PROFESSOR: OK. First of all, I was in touch by email with George, and he asked me to flash this up again, which was from a couple of lectures ago, because he thought this was very critical and very important that you sort of understood what it says. So I'm going to just say a few things about it.

I guess the first thing to note is that it compares these two different cases, coherent imaging and incoherent imaging. And you see that, apart from that, these two, the top half and the bottom half, are really identical. So you use the same method in either case, except that, of course, in the coherent case you're working with amplitudes, whereas in the incoherent case you're working with intensities. But apart from that, it's identical.

So what it's saying is, if you illuminate a thin transparency, then you get a field on the far side of this transparency which goes into your optical system. So this is his GN. And to calculate the image, you can either do it by convolving that with a point spread function to get the amplitude out, or you can do it in the Fourier domain by Fourier transforming to get the object spectra, and then multiplying it by a transfer function and then getting the field out. And then once you've got the field, of course, you can calculate the intensity from that just by doing the modulus square.

So the difference between these two things is here we working with aptitudes, so we have a coherent point spread function and an amplitude transfer function, whereas for this one, the incoherent case, we have an intensity point spread function, which is just the modulus square of this. And here we have an optical transfer function, which is the autocorrelation of this. So that's really all there is to do, all there is that you need to know about calculating images.

Any questions about that? So I'll just-- as I say, I'll just flash that because George said he wanted me to. And now we'll go on to where we got up to last time in the lecture.

So George suggested-- he sent me some more material for today, and he also asked me to add some extra to it of my own material. But I don't know how far we get with this. So we've still got some left over from last time, actually, and then the bit that George sent.

You might be fed up by then. But anyway, if we've still got some time and you're not fed up, I'll carry on with some stuff that I put in there which is to do with focusing. But it might be-- you might think it's sort of rather higher level than the rest of the course. So this is not for-- this bit that I've added is not for examining or anything. It's meant just for your interest.

Anyway, this is where we got up to, the defocus ATF. So this, again, you remember-- let's make it bigger. Now, what's going on? I've gone the wrong thing. I should [INAUDIBLE]. Sorry. I haven't got my glasses on. I can't see what I'm doing. And it's all in a different place from across the version of PDF that-- PowerPoint that I use.

Anyway, here we are, the defocus ATF. So you remember, we introduced this last time. We said that if you've got a defocused object, then you can work out the imaging of it by using an amplitude transfer function just like before, but it's not just the normal one. It's modified by the effects of defocus.

What happens is that you have to multiply the transfer function by complex exponential, e to the i, this thing. It therefore becomes complex. And you can divide that up into two parts, a real part and an imaginary part. The real part will image the real information of the objects, which is basically the absorption of the object. And that's what this shows here, the real part of that transfer function, which is cosine of this thing.

And so this is the cutoff of the optical system here. Sorry. No, it isn't, is it? No, sorry. The cutoff, I guess, is some arbitrary place I guess it's not actually shown here.

But anyway, the fact is that as you change the amount of defocus, this will be scaled accordingly relative to the cutoff of the transfer function. So you can imagine if the transfer function was right in here, it would be almost as though it was the same as an in-focus system. And then as you make the cutoff wider and wider relative to this cos curve, you see you getting more and more of these wiggles. So this was a case where it's not changing very much.

So there we are. There's the region where it's non-zero, and so it's not changing very much in that region. And so we said that corresponds to multifocus.

And then the next case we looked at was now when it's got a lot of defocus. So you can see here, it starts really wiggling. Remember-- I said this last time-- this is actually cos of this thing squared. So that's why it doesn't look like an ordinary cos function, because of the x squared. So it scales it in this nonlinear way.

And so here's the and it cut off again in this case. Now, in this case you've got some wiggles, and those wiggles are going to lead to artifacts in the image. Firstly, because of the fact this goes negative, you can see there's some spatial frequencies will be imaged with the wrong sign. So it's just like if you've got a Fourier series for, let's say, a square wave, if you take the Fourier series for a square wave and you change the sign of some of those components, it doesn't look like a square wave anymore. So you get a very poor image in that case. So this is really bad news for the image.

And the other thing that George mentions up there is that, of course, there's regions where the value of the transfer function is very small. So this spatial frequency here, for example, is not going to be imaged at all. So you'll get some gaps in the response of the different transfer of spatial frequencies.

So the significance of depth of focus in imaging-- so here he shows a 4F system. It shows here the NA in and the NA out. And you remember, of course, that as you changed the NA, you also change the magnification. You change the lateral magnification, and you also change the longitudinal or axial magnification. That goes as the NA squared.

And so we define these two terms. Depth of field is basically that the defocus tolerance in the object space, so far you can move the objects without it going too much out of focus. And so this gives some idea about the value of that, then. The depth that depth of a field is given by this lambda over 2NA squared. You get this NA squared because the longitudinal magnification also goes as NA squared.

And then on the other side, in the image side, you have a depth of focus rather than a depth of field. Notice in the title, he calls it DoF, which can mean either, so Depth of Field, Depth of Focus. You'll find that, of course, in general terminology people often use these phrases almost interchangeably and are not very careful about using the right one for the right space.

But this is what it really means. Depth of field is how much you can move the object. Depth of focus is how much you can move the screen here without it going out of focus.

And you can see that this is expressing the depth of focus now in terms of numerical aperture. You get the same expression, but it's now this NA rather than that one. But of course, the magnification also goes as NA squared as well. So the image of the bit that's in focus, here, of the object is, of course, the same as the depth of focus of the system.

Is that all clear? Yeah? Good. OK.

Even though our derivations were carried out for spatially coherent imaging, exactly the same arguments you can use for the incoherent case. So the idea of a-- we showed the defocused amplitude transfer function, in the same way you can also have a defocused OTF.

So let's think of now what happens if we've got some objects which has got some height variations. Here he's taking the case of a surface, and he's called this a 2 and 1/2 D object. I think he's probably called it that because I sometimes use that terminology. I'm not the only one, really. I didn't invent it.

But people often use this 2 and 1/2 D to represent something which is not a 2D object. It's not flat, it's got some depth. But it's not a true 3D object in that it's got no internal structure. It's just a surface.

So anyway, in this case, this is the object and this is its image. And you can see that if this object is quite deep, it's going to be deeper than that depth of field of the system. And so some parts of that are going to be out of focus. And so this is a region, for example, which is a long way away from the focal plane, and therefore it's going to be imaged out of focus. And here it corresponds to this region here.

Notice, if you look at this image, it's inverted in x. But you see that in the z direction, a distance which is moved this way here is also moved this way here. So there's no inversion in the z direction. Did I hear a question over there or was it just someone shuffling around? No. OK.

So as George points out here, what this means then for this particular object is the point spread function is not now spatially invariant. It's a spatially shift variance PSF because different parts of this surface are going to experience a different point spread function because they might be in or out of focus. So yeah, and the part, just like before, the part which is in focus is given by this depth of field, which we just had before. So you can either think of it as a depth of field or a depth of focus, the same expression.

So this is an example of this. Let's take an object which is a letter M. I guess he didn't use MIT this time because he wants to use the same for each so that you can see how each of them is affected by the different point spread function. So this is in the focal plane, and then this is to depth of fields out of the focal plane. This is for depth the fields out of the focal plane.

So when you image that you can see that it's going to become-- this one's going to be in focus. It's still a bit blurred, of course, because of the refraction effect in the imaging. So this is convolved with the point spread function of the coherent optical system. And then this is for the two defocused ones, getting more and more out of focus.

So right now, once we've got that blur, what can we do about it? And there's been quite a lot of papers that have looked at this question of whether you can actually undo the blur. Deblurring, they call it, or sometimes deconvolution because if you think of the blurring as being a convolution, you wanted to undo that, you want to do a deconvolution.

So what we actually get then, in the Fourier domain the spectrum of the image is equal to the spectrum of the object multiplied by some transfer function. And this would be true whether it's coherent or incoherent, in general. They'd be the same sort of effects, but you'd have to work with intensities rather than amplitudes.

And so what we say then is that if we-- this is the blurred thing, so we can actually blur it by just dividing by that to get rid of it. So if we get our image, we multiply by 1 over the transfer function, and that gives us our deblurred version of the image. So as he calls it here, you can call it inverse filtering, or sometimes called deconvolution.

But that is actually-- probably you all thought of, that can't work. It doesn't sound as like it ought to work. And of course, it doesn't work very well. We're going to show a bit more about this in a minute. In particular, it really goes wrong if this thing goes to 0, obviously. So you can't know very much about that.

And then also, if it goes to something which is even very small, you have to amplify it by a lot. And when you amplify it, you're going to amplify the noise. So you'll end up with a noisy reconstruction after you've done that. So this sort of inverse filtering invariably increases the noise in the image.

And so how can we get around that? Well, we just get around the 0 by making it so that there isn't a 0 anymore. So this is our new version. You can see that if mu were 0, this would just become the same as before. This is H times H star, so the H stars would cancel, and this is 1 over H.

So if this is being compared with that, then it becomes the same as before. But if this becomes small, then we're just left with this, and it avoids it blowing up and giving infinity. So if mu is 0, it reduces to the direct filter, and we just described that that's not a good thing to do. So the more noise there is there, the bigger you need to make this to avoid getting too much noise in the image. So as the noise increases, you would tend to increase the value of mu.

Yeah, so here it's saying that if the signal and noise both obey particular type of statistics, then you can find what the optimum value of mu is. And we find then that it's equal to 1 over the signal-to-noise ratio, and this is what's called a vena filter. So this N comes down to be the same as a vena filter, which you might have come across, some of you, in other things.

So what we are doing, really, is you're amplifying the spatial frequencies that have been pushed down by the imaging process. But if there is if they become smaller than the noise, then you don't do that anymore.

So he's going to give an example of this. So this is the convolution using this second of regularized inverse filter. And yeah, OK, he's saying here that he's assuming he knows that depth before he does this. So obviously, that is something you wouldn't always necessarily know.

But he knows the depth, therefore he can calculate what the transfer function was on the imaging side. And therefore he knows what filter inverse filter to apply in order to get the best reconstruction. And you can see that it works extremely well for this one. This one is slightly blurred, slightly a bit more blurred. And you can see also that the background here-- or at least, I can see on this screen, but this one doesn't show very much-- the background is becoming not black anymore.

And then this is what happens if you've got a noisy image, if you add noise to the original-- to the image before you do the deconvolution. Now you see that that actually it's not working nearly so well. And this one here, you can't really make out what there is there. So this is showing that these methods can work, but you can see that there are also limits to how well they can work, and especially in the presence of noise.

Actually, I should mention that, of course, at MIT in the Media Lab, there's a lot of work going on in this area of what they call computational photography, which is all to do with changing the properties of the camera in order to make it so that this sort of reconstruction is more efficient. So it turns out that if you're clever about the design of the imaging system, you can actually make it so that this all works a lot better. But we weren't going to go into all that.

Oh, sorry. I've got this wrong. This one was the-- sorry, this was the-- oh, yeah, that's right. This is the image before you've done the inverse filtering, and this is the image after you've done the inverse filtering. So it's actually doing quite well. This one here is pretty well not very different from that. And this one, as before, is not as good, but doing pretty well.

Of course if, you didn't know what the height of these things is, then you might have to use some iterative algorithm in order to come up with the best reconstruction. You can imagine also that-- so this is for this so-called 2 and 1/2 D object. You could imagine an object which really did have structure inside itself, so maybe these three Ms would be on top of each other and semi-transparent. And then the trick would be to try to actually produce a reconstruction of the 3D object from the 3D image, and there are actually computer packages that you can get to do just that very sort of thing.

So that was what we should have done, and we'll now go on to today's lecture. And so today's lecture is, first of all, about polarization, which we've been talking about in the lab today-- or yesterday, wasn't it? So that's a bit topical. And then George wrote these. So the bit that I've added extra is basically addressing these points here, effects of polarization on imaging in a microscope objective or a lens with a high numerical aperture.

So we'll start off with the polarization. I'm not sure what he said about polarization, first of all. But basically, then, the wave equation for an electromagnetic field is basically a vector wave equation. It's a function of electric field here. And we had then that the polarization is some function of the electric field and how the relationship between these depends on the sort of material. And so this is giving some different examples.

The simple case is this one, which we did look at earlier, is where the polarization is proportional to the electric field. And you remember that you get D by adding the P to the E. And so here, this is called the electric susceptibility, and P equals Chi times E. So if this is a constant, this is just a linear relationship, which is how we described it before.

But there are two ways that that can break down in practice. And usually, actually, for materials where it breaks down, it breaks down for both at the same time, unfortunately. So it becomes very complicated. But here we've written it with these two effects separately. this. Is where actually the relationship is still linear, but P is related to e not by a simple scalar relationship like this, but by a tensor relationship. So the direction of P is not in the same direction as the direction of E. And so you've got this 3 by 3 matrix, which describes the properties of the material.

And this will depend on-- well, for a isotopic material, like a crystal or something like that, it would depend on the crystal structure. Of course, if you've got some complicated varying objects, then it can be all sorts of even more complicated things. So that's the first case. So this is what's going to lead to some interesting polarization effects as you put that electric field into a crystal, for example.

And then this is the other one which is now looking at isotopic material, but nonlinear. So this is now no longer a constant. And so what that means, you can you could expand it as a power series. And you can see that there might be some extra terms here. This is what's called the Kerr effect. There's effectively a susceptibility which is dependent on the intensity of the light which shines on it. And of course, you might expect that there will be higher orders of this as well, of course.

These sorts of effects, you're only going to excite these if you have a large value of E. It's a bit like, if you imagine how this all comes about, this comes about because your material is made up of atoms which are all bonded together and joined to each other with springs. And when you apply the electric field, it makes these atoms vibrate against the spring. And if the vibrations are small, then it's all going to satisfy Hooke's law, and it's going to be linear.

But if you increase the forcing term, then the vibrations get bigger, Hooke's law breaks down, and the vibrations become nonlinear. That's the basic physics of what's going on here. So this is what's called the Kerr effect.

There is actually another effect which George hasn't mentioned here. There's another-- you see this is actually E cubed here. But you can have an effect where this is E squared, which is called the-- what's it called? The Pockels effect. I was really only asking because I'd forgotten the word for a minute. But my students here know that I'm always forgetting names, and they usually come back eventually.

Did you see what that did? Oh, yeah, it did this, the relationship between the susceptibility and the refractive index. And this is saying that the phase delay depends on the polarization. And in this one, the index of refraction depends on the intensity. So this is nonlinear optics. And as I say, those last two effects very often come together.

So this is now talking about polarization of light. I think that probably George must have mentioned this to some degree earlier on. But what this is showing, then, is a plane wave of light. It's part linearly polarized in the x direction-- sorry, the y direction. This is taken as y. So the electric field is in the y direction, which actually means that the magnetic field is going to be in the x direction because they're at right angles with each other, if you remember.

And you remember that you can think of the wave propagating in two different ways. One is to think of it other at a certain instant of time, so what we call here it's a snapshot, and then we plot what happens as a function of z. So of course, we have to do this because it's difficult to plot it against space and time at the same time because it's difficult to visualize what's going on.

So you can see here this electric field is always in the same plane, and it's a maximum here. It reduces, it becomes 0. It then increases in the negative direction, comes back to 0 again, and so on and so on. So this is linearly polarized light as a function of space.

And then this one shows the same thing as a function of time. So we're looking at now one particular value of positions that equals 0, and we're looking at how this varies as a function of time. Exactly the same diagram, of course. You probably saw when I flicked to the next slide, it didn't change at all. It just changed the caption.

And this is what circular polarization does. Circular polarization, you can see what this is trying to-- it looks a bit complicated here. But what it's really doing is actually quite straightforward. It's actually rotating. This electric vector is rotating either with distance or with time, according to which way you do the plot. And the length of this line stays the same. So this is what we mean by circular polarization.

And this is showing how, actually, you can think of circular polarization as being made up of two components-- one linearly polarized in this side in the y direction, one linearly polarized in the x direction. And so you can see that there are also out of phase by 90 degrees.

So you can see now that in this distance, position, this is plotted against z again. Here this one is a maximum, this one is 0. So the electric field is in the x direction. And then as you go, as you travel along, this gets smaller, but this gets bigger. So this position here, the electric field is now polarized in the y direction. And then this one goes back down to 0 again, and this one gets bigger. So this is exactly equivalent to the case that I showed before, the circular polarized case.

I might add that electrical polarization is where the amplitude of this component and this component are not equal. So if that was the case, it would-- again, you that you'd think of this electric field vector as rotating, but the length of the vector would also change as it rotates. And of course, you can have two cases. You can have it rotating clockwise, or you can have it rotating anti-clockwise. So these are called left- and right-hand circular polarization.

So this is showing an end-on view. So this is x and y. So this thing is pointing out towards you, isn't it, that means. No, it means it's going away, isn't it? But is it right? Because it's left-handed coordinate system. If that's the case because x-- oh, no, it's not. No, it's right. It's right. It's the right-handed coordinate system. Yeah, x cross y goes away from you.

And so this is showing, then, end-on, what happens to this electric vector. And the wonders of modern science have in fact changed the circular polarization to electrical polarization without me doing anything.

Now, so next is to introduce some-- sorry. I missed that one. There were lots of these things that my computer misinterpreted from George's computer. And I changed most of them, but I missed one there. That should say left lambda over 4 plate. I don't know what the difference is what he put in there. Do you know what he's put in there to make it go like that? Anyway, he did this on his Mac.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah. Yeah, but it's a funny font as well. Anyway, so this is showing what would happen if you put some birefringent-- a special birefringent material here, which is what's called a lambda by 4 plate. And what this does is it changes the relative phase of the components of electric field in two directions by a different amount as it transmits through the crystal. And so if you arrange the orientation of it correctly relative to the light, it has this effect of changing the linear polarization into circular polarization.

The way it really does that is that this-- you would say that this you can think of as being made up of two linearly polarized components, and this you can always think of as being made up of two linearly polarized components as well at 45 degrees to that. So you just resolve that into two components. So the circular polarized, you get that very simply by just changing the relative phase of those two linear polarized components. That's how it works, but I think here he's just trying to show you what it does.

And then the other important device that people use is a is a lambda by 2 plate rather than a lambda by 4 plate, and that has the property of changing x polarized, linearly polarized light into y polarized light. And again, you can see how this is working by thinking in terms of resolving this in so two components at 45 degrees to it, and then changing the relative sign. Maybe I should do that.

This doesn't [INAUDIBLE]. Is it switched on?

AUDIENCE: [INAUDIBLE]

PROFESSOR: So maybe I can just do that quickly. So what I'm saying is that this, you can think of it--

AUDIENCE: [INAUDIBLE]

PROFESSOR: This one you can think of as being resolved into these two components. And then the crystal is then changing the relative phase of one of these relative to that, making that. And so now the polarization is like this. So that's how it's working.

For the lambda by 4 plate, what it was doing was it was changing the relative phase of this and this by 90 degrees, not by 180 degrees, which would give you a circularly polarized light. But in order to get this effect, you have to get the right orientation of the polarization relative to the crystal axes. So that part he hasn't actually said anything about here.

The other thing-- maybe you'd be interested to know this. Of course, if you put linearly polarized light in and get circularly polarized light out, what would happen if you went the other way? Well, in fact, it does work the other way. You can put it in circularly polarized light and get out linearly polarized light. So you can do either of those two things.

The other thing that you might worry about is it's interesting that if we change the phase of this light, then it's going to change the phase of that light, isn't it? And yeah, maybe that's not leading anywhere. I won't say any more about.

OK, think about this. Yeah. Oh, well, there we are. I've already-- I've jumped the gun a bit. So he's saying, what happens here then? So we've got incoming linearly polarized light, the lambda by 4 plate. What does the lambda by 4 plate do? Circular polarized.

So we've got circular polarized light on here. It strikes a mirror and it reflects. What do we get after it's reflected? That's a bit hard now, isn't it? Opposite circular, right.

And then this gets back to here and it goes back through this birefringent plate. So what does it do?

AUDIENCE: Same.

PROFESSOR: Same polarization. Is that right? It rotates by 180 degrees. This is slower, actually. It's gone through a lambda by 2 plate then, isn't it? Well, I mean--

AUDIENCE: [INAUDIBLE]

PROFESSOR: Oh, yeah. Yeah, sorry. Yeah, you're quite right. Yeah. Yeah. OK everyone, about that? I guess the next thing is to say what would happen if you change that mirror by a phase conjugate mirror. But maybe that's too difficult for us to think about.

And then remember that-- I think we had this before as well-- that when you have lights interfering, the electric field, it's only the parallel components of electric field that interfere. So if you've got two waves where the electric fields are at right angles to each other, then they don't interfere.

So this is showing this case here. We've got two plane polarized waves. They're polarized at right angles to each other, and so they won't interfere. And therefore, if you look at the trying to get an interference pattern, you won't see anything in this case. And this is showing them in the same plane. And then they would interfere , and you would get some interference pattern with maxima a minima like this.

Another interesting one I guess I might draw over here is what happens if you-- this is our screen. And let's say we've got two waves coming in. One, let's say, is polarized like this, and the other one is coming in here and it's polarized like this. You see? So we've got these two waves coming in.

And you can see that their electric field vectors are at right angles to each other, and so this is another example where they're not going to interfere. So you wouldn't see any interference fringes in this case. So this corresponds to what's called p-polarization.

So if we did this the opposite, if we did it with the electric field in this direction, and the other one's coming in also with the actually field in this direction, then you can see these are now in parallel directions. and so you would see fringes where these two plane waves interfere. So I've drawn this for the particular case. This is supposed to be 90 degrees.

Right, OK. That's all George's material. That was the end of George. It didn't end with a particularly startling thing. So that's really the end of the course as George wrote it.

But he did ask me, as I say, to do a bit more. So what I think I'll do, why don't we have a break now for five minutes? And then I'll start again in five minutes' time. So this is-- now it's going to be talking about focusing of light.

And as I said before, this material I'm going to do now is just for interest. It's not going to be examinable. And if you're not interested, you can go and have a cup of coffee instead. I don't really mind. But some people might be interested in learning a bit more about focusing of light.

So I was going to say a bit more about focusing of light, and in particular some of the funny things that happen when you look at focusing of light with a microscope objective, so the tight focusing of light. And this is actually a picture of the intensity in the focal region of a lens, not high numerical aperture in this case. This is according to paraxial optics.

This is a picture taken from [INAUDIBLE] Wolf. And it's similar but not identical to the one that George showed on the last lecture or the lecture before.

So this is the focal plane. This is the optic axis. And of course, what you're really seeing, this is actually going to be radially symmetric. So if you imagine this was being revolved around the axis, you'd have the shape of this complete surface. And you can see what happens is so this is the focal plane where you got the Airy disk. And then as you go out of focus, then you start getting this strange sort of almost like a sort of butterfly shape.

So this calculation that's given in [INAUDIBLE] Wolf, it first of all, as I've said, it's using this small angle approximation. You assume that sine theta is the same as theta. And it also makes an assumption, which is called the Debye approximation. So I'll say a bit more about the Debye approximation.

But before I carry on, I'll say another thing. I'd forgotten I'd put this in. And that is that there was a really neat paper published way back in 1964 by McCutcheon, where he showed that you can postulate the idea, the concept of what he calls the 3D pupil. And this is the cap of a sphere where it's cut off at this angle alpha, the angle of the lens.

And if you do love the 3D Fourier transform of this-- ta-da-- you get back to that. So isn't that neat? So I guess, if you think a bit more about it, it's probably-- I think you can satisfy yourself that perhaps it's right. This is in k space, effectively, then.

And what it's really saying is that because the light satisfies the wave equation, the modulus of k is constant, and therefore the locus of this is a sphere of radius k, so very simple. So light propagating-- this convergent light propagating will have a k vector which lies on the surface of this sphere.

So I just added that just because I thought people might find it a bit intriguing. You can sort of generalize a lot of imaging theory into 3D using this same sort of approach. But I'm not going to say anything more about that now.

What I'm going to do now is talk about-- so this is this Debye approximation again. You remember, we said there's two ways of calculating to two different optical systems, which give virtually the same or similar sort of results. This one I'll look at first of all.

This is where you got the aperture stop in the front focal plane of the lens. This is a distance f. This is a distance f. And then in this plane here, you will get the Airy disk. So the pattern I described a couple of slides ago for the field is in 3D space would be in this region here. So in this plane here, you get Fraunhofer diffraction, and then in the regions around it you get Fresnel diffraction.

But you can get quite similar results, if you remember, by using this sort of system. So this is where the aperture stop is not actually in the front focal plane of the lens, but it actually in the plane of the lens. Or you could think another way of doing this experiment would be to get a convergent beam and place a circular aperture in that convergent beam. So the question is, what do you get in the focal region in that case?

And it turns out that if you satisfy this condition, that a squared over lambda z is very large, where is a is this radius here-- this is called the Fresnel number. And if that quantity is very large, it means that what you'll get in this focal region is very similar to this. Actually, there's a parabolic phase term that we described when we did the Fourier optics that makes it different. But the intensity is going to be the same as for this case here.

So these are two cases where you can get this Debye approximation. So in this case, it's exact for this particular geometry. In this case, it's an approximation which is obeyed if this Fresnel number is very large.

And then, not that long ago, people realized that this was an approximation, an assumption that was made. And there was a very nice paper published by Li and Wolf, so this paper here-- 1984, quite a long time ago now, but not that long, but anyway, as history goes, Li and Wolf. And they showed what would happen if this Fresnel number was not large but was of some intermediate value. In this case, N equals 10.

So you see what happens is that this pattern has got distorted. You can see that it's tending to sort of fan out. And the physical explanation of that is that, effectively, the light is being distracted by the aperture and spreading because of diffraction effects. And that is superimposed on the focusing effect. So the lens is trying to focus the light down, and the diffraction is tending to focus the light out. And so you get a sort of balance between these two things.

So in this focal plane here, you get the same intensity as you would if the Debye approximation was valid. But in the region elsewhere, it's a bit different. It's scaled according to some sort of affine transformation, I think it probably would be. What do you think? What do you think, [INAUDIBLE]?

But one very interesting thing that happens is that it turns out-- you can't really see it very well from this diagram, but the maximum intensity is no longer in the focal plane, no longer in the geometrical focal plane. It's shifted towards the aperture. And that is called the focus shift effect, and there's been a lot of papers that have been concerned with that and calculating what happens for lots of different cases. So that's the Debye approximation. That shows what happens if it isn't true.

But now I'm going to go back. And all the rest I'm going to carry on now is talking about what happens if the paraxial approximation breaks down, but the Debye approximation is valid. The case when the Debye approximation is not valid and it's a high numerical aperture is very complicated, and there have been very few papers that I've treated that case.

Anyway, tight focusing of light. We get light, we put it into a microscope objective. The light then-- a microscope objective, a dry microscope objective might have a numerical aperture of 0.95, which means that the light is converging something maybe over 70 degrees relative to the axis, so very-- you can't really think of 70 degrees as being a small angle, I guess. So the normal paraxial approximations break down.

And this is some sort of areas where we might want to do this. The one, of course, that I'm most interested in, microscopy, but also laser micromachining, optical data storage, optical lithography, laser trapping and cooling, looking at the physics of light-atom interactions, and cavity quantum electrodynamics. So there's a lot of basic physics that's being done nowadays that uses these sort of principles, too.

So this is how we model the light. Again, we've got-- just like before, we think of the aperture as being placed in the front focal plane of the lens. Of course, the lens is very complicated. It's got loads of pieces of glass in. But we think of it just as like a black box, and we describe it just by this surface here, which is called the equivalent refractive locus. And what that is, is the locus of points where this ray and this ray intersect.

And actually, if you've got a well-designed system that is satisfies what's called the sine condition, this surface will be a sphere. But I'm not going to go into the details of that. I'm just going to assume it's a sphere in what comes next.

So in this case, you've effectively got lots of plane waves incident in the focus. But unlike the paraxial case, we really have to take into account now the polarization of these electric fields and add together the electric fields in the focal region.

And so this is what happens after the light's gone through the lens. It starts off as being plane polarized. So in my diagram here the red represents the electric field, say. So we start off with the electric field being plane polarized in the vertical direction. And then after it's gone through the lens, the rays are focused on the focal point, which is at the center of this sphere.

But if you actually calculate what happens to the electric field vector, this is what it does. After you've focused it, the electric field lies on the surface of the sphere with this geometry that-- you see that these lines all go through a point on the far side of the sphere. And the same is true for the magnetic field as well. If you calculate what happens, this is what you find happens.

And you can show that this is actually exactly the same field as you would get if you placed at the center an electric dipole along this axis and a magnetic dipole along the y-axis. So if you add together those fields, you would get this.

And one more thing to say is, you can see that if you've got a small aperture system, low NA, you're in this region here. And you can see now it's becoming very close to being linearly polarized again. It's only when you get to big angles that it starts becoming a bit more complicated.

So Richards and Wolf in 1959-- so this is the same Emil Wolf who wrote the book. And in their paper-- I haven't got the reference here, I'm sorry. But anyone remember what it was published in? I've forgotten. No, it doesn't matter. *Proceedings of the Physical Society,* I think, maybe. But anyway, it doesn't matter. You can easily find it on the web if you're after it.

They showed-- you found it?

AUDIENCE: [INAUDIBLE]

PROFESSOR: OK, *Proceedings of the Royal Society of London* in 1959. They showed that the field in the focal region could be written in this form. So rather than just have one diffraction integral, we now end up with three for the three-- defining these three integrals. And then these are the electric fields in the focal region, written in terms of these integrals-- i0, i1, and i2. And these things are very similar to the paraxial things. The normal paraxial one is a bit similar to this, with the j0.

And this is an extra bit which comes about because of the system satisfying the sine condition. And I'll draw it for you, what's happening. Basically, if this is this aplanatic sphere surface, if you put in light, plane polarized light, into this system, you can see, if this is uniform intensity when it goes in, you can see that the amount of energy that's gone into here and the amount of energy that's going to go into this angle, which is very much bigger, is going to be the same. So you see what I'm saying? So this amount of energy is going to be squeezed into this angle, whereas this amount of energy is going to be spread out over this angle. So this produces this factor here.

And then you can calculate what happens. So this is-- yeah, I'm sorry about that. You see, you write things like that not so ti can be symmetric. But of course, what I mean is that this is not circularly symmetric. This really should be a circle before the modern technology distorted it. And so rather than the normal Airy disk, what you get is this elliptical focal spot like this.

And this is another calculation. These are both from a couple of my papers. This one is a really old paper, 1977, really old. And this one was from a bit more recent, 1997. So this was written with Peter Torok. And we came up here with a new way of calculating it where, basically, we look at the multimodal expansion of the field over the sphere and express it as a sum of these spherical harmonics, the modes of the sphere. And so this becomes a very actually quite computationally efficient way of calculating things.

And this is for two systems, then. This is for an angle of pi by 3, so 60 degrees, and this is for 90 degrees. So this is if you're actually the limiting case, when you're actually focusing a whole hemisphere of light onto the focus. And you can see that the sort of general shape stays much the same, but it changes a bit as you change the aperture.

The other thing I was going to say a bit about is Bessel beams. I don't think George did anything about Bessel beams in the end, but I did see it as a title on some on the contents of some of the lectures, as though it was just about to come. But maybe I just missed it.

But anyway, this is a concept that is quite interesting at the moment. Basically, there are three ways I'm showing here of doing it. One is what happens if you put in your-- this is the aperture stop of this lens, and we've put this central obstruction so that it becomes just that annular region. And we make the width of this annulus very, very narrow.

Then, you remember, just in the last lecture, we looked at what happens for this annular lens. And we know that what happens is, compared with the Airy disk, the central lobe gets narrower, but the side lobes get bigger. Well, it turns out that in the limiting cases, this becomes very narrow, you get the Airy disk becomes, given by this very simple expression, just j0 squared, where j0 is a Bessel function. So that's why it's called a Bessel beam.

And it has this very nice property that actually this is it this satisfies the wave equation. So it's like a sort of mode of free space. So this Bessel beam propagates with an infinite depth of focus. It just propagates without spreading, and so that's why these beams have become of quite a lot of interest.

Actually, the concept is really quite old. There was a device proposed by MacLeod in 1954 for producing something rather similar, which is called the Axicon. An Axicon is like a conical piece of glass. It's like a conical prism. And you can see, then, that a prism refracts the light parallel, like this. And so you can see wherever you are along this axis you can see that what you see is light that's coming at the same angle, and so you can see from this how you get this increased depth of focus of this system.

And another system was proposed, again, a long time ago, by Dyson, to do this not with a prism but with a diffractive optical structure. So you make a sort of grating, which is a bit like a zone plate except the zone plate-- the rings on a zone plate go as R squared, whereas-- sorry, the width of them, the spacing goes as 1 over R squared, whereas this one, the rings are linearly spaced. And in this case, because the grating is linearly spaced, all these rays go at the same angle, and you get the same effect as this. So that's how you make Bessel beams.

So I've been interested in them for a very long time. And this is from another paper of mine going back to 1978, very long time ago. And this here was one of the lines of that paper, as it says. "A wave with zero-order Bessel function radial distribution propagates without change." And so what this is showing is what happens if you, instead of thinking of the paraxial theory, what happens to this Bessel beam if it's not a paraxial case, if you put, say, this annular mask in front of a high numerical aperture objective lens.

And this is what you get. So this is for different angles-- 30, 60, 90, as before. I've even gone further here. I've gone to 120. And you might wonder what that could possibly mean, except it means you can see that you could do it by using a mirror, the parabolic mirror that George spoke about spoke about earlier. So you can have light coming in this way, say, a narrow annular beam coming in here, so this might correspond, say, to this case. And then the 120 case would be when the light's coming in like this. So you can actually get a bigger than 90 degrees by using a mirror sort of structure.

But anyway, what you see is what happens, if the angle of the focusing is small, then it's pretty similar to what I showed earlier for the full lens, full aperture. But as you make this angle bigger and bigger, eventually very nasty things start happening. This thing starts splitting into two spots rather than one, and you can see that it does more and more of that as you make it even bigger.

And so this is all assuming that we're putting plane polarized light into the lens. So the conclusion we come so is that putting plane polarized light into that lens is not really a good idea, not good news.

And so this is explaining things a bit more. First of all, we look at the black curves. So these correspond, then-- so a normal Airy disk is this one, and this other one here corresponds to the narrow annulus. So you can see this is the Bessel beam. You can see it sharpened up quite considerably but with some rather larger side lobes.

But then if you go to the high NA case-- so this is calculating it for a numerical aperture of 1.4, so for an [INAUDIBLE] immersion objective. And you'll find that for the annulus now, it's actually no narrower than for the full lens. So the polarization effects are really basically getting rid of any advantage in this sharpening of the spot that we get in this case.

But it turns out there is a way of getting round that, and that is to not use plane polarized light going in. And what we use instead is putting radially polarized light. If we put in radially polarized light-- this corresponds to this magenta color, so as it says here, radial. I've also called it here TM0 because it corresponds to a transverse magnetic mode. I'll show you how that comes about in a minute.

And now you can see for that case we're back again to getting something very close to what you would get for the paraxial case. So there's been a number of papers which have exploited using radially polarized input into your lens in order to get a smaller focal spot. So yeah, this is how it comes about, why it's called transverse magnetic.

So this is radially polarized light. Sorry, yeah, the red is the electric field. So that's in the radial direction. If we had a small NA, it would just be still radially polarized. But you can see that if you look on the complete sphere, it starts becoming like lines of longitude-- latitude on the sphere. And you see the magnetic field are these lines, which are purely transverse, so that's why we call it transverse magnetic. And the 0 corresponds to the fact it's the lowest order mode that's like that.

There is another case which is azimuthal polarization. This is where the electric field goes like this and the magnetic field is like this. So this corresponds to TE, but I'm not going to say anything more about that because that doesn't produce anything useful.

And this is-- actually, I said there were three ways of making these Bessel beams. There's actually a fourth way, which some people together with me at the Data Storage Institute here in Singapore have been working on, and this is to use a lens with some sort of binary optic structure. And actually, this was from a paper where we were simulating what happens if you put radially polarized light into this sort of system and you optimize this binary mask in order to get close to a Bessel beam. And these are the sort of results you get. This is along the axis you can see that you get something which really propagates without spreading over quite a long distance.

This is showing how there are various ways of specifying how the performance of a focusing system. And one way is to calculate or measure the electric field at the focal point for a given amount of power going in. And what you want, of course, is for the power of the focus to be-- the electric field at the focus to be as big as possible.

And so this ratio is this thing called F here. And what this is showing is that this one here is actually for our plane polarized light going in and showing what happens to this ratio as you increase the aperture of the lens. And so you can see that as you increase the aperture of the lens, so, of course, it focuses better and better. And therefore for the same amount of power get getting going in, you get a higher field at the focus. So that's why it rises like this.

But it, you see, only gets to eventually-- the maximum it ever gets to is a half. And the reason it gets to a half is because-- I showed that this polarization is actually equivalent to an electric dipole plus a magnetic dipole. And so only half of the power is going into the electric dipole, and it's the electric dipole that gives an electric field at the center. The magnetic dipole gives a magnetic field at the center, which we don't want. So only half of the power is going into the right thing. So that's why that does that.

And then you can see there another two curves here. This one is our radially polarized light. So you can see that it's very slow to get started here, and that's because with the radial polarized distribution, you've got very little power near the axis anyway. So it eventually starts taking off, and you see that at some angle here, which is really quite big-- so this is 90 degrees down here somewhere, so this is going to be maybe 80 degrees or something like that-- it overtakes the plane polarized one and becomes better.

But you can see there's another one, another curve here, which is for just an electric dipole on its own, so a transverse electric dipole. So rather than try and give this field, we just might try to get this polarization. And if we do that, then we see we do even better. So in this region here, you can see this is doing better than either of these other two.

So this is what we call an electric dipole wave. And this is what this looks like, then. This mixed dipole, which is the electric plus a magnetic dipole, is equal to an electric dipole plus a magnetic dipole. So this is the electric dipole field. So what we need to do is to make it so that after we focus the light, the electric field of the focused light is on the surface along these lines, like this, basically.

You see that if you look in the region near the axis, all these behave the same. They're all we're all going to be virtually linearly polarized. It's only when you get so big apertures that these become different.

And so this is now looking at what happens with Bessel beams. So this the so-called mixed dipole, which is the plain polarization case, for 30 degrees, 60 degrees, 90 degrees 150 degrees. This is for the electric dipole case, 30 degrees, 60 degrees. Now, you see here this is a bit better than this one. And this one, this spot is divided into two, but this one hasn't divided into two. But it has got these rather large side lobes developed here.

And then if you look at this other curve, these other ones here, this is what I call a T1. So this is another mode I haven't mentioned before. It's another TE mode, but it's the TE1 rather than the TE0. And you can see this one-- actually, ti's the same in all these diagrams. It's independent of the angle now. So you can see, actually, this one turns out to be slightly smaller than this one even.

And this shows what the TE1 looks like. This is TE1. So again, the electric field is transverse. But it's not going all the way around like TE0. It's going this way on this side and this way on this side. So if you could make this polarization, this would produce that results I've just described there. So this shows all these different types.

So what do you have to put into the lens to get that? So this is what you have to put into the lens. So if you put it in plane polarized light, you get this what I call the mixed dipole case. In order to get electric dipole polarization out, you have to put in light that looks like this.

For the magnetic dipole, it has to look like this. See, here it's going-- it's pointing in, and here it's pointing out. And for the TE1, it's like azimuthal. And for the TM1, it's radial, like this. So by altering the polarization of the light that goes into the lens, you can generate all these different mode patterns.

And this is what it does. This is calculating the area of the focal spot as a function of angle for all of these different cases. So here there's a whole range of them for a full lens, full circular aperture. And this is for a few different cases for the Bessel beam. And you can see what happens is we want this area, of course, to be as small as possible.

And you can see that actually in this region here, the electric dipole one is the smallest, and then the TE1 is the smallest, and then eventually the radial one becomes the smallest. But this doesn't become bigger until you get over a numerical aperture of 0.89.

And then this is for the full circular case. I've got lots of different cases here. This one corresponds to a paraboloid mirror. This one corresponds to basically a different apertization of the light, which you can see it actually is doing better here than the others, but eventually runs out of steam when you get to big angles. And you can see here, the TE case is the smallest, and then eventually the radial case is the smallest after you get above a numerical aperture 0.91.

Now, the other thing you'll notice, though, about all those spots I was showing is they're all not symmetrical. So what can we do to make them symmetrical? Well, this fits in a bit with what we were talking about before, the polarization case, the circularly polarized. The TE and the TM0 are already circularly symmetric, rotationally symmetric. This one corresponds to our normal radial polarization.

But these are the ones, we know that x polarized plus iy polarized equals circular polarized. So we can do similar sorts of things with these TEs or the electric dipole polarizations. We can add we can add one of them along the x with one of them along the y with an i there, and that actually is now going to give something which is rotationally symmetric. And it actually, if you think a lot more about it, it corresponds to azimuthal polarization with a phase singularity, with a vortex, a phase vortex superimposed on it.

And this one corresponds to some elliptical polarization with a phase singularity. So what I mean by this? This one, azimuthal polarization with a phase singularity, what I mean by that is at anywhere around here, it's polarized in the azimuthal direction. But as you go round here, the phase changes from 0 right around to 2 pi. [INAUDIBLE]? No? OK.

So it changes. As you go round here, it changes from 0 to 2 pi. So where it gets back to here, it matches up again. But here it's actually changed by pi, so it's effectively in the opposite direction. So that's what I mean by that. There's a continuous phase variation as you go round there.

And on this one, elliptical polarization with a phase singularity, what it means is that the light is-- anywhere on here, it's elliptically polarized, like this, say, but at any angle it will be a similar sort of ellipse. But the phase of this ellipse will change with a face singularity as you go around in the circle.

And so what happens if you do that? Well, this is what you get. So this is looking at the cross section for three different angles for all these different types. And so this is for Bessel beams again. They're the simplest to calculate.

So what we find is that for small angles, the radial polarized is no good. It turns out that the reason why radial polarization works, actually, is because-- sorry. I should have said this earlier. I don't know why I didn't, really.

Why it works is because the light that comes down here has got some polarization in this direction here, which means it's got a component in this longitudinal direction. So actually, the radially polarized light put into the lens produces a longitudinal electric field in the focal region. And that might seem very strange to people, first of all. You can actually generate a wave that's moving along which has got a longitudinal electric field. But that seems to be what is true.

But the point is, of course, if this angle is small, then this longitudinal component is going to be very weak. And that is why this racially polarized one here doesn't work well, because the transverse field is dominating over the longitudinal field. But as you increase the aperture, you see it gets better, like that.

And then I also show here the electric dipole. All these are all the same on this, for this small angle, but they for the larger anguish they become different. And this is showing how-- sorry. This arrow has moved itself for some reason.

I don't know why it's moved itself, but that this was supposed to be pointing to this blue line here and showing how, in this case, the TE1 case is the narrowest. It's actually even narrower than the radial in this case here, but with some rather slightly-- some bigger side lobes here. And this shows also the electric dipole case as well. And this one's showing how the plane polarized case is really very bad, as we showed earlier, once you get to really big angles.

And then this is comparing them a bit more. This is actually plotting the normalized width of that rotationally symmetric spot that you get, again, for all these different cases, and showing-- so these are the cases of a full aperture. These are the cases of a Bessel beam. And you can see that, actually, it turns out that the TE annulus is the narrowest, the TE Bessel beam, all the time, for any angle. It's always narrower than the radially polarized one. And for these cases here, for the full circular case, the TE one is always narrower than the green one until you get to this point here, which corresponds to a numerical aperture of 0.98, which is higher than you would ever get to in a practical system.

So what this is all showing is that, as I say, putting plane polarized light in is not the optimum. It turns out that very often the optimum can be different according to the circumstances, but it can be this electric dipole polarization or it can be the radially polarized input or it can be the transverse electric. And I think I-- yeah, this is showing the strength of the side lobes and showing how the transverse electric might be narrower, but it might have bigger side lobes. So here the radially polarized has smaller side lobes, weaker side lobes, than the TE.

So and there's our conclusions about this, then. So if you focus plane polarized light, you'll get a big focal spot. So you can improve that by using radially polarized light, and that comes about because you've got this strong longitudinal field on the axis. But we find that electric dipole polarization actually gives a higher electric energy density at the focus, as I showed earlier, higher than you get with radial polarized light.

And then I showed that the TE1 mode polarization gives the smallest central lobe, and it's smaller even than you get with radially polarized light. But the TE on its own is asymmetric, but you can make a symmetric version of it, which is equivalent so azimuthal polarization with a phase singularity.

So that's about it for what I was going to say about this. So this is finishing up the course with something which is really sort of ongoing research and something that, I think, there's a lot of people around the world working on these different problems and trying to optimize these systems to get the best possible focusing properties that you can get with a high NA lens. And I mentioned before all the many different applications that there's this could be used for.

So I'm going to stop, then. Questions? Any questions about this? Any questions about earlier? There's lots of questions about this are there? Yeah, maybe.

There's a lot of information there, I know. I think you're probably all suffering from overload. But it's all very complicated, actually, but very beautiful as well, I think. It's nice physics, I think.