Digital Image Space: An Infinite Finity

Ruminations on the space of all possible digital raster images

Ray Gardener, Daylon Graphics Ltd. (rayg@daylongraphics.com)

Last updated: July 31, 2015 (10th edition)

Imagine that you could search through the set of all possible digital images, at, say, HD resolution (1920 x 1080 pixels), in 24-bit color. This is a tempting idea because, on its face, a digital image has a finite resolution and therefore the number of possible images is also finite.

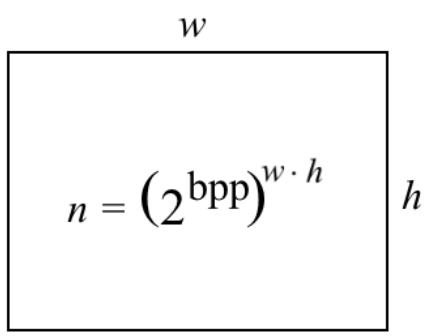

For any digital image of resolution *w* by *h* pixels, the number of possible images (*n*) is two raised to the bits per pixel raised to the power of the image's area. A 4 x 3 black-and-white space, for example, would have 4,096 possible images.

The second part of this temptation is that in that set, there would not only be every picture you've already seen, but every picture you've still yet to see. In fact, *every picture* — for all times and places — is in that set. Not just real photos either, but every painting, illustration, and otherwise fictional bit of imagery as well. There would be pictures of people and places both real and imagined, in nearly countless variety. Pictures of text, too, so every page of every story is also available; every math formula, every scientific discovery. One tiny part of the set has every variation on you blowing out the candles on your tenth birthday cake. Countless other parts are near-endless variations on every other event, important and mundane, real and made up. Every possible picture including its inverse (and filtering by every other means) is in there too.

Incredibly, most of this vast image space is just random noise, because the combinations of pixels that are informationally meaningless is far larger than those that are meaningful. And yet, the meaningful part is still mind-boggingly vast.

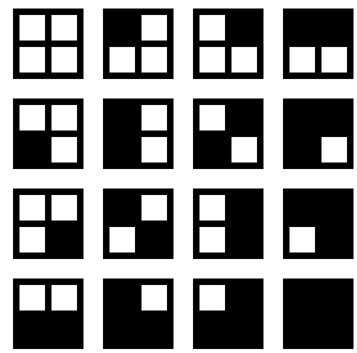

The full set of sixteen unique images for a 2 x 2 black-and-white (1-bpp) image space. At the same resolution, using 8-bpp grayscale would result in over *four billion* images.

How vast? Imagine the set of all possible images for a tiny, low-resolution 5 x 5 pixel 8-bit grayscale image. Each pixel can have 256 different shades of gray, and there are 25 pixels. So the full permutation of pixel content is 256^{25}, or near 1.61×10^{60}. That's in the neighbourhood of the number of atoms in the universe. If we think switching to 1-bpp can lower the count, we still have 2^{25} or 33,554,432 images. A consumer-grade computer can brute force the render of every such image, but evaluating thirty-three million images for usefulness would tax the patience of even a volunteer army, and in the end, it would be for pictures having a resolution and color depth that are useless even for tiny icons. Every possible glyph of a tiny 5 x 5 font would be available, but the resolution isn't enough to show words, let alone stories or formulas. And simply adding another column to explore a 6 x 5 space enlarges it to over a billion images. Long before anything practical can be obtained, the space becomes too intractable.

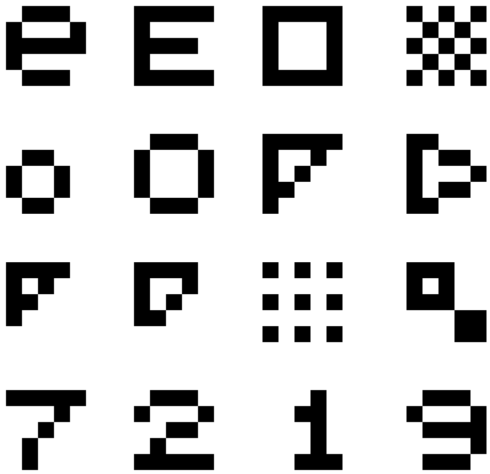

A tiny sample of the thirty-three million possible pictures
in the black-and-white 5 x 5 pixel image space.

But the fact that the bound is finite tugs at the mind. It beckons, as if someday we could explore a larger space, and then make all manner of incredible discoveries. Somewhere in the larger HD color space is every frame of every movie, which also means that there are movies of every way that your life could have turned out. And not just your life, but everyone’s life. Even the lives of fictional characters.

Wrapping Our Minds Around It

The human mind is ill-equipped to handle such numbers, even if we can calculate them and write them down. It’s absurd how large the space is. But it has to be, precisely because it’s finite. By definition, it has to be practically infinite despite being digitally bounded.

To sum up, then, the formula for the full color HD image count is $(2\char94 24)^{1920 \times 1080}$, which is $16{,}777{,}216^{2{,}073{,}600}$, which the Windows calculator app can’t handle. Fortunately, the Wolfram Alpha website can, and it works out to just under $1.5 \times 10^{14{,}981{,}179}$, which is a number with almost fifteen million digits. We could reduce the number by limiting ourselves to grayscale, but you get the idea. Simply writing out the number in longhand would take fifteen megabytes of storage.

Another way to appreciate the hugeness of the space is to imagine a picture, and then all the pictures that are exactly the same except that a single pixel is different. For an HD image, there are 1920 x 1080 = 2,073,600 such pictures simply by positioning the one errant pixel anywhere. If you also allow the pixel to be any color, we multiply further by 2^{24}, giving us almost 34.8 *trillion* images which differ by only a single pixel. In some of these images, you might spot the errant pixel while, in the rest, you would never notice the difference.

It gets worse. How about all the images that match your picture but are off by just two pixels? The number of possible two-pixel arrangements is $2{,}073{,}600^2$ - 2,073,600, or 4.3 trillion *just for a single color*, or 1.2 *octillion* for all colors. And this is just for two errant pixels; there are countless other images that differ except for a stray hair in front of the lens, a bit of dust, a scratch, or some other noise that covers several pixels. Such is the vastness of the space that it can casually use up insane quantities of images for such redundant similarity.

Clearly, we’re incapable of searching through such a space by brute force. Even with every bit of matter and energy in the universe at our disposal, conventional computers running full tilt in the universe’s available lifetime wouldn’t make anything but the tiniest dent. If the space were mapped out to the size of a football field, the amount we could search would be microscopic (and I suspect it’s actually much smaller than that).

Even with a quantum computer, someone (or more likely some large group) has to sift through the images and decide which ones are meaningful. In the meantime, we’ve consoled ourselves with practical alternatives such as drawing and painting, image libraries, photography, modeling, and procedurally generated art (towards that end, I’ve prototyped a procedural art website described in another paper).

Of Images, Image Values, and Encodings

Each unique image (pixel pattern) constitutes a unique **image value**, which is the number that, when interpreted as pixels, renders the image. For example, the binary number 01110 10001 10001 10001 01110 (e8c62e hexadecimal or 15,255,086 decimal) is a picture of a rounded square in the 5 x 5 black-and-white image space (see picture below).

When two image values are the same, the images are identical. Once the interpretative context has been chosen, the image value is the image. One is the numeric data comprising the image and the other is its visual impression.

The interpretation of an image's value number depends not only on the chosen image space (width and height), but on the choice of **encoding**. In the above example, each bit of the binary number is applied to one pixel from left to right and in the top row and working downwards, starting over at the leftmost pixel as each new row is reached. A bit value of zero is white and a value of one is black.

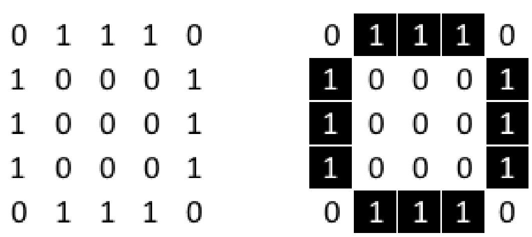

Since half of any image space is a perfect color inversion of the other half, the choice of color interpretation tends to be a perceptual convenience.

To maintain the strong 1:1 value/image correspondance, the encoding must not allow multiple values to represent the same image or vice versa. The same is true of the display system. For example, a video card that simply ignored every *n*th digit of an image value would allow many different values to appear as the same image. The same can also happen in an indexed-color space that uses the same color for different palette entries.

While the row order encoding is traditional, an encoding could just as easily be column-oriented, applying the value digits top to bottom (or bottom to top) first, and then working left to right or from right to left. For the purposes of display efficiency however, one favors a traditional encoding as the image value can be sent more or less as-is to the rendering system.

Other encodings could be byzantine. To generalize, we define an array of 2D points, of *w* times *h* length, that tells which pixel to color as each value digit is read in linear fashion (e.g., from an input stream). As long as each element of the array specifies a unique pixel location, the encoding is valid. If the display process is slowed down, one would see the image “fade in” using something akin to a random dissolve, although countless numbers of non-random fade-in patterns are available.

While the choice of encoding is arbitrary, what matters is that for any given set of different image values, their uniqueness when rendered requires that they share the same encoding. Since the image itself can be made to appear different by changing either the image value or the encoding, it’s easier to use a direct traditional encoding.

However, it may be of interest to some researchers to see how many different images can be made to appear by applying the same image value through different encodings. An obvious realization is that for image values where each pixel has the same color (perfectly flat images), the encoding is irrelevant; such image values can be said to be symmetrical about the encoding axis, or invariant with regards to encoding. As more and more value digits shift from the principal color, the image value becomes noticeably less invariant, until changes to the encoding cause massive alterations in appearance.

“Uneven” Color Quantities

In our equation for the number of images in a space, the term 2^{bpp} is the number of colors per pixel. We raise two because binary encodings are traditionally used. For example, when 16-color images were common, they fully utilized four bits per pixel. But the term can be generalized to numbers which are not powers of two, such as three.

What does a three-color image space look like? Below are the 81 possible images in the three-color 2 x 2 space:

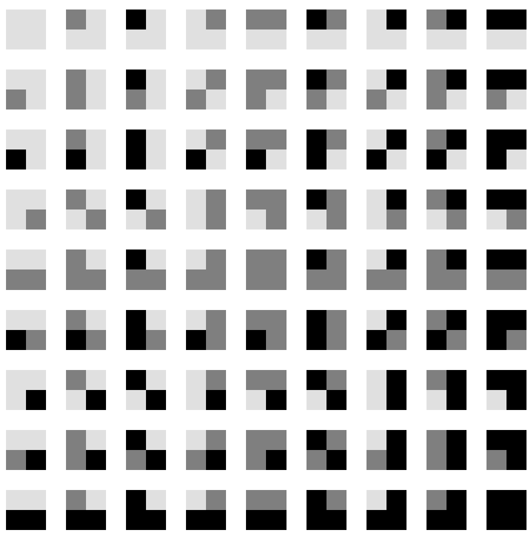

Despite being the same resolution as the 2 x 2 black-and-white space, simply using one more color increases the number of images over five times from 16 to 81. With four colors, the space would have 256 images, and with sixteen colors (the minimum to give the illusion of continuous tone), there are 65,536 images.

Encoding a three-color space can be done by using two bits per pixel, but never using a pixel value of binary 11, limiting the values in each pixel to 00, 01, and 10. All image values in the space are still unique and thus satisfy the encoding correlation rule.

While processing such image values is easy because each pixel falls cleanly on two-bit boundaries, one cannot cycle through the images by merely incrementing an image value with regular numeric addition but must take into account that the integer overflow (and subsequent wraparound) occurs when incrementing 10 and not 11. And since we are only using three of four possible pixel states, our image values are longer than they need to be.

If we solve for *bpp* in our equation, it works out to just under 1.585, which means we would need only about 128 bits instead of 162 to describe any image value in our example space.

The Impact of Using More Color

Here are the 256 possible images in the four-color (two bpp) 2 x 2 image space:

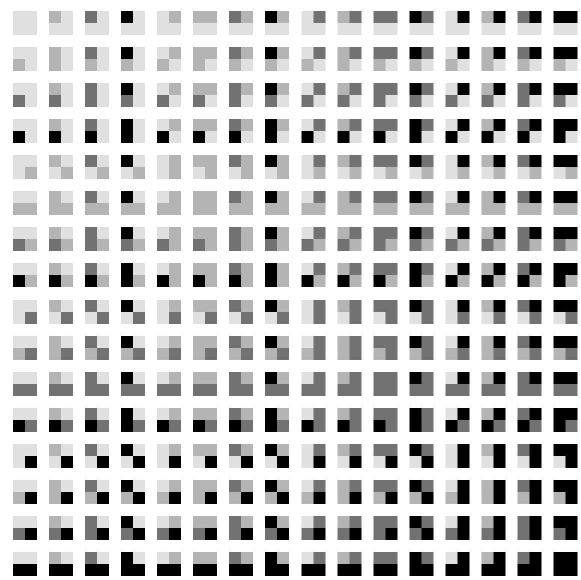

It should be easy to appreciate how an increase in the number of colors greatly increases the number of possible images. We now have over three times the number of images in the three-color space, and sixteen times the images of the black-and-white (1 bpp) space. If we use five colors, we’ll have 625 images.

It’s somewhat ironic to note that the need for extra color is more necessary in low-resolution spaces to provide antialiasing. However, this also means that in high-resolution spaces, color becomes less important, and thus some savings can be realized.

Enlisting Genetic Algorithms

One way we could search for meaningful images is to gather, via statistical analysis, the properties of known meaningful images (such as a large photo library or the frames from movies), and then iterate a loop that spawns various procedures to randomly alter a picture.

The properties could include things such as color difference between adjacent pixels, PSNR measure of blurred vs. non-blurred image, pixel count of edge filter result, edge length, waveform complexity of 8 x 8 subblocks, etc. For the properties, the full gamut of min, max, average, and distribution metrics would be collected.

The loop applies a fitness function that tells which procedures are moving the image closer to having the desired statistical properties, and those procedures are allowed to live and mutate. The process stops when the picture satisfies the properties or comes close but is not making any further progress.

The choice of initial image lets one apply dominating characteristics such as "mostly red", "having vehicle shapes", etc. In an idealized scenario, one could feed in a pencil sketch of a car and get back a reasonable photograph.

As is typical of genetic algorithm approaches, the process would be computationally expensive, requiring millions or even billions of iterations. Fortunately, some of the operations (such as filtering) can be accelerated via GPU.

Where a certain “thematic” result is acceptable, it may also help to gather properties of a small, select group of related images instead of a wide disparate group, as this will narrow the search space.

Enlisting Quantum Computers

Since the space is vast, quantum computation is necessary for the generation of images. Each image pixel would require n qubits, where n is the color bit depth. For a grayscale HD image, we need 8 x 1920 x 1080 or 16,588,800 qubits.

Thus equipped, all the images in that image space can be made simultaneously available to quantum logic processing. For the purposes of discussion, we assume that the difficulties of keeping over sixteen million entangled qubits in a coherent state does not violate the laws of physics (such a quantity of particles, of course, is enough to define more than a few molecules of matter, so we are asking quite a bit from the laws of physics).

Sadly, this is only a first step. The key to practical full image space search is to automate tasks as much as possible. Of critical importance is to automate the judgment of which images are meaningful. While a simple algorithm can decide if an image is merely noise, more sophisticated and intelligent algorithms will be needed to decide which images contain interesting geometry, scenes of people and places and other objects, etc.

Images that are mostly noise, for example, can be detected by trying to compress them. The poorer the compression ratio, the noisier the image. Conversely, images that compress too well contain mostly “flat” areas and are generally devoid of interesting content, so the use of such an algorithm would be to find images that compress somewhat well, but not too much. Other, more sophisticated edge-detection and machine-vision algorithms would then further cull the boring pictures.

However, even if only the meaningful images can be efficiently found, that will still result in far too many images for people to pick from. What’s needed is a way to search through such results and to prune them further. The needle in our haystack is astoundingly small.

The Quantum Observer

What we need is an automated system that, given specified parameters, can identify matching images directly from image space. Since such a system would require quantum computation and marked machine vision abilities coded for quantum logic image processing, and tied to a natural-language front end, I call such a system a *quantum observer*.

An example of how it would work has users submitting queries such as “Blue car, convertible, facing a sunset and mountains, with a young adult male in the driver’s seat and a young adult female in the passenger seat.”

Even for a query that specific, far too many images would result. With an image count, however, the user can fine-tune the query until the result can be displayed as a set of thumbnails. For example, as a first step he could specify the precise direction in which the car must be facing and the exact camera angle from which he is looking at the car.

Some queries are simple to express and can specify precisely one image from the full set. Example: “Perfectly uniform background of 100% magenta.” There is only one RGB value that satisfies this constraint, and only one image that can use this color for every pixel. Of course, such images are not very useful. Another reason the search can be so narrow is that even an image space of a single pixel could satisfy such a request. In essence, one is underutilizing the space. As we want our results to contain more detail, so must our queries become similarly detailed.

It’s important that the query system understand high-level concepts such as “house”, “car”, “Ford Mustang” and the like, because if the user must specify objects in exquisite detail, it will be more efficient to use a modeler, because a description of the search parameters effectively becomes an exercise in modeling. There is an analogous situation in computer programming where if one can describe a problem in sufficient detail, one has essentially written the program.

Minimizing Search Space

A worthwhile question is how low can we set the parameters of image space and still obtain useful images? We’d probably want enough resolution to display a short sentence of text. In addition to letting us sift through every imaginable pun and perhaps some wonderful snippets of poetry, we could discover useful short formulas and equations. However, these fragments of text are precisely that; lacking context, it could be difficult or impossible to determine the real value of each phrase.

Increasing the resolution to allow multiple lines of text to appear does not necessarily solve the problem, because a legitimately valuable text will have a near-infinite set of bogus copycats, and there can easily be no indication which is which. Every misspelt word and erroneous formula reside in image space too. Scientists, engineers, and even poets might be better off using their search results as inspiration instead of truth.

Text would actually be better searched by a quantum observer that explores the set of all possible *letter* arrangements, which is a different space entirely. For that, one would choose an alphabet (e.g., all lowercase English letters and the digits zero through nine with perhaps some punctuation characters), and some upper limit on phrase length. It’s the proverbial billion monkeys on a billion typewriters, but better. The bit depth of each “pixel” (character) is only five or six, and the number of characters only need be a few hundred or a few thousand. A useful digital text space could be, say, 32^{200} or 1.07×10^{301} possible arrangements. Our algorithms are also better at finding meaning in text than in imagery. A quantum computer would need only a thousand qubits.

The Minimal Image

For actual pictures, What if we want to more efficiently search for pictures instead of text?

We mentioned reducing the pixel bit depth to eight bits to use grayscale. Staying at HD resolution, that gives us $256^{1920 \times 1080}$ or $2.47 \times 10^{4,993,726}$. It's far fewer images than the color version, but five million digits is still infinity. We could sacrifice some quality and use 4-bit grayscale, as most people have trouble discerning more than sixteen shades of gray anyway. So now we're down to $16^{1920 \times 1080}$ or $1.57 \times 10^{2,496,863}$.

The combinatorial explosion is driven primarily by the image resolution, so we clearly need to reduce the number of pixels. What if we settle for a resolution comparable to NTSC video, say 640 x 480? That gives us $16^{640 \times 480}$ or $4.5 \times 10^{369,905}$ possible images. Now we're getting somewhere.

How about the ancient PC monitor resolution of 320 x 240? If it was good enough for early video games, maybe we can live with it. That gives us $16^{320 \times 240}$ or $2.6 \times 10^{92,476}$ possible 4-bit grayscale images.

Some early Web video was even smaller, so if we really wanted to, we could try a postage-stamp resolution of merely 120 x 90. So now we're down to $16^{120 \times 90}$ or $3.1 \times 10^{13,004}$. And if we also ditch grayscale and settle for pure black and white, we can get a space of just $2^{120 \times 90}$ or $1.3 \times 10^{3,251}$. A tiny wonderworld of high-contrast icons, stick figures and the like.

Once again, the numbers beckon. Three thousand digits is so much fewer than fifteen million. One can almost hear the siren song of the low-resolution image space calling out to us, begging us to somehow explore it. Maybe we can; how hard can it be to entangle a mere 10,800 qubits? What's heaven for, after all, if not so that our grasp should extend beyond our reach?

Reflections on Users

An interesting side effect of accessing full image space is that one would come across pictures whose content only makes sense when viewed by those of a different language, or even a wholly different culture. When considering pictures of book pages, those pages will also appear in a near-infinity of different languages. Shakespeare in hieroglyphs? War and Peace typeset in the alien Forerunner (Halo) font? Oh yes. There will be a multitude of interesting images that will baffle explorers, intrigue linguists and delight artists.

There will be pictures of machines whose function will be utterly baffling to humans, mountains that hang sideways in alien skies or oceans, entire encyclopedias devoted to porn requiring twenty genders, traffic systems for amphibians, cube-shaped planets inhabited by living icebergs, and so on and so on.

The amount that makes no sense whatsoever will dwarf that which does, but it will never fail to stimulate discussion, and in the long run, that's a good thing: we'll never get bored. For those who don't want to waste time on such mystery meat, the quantum observer will need to recognize it and focus on what its users can understand. It could be considered cultural or "alien noise." A possible algorithm would be a machine-vision system that identifies known objects, and then evaluates what they are doing for violations of physical law.

The Distribution of Useful Images

If we use a typical raw digital image encoding, numbers that increase in value cause the bits of a pixel to change in color, and then as the number overflows what a single pixel can represent, the next pixel in the scanline begins to change color. If we start with a value of one and pad our number with zeros on the left, then the pixel in the bottom right corner will be RGB[0, 0, 1] (which is the darkest possible shade of blue) and the rest will be black. As the value increases, that pixel cycles through its full spectrum and eventually the pixel to its left will turn dark blue. When this happens, the first pixel will cycle all over again until the second pixel increments to the next lighter shade of blue, and so on. Since there are over sixteen million different colors per pixel, the first pixel must fully cycle sixteen million times before the third pixel can start being colored. Even the first scanline of a color HD image would take an ordinary computer forever to render. Once the pixel in the top left corner is reached and finishes its cycle, the entire space has been rendered.

An obvious observation is that the choice of encoding is poor: it starts off trying to render the bottommost scanline and works from the bottom up, leaving the rest of each image black. If we're going to traverse image space, we almost certainly want to pick different values or change the encoding. For example, if we want the top left pixel to not be black, but somewhere closer to, say, 75% gray, then we should start with a value that is very, very large and much closer to the maximum possible value.

This raises the question of how images are distributed through image space, and how a particular encoding or "value hop" affects the distribution. If we could find an encoding where each change in value always leads to the next interesting image, we could find lots of useful content even with a traditional computer. We could experiment with different encodings on small spaces such as our 5 x 5 black-and-white, but even if we find an encoding that "hops" quickly between useful content, there's no guarantee that the same encoding can work in larger spaces. As the space increases in resolution, a suitable encoding could become exponentially difficult to find. But… what if it exists, and we got lucky?

To help get us started, we can examine the numeric values of existing images, and see where they lie along the value axis of the image space, and how they cluster. By converting image values to use different encodings, we can see if they cluster better or worse.

The Image Lottery

An interesting game to play is "image lotto": One picks a number from zero to $(2^{bpp})^{w \cdot h} - 1$, and the number is rendered as an image using any multitude of encodings. If the image isn't random noise, we have a winner. Since typical numbers will contain several million digits, players will opt to have them randomly generated, maybe using a favorite short number as the random number generator's seed value. A central server keeps track of all played numbers to avoid replaying the losing picks.

Image lotto may be worthwhile to play for small image resolutions as the numbers are smaller. In the 3 x 3 1-bpp image space, for example, there are these 512 possible images:

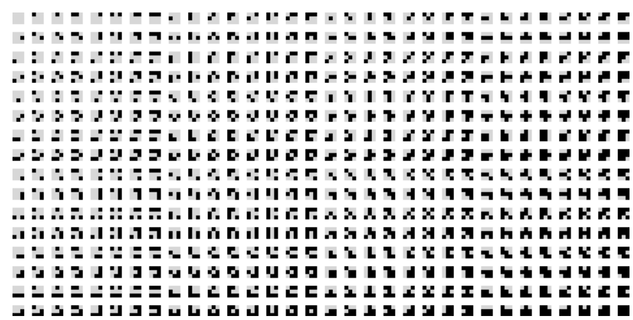

The odds of randomly picking an interesting image are 276/512, or almost 54%. In a tiny space, there are more interesting images because their iconic appearance easily maps to things humans find noteworthy. Also, one of my main criteria was symmetry, and small spaces have more of their images exhibiting symmetry. Some of the useful images were more meaningful than others, however, so the odds go down depending on one's threshold. Some images were difficult to include; I found myself asking "What is art?"

That kind of luck won't hold in larger spaces. As resolution increases, the chances for an image to be just a random arrangement of dots goes up. Also, for each good image free of defects, there is a staggering number of images that look the same except that they contain noise, filtering, deformations, etc. Although one should feel fortunate to pick any image containing anything remotely recognizable.

If we generate image values with a traditional random number generator, they will produce noise images if we use a standard encoding, because in normal images, neighbouring pixels tend to vary only a little. If we want to keep a standard encoding, then the value generator must incorporate a less random element.

The more non-randomness we use, the more our value generator becomes a procedural art generator. It's an easy slope to go from "pictures contain similar colors for neighbouring pixels" to "pictures sometimes have edges" and so on, adding knowledge about geometry and form.

The challenge is for a procedural system to not get trapped in any tiny subset of the space. A Mandelbrot set visualizer, for example, can generate a vast quantity of images but they are all focused on that set. Cellular automata systems create interesting local variations but larger areas tend to be random or abstract.

If one has an art generator, one is simply creating images directly. The goal is to pick a meaningful image value without setting any expectation on what the image might be. With an ideal encoding, random numbers would often pick meaningful images.

We can also choose to pick instead of generate. For example, a random normalized number can be multiplied by the image count. We could also iteratively construct a value using smaller random numbers, and then evaluate the value for noisiness using a compressor, and stop when the compression ratio indicated a suitable level of non-noisiness.

Image Surfing

A better interface to the quantum observer would be to give it a starting image (even a rough sketch would suffice), and then issue high-level verbal instructions on how to modify it. Thus, the system winds its way through image space as you tweak your original input.

It would feel as if the system is a powerful artist able to paint up your modifications immediately, but behind the scenes the quantum logic is finding matches. Like Al Pacino in the movie *SimOne*, you could start with a woman's face and then ask "longer hair, blonde, dirtier blonde, wider mouth, blue eyes, more tanned, cheek mole on the right" etc.

A great deal rests on the quality of the machine vision component. If it's even a little off, it will accept image space matches that can be frustratingly distant (if even by a little) from what is expected.

Exploring the Spaces

The Digital Image Space Explorer tool is a web page where you can explore some of the small black-and-white (1-bpp) image spaces.

It's available by visiting http://www.daylongraphics.com/support/image_space_explorer/index.htm.

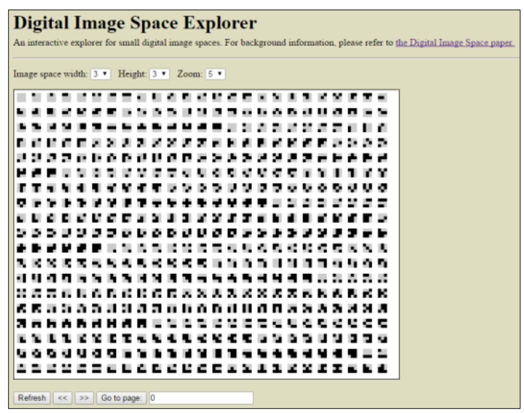

If it doesn't work, make sure that Javascript is enabled in your browser.

The Rise of Noise

As the image space dimensions increase, the number of images that are random noise increases exponentially. Finding a useful image simply by browsing becomes a treasure hunt suitable only for the extremely patient. The black-and-white 6 x 6 space alone has almost 69 billion images.

The Digital Image Space Explorer tool uses a traditional encoding, so black pixels appear starting on the top row of the image, working their way down as the image value increases. We can save significant time by starting our search to images that have the last pixel (the one in the bottom right corner) black and all the others white. This partitions the space into two equal regions with the larger-value region being more interesting, since as we go forward, all of its pixels can be black. Of course we skip over images that are smaller than the space (such as a box smaller than the maximum possible box) but that's of minor consequence. One can still find such images next to a dark corner pixel anyway.

An interesting question is: what is the proportion of noisy images to meaningful images? How does it change as the image space dimensions change? And what constitutes a meaningful image? Is beauty in the eye of the beholder, or is there some mathematical definition? In the 7 x 7 space, for example, there are enough pixels that images which would otherwise be meaningless blobs can be accepted as icons of animals. Game designers can identify classic Space Invaders aliens, and so on.

One thing is certain: the percentage of meaningful images in the HD color image space must be vanishingly small. This is both good and bad; it cuts down the search space, but it means we need to develop tools that can mathematically quantify what "meaningful" means, otherwise even a quantum observer will spend eons on the wrong regions.

Symmetry is one obvious metric, and easy to code for. However, many noisy images are also symmetric. As mentioned before, enlisting compression schemes to evaluate noise is useful. Another metric would be to skip images that contain disconnected regions, at least in small spaces. Some of these methods need to be optional, because they can also reject images that could be considered meaningful. Another noise detection routine is to perform an edge detection pass and then count how many edges result. In a meaningful image, edges are long and low in number.

The Iconic Virtues

The smallest spaces, by virtue of being the most iconic (minimal) may have something to say about the most general qualities of reality. If we imagine our universe as being an image space, then the pixel patterns constitute the available arrangements of states that the universe can inhabit. This might even be true insofar as the holographic principle is valid, where it's theorized that a volume of space is encoded on a two-dimensional boundary. Small image spaces were also employed (in a different way) by Eastern concepts such as the I Ching.

What does the smallest space have to say? It's a single pixel in size, and can be black or white. A universe that can be dark or light? Good or evil? Something or nothing? “To be or not to be?”

In the 2 x 1 space, those possibilities are joined by two images which are black-white and white-black, for a total of four possible states. Evil corrupting good and then good triumphing over evil? Life and death are normal? Or maybe the simplest way to state that the only constant is change? Overall, not a bad place to look for life's basic truths.

The 3 x 1 space is interesting, and its image count (eight) matches that of simple Eastern hexagrams. From start to end we perhaps have “life”, “first/best”, “average/neutral”, “mostly good”, “last/worst”, “balance”, “mostly bad”, “death”. The “balance” icon could also be interpreted as “struggle” or “conflict.” It's also the first icon that can represent the concept of separateness (the “average” icon can also do this, depending on how you interpret black vs. white).

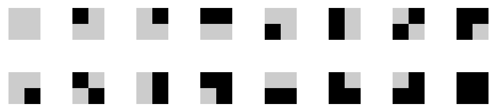

The 2 x 2 space is about the limit of minimal interpretation, although the I Ching is four times the size. The two diagonal icons could represent forces arrayed in opposition, while the four bar icons could mean forces working together. Their different orientations introduce the concept of orthoganality and rotation; indeed, the 2 x 2 space is the first one to be two dimensional. The space is very symmetrical if we allow for reflections about the diagonal axes, perhaps alluding to the high symmetry of physical forces when the universe was young.

The icons are mute, of course; whatever meaning they contain is only what we abscribe. Insofar as they are art, their value as art is up to us, and says more about us than them. It might be comforting to think that there is some singular universal interpretation, but it's also nice that we get to choose. By being so small, the spaces shoo away any thoughts of details, they force us to focus on epic generalities, the things that are most important.

It's akin to meditation: we stare into the void and it stares back at us, asking us who we are, asking us to say what is meaningful, and leaving that choice to us.

Consequences of Success

What happens if we're able to quickly search image space and find what we want? For starters, we effectively turn everyone into an artist. Images on demand; no art skills required. It's admittedly not a useful thought exercise today, but for our descendants, it gives a tantalizing glimpse into their possible future.

Almost as quickly as you can think of a picture and provide a reasonable search query, you could have it on your computer screen. The only thing faster would be a machine that can see the image in your mind and copy it (and given the difficulties of the quantum approach, a mind-reading machine might actually be more practical to construct). Having both machines would be ideal, as it would let one immediately find and sift through useful similar images just by thinking about them.

Communications would be transformed. Since a picture is a thousand words, people would increasingly opt to exchange ideas using imagery. Presentations and reports would become much more visual. The cost of storyboarding and other visual tasks involved with motion picture production would plummet. The age-old problem of other people "not seeing what you see" or "not sharing your vision" would fall away. People could actually understand one another at a level we can today only dream about.

The full usefulness of image space is, somewhat paradoxically, to not find what you imagine, but rather to find the meaningful images that lie nearby, that share various similarities. One's thoughts are merely a starting point. With an ability to alter parts of queries ("What if the car was red instead of blue, or it was a van?") we can use a quantum observer to quickly search interesting scenes from the multiverse. Engineers and architects, for example, would want to explore different design options for their creations.

Therein lies a danger. As a practical infinity, the space can lure the unwary into "search addiction." People might aggresively compete to find the most useful or interesting content. Designers may experience increasing depression, agonizing if they should have kept searching a little longer, if they had merely settled for images that were "good enough."

Finding perfection would ultimately depend on how well one can think and express oneself to the quantum observer. Even with such a wonderful tool, there might still always be — as with all tools — someone who could use it better.

And even that person, as gifted as he or she is, allied with every member of her massive population, searching furiously every moment, and even immortal enough to live for the billions of years the universe has left — would still be wondering at the immensity of what had to lay unseen. At the very end of time, countless highly-evolved beings — so evolved as to be almost gods — would quietly weep.

And if you imagine such a scene — no matter how you do so — that picture is also in the set. Perhaps our descendants will come across such an image, and wonder what it's about, not realizing that they are looking at themselves.