Berkeley Herbert Wertheim School of Optometry & Vision Science

Young Vision Scientist Podcast, hosted by Reem Almagati, OD, MS

Episode 1: Computational Neuroscience and the Problem of Representation with Christian Shewmake

Published: December 7, 2022

[intro *Familiar Skies* by PNFA]

Reem Almagati: You're listening to the Young Vision Scientist podcast. My name is Reem Almagati. This is an initiative by the Vision Science program here at Berkeley at the Herbert Wertheim School of Optometry. In this podcast, we talk to current PhD students and postdocs, where they share with us research that they're doing currently at Berkeley.

Today, we have Christian Shumake. He's a rising second year here. How are you Christian?

Christian Shewmake: I'm doing well, Reem, thanks!

Reem: Thank you so much for being our guest for this very first episode. I want to start by asking you to talk about yourself and your background. What were you doing before coming here to Berkeley?

[*Familiar Skies fades out*]

Christian: Cool, thank you so much for having me. This is such an incredible program and I'm excited to be a part of it. So I got my start in biomedical engineering., I was super interested in math and physics and computer science, and really couldn't choose between them. So I chose biomedical engineering. I was sort of obsessed with the brain wanting to understand, uh, neural decoding, things like that. You know, it's the more flashy amongst the BME students.

So while I was in undergrad, I studied with people who were working on bringing computer interfaces and all kinds of other cool things about the mathematics of nonlinear dynamics and control theory, So I think I, I grew a lot in that time in my appreciation for mathematics in particular.

And then after that, I joined this wetwear company in Berkeley, which was building these chips, uh, pleading neurons on these like micro electro arrays, with like microfluidics. And the, you know, sort of dream was to build biological computers, right? These machines which could compute based on neural interactions if we only knew how to pattern them and, you know, sample from them and decode them. And all that sort of started with the problem of olfaction. So we were making these olfactory cameras by decoding activity in these neurons.

Reem: Can you talk more about olfactory cameras?

Christian: Yeah, sure.

Reem: That sounds very visual, but it's not.

Christian: Yeah. Well it's funny, it's, you know, it's very similar to the problem of vision in a sense, which is that you've got, you know, some big world out there with all kinds of things happening that we can't immediately detect, and so we proxy it with the things we can measure and try and make the best inferences we can.

So the olfactory system also is subject to a lot of these same kinds of information, theoretic limits on sampling and so on. So in our case, we were essentially giving each neuron the ability to detect certain compounds in the air with these specialized proteins, and then their pattern of activity would tell us something about the various compounds in an environment and their concentrations. So you can use them to detect things like bombs or to detect various diseases, for example, that have markers, biomarkers in particular.

Reem: During your time working with olfaction, were you looking at cue combination and did you work at all with visual? Because when the brain is presented with any kind of uncertainty, it can combine multisensory input as a form of cue combination to infer more about the real world.

Christian: That's a brilliant point. We should have hired you as our, uh, as our engineer. Yeah, at the time, no, it was very simple. You know, we really were just trying to get, get the, you know, the biological components to behave in the way, you know, you sort of, it's already a bit of a challenge, an engineering challenge just to get them in that configuration. But you're right, I think many of these fields, like brain-computer interfaces, would benefit from cue combination.

If I can't decode what's going on from your brain data alone, if I strap a camera to you, then now maybe I can make better inferences. So that, I think that's brilliant and I think, uh, it'd be cool to see work go in that direction.

But yeah, I mean, so working on, working on this problem of neural decoding, and these sort of synthetic systems got me really excited about just how you encode information and populations of neurons in the first place.

So essentially, when I left this startup, I went to go work at Wash U in St. Louis for my master's, just focusing more on mathematics and dynamics and trying to get my foothold into the fields that later fed what I'm working on now.

Reem: And how did you move from that to the visual system?

Christian: Yeah, so, a long path from there to then working with Bruno Olshausen here. You know, sort of fell in love with this problem, this idea really of, you know, manifold structure. This idea that we have these high dimensional curved surfaces, which our neural data is sort of moving along on, you know, the neural activity is moving along on.

And, you know, I couldn't put it down — so I was working on other things for a few years after the master’s and I was just reading papers at night and, you know, just fell in love with it. So I was, I was really fortunate actually. I happened to bump into Bruno Olshausen at a talk on Berkeley's campus. And I just really dove into all of his work and I felt like vision had all this, all the important components of what I was interested in with all these really neat problems of decoding and inference, but also dynamical systems and control, and, also touched into these artificial systems. So I think I was very lucky to find that intersection and it's been a blast getting to work with him.

Reem: I want to give a suggestion to our listeners: You can look up YouTube videos or lectures of Professor Olshausen. He has a really good way of explaining things and you don't really need to have a background knowledge of vision science. I think he can break really complicated problems into simple words. And it's always a blast to listen to him talk.

Christian: Yeah, we really are blessed. I mean, we have some special humans who've gone that route and have discovered some very deep, deep insights, and you can actually bring that up so that anyone can feel amazed about the natural world. And Bruno is just very special in that regard.

Reem: Yeah. So you've been here at Berkeley for exactly a year now at the Vision Science program.

Christian: That’s right, crazy!

Reem: And well, the first year was a lot of exposure. And I want you to share something that is not related to your research now, but something you saw maybe during lab rotations or, you know, lab visits, that was really fascinating.

Christian: You know, I was actually sort of new to vision and even to visual neuroscience, so I knew very little when starting the program about the visual system. Uh, like surprisingly little and came more with just a passion for this problem.

So I think the thing that's been just most incredible to me is I get to interact with faculty like Suzy Fleiszig, who works on these like cellular mechanisms at the level of the cornea and, like, how that plays a part in the larger story of the optics of the eye and immunity and how, um, our body's immunity — and how all these systems interact to actually, like, support vision.

So seeing just, like, a lot of conversations with her talks with Austin Roorda, looking at the retina and actually like studying — it's crazy, we can image these things and study function at that level. And to talks by people like Michael Silver that are actually looking at, oh, like, what are these, you know, either early vision like Bruno works on — so like V1, V2 and all the way up to like higher vision and visual attention.

So I think what I've been most blown away by is, like, how there's this massive map of what it is to do biological vision. And we have people who are just really deep in each of those areas. So I think I just have a lot more perspective than I had when I started that completely reshapes what questions I ask and how I ask them.

Reem: I remember being so excited and looking at you and seeing how blown away you were when we visited Dr. Roorda's lab and just seeing in real time, him sitting there, and we can see the individual photoreceptors in real time. That was really fascinating.

Christian: Unbelievable.

Reem: Unbelievable.

[music interlude *Familiar Skies*]

Reem: So let's talk about the topic that you're doing in your research. I know you're doing computational neuroscience and just, if you want to talk in general about the area that you're researching.

Christian: Yeah, totally. So I think, you know, there are many different problems in vision that need to be solved — all of them extremely interesting. The one that I'm sort of most captivated by is the problem of representation.

So representation is sort of asking the question about, you know, if I have things out there in the external world, you're not actually experiencing, let's say 3D objects. You're not sensing, directly, their 3Dness. You know, you are taking these impoverished measurements from the world through, like, light bouncing off of them and hitting your eye, or sounds that you hear. And somehow we have to take all these different, sometimes conflicting, signals and try and make sense of what's out there.

So I think the part of that whole sort of constellation of sub-problems that's most fascinating to me is — somehow you have neurons, like, wet bags in your head that are connected together in such a way that you can actually store and represent entities in the world and manipulate them in your head in a way that's consistent with the world.

And I don't think we have any theory that really explains how that's achieved, even theoretically sort of models of this, but also in like, wet living systems. So this problem of representation it's not just relevant for vision, it's also relevant for people who study, you know, motor control.

How do I represent the states of the joints of my arms? How do I represent, sort of, words? You know, what is the space or, like, the mathematical space that we can think about that the brain is using to embed those signals?

Reem: Why do you think it's important or what is the value that understanding representation adds to our understanding of vision?

Christian: Oof. Yeah.

Reem: And just to clarify a little bit on the question… a lot of people are really interested in understanding visual representation, you know, to use in AI models and computer vision, but that's not quite what you're doing. So in terms of understanding the biological visual system, what is the value that understanding visual representation adds to that?

Christian: Well, I am actually very interested in both, in sort of biological vision and artificial instantiations of visual perception.

And I think, I guess at a higher level, maybe by analogy: Suppose it's, you know, the 1600s and you see these birds flying in the sky. Or you see leaves that are cast about in the wind and never seem to drop. This would've seemed like magic at the time – how could we have, you know, the fact that we have come from seeing these sort of mystical, floating creatures to being able to design very precise flying machines of different varieties, I think, boils down to us understanding the principles of flight.

So the fact that we understand Bernoulli's equations, we understand, like, pressure and how lift is generated. And once you understand those principles — like, what are the principles that govern flight? — I think you can then better design flying machines. So, I think by analogy, like what does theoretical neuroscience or representation have to say about vision?

I think if we understood the principles of what the visual system, of what it was, what you needed to solve in order to solve visual perception, then I think we would better understand both our brains and, like, the wiring that was chosen — specific algorithm chosen by the brain instantiated in neurons — but I think we would also better understand how to design, you know, perceiving systems, like convolutional neural networks, for example, or other vision models.

So I think, yeah, I think very concretely, innovations in new theories and new directions to explore result in new architectures when it comes to, you know, the problem of perception in computational systems. But also new questions that we might ask about what's happening in V1 or the "top down, bottom up" interactions in the visual system.

Reem: I remember during my master's program, we had a guest speaker from Vienna and in their lab, they're doing work on retinal prosthetics… Even though we understand the representation of visual information in the retina, somewhat — we can't really feed the brain what it needs because in the visual system, you have signal coming from the retina but also you modulate the signal. So what is the kind of signal that we need to send to the brain to be able to kind of restore our vision when we have the neurosensory retina damaged or diseased?

And I just want to know your thoughts on this and how that relates to computational visual science.

Christian: Totally. That's a fantastic question. And I think this is a deep, I mean, a really deep and interesting problem because there's sort of this handoff at every point. There's, you know, photons have done the job of encoding 3D objects and their surface properties into photons coming and streaming towards the retina. So that's one conversion that's already happened.

So what's this new format? Well, the retina receives that format. And what are the properties of that signal when it comes in? How is that information stored and represented?

And even, let's say, when you have all the circuitry of the early… you know, the retina itself — it's amazing that we have people even in the department who have actually mapped what a lot of these circuits are doing in terms of how they transform information into some new format that they can be, you know, sent down the optic nerve.

So, yeah, I mean, this is… it becomes really clear, I think, like, how important information theory is in that problem of how much information are you losing through this process and are you sculpting some new representation that's good for some downstream task?

[music interlude *Familiar Skies*]

Reem: Moving from here, I want to talk about the current research you're doing. What are the specific questions you're attempting to answer? So, currently, what research are you doing?

Christian: Yeah, great question. So maybe I can break them broadly into two categories. The first category is more mathematical questions, not necessarily specific division about how we represent, let's say, something called group structure or manifold structure in data. And that turns out to be relevant for vision. We'll talk about that in a second.

So that's sort of, you know, broadly developing machine learning methods for discovering or using symmetry structure or group structure, which we'll talk about, from data for certain downstream tasks, and applying that to machine learning problems — thinking about what that says about neural encodings. That's sort of one direction.

And the other direction is looking at how these ideas from group theory or Riemannian geometry inform our understanding of V1 or hierarchical interaction in the visual system. So looking at sort of models of sparse coding that incorporate group structure, in that sense. So those are sort of the two general categories.

So maybe within the first category, with my collaborators, Sophia Sanborn, Chris Hiller, and Bruno Olshausen, of course, we just submitted a paper to *NeurIPS* — which is *Neural Information Processing Systems*, it's a machine learning conference — that's on something called bispectral neural networks, which essentially is trying to pose a new way of thinking about the quote "invariance problem," or it's also been called the "problem of universals."

And so the idea behind it is, you know, our visual system is really amazing in ways that our artificial systems still struggle, which is — suppose I take this, you know, this pen, let's say, and I … no matter where this is in space, if I take this object and translate it, you'd still understand that it belongs to the category "pen." Or I could rotate the object and you could say, "Yes, Christian, you're still holding a pen, why are we doing this?"

Reem: Or if it's presented in clutter and it's part, or a lot of it is covered…

Christian: Yeah. There's some really phenomenal — this huge range of transformations I can make to this object for which you'd still consider it a pen. I could change the color, the size, some of the shape, you know, and this sort of — you can imagine that space of transformations on the pen and sort of this collection of pens that it describes kind of define the category of "pen" itself. So it's not just this one instance, but sort of that collection and what transformations are valid transformations of a pen.

So the… right now it's really hard to design artificial neural networks, which have this kind of robustness to this wide variety of transformations, which means that you have to train them on large amounts of data that include many different instances of this thing called "pen."

Reem: And that's when you incorporate Bayesian inferences, right?

Christian: Yeah. So Bayesian inference plays a huge role there and also group theory. And there are different ways of tackling depending on the strategy you want to use for the problem.

So in this case, we look from the direction of group theory. So essentially there's a way in which you can define a neural network layer, that we call it, or just the transformation of an image of some object using ideas from group representation theory. Which makes it such that if you were to translate this pen anywhere, the output is sort of constant, the output doesn't change, which is good — Okay, it shouldn't change what the category is because you're transforming it. But also it's unique.

So it's unique: Only “pen” will have that representation in the output space. So those numbers only change if you change what the object is itself. So I guess maybe the main contribution of the paper is — and people have shown that you can use these sort of group theoretical ideas to get this invariance — but I guess what we've shown in the paper is that, just from data alone, without knowing what group actions, what is it, translation, rotation, and so on, without knowing what group actions are relevant to the problem, you can actually discover symmetry structure just from data. So you can discover that translations don't change the category or rotations don't change the category and actually know that that's the transformation: rotation or translation. So that's sort of the main result of the paper.

Reem: And, it's training a neural network?

Christian: Correct.

Reem: That leads me to this question of what is really special about the human visual system… Is that you can present transformations of an object or presenting it in an ambiguous way or viewpoint that you have never been exposed to, so your brain has not been trained to as a neural network, but we seem to not have any issues with that. And I think that's really fascinating how we can get to that.

Christian: Truly. I mean, this is really remarkable. And I think it's also… I mean, people in applications struggle with this. The self-driving car community is racing for better perception models that are more robust in the ways you just described… Because you can't see, you know, you can't collect all the data, which gives you every instance of, let's say, "pedestrian" in every pose in all lighting conditions. And if someone trips, you know, on the side of the road, then maybe your classifier is going to say that it's something else, let's say, "speed bump" and not "person." So, it's really important, I think, and really remarkable that our brains have this ability.

Reem: I want to talk about neural networks that are used as classifiers because the problem of vision is not just to classify objects.

Christian: Oh, completely.

Reem: And I just want to know your thoughts on that because I know this project you're… most applications of computer vision are, you know, mostly concerned with classifiers because they have so many practical applications. But vision is not a classifier.

Christian: Yeah, absolutely. I mean, this is… I even was sort of misled about this prior to joining the Vision Science department. It's one of those things that's really changed how I think about what vision is in the first place.

And it's also why I think representation is such an important problem. Because let's say the task you're interested in doing, in solving, is classification: "What is this object?" Then, if you have some representation of that object, which is independent of its pose, which sort of factors apart its color and other things, that's sort of based on its shape, then you can answer those kinds of queries rather easily.

But suppose your task instead is tracking — "I want to track this object," or, let's say, you want to understand something about the visual scene; make predictions about what's going to happen in the visual environment.

These are very different tasks, but the same representation may support a whole variety of those tasks. So the idea that if you first form a good representation of the world, you can use that for problems in, let's say, memory of "how do I want to store this efficiently" or… "how do I want to perform visual analogies?" Like, how this object transforms into the next one or things like that.

Reem: I really like Bruno's approach to vision: What are the problems that evolution tried to solve for us to have this particular visual system, which I think is a really good way to think about vision.

[music interlude *Familiar Skies*]

Reem: You said you just published a paper and..

Christian: Submitted, submitted!

Reem: Submitted, okay. Hopefully fingers crossed for it to be published!

Christian: So in our paper, I think many, many readers aren't familiar with this object called the bispectrum. Many people have heard of a similarly inspired object called the Fourier Transform, which are related to two point statistics, is what it's called, and the bispectrum is sort of a third order version of the Fourier Transform, which has all these really nice properties with respect to transformations.

So because that's something that most people don't have lived experience with either in, you know, living or in, you know, mathematics, a lot of the paper is dedicated to just explaining — What are the properties of this object? Where did it come from? Why is it interesting? So part of the background is sort of explaining the mathematical background and trying to make that accessible.

And then we, sort of, proceed to show, "Okay, here are the properties of the bispectrum. So this is what you'd expect. Here's the problem itself of representation. Here are the, you know, we want the output of this thing to be invariant to group action…" and you can sort of list these properties.

And basically the work is just showing here are the guarantees this mathematical object makes, and here's how we show, like, "Yes, it does satisfy these guarantees on data in a format where it's completely learned." So, we don't actually, you know — if you're a mathematician, you would solve what the bispectrum coefficients should be; you'd write them all down, you'd code them into your computer, and then, yes, you'd get this property, of course, because that's the sort of — mathematics tells you that.

The really remarkable thing is, we show that you don't have to be a mathematician. You just have to have data, which has transformation structure, and it actually learns the coefficients that end up resulting in the bispectrum. And we basically just sort of try and show people that, in fact, if you set up this problem, which is very natural, given the object, the mathematical object, that all these nice properties roll out for free, and it's very interpretable. You can understand what the network layer is doing, and that's, I think, really important for people in the neural network community.

Reem: So when talking about the bi-spectrum model, I want to know, if you're implementing any of the top down influences of the modulation of, representation of the visual information in your model. And I know that you're using images that are more simplistic and less complicated and don't necessarily require prior knowledge, but I wanna see how those two tie together.

Christian: Yeah. So in this particular model, we don't address the question of top down, bottom up interaction, which I think is one of the most interesting interactions in the visual cortex. However, we do address, I think a very specific question, there's been this open question for a long time, which is, if I have several neurons that I, of a certain, kind of a certain variety that I can link together to perform certain computations, how should we connect them in order to get this invariance property?

So if I have some object transforming in the world, how should you connect your neurons so that the output is the same no matter how that object transforms? That's been an open question for quite a long time, and there have been several attempts to address that question through things like pooling, which is something from the machine learning community, or even the idea of the energy model of complex cells in vision science. And so what we say is, here's actually a third way to do it, and this way actually comes with all these mathematical guarantees, and it's based on the properties of this object called the Bi-spectrum.

So essentially what it says is there's a way to connect your neurons together to be invariant to group actions, that is directly derived from the properties of that group and the kinds of neurons you're using.

So that's the statement of the paper. And I think what the most interesting upshot from the paper is that instead of having to specify what transformations are important for that task ahead of time, we show that there's a way to learn the way to connect the neurons to perform this task, and you show that what you learn ends up approximating this mathematical object called the bispectrum. So that's, I think the, the main contribution of the paper.

Reem: So, in other words, in a way, you're looking at a way to connect the neurons, that guarantees performance that's invariant of the changes in the image.

Christian: Yeah, exactly. Exactly.

Reem: And it does still work as a classifier…

Christian: Right, so it itself is not a classifier, but the output is a new representation, a bunch of new numbers. They're not pixels, but they're new numbers, which are neural activations. And if you try and feed that to a classifier, upstream, that's trying to make sense of what it's seeing, then it's an easier job for the classifier to decide if these two objects are the same or different because it's removed that transformation as an ambiguity or as a factor of variation.

Reem: Oh, that's interesting. So in a way, it's kind of, instead of learning from just a set of images, you're giving it kind of these components that's like the output of the neural representation and then the classifier can then perform better if you deliver that input, as you know, the training data.

Christian: Yeah, exactly.

Reem: That's pretty awesome.

Christian: It's pretty cool, I think. I think it's pretty cool. Well I think this is sort of the trend in the last, I mean for a long time and, and deep learning is this shift away from just focusing on training models to classify objects, so “Give me an image, what's its label?” To instead saying actually, if we want to build really robust systems that can generalize across different *kinds* of tasks, then they need to form some kind of internal representation — so that they could answer different kinds of queries, like, uh, “Where is this relative to that in the scene, or how many of them are there?” Or “What is in the scene?” for example.

And those kinds of queries depend on a very rich model of the world in your internal representation. So we sort of see our work as contributing to that thread, which is saying that here's an important property representation should have, which is this equivariance or in-variance to group transformations, and here's a way to build it in and that has all these nice mathematical guarantees — as opposed to some of these more ad hoc sort of empirical methods, which seem to work.

Reem: That's really interesting. Moving from that, I just learned that you're organizing the NeurIPS workshop, and it aims to bring topology and geometry to bear on neural systems and biological systems. Can you tell us more about this workshop?

Christian: Yeah, sure. We're pretty excited about this. So, I'm lucky to be organizing this with four really awesome collaborators. The first is Sophia Sanborn, here at Berkeley, between here in Berkeley and Santa Barbara, Nina Milan, Simone, and Ariana — they're two collaborators from Europe, so we've been having a blast getting this together.

I think there's this really incredible convergence between ideas, both in, uh, these really deep fields of mathematics, things that have been applied to mathematical physics. The same mathematics that we use to talk about space time and particles and field equations and things, are popping up in places like neuroscience, you know, are popping up when we look at, uh, how neurons represent structures in the world, in the brain. So this is, you know, incredibly exciting and I think anybody who's been involved in any one of those fields, when they see this, is pretty struck.

So essentially, you know, my role in all of this is to like beat the drum, you know, essentially is to like make a gathering place for people who have that deep background and knowledge in the mathematics, but also neuroscientists who are completely immersed in the field and immersed in the neural data, as well as people who in machine learning, like we talked about before, are interested in these kinds of representations for artificial systems.

So the goal of the workshop is to basically just bring all these people together and get them talking and help them elevate their work to the broader community. It's gonna be held in early December in New Orleans at the end of the NeurIPS conference. So, we'll have a one day workshop at NeurIPS with various talks from people who are doing really interesting work from all of those different angles and some discussions, like panel discussions and things, from people that I really admire in the field. Uh, and as well, we're going to be releasing, like, a proceedings, so like a publication. We've taken something like 95 paper submissions and gotten over a hundred reviewers to go and take time and peer review these pieces. So we'll be releasing a PLMR publication with this collection, as well.

Reem: And is the workshop open for everyone? Do you have to apply or do you just RSVP?

Christian: Right. So it's a, it's a workshop as part of NeurIPS, the conference.

[music interlude *Familiar Skies*]

Reem: Moving from the work you're doing now: What questions have the work you've done led you to ask and plan for future research.

Christian: There are many really interesting directions you can take. This object, the bispectrum and its learned version — you know, many people in machine learning, I think, will be interested in using it as a primitive in larger neural networks. So how to sort of learn representations that are invariant to different kinds of transformations through the hierarchy of the network. So that's definitely a subject of future work is how can you incorporate these into larger neural systems, like artificial neural systems and can you… Basically, there are certain guarantees this object has, which are very nice compared to how people deal with this problem currently, which is something called max pooling. So we sort of compare it to those.

So I think comparing to things like max pooling, comparing to existing layers people use for invariance will be really interesting. And one of the problems we'll have to solve. There are many problems to solve, to do that.

I think the other, which is really interesting is... Well, there's some subtle ones about the mathematics, which is — there are different kinds of groups. So there's something called commutative group, which is where if I move something up and then to the left, that's the same thing as moving it to the left and moving it up. But if… there are non-commutative groups as well, where that doesn't hold. So you can imagine, if I were on a sphere and I move along a certain path…and then I sort of move. Yeah, they're the… we could go through an example, but… So non-commutativity is also a very important part, like 3D rotations are a non-commutative group.

Which are very important for vision. So right now, the work that we show and the results we show are for commutative groups. We have early results for non-commutative groups, but that's, I think, a really important extension and we think that object will be able to push past that barrier that has stopped many other groups, too.

And maybe the final thing, which I think is really interesting, is… so Sophia Sanborn, my collaborator on this project, has also looked at applying these ideas to the construction of complex cells and the visual cortex in D one.

So, in particular, you can make predictions about what kinds of connections you would expect or what invariances you would find in complex cells based on the sort of the bispectrum and what sort of operations have to happen. And it turns out it's neurally implementable. So you can actually take this algorithm and you can come up with a neural circuit, which implements it as a way of connecting neurons. And she also shows that it replicates things…

Reem: It's really interesting!

Christian: Yeah, it's really interesting! And she also showed that it can replicate things like end stopping, which is like a — I mean, there are many other really interesting phenomena that come out. So even looking at connections to what goes on in neural circuits, in the human brain.

Reem: If you can share with us some tips or difficulties that you had coming into the program, or, you know, in your background work before coming here — skills that were really valuable in what you're doing now and a little bit, if you're comfortable with it to share, you know, what skills you wish you had before coming here.

Christian: So, once I think — at least if you were doing computational work — like some background in mathematics is really helpful. It's been helpful for me, and I think other people in the lab, I think, have benefited from kind of, like, leveling up in that way for those problems.

So particularly Bruno works a lot with questions about probability for Bayesian inference. There's a lot about sort of property, like linear algebra essentially is extremely important in any kind of applied numerical topic. So those are really helpful.

Being able to program, like write in Python and be able construct, you know, complex programs to do, to work with these different problems. All those skills have been super helpful and have really accelerated all of our work. We couldn't ask the questions we ask if not for those tools, I guess.

As far as things I need to develop, there's so many! One that I'd really like to work on pretty diligently is how to … how to take an idea for a project that can go many directions and maybe could go on for many years and how to decide what the most important path is. Like, how do you break up a research problem? And make progress and then, you know, look from there. So I definitely, I'm looking forward to growing in that way.

Reem: Well, it has been a pleasure talking to you, as always. And I want to thank you for coming here and taking the time to have this conversation and, you know, supporting me in this endeavor as well.

And I think it's important to share things about vision science, because it is really cool! Thank you. So thank you for coming today.

Christian: Oh my gosh! No, thank you. This was a lot of fun!

Reem: Thanks, everyone!

[outro *Familiar Skies*]