EDWARD BRIGNOLE: My name's Ed. I'm a postdoc in Cathy Drennan's lab, and previously, I had worked at the Scripps Research Institute with Francisco Asturias. So some of the work that I did there is what Liz and Joanne like, and we'll talk about that.

So I thought I'd start with just finding out, has anybody here done electron microscopy. You've done some EM. OK, on--

AUDIENCE: Gold nanoparticles with a [INAUDIBLE] spirit thing.

EDWARD BRIGNOLE: OK. Over here?

AUDIENCE: No. In St. Louis.

EDWARD BRIGNOLE: In St. Louis, OK. So you sat at the microscope and worked on obs and--

AUDIENCE: That was my favorite [INAUDIBLE]

EDWARD BRIGNOLE: Yeah. Anybody else? So how about a light microscope? You guys used? High school biology, maybe?

OK. Everybody's used a light microscope. All right. So that's good.

And then I guess at this point, you guys have had two lectures on fatty acid synthesis, so you sort of have some feel for the enzymes and who's involved and what they do. All right, so I thought we'd spend the first 20 minutes talking about electron microscopy and what it can do and how it could be used.

And there is actually quite a bit that's changed since this paper in 2009. There's a lot that's happened in the last few years. And we can talk briefly about that. And then we can move into fatty acid biosynthesis and tie that into what you guys have learned already.

And then there was a bonus paper at the end. If you are really excited about this, there's some polyketide synthase structures that have come out in the last year or two. And those are pretty interesting.

So if you've got the handout handy, there's some questions. These are what I thought we'd focus the conversation around. The first part's about fatty acid synthase in EM, and then there's these bonus ones at the end.

So when you guys were looking in the light microscope, you're probably looking at biological samples, I would guess. So you were probably looking at, say, a cell. But all the bits and pieces of the cell that perform all these interesting functions, we want to understand these. And so being able to actually see them and see them in action allows us to understand how they work.

So for instance, if you pick out this piece of machinery here-- and you can see that it's got an active site where it binds substrate and maybe moves it around or acts on it in some way. And you might have allosteric subunits, and you can find that it's got four subunits that are round like wheels. And it can move from one place to the other. This is just an analogy, but same thing would go for, say, motor proteins transporting cargo, or in this case, fatty acid synthase.

So if you've used a light microscope, the electron microscope is conceptually very similar. You've got the light source at the top versus an electron source. Condenser lens will focus that on the specimen. The objective lens forms the image.

It's magnified by the projector lens, and what you get out is this enlarged image. So you can see things that you couldn't see by eye. Light microscope, you can get up to about 1,000x. And in an electron microscope, you can go up to 500,000x or even beyond that.

So maybe an interesting place to start here is why can electron microscopes do this but light microscopes that. Why can you get only this magnification with a light microscope? Guesses?

If you look back up at the top, what are the sources? You're using light versus electrons. Why would you be able to get a higher magnification image using electrons than light?

AUDIENCE: Diffraction limit.

EDWARD   Yeah. Do you know why? Why would light have a diffraction limit that's in the micron range or

BRIGNOLE: nanometer range versus electrons in the actually like picometer range?

AUDIENCE: Since your wavelength could be--

EDWARD BRIGNOLE: Exactly. That's what I was looking for. Yeah, wavelength is what I was looking for. So either visible or even if you go to a UV light source, you're talking about nanometer-sized waves versus electrons, at the typical electron acceleration voltages that are used, are in the tens of picometer wavelength. And so that's the main reason.

You can magnify these images further. But you're not going to get any higher resolution versus an electron microscope. But you don't often hear about tens of picometer resolution images by EM.

So I guess maybe I'll flip back to this slide. So this is differences in the source. But you could actually theoretically go 100 times or more beyond these magnifications by EM. Why do you not typically hear about that? What else could be limiting resolution as you go down through this path here?

A number of you wear glasses. Do they perfectly correct your vision? They don't for me. Yeah, so same thing with these lenses.

No lens is perfect, and you've got different aberrations. So the way light or electrons that are coming in are bent is never perfect. They're not going to achieve-- and there's correctors that you can use to compensate for this.

Also the wavelength of the light, having it perfectly tuned to a particular energy of either photons or electrons, there is going to be some distribution. And so you're going to have some that are a little more redshifted or blueshifted, higher energy or lower energy. And those are going to also not come to complete focus there. And so this is for the lenses in the source, why you would typically be limited to about an angstrom resolution unless you buy some fancy correctors for your microscope to correct for spherical aberration or energy filters to correct for chromatic aberration.

So what kinds of things in the cell could we look at by electron microscopy? Maybe you guys have seen images in papers, probably in your textbooks, EM images of-- what? Help me out.

You see in sections about, say, muscle where there's a section of some muscle fiber where you can actually see some of the proteins that are involved, the filaments. So you've seen

things like that. Tissue sections. You could look at tissue sections by EM. You could look at individual cells.

Could you look at an elephant by electron microscopy? Have you ever seen that? No. So why would you not image an elephant in an electron microscope?

AUDIENCE: Simply because they're too [INAUDIBLE].

EDWARD BRIGNOLE: OK, yeah. Exactly. You'd have a really hard time preparing that elephant even if you had a microscope that was big enough.

But you could, say, take an x-ray of an elephant. Right? But what is it about electrons, maybe, that you wouldn't have to deal with with x-rays?

AUDIENCE: Killing the elephant?

EDWARD BRIGNOLE: I mean-- [LAUGHTER] Right. So why would the elephant have to be dead to image it in an electron microscope?

AUDIENCE: When you're shooting it with electrons, even just for the cell, it'll kill the cell. [INAUDIBLE] you do it to a--

Oh, wait. Don't you have to [INAUDIBLE]?

EDWARD BRIGNOLE: Yeah. So in some cases, you might negatively stain. Typically, you would stain a sample in some way or another.

But what I was getting at is the vacuum. So the microscope is under high vacuum. So electrons have mass, and they're going to interact strongly with the matter that they're going through. You actually couldn't get an electron through an elephant. You could get x-rays through an elephant though.

So thickness is one issue. You'd have to cut really thin sections of your elephant, so about 200 nanometers thick. You can go thicker than that, but then there are some other issues with resolution that occur. So this is about the high end of what you would want to be for a good EM specimen.

And then if we're looking at, say, individual macromolecules, the size of those macromolecules, to be able to look at the image and pick them out, would have to be-- if there

are single particles floating around like a virus particle or something, you can usually do that. Because they're much bigger than 100 kilodaltons.

Probably in your textbook, you've seen EM images. Or even in newspapers, you pick up the New York Times, and there's an article on Zika virus or something, and there's an EM image of it. So those are much bigger than 100 kilodaltons. But this is about the lower end for individual particles.

So I thought I'd just throw this up so you could look at the different bits and pieces of an ant in a light microscope. An electron microscope largely overlaps with the high end of the light microscope, where you could look at cells or sections of cells if the cell is a micron or more thick.

Bacteria, bacteria-like viruses, bacteriophage in this case. Electron microscopy has resolution down to this range. But in order to visualize things like this, you'd have to be able to pick them out of your image.

So you could assemble 10 kilodalton particles into, say if it's actin or something, into a large polymer. Then you can pick out the polymer, and in the process of reconstructing it, identify, say, 10 kilodalton-sized subunits. But to look through an image and pick out a 10 kilodalton piece, that would be impossible. And then x-ray crystallography and NMR are typically imaging structures of about this size down to resolutions in the angstrom range.

All right. So we've sort of gone through the different kinds of things you can see. And then the one last thing I wanted to say is, there's lots of different kinds of cellular structures you can look at in an electron microscope. And I hinted at, say, if you can assemble smaller pieces into larger structures, then you can image them.

So if you can coax, say, a G protein-coupled receptor that you're interested in into a two-dimensional array, then you could visualize these small very interesting proteins as part of this 2D crystalline array. Or in the case of actin, polymerized into a filament. And so there are different ways to reconstruct molecules that arrange themselves into arrays by, say, electron diffraction or in filaments because each unit is related to the unit that's before and after it in the filament.

But the brand of electron microscopy I'm largely going to be talking about today is what we call single particle EM, where you've got these freestanding proteins or virus particles in solution.

And you're going to try to pick individual ones out and figure out what their 3D structure is. And each molecule is independent and not necessarily related to the other ones that are around it.

All right. So we talked a little bit about why an elephant wouldn't survive in the microscope. And that had to do with specimen preparation.

So electrons, because they scatter strongly off of the matter that it's traveling through, if you have gas in your column. Then the electrons are going to scatter off of that before they get to your protein. So the really good microscopes have really high vacuums, and the specimen has to be preserved somehow to survive that.

So you probably wouldn't want to just put your protein in buffer and stick it into the microscope because basically, all the buffer would evaporate, and you'd just have a dried out protein. So you need some way to either, if you're going to dehydrate it, to stain it, which is what we're going to talk about in this paper. Or you can cryogenically preserve it and then keep it at liquid nitrogen temperatures while you're imaging it.

And then I think one of you guys also mentioned radiation damage, that the elephant wouldn't survive being bombarded by radiation. And so at the specimen level, these radiation damage doses are equivalent to an atomic bomb going off if you scale it up. And so basically, this is what you're doing to your specimen while you're imaging it.

And so in your case, you're looking at gold nanoparticles, you had said. And so you can hit a hefty dose on a gold nanoparticle. But on biological specimen, you'd be breaking carbon-carbon bonds, and your protein is rupturing as you're imaging it.

So typically in electron microscopy, we'll only expose the area that we're going to-- and for biological specimens, just as we'll focus adjacent where we're actually going to expose and then expose the area. So the first time that area sees a decent dose of electrons is when you're actually acquiring an image of it.

And I guess one last thing I could point out about this 30 electrons per angstrom squared is, even by 30 electrons per angstrom squared dose on your specimen, a large amount of the high resolution signal is already lost. So the first five electrons per angstrom squared has most of the high resolution information. But it doesn't have enough information in it to actually visualize your whole structure. So you want to give it enough dose that you can see the whole thing but not so much dose that you've destroyed the whole thing.

And I guess one other thing that limits what we can see in the microscope is, if you want to image something at atomic resolution and the stage that's holding the specimen is moving by a few angstroms at the same time, then it's going to be blurry. The features you're looking for are blurred out.

I mentioned at the beginning that this paper was in 2009 that we're going to talk about. So in the last two to three years, there's some new detectors that have come online. And these are revolutionizing the field. So if you look at structures by single particle EM that are at less than five angstroms resolution, it went from, around the time of this paper, there were one or two to now there are tens to even 100 a year in the last-- like in 2015. So there's a whole mess of developments that are responsible for this, but the one that's the most important of these is these direct electron detectors.

So did you guys have a chance to look at, say, the figures of the paper we're going to talk about today? Did you have a chance to look through the methods at all? Did anybody notice how the images were acquired?

So this predates direct electron detectors. So what sort of detectors were used? Anybody notice? Going once.

OK, so some of the images were collected on CCD cameras. And some were collected on film. So can you think of an advantage of one versus the other?

Anybody here into photography? Friends who are into photography? Does anybody still shoot their images on film? Maybe some purists of image quality, something?

Anyway, but why do most people use digital cameras these days? You don't have to go and develop your film, for one. Right? You probably don't even know about having to go and develop, though.

So that's a distinct disadvantage, is the throughput. You can snap 100 pictures on your camera. You don't have to wait a couple of days to see the results.

So the same thing would be an electron microscope. So if you're imaging your specimen on film and then you have to take the film cassette out and go into the darkroom and develop your film and then realize that there was some parameter wrong or somebody didn't change the developer recently, the whole batch would be gone. So throughput with film is low, but the

signal-to-noise ratio and the point spread function of detecting the electrons where they strike the film is good. So the image quality is better with film.

And also, the area that you would expose is bigger on film too. So typical CCD cameras are, say, 4K by 4K pixels. Film would be like 10K by 6K. So you'd have a much bigger area, which means more particles per image. So that would be the advantage of film.

CCD cameras, I mentioned, are a little worse performing. So what they have is a scintillator layer. So it's like a phosphor layer.

So the electrons would come down, and some of them hit the scintillator and bounce off. Some of them will hit the scintillator layer and go through. Some of them will hit the scintillator layer and zig around for a little bit and then give off some photons.

So you can see what the disadvantages are if you're limited to, say, a 30 electrons per angstrom squared dose. If a decent number of your electrons are being lost or not detected accurately-- in this case, you have a point spread about the area where that electron struck, where you're actually picking it up. So this is going to cause a blurring of your high resolution signal.

And then this is just to convert the electrons to photons. Then you would typically have some sort of fiber optic coupling, where you would also lose some signal and also has a point spread. And then this is connected to the actual detector, like what's in your phone, basically.

So this is how some of the images were collected in the paper. Basically, the nice thing about the CCD camera is its high throughput. You can get lots of images really fast. But for the data that was used to generate the 3D reconstructions, that was collected on film.

And then I guess I'll just say one more word about then, these direct electron detectors. Basically, they cut out all this extra business. You basically detect the electrons directly.

So they come in, and actually, each pixel has the ability-- on some of these, the newest top-of-the-line detectors, can actually figure out which quadrant in the pixel the electron struck and can actually count each electron event on each pixel as it's happening. So you have some electronic noise in here, so there's little bits of noise. And on a typical CCD camera, you would integrate this whole signal over time to come up with-- this is down here. So you'd integrate the charge that's accumulated over time.

But in these counting detectors, you could say, here's my threshold for an electron event. And you can filter out all this noise. So you can say an electron struck here, an electron struck over here, an electron struck over here. And so you've got better signal-to-noise, much tighter point spread than you would have in a CCD camera.

And this is what's allowing, say, in the last-- there was a really nice structure in *Science* a few weeks ago of p97, which is a AAA-ATPase with the end domain. So this sort of ties back in, I think, to maybe some of the proteasome stuff that you were doing before. P97's not a protein degradation machine, but it's a AAA-ATPase. And that was at 2.3 angstroms resolution. And basically, what's making this possible is these developments.

AUDIENCE: Can you explain what the fiber optic [INAUDIBLE]?

EDWARD
BRIGNOLE:
Sorry, I didn't label anything up here. So this is your phosphor layer. It's like a scintillator. Phosphor scintillator.

And then down here, you've got your detector. And then the fiber optics is basically coupling the photons that you see here, channeling them down to pixels in the detector. Some detectors just don't have the fiber optics. They'll have a lens of some sort here.

All right. So any other questions about EM before we will dive into fatty acid biosynthesis? OK.

So if I show up this cartoon with lots of different two- or three-letter colored short versions for these enzymes, do these names, I guess they look familiar to you now probably. OK. So this is the scheme for the eukaryotic cytosolic fatty acid synthesis. There's some differences in the bacterial system and the yeast system is a little bit different also.

But basically, to sort of-- I don't know-- to help me remember what all these enzymes do, I like to group them into the enzymes that are responsible for chain elongation and the enzymes that are responsible for chain processing. So basically, the malonyl acetyl transferase-- so in our fatty acid synthesis, we have this bifunctional enzyme that can transfer both malonate from malonyl-CoA onto the carrier protein or acetate onto the carrier protein. In different systems, so in bacteria, they've got a malonyltransferase, and then they have a specialized ketoacyl synthase that picks up the starter unit. So there's some differences like that.

But basically, these are the enzymes that are responsible for collecting the starter unit and the elongating unit and joining them together. And then you've got these three enzymes here, the ketoacyl reductase, the dehydratase and the enoyl reductase that are responsible for

processing this beta carbon. So you've got the hydroxyl, the alkene, and then the saturated chain. And then it goes around again.

All right. So I mentioned that the different organisms have different systems. So in our mitochondria and in plants, chloroplasts and most bacteria have a system like this, where the individual enzymes are the dissociated players. In fungi, some of these enzymes are joined into one of two different polypeptides.

And some of the names here might look unfamiliar. So like this malonyl palmitoyltransferase. So in this case, it's got an acetyltransferase to select the starter unit and a malonyltransferase to select the elongating unit. And then the palmitoyltransferase, which transfers the product back onto CoA. And this is a bifunctional in this case.

And then in our cytosol, we've got this giant monster enzyme that's got all of the catalytic domains fused into one humongous polypeptide. This is what attracted me to this project in the first place, just how bizarre it is to have all of these enzymes all tied together. And then we had this one section of the protein. It has homology to methyltransferases, and we called it the structural domain in the paper.

So there is a bonus question in the handout, which is, where did this domain come from? Why do we have this non-functional methyltransferase domain in our fatty acid synthase? So think about it. If we have time, we'll come back to it at the end.

Then the other cool thing about this enzyme is it has to dimerize to be active. And so you end up with a 550 kilodalton monster protein. So I mentioned that the enzyme's responsible for elongation and for processing. And the cool thing is when you look at the sequence of the protein, you've got the elongation enzymes clustered at the N-terminus, processing enzymes clustered together in the middle, the carrier protein's way out here at the end, and the thioesterase is there.

So there was some decades of controversy about how the acyl carrier protein, which is way out here, would be interacting with the enzymes, which are way over here at this end. And so a model was proposed where the enzymes sort of come together in a head-to-tail fashion. So you would have one going this way and the other one going the other way. But then, that didn't jive with some of the biochemical results, which said that this acyl carrier protein could interact with the enzymes on its own chain. And so there was this controversy in the field, which was

resolved in part by this crystal structure.

So now we could see how the two subunits associate with each other. So one of the chains is just colored in white. The other one's got the catalytic domains all colored in. And so you can see the cool thing here is the elongation enzymes are all clustered together down here like in the legs. And up here, in the torso and arms, you've got the processing enzymes.

And the other cool thing about this structure is if we cartoon in-- we know the acyl carrier protein has to be tethered to the C-terminus of the ketoacyl reductase by a 10 amino acid linker. So that puts the acyl carrier protein right here, and it would be completely surrounded by all of the catalytic domains that it would need to contact. So that's kind of cool.

Oh, yeah and then the thioesterase has a 25 residue linker. And so it would be somewhere around here. And so basically, all of these enzymes are just sitting there in a chamber, and the acyl carrier protein just has to bounce around to the different things. So if I make a cartoon version of the acyl carrier protein with its phosphopantetheine arm docked into each of the catalytic sites, you can see where the acyl carrier protein would have to go on this reaction chamber. And then the same thing would have to happen on the other side.

So now I've got a question for you. What happens if we make a mutant heterodimer. So this is actually an experiment that was done, but if we make a mutant heterodimer where we knock out the ACP on this subunit but leave this other subunit intact. So if the wild type has 100% activity, how much activity would this mutant have?

Any guesses? What would you think? It's firing on one of its two cylinders.

AUDIENCE:	50%.
EDWARD BRIGNOLE:	50%, exactly. So that's all good. That makes sense. What if we do another experiment where we knock out the elongation enzymes in the other reaction chamber? Now what do you think?
AUDIENCE:	Expect it wouldn't be active.
EDWARD BRIGNOLE:	You'd expect it wouldn't be active at all. But the experimental results show that it had about 25% activity. So the only way that could happen is if this acyl carrier protein can elongate with the enzymes from the opposite chain. Right?

So that looks like a pretty long reach, but let's figure out how far that would be. So a 10

residue linker to the acyl carrier protein would be about 35 angstroms. The acyl carrier protein itself is about 23.

Then you have the phosphopantetheine arm, which would be these black things coming off of our acyl carrier protein. So if we draw that to scale from the C-terminus of the keto reductase to the end of the red sphere, it would be about 60 angstroms. So if we draw how big that would be, that's this gray sphere right here. So this is how far the acyl carrier protein can reach, and you can see these are clearly out of range. And actually, even the elongation enzymes in its own side are also sort of at the limit of what the acyl carrier protein can reach to.

So it's hard to imagine what would happen, but you would need to have some sort of conformational change to make these things happen. And I'll point out one other difficulty, which is the access to the enoyl reductase and the dehydratase are sandwiched in the space between them there. And so you'd need to have some other separation of these domains, possibly, to get the acyl carrier protein in there.

So we wanted to look at this EM. Do you know why that would seem like a good idea based on what we talked about with EM so far? Does fatty acid synthase seem like it would be a good target for EM? I mean, there was already a crystal structure of it. So should we have tried crystallography, say, to answer questions about conformational changes?

AUDIENCE: You might not be able to crystallize it in the conformation you wanted?

EDWARD   Yeah, as it was, that was difficult molecule to crystallize. There were crystals of it from back in
BRIGNOLE: the maybe '70s, '80s, but it wasn't until mid-2000s that they had actually gotten-- they solved it initially at, I think, six or seven angstroms. And then this structure was, I think, also not the highest resolution, somewhere in the three to four range. So yeah, you would have to find ways to trap the conformations that you want, lock it in, and cross your fingers to get crystals.

Why look at it by electron microscopy? Is it big enough? It's 550 kilodaltons that you can see individual molecules.

Possibly we could even see them in different states, and we might even be able to perturb those states if we threw in some substrates. Then we had a whole panel of mutants that our collaborator had meant. The experiments that I had described about knocking out the ACP in one chain versus the elongation enzymes in the other, there was a whole battery of mutants

that we had available to us.

All right. So to do the electron microscopy, we need to put our protein on something that we can stick into the microscope. And typically, that's a metal mesh with a cart that's supporting a thin carbon film. And then we stick the protein onto the thin carbon film.

So this is about three millimeters across, this little grid. You can put about five microliters on it. And then to get a good dispersion of particles on the grid, you need about 15 nanograms per microliter. If you go too much above that, you get protein everywhere, and you can't pick one particle from another. And if you go much below that, then you have to collect lots and lots of images to get a few particles.

So this is sort of the sweet spot, in the 15 to 20 nanogram per microliter range. This is one limitation for EM, the concentration dependence. So if you have a molecule that falls apart, it has a high Kd and it falls apart at these concentrations, that could be difficult to work with, for instance.

All right. So there's a couple of different ways to prepare specimens. I think we already talked about staining the specimens or cryogenically preserving them. So the way that would look like for a stain experiment is, you've got your thin carbon film, you put your drop with your protein molecules on it, you blot off the excess solution, replace it with a heavy metal salt solution-- typically a uranium salt-- and then you let it air dry. And the specimen is then embedded in this heavy metal.

And that's why we call it negative stain because what we're imaging is, you've got your protein molecule, and it's embedded in this dense stain layer. What's scattering the electrons most strongly is the material around your specimen. And so you're imaging where your protein isn't, basically. Or the stain excluded area is what you're imaging. So what you have, in this case, is a dark background, and your particles look light.

The other way to prepare specimens is to cryogenically preserve them. So the first part starts out the same. You would put your proteins on the grid. And sometimes, you could have a grid that's got little perforations in the carbon, so you actually would have your protein suspended in these perforations when you blot it.

And then you plunge it into liquid ethane that's cooled to just about to liquid nitrogen temperatures. Here is a picture of the dewar with the liquid nitrogen. And then there's a little

cup in the middle with the liquid ethane.

So why not just plunge it directly into the liquid nitrogen? Does anybody know? Does anybody do rapid freeze quench for any of your experiments or anything like that?

So have you ever messed around with liquid nitrogen that any splashed onto you? Did you get burnt? No.

So the reason is liquid nitrogen has a lower heat capacity, so if it touches you, basically, there's a layer of gas between the liquid and your hand or whatever it spilled on. But with liquid ethane, the heat transfer-- basically, this grid will go in there, and it'll freeze so fast that ice doesn't have a chance to form crystalline ice. So basically, everything is, on a microsecond scale, frozen.

So now you've got this amorphous ice with your protein embedded in it. Can you think of some advantages or disadvantages? If this gives you something preserved and it's happy in its buffer, why wouldn't you always use that? Why would you use stain? Can you think of some advantages? Maybe the obvious thing, why don't I ask you for some disadvantages. Why would you not want to use stain?

AUDIENCE: The stain used could possibly disrupt your specimen.

EDWARD BRIGNOLE: Yeah, and that does happen. Sometimes people have to play around with different stains. The uranium salt stains, the uranyl acetate, for instance, is a low pH. And if you try to pH it, it crashes out. So if your protein isn't happy in that low pH stain, that could be a problem.

Also the stain layer is dried, and so your specimen is dehydrated and dried out here. And typically, that-- I drew my specimen like this. But when it dries out, it flattens out like this. So that's a disadvantage.

What about contrast? So if this is my amorphous ice, my water layer with my protein, do you know what the difference in the density of protein versus an aqueous buffer is? They're pretty closely matched, actually. Protein's like 1.2 or something like that.

So basically, you have pretty weak contrast in a frozen hydrated specimen because here, you're looking at the difference in density of your protein versus the buffer around it whereas here, you're imaging the difference between the density of, say, your protein and uranium. So you get a lot better signal here. But you've got some specimen distortions.

And so we basically just went through these. There's one other advantage I'll mention to sticking your protein onto a carbon surface as opposed to freezing your protein in a hole. And that is that most proteins tend to have a preferred orientation.

Many do, and in the case of fatty acid synthase, it's sort of this. It looks like a headless person that's got arms and legs. It'll very rarely hit the grid and stand straight up. It usually falls back onto its back. And so in some circumstances, that could be an advantage, and other circumstances you would actually want to have many different views to make a 3D structure.

So I listed that both as an advantage and a disadvantage, the preferred orientation. Depends. You could use it to your advantage. In other cases it would be a disadvantage. All right.

So you said you used an FEI microscope. It might have looked like this one.

AUDIENCE: Yeah.

[INTERPOSING VOICES]

EDWARD BRIGNOLE: OK. Yeah, this is an F20. This is the microscope that all the images in the paper were collected on.

So there is a specimen port on the side. The electron source is up here at the top. The column with the lenses and apertures in it is here. There's a phosphorus screen here that you can look at through the binoculars to see what's going on. There's the knobs that you can use to control the microscope, focus, move the stage around.

And then the camera is right here below the column, right where you can knock your knees into it when you look in here. Yeah, I mean, you put a half a million dollar detector on there. And you can knock your knees into it.

Actually, the newer microscopes these days actually look more like giant refrigerators. And basically, all of this is housed in this environmental chamber, and you operate the microscope from the room next door.

So we put the grid in the microscope. At low mag, you can get an image like this. Little higher, just zooming in on one of these squares here, you can get an image like this. This is negative stain specimen so there's little chunks of stain around.

If you ever happened to do some negative stain experiments, I usually like to look for areas that have this smudgy appearance. It looks like little pencil lead shavings that somebody wiped their hand across. That's usually a good sign. And then if you zoom in another tenfold, you can get an image like this. And if you look carefully at it, there's all the individual 550 kilodalton fatty acid synthase molecules.

So how do you get any information out of that? Any ideas? You can pick out the individual molecules here. If you squint at it, can you maybe make out the legs and arms, the processing portion, and the elongation portion? Maybe?

OK, it's tough. Does anybody here do spectroscopy?

AUDIENCE: No.

EDWARD BRIGNOLE: No. So based on what you know, electron microscope images can have potentially high resolution information in them. But you're limited in dose you can apply to the specimen before radiation damage becomes a problem. So what we have is a signal-to-noise problem here. You've got high resolution signal buried in lots of noise.

It's like having a low exposure image of something. What could you do to boost your signal? If you're going to take a picture at night, what would you do?

You need a really, really long exposure. Right? But you can't take a really, really long exposure. So what would be a different way to do it?

Say like, in the case of spectroscopy, if you had a sample that's damaged every time you stuck the cuvette in the area, but you could say, take a cuvette and take a spectra, take another one, take a spectra, take another one, take a spectra, and you can average lots of them together, that would boost your signal-to-noise. So that's what we have to do here. We have to extract all these particles out and find a way to average them together.

So if we put soccer players on a EM grid-- if any of you are soccer fans-- and you collect an image of them. You get this noisy image like this. In the computer, you can go through and pick the particles out. And the computer can do its best to line them up for you. And if it does a good job, and you get lots and lots of particles, when you average them together, you get your high resolution signal out.

So that's all fine and good, but not every protein is going to land in exactly the same

orientation. And in the case of soccer players, you probably would have a hard time finding soccer players that are in exactly the same conformation every time you image them. So in this case, a soccer player might prefer to kick with his right foot or left foot or might have his right arm or left arm up or down. These are just a couple of different conformations maybe that you would observe.

So now what do you do? You've got these averaged together, and you're like, I got an insect. Does anybody here-- have you looked at, say-- you could do this by spectroscopy. But it sounds like nobody here does spectroscopy.

So you've got different sorts of things that you want to categorize, basically. So say, sequence alignments, that would be analogous to this, where you've got sequences that you've lined up. Here we've got images that we've lined up.

And then what would you do? You'd look through columns of residues or the computer would do this for you and say, this cluster of sequences all have these particular residues. So I'm going to put them into one bin. And these have a different sequence, and I'm going to put those into a different bin.

The computer can do the same thing in this cage, basically. It'll look at these images and say, some of them have a density here, and some of them have a density there and split them up based on differences in the intensities of these pixels. If this is a dataset of 100 images, then you split it. Now you've got 50 and 50. You might have 25, 25, 25, 25 if everything's evenly distributed.

And the one thing you'll notice as you split things down further, you're averaging fewer and fewer particles together. And so your signal-to-noise is getting worse and worse. So you can split these down.

And the way I typically do this is a little bit empirical, but I'll split it and then split it some more and split it some more and look till I get to a point where I'm not seeing anything new. Because if you split this image more, it's basically going to be split on the basis of noise because there's no other conformational change. The same thing would go for orientation. If you put this on the grid and it landed in three different orientations, you would want to separate things out using the same strategy. All right, so this gets us to the averages, like the averages that you see in the fatty acid synthase paper.

There are also 3D structures in the paper. So how do you go from information like this to a 3D structure? So I mentioned one way is if you've got lots of different orientations of the molecule on the grid and you've got each one of these averages is a different view, you can use the computer to try to put those different views together to come up with a 3D structure. And there's lots of ways to get that wrong. In this case, we've got a preferred orientation where they're all lying on their backs. We don't have lots of different views.

So what would you do instead? Yeah?

AUDIENCE: Get it from the sides?

EDWARD BRIGNOLE: Yeah, exactly. So that's the thing to do. So now you've got this stereo view of your molecule where you've got a tilted view and an untilted view.

An extreme example of this would be if every particle-- let's say if you're looking at cells, no two cells are the same. You couldn't do these averaging methods, but to get a 3D reconstruction of a cell, what you would have to do is take the stage and tilt it by a degree, tilt it by a degree, tilt it by a degree, go up as far as you can one way, and then do it again the other way. So you'd have up to maybe plus or minus 70 degrees. And that would be equivalent to the way a CAT scan or something might be done, where you've got images of your broken leg or something like that from all the way around. And then you can have the 3D reconstruction of it.

Right. So we've got these two images of our specimen. They're related to each other by some tilt that you know. If you take out these particles from this image and you line them up, that tells you what view you've got of them in this.

So take, for instance, this molecule here. If you have to rotate this 90 degrees clockwise, that tells you you're looking at him in the tilted view with his feet up in the air. And this one here had to go 90 degrees counterclockwise. That means that you're looking down on him in this tilted view with his head up.

And so you can take the alignment information from this image and apply that as a projection parameter for these images. And so you could take now these tilted views of these soccer players, and you know which view they are. And you can come up with the 3D reconstruction that way.

So this is actually a fairly old method. I think it's still widely used and very elegant because it's

basically just two images of the same thing. And then you can get a reconstruction that's pretty easy to get out.

There is one disadvantage to this approach, one major disadvantage, which is that you can only tilt the stage so far. So if you could tilt the stage up to 90 degrees, then you would have views exactly all the way around. And you could have a reconstruction that's fully complete.

In this case, you can only tilt to 70 degrees, and so you've got a missing cone of information in the reconstruction. And so basically what that means is you've got better resolution in x and y than you do in z. Yeah, sure?

AUDIENCE: Is there some graphene packet thing that came out of the [INAUDIBLE] lab where it's this packet filled with solution that you shoot your EM at the--

EDWARD BRIGNOLE: Yeah, to keep your protein hydrated, basically. So you encapsulate it in some sort of graphene tube or something.

AUDIENCE: [INAUDIBLE] exactly [INAUDIBLE]

EDWARD BRIGNOLE: I vaguely remember seeing something like that. I think there are groups working on things like that, but it's not widely adopted or used yet. But yeah, there are some pretty exciting things like that that might allow you to directly image your molecule while it's tumbling in solution, isolated from the vacuum on the microscope.

AUDIENCE: But then there's some limit with what your computer can reconstruct. I mean, it's just infinitely many tumbling orientations or something.

EDWARD BRIGNOLE: Yeah, it's a tough experiment to do. The other issue is compressing all your dose into a pretty short amount of time so that you basically obliterate the molecule in this field of view but you capture the image of it faster than it's tumbling, say, or something like that. So I don't know. Maybe it depends on its tumbling rate. Yeah, something like that.

But yeah, I think it's an exciting time for EM right now because now there's new detectors. There's actually some other examples of advances that I put on that slide that allow us to look at smaller things, potentially specimens that are still hydrated. Yeah. I don't know if my email address is on this, but if you come across that paper or see anything like that, feel free to bounce it to me.

All right. So through the methods that I just described to you, basically from the fatty acid synthase image that we looked at a moment ago, we can sort out some images. And if you look at this image, it looks a lot like that crystal structure I showed you earlier. There's the legs you can clearly see in the average and the processing enzymes in the upper proportion.

Then we could sort out a whole bunch of other different classes. And these puzzled us at first. They're kind of fun to look at because we thought maybe it's winking at you. It's got one eye open and one eye closed or like the other one. And we described these as different views. This looked like it had a pirate's hat on or something.

But one of the other things that puzzled us at first too is, this lower portion of the structure here, it looks like maybe we were getting some sort of proteolysis and these malonyl acetyl transferases at the legs, we thought maybe they were getting cut off. So we were relieved when we generated the 3D reconstructions then, that these actually weren't getting cleaved off. They're just rotated 90 degrees on the grid and coming out towards us.

I'll say a quick word now about-- have you read any crystallography papers yet for class when they talk about resolution of the structure, what kinds of things you can see in the structures? So in crystallography, you have a defined resolution limit based on the highest angle of scattering data you collect. So it's defined in the experiment.

In EM, we don't have that. We just have images. So the way we calculate resolution in EM is we would take our dataset-- so the data that went into this reconstruction here if there is 1,000 particles-- we'd split it into two subdatasets, one with randomly selected 500 particles and the other with another randomly selected 500. And we'd generate a reconstruction from both, from each of those. And we'd compare those, the reconstruction of this half of the data to the reconstruction from this half of the data and see how similar they are to each other. And that's how we would figure out resolution by EM.

So there's some problems with that. Can anybody think of one way this would be biased or any way that this would be biased if you just take your data, split it into two halves, reconstruct it, compare the two?

For one, it's sort of like if you have one person do the experiment and they do it again, but nobody else can do it. There is a bias in-- you're taking the exact same approach to initialize both of these experiments. They're both going to converge to the same local minima. So you could be precisely wrong and have a false high resolution.

So typically, what you'll see in EM papers is a curve like this, where you basically are comparing the two half reconstructions to each other, one from first half of the dataset, the other from the other half of the dataset. You compare them, and then you look at, say, at low resolution, how similar are they. Add a little bit higher resolution, how similar are they? If you go to, in this case, 20 angstrom resolution, how similar are they? And in this case, they've got a correlation of about 10% or 15%.

So in the case of this paper, the way we reported the resolution was when the correlation between the two halves of the data fell to about 50%. So we reported a resolution of about 30 angstroms for these structures, but like I said, that doesn't necessarily mean they're right. And in our case, the advantage we have was that there was a crystal structure. So if we just dropped the crystal structure right into the EM reconstruction, that looks like a pretty good match. And one thing that we didn't do for this paper, but people sometimes do is, instead of comparing half of our data to the other half of our data, we could have compared our data versus the crystal structure and come up with a similar curve to compare our data to this high resolution data and see where things fall off.

So in these reconstructions, there's lots of different conformations. You can see the lower portion swinging back and forth, the upper portion twisting relative to the lower portion. And then, you see this other conformational change. So we take the arms off of the end and look at what's happening in the middle here.

The enoyl reductase and dehydratase are rotating like this relative to each other. So one side opens up while the other side closes. They sort of cross over each other like that. And so when one side rotates, they sort of rotate, but then one side tightens up while the other side comes loose.

So we wanted to know how these related to catalysis. So I guess if you've looked at the paper-- I'm going to try to speed up here a little bit-- what we did was we looked at some mutants in the presence and absence of substrates to look at how this-- we've got all these different conformations-- how the frequency that we see different conformations changes. And I'll sort of cut to the chase. I think in the paper, there is histograms, but I think the pie charts are a little more telling.

But basically, what happens is if you add substrates, the conformation that becomes most prevalent is the one that's represented in blue here. So if I go back, you've got these,

basically, four different categories where the lower portion is perpendicular or parallel and then whether the upper portion has this asymmetric appearance or not. And the one thing that jumps out right away is that add substrates and you get lots of asymmetry in the upper portion and the lower portion in parallel with the bottom portion.

All right. So why would that be? So if we look at one of these conformations where the lower portion's parallel to the upper portion, the upper portion has this asymmetric appearance. What might this reaction chamber be good at doing?

So these enzymes in the lower portion, the ketoacyl synthase and the malonyltransferase, they've come up close to where the acyl carrier protein would be. And at the same time, this side is closed off, so it would have a harder time doing processing on this side. So at the same time over here, this side's opened up. The acyl carrier protein can easily get in here to do the processing, but these enzymes over here are out of reach of the elongation enzymes.

So what this means is one side could be elongating while the other side's processing. And then the structures are symmetric. So if you flip it around, then this side, once it's done elongating, could process. And then this side could elongate. So it's kind of cool because it can sort of balance out what it's doing from one side to the next.

And then we have these confirmations where the lower portion is perpendicular. And remember, at the beginning I had said that we know that this acyl carrier protein can elongate with the enzymes in the opposite portion. So because of the symmetry in the system and also the resolution of the structures, we can't tell the difference between whether the lower portion is flipped 180 degrees relative to the top or not because it would look the same.

But the fact that we can see these go 90 degrees is suggestive that it could probably unravel and go the rest of the way around. So in the crystal structure, the way it had them is, they were coiled like this. And so it's pretty easy to imagine that they would just uncoil.

One other line of evidence that I think is telling-- so our collaborators made a mutant that has all of the active sites and the acyl carrier protein knocked out of one subunit. So there is one subunit totally wild type and the other subunit that's totally dead. So the interesting thing about this mutant is it has to do the condensation reaction, the elongation, in this conformation, sort of crossed over.

But to pick up its starter unit or elongating unit, it has to coil back around for this. And we know

based on the rate of this enzyme, that this probably happens about 100 times per minute. There's a functional catalytic event that happens 100 times a minute. So it probably is sampling these much more rapidly. Sure.

AUDIENCE: So this isn't compensating for when you knock out one half that's naturally [INAUDIBLE]

EDWARD
BRIGNOLE: Yeah, that's what we think. It's sort of naturally sampling both sides. It's interesting to think about because let's say this side picks up acetate and-- well, let me think about this for a second. But if one side is ready to elongate and the other one's got a starter unit, and this side over here is already loaded with-- so basically, this will pick up an acetyl group and transfer it to the ketoacyl synthase.

And then it comes back, and then it picks up an acetylic group again. So then it would be stuck because it would be trying to extend an acetyl with an acetyl. But it needs a malonyl.

So what it could do is flip around, sample the other side, which might have a malonyl group, to continue on on that side and then come back around. And it could transfer the acetyl group on. So it'd allow maybe one way that it doesn't have to necessarily go backwards, though the malonyl acetyl transferase can function in reverse.

So that's another thing, is it won't get stuck if everything's all loaded up with acetate because the malonyltransferase will run backwards to cut the acetates off. But this would be one way to maybe not. It wouldn't have to necessarily rely on going backwards to get unstuck. It could just twist around.

All right. So I think we'll just finish up with a quick movie that shows these different conformations. So the bottom parts picks up a substrate, elongates it, goes up here to do the processing. Meanwhile down here, it's elongating. You can see how up here in the upper portion, the separation of where the dehydratase and enoyl reductase is so the acyl carrier protein can fit in. Here we go.

So then there was that bonus question at the beginning. I know you guys probably have to run. So the bonus question at the beginning is, what is that structural domain that's-- the methyltransferase, where did that come from. Any ideas? You guys have talked a little bit about polyketide synthases yet?

A lot of them have domain architectures identical to our fatty acid synthase with functional

methyltransferases. So it seems like we probably picked up our fatty acid synthase from something like a lovastatin synthase or something like that. So it's interesting to think about. And then we didn't need it, so it's now not functional.

All right. Cool. Thanks guys.