Bitesize Physics - All About ~~Vauxhalls~~ Voxels

By Matthew Benbow, Superintendent Radiographer CT & MRI, Royal Bournemouth Hospital
BAMRR News Editor

Me in my 1978 Chevette, Truro, 1986

Around 20 (or so) years ago, my sister bought herself a green Viva, my brother in law had a red Cavalier, my wife-to-be drove a blue Nova and I myself became the proud owner of a beige Chevette. For those of you too young to know what I am blathering on about, these were all Vauxhalls. Whilst there are many fine tales to tell about these beasts, this article is about an altogether different homophone that I certainly hadn't heard of back then – voxels.

Patients of course come in three dimensions. They have length, width and thickness. As medical imagers we need to produce images that demonstrate the internal structures and pathologies of these three dimensional beings, but so far we have not yet cracked the ability to use 3D imagery such as holograms to do this. This leaves us with the slightly challenging situation of viewing three dimensional subjects on two dimensional computer screens. To address this, volumetric imaging modalities such and CT or MRI produce a great number of two dimensional cross sectional images, or slices, which together tell the story of the whole volume being examined.

2D images are composed of rows and columns of picture elements, or pixels. The more of these there are, the finer the detail and the higher the image resolution. High resolution ought to be preferable, but it I always at the cost of image noise. With photography, if you set your digital camera to a higher megapixel resolution setting, the amount of light reaching each pixel on the light detector (CCD) is reduced, resulting in increased image noise. It is the same with MRI – higher resolution leads to increased image noise. One photographic solution is to set a longer exposure time to allow in more light, but this raises the risk of camera shake and movement unsharpness. With MRI we can also increase the scan time to improve our signal, but similarly this is at the risk of patient movement, not to mention shortening our tea break.

But photography and MRI have a fundamental difference Each MRI image is of course two-dimensional and so made up of pixels, but unlike photography, each pixel represents a certain thickness of the patient. So how does this work?

In previous editions of Bitesize Physics I described the process of spatial encoding. It is too long to go into again here, but suffice it to say that it enables signals received by the coils to be spatially located accurately in the final image. Each signal originates from little cuboids of tissue from within the patient called voxels. Each voxel has a pre-decided size, determined in the scan sequence parameters, which are set to achieve the desired resolution required to demonstrate pathology. The scan operator is able to adjust the dimensions of these voxels by manipulating the scan parameters. The front face of the voxel can be modified by changing either the field of view or the number of phase / frequency encodings. The voxel thickness can be adjusted by varying the chosen slice thickness to be scanned . Many scanners inform the user of the actual size in millimetres of the prescribed voxel. Once the signal from a voxel is received by the coil, the slice thickness dimension is flattened such that the 3D voxel is portrayed as a 2D pixel in the final image on the display screen

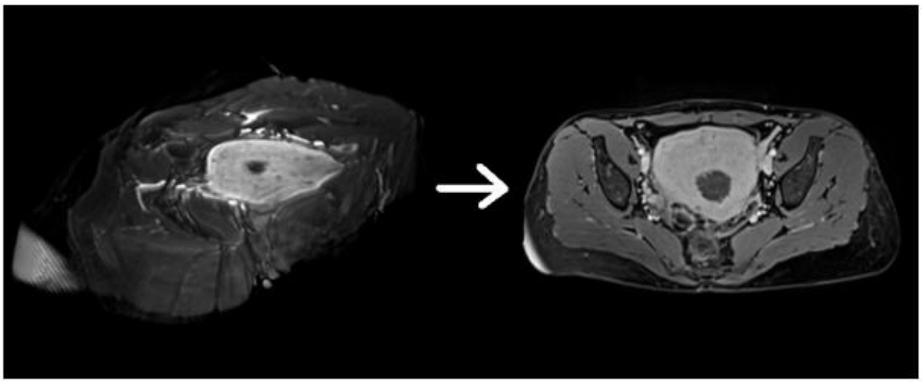

3D block of voxel data flattened to 2D image made up of corresponding pixels

So what are the consequences of making voxels bigger or smaller?

Changing the size of the face of a voxel in either direction will affect the resolution of the resultant image, but this should always be done with regard to the field of view. For example, if you double the field of view without also increasing the resolution (the number of voxels), then each voxel will also double in height and width, i.e. the same number of voxels must now fill an area four times larger.

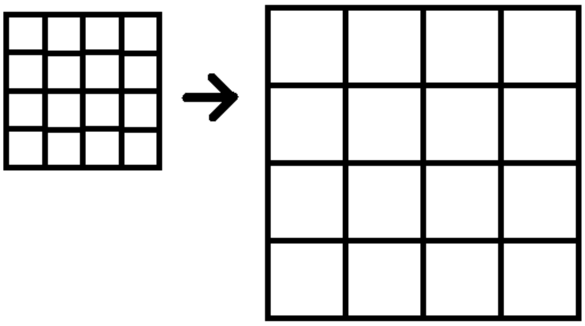

Increased FoV, decreased resolution

The corresponding image pixels will therefore also double in height and width and this may well be noticeable in the final images such that they are simply not detailed enough. The effect on pixel size is the same if you keep the field of view the same but decrease the resolution.

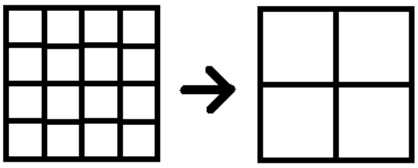

Same FoV but decreased resolution

Alternatively you may choose to reduce the field of view and keep the same number of voxels, or, for the same field of view increase the number of voxles. For both of these situations each voxel face will reduce in height and/or width and so the image resolution will increase. If pushed too far however, the noise will be intolerable and the resultant image unusable.

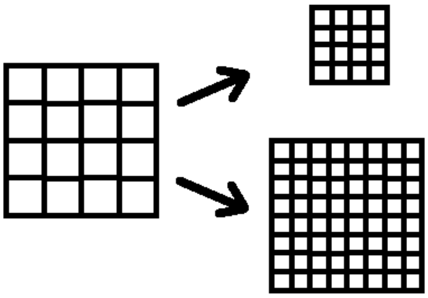

Higer resolution, but reduced SNR

So what about the slice thickness?

Reducing slice thickness will also result in a reduction in the volume of each voxel and so whilst the ability to resolve smaller structures may well improve, noise levels will again increase. Increasing slice thickness it will increase the signal to noise ratio, but may cause partial volume effects to hide subtle lesions within normal tissue.

So why does changing voxel size affect the image noise level?

Well, this is where I would like to introduce a term first told to me some 20 years ago by a fellow MRI radiographer Steve Ross – '*not enough meat in the box*' (thanks Steve, enjoy your retirement mate!). Think of signal as the meat. We require enough signal to make an image that is fit for purpose. The signal (or meat) needs to overpower noise, which is always present. The way to ensure this is to make sure the box (voxel) is big enough to hold a good amount of signal, or meat – i.e. enough meat in the box. Increasing the box size (or voxel dimensions) allows room for more meat (or signal) and so the detrimental effect of the ever-present noise is lessened. However, when we require high resolution imaging and cannot avoid the need for a small box, we will face the challenge of increased image noise, or more correctly, signal to noise ratio. One solution is to scan for longer, perhaps by increasing the signal averages. In this way you can cram more meat into the box and improve the signal to noise ratio, but this will of course take longer and there is always a limit as to how much meat you can actually fit in.

So are these signal to noise ratio changes quantifiable?

Yes, and many scanners will also calculate this for the user when adjustments are made to the in-plane resolution, filed of view, slice thickness, bandwidth, phase oversampling (no phase wrap) or signal averages. In terms of variations made to the voxel size, the maths is fairly straightforward. If you double either the height, width or thickness of a voxel, then its size will double and so will its signal to noise ratio.

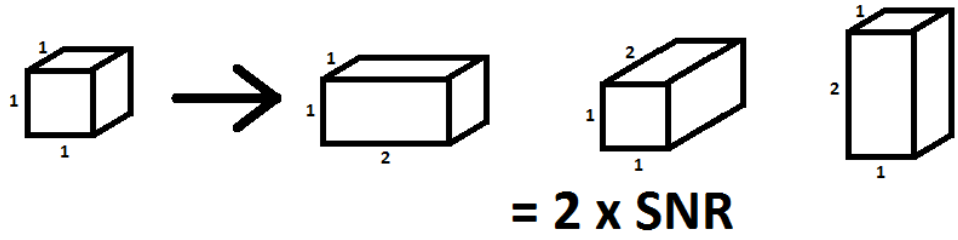

Alternatively you could double its volume with a combination of increasing more than one side by a smaller amount (figure f). By increasing all sides of a cube from 1mm to just √3mm (1.26mm), its volume is doubled and so therefore is the signal it produces.

How you proceed in setting up scan sequences therefore depends on what you need the images for. To depict small structures such as perhaps the ligaments of the wrist, then small image voxels would be the order of the day. To ensure good signal, dedicated multi-element, close-fitting receive coils are used. With 2D imaging, signal can also be maintained by compromising the slice thickness resolution slightly rather than the in plane resolution. Voxles acquired for 2D imaging therefore usually have dimensions similar to a ship container.

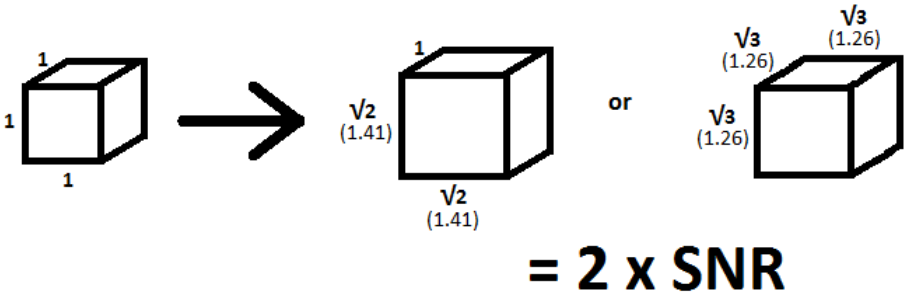

When 3D imaging is used however, then the viewer often needs to retrospectively need to reconstruct multiplane reformat images or perhaps three dimensional images such as maximum intensity pixel projections in a range of orientations. For this reason it is desirable for the resolution to be not only similar in all three planes, i.e. cubic, but also small, usually under 1mm or even approaching 0.5mm. This is generally referred to as isotropic imaging and can often lead to challenges in maintaining good signal to noise, especially when using smaller fields of view.

So how can we apply this knowledge practically? Consider this example.

I was called by a radiologist a few years back who was scanning at a nearby site. They were just setting up a prostate service and he was getting unacceptably noisy high-resolution T2 images. Their scanner, was not dissimilar to our own and the coils were being used correctly. The scan time was as I would have expected it to be, so it was clear that something in the parameters was not right. I spoke with the radiographer who told me that they had modified a pre-loaded manufacturer's pelvis T2 sequence as follows:

They had reduced the field of view from 320mm to 200mm and the slice thickness from 5mm to 3mm as they were looking to scan with high resolution. The resulting images were dreadful, so in an attempt to 'buy back' some signal they had increased the slice thickness and were now up to an undesirable 7mm! I thought about this and could see that the original protocol had voxels with dimensions 0.6 x 0.6 x 5mm, the changes that were made resulted in them now being 0.4 x 0.4 x 7mm. They were not scanning ship containers, they were scanning railway carriages!

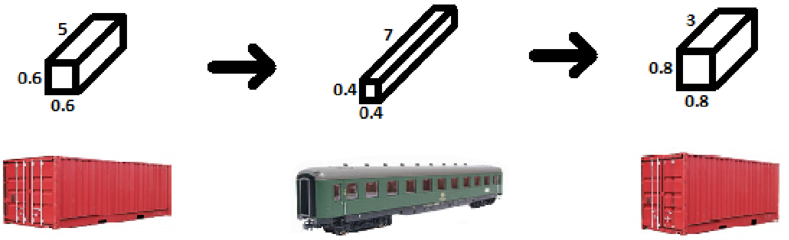

Good signal

55% signal
unhelpful voxel dimensions

Good signal again
ideal resolution

The error lies in that when they reduced the field of view, they did not also reduce the image resolution, but left it at 512 x 512 mm. The result of this was to produce voxels with a front face of 0.4 x x0.4 mm and hence 39% of its original volume - there was now not enough meat in the box! Their attempt to increase the box size had some logic as the newly increased slice thickness of 7mm resulted in a voxel of volume 55% of the original, but it was not really enough, plus the voxel now had very strange dimensions – far too long but unnecessarily thin.

I therefore proposed the following:

Maintain the FoV at 200mm (as required). Reduce the resolution to 256 x 256. At the smaller field of view this would still result in sub millimetre pixel size. Reduce slice thickness to 3mm (as required. This resulted in a voxel dimension of 0.8 x 0.8 x 3mm.

The smaller field of view meant that signal was now 72% of the original sequence, but the reduction to 256 phase encodings meant that the scan time had all but halved. Therefore the signal could be brought back to 100%+ quite easily by adding oversampling (no phase wrap) or by adding a signal average.

So in summary, it is important when adjusting the field of view or the image resolution in any of the 3 scanning dimensions to be mindful that there will always be an effect on the signal to noise ratio. Resolution of course needs to be appropriate to the pathologies being looked for but the 3-way relationship / battle between voxel size, signal and scan time will always be present.