Human-in-the-Loop-Autonomy

featuring James Edmondson interviewed by Suzanne Miller

Suzanne Miller: Welcome to the SEI podcast series, a production of the Carnegie Mellon Software Engineering Institute. The SEI is a federally funded research and development center at Carnegie Mellon University in Pittsburgh, Pennsylvania. A transcript of today's podcast is posted on the SEI website at sei.cmu.edu/podcasts.

My name is Suzanne Miller. I am a principal researcher here at the SEI, and today I am pleased to introduce you to James Edmondson, who is basically the SEI's middleware expert and actually is one of the more recent additions to our staff. James builds middleware for distributed artificial intelligence. He specializes in real-time systems, control, and distributed algorithms. Fun stuff.

James Edmondson: Fun stuff.

Suzanne: In today's podcast we are going to be talking about his work on autonomous systems. Welcome, James.

James: Thanks for having me, Suzanne.

Suzanne: Absolutely. So, let's start off by having you explain exactly what an autonomous system is. That's a big word, but I know in the common literature let's talk about Also, as I understand it, this is a new area of work at the SEI. What prompted the SEI to begin research in this field?

James: Well, I'll start off with the autonomous systems, if you don't mind. "Autonomous systems" has a lot of different definitions. And we really focus, in my group, on robotic systems and autonomous systems for robotic systems, but there's a lot of different areas concerning software. But, essentially I would probably define it as "a computational system that performs a desired task, often without human guidance." And there really is a scale of autonomy involved. You can have between partial autonomy and full autonomy, and we focus in the area of partial autonomy. So, we're really trying to build systems that complement human users and not try to replace them in a certain task.

Suzanne: So, things that reduce the cognitive load on a user or that do a complex or dangerous task that might not be something you'd want to risk a human for, for example.

James: Absolutely. Or it could just be—and we'll talk a little bit about this later—it could be a situation where you only have so many people in the area, and you have a lot of area to cover. So, you really need to extend the reach of the human operators to give them more range within a shorter timeframe than you would be able to do with that person by themselves.

Suzanne: So, that's important in a lot of different emergency-responder, first-responder contexts, as well as military contexts. So, a lot of different applications for this kind of thing.

James: Absolutely. It's a huge field.

Suzanne: And so, what got us started in this here, at the SEI? Why the SEI to do this work?

James: Well, I've really been researching this kind of area for quite some time. One of the things that you see in literature from various agencies, government agencies, from DHS [U.S. Department of Homeland Security] to FEMA [Federal Emergency Management Agency] to military, obviously, documents is that they believe that autonomy has a large role in the future of our systems. So we really started investigating partial autonomy because we believe it is the future of software, and that really we need to start focusing on these smart devices and smart software that form synergies with human operators, and try to enhance what they're able to do.

Really, we believe that partial autonomous systems are the better way to go for especially government agencies. There's a lot of pushback on having something fully autonomous. There's a lot of scary stories about, "We don't want to give them full control." If it's a drone, for instance, flying over skies, "We don't want them to have call for fire. We don't want them to have the ability to do that." There needs to be a human person that makes that decision.

Suzanne: You want human-in-the-loop, is what's typically talked about.

James: Exactly. So, we focus on partial human-in-the-loop autonomy. We think that is a huge part of not only government work but also commercial work. No one wants a system that has no ability for a human to feed back into the system. You always want to have the ability to at least look into what it's deciding and how it's…

Suzanne: And understand how the decision was made and all those kinds of things.

James: You need a lot of forensics. I mean, for both legal reasons and also for just operational reasons. Absolutely.

Suzanne: So, you work out of the Advanced Mobile Systems (AMS) Initiative, which aims to help soldiers who use smartphones, emergency workers who use them when they're responding, roles that we call "on the tactical edge" here at the SEI. I can see some ways that this would fit with autonomous systems, but why don't you say how you see this fitting in with the Advanced Mobile Systems Initiative?

James: Well, I won't speak for Ed Morris, who leads the group. He's got his own focus on what the group should accomplish. But, I think that it's key to remember that Advanced Mobile Systems is really about the future of mobile computing in either tactical-edge environments or in any environment.

I know you guys recently talked with Soumya [Simanta], who talked more about the ability to use cyber foraging, cloudlets, and the ability to use edge analytics and things of that nature, for not only military applications. That's really not the focus of AMS, I would say. It's really about any kind of tactical-edge environment. We consider that to be an environment dictated by bandwidth, by resources—whether that's human or technological. What I mean by that is, you might not have enough people to search an area, like we said earlier. Or, technological where you're in the middle of an earthquake kind of situation, and you only have one bulldozer or whatever, and you need to know where should you send that to help save lives? You could start clearing the whole area and get there after four weeks. You might save some people; you might not. If you can build systems that help people make better decisions to triage situations, when they know they have a lot of stress on the human operators, they have very little time to make decisions, and they have very little bandwidth to work with for communication. We want to help solve problems in that area. We think that autonomy is a great way to augment people and help them accomplish more in less time.

Suzanne: There's a lot of technical challenge there because autonomy also involves some fairly complex software, some very complex algorithms that may also be constrained at the edge. So, how do you deal with that in your work?

James: I think we're in the middle of a major shift in computing from the sequential processor and single tasks, single computational units to accomplish something, and really looking at parallel computing. Granted, we've been doing parallel computing for decades, but there really isn't a focus on it. When I was at Vanderbilt, you weren't required to take parallel computing. You weren't required to take a lot of these high-performance computing classes. They had their base requirements, and that's how most undergraduates are taught right now, is "parallel, if you want to, as an elective."

Suzanne: For those who don't understand necessarily that particular discipline, just give a couple of words about what's different between parallel computing and sequential computing.

James: Absolutely. When you program for a parallel or distributed system, you really have to take into account the perspectives of each agent. So, you can't just program—this variable interacts with this other variable in this way, kind of in a linear fashion. You have to be a little bit broader-minded to understand that you have asynchronous events. You have events firing off at times that you're not expecting, and you have to handle those events. Somehow, it has to make sense and have a cohesive kind of plan for solving that in a parallel way. What I mean by that is, with heterogeneous cores, or with CPUs [central processing units] nowadays, you have, like, eight cores on them, for instance, with the i7 [4th Generation Intel Core i7 Processor]. If we do the kind of program we've always done, you can take advantage of one of those cores and maybe, you can launch 20 applications at once. The eight cores will kind of go between those…

Suzanne: They'll balance it.

James: Exactly. They'll do load balancing on that to have the eight cores work as best you can there. When it comes to working with a swarm of drones like we're doing in this project, you don't have the ability for a compiler to really optimize how these things are going to work together. You can't just launch eight of the applications and hope that one of the drones does that. That would be a choke point on all the communication, all the computation, everything else in the thing. You really need to understand how to program in a way that harnesses all of the hardware that you have available and the interactions between them. It really does make it more complex.

Suzanne: That's part of what makes it fun.

James: That's what makes it fun. That's why also we haven't made as much progress as you might think that we should, and be in the state that we should to handle this. Like we had talked about a little before the interview, the FAA [Federal Aviation Administration] is going to open up our airspace in 2015 for commercial aircraft, drones, whatever you'd like to call them.

This is going to help out a lot of different industries, from agriculture to…. That's actually considered to be the largest buyer of drones when the airspace is actually opened up for the public will be farmers who want to irrigate their crops, or they want to spray pesticides in a way that doesn't require them to keep hiring a pilot to have one plane that goes with this. If you think about it, if you have a swarm of these drones doing that, you could have 20 or 40 of them doing your whole field every week in a fraction of the cost you would pay someone to do this in a large plane, with less refuels and everything else.

Suzanne: As long as you've mastered some of those issues...

James: As long as someone has. I don't think the farmer is going to have to program this. A lot of what we focus on this project is having the farmer just be able to select a region they want the swarm of drones to operate in, and then have that swarm do something intelligent to accomplish the mission that you want.

Suzanne: So, that's the autonomy aspect.

James: Absolutely.

Suzanne: For the drones to be able to make some of those smart decisions about how they're going to operate within that region, for example.

James: Absolutely. It's a big leap from our current state of the art in that the way we currently control any kind of swarm or group of entities is to select them individually and make them do something. So, if you can imagine a farmer trying to tell 20 drones what they would like to do with their crop. They would select individual waypoints along the edge of their farm, and they would do each row individually like this. That can be a lot of time.

If they mess something up, who knows what that might mean. It might mean the loss of a crop. It could be various other things. What would be nicer is if you could actually code in the intelligence for what the farmer wants and all they have to do is select a region and say, "Go."

That's the kind of thing we're kind of looking for to help out, whether it's search-and-rescue people, ground crews, for them to be able to say, "We know that buildings are in these locations. We think human survivors might be there. Drones, do what you do. See if you can find someone with your thermal cameras, with your regular cameras. Let me know if you find something interesting, and report it back to me. And, I will decide if I should send more people over to help out there and to see if there are people under the rubble," for instance.

Suzanne: So, the communication aspect of this is critical. So, that leads us to some of the work that you're doing on throwable, wireless access points. Why don't you talk a little bit about that? How does that contribute to this whole emerging area of autonomy?

James: Absolutely. Well, one of the things about drones that you find, especially for electrically powered drones, is they have a very short uptime. They have a very short flight time: 20 minutes, sometimes there's less, 10 minutes. It depends on how many sensors, how much CPU, and how much battery drain you really have on the drone.

So, we realized pretty quickly that, if you want to support an infrastructure, a wireless infrastructure, or a new communication medium between groups in the field—for instance, if you've had an earthquake and all the cell towers are down, and all the telecommunication lines are down—you may actually need something in place quickly, rapidly that can actually support an infrastructure for communication. So, we think that these wireless, throwable, thermal acoustic sensors—that have built-in wireless access points for ad-hoc mode—that they can actually form their own infrastructure and allow you to send messages, whether it's video, whether it's voice over IP, whether it's just simple radio chat. It could also be instant messaging, to be able to communicate between these, to say that, "We've found someone over here, we need five more people" or whatever. It could be important messages. It could also just be possible locations they should look at, so pictures that are disseminated among the group. Now, this has a lot of different uses.

This could be used in search and rescue, obviously, is our main focus, but squads and militaries would also use this for establishing a perimeter around themselves. You could use these for establishing a perimeter around a building. So, we've been talking with the Department of Homeland Security and the Department of Energy trying to find some grants for that kind of work, where we're looking at—not only sensing people but also maybe sensing other drones. So, the thing about opening the airspace in 2015 is that we not only open ourselves up to people using it for surveillance of various things, like law enforcement or whatever, but we live in a post-9/11 world. We also have to consider that people will use this inexpensive technology to target us, whether it's for "reconning" a potential target like the Boston Marathon or something like that. We really should have some measures for defending ourselves. So, we're looking at various aspects into that, and that's some of the things we might talk about later.

Suzanne: Now, this kind of communication—I'll take your example of the earthquake where we've got the cell towers and everything down—lots of different people are going to need communications resources. There's going to be a lot of data that comes at this. I'm imaging this sort of fairly underpowered—in comparison to a normal cell tower—wireless network. So, how do you deal with filtering out, and how do you deal with the volumes of data? Is it just more access points, more throwables, or do you have strategies for kind of dealing with the data volumes and filtering that's going to be needed to be able to make this a useful resource in a limited infrastructure?

James: Well, I won't claim that our project goes to solve all those problems. In the Advance Mobile Systems initiative that Edwin Morris runs, we really have a lot of different technologies that are really aimed at tackling the tactical-edge problems. So, the Wi-Fi ad-hoc mode networks really only support hundreds to maybe thousands of devices. You really can't scale it much more than that. So, if you do have a situation where you're trying to support millions of people, for instance, in a displaced kind of way—a major city is hit, like New York or something—you're going to have blackouts, just as well as the cell towers. But, you might be able to use other technologies we have at AMS, like the cloudlets, like other kinds of forward-provisioned elements that would be able to—you'd have Wi-Fi ad-hoc kind of networks that are connected through wired connections. That way you can scale up the number of devices you support without actually everybody being on the Wi-Fi ad-hoc mode kind of thing.

So, one of the things we do in the project, though, to try to do this, is—to support the volume that we think is going to be in this—is that we try to operate in reduced-information modes that maximize the utility of the bandwidth that we have available. So, like for on the quadcopters we're using, we have a control plane and a data plane. What I mean by those is they're just radios. They're broadcast radios. The control plane really has a very long distance that it can communicate over, but it has a very low bandwidth. That's just kind of the tradeoff you take with radios.

So, we try to use that medium for the collaboration between the quadcopters, to be able to say, "I'm going to go to this location. Maybe I've found someone over here. This is who you should look at," something like that. Then we have the data plane, which is the high bandwidth but short-range Wi-Fi access point. Those are more for transferring video or images or thermal images or audio or whatever else you might need that needs that kind of megabits-per-second kind of connection. So, we try to make smart decisions about what we send. By default, for instance, if we find a person at a location, we'll only send the metadata that locates that there's a GPS ping at this location, for instance, that has a likely person there, with a certain probability attached, that you're going to find a person at this location. Or maybe you might find 20. There might be a building that you're at that has a lot of people that were there and need help.

We try to send that data, and then the user would request a higher feed from that location. They would try to set up a network bridge back with throwables or with more drones. The drones can form their own network bridge between each other, and that would dictate the bandwidth available back to the person.

Suzanne: So, you're using the control plane as a filter to try and minimize the number of places that you've got to have the high-bandwidth access?

James: I think that's a fair way to characterize it. We do have other kinds of filters on there that would—so, I mean, "filter" is an over-used word, really, in our field—but we try to do a lot of filtering on it to make sure that we're utilizing the bandwidth as best we can.

Suzanne: So, that's a way that you're trying to maximize the resource. And, that's actually part of the parallel aspect of all this.

James: Absolutely. Absolutely.

Suzanne: This is very cool stuff. I think our listeners should be very interested in this research. I also understand that this research naturally involves several collaborations. Do you want to tell us a little bit about the people that are collaborating with you on this project?

James: Absolutely. I think we have some of the top minds in the field, really, working on the electrical and computing engineering, especially, aspects of this. Dr. Kenneth Mai and his graduate student, Tom Jackson, from the ECE Department. That is Electrical and Computer Engineering Department at Carnegie Mellon [University]. They're working on that throwable thermal and acoustic sensor, with the access point built into it. Right now, we're coming up with the first prototype, so I don't know if "throwable" might be the right word for it. "Rollable," maybe, or "placeable" for the first prototype? We're trying to be ginger with it.

Suzanne: Do you know how big they're talking about?

James: Right now, it's about at the size of a softball, but that's just the first generation. It will come down in size. They're actually doing most of the hard skinning of it by hand, rather than fabricating it, so that's going to be a lot bigger.

Suzanne: A little bit bigger.

James: A lot bigger. But yes, Kenneth Mai and Tom Jackson are working on that. Then, we're also working with Anthony Rowe and several students on a project called Drone-RK, which is an open-source, middleware designed for the Parrot AR Drone 2.0, which is a very popular mass-produced quadcopter. It's like $300. You can buy it on Amazon. So, we wanted to stay in a price point that supports search and rescuers.

One of the hardest things to do is to find a platform that does everything you need to do at a price that people could afford that would be using this. If you're talking about a fire department—an all-volunteer fire department—their yearly thing here cannot be buying 14, $50,000 quadcopters or whatever. There are those from Lockheed Martin and various other places who make some very robust quadcopters and fixed-wing airplanes that would be able to support these guys, but they're very expensive. Fire departments can't afford that; wilderness departments, you know, emergency management services, various other people, they can't afford these kind of price points. So, we are trying to keep it within the footprint we think people can afford. That causes a lot of problems with development.

So, Anthony Rowe, he's working on the side of trying to make these commercially, mass-produced, inexpensive drones work in a way that can be useful in a swarm. And so, adding GPS, adding thermal sensors onto the drone, various other things like that, to make it so that it will actually accomplish the mission task that you want them to do. We've got several students. We have interns right now from Vanderbilt University and also from Carnegie Mellon. There's five different students working right now in the summer. Then we'll have Ph.D. students and master's students also coming in the fall to continue the project. So, there's a lot of good stuff there.

Suzanne: So, lots of different ideas coming into this?

James: Absolutely. This is, like you said, we think this is a very interesting project. We think that a lot of people have responded very well to it. We have had interest from Tim Voss, from Lawrence Livermore National Labs, to work on a security aspect of this, where you really don't want people to have unauthorized access to send whatever commands you want. So, you know, even for like a farmer. He doesn't want his next-door neighbor to, for instance, take control of the swarm and water his own crops or something like that.

Suzanne: That would be bad.

James: So, Tim Voss, he has a lot of experience, and his team has a lot of experience, over at the HOPS laboratory, with security layers and authentication schemes and things of nature that we think might be able to integrate into the swarm idea and have a semblance of security. There's a lot of interesting things. There's a moving target software defense kind of layer that has it so that the swarm—one of the things we have planned for 2014 to investigate is that the swarm dynamically responds to what it believes are intrusions. So, it might change the complete algorithm and what variables it's looking at between the swarm and various other things to accomplish the same mission task, but in a different way that can't be predicted as easily. And then change that frequently, so that they can't anticipate it. And we think that's…

Suzanne: Sort of like changing your route from home to work to be more safe.

James: Absolutely. So, if you think of it like that, on your way home, you might take the bus one day. You might take a car. You might walk. You might do whatever. And, you might take a different path to each one of those. So, try to strategize those different ways of going home and make it unpredictable. And this can be done with cryptography. This can be done with a lot of different ways. At the core of it, you really have to have a flexible way of programming the drones that can understand that it needs to accomplish a mission task, and it doesn't necessarily matter which one it uses. They are certainly more efficient ways to do certain things, but they might be able to change this and be able to actually do...

Suzanne: The most efficient isn't always the safest.

James: Absolutely. Absolutely. The middleware we're developing on this project for real-time artificial intelligence is actually available as open-source, as are the area coverage algorithms and the network bridging algorithms that we're implementing in the MADARA [Multi-Agent Distributed Adaptive Resource Allocation] system. [The algorithms are available at the open-source SMASH-CMU Google code repository.]

Suzanne: And we'll include lists of some of those links on the transcript for this podcast. Well, James, this work I think does promise some amazing breakthroughs, both in the use of autonomy, in partial autonomy with human-in-the-loop, as well ad-hoc mobile communications. And, I look forward to seeing how this work evolves. We'll probably have you back in a little while, maybe next year, see what's going on with the Lawrence Livermore project. But I really want to thank you for joining us today. This is very exciting work.

James: Thanks, Suzanne; it's been a pleasure. Thank you for having me.

Suzanne: If you would like more information, on James' work on autonomous systems and ad-hoc mobile communications, you can check out his blog post. Just go to blog.sei.cmu.edu. In the right-hand column, click on James' name, James Edmondson, or the Artificial Intelligence tag. For more information on the research James' team is doing in pervasive mobile computing, please see their "Our Work Site," at sei.cmu.edu/mobilecomputing/research/index.cfm. This podcast is available on the SEI website at sei.cmu.edu/podcasts, and at Carnegie Mellon University's iTunes U site. As always, if you have any questions, please don't hesitate to email us at info@sei.cmu.edu.

Editor's note: Since the time of this interview, the collaboration with Lawrence Livermore has been postponed. Instead, James's group of CMU professors, students, and SEI researchers will be focusing on new algorithms for area coverage and network bridging, quality-of-service in swarm artificial intelligence, and software model checking of distributed algorithms that focuses on formally proving collision avoidance in a swarm of drones.