conventional way.

So he adopted an original idea. In his 'Texturologies' series, he painted like Jackson Pollock, throwing paint on to the canvas, but he also added local soil, so that Vence literally became a part of his painting.

There's a good photograph of him surrounded by plastic buckets of local soil, using them as a palette.

Rather than dig up my garden, I decided to use pixels instead of soil, and dug through my photographs instead. I chose images that would otherwise be rejected as too dull, in shape or in colour.

3. How to combine images

This led me to a fortunate discovery. In cryptography, the art and science or producing and using codes and cyphers to conceal confidential information from hostile readers, you might distinguish two broad approaches. One way is to build an algorithm. This might be a very simple one, such as the well-known Caesar alphabet, in which every letter is replaced the letter one number after it in the alphabet. (Or 2 or 3 numbers after, or whatever you want.) So 'abc' becomes 'bcd'. If your correspondent knows the algorithm, all that is necessary is to replace the letters in the cypher text by the ones immdiately previous (or 2 or 3 numbers before, etc.,as agreed), and out pops the original message. Apparently this simple algorithm baffled the Romans, though it would not withstand the NSA or GCHQ for long. The algorithm is deterministic – it must always give the same result – or it wouldn't work as a cypher because you couldn't decypher it.

The simpler alternative, is the 'one time pad'. This is theoretically the most unbreakable cypher in existence, but there is a cost. If you want to send a 500 character message, you also need a 500 character key, a collection of random letters. You add the one to the other. (Using some simple standard algorithm). At the receiving end, your correspondent must have the exactly identical key, and can then subtract this key from the message and the plain text is left. Usually the key is simply a set of random characters – numbers or letters. Generating them is time-consuming, and not easy: randomness is a hard concept!

The problem is made worse because you need a lot of key material. You can only use key material once. Both sender and receiver must have identical sets, for every message that is sent. Generating, securely sharing, and securely storing the key material is difficult, especially if you consider the amount of material that governments and corporations now routinely need to encrypt. If you use the same key material twice, the cypher is theoretically vulnerable. (And famously was actually broken, in the well-known case of the Venona traffic, which enabled the West to read encrypted Russian spy traffic relating to atomic spies in the 1940s, and to the Cambridge spies Kim Philby and Donald Maclean.)

4. The maths of images

However, we are interested in art, not spying. Luckily we have an excellent and virtually endless source of 'one time pad' key material that can be easily passed and used to interact 'randomly' with images. This source is, quite simply, other images.

Each image is a list of lists, a two-dimensional array of numbers. It's easy to make them exactly the same size. You can specify the exact size, in pixels, and layout (height vs width). In fact these are often standard: two images created the same way will often be identical in size.Two images can be combined in various ways, using the first as the 'message' and the second as the 'key', resulting in a 'cypher text' which is actually a third image. This can then be combined with another image as 'key', to produce a further 'cypher text', and so on. Each new text is a new image.

Visually, each pixel on a screen consists of three 'bits' of data, each shown as one point of colour. For instance, a white page on my screen shows up under a microscope as:

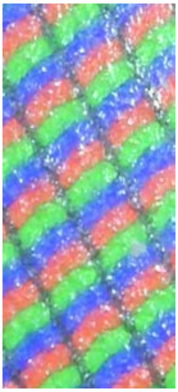

Red, green and blue together form a pixel. If each is 'on' at full luminosity, the result appears white from a distance, as in the right hand image. These are of the same screen: the only difference is the degree of magnification.

A black line on the screen shows up as no colours in the pixel:

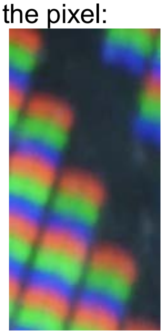

Here we have the additive model of colour mixing: each time you add any light, of whatever colour, to the pixel, you get nearer to white. Black is the absence of any light from the three pixels. (As you know, the subtractive model of colour mixing, which painters and printers use, works the other way. Each colour you add darkens the mix until you reach black).

By way of comparison, these spots of colour are small. How small depends on your screen, and on your system settings. Here is a human hair on my screen to show the size of these settings.

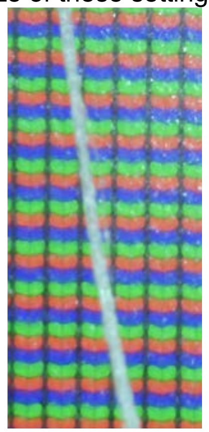

The pixels are small, but not extremely small. There is a threshold below which the eye cannot distinguish them as separate items, and forms an overall perception of a single colour.

There are two key points here. The whole system depends on the human ability to interpolate. We cannot see infinite detail, so we get as near to it as we need and ignore the discontinuities that are actually there. This is the same when we watch a movie: provided the frames change at about 14 frames per second, we see it as a continuous moving object. We cannot actually 'see' colours like yellow. The receptors in our eyes can only see red, green and blue; our brains combine their signals to calculate a value for 'yellow' when we look at a lemon, and this is the colour we 'see'.

Secondly, the critical thing to remember is that these colours we see are just numbers. On a screen, each point of light simply embodies a number between 0 and 255 – eight bits of data.

5. Combining pixels

When you consider a colour just as a list of three numbers, and an image just as a list of these lists, then you can use all sorts of mathematical techniques to combine diffferent pixels, and therefore the images that they make up.

An example of simple blending shows an effect similar to a 'double exposure', when you forget to wind on the film in an old analogue camera. This:

plus this:

gives this 'double exposure'.

In this case, half the simple numerical value of each pixel comes from one image and the rest from the other. The result is as if each shows through the other.

6. Complex transformations

Typically, images are stored as three (or sometimes four) 'channels' of values: all the red values, all the green values, etc.

It's possible to take one channel from one image and mix it with two channels from a second image

Here we start to get interesting colours, but the underlying shape is still visible, as in the previous example of a simple combination.

An alternative is to read the numbers as digital bits and then to compare each set. Formulae for doing this include:

<u>Bitwise OR</u> Each bit of the output is 0 if the corresponding bit of x AND of y is 0, otherwise it's 1.

<u>Bitwise XOR</u>: Each bit of the output is the same as the corresponding bit in x if that bit in y is 0, and it's the complement of the bit in x if that bit in y is 1.

This produces very interesting effects. Here is an image of ice on a glass roof:

Here is the interior of a heated greenhouse, seen through condensation

Blending the two together (using bitwise XOR) gives:

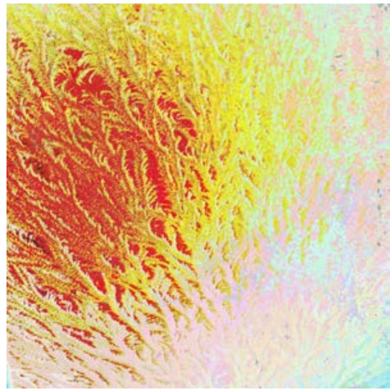

Taking a third image, almost monochrome, of a tree trunk

And blending this with the first blend above, again using bitwise XOR, gives:

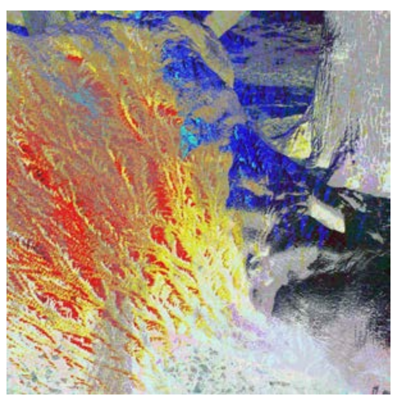

All the colours seem to have sprung from nowhere, yet they are intricate, new, bright and to me at least aesthetically pleasing. Huge detail creates interesting interpolation effects.

In fact, the duller the original colours, the brighter seem the results. Here is another set of examples: an office building wall.

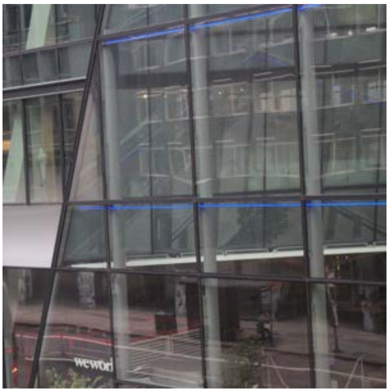

Some flexible pipes and wiring conduits (beside London Underground railway tracks)

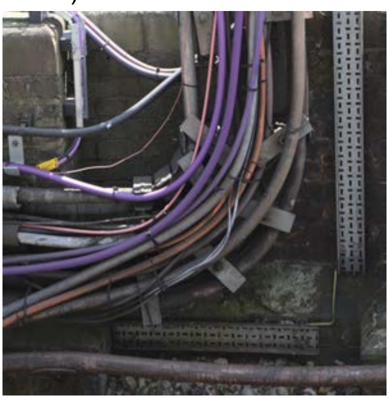

XOR Blending the two produces:

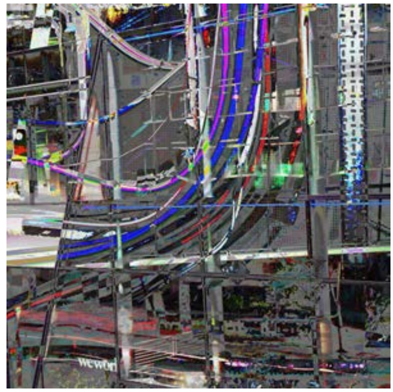

Now take an image of logs:

and XOR this with the previous blend:

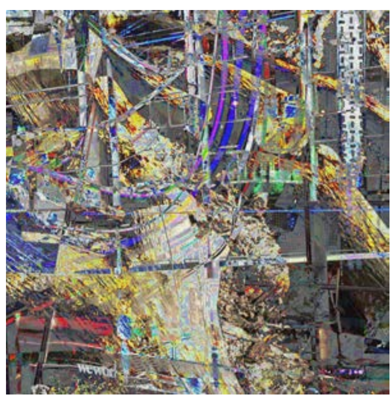

Once again, an image emerges which has

1. new colours, much more varied than the originals
2. new shapes mixed in with the old ones.

Continued experimentation with other images and combination methods develops further variety

7. Is this generative art?

This process is not entirely algorithmic. Much of the result depends on my own selection of images and combination methods.

In addition, what I do is deterministic, in the sense that combining the same two images in the same way should produce the same result.

I find that choosing the images and combination methods gives a degree of creative control which I find useful. However, it would be possible to choose images at random, and to choose combination techniques at random.

As Andy Lomax said in a recent paper, 'Hybrid Creativity', “the computer can become an active assistant in the process of discovery as well as being a medium to work with, enabling creative exploration with systems that the author previously found overwhelming”. (Lomas 2018)

My own creative exploration at the moment is focused on the complexity of the result, producing and juxtaposing new colours.

For example, take this image of clouds seen from an aircraft.

I rotated it by 90 degrees, re-combined it with the original,

and then used a filter, which can 'sharpen' or blur images, in this case give an embossed effect, as follows:

and finally rotated and combined it again.

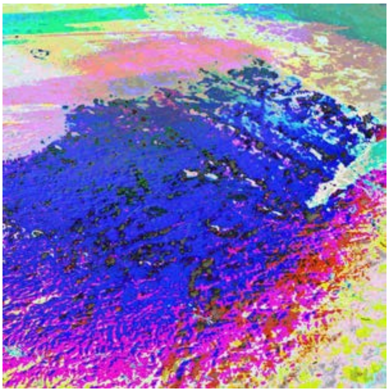

The process of creative exploration involves, first of all, taking suitable photographic images. These have to be fairly large. At the moment I aim for 3000 px by 3000 px square images: some 9 million pixels, each containing 3 or 4 8 bit values, either 27 or 36 million numbers. The detail is important if they are to be printed out at any decent size. The square format is so that they can easily be rotated.

I've found that the most useful images are often the ones you wouldn't look at twice: there are no predominating shapes, few bold lines, and a great many subtle variations of colour, often quite dull colours. Mud, water, condensation, clouds, are ideal.

For another type of picture it can be quite fun to combine images with many straight lines – such as blocks of flats and office buildings.

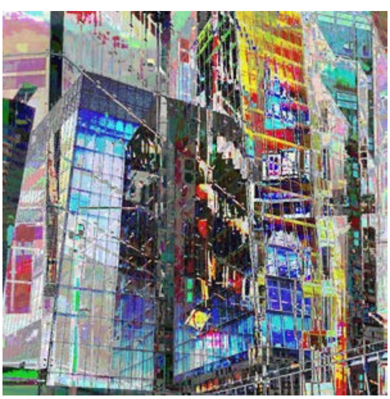

Or to work with very 'busy' images:

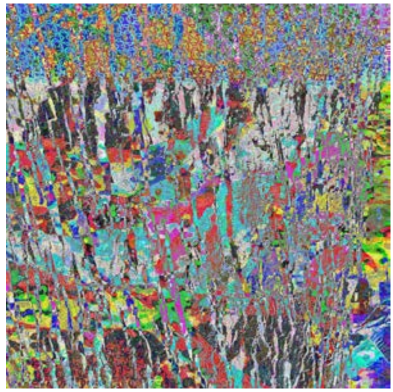

or to combine broad sweeping patterns with large amounts of detail.

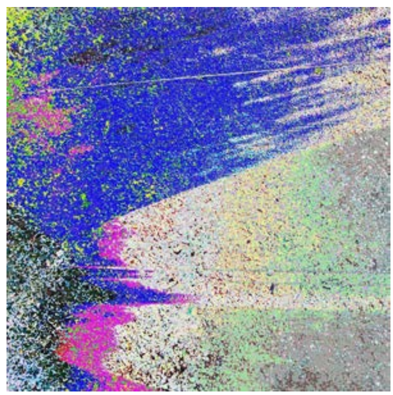

I'm not sure where to go next. At the moment I am trying to build a 'front end' which will allow me to select images either by hand or at random, and then to select and combine techniques.

I'm also experimenting with printing techniques, using commercial art quality printers.

For the technically minded, all programming was done in Python using OpenCV, and the Python Imaging Library (PIL or PILLOW) all of which are free add-ons. Both are well documented on the internet, and OpenCV is also well covered in 'Mastering OpenCV4 with Python' by Villan, Packt 2019.

As Prof Soddu says, “Designing is, in fact, this: activating a logic of development capable of controlling the evolution of the system towards a goal. The difficulty consists in the fact that we do not yet know this objective. It is true that we know some of its attributes, such as to define, in negative, the degree of quality, but we do not know how these attributes can be expressed in the artificial form that we are creating.... To design is, therefore, to control a dynamic process of development without knowing exactly where this development will lead, (Soddu and Colabella, 2020, p 181)

Summary

In summary, I've tried to produce a system which takes the normal and the everyday and turns it into the exotic and the intoxicated. I enjoy colour and I like to be regularly pleased and sometimes even amazed by the colours this system can produce

References

Lomas, A: “On Hybrid Creativity” Department of Computing, Goldsmiths College, 2018; https://doi.org/10.3390/arts7030025

Molnar, V: interview from https://vimeo.com/372579247, from MuDA (Museum of Digital Art in Zurich)

Siimon, J F: “Every Icon”, 1996: see https://v2.nl/archive/works/every-icon

Celestino Soddu, C and Colabella, E: “GENERATIVE ART & DESIGN, Theory, Methodology and Projects”, english edition 2020 Domus Argenia Pub.

Townsend, J H: “How Turner painted: materials and techniques”, Thames and Hudson 2019