AFRL AUTONOMY TRANSCRIPT PART 2

Jim Overholt
Hi! I'm Jim Overholt, I'm the senior scientist for autonomist systems at the Air Force Research Laboratory, today I'm joined by Kris Kearns my friend and colleague at AFRL and were going to tell you about the S&T strategy for autonomous system.

Kris Kearns
And like Jim said, I'm Kris Kearns I'm the portfolio manager for AFRL and about two years ago we set down on a journey to establish what is our vision, and what is, what are the goals. And so our vision is all about creating intelligent machines that are seamlessly integrated with people, with our operators. And we broke that down then into four goals. The first goal is about highly effective human-machine teams. The second goal is about teams and machines working together. The third goal is about working in complex environments and performing their mission in complex environments and then the fourth one is all about how we test, evaluate verify and validate. That these are behaving in the way we want them to behave. Today what we are going to do is dig down deeply into the human machine teaming goal and give a little bit more understanding about what do we mean, what do we say, how do we say and how do we envision this human machine team coming together. And so, what we've done is we took each one of these goals and we've identified what are the enduring problems. And these enduring problems become long term challenges. There not, there not technology programs that I can put into place in the near term, I can fix it, I can solve it. But these are things that we are going to have to deal with over and over again. And were going to have to spend time figuring out and improving the solutions that we came up with. So, under these highly effective human machine teams, we came up with three enduring problems. The first problem is about what we call enable calibrate trust. And really I think what we clue on is it's really about calibrated trust. So it's making sure that people can trust the technology, trust this intelligent machine, based on what they want it to do, what it's supposed to be doing in the particular environment in which it is operating. So it isn't about just having trust and using it technology but it's having the right amount of trust in the particular situation it is operating in that mission context.

Jim Overholt
And that, and that's a great point and its one of the key areas that researchers are looking at now in terms of trust and there, and there's two sides of that. There is the lack of trust in systems, will this system perform in a way that that I'm expecting it to do and certainly that's a concern especially as we start rolling out these kind of systems. An equal concern and in many ways potentially even more dangerous the over trust of systems. So this notion of calibration that Kris talks about is absolutely critical cause we have to have the right amount of trust for the particular situation were in, and this is one of the key enduring research problems that we've looked at in the human machine team area.

Kris Kearns
The second then enduring problem is about creating this common understanding and shared perception. So if we want humans and machines to work together what they have to be able to do is understand what is going on in their environment be able to communicate that to their partner

and do that in a means that they can effectively communicate and share what's going on and understand what's going on in their environment.

Jim Overholt
You think about humans, you know you and I working together, we start to learn each other's traits and tendencies, and trying to get machines to start understanding how humans operate, what they, what they mean, and their body motions, their little jesters. This is a real critical area. And if you're going to have humans very tightly coupled with an intelligent machine capable of making decisions they have to come to some common understanding about what are those factors about the particular time in order to make that decision. So if there talking two different languages and we've all had to deal with working obviously with people from a foreign country, we don't know the language, it's extremely difficult to try to come to some common understanding , common consensus. So it's also a problem when it comes to computers, systems, intelligent systems, as well as humans.

Kris Kearns
And so the third goal then is about, and you kind of led into that a little bit, it's about creating an environment. So that you can have this effective and shared decision making. And I always like to go back to the example, today the way we interact with technology is we are taught how to use keyboards, we are taught how to use mice, out pilots are taught how to use a joy stick, not necessarily the most intuitive manner in which I interact. In so, I use my hands, I nod, I give you a lot of nonverbal communication. I use, we use voice, we use touch, we have lots of ways we communicate. Creating this ability to interact with a machine in the same manner, will I think improve the ability for us to operate together. People to operate with machines, these intelligent machines.

Jim Overholt
And you know it's interesting from a robotic standpoint one of the things that's really interesting of course. We know about Seri, we know about the different things that understand what a human is saying. Some level, but there are other ways of communicating like you said so were interested, can we give hand signals, can you notice something that's going on with a human in terms of maybe raising an eyebrow, or the change in their body posture. It's also in this environment for this effective decision making has rapidly got to take in all the different data streams all the different attributes of a scene and it has to be able to coalesce those in such a way that the two entities or the multiple entities making a decision have to rapidly be able to, if you want to call it digest all that information and make the decision. That isn't easily done. Think about a bunch of people in a room trying to come to a consensus, when they are all looking at different papers, and there looking at different types of data sources. How do we do that now with machines and humans and making sure again there is a trust between the decision and making sure that we have this flexible effective decision making?

Kris Kearns
So what we've done is we set out and tried to explain about these enduring problems what our vision is, but let's go a little deeper. So, if you look at the world today. Today we use machines as tools. They are more advanced than a screwdriver, but they are essentially, they are a tool. And it's because of the things that we've talked about. I don't know what it's doing. I can't

understand what decisions it's made. I don't' understand what it's taking in. And I have to work and interface through with a keyboard, with a joystick, with a mouse. And so, in the world of today we have machines as tools. What we envision is a world of tomorrow, where machines are our teammates. And it's a seamless partnership between the machine and the person. So we have things like more institutive ways to interface, we can use things like, I can talk to it, I can touch a screen, and I can say go here. It can pick up on my nonverbal communications that I give a hand movement and it means something. So those kinds of interactions. We start sharing an understanding wants going on in the environment. It doesn't mean that we have the same understanding, but we start sharing what our understanding is. Between the two we can make more effective decisions.

Jim Overholt
You know you're hitting on one of the parts you'll see us talk about more and more. When you start coupling that machine, and you start coupling the human together. There needs to be, we like to refer to bi-directional information flow. This bi-directional information flow idea is currently we get most of the information from the machine. We see a future in this vision that's says the machine is able to start pulling information off the human. And that's going to be as we go through the rest of this segment a real key critical piece for where were going. Understanding of the physiological, behavioral, emotional context of the human, and how do we couple that along with the information processing with the machine. These are key areas. And they are very unique in terms as we start launching out into our vision in terms of implementing human machine teaming. The one thing that we really try to highlight then is, this the notion that this effective decision making shared common understanding is an absolutely critical piece. It's that enduing problem space that Kris talked about, that we need to actually work on over these next years in order to see what that vision is. In order to see bi-directional flow, in order to see this kind of shared perception between the two. These are key-key areas. That we are working on.

Kris Kearns
And so, we think the way to get two effective autonomous systems is through this human-machine teaming that we've talked about. And when we say what does that mean. It includes things like collaborated trust that we talked about. It includes training together. It includes that common understanding shared perception, and this bi-directional flow of information. And when we think about that in the context of the Air Force mission it isn't a particular application, it isn't a particular mission that the Air force executes. It cuts across the entire domain of what the Air Force does. So things like these UAV unmanned vehicle operators. It includes command and control aspects our satellites, and our space operators, our weapons operators, our ISR analysts. The guys that are taking in the intelligence and making information out of it that you can make decisions on. So, there's a broad spectrum of just about everything the Air Force does that's an application space for this human machine teaming and bringing the two together, in a collaborative manner.

Jim Overholt
It really hits on one of our earliest points autonomy is not just about platforms. It is about all these different domains, and all the kinds of missions and all the kinds of jobs that Airman have to do currently and in the future.

Kris Kearns
So we have to then say this is all great. This is all nice to hear about, but what really do we need to develop from a technology perspective. And so we have taken this and we have broken it into I think we have five key areas. That we think are important to focus on now. That will enable us to make progress towards these human machine teams. And so the first one is about the human state sensing and assessment, and I think on the surface this one will seem odd to us. But let's think about it this way, I'm driving my car down the road, and I'm using the navigation system on my car to get me to a location that I've never been to before. And as I'm within the last thousand feet, my car, pops up on my navigation screen change your oil in two thousand miles. Well, wouldn't it have been great for my car to know she's stressed out, she's trying to figure out where she's going, maybe I can wait till she's not in this situation to say, you need to change your oil in two thousand miles. In order to make something like that happen, we think what we have to do is, we have to provide data on what's going on with the person. So if the machine could sense I'm under stress. If the machine knew that I maybe I'm even sweating a little bit, my heart rate is up. These are the kinds of sensors that we could then tell what's going on with the person. In some Air Force missions what we would also do then is incorporate how are they performing duties that they need to be performing.

Jim Overholt
So the assessment piece.

Kris Kearns
So if we can sense them, and then we can say how are they doing? Then the machine will be able now to team with me, and be more able to support me in what I'm trying to do.

Jim Overholt
You know this is not something that's science fiction or something that's brand new. You look at what's going on in personalized health rate now and you see people taking different devices and their measuring their heart rate, their measuring a variety of different data. Their trying to use that information in such a way that they are trying to improve their health. And these are the kind of things that we are looking at harvesting, researching to see where these kinds of devices could potentially give us some kind of information on the physiological aspects of the human, looking at the behavioral aspects; these are the kind of things that that particular area really dwells on.

Kris Kearns
Right, so we can pick up these technologies now, and we can militarize that we can put them into a military environment and then we can figure out how we can access what's going on with our operators. So that now we can interface the people and the machines, which brings us to our second area which is how do we effectively interface humans and machines. We talked earlier about we all learned how to use keyboards, mice and joysticks. Well, we can do better than that. Technology exists today that I can touch a screen and I can point to things on a screen and I can interface and interact with a machine that way, being able to communicate by voice is coming up to where we might be able to utilize it in our environment but in the future, think about 20-30 years from now absolutely we ought to have systems that I can talk to, and I can direct by verbally. In addition to we're starting to see things coming out of the gaming community on

being able to understand what this means. So improving that interaction between the human and the machine is definitely a research area for us.

Jim Overholt
We don't look at these thrust as being so unique or orthogonal to each other, in other words this idea of sensing the human ties beautifully into how do we, Okay, start developing interfaces between humans and machines. And this is exactly where this is going. Monitor and access the human, have him integrated with the kind of command and control that he's got to do for some kind of intelligent system. Really key.

Kris Kearns
Right, so then, in order to do that the next step becomes if I want to share my work, if I want to allocate tasks, then there's a fundamental underlying science that allows us to then say what are the tasks, how do I break them down, how do we think through things, and then how do you put a human and a machine together, so the that the human can say you do that, and I'll do this. The machine to be able to say I can't do that, and being able to share the tasks and cognitive work that needs to be done in most of our missions.

Jim Overholt
So that cognitive tasks analysis, cognitive science starting to like you said, look at work how work flows, how the humans interact with the tasks that they are being assigned, and the cognitive science part which is starting to understand how humans process information. This is a science that's been around since the 60's at the same time artificial intelligence kicked off, and the efforts right now in terms of incorporating in models of the brain, and models how we take in information, be able to process access memory. All these different things we are heavily exploring those areas to be able to see what can we harvest from those kind of research fields, to be able to put them into machines, to start trying to think and reason like humans.

Kris Kearns
So then our fourth area is about what we say is human learning, its machine learning, and then its human and machine learning. So then if you break that down, its humans learning about machines and learning how to interact with them, machines learning about that person and what are their preferences, what are the things they would do. And then now let's put the human and the machine together, and let them train together and become an efficient and effective team.

Jim Overholt
And this is such unique area as well because if we look at a future where an Airman is partnered with a potentially intelligent system, then this intelligent system we see as being personalized potentially being part of that Airman through his career, in which means the intelligence system has to have the capability to be able to know his tendencies, know his likes, understand when he's at top performance, these kind of things

Kris Kearns
Finish his sentences

Jim Overholt
Finish his sentences. There you go, these kind of things but your better than a machine. So, but these are the king of things, so when you look at this idea of the two systems coming together and learning together that is an absolutely unique new area. That again we think is absolutely necessary if we are looking at this vision of autonomy in the future.

Kris Kearns
And so underpinning all of that then is a whole lot of data and a whole lot of understanding of that data which is our fifth area. So pulling together the data that people understand, the data the machine will understand, being able to put it together, knowing that we as people understand at different context time format, machines have their own context time format, how do you integrate the two so you have this shared understanding of what's going on in the world and with a team mate.

Jim Overholt
And so we then, then emphasis heavily then depends on not only bringing the data, but varying the data in such a way that it's understandable by the entities that are going to be looking at it, being able to rapidly do that, and it still ties back to these notions that we had of sensing the human, what is the human feeling right now, how do I have to change the data in such a way that maybe they will process it in a much faster way. So these kinds of things all tie together, there not individual swim lanes, but they 're really critical it's what we feel are the absolute five areas we need to really look at and spend time and research and development over these next 20-25 years and out into the future to really realize this vision of human machine teaming.

Kris Kearns
So I think hopefully what we've done is we've given a pretty clear understanding of what we mean by human machine teaming and the kind of effective human machine teams that we want to develop. So, in future podcasts I think what we're going to do is we'll go through each of the other goals that we have in our strategy

Jim Overholt
Yeah, in a matter of fact the next one were going to talk about scalable teams of autonomous systems so this notion of taking multiple intelligent systems autonomous systems together and being able to see how they operate to handle tasks and deal with humans.