[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: OK, so you'll remember, in the last lecture, we discussed the technique called schlieren. And schlieren, we were relying on the fact that the plasma has refractive index gradients inside it which cause a deflection, and that deflection angle of the rays going through the plasma is going to be something like 1 over 2 times the critical density times the integral of the gradient of the electron density along the path of r probing b. And what we did with schlieren is we filtered the rays using a lens and a stop at the focal plane, and we we're able to filter out rays with certain angles.

We could let through undeflected rays. We could block undeflected rays. We could do it in the isotropic fashion. We could do it with a knife edge. And this enabled us to image effectively the gradients in the electron density inside our plasma, and this is particularly useful if you've got something with sharp density gradients such as shocks.

But when we were doing this, we actually made a pretty big assumption that I kind of buried here. We assumed that we have no displacement of the ray, that the rays pick up some angle as they travel through our plasma. So I'll draw a little plasma here.

We assume that the rays picked up some sort of angle but they didn't themselves get displaced. Whereas, for some realistic extended plasma, they would be displaced inside the plasma as well as picking up some angle. And so the rays themselves would not only have an angle, they would have some displacement, and when we use our imaging system with a lens, the rays should end up somewhere else on our object plane, not in the same place that they started with.

For schlieren, we assumed that there was no ray displacement. And this is actually equivalent to assuming that our object was very, very thin. So if instead of having this extended object like this, we have an object that's incredibly thin like that, and you have rays of light coming in that then just get instantaneously deflected like that-- this works out OK because there's no sort of path along inside the plasma for these rays to become displaced. Whereas, they are in an extended object.

So we made this assumption, but we're now going to relax that and have a look at what happens when we do true shadowgraphy, which is where we're interested not just in the angles of the rays but their displacements as well. So I'm going to try and draw a slightly better version of this image, hopefully with some nice, straight lines. Let's see how it goes.

Here's my plasma again. I'm going to have a lens here. I'm going to set up my lens such that it brings into focus rays from this object plane onto some image plane back here.

And what I'm going to have is the undeflected rays. So imagine there was no plasma here. The rays would just go through here. And they would all go through a focal point back here.

So this again looks like the setup we had for our schlieren, but the difference is now if we consider the rays being deflected inside the plasma, they're going to pick up some sort of displacement. So for example, this ray is going to be displaced down like this. This ray maybe is also going to be displaced down a little bit, and this ray is going to be displaced upwards such that when they get to the object plane, they now have some distance that they have been translated compared to the undeflected one.

And then, all that the lens does is it puts light back on the image plane where it came from on the object plane, with some inversion, other things like that. But that means because these rays have actually been moved, because the light has been moved around, they're going to show up in different places on the image plane. And this is where I really hope I can draw this properly, because this is quite a complicated diagram, so give me a moment to focus on it.

So let's start with this ray here. It's going to go down. It's going to go down below the focal point because it's being deflected downwards. And so it's going to end up here, so slightly above where the original ray was going to go.

This ray here is going to get deflected downwards, and then the lens is going to push it up. And it's also going to go below the focal point here as well, so it's going to be deflected up from its original position. And this one here is going to go up.

I'm going to bend that one slightly. Oh, no, I didn't want it to go over there. Ah, I got it right the first time. There we go. And this one is deflected downwards because it was deflected off of the object here.

So what we're seeing is that this lens is now imaging the light into slightly different places than we had before, so we have the displacement of the rays. And you might want to think of that displacement as an intensity modulation, because, of course, there aren't just three rays. There are a large number of rays, and these all make up an image.

And so initially, maybe you had a nice uniform pattern. These yellow rays would have made a nice uniform circular beam like this. And now your blue rays are going to be distorted, and there's going to be regions of very bright intensity and regions of lesser intensity and regions that are about the same. So we'll have all sorts of interesting features inside this.

And one way to think about this, imagine we took our beam-- and now we're looking at the laser beam. It's coming directly towards us through the plasma. Imagine we arranged it so we have nine little dots. Maybe we've got nine little laser pointers pointing through our plasma like this.

Now, as these rays come through the plasma and they're deflected onto our camera here, we would end up with the dots in different places. So we could have this one is moved up here, this one is moved like this, this one is moved like this, and so on. And you can just think there's a series of displacements for these dots, and if you're able to measure all these displacements, perhaps you could work out something about the plasma itself. And we'll talk a little bit about some more quantitative measurements for doing this, but this is just qualitatively the picture here.

Another way to think about this, if you prefer-- these are just different conceptual ways of trying to think about the same thing-- imagine we initially put through this sort of tic-tac-toe grid of light here. We've sort of blocked it off so we've got two vertical lines, two horizontal lines. These lines, as they go through the plasma, the light will be deflected by different amounts.

And so for example, this line here could be deflected like. This one could be deflected like this. And so you sort of see we have some sort of distortion to the grid that we're getting out of this, and our initial grid with all these nice, straight, parallel lines is now a distorted grid instead.

And that's very important because shadowgraphy does not actually produce an image. I've called this the image plane here, but we don't actually get an image out, because we have something which is significantly distorted and it no longer preserves the things that we like to preserve in images. We take a picture with a camera, we'd like straight lines to remain straight.

Here, the straight lines do not remain straight. They don't remain parallel. We don't preserve length, anything like that. But there's still clearly information stored inside these images, pictures, that correspond to the density inside the plasma, and we'll talk a little bit more about how to do this quantitatively. So any questions on this? Yeah?

AUDIENCE: So if the plasma here is acting like a lens, what do we need the second lens for? Just as a reference?

JACK HARE: Absolutely, yeah, great question. Yeah, so I put the second lens in here for a couple of reasons. First of all, I find it conceptually simpler because you can see the deflection angles at the focal point here, because when I read off the position of the rays at this focal plane, I can tell whether they've been deflected up or down, which is slightly harder to do here.

You could absolutely put your detector right here, OK, so you could put a bit of film or something out here. In practice, you can't, because it's right next to the plasma. So you need some sort of lens optic.

This technique actually looks a great deal like proton radiography, which we'll talk about later on here. In proton radiography, you don't have any lenses. You can't-- it's very difficult to make lenses for charged particles. And so you do, in fact, put your bit of paper, like your film in this case, very close to your plasma.

You might put it a little bit further back, and we'll talk in a moment about actually how important is your choice of exactly where you put your detector, because I could put my detector here or here or here, and I will still get an image forming, but the image will look different. But yeah, we do not need the lens to do shadowgraphy, but in almost all realistic setups where we're using some probing beam through a plasma, we're going to have a lens that allows us to put our detector nice and far away and safe from the plasma. So, yeah.

AUDIENCE: All right. Thank you.

JACK HARE: Cool. Other questions? Anyone from Columbia? I don't see anyone, so I'm going to keep going. Cool. So let me erase this side.

So as I just mentioned, then, we can actually change the shadowgraphic image we get by moving the position of our object plane. So that could be moving our detector, if we've just placed a detector here, or it could be by moving our lens, because as we move the lens, we change the place we're focusing on. Or we could move other-- anyway, you get the idea. There's lots of different ways of moving this to some other place.

So let's have a look here. We've got, again, our plasma. And we got our rays coming through it.

We've got a ray that gets deflected downwards and a ray that gets deflected upwards and another one that gets deflected down like this. I can draw some different planes in which we can put our detector. Let me just get them in the right places.

Or maybe I can call these 1, 2, 3, and 4, like that. And so then we can have a look at, What would we see for this very simple system where we've got some sort of focusing density gradients and some defocusing density gradients here? So this would be a maxima in electron density, and this would be a minima here, right, so that we would get focusing and defocusing. So let's draw what those patterns would look like.

So if I have intensity like this-- and let's say that this is the y-coordinate, so I'll call this y here, and I'll do the first one here-- we'll get something-- we're going to have a little bit of deficit of intensity around about here because the ray's being deflected away. So if this is the background intensity of our probing beam, we'll have some sort of drop, and then we'll have an increase because the light is being focused together in the middle here. We'll have another deficit in this region, and we'll have another little increase over there.

So we get some sort of nice little modulation, what it's looking like. If we go to 2, we can see that the rays are now getting closer together and further apart, so you'd expect for this pattern to be even more exaggerated. Maybe I should've drawn this less exaggerated to give myself more space. But let's say it looks like this, and this is bigger.

If we get over to 3 now, we see something interesting. We actually have rays crossing here, so this means that a lot of light is all being piled into the same place. So we're still going to have a little bit of defocusing, but we're going to have a very sharp spike here, which goes off my page. And then we don't really have so much precise focusing down here, but maybe there's a little bit of light, so like that.

And then for number 4, we've got away from having this crossing here, so we won't see it so strongly. It's going to look a little bit more like this, but something maybe a little bit like that. Now, one thing I haven't really exaggerated enough here is the fact that as these rays move outwards, the positions of these maxima and minima are changing. So in reality, maybe this one is a little bit closer in, and then I'll make these ones go further out. So you can see that actually the position of where we get peak intensity is going to change depending on how far away we are from the detector, because the rays have traveled, and so their displacement has changed as well.

So there's a lot going on inside here. You can see that there's a lot of richness, and you can see that although we can identify why there are light and dark regions inside our image, it's hard to map them directly back onto whatever's going on inside the plasma. And again, that's because, as I said, this is not an image.

You can tell it's not an image because if you move your detector or if you move where you're taking the data, you'd get a very different picture out here. So although this maxima here corresponds to this region of the plasma, it shows up at different positions on your detector at different places. So you can't really say this is a direct image.

It is bijective, so this is means that at least for these smaller deflections, we have what's a one-to-one mapping. That's certainly true for the small deflections, 1 and 2, like that, for these two slices here. Once we get to 3 and we have the rays crossing, you can see that at this point, the light which is coming from these two regions now maps into the same place on our image-- or on our detector, and so we no longer have that one-to-one mapping. And indeed, after the rays have crossed, these two cross over, and so it's very confusing looking at this and trying to work out where all the light has come from.

Although we do have this one-to-one map for 1 and 2-- so in principle, we can work out where everything came from-- it doesn't have properties that we'd normally like to have from an image. So for example, parallel lines do not go to parallel lines from the object plane to the image plane, as we sort of discussed over here, and so you might want that in a normal imaging diagnostic. If you take a picture of a square and it comes out looking like something spaghetti, that's not really an image. And also, the lengths are not preserved.

Things get even worse-- preserved-- things get even worse for 3 and 4 because now we don't even have this bijective property. We end up in a regime which is called the caustic regime, which I'll talk more about later. Caustics come up all the time in other fields. They are sometimes called optical catastrophes, and although they're very, very pretty, they do make analysis of this data very, very difficult.

So what I'm trying to show here is something as simple as moving exactly where you make this measurement makes a big difference to the data that you have. It turns out that the easiest place theoretically to analyze your data is nice and close to your plasma, but the easiest place to actually make the measurement with decent signal-to-noise is somewhere around about 3. So you can't have it all. You can't end up in a regime where you have the absolute best data.

Again, this has all been very qualitative. We're about to make it quantitative by showing, at least in this small deflection regime, what you can measure. But does anyone have any questions on this general schema before we keep going? Yeah?

AUDIENCE: Are you saying that there's a bijective map between your measured intensity distribution and the incident intensity distribution, or between the measured intensity distribution and the plasma density profile?

JACK HARE: Really the former.

AUDIENCE: OK.

JACK HARE: So between the incident laser, the laser on this side, and on this side. But I would argue that if you have that, you can then infer something about the plasma that you've gone through, because you have some idea of-- if you're measuring here and you know it corresponds to here, then you know that the chord that you took for the plasma was roughly that.

AUDIENCE: Right, you have--

JACK HARE: So you have some idea of what sort of plasma properties you were sensing. Or, alternative way around, you have some idea that if you ended up here, that you must have gone through this bit of plasma, and to get that displacement, that bit of plasma must have had some certain density gradients within it.

AUDIENCE: Right, like, some sort of average density gradient. OK.

JACK HARE: Yes, exactly. This is all line integrated, and we'll talk about that in a moment when we-- we'll make this displacement. We will work out-- I guess it's a y, isn't it, because I did choose a y-axis. We will work out what that displacement is in terms of the plasma parameters, so we'll make that quantitative.

But yeah, this is kind of what I'm getting at, is have some idea of what you're actually measuring. Once you get into this caustic regime, you don't really know exactly where everything has come from anymore. Yeah, so I saw a hand first.

AUDIENCE: Yeah, is the intensity still proportional to the density gradient, as it was previously? Or is it--

JACK HARE: I'm sad that you took away from my previous lecture that initially the intensity is proportional to the density gradient. It almost never is in any realistic situation. So that's not true, and it's still not true here. And we will derive that it's proportional to the second derivative of the density gradient, and we'll also show that is also not true in most reasonable cases.

So please take away from this that neither of these diagnostics are easy to interpret, but you will read in the textbooks or online that they're proportional to the first or the second derivative. It's almost impossible to set up your diagnostics such that that's actually true. So, cool.

AUDIENCE: So you say that if we are at, like, 2, we can guess that it's coming from the first ray, but how can we say that? Or why can't we say that it's like the lower ray, incredibly deflected off? So how can we know what the caustic regime is?

JACK HARE: We will-- yeah. Yeah, so that's a great point. So it is very hard to find the caustic regime, but the caustic regime is defined by extreme intensity variations. So if you see only small intensity variations, you can't be in the caustic regime.

Theoretically, this is a spike to infinity, right. Fortunately, the universe doesn't allow that to happen because your optics aren't perfect and your detector isn't perfect, but this is very, very bright. So you can tell by looking at an image, if it's got no hugely bright regions, you're not in this regime.

And then your second question was, yeah, but still, it could come from somewhere else. You're right. When we talk about some of the advanced methods for, I guess, deconvolving or processing this data, we will come across some long, exciting-sounding phrases like optimal transport and Voronoi diagrams, and there you're trying to minimize how ridiculous your density distribution has to be to give you this result.

And so there are very mathematically grounded ways of trying to put this back, but if you're just staring at an image and it's got small intensity perturbations, you're probably going to be like, hey, it's most likely to have come from up here. It's unlikely that it went, whoop, like that.

AUDIENCE: So if we image way far to the right, then it would start to look more reasonable, right? But then the image would have--

JACK HARE: Well, yes. What I haven't drawn here is, if I draw in more rays, they will actually be more-- oh, can do it easily here? We'll find out that there are more caustics coming in later on. So once you've gone past this point, you will always have caustics in here.

I just haven't drawn enough rays to make that point really clearly. But if you can imagine drawing more in, you'll see that they will-- like, maybe there's one that doesn't cross this one here, but eventually-- no, that doesn't work because it's not a straight line. Like that, OK, and then this one, crop this here, we'd have a caustic in that region.

AUDIENCE: Right. So they eventually just start processing more and more and more.

JACK HARE: They will eventually cross. And so I guess what I'd say is there is always a point, some place where you can put your detector back here where you will end up in the caustic regime. And there's indeed a dimensionless parameter that tells you whether you're in the caustic regime or not, and it's to do with the deflection angle and this distance.

Thank you. Cool. Any other questions? Anything from Columbia?

OK. So let's try and make this a bit quantitative, because I can see that folks want to get some numbers into this. I think here I'm basically following Hutchinson, so if you need to look up the equations in more detail, this is where you want to head. We're going to be working in a regime with small angles, so small values of theta.

And remember that theta is going to be equal to, I guess-- I put d dy here, though I want to make this kind of two dimensional, so I'm just going to write "gradient" here. And you can think of theta as a vector which contains the angle with respect to the x-axis and the angle with respect to the y-axis here.

OK. This is the gradient of the line-integrated refractive index. So I'm still going to work in refractive index here because it's a little bit more compact, and also because this applies to any inhomogeneous medium, not just the plasma, so you could use this for air and other things like that. So we're going to write it just in terms of N, and at the end, I'll turn this into plasma density so you can see the final result.

So we're going to assume that we've got our rays of light, again, incident like this onto some plane that we're going to call-- that has coordinates x and y. We're going to have some initial intensity profile that's incident on this plasma here. So our plasma is just past the plane we're doing this. And of course, this could be something like in uniform, or it could be Gaussian. It could be whatever you want, so whatever you can actually come up with for your laser probing.

And our rays are going to get deflected by some angle theta, like this. And then we're going to have our detector, and that's going to be in a prime coordinate system x prime, y prime. And of course, if don't want to put my detector just here, I can always put the lens-- no, I can put my detector here, and that would also have x prime, y prime, with maybe any magnification that the lens does, but that's not really relevant to this. That's just optics.

So we're trying to work out how we get from the intensity initial to the intensity-- what am I going to call it-- I detector, x prime, y prime, like this. OK. So we can just stare at this and do some simple geometry. We can say that if we're just talking about coordinates, x prime, y prime, is going to equal wherever we started out, x plus-- ah, this is important. We need a length scale.

We're going to put our detector at distance L from the plasma here. We're assuming the plasma is pretty thin still. So this is going to be Ld dx of the integral of N dl, and this coordinate in y is going to be y plus L times d dy integral of N dl.

So again, I said we're using a small-angle approximation, so we've taken the approximation that tan theta is approximately sine theta is approximately theta, so this is just a simple linear relationship here, where this is L times the angle in x and this is L times the angle of the line, the theta line. And so we could write this more compactly in a sort of vector notation as some vector x prime is equal to some vector x plus L, gradient operator on N dl, like that. So this is just another representation of this.

OK. Now, one thing that we need in order to make some progress here is we need to assume that the overall intensity, so the integral of this over x and y, is equal to the integral of this over x and y. So we're sort of conserving our intensity, so we could write that down as I on the detector d x prime d y prime is equal to I incident dx dy. So the plasma is not absorbing, and I guess it's also not emitting any light in this wavelength equation. Yeah, that's a pretty reasonable approximation.

And then we can skip ahead, and we can say that the light which is incident divided-- the intensity which is incident divided by the intensity on our detector is going to be equal to 1 plus gradient squared-- sorry, 1 plus L times gradient squared of the integral of the refractive index along the path. And if we work with the relatively small value of this, so if we assume that this is much, much less than 1, we can rewrite this in terms of a change in intensity. So this would be the change in intensity in our image delta I normalized to our initial intensity, and that is going to be equal to 1 minus L over 2 critical density, delta squared integral of any dl, where in this last step I've substituted out the refractive index for the expression that we had for Ne much less than N critical.

So again, we've made an assumption that we have small intensity variations. We've also made the assumption that Ne is much, much less than N critical, so we can use this nice, linear formula for the refractive index of the plasma. That simplifies it.

So this is kind of like our final nice result here, so for very small intensity variations, you do indeed get an intensity variation which is proportional to 1 minus gradient squared of your electron density. So using this formula, where do we expect to have bright regions in our plasma? So where-- or, sorry, where do we expect to have bright regions in our interferogram? What do they correspond to? Yeah? Anyone?

AUDIENCE: The bright regions would be places where the second term is small, so the second derivative being small, the point of inflection of the density.

JACK HARE: OK, so only where it's small?

AUDIENCE: Like, where it's negative and big?

JACK HARE: Beg pardon?

AUDIENCE: Like, where it's negative and big?

JACK HARE: Right. Yeah, exactly. So the bright regions here are going to correspond to minima in the density, and the dark regions are going to correspond to maxima. And that's because, very roughly, when we think about what the plasma is doing, we see that minima in the electron density act as focusing lenses and maxima act as diverging regions. And so that's kind of what we saw around about here. OK.

AUDIENCE: How did you get the-- how did you get the I initial over Id formula, with the Laplace--

JACK HARE: I skipped a few of the steps.

AUDIENCE: OK.

JACK HARE: Yeah, it's not obvious, but you can go to Hutchinson's book and see if you can follow the derivation there. But yeah, there's a little bit of magic to do that step there. The thing that you want to recognize is that we're doing something that looks like it's got a Jacobian or something like that involved there, so it all eventually ends up working out. There's also a paper by Kugland that I'll mention later that's really good for this stuff, if you want to see an alternative derivation.

AUDIENCE: OK. I'm wondering, how do we-- what is the precise meaning of the gradient of a line-integrated quantity? Is it like we're changing the limits of integration by an infinitesimal amount? Like, what does that mean?

JACK HARE: I mean, mathematically, you can just work it out, right. There's nothing wrong with this, because, for example, if we are doing any of x, y, and z integrated dz, like that, then you can still take the derivative of this. You just won't have any z components anymore. You'll just have components in x and y.

AUDIENCE: OK.

JACK HARE: In reality, when I see this written down, I often see people-- and I do the same thing-- who sort of sneakily move this integration sign to here, and then it makes a little bit more sense because you're actually looking-- for each step, dl, that the ray goes through the plasma, you look at what the local density gradient is. And I think these two things are different, kind of obviously, but they are pretty close for thin plasmas. And what most of the time people are doing is making approximations that the width of this plasma, which might be a or something like that, is much, much less than L.

So this is effectively assuming that the actual path that the ray takes through the plasma is unimportant. We only just care about what path it takes through free space afterwards, which is going to be in a straight line. If you don't have this condition, you actually end up in the caustic regime more easily, and again, I'll point you to some references later which talk about this in a bit more detail.

For the-- in the case where the thickness of the plasma is much less than the distance between the plasma and the detector, it doesn't really matter which way around you do this operation. So, yeah, good question. Yeah?

AUDIENCE: When you say "the thin plasma regime," we're sort of saying that the path inside the plasma doesn't matter. Isn't that just like the schlieren?

JACK HARE: Yeah, and indeed, this effect pops up in schlieren as well. But in schlieren, the bigger effect you get is by putting the stop and blocking out the rays, but you will have shadowgraphic effects inside your schlieren imaging system, too.

AUDIENCE: OK, I see.

JACK HARE: So, yeah, I just introduced the schlieren first, and now we have the shadowgraphy. But these are both present here, and they're also present in some of the interferometry we'll talk about later. But in each of these, there's something that causes the biggest intensity modulations. In the schlieren, it's the schlieren stop, clearly, but here, again, we're only looking in this regime with small intensity modulation for the moment. But without the schlieren stop, this is the effect that shows up. Mm-hmm?

AUDIENCE: Our signal is proportional to the Laplacian density, and so why is it that maxima and minima of the signal correspond to maxima and minima of the density rather than maxima and minima of the density gradient?

JACK HARE: Hm-hm, hm-hm, hm. Yeah, I see what you're saying.

AUDIENCE: I think it's because it's the Laplacian of the line integral of the density as opposed to the Laplacian density.

AUDIENCE: Oh, like, the integration over space brings you back a level?

AUDIENCE: Yeah.

JACK HARE: I'll have a look at that and check. I do see what you're saying, and I see what you're saying. And I don't know what the resolution is right at the moment, but yeah, I'll have a look and see if I can work that out.

AUDIENCE: OK. Thank you.

JACK HARE: Cool. So again, this looks really nice because it looks like you can get out the second derivative with Nz, and then maybe you could double-integrate that and you could get out the actual line-integrated electron density. But in reality, the assumptions we'd have to make to get here mean this is really hard because the actual signal term, we have assumed, is much, much smaller than 1, which means that it's really, really hard to measure.

So for any realistic system with a realistic signal-to-noise, we don't want to have this limitation. We don't want to be working at position 1 or even position 2. We're actually going to get the best signal when we go to position 3, where all of this no longer applies and we can no longer get this nice result.

So this is what I'm saying where you will find people saying shadowgraphy is proportional to the second derivative of the electron density, and that's sort of true. But I've never seen anyone do it. Like, it's not actually possible to do that measurement in a meaningful way. You have to work in a regime where you can actually measure the modulation. And I'll show you some example pictures of schlieren and shadowgraphy towards the end of this lecture so you get an idea of what it looks like in a plasma.

OK. I think I've kind of said a lot of this already. If we end up in this regime here, where we have these caustics, we've clearly lost information. We can no longer do this mapping here because it's no longer unique.

We've got raised crossings, so different bits of x prime and y prime are mapping onto-- or, one place in x prime, y prime, might map onto multiple x and y, and so we can't do this simple thing. And it's kind of obvious, actually, that you'll be losing information, and even a more advanced technique isn't going to be able to do the reconstruction. So I want to talk just now a little bit about some of these advanced reconstruction techniques which go beyond this, and then I'll show you these examples.

So the first paper I've seen that really tackles this very nicely is Kugland, et al., in *RSI,* 2012. So Kugland points out pretty quickly that shadowgraphy is a direct equivalent to proton radiography. So mathematically, they both deal with the same quantities, which is a sort of deflection potential. In proton radiography, this deflection potential is to do with electric and magnetic fields-- and we'll get on to proton radiography later-- and in shadowgraphy, the deflection potential is to do with density gradients.

But once you work in terms of this deflection potential and forget where it came from, you get exactly the same results mathematically, and he has this nice geometric approach. And so if you've found the derivation I just did not very convincing, you can go and have a look at this. It's a little bit more rigorous, and also, he then extends it into the caustic regime and beyond, and shows what you would expect to get from a plasma where you have caustics.

So this is nice, but he still doesn't really tell you how to analyze it. It's really talking about the problem where you go from knowing the density of the plasma to predicting what you're going to get out. Going the opposite direction to going from your intensity variations to the density isn't particularly well developed, so this is what some people call the forward problem. And the forward problem is going from Ne of x, y, and z, to intensity on your detector in x prime and y prime-- useful, but not exactly the solution.

So then in 2017, there were two almost competing papers that came out about this. There was a paper by Kasim, et al., in *Physical Review E,* 2017. What Kasim did is he tried to reconstruct this deflection potential that Kugland had come up with. And he did this using a technique-- that was borrowed from, some would say, computer graphics, or at least some fields of applied mathematics-- which used a Voronoi diagram. Has anyone come across Voronoi diagrams before?

AUDIENCE: They're huge in robotics.

JACK HARE: Sorry?

AUDIENCE: They're huge in robotics.

JACK HARE: OK, cool. Right, so there's like-- they basically-- I think this worked well because they took some work that other people have been doing in different fields and applied it here. So a Voronoi diagram-- ha-ha, please don't shout at me if I get this wrong-- is roughly, if you have a series of points which are randomly distributed, how do you draw quadrilaterals around them such that every point inside the quadrilateral is closest to this point and not to any of the other points? So it's a way of tiling up a space.

And from this, you can imagine that these tiles that you've produced in your shadowgraphy can then be related back to a more uniform grid of tiles which all have the same shape and the same intensity, which is your intensity beforehand, and this is your intensity at the detector. Again, this is just a very hand-wavy sketch of what they did. And they came up with an algorithm to do this, and this enabled them to do the inverse problem, which is intensity at our detector in x prime, y prime, going back towards density in x, y, and z.

And as we discussed before, this is an ill-posed problem. There are a large family of possible density structures that produce the same intensity. But this Voronoi diagram is making conservative assumptions about where light has come from in order to be able to put stuff back. At almost exactly the same time, Bott-- both these groups are at Oxford. Yes?

AUDIENCE: So would this be applicable to shooting an array of lasers through, or shooting in a grid?

JACK HARE: Oh, so no one does that, but they should. So they did it in proton radiography in the early days. They actually had little beamlets, and then you can uniquely work out where each beamlet is deflected. But of course, you only get, like, n beamlets of points of data, and it's like, it doesn't look pretty, you can't put it in nature, that sort of thing.

So people moved very quickly away from that technique, so now we have data that's impossible to analyze but is very beautiful. Whereas, we used to have data that was entirely analyzable but not very beautiful. So I have opinions.

AUDIENCE: So this actually would be just a finite number of rays that should be--

JACK HARE: No, I mean, this is entirely to do with-- this is not to do with the finite number of rays. This is still to do with our nice, initially uniform beam. It's just the way that we segment up the final image.

These dots don't really exist. You're actually-- in reality, you're trying to find-- ha, this is where I'm probably going to get it wrong-- you're trying to find regions which contain the same intensity as one of these initial squares, and you're trying to sort of tile them together in such a way that you don't have to have something that looks like a congressional district, right, where it--

[LAUGHTER]

OK, because that's obviously silly. That's unlikely to happen. So we're trying to-- this is-- you could maybe use this algorithm to sort out a lot of problems in this. Anyway, so at the same time, Bott was working on an algorithm that ended up doing the same thing, and this is in *JPP* in 2017.

This paper is something like 120 pages long. It very much helps to have your thesis advisor be the editor in chief of *JPP* if you want to publish a paper with them. And they use a very interesting technique called the-- I'm probably going to say this wrong, and I certainly handled all the accents in my notes here-- Monge-Ampere optimal transport. And this optimal transport algorithm had actually won someone the Fields Medal only a few years before, Cédric Villani, who wears these incredibly huge bow ties, and he has a wonderful book called *The Life of the Theorem* where he discusses how great it is to be Cédric Villani.

But, so this algorithm was derived not at all to do with proton radiography, but it is to do with how we-- what is the most conservative way to map one function into another function like this. And once you derive that, you can then put it back where you started from, so this also enables you to do this same inverse problem like this. And I'm not even going to slightly go through my guess at what the Monge-Ampere equation does because I don't have one.

But it seems to work. They give similar results. It turns out this is much faster. I think most people use some version of this at the moment, but this one is maybe easier to understand.

So both of these give similar results with slightly different techniques And I think, in the end, Kasim wrote a code based on Bott's paper that was faster than what Bott had done, so I think people have converged on using something to do with this. And again, these techniques were mostly invented for proton radiography, but now, as we said, the mathematics is the same for shadowgraphy. I have not seen anyone use either technique to properly analyze shadowgraphy, but it should be possible. Yes?

AUDIENCE: How does this problem differ from more general tomography problems we encounter? Right, like, tomography or when people do tomographic reconstructions--

JACK HARE: There's no tomography here. We only have one line of sight. You can't do a tomographic reconstruction--

AUDIENCE: OK, so there's only one line of sight.

JACK HARE: Yeah, I mean you can-- OK, so now you can ask yourself, If I have multiple lines of sight, can I do tomographic reconstruction? Which, like, yes, obviously, but it's also hard. But, you know, it might be possible.

But this is a single line of sight, so we're not trying to-- ah, thank you. This is the mistake I've been making. We are not actually getting this out. We are getting our best guess at this out.

Right, so this is not a full three-dimensional reconstruction. This is still a reconstruction of the line-integrated electron density here. So, yeah, yeah, that's a good point. I forgot about that.

And it's sort of obvious that you should be able to do that. But there still could be multiple profiles that still produce the same intensity distribution, so it's still not particularly well posed. If you do have some caustics inside your protected image, none of these work, right, so we no longer are able to do the reconstruction.

You can actually have a go at doing the reconstruction if you have some strong priors. So if you put-- you have some optimization algorithm that thinks, like, there's a shot here and that shot is going to cause caustics, and I think the caustics will look like this, you might be able to do it.

But of course, you'd obviously have very strong priors. The techniques are very line integrated, as I just remember them pointed out here, so we're not getting a full 3D structure. That's maybe a bit too much to ask from a diagnostic which is clearly line integrated, but it's still a limitation.

And then the final problem that in proton radiography is particularly profound is actually how reproducible this is, how reproducible your initial-- I'm going to run out of space-- your initial intensity is, because they-- before you do the experiment, you fire your laser beam through the chamber, and you measure that beam profile. But then when you actually do the experiment, that beam profile changes. You know, lasers are not completely stable.

The beam profile changes from time to time. And so that means what you think you're mapping from here to here is actually slightly different, and that's going to introduce some noise. This is very important for proton radiography, where it's very hard to measure the beam.

For a laser shadowgraphy setup, you actually have more of a chance. You can put a beam splitter for the plasma and sample the beam itself, so you actually simultaneously measure this quantity and this quantity. So there's a lot of scope for doing some really cool stuff with shadowgraphy. Questions? Yeah?

AUDIENCE: Are there obvious practical reasons why you wouldn't be measuring the dynamically helpful distances, or even if you have--

JACK HARE: I think it's a great idea. No one's done it. I want to do it.

[LAUGHTER]

Yeah, absolutely. So it seems to me like if you have these images of different places, you should be able to reconstruct the trajectories of the rays, and that would give you more information.

AUDIENCE: Yeah.

JACK HARE: And in fact, in some of the first proton radiography papers, this is discussed, from looking at the position through multiple stacks. But I haven't seen it actually done in practice.

AUDIENCE: You could kind of do it.

JACK HARE: OK.

AUDIENCE: But not fully, because it's-- the proton-- typically, all of your energy gets deposited first--

JACK HARE: Yeah, so I think that's the problem, but I like the idea. Like, it's a cool-- if it works, if people used this, yeah. But you could definitely do it with shadowgraphy, relatively straightforward.

Yeah. The other thing you could do is you could put multiple lasers at different wavelengths through, and they'd be deflected by different angles. And then you could use those different deflections, like when you have your proton radiography and you use different particle energies to determine between electric and magnetic fields. Here, we don't have electric and magnetic fields.

There's only one thing that can cause the deflections, which is density gradients, but those different colors would enable you to. So you could imagine that one of your other rays would get deflected less if it had a shorter wavelength, and it would take a trajectory like that. So by comparing where it ends up in one wavelength to where it ends up in the other wavelength, you should be able to actually just precisely measure the angle that's made and therefore what the density gradients are inside, but I haven't seen anyone try that yet. Yeah?

AUDIENCE: When you responded to John's question, you mentioned that you only have one line of sight here. But if your detector is able to do x and y positions of your intensity, could you consider each pixel as a different line of sight?

JACK HARE: But it's only a chord through a certain bit of the plasma. It's not a line of sight through the same bit of plasma. When we do tomography, you know, you imagine you've got some cloud, and you're looking through the same bit of plasma from multiple angles. Here, you're looking at different bits of plasma, so you can't topographically reconstruct the density there because it's literally a different place in the plasma.

AUDIENCE: Sure, OK, yes.

JACK HARE: Yeah?

AUDIENCE: I would've expected a smaller wave to be reflected more because the length scale of the plasma this large wavelength--

JACK HARE: We are doing geometric optics, so we don't actually care about the wavelength.

AUDIENCE: Oh, oh, OK.

JACK HARE: Yeah, so in systems where you're not doing geometric optics, where the wavelength is comparable to the size of the plasma, that would be a diffraction effect, and that would be more important.

AUDIENCE: Got you. OK.

JACK HARE: But we are actually doing geometric optics where that isn't important, but you're right. So all of the stuff I've been doing, the reason I'm drawing straight lines everywhere is I'm doing geometric optics and doing ray optics. So in that case, a different deflection angle comes about because for shorter wavelengths, the critical density is higher, and so this quantity becomes smaller.

This is just a little bit of a lighthearted little picture show. And we may finish off the lecture with this, or we might get started on interferometry, depending on how I feel. But here are some nice pictures.

Shadowgraphy is absolutely everywhere. You have already seen many shadowgraphs before. If you've ever seen a mirage, that is a shadowgraph. That is the natural focusing of light by refractive index gradients in hot air, right, and you see that shimmering. You see the fact that there appear to be mountains below where the mountains actually are.

That's because the rays of light have been bent by the hot air back upwards into your eye. And so this is what I mean when I say that shadowgraphy is not an image. Whenever you take a shadowgraphy thing, what you're really seeing is some sort of mirage.

So the first person to ever, we know, study geography was this guy Jean-Paul Marat. That's a portrait of him. Here's his shadowgrams. He actually drew these by hand because he didn't have cameras back in the day.

This is the guy. Here he is in 1793. He was actually deeply unpleasant. He was a Jacobin. He was responsible for the deaths of hundreds of thousands of people in the French Revolution, and he was eventually murdered in the bathtub.

[LAUGHTER]

But before he was murdered, he had a very famous guest. He had a very famous guest who came and actually sat, and he sketched the shadowgraphic effect of this-- the effect of this guest's bald head on the air around. So does anyone know who this is? That is, of course, Ben Franklin-- so there we go-- who had made a habit of hanging around in France. But these days, we usually do things which are much more exciting than Ben Franklin's head.

So here are several different models. These are models for the Gemini capsule that was part of the American space program, and they want to understand what the shockwaves were around it. So you can see that this capsule is coming from right to left. It's got this blunt end here, and we have this very well-defined bow-shock structure here.

So we think that that's a caustic? It's a big intensity variation, right? It's extremely black here, and it's very bright around the outside here.

Behind it, what do we have? We've got a set of shocks there, a shock here. There's another set of shocks coming off here and here, which interact with the outer shocks, and then we have this beautifully turbulent flow behind it, right, and the same thing here on this more hemispherical object.

And so don't know how big these were, but you can do this in a wind tunnel. In this case, it was probably not a wind tunnel but a static tube of gas that they fired these through with a cannon. And then they would have used some sort of bright light source, probably not a laser, in order to do this.

Also, when you start looking at the literature, you get lots of beautiful pictures of bullets. Bullets are particularly good sources. So I believe, in this picture, the gun is just here, and so you can see the cloud of exhaust vapor coming out of the gun barrel.

You can see this is the sound wave of the shot going off, and then supersonically moving away from the gun is the bullet. We can see the trail with the defined structure behind it and these very clearly defined shockwaves here. And again, these shockwaves have these light and dark regions corresponding to changes in refractive index. And this is, I think, a zoom-in of that photograph, but I can't be sure-- it looks like it is-- of that.

So you can see you can get an extraordinary amount of qualitative detail. What you can measure from this straight away is the shock opening angle, so you can measure the Mach number. What's going to be a lot more tricky looking at these pictures is to work out what the density and temperature of the air is everywhere inside this picture.

That's not really going to be doable because we're not in this small-intensity-variation regime. We've deliberately gone into a regime where we get caustics, which means we also get a strong intensity variation, so we can actually measure. If we were looking at one of those small-intensity-variation shadowgraphs, it would be very boring. It would mostly be gray with very, very tiny modulations to it, so. Yeah?

AUDIENCE: On figure B, why are the light and dark regions above and below swaths?

JACK HARE: Pass. I'll put it on a little problem set.

[LAUGHTER]

I don't know immediately why that is. The thing is also to remember, if you end up in a regime where you have caustics and you have strong deflections of your rays, you may actually end up in a regime where the rays don't go through your first optics, that they may be deflected out of the deflection volume of your first optic. And then they would just show up as dark regions.

So you can also have, overlaid with the shadowgraphy effects, what are effectively schlieren-type effects, where we are rejecting rays by their angle, but that's just due to the physical size of the optics we use. So I don't know if that's the case here, but it does complicate the interpretation further. Yes?

AUDIENCE: In the previous pictures with the Gemini capsule, there's the effect of pressure waves, shockwaves, and there's also supposedly intense heating that should be, like, on the surface of the sphere, and that's also going to change the refraction index.

JACK HARE: Yes.

AUDIENCE: Which effect of those-- like, how could we distinguish if something is a thermal refractive index change or a pressure--

JACK HARE: In air.

AUDIENCE: In air?

JACK HARE: Yeah, I think that's very difficult because they both are just the change in refractive index.

AUDIENCE: Yeah?

JACK HARE: Yeah. I don't think it's possible to tell the difference between those two, yeah. There's also some funky pictures in this book. This looks pretty straightforward until you realize the bullet's actually flying backwards.

I don't know why they did that, but there we go. This is a great book, by the way, Settles' *Schlieren and Shadowgraph Techniques.* It's got wonderful pictures inside it if you want to know more about this. It's in the bibliography on the syllabus, and it's a great read. But I appreciate most people are not going to be using shadowgraphy and schlieren, but it's still good stuff.

But you don't have to fire bullets at things. This is actually a picture of the author of this book writing his book next to his heating unit. There he is. His head is not quite as impressive as Ben Franklin, but there's his computer as he types away.

And you can see that you can actually make these measurements even in relatively benign conditions, and that's because, again, if we put our detector further and further back, even small variations in the angle going through the refractive medium are going to be mapped into large intensity variations. So people use this technique to look for flows of air. You can look for flows of air for all sorts of reasons. A pretty benign reason would be in the HVAC industry, if you want to see whether these things are working.

So there are applications for shadowgraphic techniques and schlieren techniques just in very benign conditions like this, but we, of course, are interested in plasmas. So we talked about this a little bit more on the last lecture, but remember that we need a very bright light source to overcome the self-emission from a plasma. So that means we really have to go to a laser. We just don't have any light sources that are great.

Lasers are actually not ideal for any of these techniques. You really want to have a nice, large focal spot. That turns out to be true for shadowgraphy as well, but we needn't really go into why.

And so this small focal spot gives us quite limited dynamic range, not to mention the fact that we are assuming in all of this that we don't have any coherence effect so we don't have interference, and we'll talk a lot about interference in a little bit. And so really, we don't want to have that coherence when we're doing schlieren and shadowgraphy, but if you've got a laser, we tend to have coherence that we don't want. So these are not great, but you can still get some nice images out of it.

Here's a device called an X-pinch. It consists of two wires that are crossed here. This is only 1 millimeter across here, so this is a pretty small object. And these are X-ray, gated X-ray images of the X-pinch.

We put, in this case, 200 kiloamps through each wire, and it forms a plasma here, which pinches. The fields here are maybe a hundred or a thousand tesla. And it pinches the plasma inwards like this, compressing it up so it becomes extremely hot, and it emits a burst of X-rays, which you can then use for imaging things. So when we get on to self-emission diagnostics, X-ray diagnostics, we'll talk a little bit about this.

So these are the X-ray images, but this rather beautiful schlieren image was captured of this X-pinch in 2008. Here, they used a dark-field schlieren system with a circular stop. You can tell that because it looks up-down symmetric, so we're not-- we don't have a knife edge, but we have some distinction between the different directions. And you can tell it's dark-field because outside, where there is no plasma, it's dark here. If it was light-field, this region would be filled with laser light, and we'd have darkness wherever we have lightness here.

And you can see a beautiful amount of detail. This projector isn't really doing it justice. You can see in the center here, this is the pinching region.

It's far too dense for all of the laser light to make it through. It's much above the critical density, and so there's very strong refraction of the laser light outwards, dark in the center.

There's jets of plasma going up and down out of this compressed region. And there's also ablation streams coming off each of the wires which have this beautiful modulated pattern, which is actually due to a instability in the wire ablation process. So this is extremely rich. There's a lot of information you can get out of this, even though it's not quantitative.

Another thing that's been done is using schlieren imaging to image shocks in what's initially a gas, but quickly becomes a plasma. So what we had in these experiments was a metal liner. This is only about 5 millimeters across, so it's still pretty small. And it was filled with a gas, 8 millibars of argon, 15 millibars of nitrogen.

A current was put up through this metal cylinder. Again, the cylinder is only about that tall and that wide. And as it does so, it heats the outside of the cylinder and launches a shockwave inwards.

And that shockwave couples the gas, and it launches this first shockwave in. And as the current continues to rise, you actually get above the melt point of the metal, and that launches another shockwave due to material strength that starts coming in. So we get these converging shockwaves.

And the beautiful thing here is that in argon, we had a beautifully circular shock, but in nitrogen, for some reason, we got this hexagonal shape of shock which has never been explained, absolutely bizarre, because there's no sixfold symmetry in this system. It shouldn't happen. So there's some instability which is giving it this really bizarre shape.

And again, this was dark-field with a circular stop. That's a pretty standard configuration. Dark-field is a bit more sensitive than light-field because if you see any light at all, you know that it's being deflected, and that's really what you're looking for here.

And then, shadowgraphy, this is an image that I took of an imploding wire array. So again, this thing is only about 16 millimeters tall, 16 millimeters in diameter. We've got eight carbon rods here. Current goes up through the rods, ablates plasma off them. J-cross-B force accelerates the plasma inwards, and you get a Z-pinch column in the center here.

And looking from the side using a green laser beam, we can see that we've got shadows corresponding to the four wires. So there's four wires on this side, and they're blocking the four wires on the other side. And you can see the column of plasma in the middle here, and what you can see is there's these very strong modulations to the intensity.

These are caustics, OK. We also see modulations sort of flow, as we saw in the X-pinch. And these caustics mean that this data, while very pretty, is pretty useless. There's not much we can actually do with it because we can't do any decent analysis.

But it does tell us, because the caustics are on a large range of different spatial scales, that we must have density perturbations inside the plasma, and a lot of different spatial scales. And so this means that plasma is likely to be turbulent. So this is a little turbulent Z-pinch inside a pulsed power machine.

So is that it? Oh, and you can also do some pretty good 3D simulations of these things. And then you can spend a lot of time doing Monte Carlo ray tracing, tracking rays through them, and seeing what the shadowgraphy looks like.

And don't think this is a particularly bad match between what we saw from our simulations and what we saw in our actual data here. So it is possible to use computational tools to work out what we would predict. So that's it. That's all I've got on shadowgraphy and schlieren. Any questions on that? Yeah?

AUDIENCE: Even if you have to worry about caustics, and you don't feel like you can get all the way back to initial distribution, can you at least analyze it for frequency distribution or something to get, like, oh, I must have had this many spatial scales in my original plasma or something, or--

JACK HARE: Yeah, so it's very tempting, when you have an image or when you have a time-series bit of data, to Fourier transform it and look for spectral content. And in particular, when we're talking about turbulence, we might do that, and we might look for power spectra corresponding to some turbulent density fluctuation spectrum, so, like, Kolmogorov "K to the minus 5/3" distribution.

And so I did this. I did this, and of course, you get a really nice K to the minus 5/3 on this. And then I took the background image, the one without the plasma, and you also get K to the minus 5/3. And then I took a photograph of the experimental apparatus and Fourier-transformed the photograph, and you also get something like K to the minus 5/3.

The trouble is, when you're doing Fourier transforms on images, you've got to think, How many pixels have I got? You've probably got, like, a thousand by a thousand pixels. And so when you're doing your Fourier transform, your dynamic range is only going to be about 10 plus 3.

But at those large scales, at the smallest K, it's going to be like large-scale structure, so you wouldn't fit a power spectrum there. And at small scales, it's going to be down at pixel noise, so you wouldn't fit it there. So actually, you've only got maybe an order of magnitude, and you can fit any straight line you want to a curve and claim that you've got K to the minus 5/3 or K to the minus 3/2 or whatever, because when you do turbulence theory, they all turn out to be roughly the same.

So that's one reason it's really hard just to Fourier-transform these. The second reason is, as we've talked about before, this is a mirage. It's not an image.

So if I see a region like this black region here, or another region-- maybe it's easier to point out on this one. You see these sort of black voids here, and you think, OK, I could just be like, hey, this is about 2 millimeters long, this is 1 millimeter long, make a histogram, fit a power law or something to it. But these don't represent an object that is 2 millimeters long. They are a defocusing mechanism.

That could be a really tiny region of very high density, that defocus, and then it's projected out into this larger region. So we can't-- there's no spatial information properly left inside this image. There's some spatial information.

This sort of structure here corresponds to this sort of structure, and you think, it's about 5 millimeters across, it's probably slightly de-magnified, because it will actually make a larger image. But each of these individual voids no longer has the same spatial size as the structure that produced it, so it's very hard to infer things about turbulence from them. But, yeah, it's a good question. Any other questions? Mm-hmm?

AUDIENCE: Can you get back into the difficulties of using Bott's base? If you reconstruct Bott's example, even though there's the usual caustics, you still get some information where one--

JACK HARE: So don't think you can. I mean, when you run this algorithm on your data, shadowgraphy or proton radiography, one of these Monge-Ampere optimal transport algorithms, it will always give you an answer, right. So it returns it returns a solution, right.

But we know that these reconstruction algorithms don't work when we have the caustic regime, and so that solution is very suspect. Right, like, we don't believe we should trust it. So I don't think there's anything you can use from that solution to help you reconstruct the actual thing.

I think, at that point, if you've got caustics, your best bet is having some very strong priors as to the sort of plasma you think you've got, and then pushing Monte Carlo rays or protons through it, doing the forward problem, and adjusting the forward problem until it matches some of the features you see on your actual data. I don't think you can do the inverse problem easily.

But there have been a few papers where people have tried to do this, because we almost always end up in the caustic regime. So people have all this data, and they want to use it. Like, it's reasonable to try and do something with that data. It's just very hard, so, yeah. Yeah?

AUDIENCE: So for those images of, like, the heads and the room, those are not super small-scale effects. So to see those, do you just push your imaging surface really far away?

JACK HARE: If you take a laser pointer and make it diverge by taking off the little lens at the front of it, and you take a candle flame and you project it onto a wall, you will see this. Like, again, you see mirages just using sunlight, so, yeah, this is not a hard thing to observe. Yeah.

AUDIENCE: OK, so if you see something with your naked eye, it's either a really small object or something really far away?

JACK HARE: Yes.

AUDIENCE: OK.

JACK HARE: Yes, exactly, and usually, it's quite far away, right. I mean, if we-- you can often see even the heat rising from a vent or something like that, but not when you're right up close to it. Give it a go next time you see something, if it's safe. [LAUGHS] Put your eye right up next to it and see if it disappears, so, yeah.

AUDIENCE: OK. Cool.

JACK HARE: Please don't, anyway.

[LAUGHTER]

Any other questions?

AUDIENCE: Can you give one more example of when you can see-- like, in this exact duration, is like if you have sunlight coming through a window and over like a heater or something like that.

JACK HARE: Yes, onto the far wall.

AUDIENCE: Onto the far wall. OK.

JACK HARE: Yes. Yeah.

AUDIENCE: Yeah, exactly. That's what I'm saying.

JACK HARE: Another place-- the bottom of a swimming pool. So when you have waves on the top surface of a swimming pool, you get those bright lines on the bottom. Those are caustics, right, so exactly the same sort of physics is produced as these.

So there is a beautiful book called *The Natural Focusing of Light* which tries to analyze caustics but mostly does it in a way that I don't think is practical. It's theoretically beautiful, and one of the things that they point out about caustics is that every time you get a bright region, you get a dark region on one side but not on the other side, and that tells you what direction the caustic came from.

And so you might be able to trace back a series of arrows around one of these caustics here and work out where the point was that the caustic originated from, but it's very tricky to do the analysis of it. But the point is the book is called *The Natural Focusing of Light,* so people refer to this field as natural focusing.

No one has tried to make a lens. No one is trying to do any focusing. Our medium, with its inhomogeneous refractive index, has just done it for us, and then we might try and work out what we can learn about the medium from looking at the light that's gone through it. Yes?

AUDIENCE: One more thing, that-- to make sure I don't try this later. So if you have an image like the one with the bullet, for instance, there are some really clear caustics, then parts of the rest of the image look like they might be fine, if you want to analyze the rest of your image, is that--

JACK HARE: Oh, yeah, you can cut out the bit with the caustics.

AUDIENCE: OK.

JACK HARE: Yeah, yeah, that's fine. So, I mean, almost none of this image is actually suitable, but maybe this bit would be more-- because it's not small intensity variations. You can see, if you think about this gray as, like, 0.5 and you think about the white as 1 and the black as 0, you can see that you're getting modulations on the order of 0.5 inside this image, so it's clearly not in that small regime.

But it's not clear that these are caustics, so you may still be able to use one of the complicated Monge-Ampere style reconstruction techniques. You just won't be able to use the nice formula that we wrote down analytically, so.

OK, you've successfully timed me out. I was going to start talking on interferometry. Well done, everyone. So we'll leave it there, and we will pick up on interferometry on Thursday. Sounds good.