NETEVENTS

EUROPEAN MEDIA SPOTLIGHT - INNOVATORS IN CLOUD IoT, AI & SECURITY

DRAFT

Conference Debate V: Practical Applications of Artificial Intelligence from Cybersecurity to Cloud to the Internet of Things

Analyst Chair:

Rik Turner, Principal Analyst, Ovum

Panellists:

Paul Ferron	Director Solutions Sales, CA Technologies
Dean Bubley	Analyst & Founder, Disruptive Analysis
Jan Guldentops	Director, BA Test Labs
Scott Register	VP, Product Management, Ixia
Jason Steer	Solutions Architect, EMEA, Menlo Security

Rik Turner

Hello all. Yes, Rick Turner, from Ovum. I took a little bit longer there, because I thought I'd let you listen to that lovely music. But we're here to talk a little bit about artificial intelligence/machine learning. I'm going to let the eminent members of the panel describe whether or not there's a difference between those two things.

But I want to give a few words first. Let's see if I can do this. Ah yes, there's me. See, I'm now a Principal Analyst, not a Senior Analyst. I got it just two weeks back, I was finally promoted, but too late for the documentation here.

Anyway, I wanted to talk a little bit about the situation that we find ourselves in cybersecurity and why therefore application - oh sorry, artificial intelligence might have a role to play and what role it might be, at which point we'll throw it over to these folks.

So here we go. Essentially, over the last two decades, the IT security industry, both the vendors and the practitioners, can be said to have suffered a fairly drastic defeat, and I don't mean that in any way, shape or form to suggest that that's entirely their fault.

But if you went 20 years ago and talked to folks, as I did, in the IT security space, they were talking very commonly about preventing. I mean, prevent was the dominant word in their vocabulary. We can prevent breaches. We can prevent things - bad stuff happening.

Today, they're basically talking about detecting, mitigating and remediating, I've got responding. That is inherently a more conservative stance, I think you'd have to agree. And I mean that with a small C, that conservative. I know we're in election week, here in the UK, but I'm not making any advert for any particular party here.

But in other words, it's rather like being driven back - if you were defending Byzantium, the walls that were right round the outside, finally they got breached, so suddenly you had to go back into the next layer of walls, right. You've now adopted a detection stance. You know the folks are getting in, so now you've to do is detect it as quickly as possible and mitigate, where you used to, 20 years ago, be able to talk about being able to prevent. So that's already some kind of a defeat, right.

Back in the good old days of signators, well, basically, we just - we had to find a patient zero, the unlucky so and so that was indeed affected, who might be, let's say, a Symantec customer. Pick on them. They're the largest. He would then say to the company, look, this has happened to my machine.

They would take that code, they would analyse it and they would come up with an antidote - or to be more exact, actually - I've used the term antidote - perhaps more correctly, we should say, is anybody here old enough to remember identikit pictures? It was - what, actually, are signatories is a picture of the code that is attacking. And you say, right, if anything looks exactly like this, block it. That's how those signatures worked. And then they just had to disseminate it to everybody else in their network.

So anybody remember turning their computer on at home and then having to go and have a cup of coffee for 20 minutes or half an hour while McAfee downloaded it's - all of updates for that morning? That's where it got to by the mid-noughties, because there were so many attacks, so many signatures had to be updated all the time.

And of course, now we know, thanks to Symantec - I mean, it's some two years back in the Wall Street Journal that told us that 45% of their signatures were only picking up - or rather, their signatures were only picking up 45% of the malware that they were - that was attacking on a daily basis. And that's only gotten lower. I mean, I'm using 30% to 40%. Maybe it's down to 25% by now. It's on the wane.

In other words, signatures - it's true, AV still has a value, if only because it gets rid of the common or garden junk that's coming at you on a daily basis. I'm not advocating

the end of AV, I'm just saying, AV is just to reduce the noise. Then you have to worry about the stuff that evades all of the signatures.

So, surprise, surprise - why is that, first of all then. Let's ask a bit about that. I mean, and why I say it's not entirely the industry's fault.

Okay, first of all, there are just so many more malware actors out there, compared to 20 years ago. I mean, the number of people that are actively employed in the malware industry is huge. Be they - and primarily criminals, we've heard earlier today. Vast majority of folks are out there because they can make a buck with it.

But we've also got the hacktivists, let's not forget, Anonymous and the al-Qassam freedom fighters, or whatever they were. Cyber fighters, I think. The state-sponsored groups are clearly very powerful institutions. There's so many more folks launching. I mean, there are - the labs that are out there are talking about 500,000 new pieces of unique malware daily being attacked - are being found, and therefore launching attacks on a daily basis. It's astronomical.

It's also, there's a very active market on the dark web. You can go on to the dark web these days and buy a rootkit for like 10 bucks and do a couple of little tweaks so that it doesn't meet any particular signature - it's therefore, basically, gone invisible, and you can launch an attack with it.

I mean, it's not even the case that there's - everything is no so sophisticated. No. You can get a cheapo little rootkit and launch an attack. You can be an idiot yourself. It doesn't - you don't have to do anything particularly sophisticated. It's just the barriers to entry have been reduced so much.

And of course, the cloud makes it so much easier. I mean, the great Dyn attack, as we all know and love, basically, that was the botnet of things, right. That was all of those CCTV cameras that could have something put on them that then, wham, we launched the Dyn attack and took down half of the Internet of the United States for 24 hours, or whatever it was. So it's gotten a lot easier, right, and therefore it's a lot more difficult to defend.

Surprise, surprise. With the wane in signatures, up have turned brand new acronyms. Brand new the next best thing this week. Sandboxing was really hot. It did great things for FireEye. They had a lovely IPO and everything was - and it was going to be Sandbox was going to solve everything. Well, of course, it didn't. And so now we've got the next big things coming up.

Most of these are Gartner acronyms. I make no apologies. They are my far larger competitor. But hey, they create the acronyms, I'll go with them. I don't care.

User and event behaviour analysis. So that's all sitting on a network and watching what everybody and everything does on that network and saying, that's anomalous, that's strange.

Endpoint detection and response. Clearly, if stuff has got in, if we haven't picked it up, prevented it at the door, then it's on. Let's try and pick it up and see if we can respond as quickly as possible. It's that whole move that I spoke of earlier, about the move from

prevention to detection and - well, response, which can be in mitigation, in the first instance, remediation beyond that.

And then finally, AI machine learning. I make no apologies here for mixing those two. The industry does it quite barefacedly that's the cybersecurity industry. You walk round the RSA conference or tomorrow's info security, everybody will say they're doing AI. Everybody will say they do it - machine learning, whatever they mean by it. And good luck to you - you'll have to sort it out and sift through and find out who's really doing what.

What I think - oh, and of course, finally, sorry, threat intelligence. This is another buzzword for the last few years. This is essentially I no longer know who's going to rob my bank, therefore please provide me with photographs of every known bank robber on the planet.

Now, my bank is in London, so please could you then narrow it down so that all - it's only the bank robbers who operate in the UK - that's - we're starting to move towards threat intelligence here, rather than just threat feeds. And then, beyond that, could you also make sure that these are only the bank robbers who are actually out of jail at the moment, and so can rob my bank?

That is a good approach. It's very nice, except for the fact that of course, if tomorrow, Joe Schmo, who's never robbed a bank, decides to walk in a bank. This - the famous unknown unknowns of Donald Rumsfeld - God, I hate quoting Donald Rumsfeld.

But in any case, those - all of this stuff is going on. And of course, threat intelligence can also be thought of in a way as big data applied to cyber security. Loads of information, which I've got to sift through, to try and find out what's going on, because my signatures are stopping working.

Right. What do I see as the role of AI/machine learning? They'll tell you what's what. But what do I see as it? Well, first of all, it's got to help me model what is normal so that beyond that I can then - it can alert me as to what is abnormal, what is anomalous.

Also, it can help me threat - filter out the noise in threat intelligence. I think that's a point that Scott was talking to me about earlier, and he's going to probably pick that up a little later.

Filter out the noise is going to be so vital. I was talking to some banks recently that were getting 200,000 alerts every day. Isn't no way you're going to be able to filter through that with anything other than automation, for a start-up. And of course, hopefully, to find all the hidden threats.

That is the disclaimer, at which point, please could somebody put up the list of all of these folks so that I can move on?

I'll start off just giving - if anybody hasn't been already on a panel, please put their hand up. Yes, there are one or two. So I'll go - we'll go through again and do the panel first, then I'll start firing some questions off. So…

Paul Ferron

Very quickly, Paul Ferron, Director for our Digital Identity Strategy in EMEA.

Dean Bubley

Dean Bubley. I'm an Analyst and Futurist and I run Disruptive Analysis.

Jan Guldentops

Jan Guldentops. I run a small test lab called BA - not to be mixed up with the real BA.

Dean Bubley

They did an extensive test this week.

Jan Guldentops

Yes.

Jason Steer

I'm Jason Steer, CTO for Europe for Menlo Security.

Scott Register

And I'm Scott Register, VP of Security Products at Ixia.

Rik Turner

Excellent. Well, I think I've got to start with anybody who's got futurist in their title, so Dean, could you please give us overview of what about the applications of AI in the first place, more generally, and then if you want to go on to a little bit more about and what's that - what's the difference between that and machine learning? And then finally, what do you see as the applications of it within cybersecurity?

Dean Bubley

Right. A very, very quick overview. Artificial intelligence is a pretty much umbrella term, which covers a whole range of techniques and applications for using compute capabilities to spot patterns.

Now, there's a bunch of both methodologies and examples of how that gets used, so some are a little bit outside of this - things like machine vision, natural language recognition, which you might see in chatbots, for example, or machine vision increasingly where you've got computers categorising images so they can actually spot the cat, rather than a lion, for example.

But I think the most interesting thing over the last few years, which is actually underpinning a lot of what's being talked about now, is deep neural networks, which is made famous by Google's DeepMind acquisition. And essentially, it's a way of using the unstructured - or finding patterns in unstructured data, by the machine essentially teaching itself in terms of observed data.

And one of the characteristics that is often applied in - whether that's cybersecurity or other use cases - I look at the uses of AI in the telecoms industry, for example, is having a large data set to do what is called training for the machine, on a whole range of different variables.

So that could range, in this case, on threats, or elsewhere in the telecoms industry it could be looking at the patterns of the heat in cellular towers to look out for impending failure of machinery - so predictive maintenance, in a whole set of sectors is an important use case. Or it might be looking for patterns in usage and behaviour, either to look for security, or just because they're trying to predict churn in a particular mobile user.

And often you'll find that algorithms are referred to as being trained in the cloud, and then used on the edge. So you basically develop the model for a particular thing you're looking to apply AI to, in the data centre, using a data lake of the existing - pre-existing data, but you actually push that model down to, well, whether it's an endpoint of one form, or increasingly, it could be even a small IoT device, if they are capable of supporting that level of compute.

I think, quickly, just to say on security, cybersecurity, there's a - clearly an awful lot of things around the pattern recognition side for malware. That's outside of my normal focus. But one of the areas that I'm conscious of at the moment is data integrity protection, which is a new class of security risk, which is not so much about interception of data, or - but so much of the risk of it being changed by someone.

I don't particularly care if someone hacks Fitbit or even my account, and finds out my heart rate is currently 69. I do care if they change it to 169 and my insurance company takes a dim view of it.

So there's various techniques that we are looking at, not just with AI, but also blockchain technology, to bring in another buzzword, to protect and give what's called non-repudiation and chain of custody of data, to make sure that whether it's a data point like that, or whether it's a piece of software, is exactly the same as what it's supposed to be, and is not a subtly altered version of it.

I'll leave it at that for now.

Rik Turner

Scott, could I quickly ask you, because there is the case, definitely, that a lot of the vendors on the floor of RSA, virtually all of them, were quite readily mixing machine learning and AI as if they were interchangeable synonyms, and I'd like to get that from you, what you…

Scott Register

Yes. I would certainly argue that what we're doing now, in network security is not really artificial intelligence, it's machine learning. I mean, it's pattern recognition, and not that there's anything wrong with that, but facial detection is really based on machine learning.

Faces don't really change often. I mean you get moustache or not, but everybody has two eyes and they're horizontally aligned, et cetera.

What we're doing, but network security, cybersecurity, that doesn't really work as well. I mean, that thought just because things change very quickly, like attacks today look nothing like really, attacks 10 or 15 years ago. So there are certainly applications, and they're good, for pattern recognition.

I think a lot of - there are a lot of endpoint tools now that can take basic analytics information, telemetry information off of endpoints and very quickly say, oh yes, this looks suspicious. This like a breach. And that's detecting a pattern of other system calls or network communications, but they're not really - nothing is really good at very novel things - an attack type you've never seen before. And that's - if we can get there that would really be artificial intelligence.

Rik Turner

Right. Any other thoughts on that one? Anybody? Right now? I mean, obviously - yes, carry on.

Jason Steer

Well, certainly, when I worked at one of the vendors you've mentioned earlier, Rik, FireEye, yes, the sandboxes were claiming to be the future, but they're based on algorithms and development and understanding of known good and bad, and I think this is where even machine learning, it's still having to be modelled on something.

And as we know, every single day, modelling only gets you so far and it's fail to plan, plan to fail is - I think there's - again, going back to the rhetoric of a previous discussion here on vendors overselling is customers feel that they have a bulletproof jacket on now that they've bought the latest and greatest endpoint solution, therefore I can move on.

And is that, I guess, absolution of responsibility to somebody else now I can go on and solve another problem, because I've bought a product that solves it. And I think here's - the problem is, no product in its own is going to solve the problem. It's going to give you the flak jacket that says, we're not going to be breached.

Then, without saying any more and boring everyone, is - that's the mind-set we've got to get to, is we're going to be compromised, just like we saw with BA, we've seen with NHS, is what's our plan? Because this is going to go wrong at some point, and we don't have a plan.

Paul Ferron

It's the general problem. There's no such thing as absolute security.

Jan Guldentops

No.

Paul Ferron

And, I mean, that's what customers expect, because the customer's also in fault. He really wants to buy that magic bullet, while in fact, he should be buying something that's going to help him in 99-point-whatever% of the situation.

Jason Steer

But this comes back to a broad question, which is, are we secure? You may talk to many CSOs, CIOs, the Board is saying, are we secure, yes or not? It's a binary question. It's a binary answer. But in the world we live in there is no binary of yes or no security. Security is, what's my budget? What level of risk are we happy to accept?

Jan Guldentops

I mean, I fully agree with that, right, and part of it, it's the vendor's fault, right. We've done a really good job at selling it as the magic bullet that's going to solve everything. But the reality is that - and if you go back in history, it's all explainable and why, but the reality is, it's very black and white.

A lot of the stuff that we do is black and white. It's ones and zeros. It's either it's secure or it's insecure. Either you're on the network or you're off the network. And that state of mind is just no longer maintainable and it needs to change. So we need to go to a more granular approach.

And I think that area is where machine learning or AI, or any of those mixtures - mixes can add a lot of value. Because the more information I have, the better I can make a decision. That's universally true.

Dean Bubley

I mean, the big danger, as well, is that there are poachers and there are gamekeepers. So at the same time we're improving the tools for security, you can bet that a similar set of techniques are being used by the bad guys as well.

Whether that's in terms of looking for patterns in data to try to draw inferences about - it could be - we were talking before about passwords or anything like that. They might not even necessarily need to have the actual data. They might be guessable, after you spot a pattern.

You might find that AI has all sorts of other strange exploits - it creates exploits as well. How long before someone - one of us gets a phone call saying, ah, I'm just calling about the bank loan you agreed last week. What bank loan? Oh, this is the recording we've got of it.

When you start having language recognition, you may well find that never mind fake news, you have fake everything, and you start having to try to prove the authenticity of data, of recordings, of anything else.

Jan Guldentops

It's one of the strange thing we see. Everybody here talks about innovation and solutions. For the other side, it's also innovation. It's also research. There's nothing more fulfilling than finding a security bug in a solution you're researching.

So it's going to happen. Artificial intelligence, whatever you're going to use, as a security product is going to fail. So that's what we need to work on. And one of the things I always strike is I have a very bad relationship with ISO certifications and documents. I hate them, to read them. But in a way, every manager, every CEO should read a 27000 standard, as common sense, and it's not happening. And that's why we end up in the zero and one situation.

Rik Turner

So in other words, security is not like pregnancy - it's not a yes/no answer. It's much more complex.

Jan Guldentops

No.

Rik Turner

That's fair enough. Makes sense. So are we then saying that artificial intelligence or machine learning - certainly machine learning, from what you - from what Scott's said already today, which is pattern recognition, but further down the road, a broader understanding, say - the more broad understanding of artificial intelligence is going to have an increasing role in security? Or do we think that it's - I won't say a passing fad, but it's only got so much that it can do? What do you think?

Paul Ferron

I think it will be more increased focus on it. I mean, the ability to have more context on - because in the end, it comes down to two things, right. There's a user trying to connect to information.

Rik Turner

Right.

Paul Ferron

That's the end of it. That's all there is. There's a transaction at that moment in time between a user - it could be a service in service of that user, but the user and the information. And the more context I have, the better I can decide whether I want to allow that, yes or no.

So the more analytics and the more capabilities I have to understand that context, gives me capability. But the key thing is not just having that constant, but also the ability to

then take the appropriate next step, right? Because saying, I'm not to connect the user is just another black and white answer. That's also not a good answer.

Scott Register

Yes, I was going to say, certainly I think it will be more - and at the end of the security decision there's typically a person, right, so deciding, oh, that looks bad - I need to go clean it up. The holy grail is having it all automated. That will probably still always be the holy grail, right.

Rik Turner

Yes.

Scott Register

And so what you're really trying to do is make that person much more effective. And we have all these scale problems driven by OT and cloud, like the number of things you have to look at are much larger, so machine learning is really a tool to help you focus quickly on where you need to go.

So there are things like standardisation of input data. Common standards like, does this log come in that format? And IPFIX and things like that, that - yes, maybe Google can afford unstructured data analytics, but for most in the enterprise you really want structure and everything looks the same and you know how to deal with it. Then filter out junk noise and that - and machine learning can be very helpful there in helping you get to the end state of, oh, that looks bad - I need to take care of it. And so yes, I think that will drive it forward.

Dean Bubley

We haven't referenced IoT that much. It's also worth thinking about looking through not necessarily large data flows, but data flows from a large number of objects. And so for example, we may well see this in the consumer space with the next generation of the mesh Wi-Fi, like the Google and Amazon products we have with multipoint Wi-Fi that are managed from the cloud.

And I think there's a big problem with IoT security in that it's almost impossible to get, particularly consumers to update the firmware in their toaster, right, or their washing machine. And so there needs to be ways of spotting anomalies and shutting them down, because you're not - you simply will not get people to update the software in their various IoT devices.

And if you can essentially get the - or the Wi-Fi in someone's house, or the home gateway, to look, globally and go, why are 5 million toasters trying to access the Bolivian government's website - that doesn't look right. Flip the switch on that address, that MAC address range. I think that's - that type of level of intervention is also going to be feasible, and possibly mandated in law.

Rik Turner

Could I make - ask if there's anybody with a burning question they'd like to put to the panel? If so, please - hand up there. Thank you, sir. The mike is on the way.

Audience Q&A

Oliver Schonschek, Insider Research

Thank you. Oliver Schonschek from Germany. Well, you just said correctly when, let's say the toaster now starts to communicate with a website it doesn't look good, and we should shut it down. But from the user perspective, actually, the user just wants to have his toaster in function, so he's not interested if the toaster sends spam to someone.

So there must be a solution that the toaster still is working, while being insecure. So the artificial intelligence then would be to decide that, let's say, it can go to the smart home hub, but not to the banking site. So to set rules, depending on the function of the product.

Jan Guldentops

Yes, I sort of agree with that, because the switching off of the toaster is that black or white answer, right? It's either it's on or it doesn't work. That's - to me, that's not the approach. That is just not going to be accepted.

I've recently been talking to some large enterprises where the discussion on authentication was exactly that, right. We see patterns of misbehaviour, but the business is just saying, you will not block our customer from connecting to our systems. It's just a no-go. Right, so how do you deal with that?

So I believe it's more the - that middle ground that we have to find. The ability to take that next step. Okay, what does that mean? And if the update example of the toaster - if you look at Tesla, Tesla is doing a great job. If the car has some malfunction they can do it overnight - not to one car - they can do it overnight to every car. So, why won't I be able to build a toaster that can do exactly the same thing?

Paul Ferron

There's also another way of solving this problem. You can - that toaster doesn't have to communicate to a limited set of things. We always have an everything allowed policy these days. Right, so you can fix that quite quickly if you know what you're doing, in the industry.

Scott Register

One - I think in part the answer to that question may be though, in - like for why is Tesla on top of this and the toaster is not? For automotive - for the automotive world there's a rich body of legal precedent and legal responsibility. If your car is unsafe you get sued and you can - but in the toaster world, or the IoT world, like in the Dyn attacks, no one

was responsible. No one was liable, right. And every - not the service provider, not the I T manufacturer - and so - we really don't have time, I think, to go down that trail today.

Rik Turner

No.

Scott Register

But clearly there is a need for - and there is evolving legal framework to deal with some of these issues. We're not there yet.

Rik Turner

Exactly.

Scott Register

But at some point, holding someone responsible becomes really important in getting them to fix it. They won't do it out of altruism.

Paul Ferron

Which leads us back to GDPR.

Scott Register

Yes, quite.

Rik Turner

All right, folks, listen. We're going a bit - going any further, I'll just say that all this talk of toasters and spam, clearly folks are getting hungry, and time to cut for lunch. So thank you very much.

Manek Dubash

Thank you very much for that. Thank you to the panel. I've only one thing to say. Freedom for toasters!

No, actually, I only have two more things to say. Hold on. Don't get up yet. First of all, when you leave here, the tables will be cleared. This room will be used for the meetings this afternoon, so please take all your stuff with you.

And the second thing is that for the journalists amongst us, there's a treat. Oh yes. You can win a bottle of wine on our evening Thames dinner cruise this evening, if - if you take with you - you pick up a spinner, a fidget spinner, which I understand is something that people do with their fingers. I don't really get it.

Anyway, if you pick up one of these things at lunch and you fiddle with it in front of a NetEvents banner or in front of - with a NetEvents branded thing and you upload it to tweet on Twitter, to tweet it to @neteventstv with the hash tag, NE - for NetEvents -

nefidgetspinner - all one word, and you have a chance to win a bottle of wine this evening.

So, hey, groovy, hey? We're with it.

Unidentified Female

If we're really good, a bottle of champagne.

Manek Dubash

If we're really - if it's really good, you get a bottle of champagne to carry home. To carry home.

Right. Okay, good. I think that's the end of the plenary session this morning. I hope you have a productive afternoon, and hopefully see some of you just before we go.

[end]