AI in Humanitarian Assistance and Disaster Response

featuring Ritwik Gupta as Interviewed by Andrew Mellinger

Welcome to the SEI Podcast Series, a production of the Carnegie Mellon University Software Engineering Institute. The SEI is a federally funded research and development center sponsored by the U.S. Department of Defense. A transcript of today's podcast is posted on the SEI website at sei.cmu.edu/podcasts.

Andrew Mellinger: Hi, I'm Andrew Mellinger. I am a software architect for the Software Engineering Institute's Emerging Technology Center [ETC], and today I'm going to talk with Ritwik Gupta, who is a software machine learning research scientist at the ETC along with me. We are going to talk about the use of artificial intelligence in Humanitarian Assistance and Disaster Response projects, which we currently refer to as HADR.

Andrew: Hi, Ritwik.

Ritwik Gupta: How's it going, Andrew?

Andrew: Good. Can you tell me about kind of how you got to the SEI? How you got to the ETC, and what you actually do there?

Ritwik: I graduated from the University of Pittsburgh—shout out to Pitt—in 2017 with my degree in computer science, my bachelors. I was looking around my life in general about, *Should I go to grad school? Should I just go work full time? What should I do?* I really wanted to come to CMU for a PhD. That is when I talked Dr. Andrew Moore, who was a dean at the time. He told me, *You are so young. Go work for two years. Come back. Get some money. Get some experience, money mainly. You'll have more scope in your life to go do stuff*. So, I agreed, and I said, *All right, I'll put off the whole grad school thing, but I'd still like to work in a place where I can still do that academic work and still do some research*.

As it happens, by chance, the SEI found me while I was giving a talk at a high school about machine learning. They were like, *You should come work here*, and I said, *Sure*. So, I ended up here, long story short. The team that we work in is called the Emerging Technology Center. What we focus on is applied AI and ML in all sorts of domains, human-machine interactions, machine emotional intelligence, and then advanced computing; those are the three pillars.

My background has been traditionally artificial intelligence, ML, and health. I used to work at this awesome department at Pitt called the Department of Biomedical Informatics. They used to do all sorts of really cool cancer pathology, whole-genome sequencing, whole-genome causality modeling stuff. That was really fun, but then I've also worked on more electronic health record side of things and just epidemiology a little bit, some whole-side imaging. My background is really, *How do we take this really cool statistical machine learning or these AI capabilities and actually have direct impact on human health or human lives?*

When I came here, my goal was really open-ended. First of all, it was, *Let's go to grad school*, which if you're working at CMU I feel like that is pretty easy. The goal was, *I want to build a body of work that lets me do that, lets me do more of, how do we take AI and ML and impact human lives?* While I have been at the SEI, I have led a small research effort focused on robotics and specifically how robots understand anomalous behavior, not only in their environments but also as they interact with different humans in their surroundings. That has been a lot of fun, so we've done a lot of work with inverse reinforcement learning to understand autonomous robotic behavior.

Separately, we've done a lot of work with remote sensing, EO [electro-optical] synthetic aperture radar imagery and understanding what are different parts of the world which are damaged or which are about to be flooded or whatever, then some auxiliary work and just like adversarial machine learning or some other cool stuff. Really merging all those gives me, I think, a broad platform to talk about how we can take AI and ML and directly use it to help human lives directly, whether that be directly through their health. It is actually impacting the way they spread diseases or the way they understand neural sheath damage for example. Or, more abstract, like, *How do we assist first responders and disaster people to better understand what parts of the world are damaged after disaster, so that they can better plan logistics and resources,* which then goes on and impacts human lives.

Andrew: Right. You mentioned different modalities like computer vision and synthetic aperture radar. Can you talk about the background a little bit before we get into actually the HADR part? What is available? What is the state-of-the-art, the state of practice in those areas right now?

Ritwik: The world of HADR…So computer vision is a specific field of machinery, right, focused on how do we use computers to do vision tasks, anything with imagery. Then SAR [synthetic aperture radar] and EO would be modalities within that domain. But speaking generally, HADR…We call it HADR I feel that people call it HA-DR. I prefer HADR just because, like, rolls off the tongue, and people think I'm talking about.... HADR is a very broad field. It can encompass things like computer vision, like, *How do we use computer vision to do damage planning or something like that? How do we use computer vision to do routing across constructed roads?* By all means, HADR can also involve things like rescue robotics like, *How* *to build robots that can crawl around the rubble and find people*. It can involve things like, *How can we stop the spread of misinformation on social media during and after a disaster*, which is very common. *How do we understand disaster forecasting? If there are floods, how do we forecast where the high-water mark will be? How do we forecast what will be impacted? What critical infrastructure is impacted?* The field is very broad like NLP [natural language processing], computer vision, robotics, game theory. It is the whole gamut of the AI world.

I specifically focus on computer vision just because that has been my background. But that doesn't mean that, like, we are limited in what we can do. My interests are then combining that with robotics, so not only can we make predictions about the role, but then we can actually do something about them via robotics. Either way the goal is to make sure that less human lives are at risk when we are responding to these things because they are really dangerous.

Andrew: We talked before about some other projects you have worked on like the Defense Innovation Unit's xView that kind of led up to this. Can you talk about that a little bit?

Ritwik: Yes. There is this team of people in the DoD, and they are called the Defense Innovation Unit. They are based out of SF, and there is a small team in the Pentagon, I believe.

Andrew: SF?

Ritwik: San Francisco and Mountain View. What their goal is, is to basically take some really, really cool commercial technologies that are available from startups or wherever in the world right now and transit them over to the DoD space to be highly applicable and impactful. One thing that DIU focused on last year, back when they were still called DIU experimental, because they weren't sure if that unit was going to do a great job, and they are.

They said that this whole world of operationalizing computer vision stuff that we do with satellites is really important. It is not sufficient to just create prototypes and say, *Yes, we can detect buildings or something.* It was like, *Let's get that in the hand of someone who can actually do something about it; who needs that information?* Last year, DIU launched this thing called xView 1 competition. They created one of the largest overhead satellite imagery datasets where they had millions of buildings and just millions of objects overall that were very fine grained. They had everything from a pick-up truck to a pick-up truck with a trailer, except like these objects could be really small in the image.

It could be, like, identify this as a pick-up truck versus a pick-up truck with a trailer, and the thing is, like, five pixels wide. It could be really tough. There could be over a thousand objects in one small image. It could be really, really dense detection. The issue with xView 1 is that they weren't able to do the whole operation part of it. It's a really cool regional dataset. It is really hard. We competed in that, and we got fifth place in the world.

In my opinion, our metrics weren't anything great. It was OK. We got a mean average precision of like .28 or whatever, which again on comparable datasets that aren't this difficult, the same model that we used could get a mean average precision of like .8. Just for a comparison of the gap, there is a lot of work to be done in this area. It is a really good academic research dataset, but it wasn't clear how to, like, make that useful for anyone doing operational work at the moment.

DIU then reached out to us and said, *Hey, you guys did really well in this challenge. Would you want to help us figure out what we want to do for the next one?* That is a little more operational. This is a little bit more impactful but still involves computer vision in some way. We went through a whole bunch of ideas. We talked with USGS [U.S. Geological Survey]. We talked with FEMA [Federal Emergency Management Agency], NASA [National Aeronautics and Space Administration]. We talked with all these people who have all the data and all these amazing problems. I got to realize for the first time in my life these are problems real people have to solve every day, so that we don't get completely screwed when the next hurricane hits or whatever. We ended up deciding on creating a challenge that lets people understand or lets models detect how damaged a building is after a disaster.

What happens is—I'll use California as example because they do a pretty good job of it. After, let's say, a wildfire in California, they have about 24-to-48 hours to make a presidential declaration of emergency. That lets the National Guard militarize. It lets them do all these things that enhance their response to these fires. What they have to do is they have to go in and give the declaration, like, *We have 17,351 damaged buildings or whatever. So, we need help*. That is a really exact count. The way they do that is they literally send people in after disaster within 24-to-48 hours and someone has to sit there and count buildings.

Andrew: So, this is manual.

Ritwik: It's manual and you can imagine how dangerous. Let's say it was a flood, and there are downed power lines in the water, and there are, I don't know, alligators roaming the streets or, like, looters, I don't know, shooting places up. You don't know, right? So it's really dangerous to send people in after a disaster to go and do this task just to declare emergency, so you can get more help. The idea was this happens everywhere. This happens in California, North Dakota. It happens around the world in Nepal, everywhere. You need to be able to figure out, *What is damaged? How damaged is it?* So we can get not only declare an emergency, but we can also then respond accordingly, right? *If I don't know this village is damaged in this remote part of Nepal how am I going to get help or logistics or resources there? I'm not.*

What we did, what we have gone about and done is via our partners at DigitalGlobe, which is now Maxar [Technologies]. They have this amazing repository of satellite imagery from before and after disasters. We have gone in. We have collected that data. We have annotated it with not only where those buildings are, so polygons over buildings. Also, we worked with partners at Cal fire, California Air National Guard, FEMA, NASA to create this scale that lets people assess on a scale of 0-to-3 how damaged a building is. Is it *not destroyed*, *minor damage*, *major damage* or *destroyed*? Because right now they are using binary. They were like, *Is it damaged or is it not*? That is really binary and doesn't really give as much context as you would like to have when responding to these things. So, we went in and created this scale. We got this whole dataset annotated. We are going to be releasing a prize challenge coming up in the next week or so that takes this massive dataset. It is over 30,000 square kilometers of annotated satellite imagery. The challenge is to create machine-learning models that can find buildings and then tell me how damaged they are from electro-optical satellite imagery in the visible spectrum.

That is really exciting. That is called the xView 2 challenge. If you go to xView2.org, you can find out more about it. But the prize pool is $150,000, so it's a massive prize challenge as far as machine learning challenges go. The idea is that we will get fully Dockerized containers back from competitors and whoever wants to submit a model. The idea is that we will then help transition those models over to the Cal Air National Guard or UN [United Nations] or NASA or who can actually directly use this technology to actually do their work. We are pretty excited for it. It should be a huge force multiplier for damage response, disaster response.

Andrew: Are you doing any AI yourself for this, or are we are just supporting that?

Ritwik: We are. So not only are we creating the dataset, we are also creating a baseline model—we will call them reference models—so that people can get started with the challenge. Not a lot of people have experience working with satellite imagery or like change detection stuff like that or just advanced machinery methods. We are creating a library of code that lets you do localization and classification or ordinal aggression—is what we are going towards—over this dataset. Then we are going to release it to public when the challenge comes out, so that they can start working on top of that and hopefully beat us and do better. People don't have to use that if they don't want to. Some people are very capable, and some people who are competing are part of major companies. They already have, like, well-established ways to do this, or [they can] invent new ways to do this task, which is what we want. If we were able to do it already, then we wouldn't launch the challenge.

Andrew: Fantastic. What's your role after this then? So, you are going to get the data back and we adjudicate it, or how does that work?

Ritwik: Yes, again. We are working with the Defense Innovation Unit. We are working with some other partners and contractors to run the challenge. Our role is going to be basically to keep on continuing to deploy better baseline models, so that people can keep continuing to work against a better and moving target. In general, though, our role is to continue doing some good research: find better ways to do this task and find other ways that we can be impactful in the HADR space. Maybe we figure out that electro-optical imagery isn't the best or isn't the most useful for this task. Maybe we need to go over to synthetic aperture rate or SAR imagery. We are going to be working on finding what is new and what is next and just being ambassadors for bringing this world of AI and HADR together.

We've been trying, and we've been working with a group of some excellent faculty at CMU's campus and at Berkeley to create basically a center of excellence, a center of research for HADR work. Dr. Howie Choset, Matt Travers at the CMU Robotics Institute have been working really hard over the past year or so to create, like, a center at CMU that lets faculty that focus on HADR work come together and do all sorts of cool research together and publish under one arm. Equivalently Trevor Darrell and his team at Berkeley AI Research Lab have been doing the same thing. And, they have been working really hard to make sure that the work that they are doing, which they already apply to HADR, is organized and under, like, a common umbrella. We have been collaborating with them to see how we can make their work more impactful.

We will be hosting in NeurIPS [Thirty-third Conference on Neural Information Processing Systems] this year. NeurIPS is a very large machine learning conference in Vancouver this year. We will be hosting the first ever AI + HADR workshop at NeurIPS December 13 that is focused on bringing members from the AI world—whether that be robotics or computer vision, NLP or whatever, and people from the HADR community—whether that is people who do epidemiology research all the way to people who would do, like, from the Cal International Guard who do fire response or fire intelligence collection—together into one room, so that we can understand the powers and limitations of AI so that the HADR people understand that. More importantly I think for the scope of the venue, which is NeurIPS, so that AI people can understand what is lacking in their research or what is limiting in their research to be perfectly operational. And get new ways to guide their research to things that they may not have ever known about. We are trying to make that happen. We will make it happen on December 13th. Really we are hoping that will be a good event. We will also announce the winner of the xView2 Challenge at the workshop as well.

Andrew: Along with a lot of the image recognition things in a very abstract way, can you give me some examples really quick about, like, what is actually discovered out of using these techniques? Imagine a disaster. Besides saying, *Hey I've got a damaged building*. What else is on the horizon for things we can do? Things, like, *Can I actually find people in the rubble or something else?* Or what's available?

Ritwik: Yes, again, I think that is a very good example, finding people in the rubble because that comes out of CMU, so it's close to us. Howie Choset and Matt Travers, they have this thing called the Biorobotics Lab in the Robotics Institute at CMU. One of the things that they make is called the Snakebot. People have probably seen videos of it. It is literally this robot made out of multiple, independent components that chain together, and it moves like a snake. You can control it like a snake. It rolls, and it twists like a snake. Back when the Mexico City earthquake happened, about two years ago now I think, Howie's team was actually called in to go to Mexico City and use a robot to dig through rubble and assist in finding people and stuff. I don't know if they actually found anyone, but the idea was that, it wasn't one of the first times, but it was a really cool time where there was a real earthquake and a real big disaster, and people recognized the use of robots that were getting in these really tight places that humans would never have a chance of getting into and exploring and basically serving as a force multiplier for, *Can we find people when there aren't enough people on the ground to do that task?* I think it was a really cool demo of a technology that is only going to get more pervasive in the future.

Like we saw like with the recent Fukushima disaster, when the radiation's terrible we send in robots, so humans don't get completely destroyed. They did the same thing with Chernobyl. Those are extreme situations, but even just during things, like, *How can we use autonomous drones to locate stranded, isolated communities and drop them resources or communicate with them?* I mentioned the spread of disinformation. A few years ago, I forgot which hurricane it was, it was one hurricane flooded the streets. And there was this picture circling on Twitter of a shark swimming through the streets, and people were literally freaking out, *Like oh [expletive]. There are sharks in our street*. What other people are doing is, *Can we use natural language processing and social media and whatever to find disinformation like this and help disaster agencies or coordinating agencies put a stop to it before it gets to panic state.* People aren't leaving their houses because, *The shark might get them*.

Again, there are so many uses. The issue is that a lot of the HADR world is so manual. Because not only is the work to recover from a disaster completely manual—and so people really don't look beyond that—but the technologies to even evaluate those things are manual. We were recently invited down to FEMA last week to talk to the remote sensing people about, *What can we do with remote sensing to make disaster*… One of the things that I found out about it—I didn't know this was a thing was—well, I knew it was a thing. There is debris. There is debris after disasters, but I don't really think about it. It just goes away right? I never looked into what happens. Someone has to get rid of that. There are contractors hired by the state, city, whatever, that take care of debris for them. How much debris there is, where that debris is, something that people just have to go find. Understanding how big debris piles are is a challenge that exists in the world and something I was completely unaware of. So Argonne National Laboratory is actually building this whole thing with drones, which can do volumetric LIDAR scans of debris piles and give you an accurate measurement of, like, *You have 35 cubic tons of debris in this area, and this how much effort it will take to move it, and this is how much you should be billed for that to happen.*

Andrew: Alright. Well, thanks for coming by. We could talk about this all day I am sure. I think afterwards they will have links in the transcription, other resources available. They can follow us online.

Ritwik: Yes.

Andrew: Thanks for coming.

Ritwik: If I may add too, if you are interested in more resources about HADR and how HADR is being applied to AI—it currently is being used for the workshop—but we have a URL called HADR, hadr.ai. If you go to hadr.ai, you will be able to find some really cool resources that talk about how to apply HADR to AI and what is possible and the people working in this area. It's a pretty good resource if anyone is interested.

Andrew: Hopefully, we can talk about xView2 after the challenge is done.

Ritwik: Absolutely, absolutely yes. Thank you.