Dynamic digital phantoms

In radiation research the term phantom is used to describe an inanimate object or system used to tune the performance of radiation imaging or radiotherapeutic devices. A wide range of shapes and size of phantoms are used for quality control and quality assurance in medical physics. See http://www.pnwx.com/CIRS/ for an sample of the vast selection currently available.

The construction parameters and the composition of phantoms are strictly specified. Hence comparing the results from repeated imaging of a phantom with different imaging parameters allows isolation of the effect of the varied parameter. With this knowledge one can assign settings for optimum imaging of that particular phantom. Ideally the phantom would be exactly the same for every scan to allow easy comparison. However in some cases natural degradation of the materials involved means that the results will change slowly over time and designers of phantoms often attempt to select materials to reduce this degradation. However the nature of the physical processes involved in the imaging modality limits the range of possible materials available for construction, potentially making it impossible to eliminate degradation. This is especially true in the case of multi-modality phantoms in which the material choice may be very severely limited. However, if the degradation is very slow, is measured before use or is factored into the image evaluation, then it can be managed.

Physical phantoms are often roughly grouped into two types, those shaped as a composition of a few geometric shapes (here in referred to as simple phantoms) and realistic anthropomorphic phantoms which reflect anatomical features. The former is both simpler to manufacture and allows easier interpretation of their results whilst the latter is more relevant and representative of the actual imaged systems.

As an alternative to taking measurements with a physical system, computational techniques such as Monte Carlo modelling can be used to perform *in silico* measurements. This power of computational models is that they can be left running in the background to test hundreds of slightly different sets of parameters. As the computation is digital, the inter simulation variation can be eliminated. Coupling the reproducibility with the increased number of experiments/simulations can be completed means that computational models can often pick up very small effects. These effects may have been hidden by experimental noise or within inter experiment variation in physical experiments. Of course any computationally uncovered effect will need to be tested experimentally. However testing the existence of a suspected effect is practically a lot simpler than deducing the existence of the effect in the first place.

Computational simulations require phantoms for the same reasons as physical measurements. There are however important advantages of digital phantoms, including their absolute lack of degradation over time, the ease by which they can be entered into different computational models, the lack of limitations on materials used in the ‘construction’ and the ease by which they are transferred and copied. However their form still needs to be described. Often a physical phantom will also have a matched digital version to allow prediction of experimental results using the physical phantom.

A common example is the verification of treatment planning in radiotherapy. The dosimetric result of a complex radiotherapy plan is recalculated on a digital phantom corresponding to a simple physical phantom. The results of the recalculation can then be cross-checked by measuring the dose recorded in the physical phantom when the complex plan is delivered to it. This provides good confidence in the treatment plan before it is actually applied.

Very basic phantoms can normally be well described by geometric shapes and thus can be described by using very few parameters. As an example consider a spherical phantom which may be absolutely specified by 5 parameters (3 positional, one for size and one for composition.) Computationally specifying geometry like this is an excellent framework as the computation program as they only need store a small number of parameters to describe completely the whole simulation environment. This allows flexibility when writing the equations governing interactions between and within the regions (specified by the geometric shapes) allowing complex behaviour to be modelled accurately. The geometric modelling approach can be referred to as a 'continuum model' as within each region the volume parameters are assumed to be continuous. An important limitation of continuum models is that they are only approximations of most situations and lack the flexibility to describe very small scale fluctuations. Historically these models were quite crude but they have become more and more complex and realistic as can be seen by a chronological viewing of previous continuum models. (See http://www.rpi.edu/dept/virtualphantoms/phantoms.htm for a comprehensive list of computational phantoms used in radiation research)

For describing very complex geometry we must turn away from a description based on simple shapes to full voxelized models. Voxelized models divide the environment into small cuboids (voxels) and assign data to every voxel. The only limit on the accuracy of their sample description is the size of the voxel. The easiest way to acquire a voxel model is merely to use some existing volumetric image data of sufficient resolution. Selecting the resolution is not trivial: smaller dimensions and thus voxels describe the sample more accurately but this process greatly increases the number of parameters in the simulation. Even a relatively coarse voxelized model (100 x 100 x 100 voxels) is associated with up to 10^6 parameters that need to be entered into the simulation. In finer models, simulations based on solving equations linking the voxels will be inevitably slow. Simple equations and care will be required to avoid numerical noise adversely affecting the results. In practice the optimum choice of voxel size is dependent on the nature of the site being simulated and the nature of the simulation itself.

The problem with using raw image data is that the ground truth regarding the composition of each voxel in the model is essentially unknown. What exactly is being represented by a voxel is only specified as far as the imaging modality allows. So for example, say we have a voxelized model provided by high resolution MRI which we want to import into a Monte Carlo X-ray CT simulation. Conversion of the MRI voxel greyscale values into values relevant for CT (electron density) is required. A straight one-to-one mapping will not give sensible results due to degeneracy and noise in the source image. More complex mappings which take into account factors like the voxel location in the image and its neighbour's values are better but still are very problematic and will likely require human 'tweaking' afterwards.

These issues are well known from research into computational outlining of tumours. This problem is still awaiting a solution which is why the majority of outlining is still tediously completed, or at least tediously adjusted after a computational pass, by humans.

The 'natural' way to divide up medical images by hand is into organs as this allows prior anatomical knowledge to be used to help guide the organ outlining. Hence, rather than trying to outline electron density values, it is instead easier merely to outline organs relevant to the simulation. After outlining the voxels relating to a specific organ the organ volume can easily be set with whatever values are relevant for the imaging modality being considered.

The choice of whether to use an outlined voxel model or a continuum one for a digital phantom is highly dependent on the geometry of the area being simulated and the nature of the computational simulation. Certain parts of the body with fractal like geometry such as the lungs are especially tricky as fractals have both numerous interfaces and very small features making both continuum and voxelized descriptions quite inaccurate.

As both model types have matured some of the initial limitations are starting to be overcome. For example modern continuum models are now so complex that they are often an extremely good approximation at very least as the organ scale. Naturally this increase in the accuracy of the geometric description could have slightly decreased the complexity of the equations that can be utilized but increases in the availability of computational power have more than offset this limitation. Similarly, well used voxelized phantoms now are composed of voxels that are extremely well characterized and accurately segmented into organs. Due to the growing complexity of geometry in continuum models and the increasingly improved descriptions in voxel models both approaches are converging and complementary.

Actually there is no need to be limited to one approach alone as conversion is possible. For example a phantom called Moby (created by Dr. Paul Segars of Duke University see http://deckard.mc.duke.edu/xcatmobyrobyphantom.html) has aspects of both continuum and voxelized approaches. The user specifies approximately 100 parameters which are used to generate a dynamic continuum model. This model is then volumetrically resampled onto a grid at set time points in the respiratory/cardiac cycle to produce a series of voxelized models. This allows the benefits and ease of working with voxelized models while the base continuum model provides 'free' organ segmentation, well defined motion of the organs and a way to easily generate similar voxel models at different voxel sizes if required.

By using such a digital phantom we have the rare ability to simulate the entire process of imaging and treating a moving object with reference to a known ground truth. This model is thus idea for research into motion issues in radiotherapy and molecular imaging.

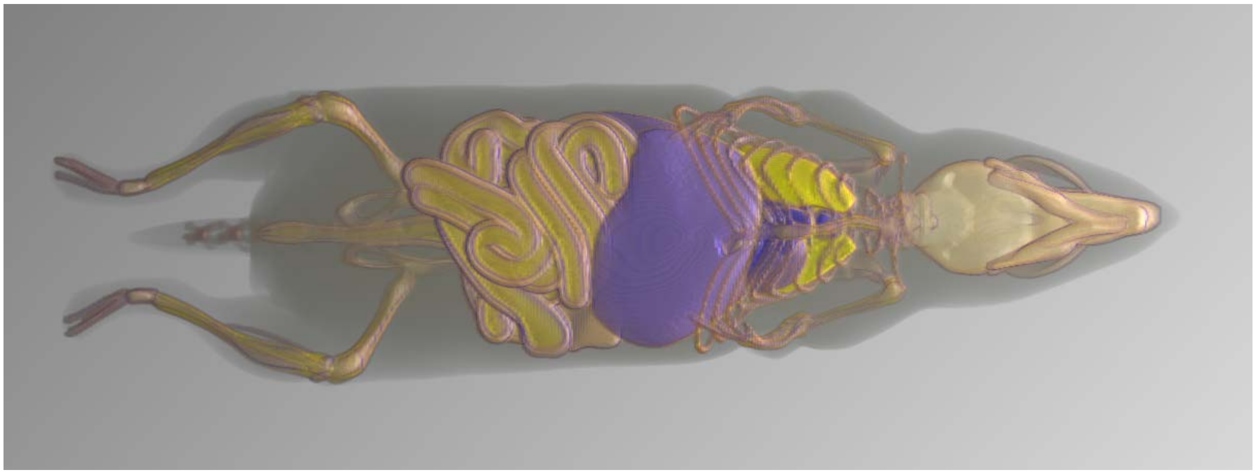

An example of one of the voxelized arrays produced by the Moby package and rendered by the Voreen rendering engine.

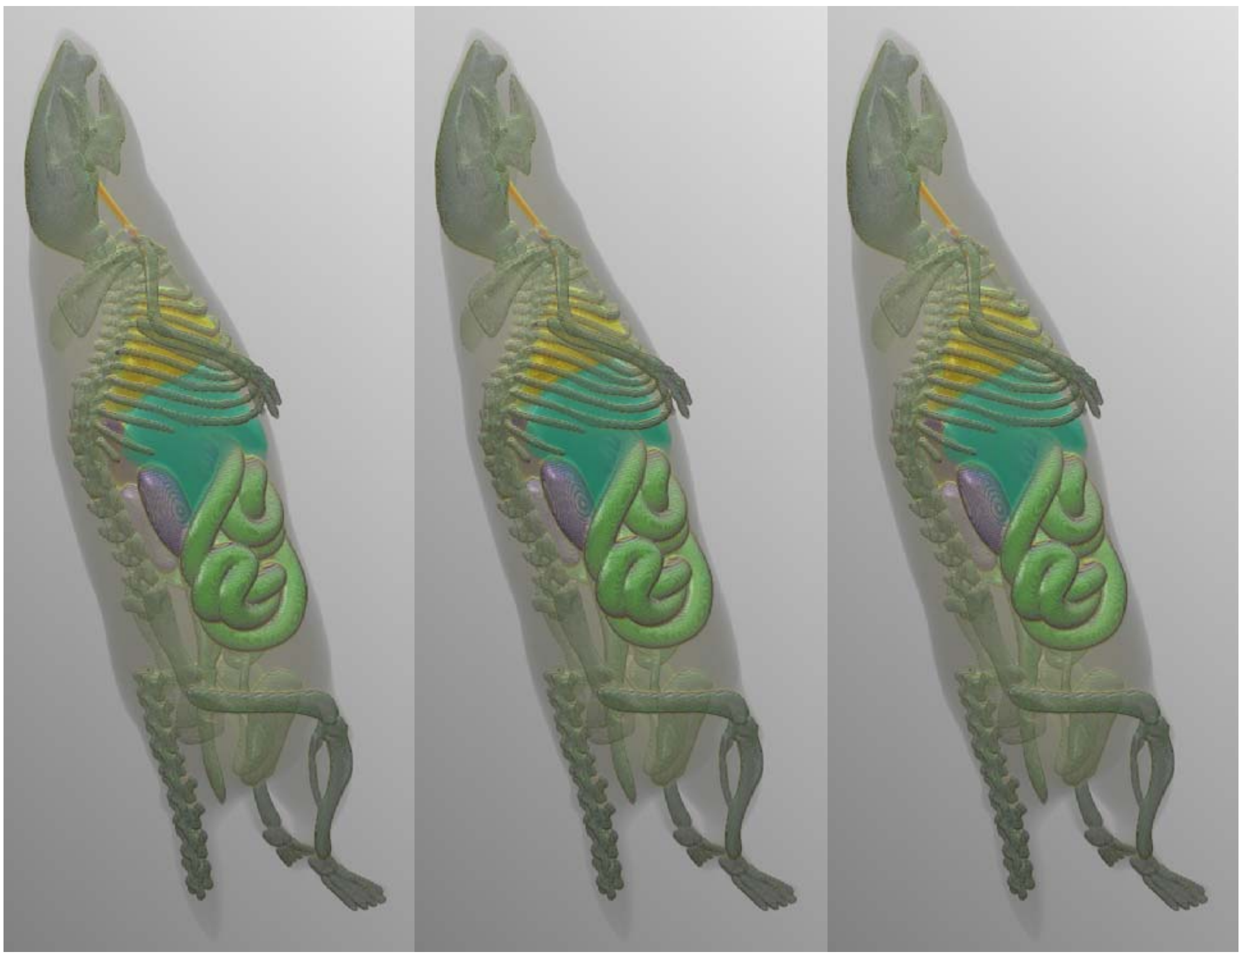

Images of a series of renders of the voxelized mice. Each of these renders comes from a well-defined time point in the respiratory cycle.

This note was prepared by A Kavanagh and B Vojnovic in October 2011

We acknowledge the financial support of Cancer Research UK, the MRC and EPSRC.