MIT CSAIL Alliances | Justin Solomon Podcast Export 1

Welcome to MIT's Computer Science and Artificial Intelligence Labs Alliance's podcast series. My name is Steve Lewis. I'm the assistant director of Global Strategic Alliances for CSAIL at MIT. In this podcast series, I will interview principal researchers at CSAIL to discover what they're working on and how it will impact society.

Justin Solomon is an associate professor at MIT. He received his PhD, MS, and BS in computer science as well as a BS in mathematics from Stanford University. Beyond his published research papers and articles, Justin has also published a book called *Numerical Algorithms.* His research focuses on geometric problems in graphics and machine learning using tools from large scale optimization optimal transport in differential geometry.

Justin, thanks for your time today. Can you explain to our listeners the focus of your research and some of your aspirations?

Sure. So I'm the PI of the MIT geometric data processing group here in CSAIL. That name is purposefully vague. That's because we study a lot of different problems that all have a geometric flavor. So really our research kind of congeals around two broad themes.

One of them has to do with the processing of geometric data. There are so many different applications out there that we study that really do involve shapes, whether they come from the natural world or even the virtual world like in computer graphics. So some of our collaborations include working on algorithms for computer vision and the autonomous driving space to coming up with algorithms that can support artists and engineers in the computer graphics domain. These are areas that really naturally work usually with three-dimensional shapes, trying to process, understand, edit them, and so on.

The other half of our research group works on geometric approaches to analyzing data. So a lot of times when you have some more abstract statistical problems you can view a data set or some collection of information as actually cutting out a geometric feature in some high-dimensional space.

And one of the really exciting and fun aspects of the research that we do is that a lot of the insight that we've gained over many years of working on these very low-dimensional, tangible, and easily visualized problems in computer graphics and vision actually helped give us a little bit of perspective and insight and intuition for these high-dimensional problems in data analysis.

So while some of our collaborations really are working on these pretty concrete sort of 3D-style problems, other ones are much more abstract and working on but you might consider to be more statistics, probability sorts of things. The mathematical and algorithmic toolbox between both of those is remarkably similar.

Great. Thanks for that background. What progress has been made in applied geometry in the past decade that you believe is the most innovative?

That's a great question. So from my perspective geometry for a long time was not really considered an applied discipline in itself. So a lot of times there were different geometry faculty who might work on problems specifically in medical imaging or geometry for computer graphics or geometry for data analysis. Somehow it was the sub-discipline in a lot of different areas.

To me what's one of the more exciting and interesting developments in this community is the realization that we're all working on the same problems and we can actually come up with a pretty general purpose toolbox that can help us understand and cut through many different applications.

Some of the nice side effects of that is that the folks who study geometry don't really regard it as just a mathematical or a theoretical exercise anymore. It's really focused on developing a toolbox can be widely adopted both in pretty practical applications and in more exploratory settings.

And in some sense that's made the community a lot more inclusive, as well. It used to be that if you worked in this space you had to have background in differential geometry or specialized algorithms and so on. A lot of times now we've really made it so the geometry toolbox is something that you can actually download, work with, and apply to the disciplines that you're working in without the sort of detailed and specialized knowledge and still have some pretty effective results.

So it's the tool set, in your opinion?

That's right. I think what's most exciting and innovative is the idea that you don't need to hire a specialized expert that is an expert in applying geometry specifically to the problems in your domain area. These days what we've developed are useful and intuitive tools that people can use at home without having to have all the detailed knowledge of every piece of mathematics.

Can you give an example of how those tools might be used in a practical application in any domain?

Sure, yeah. So for instance, we have a lot of collaborations in the autonomous driving space. Maybe that's a good, tangible place to start. So in autonomous driving geometry is pretty obviously a key factor in how your car's perception system should work.

So if you think about the environment around a car as it drives down the road, there are many different three-dimensional inference problems that you have to solve. There's other pedestrians, other cars, obstacles, and things that you need to understand in order to navigate around them.

And of course doing that requires some really complex reasoning about the environment. The first thing that you have to do is to find all of the objects in the scene and maybe predict their dynamics, how they're moving around. And then of course later on there are interesting planning problems that have to do with actually interacting with those objects after you've understood the scene.

So these days when we work on geometry problems we're developing tools that can, take for example, a data set of scenes that maybe are labeled or unlabeled, and can have the perception system automatically tell the engineer precisely what's there.

Now, in order to do that, there are some pretty complex mathematical and algorithmic things that have to happen. In order to understand the 3D world, we have to take input from some sensor, like a lidar-style system or maybe some kind of stereo and go all the way from that to pretty semantic information about the world around us. It's a pretty complex series of operations.

But these days there are actually pretty reasonable packages you can just download, some of which were developed from our team, that have pretty nice performance on these sorts of tests. So there are a lot of really exciting technical developments behind the scenes that are enabling these sorts of things, but we also see them pretty widely adopted in academia and industry outside of just our specialized area, which is pretty exciting.

So for example, several years ago my PhD student, Yue Wang, developed a tool that he called the Dynamic Graph Convolutional Neural Network. The naming is a bit of a mouthful but the algorithm itself turned out to be one of the state-of-the-art tools for understanding point clouds, which is a particular expression of a 3D shape.

And what's really interesting is the speed with which the community has adopted that and other tools in this space. And a few years later now Yue's work is widely cited and used and incorporated into different systems and it makes working in this area really exciting and fun. We're not just working on abstract problems anymore. It's really something that can have a pretty concrete effect on users in other domains.

So that's a great example how your research is affecting the way you engineer autonomous driving systems. Can you tell us how your research group interacts with industry?

Sure. So our group interacts with industry in a lot of different ways. Of course when you work in an area like geometry, as I've mentioned a few times, there's so many different applications that we use our collaborations with industry as inspiration for interesting research problems. So in many ways we end up with a model where we're almost consultants.

We have friends from different application areas and domains who maybe don't understand the theoretical aspects quite as much but can come to our group and say, we think we have some interesting geometry problem to study, we don't know how to articulate or precisely what tools are out there. And we've had many different successful collaborations that start that way.

In terms of actually how interactions work with our group, that takes many forms, everything from sending students off to internships for the summer and then they come back with some inspiration or knowledge of domain areas to longer term research engagements, even some configurations where engineers are actually placed in our research group and maybe they're the domain expert but they're also interacting day-to-day with the team.

So for example, in autonomous driving our group has a long-term collaboration which is supported by the Toyota Research Institute, which is focused on 3D computer vision for autonomous driving. We have a second Industrial relationship with the MIT IBM Watson AI Lab, which is positioned right across the street from our office and actually in that case our collaborator, Misha, is in our office once or twice a week and interacting really closely with our graduate students on optimal transport problems which are related to some of the data analysis problems that we study here.

That's great. And for the folks who don't know it CSAIL the Alliance Program has a visiting researcher program where researchers can be embedded side-by-side with PhD students and researchers in the lab at CSAIL. Let's talk about machine learning for a little bit. So what types of machine learning problems can be approached from a geometric perspective?

That's a great question. So we've already talked quite a bit about geometry problems that can be approached with machine learning, things like understanding 3D shapes for 3D modeling or autonomous driving. The other kinds of problems that can be approached from a geometric perspective involves analyzing data as itself a geometric object.

So if I go out there and I collect a big data set like survey responses or information about people who are using my website, you can think of each of your data points as some position in a very high-dimensional space. Like every single piece of information that I have about my data point is some other dimension. And oftentimes when we talk about data we do tend to use geometric language. Some data points are closer to each other than others. Some are far away, similar, distant. Even notions like curvature sometimes come up.

So in our research group a specific topic that we've studied related to that is an area that has morphed from an area of economics to one in mathematics and now one that's largely in computation, which is known as optimal transport. Optimal transport is a theory of probability that links probability to geometry.

So oftentimes when we talk about a data set, we can think of it as sort of a random sample from some bigger space. Like if I go out there and I collect the height of a few people in my classroom, that's maybe a sample from this continuum version of all the heights of all the people in the world or even some generative process that underlies that.

When I think of my data set as a sample from a probability distribution like that then what optimal transport tries to do is to give you a geometric way to compare to probability distributions. So like if I go into my classroom at MIT and I collect a data set of everybody's height and eye color-- eye color is probably a tricky one-- but in any event, then I go to a classroom somewhere else or maybe everybody's taller then I want to somehow come up with some inference to be able to compare these two different distributions.

The way that optimal transport does that is by treating your distributions kind of like piles of snow, a pretty familiar story here in Boston. So if I have two data sets then essentially what I would do in this mathematical theory is to measure the amount of work that it might take for me to displace or transform one into the other. And the more similar they are, the easier that task can be.

So this sort of optimal transport problem originally came from the world of economics. I believe Kantorovich in the 1960s or so was studying these problems for moving the Russian troops to the front. Many people in France work on these problems. They often talk about it as delivering flour to bakeries and what the most efficient ways to do that.

The reality is these days we apply transport problems to more abstract settings and data analysis. Where for example, maybe I have a computer that's trying to learn a model of the world. What's generating the data? What's the process that's underneath the information that I observe? Then there's a really natural transport problem that appears there, which is to try and understand the distance or the similarity between the model that I have of the world and the data that I have on hand.

And a lot of different problems in assorted application domains can be phrased really naturally in terms of that language. Everything from natural language processing, document retrieval, these sorts of things, to Bayesian inference and generative modeling. So can I come up with a software that randomly generates data points that look like the data of I've already observed? So for example when you read about deep fakes in these sorts of image synthesis techniques they kind of fall into that area.

So in each of those cases the geometry is not just some 3D shape but rather trying to understand how my computerized model of the universe aligns to a data set or some information that I've gone out and collected. And so essentially what our group is focused on is taking this abstract mathematical theory, in this case this idea of optimal transport, and really making it practical and tractable on a computational system.

And that involves all kinds of fun perspectives from large scale optimization to inference and numerical methods. And really we're right at the forefront of being able to translate these sorts of abstract geometric pictures to algorithmic reality. At this point we really can solve these problems and apply them not just to small problems in economics but to these extremely large scale learning and inference challenges that really are naturally phrased in this sort of shape analysis language.

Fascinating. Can you talk about what other application areas that your group interact with?

Sure. We have so many different collaborators in so many different spaces. It's really a lot of fun. My personal background is largely in the computer graphics domain and we have many friends in that space and continue to publish quite a bit. So in computer graphics and also in scientific computing geometry really shows up in many different places. The two kind of main ones that come to mind are in computer animation, like taking a 3D shape of a digital character and deforming it and understanding how it moves over time, and also in some scientific computing and physical simulation style problems.

Oftentimes you know your favorite movie effects come laden with hair and fluids and solids and rigid bodies. And all of these different objects are really 3D shapes that are interacting with the world. So many of our graduate students and members of our team studied problems in the space of geometry processing, which is the area of computer graphics that has to do with synthesizing and editing three-dimensional shapes. And every year we get to share some of the results that we have with that community.

So another space the geometry really affects the application domain is in medical imaging. And indeed, we're really lucky to share a graduate student with Polina Golland's group here at MIT who studies medical imaging problems. In his case he's studying medical imaging of the placenta, which is an interesting organ because it's extremely non-rigid. It's something that deforms quite a bit both in the human body and eventually outside.

A lot of the geometry processing techniques that we've developed in the computer graphics domain are providing some insight into the medical imaging space as well. And that relationship goes both ways. We have a lot to learn from how people process modalities like MRI within the context of computer graphics and vise versa.

There are many other application domains that we've studied over the years. One of the really fun ones that we studied a couple of years ago was in political redistricting. Of course, the shape of your political district has a big bearing on how you elect candidates to Congress. And one of our collaborations that went for several years involved how to analyze and understand differences between different districting plans, which is another fun and pretty impactful problem. And some of the algorithms that we developed there actually are getting used in the analysis and the redistricting cycle that's just coming up.

So overall, I think geometry is this really exciting area because of just the broad set of spaces where it gets to be applied. And somehow we get the pleasure of interacting with people in all these different areas and learning all of the interesting challenges that they run into, many of which are unanticipated and help us inform the different projects that we have here.

It's pretty common for the students we have working on graphics to step in and help our machine learning students or to get some insight into the medical imaging problems that our student Maz is studying, for example. And it's really these collaborations that drive our research forward.

That's one of the things about CSAIL in general with these communities of research concept that exists at CSAIL where you have all these interdepartmental collaborations and lots of progress can be made when you're sharing information. So can you talk a little bit about your book, *Numerical Algorithms?*

Oh sure yeah. Right, I guess this book came out several years ago at this point. So it's a textbook on numerical methods. Numerical analysis really underlies basically all of the research that we do. The phrase numerical analysis refers to computerized techniques for number crunching where the numbers are real valued, like they have a decimal point. So you can contrast that with like discrete mathematics where you're doing counting and that style thing.

So this is a textbook that was intended to give more of an applied perspective on numerical problems. So basically every single chapter in that book, whether it's studying numerical linear algebra or optimization or solving differential equations, tries to place the problems in the context of some applied space and show where these different things get used.

Really much of that is lifted from my own research and research from my colleagues in these different areas. And it's a lot of fun to teach this kind of material. So hopefully in the coming year or two I'll find some time to come out with the next edition but of course, there's a lot of work so we'll see. TBD.

Well great. We look forward to that. And we know you're very busy and we appreciate your time today on the CSAIL's Alliance's Podcast. Thank you very much, Justin.

Of course.

And if people want to find out a little bit more about your work, can you point them to a URL?

Sure, yeah. So my research group is the MIT geometric data processing group. So if you go to gdp.csail.mit.edu can find our group web page there. My personal web page is people.csail.mit.edu/jaysolomon. I'm sure we can put that in the notes for this podcast somewhere.

I feel very strongly that we should share all of our research as openly as possible. So on my personal web page you can find a PDF of every research paper that comes out of our group. And so you're invited to take a look and of course, reach out if you have interesting applications, ideas to share, or problems we should be studying.

Great. Well thanks again, Justin, we appreciate your time today.

Of course. Thank you for having me.

You're welcome.

[MUSIC PLAYING]

If you're interested in learning more about the CSAIL Alliance Program and the latest research at CSAIL, please visit our website at cap.csail.mit.edu and listen to our podcast series on Spotify, Apple Music, or wherever you listen to your podcasts. Tune in next month for a brand new edition of the CSAIL Alliance's podcast and stay ahead of the curve.