

Thanks. Good afternoon. Thank the organizers for inviting me here today. It's a pleasure to be here. I'm a physician. I take care of patients mainly with lung cancer and breast cancer, and we've been doing studies of functional imaging in the clinic. I'd like to just briefly review the two main studies that we've been doing and some of the pitfalls and benefits and the risks and that sort of thing. So this is the traditional approach to 3-D information. Dose distribution is reduced to a DVH, and when one tries to relate metrics from the DVH to the outcome, and this always didn't seem quite right to me because it ignores all this stuff. It ignores all the anatomy, the physiology; all the spacial information is discarded obviously in a DVH. So we've tried to take some of this functional information and generate what we've called a dose function, or DFH instead of

a DVH, and then try to relate this to the outcome, and then try to also include some things related to compensation. And I'll show some of these things as it relates to normal tissue mainly. The same schematic can be used for tumor or for different normal tissues. What I'll spend most of my time speaking about is our studies in lung. We'll go on, we'll show some of the data that we have in the heart, some information regarding tumors. And the theme through all of it is this functional business, the DFH versus a DVH, and I'll try to highlight the frustrations and limitations of this approach. So in the long, the two main areas we'll look at is, we'll look at limiting dose to the good lung, and then studying regional injury in the lung, and this will be obvious when I show the examples. This is the modality we've been using mostly. This is a SPECT lung perfusion scan. We chose to

use this because this appears to be the most sensitive indicator for regional lung function. You could do the same studies that I'm about to show you, you could do them using ventilation imaging, which colleagues at the Netherlands Cancer Institute have done, both perfusion and ventilation. But this is a typical pre-treatment normal SPECT perfusion lung scan, and it's oriented just like a CAT scan, so from top to bottom and obviously the normal lung is black, that's blood flow, and the center is white. That's where the mediastinal structures are, the heart, blood vessels, and whatnot. So this is a normal pre-treatment scan. So here is a patient who has a lot of tumor and a lot of emphysema involving the left lung. So this is a set of beams, this is an anterior, posterior, and an oblique beam. And I have the isodoses drawn out here, and you can calculate the DVH on

the CT-based lung volumes, or you could then calculate the dose function histogram, where this is now, instead of percent of volume, it's percent of counts, percent of blood flow getting each of the doses. And as you can see, depending on the case, the DVH and the DFH can be very different, and similarly the equivalent uniform dose that you compute can be very different. This is actually one of our first patients. I still remember doing this case with David Spencer about ten years ago – a patient with a small cancer here in the left lung, and this is the pre-treatment perfusion scan. As you could see, on the CT, this part of lung looks just like the rest of it, but on this perfusion scan, you could see that there's a lot more blood flow here and here and here, and there's a relative positive blood flow here and here. So if you did an APPA field versus an oblique, you would get

these two different DVH's, but if you compared them with regard to the DFH, you could

see that there is a greater difference between the oblique versus the APPA beam. And I guess at this point, we called it a DV subscript VH. So you can see that the functional information might lead you to choose a different beam arrangement than you might have used otherwise if you just based your decision on the DVH. And I should say, honestly, this was a surprise to us, the idea that the pre-treatment SPECT scan is going to have these dramatic heterogeneities, we didn't expect. When looking back on it, it seems obvious. Of course there would be heterogeneities, but at the time we were doing this study only to look at changes in SPECT perfusion over time as a result of the radiation. We expected the perfusion to be uniform before the radiation, and in retrospect, that was

naïve, but this was, I think, patient number two or three. And we said, "Gee, this is really neat. You could actually aim the radiation beam to try to miss the functioning lung." How often do you actually do that, or how often do we actually do that? In short, you don't do it that often. Okay, most of the lung cancers tend to be big – big, central mediastinal masses, and usually giving most of the dose from AP and PA still winds up being best. And in order to use this approach, you need ability and desire. What I mean, you need the ability, you can only use multiple beams oriented from strange directions if the lesions are small. And you need to have a desire to do it. The lung function's got to be really bad. So there are the occasional patients, and we've estimated it's about 10%, maybe 15% of our patients, where the tumors are small, giving me the flexibility to adjust my beams and

the lung function is poor enough that I want to adjust the beams, that I actually do orient the beams based on the SPECT information. Now, I tend to treat elective nodes, and there's clearly a movement among physicians to do less nodal radiation in lung cancer, so as you get away from nodal radiation, the target gets smaller, and therefore, maybe you have more ability to use these unusual orientations. And also IMRT gives us much more flexibility and where to dump the incidental dose into the poorly perfused areas of the lung. So the major study we've been doing at Duke for about the last 10 or 12 years or so is doing a prospective study, a looking at radiation-induced changes in lung function, and you can break up the end points into regional and global, and also clinical and subclinical injury, and we've focused mainly on this box, and I'll show most of our data relating to

this box, looking at regional injury, which is a subclinical injury, obviously. So what we do is we take local dose. We try to relate it to local changes in function using this perfusion lung scan, which, by the way, is a microemboli study, and then we try to sum up the local effects to see if this relates to changes in whole lung function, pulmonary function tests, development of symptoms, and whatnot. So again, we start with the same perfusion scan. So we've done this now, and over 300 patients have been enrolled, and most of them have lung cancer, and we do serial post radiation evaluations, typically 6 months, 12 months, 24 months, 36 months. And this is lung cancer, so you have tremendous falloff of patients. The one-year survival rate is somewhere 20 and 50%, depending on the subgroup. So we have very few patients at the longer time points. But

again, we look at imaging, as well as function. So this is a typical patient. This is a pre- and a six-month post radiation SPECT scan. As you could see, within the 60 gray isodose line, you could see a marked reduction in perfusion, and you can't quite see it, but at the

lower isodoses here, there is reduced perfusion, but not quite to the degree that you have at the high-dose area, and this is a dose response curve for this individual patient based on this patient's data at 6 months. And you get this very nice sort of dose response curve here. That's repeated now in many, many patients. This is the data, I believe, on our first 20 patients, and we did this at that point, at some shorter time points as well, 1-1/2 months. And basically, this is the short time points. These are the longer time points. And it doesn't change much after about 6 or 12 months. These curves of two and three years

look about the same. All right? So you get this sort of gradual dose-dependent reduction in regional function after the radiation. Well, I should say that we were surprised by this, because the classic clinical teaching is if you give anything less than 30 gray, the lung still works, and if you get much above 30 gray, the lung stops working. So we expected a very steep dose response curve at about 30 gray, and we got this. And we got worried, you know, we're doing something fundamentally wrong. And then we have this wonderful publication from Joos Lebesque and colleagues. They have published this data several times. This is just one of their figures. Their data is the open symbols, and the Duke data are the closed dots, and as you can see, the data is almost superimposable on each other. So, this is a real phenomenon, true in Europe and in Durham that their dose-

dependent reductions in perfusion that can be detected by SPECT. But so what? All right? We proved this one. But so what? Patients don't care if I tell them their scan is abnormal. They care about how they feel. Can they walk? What's their pulmonary function? So we've tried to then take these local effects and see if we can predict changes in total, global lung function, in PFT's, which is the pulmonary function test, and then I'll show some data afterwards for symptoms. So essentially, you can call this the overall response parameter, which is how Joos Lebesque likes to call it. It basically, is the integral injury. You take the dose response curve here, and you take your differential dose volume histogram or your dose function histogram more accurately, and you essentially sum up the effects at every dose bin, and see if we can predict the change in

whole lung function based on the abnormalities seen in regional perfusion. It sort of makes sense. The lung is an organ. It's structured in parallel. The change in global function should be related to the sum of the regional injuries. So here's the data, you know. This is eight years of work, and this is data on the first 90 patients or so that are valuable six months out. And as you can see, there's terrific scatter in the data. The prediction based on the summing up, the integral injury, and the actual reduction in pulmonary function, the DLCO is just one of the many tests that are part of a pulmonary function test. As you could see, here's the line of unity, and if it was an ideal world, all of these dots would lie along a line, and obviously, they don't. There is a relationship. They are statistically related to each other, but the correlation coefficient is very bad. The R

squared is 0.09, which is a terrible R-squared value. So what's going on here? Well, it appears the tumor location matters. Here's a patient with a tumor here in the right hilar area, and you can see this thing, which we've called hypoperfusion adjacent to the tumor volume. And you can see this, and you can tell on the CT, there really isn't too much remarkable about this piece of lung compared to the rest of the lung, but you have this

reduced perfusion here, and we've hypothesized, and then think I'll show in a minute that this matters, because as this tumor shrinks, you're going to have some reperfusion of this part of the lung, which is going to compensate for the injury that the radiation caused. So if you look at that correlation coefficient again, this is in the subgroup of patients with lung cancer, so the number of patients has changed, but if you look overall, the R-value

was only 0.4. But if you then throw out the patients who have this adjacent hypoperfusion, the patients where the tumor itself is causing changes in the blood flow, if you take them out, the R-values get a little bit better, especially when you look at the patients who have one-year follow-up, but when you get out here, you get to very few patients. The number in parentheses are the number of patients. So you can see, in the patients with longer follow-up, if you throw out the patients with this hypoperfusion adjacent to the tumor, there actually is a fairly good relationship between the sum of the regional injury and the change in whole lung function. So there is hope here, but the correlations are by no means good, and one wouldn't want to make great clinical decisions based on these sorts of correlations. Again, here's the comparable data from

Joos Lebesque's group in the Netherlands. This is in their patients with breast cancer and lymphoma, so if you take away all the patients with lung cancer, and say, "Okay, these patients should not have tumor effects on their blood flow, this might be a better group." Even in this "better group," their R-value is still only 0.6. All right, so this approach might work, but it's not really a home run. Well, why isn't it a home run? Well, let's look at the surgeons' data. This is lung that is not irradiated. This is lung that is taken out, okay? If I take out a certain percentage of the perfused lung, what's the percent decline in pulmonary function after surgery? All right? And how good are the surgeons in predicting changes in pulmonary function? Well, here's the R-values that they get for DLCL, which is what we did. We had R-values that were, you know, 0.6 or so, 0.5, not

that different than what the surgeons get, all right? So here you really do have a step function. You know, I love this. In the patient, or the lungs out of the patient. None of this gradual dose response business. This is lungs in, lungs out. And even here, under the best of circumstances, the surgeons are getting things in the range of 0.8, so it's not that surprising that with radiation and a dose response curve and tumor shrinkage, and all these other effects, that we're only getting things like 0.6, 0.5, and maybe in the light of how good or how bad the surgeons do, getting a 0.6 or a 0.5, maybe that isn't all that bad for a correlation coefficient, given our complicated a problem this is. Here's to show some of the data regarding the central tumor business. Here we are predicting declines in lung function when indeed fully 40% of the patients who have these central tumors with

hypoperfusion have improvements in their breathing tests. Remember, I showed on this graph all these dots that are below the line. These are patients whose breathing actually got better after the radiation, not worse. And those patients are these patients, who had the adjacent hypoperfusion present at the start of the radiation. So what about predicting for symptoms? This is some data that Moyed Miften has recently compiled. The patients with pneumonitis versus those without pneumonitis. Here's the FEUD or the functional EUD versus the EUD. This could also be moon lung dose and moon perfused lung dose,

any metric here that reflects the SPECT information. And it turns out that the patients who have relatively poor perfusion in the lung that's being irradiated, tend not to get pneumonitis. It shouldn't be a surprise. If you irradiate the functioning parts of the lung,

you're more likely to get symptoms than if you irradiate the nonfunctioning or non-perfused parts of the lung. So the lung, I think we've shown that this approach is a little bit better than this approach. The clue at dose-dependent changes in regional perfusion, there is an association between regional and whole lung changes, and there is some effect going on with central tumors responding that impact how patients' whole lung function changes after radiation. So what did we do in the heart? Exactly the same thing. We took local dose, related it to local perfusion changes in the heart, and then we try to relate that to changes in heart function. Now we were talking about wall motion, ejection fraction, episodes of chest pain. So, same sort of thing. The study started several years later, so we don't have as many patients. The good thing about the breast cancer patients are they live

a lot longer, so we have many more patients in follow-up. The data I'll show are based on 90 patients who had normal pre-treatment perfusion scans of their heart. These patients all got photon tangents to the left chest wall using conventional fractionation 2 gray per day, 46 to 50 gray, and I should acknowledge Patty Hardenburg, who started this study when she was a faculty at Duke in the nineties. So this is a pre- and a six-month post radiation SPECT perfusion scan. Again, this is showing blood flow to parts of the myocardium, and you can see that there is reduced perfusion in the anterior myocardium, the part that corresponds to the deep part of the radiation field. And you see this in almost every patient. In patients who have, in almost every patient that has significant volume of left ventricle in the field, I should say. So if there's a very small amount of left ventricle

in the field, you tend not to see this. In the patients who have five to ten or greater percent of the left ventricle in the field, they do get a defect. And I point out left ventricle only makes about half of the volume of the heart. So when you see a DVH of somebody's heart, and there's 2% of the heart in the field, or 3% of the heart, don't be reassured. Those patients have a 70% chance of having a perfusion defect in that heart from the radiation. I personally think this is one of the most under-recognized and underappreciated toxicities out there. All right? There clear data that we cause heart disease in breast cancer patients, and I think this is, in part, one of the reasons. And you can detect this six months after radiation using these functional imaging tools. But so what? So what to have a functional perfusion defect? Does it relate? Does it predict for a

functional change? Well, it does, all right? These are the rates of wall motion abnormalities. This is on a test that the radiologists say the heart muscle is not moving normally. Or actually, it's a cardiac radiologist. Is the heart muscle not moving properly. And in the patients with perfusion defects, something like a third have heart muscle that's not moving properly. That, to me, seems like an enormous problem. In the patients that don't have perfusion defects, far fewer of the patients have wall motion abnormalities. So in the heart, I think we've shown that radiation causes perfusion changes, and this does lead to changes in wall motion. I've not shown the data, but there are some changes in

ejection fraction in patients who have large defects in the heart, and I'd also point out that there is data for about four of the centers corroborating this data, albeit with much smaller

numbers, but there's data from other places that show essentially the same thing. I'll also just mention quickly that in the patients who have only 1% of the left ventricle in the field, a group of them had a defect, about 20%. Well, we pulled out the port films of those patients, and it turned out that the port film, the ones who had the defect are the ones with the port films systematically set up about 2 mm too deep. And as a physician, I check a lot of port films every day, and if it's within 2-3 mm, it gets signed. Right? No one makes – requests a shift for 2 mm, even 5 mm, you get criticized sometimes by the therapist. All right? But in our patients, in this group, right here, this 20% of patients who had a defect, almost all of them are patients whose port films set up deeper than was planned. All right? So what about the tumor? We've been using PET imaging to help

guide the therapy for lung cancer and for other tumors for many years. This is an early study of a Mike Munley of our first 25 or so patients where the radiation fields were changed in about a third of the patients, and now we use PET in almost of the curative lung cases to adjust our radiation fields. This is an example of one of the patients that we studied with Su-Min Zhou. The red is the GTV based on the CT. The white is the FDG PET abnormality, and we're looking at things like PET versus PET-weighted EUD's versus traditional CT-based EUD's to see if this relates to outcome. And as we start to move into IMRT, there are a lot of people looking at purposely putting hot spots in these parts of the tumor. Cliff Ling and many others have written about this dose-painting business of purposely delivering a heterogeneous dose. Here's one of our patients; this is

purposeful delivery of a higher dose to these areas. I always liked this 91 gray line, looks just like the United States, I always thought. See? Anyway, so you can see, you could scope the dose very nicely using IMRT. This is a study ... we didn't actually treat the patient with this. This is done by Shiva Das. Moyed Miften has done a fair amount of work along the same lines using IMRT, and if you optimize your IMRT plan based just on the CT-defined target versus the PET-defined target, you do get a different plan out of the IMRT software. You can see the result in dose distributions are slightly different here. And in a few patients, again in study form, (we've not formally treated patients based on this in any systematic way). We have, though, in study mode, and occasionally looked it, eyeballed it in a real patient. Look at the PET abnormality, merge it with the SPECT

perfusion scan, so that we have total functional imaging. I have beams to try to go through the nonfunctioning parts of the lung, like here, that are low perfused, and try to hit the hot areas on PET scan which we think might be the most important areas of the CTV to hit with the high doses of radiation. I think there's a huge area to be exploited here in functional imaging to use imaging as an early marker for response. Here's a summation of the data. This is probably about a year old now, or six months old, where people did PET imaging early on during therapy for a variety of tumors, for a variety of therapies. Here's, for example, after one or two cycles of chemo, here's at the end of treatment, here's after three cycles of chemo, two weeks into chemotherapy, during the radiation, so different times during therapy, and this is for the different end points,

usually free of disease. In the patients whose tumor was responding based on PET early on during treatment there was a far better outcome in all these studies than in the patients whose PET scan did not respond. So I think the day is upon us very quickly, when we're going to be able to tell a patient within one or two weeks of the start of chemotherapy or the start of radiation that their tumor is changing or shrinking or functionally changing rapidly enough to want continuing treatment because it might be a lot of these people who are not responding, it's no use wasting their time continuing with a therapy that's not going to work. And it seems to me that PET imaging and perhaps MRI imaging, functional imaging of the tumor is a terrific way to monitor therapy early on in therapy to help us tailor the therapy better for our patients. So in summary, I think I've hopefully

shown that the functional imaging does allow us to consider spacial information, and if the dose volume histogram discards spacial information and functional imaging helps us bring some of that back in, the concepts of EUD versus FEUD, functional DVH's, potentially looking at these metrics, the EUD and the FEUD and the DFH to rank plans, I think it's a terrific way to study normal tissue effects, and I've shown data looking at SPECT mainly, and a little bit with PET. I've not reviewed the world's literature for you here, but there's certainly other people who have looked at MRI. There are some imaging studies done in the parotid glands as well, which I've not shown you. So I'd like to just acknowledge the many collaborators at Duke. This involves many physicians. The pulmonary work involves a pulmonologist. The heart work, I've worked closely with a

cardiologist, multiple statisticians, protocol coordinators, a bunch of physicists here. Mike Munley did much of the early work. He left Duke. So Mike, so if you're in the audience, I took your name off the slide. I'm sorry. And very close collaborations with our colleagues in nuclear medicine. I'd also like to acknowledge Varian, as well as the NIH and the DOD for grant support, as well as Julian. All of the images I showed were PLUNC-based images, or for the most part were PLUNC-based images. We were very fortunate to have a nice working relationship with UNC to use the PLUNC software to do these studies. And thank you very much.