

I'd like to talk about attenuation correction in PET/CT imaging, and in the second part, we'll be discussing image fusion. I'd first like to give a quick overview of what a PET/CT scan involves. You have a single-patient bed, and the PET/CT scanner, which has the CT component in the front and the PET component in the back, which typically in commercial systems may effectively be thought of as two separate systems back-to-back, but the key point being that they're aligned in plane. So you begin with a CT acquisition, first acquiring a topogram, and the topogram is used to define the range of the actual scan. Then you have the spiral CT scan. And following that, you have the PET acquisition, so the patient moves further into the scanner because of the fairly significant separation between the CT and PET components. And at this point when the PET data is being reconstructed, you're actually using the CT data. This is when you're performing CT-based attenuation correction. You're going to use those images as part of the

PET data processing, actually for both attenuation correction and scatter correction, and this gives you the PET images, which you can then fuse, in a workstation display, functional and anatomical information together in a single image. So in the first part of this talk, I'd like to go over the principles of attenuation correction in PET in general and talk a bit about the features and the images you get when you choose to do attenuation correction or not, and then talk about the principles of specifically CT-based attenuation correction, and this will involve discussing the requirement of energy scaling of the CT images, and that necessitates a discussion of what are the scaling characteristics of human tissues and how do we do that. I'll describe a slightly more sophisticated scaling method beyond some of the ones that are currently employed, that is, kVp dependent scaling, which depends specifically on the voltage setting of the x-ray tube used in the CT scan. I'll then discuss the benefits of shorter acquisition times, respiration artifacts

and protocols and how they affect the situation, and finally, the use of CT contrast agents. So to quickly discuss the principles of PET itself, you have a glucose analog and you have a positron-emitting isotope in that, so you are basically measuring coincident photons, which you then reconstruct to get the PET image. So what you're building up an image of in reality is the positron annihilation event. So it's not quite the production of the positron, and here's where some of the inherent resolution issues come into in PET, in that you're not actually measuring the location of the production of the positron in the radioactive decay, but actually the subsequent annihilation and production of your coincident photons some distance away. Fortunately, for F-18, because it's a low energy positron emitter, this positron range effect is not too great, but it does contribute to some blurring of the signal. Now if that was all there was to it, you would simply acquire these coincident events and reconstruct an image and you would have exactly that,

an image that is a map of the annihilation events of those positrons, and it would, in that sense, be a real quantitative measure of your tracer distribution. However, as is shown here, what can happen is you have an annihilation event, which produces two back-to-back photons, which you have detectors there to measure. However, one of those photons, or both of those photons, may subsequently interact with the body, so what you find is that scatter and absorption of those emitted photons by the body cause the true emission signal to be attenuated. So anything you measure, even after you perform all your other corrections, you're still getting ... measuring a signal that is lower than what you want. So you need to actually perform a correction to account for that. And that's why to get qualitative results in PET, you need to correct what you measure with attenuation correction factors. So before PET/CT, there was not CT-based attenuation correction, so what you would have in a traditional PET scanner is you might have

some kind of rod sources or point sources, and you're measuring those factors by doing a transmission measurement to essentially explicitly determine what they are. You basically want to measure the attenuation through all lines of response you are interested in measuring emission data from, so you might have, as is shown here, you have three rod sources that are rotating in your scanner and you're able, therefore, to measure the coincident lines for all your lines of response, and you would first do a blank scan, so you essentially normalize to what the flux is from these rod sources, and then you would do a transmission scan as part of the PET scan. So for every patient scan, you're doing a transmission scan first, and you can calculate the attenuation of the signal because you know it from the blank scan, and you know how much it's attenuated from the subsequent transmission scan with the patient in the gantry. So

you're basically in traditional PET scanners using sources to provide an essentially direct measurement of these attenuation correction factors, and I say direct with an asterisk because these are generally noisy measurements, low photon flux, and you may actually not use them directly, but instead reconstruct them into a transmission attenuation image to control noise and maybe allow further image processing or segmentation, and then you can reproject them to get a slightly better result for your attenuation correction factors. So you may ask, "How large are these things we're actually dealing with?" Again, to reiterate, these are the factors by which your true source emission signal is being suppressed by interaction with the body, and it's a multiplicative factor. The true event rate that would be measured needs to be multiplied by the ACF. And how large are these? Well, the physics is fairly straightforward, so these interactions are well known and well described for all different atomic materials, for all different

well-known tissue compositions and reference tissue compositions that have been well studied. And as we know, the body is very similar to water. Soft tissue is similar in its attenuation properties to water. So we can look at how big are the factors for path lengths through water. And that's what this curve shows. So it's showing you how much a signal will be attenuated if you go through 20, 30, 40 cm of water. And some typical values for tissue are, 20 cm phantom, something of standard, will give you an ACF of 7. So your signal is being suppressed going through the center of that 20 cm cylinder by a factor of 7. And for 40 cm, we're up at 50, and for 60 cm, which may correspond to the largest patient population right through the center; we're finding that our signal that we're trying to measure is being cut down by factors of hundreds at that point. So you see, the reason why PET images for larger patients are degraded so much is because you really are not measuring the same strength of signal at all, and it's an exponential

rise, so the situation only gets much worse. So in the PET image reconstruction, you can choose to do the attenuation correction or not, or as may be the case, you might to choose to have it both ways. Certainly, images without attenuation correction applied are ... have been used extensively. And I show here a coronal view of the same data reconstructed with and without attenuation correction. So with attenuation correction, we are dealing with a quantitative or what we hope is a quantitative image, a quantitative map of the glucose tracer distribution in the body. Without attenuation correction, we see an image that, even though it's clearly the same patient, is really quite different in some respects. So what are the features of the non-attenuation corrected image that we see? Well, first of course, it's not quantitative. We know that. One feature that stands out is that the lungs appear hot. If we use our image on the right-hand side as our reference, as in a typical tracer distribution with some evidence of disease, we see very low activity in the

lung, and that's what is expected, whereas without attenuation correction applied, the lungs appear hot.

We also see suppression of inner activity relative to activity from the outer surface. For example, if you look on the right, you see the liver is really quite a large uniform area, which is what we expect, whereas the arrow on the image on the left shows that you actually get a decreasing distribution as you're going more inwards into the center of the body. This is also the reason why, if you look at the outline of the body, the skin appears relatively hot, and that's also a non-quantitative artifact of not performing attenuation correction. But the reason these are still considered useful is that focal uptake in disease is still very much apparent. You still see a lot of structure. You see a lot of the areas of increased focal uptake appearing in both images. And why would we expect the image without attenuation correction to look as it does? Well, you simply think, "What is attenuation doing? It's suppressing my signal, giving me a smaller signal for any regions where I'm going to get – have to pass through a lot of tissue for my photons

to get out." So anything in the center of the body you expect to be suppressed, relative to anything near the surface, because those photons near the surface have less tissue to pass through to get out, so the signal will be stronger, and if you don't take measures such as attenuation correction, that will be revealed in the image as an area of increased uptake. As well as using attenuation correction to get quantitative images, another benefit is that you can actually incorporate it further into the statistical nature of your reconstruction algorithm, use it for weighting. So here, I show the same image, a coronal view with three different reconstructions. The first is an essentially analytic conversion filtered back-projection, which is mathematically exact, but takes no account, really whatsoever of the noise structure that's present in the data. And PET is a very statistically noisy imaging modality. So, as might be expected, the image is really not all that good. You go to an iterative method, a statistical method that takes some account of this

property of the data, and you do much better. So this image might be an OSEM-type statistical algorithm, but takes some account of some Poisson nature of the underlying photon statistics. So this is a likelihood image. This is the image that fits best with the measured statistically noisy data. But you see if you look in the center of the image, you can still see some evidence of some streaking, some correlation of noise in the liver, which you don't expect to be normal. That's artifactual, and this again, because these projections are very much attenuated, so we're getting very poor statistics through the body there. And finally, on the right, is again an iterative statistical approach, but this time with some attenuation weighting, so you pay less weight to those lines of response passing through the body that you know are the noisiest. And you essentially suppress the noise contribution from them to the image. This is implemented on some systems as attenuation-weighted, OSEM statistical algorithm. And you see really quite a drastic improvement by

going to that. So you're using your attenuation correction factors, but to build the signal up to make it quantitative, and to statistically weight each line of response according to how noisy you expect it to be and incorporating that into your reconstruction algorithm. So, so far everything I've said is really been about attenuation correction in general, not specific to PET/CT. But now I'm going to describe that. So what you have in PET/CT imaging is basically another way to do it, which has many benefits. To summarize transmission-based attenuation correction in PET, you're basically measuring essentially directly at 511 keV or near 511 keV these line intervals that you need. And I say near, because some sources have a slightly higher energy, but once you're up at these high energies, the energy scaling properties are fairly simple, and you can easily scale down in energy. You may do some reconstruction to get a very noisy attenuation map at 511 keV, and you may segment that and replace it with some new

values to reduce noise and re-project it, and that's how you're getting your ACF's. So, it's based on a

measurement and then a fairly noisy 511 keV map that comes directly from that measurement, and that's re-projected. In PET/CT, what you have is you have the CT images already, for free. It's part of your scan protocol, and they may be thought of as corresponding to an attenuation map also, but an attenuation map at a much lower energy, because your x-rays are at a much lower energy. And even then, it's not that simple. That's why I'm saying about 70 keV, because you're not even using a monochromatic source with a single energy. You're using an x-ray tube, so you have a spectrum of energies. So the image does not precisely correspond to a single energy. However, statistically, it's a very noiseless image. You're dealing with CT, not PET. You're not counting single photons. You're dealing with a huge number of photons, very high flux, so very low noise. So if you can transform this attenuation map that the CT

images correspond to, to an attenuation map of the right energy, it's certainly going to be very good from the point of view of not having a lot of noise in it. The question is, are we doing the scaling right? Is it free of bias, as well as free of noise? And if you can get that attenuation map this way, then you can downsample it and smooth it, because it's a much higher resolution image, being CT compared to PET, you need to do those steps, and re-project it and again, you have your 511 keV ACF's that you need. So just to summarize some of the key points so far, for CT-based attenuation correction, for transmission attenuation correction, you're doing it at or near 511 keV. CT x-rays are much lower in energy. They're going to be lower than 140 keV. That's typically the highest tube voltage you use, so that's a kind of a limiting factor. Your x-rays are at or lower than your tube voltage, so we need to scale them to 511 keV. The transmission ACF's are much noisier than the ones we will get using CT-based attenuation

correction. And a real practical advantage from being able to do this, is that it's eliminating the need for a transmission scan, if you can confidently routinely do CT-based attenuation correction, you don't need that as part of your patient scan protocol, and if you don't need a transmission scan, you don't need a blank scan, and you don't even need the sources. So you're seeing a real benefit in terms of your system design here. Of course, you still need to do some calibration. On the PET side you still need to calibrate the PET detectors. You're not eliminating that, which you would do with a 20 cm phantom acquisition, but you are eliminating the need for the blank scan with the sources, and the need to have the sources themselves if you're doing solely CT-based attenuation correction. So we should ask ourselves in a little more detail, "What really are the CT images?" Well, CT scanners are measuring the local photon linear attenuation μ , exactly what we want in PET, but they're measuring it at the effective energy of the x-ray

beam at around 70 keV assuming a fairly standard CT protocol with 120 V keV tube voltage. So you've got some spread around 70. CT images are normally expressed in Hounsfield units, which is really just a recasting of that information. They're still basically measuring the photon linear attenuation. The Hounsfield unit is nice, because it references everything to water, so zero Hounsfield units is water, -1000 is air, and you essentially remove the need to be referencing the actual linear attenuation value μ and what energy it was at. So the other thing about CT, as I said, the reason they're noiseless is because the tolerance on these CT scanners is so high. You're talking about calibrating one of these with a 20 cm water phantom so that you have a CT number uniformity of 0, plus or minus 4 Hounsfield units over the whole range of that 20 cm water phantom, and if you think about a 10-Hounsfield unit difference, even though in CT that might be quite a lot. You're distinguishing white and gray matter in the brain by

looking at such small differences in the Hounsfield unit. But in absolute terms of linear attenuation, it's

really a very small difference. 10-Hounsfield unit corresponds around the water value to a 1% difference in linear attenuation. So on the PET side, these are small differences indeed. And because these x-rays used in CT have lower energy than PET, we will, of course, need to scale these images up to perform attenuation correction. So what are we dealing with in human tissue? This just kind of summarizes the various Hounsfield units you're going to measure. We may think of it roughly in terms of three groups. We have adipose tissue, fat, which is fairly low in terms of its Hounsfield unit, and therefore, its photon linear attenuation, we find this in various amounts throughout the body. It's both in between the organs, intraabdominal fat, and deposited around the body. Also, there is yellow marrow in the long bones, at the end of the long bones in adults, and the values ranging from -50 to -100 Hounsfield units, with 0 being water. Then we have the soft tissue group, and so there's a lot of different components here, but they all

have attenuation values, Hounsfield units very close to and above water, 0 to 80 Hounsfield units would typically encompass this whole group, so we have blood 40 Hounsfield units, the liver 40 to 50, muscle, gray matter, etc. So we have a lot of tissue in the body that's attenuating, which is similar to water, slightly higher in terms of its Hounsfield unit, its attenuating value. And then higher than that is the bone tissue group, and there's actually a lot of range in terms of the Hounsfield units for the bone tissue. You have in the vertebrae you have trabecular bone at around 100 to 300 Hounsfield units, also called spongy bone or cancellous bone, and on the high side you have cortical bone on the surface of your long bones, and that's the densest with very high values of 1200 Hounsfield units or more. It's important to remember, that even those values are very high, the actual amount of cortical bone you have is not very much in terms of the amount of pixels that actually have this value of 1200 Hounsfield units. It's fairly

small, because you only have a thin cortical shell on your bones, and the rest of it is somewhat in between these two groups of values. You have lots of different bones, which are in between the trabecular bone numbers and the densest cortical bone numbers. So what we need to do, is we need to be able to relate all these measured Hounsfield units we get for all these tissues to what the attenuation values for those tissues would be at the PET energy of 511 keV, and it would be very unfortunate if things were so complex in terms of how these scaled that it was non-unique, and you couldn't say that this Hounsfield unit maps into this value. It might be mapping into different values, depending whether on whether it was one kind of tissue or not. But that's not the case. It tends to be fairly simple to have a straightforward mapping of Hounsfield units into 511 keV values, and that's fortunate, because it allows us to routinely implement CT-based attenuation correction. We can look, using reference tissue values, at what those

scale factors are, so this graph basically shows if I take a tissue attenuation value at a CT-type energy of around 70 keV, what's the ratio between that and the value at 511 keV. So these are the scaling factors we need to use. And what you see is that if you look around 0.5 there on the graph, you find that lots of tissues have the exact same scale factor. We have the soft tissue group is all together. It's all very similar, and it's very close to the factor you get for water. You basically take the attenuation you measure in CT and cut it by half, and that's basically what you need for PET. And you also see that fat tissue is a little bit higher, but fairly close, close enough for us to deal with that. And the lung tissue is also got a value around the factor for water. It's much less dense, but it scales in the same way. Even though it's much less dense, that factor pulls out. What is different is the bone scaling. The values for bone are lower, and we see the denser the bone is, the lower the value. We have cortical bone down there with a fairly low scale

factor, and the spongiosa bone with a somewhat higher factor. But already, just looking at a graph like

this suggests how we can do this. If we can take everything except bone, we can put one scale factor in for that and scale the whole thing. And then for bone, we're going to either take some mean value around 0.4 or thereabouts that will basically work fairly well, or we could have some slightly more sophisticated model of the changing value. But basically, you're just treating everything but bone one way, and then a slightly different treatment for the bone. So here I kind of generically describe two possible ways of doing this. The first is to have a simple threshold model. We can say we want to treat bone and soft tissue somewhat differently, based on that previous curve. The scale factors for soft tissue and bone are well known. We saw that a reasonable value is 0.5 for soft tissue and maybe 0.4 or thereabouts for bone. So if we could simply separate them with a simple threshold, for example 300 Hounsfield units, we see there on the top left how we separated a CT image into its bone component and its non-bone component, with

this threshold value. We simply scale each of those with a different factor, recombine them, and there we have our 511 keV image, and we can re-project it and get our attenuation correction factors. Another slightly different approach is to take a mixing model, so this accounts a little more for the different types of bone we have and gives it a variable scale factor, depending on how dense our bone is. We're basically saying, bone is a mixture of water and dense bone, which is not unreasonable. In fact, bone is mostly a mixture of the same components you find in dense bone and various amounts of marrow and fat tissue, and so that's perfectly reasonable to assume that kind of mixing model. And so you scale things as basically soft tissue or water/air mix below, say, 0 Hounsfield unit, and of water/bone mix above. Again, you have what is key, a one-to-one mapping between a Hounsfield unit into a 511 keV value that's unambiguous. As described, these methods aren't perfect. They don't account for different kVp. Your

values, your scaling is going to be different if you have different tube voltages in your CT scan because your average energies are different, and they also make some assumptions about the location of thresholds and where we should have breaking points, for example, is 300 really a reasonable breaking point? Is it an optimum breaking point to distinguish soft tissue and bone, and in the case of the mixing model, should we really break that curve and change the slope at 0 Hounsfield units? Since most soft tissue is slightly above 0 Hounsfield units, but it's similar to water and is not really a water/bone mix, that might also be questionable. So here I just want to describe a slightly more sophisticated procedure where you go to a kVp-dependent energy scaling. We know that CT scans at different kVp settings correspond to different effective energies, so we're generalizing to scaling that takes account of that. Typical values you have in CT scanners are anywhere from 80 up to 140 kVp. Typically 120 would be commonly used for most

scans. Pediatric cases you might scan lower at 80. If you have a very, very large patient, you might want to use harder photons to get more signal through, and you might go up to the highest energy available with a kVp of 140. So one way to do this is to take an electron density phantom, which has a bunch of reference tissue inserts, for which you actually know the electron density, you know the 511 keV linear attenuation values. And you simply scan that phantom at a low kVp. And you basically get the curve like I show here where you see that as you get to higher Hounsfield units, which corresponds to denser and denser bone, you actually see a separation of this curve, this transformation curve for different tube voltages. So with this curve, which is the actual measured data, you can imagine simplifying it a bit. You could basically do a two-slope model that fits each of these different kVp curves reasonably well, and if you did that, you would find that your breaking point, the point at which you change from one slope to

another really works out to be around 60 Hounsfield units, and that's actually kind of a nice result, in that

it corresponds to what we would interpret as a real breaking point from what's observed in human tissues. If you remember, the soft tissue group kind of went up to 60 or 80 Hounsfield units, and thereafter, you're into the low density bone, so by doing a direct measurement calibration such as this, you really find you get thresholds and breaking points that are kind of more in correspondence with the reality of what you expect than maybe some other methods assume. So now that we have a way of doing this and we have CT-based attenuation correction, and it's viable, what are some of the other benefits of that? Well, one is in terms of the reduction of whole-body scan times. So you don't have the need any more for a transmission scan, and that is going to mean faster total scan times. And here I just show a progression in the reduction of scan times if we go back a few years and look at a BGO partial ring scanner, such as the ART. Now we'd have almost the best part of an hour as the scan time, and you see that half your time is

really spent doing the transmission measurement. There's been some studies in terms of what the optimum sharing of the total time is between the transmission and emission, but certainly you might think of it as roughly taking up almost half your scan time. So then you see with going from BGO to LSO, for example, faster scintillators do give a substantial reduction in imaging time, but a real benefit is to simply remove the need for that whole transmission part altogether. And so you see with the LSO PET/CT, for example, you're getting very, very fast scans, certainly a great improvement, and CT-based attenuation correction is a large part of that. So I'd like to now talk a bit about the CT respiration protocols you would use in PET/CT. During the PET acquisition, the patient is breathing shallowly, averaged over many cycles, as is absolutely necessary because of the long duration of the scan. With fast CT scans, if you were just doing CT and you were scanning through the thorax, given that with the latest 16-slice scanners, you

can cover that region in a matter of seconds. If it's possible for a patient, you would typically do that with full inspiration breath hold, because you're freezing the motion. You're eliminating artifact from the CT, and the reason you would do full inspiration, is simply it's easier to hold your breath – hold a deep breath – than to hold it at any other part of the breathing cycle. The problem with this protocol that may be optimum for CT, is that it's certainly not optimum for PET/CT, the reason being that when you breathe in deeply, you have a maximum displacement from your average position for shallow breathing. So you're basically maximizing your mismatch with the average position of your breathing in PET. And we see at the bottom here an example of an artifact that's fairly severe that can result from that is that you displace the chest wall so much by having a full inspiration CT from its mean position in PET, that it simply disappears in the attenuation correction. You've displaced it into the lung field, and you're not building

up those ACF's like you expect to. So other possibilities with CT are to basically try and get closer to the average part of the breathing cycle in PET and that would be to either also do the CT with shallow breathing or to attempt if possible a partial inspiration breath hold held much closer to the mean point of the shallow breathing cycle. And that, these are very crude, but just to give an idea, if you're breathing shallowly, you may expect excursions of certain organs in the body or the diaphragm of 2 cm perhaps – it's very crude – versus up to 10 perhaps for deep breathing. So there's really a drastic difference between the deep and regular breathing. So we do get respiration artifacts with these breathing protocols, which can propagate into the PET. This was really much worse with single-slice CT systems and PET/CT, because there you can get much more severe artifacts on the CT if you're not able to perform breath hold, and in this CT image, we see a pretty crude example. You get distortions that are geometric in the sense

that you're not simply deforming the shape of things. You're actually mis-ordering your slices geometrically, and it's connected with the fact that your organs are moving in an axial direction because of your breathing, and you're moving through the CT, and if your organs are moving faster at any point in the cycle than the CT, they can kind of overtake and fall behind and lead to very complex interference patterns, and we see that here in the CT image with this detached liver dome hanging there in the lung, and the artifact certainly does propagate into the PET images, as we see on the right. With multislice CT, things are a bit better. This gives two images, both using tidal shallow breathing during the CT acquisition, and we see on the left with a 16-slice CT, you really don't see any severe distortion. You really couldn't tell with just a cursory examination whether this CT scan was acquired with breath hold or

with shallow breathing, whereas on the single-slice CT on the right, you see that very severe characteristic geometric mis-ordering of slices that leaves you with a kind of piece of the hanging liver dome. So what this means is that with shallow breathing, you still don't expect things to be perfect in the sense that there's no guarantee that that 16-slice CT scan corresponds to the mean position of the PET breathing cycle. It may still be slightly distorted. It may still lead to some artifact at the boundary of the liver and the lung, but it's a lot better than having the severe artifacts that result from slice mis-ordering that you get with single-slice CT. So we're seeing a real benefit here with having a somewhat faster 6 or 16-slice CT in PET/CT than with a single-slice that may take much longer for a whole-body scan, 90 seconds, and may not be practical for any kind of breath hold, full or partial inspiration as the scan, CT scan goes through the thorax. So, it's less severe artifacts with a multislice-CT. What are the other types of motion

we may encounter? We have cardiac motion with a very short time scale, respiration, as I discussed, motion contractions of the muscles in the gut, peristalsis, which is unpredictable, but may occur on the order of minutes, maybe as long as 20 minutes between severe motions. We also have muscular spasms, which you can't predict, and various motions by the patient due to discomfort perhaps, that are unpredictable. And you have to be aware of these things, and here I give an example. If you look at the PET in this image, you see that you have some kind of a shadowing of this focal uptake that you might wonder, "Is that two nodes? Is that two areas of uptake? Is there something funny going on here?" And you have the CT image, which doesn't really reveal much, but if you look at the image without attenuation correction, you really see a kind of doubling of the shoulder. And what this suggests is that basically this is a double image in the sense that the patient moved halfway through the PET acquisition,

and so if we look at the CT, we see that the CT corresponds to one position we see the patient in, but not the other, and the attenuation correction kind of removes the part of the shoulder in that second position that would give us a clue. So in situations like these where you have a suspicion of motion, you can get a real benefit from looking with and without the attenuation correction applied. I'd like to talk briefly about the use of contrast agent in CT. Contrast agents are very x-ray-dense materials, typically with a high attenuating element in it. They also have to be tolerable to the body, giving effectively two choices, iodine and barium. We may have oral contrast agent that's swallowed or IV contrast agent injected; a bolus is injected, which then disperses throughout the vascular system, with a small risk of an adverse reaction. So why do we use contrast agents in PET/CT? Well, CT contrast agents can improve the diagnostic utility of the CT images. I give some examples here with IV contrast in head and neck. You have a lot of small

dense, densely-packed structures there, and in a transverse view, it may be hard to tell a vessel from a node, for example. When you enhance the vascular structure, you see a lot more detail. Things become a lot clearer. Similarly, oral contrast is used because it's very hard to tell whether something is part of a

complex bowel loop or whether it's a cyst or something suspicious for a tumor on CT. But with oral contrast, you definitely see the whole bowel separated, and it's much clearer. So these improve the CT images. They don't affect the PET, except in the sense that they may introduce biases through the attenuation correction, which are generally small. Other than that, their benefit is for the CT and not for the PET. You may look at how big this bias is. The CT contrast agent has very little effect at high-PET energies, so one can kind of correct for that bias by replacing the enhanced values in CT with water values. This is a measurement of some contrast agent in solution. Bottles were prepared with water and

contrast in the same PET activity, so PET should be the same in both, and we have a very dense scan on CT with 900 Hounsfield units in that contrast bottle, and you see that CT-based attenuation correction gives you a fairly substantially increased value for that activity, which should be the same as the other one. If we set the values to 0, we get the result we expect in PET. So here's an example of that positive bias taking place. And it's important to realize, though, that you never get enhancements as high as 900 Hounsfield units. For one reason, you don't see them. You can distinguish much lower enhancements in CT, and secondly, you don't get huge, filled bowels of that size as with this water bottle. So this is really a kind of very much an upper limit, and the biases in PET are expected to be much more modest. So I'd quickly like to talk about the second part, some remarks on image fusion in PET/CT. In terms of fusion displays, there's lots of different nice packages out there. Every vendor has their own flavor of fusion

workstation, and there's many excellent third-part packages too, so rather than try and discuss any one of them or their particular idiosyncrasies, I'll just try and review what they all tend to have in common. You're typically going to have PET/CT and fused images available in transverse, sagittal and coronal sections, and all of them will typically have link cursors, allowing you to ortho-navigate through the images, and to grab pixel and ROI values, both in terms of the PET and the CT image, and some of the color tables may be familiar. You're typically looking at a gray scale CT image, and the – an inverse gray scale would be favored for PET generally. And then you have fused images, fused through alpha blending, which essentially allows you to overlay the PET in one color table onto the CT, typically with a slider that allows you to make the overlaid image fully opaque or fully transparent, and you get the images like I show on the bottom right, using a black body color table for the PET. And so typically,

you'd be able to draw an ROI on one image and have it appear on all three, and from the CT, you'd be reporting Hounsfield units for that ROI, and on the PET, you could get the standard uptake value, depending on perhaps the body weight or the absolute concentration measured by the scanner in becquerels per milliliter. So this slide is just to summarize some of the benefits of what David Townsend called hardware fusion yesterday in PET/CT versus software fusion, for example, you have other images immediately available as part of the scan session, you only have to position the patient once, because it's the same bed, you have less internal organ movement because you're not shuffling the organs around as you move from one scan to another, even if they're done almost at the same time. You have improved registration accuracy. It's a single scan, so that's a very practical convenient factor. And you don't need to align it, in the sense that you've already performed a hardware rigid registration simply by having it in

one system and one scan. Given time, I'll just concentrate on this example. This, an ovarian cancer case, this is to just emphasize the benefits of localization in PET/CT, in this case, there was post-surgical changes in the pelvis, and on the PET you see that the suspicious uptake is actually very much focused at the location of surgery and is suspicious for tumor. And so this kind of detail really only comes about having the availability of images that are confidently aligned. This is another example where we see in the

PET, we see a small lymph node at the left neck base. However, on CT this is seen to be within normal limits. On the PET scan, you see an area of focal uptake that's indicated. However, you might wonder, you know, is this the same node? Can I really be sure? Do I really know what this is? But with the fused images, you have a very high degree of confidence that that is the suspicious node, and you can give a very direct instruction as to what should be biopsied or removed. I'd like to talk a little bit, finally, about

how we interpret artifacts in these fused images. We know that scaling, tissue scaling, as I described it, anything that's really dense on CT, is going to lead to artifacts in PET. Examples would be prosthetics, metal, bolus IV contrast, and precipitation of oral contrast, and the first key to understand these artifacts is, is there something on the CT that looks like it's causing this? And I give two examples. Here one is a subcutaneous chemotherapy port, which is very dense on CT, and you see a hot spot on PET that you would suspect is an artifact. Also, if you have IV contrast and you don't flush it or you scan too early, you can get a bolus in the subclavian vein before it's diffused throughout the vascular system. And this bolus would have very high, maybe thousands of Hounsfield units on CT, and that's going to scale and give you maybe some artifactual focal uptake. And again here, we see that it really corresponds exactly with what we see on CT, so we understand why it's there. A slightly more unusual case, here we see suspect

artifactual uptake in the stomach, which correlates with very high Hounsfield units on CT, which may be due to precipitation of the oral contrast, so it's not ... your contrast protocol isn't behaving like you'd like. However, it's a bit unusual, so you might be a little less confident about what to make of this. Well, here you can also look at the images without attenuation correction, and we see nothing there, so that really confirms our suspicion that in this case, it was an artifact introduced by CT-based attenuation correction, where there was a very, very high Hounsfield unit structure on the CT images. I'd finally just like to make a quick remark on the use of software fusion in PET/CT. We may think of hardware fusion in PET/CT versus software fusion a little too much as two opposing things rather than two things that can come together. A PET/CT certainly provides a very good overall rigid registration and more, because we're preventing organs moving around from one scan to another, but differences are still going to persist

between the PET and CT due to the motions that we can't avoid, such as respiration and complex movements in the abdominal region, GI peristalsis and gas formation, etc. So there really is a utility in applying some of the more sophisticated methods we have now in software fusion to PET/CT. I give a quick example here, where an original image PET/CT has a deformable algorithm applied to it, and it gives us a displacement map that we think might be quite meaningful in the sense that the position of the liver or the outer liver wall is something we might expect not to be well aligned, due to the respiration artifacts. And so here we're getting a result that certainly these algorithms require more evaluation, but may have real benefit in improving even further the registration in PET/CT images throughout the body. So I'd finally just like to acknowledge everyone, my colleagues and everyone at the Cancer Imaging and Tracer Development Program at the University of Tennessee, and I'd also like to thank Vitaliy Rappoport at CPS Innovations for his contributions. Thank you very much.