AFJAGS Podcast: Episode 4

Artificial Intelligence & Military Legal Practice with Colonel Frank Coppersmith–Part 1

HOST: MAJOR RICK HANRAHAN, USAF
GUEST: COLONEL FRANK COPPERSMITH, USAF

Part one of a two-part interview focuses on artificial intelligence at large, including what some of the leading minds think about AI, its history and development, capabilities, and computing power among other areas.

Maj Rick Hanrahan:
Welcome to another podcast episode from the Air Force Judge Advocate General's School. I'm your host Major Rick Hanrahan. Today's topic is the first part of a two-part interview where we discuss the impact of artificial intelligence on military legal practice with Colonel Frank Coppersmith. Here are few highlights from today's show on the power and impact of AI.

Col Frank Coppersmith:
Today, we're going to look at computing power. We see that same computing power that took 90 years to develop, now developed every single hour, and that's really hard to get your head around. Certainly, are other lawyers operating in our space are going to be using this technology. And if we're not using this technology, we're going to find ourselves falling quickly behind.

Announcer:
Welcome to the Air Force Judge Advocate General's Reporter Podcast where we interview leaders, innovators, and influencers on the law, leadership, and best practices of the day. And now to your host from the Air Force Judge Advocate General's School.

Maj Rick Hanrahan:
Welcome to another podcast episode from the Air Force Judge Advocate General's School at Maxwell Air Force Base, Alabama. I'm your host Major Rick Hanrahan. Today's topic is, the Impact of Artificial Intelligence on Military Legal Practice. The topic is broad, complex, and evolving at remarkable speed.

The interview with our guest, Colonel Frank Coppersmith is approximately an hour in content, so we broke it up into two parts. This first part focuses on artificial intelligence at large, including what some of the leading minds think about AI, its history and development, capabilities, and computing power among other areas. In part two, we hone in on AI's role in military legal practice. Last, if you like the show, please subscribe on **iTunes** and leave a review. This helps us to grow and foster outreach.

I'm excited to introduce our guest today, Colonel Frank Coppersmith. Colonel Coppersmith is a category B reservist attached to the staff judge advocate **U.S. Cyber Command** in Fort Meade, Maryland. Colonel Coppersmith served in active duty for five years before entering the Reserve Corps and has held a number of assignments through his career. Colonel Coppersmith is the current CEO of Smarter Reality, a software consulting and development company based in Austin, Texas.

He and his team of designers and engineers support small businesses and entrepreneurs with creating products for eCommerce, artificial intelligence, augmented and virtual reality, and entertainment. Colonel Coppersmith received a BS in Electrical Engineering from the Citadel, JD from Samford University, and a Master of Business Administration from the Wharton School University of Pennsylvania. Thank you for being on the show today, sir.

Col Frank Coppersmith:
Well, thanks so much, Major. I'm excited to have a chance to chat with you today.

Maj Rick Hanrahan:
So today's topic is, the Impact of Artificial Intelligence on Military Legal Practice. This is a very large topic to say the least. The news is replete with the impact that AI will have on humanity. For example, billionaire and SpaceX CEO and founder, **Elon Musk**, has called AI "Our greatest existential threat." And the late world renowned theoretical physicist and cosmologist, **Stephen Hawkins**, stated that AI could be "The worst event in the history of our civilization," unless society finds a way to control its development.

Sir, with that in mind, talking about artificial intelligence, if we could maybe take a 30, 40,000 foot view and step back, maybe you could talk briefly just about the pace of innovation from a computing power standpoint.

Col Frank Coppersmith:
Yeah. I'm happy to. I think even before I do that, I want to dive in a little bit on both of those quotes. I think you left out, actually, the one from Elon Musk that is the most impressive where he actually talks about AI as, and this is a quote from him, "unleashing the demon." I think he believes that. And I think many who are starting to study artificial intelligence believe that because of the wide ranging of impacts that AI is going to have.

Although it's not merely that it could do something evil, this isn't about building terminator robots or scary machines. It's the idea that we are building technology that is demonstrating that it can perform many tasks that have been just focused only the things that humans could do. Things that used human cognition, and human talent, and human capabilities, and find machines that can do them not just as well but better.

When you see these concerns raised by these thought leaders, by these technologist who have great visibility into what is being built, I really think what they are seeing is a world where we begin to prefer the decisions and the choices, and the operations that machines perform, machines powered by AI because it'll be able to access more information. It will be able to take more considerations into account. It won't be driven by at least human nature, human biases. And that ultimately, we're building a world where, in many ways, humanity becomes simply the boot routine for the future AI gestalt.

And that may seem a little extreme, but that is what these folks who have such great visibility are seeing. And the reason that is gets to your question around the

growth of technology, and where things have gone. So one of the things that's really hard to believe, especially if we spend—if you think about how our experience as judge advocate's been, when I entered my very first law office in January of 1995, the list of things I did not have access to is hard to believe.

I did not have access to the Internet, or the World Wide Web because they didn't really exist. I did not have access to WebFlite because it had not been built. I did not have access to a cell phone, or to effective texting because they simply didn't have wide adoption. But we all had isolated personal computers that sat on our desks and a printer—and that was seen as just terrific technology and wonderful.

Today, of course, we're all interconnected. And what has driven that has been a growth of computing power that is just hard to get your head around. One of the ways that computing power is measured is in the number of calculations that can be done for a certain amount of money. And the standard is one million calculations per second for $1,000. Now, from the time the very first mechanical computers were built, say in the 1900s, until about 1990 or so, it took us that long, about 90 years, to where we got our very first computer that can perform one million instructions per second, basically a million calculations. And at the same time, could do that for $1,000 in cost.

Today, right now in 2019 where we sit, modern computers add that much computing power every hour. So what took us 90 years from the invention of the very first mechanical computers up through the launch of the very first solid state computers, up through microprocessors, up to the very first PCs, which came in right before I became a lawyer, that entire almost century of innovation, century of advancement, we see that every single hour.

And so when you take in that kind of horsepower, when you're taking that kind of processing power, that's really where AI begins to come from. A sense that it has access to more information than we've ever had before.

Maj Rick Hanrahan:
Yes, sir. And I think you mentioned in one of your papers that by 2020 the average desktop computer will have roughly the same processing power as the human brain. And by 2050, the same computer will have more processing power than all of humanity combined.

Col Frank Coppersmith:
That's correct. And the piece that's left out of that is the other half of that. So once we have great computing power, what we have to have is we have to have information that we can begin manipulating, information that we can begin understanding. And what has changed in the world for us has been the rise of smartphones. With the arrival of smartphones, the amount of data that we're generating every day, every hour, every year is magnitudes more that existed digitally even 10, 20 years ago.

To put that in perspective, almost all the data, almost all the digital data that's ever existed, more than 90%, has been created in the last two years. So when you combine all that digital data, location data, photographs, imagery, texting, email, all the things that go along with it, with this incredible computing power, that's where AI comes from.

Maj Rick Hanrahan:
So, sir, we're going to put in the show notes a slide that you'd provided to us by **Ray Kurzweil**. It's caption, Growth of Computing Power. Can you speak briefly to this slide?

Col Frank Coppersmith:
I sure can. So when we think back to the birth of computing, we had some of our very first analytical engines, mechanical computers, adding machines created at the turn of the last century, so around the 1900s. It took 90 years. It took 90 years before we got to the basic measure

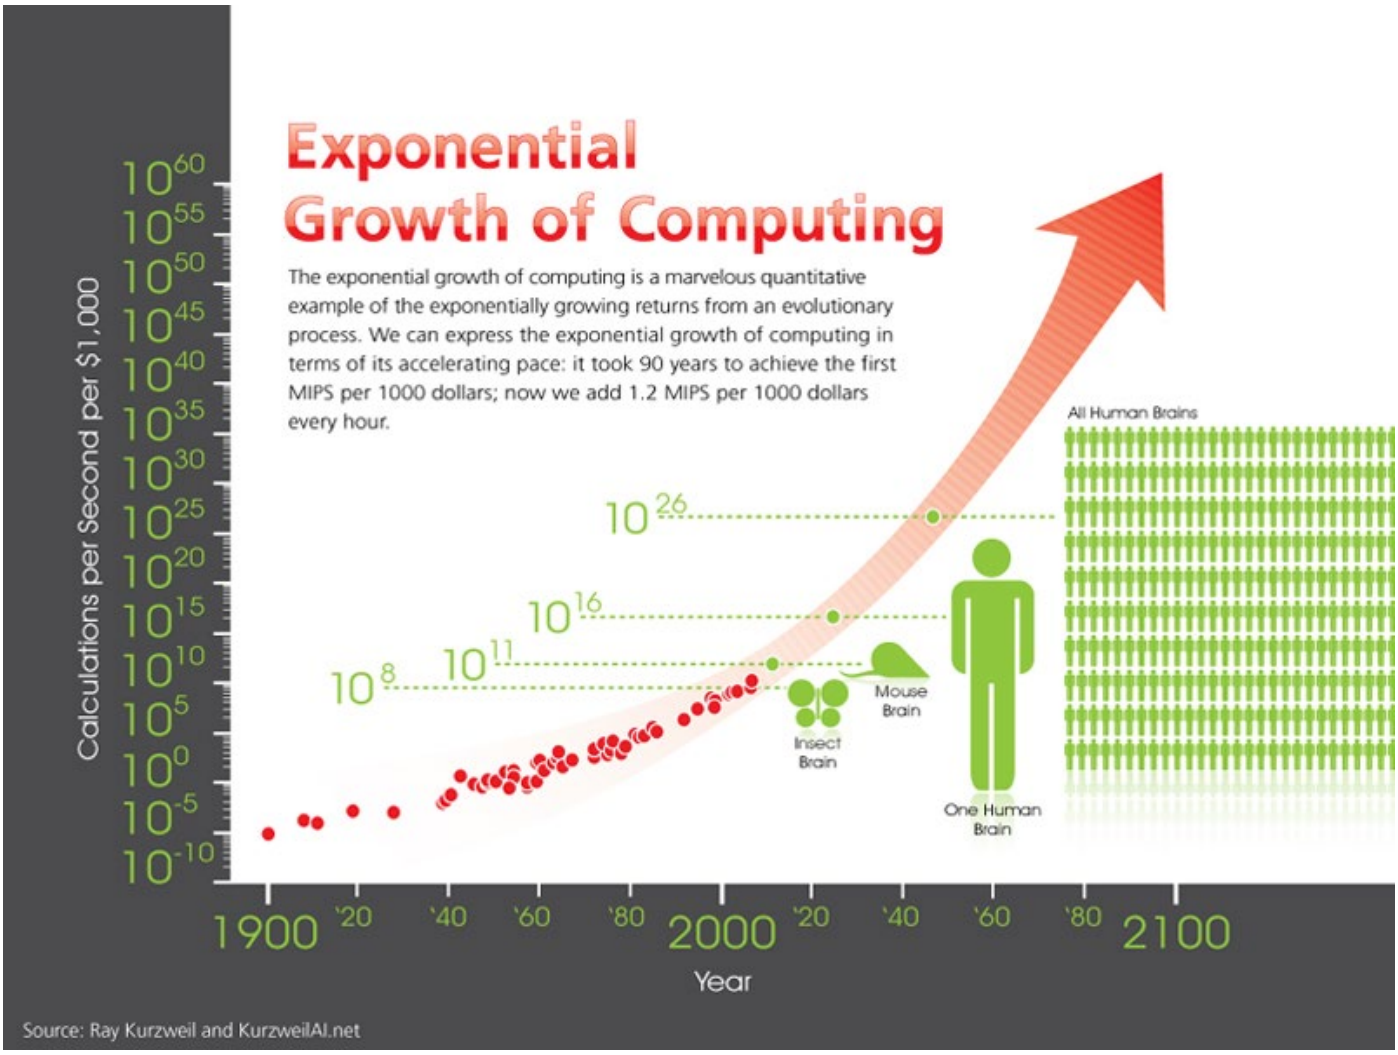

of computing power, which is one million instructions per second by spending $1,000. So 90 years to get to that. And you can see that it's a curve. You can see on the slide.

Today when we look at computing power, we see that same computing power that took 90 years to develop now developed every single hour. And that's really hard to get your head around. It's 90 years of technical innovation, 90 years of improved computing power, 90 years of improved microprocessors, and we see those improvements every single hour today.

The other interesting thing about this slide, as you look at it, you realize the scale on the Y axis scale is not a linear scale. It's a logarithmic scale. So when you look at the curve, which goes up and to the right, if it were put on a linear scale, it would be up into the right as almost a straight line straight into the sky.

Maj Rick Hanrahan:
Right. So this is exponential, an exponential curve?

Col Frank Coppersmith:
I'm sorry. You're correct. It's actually an exponential curve. That's correct.

Maj Rick Hanrahan:
Yes, sir. So with that idea in mind, what is artificial intelligence?

Col Frank Coppersmith:
So artificial intelligence is nothing more than using a machine to perform actions that were otherwise assigned human cognition, or human ability. So it's a sense that rather than you or I, or a person having to make a decision, or able to turn that decision over to a machine. And so that means artificial intelligence can be something as simple as a toaster with almost something as basic as a timer on it, all the way up to the modern deep learning functions of today.

Maj Rick Hanrahan:
And is it fair to characterize artificial intelligence into different types? And if so, what would those be?

Col Frank Coppersmith:
Yeah. So there's a host of—when you think about artificial intelligence, that's a very, very large bucket. You typically break it into a couple categories. But first, the most basic category is one that pulls in information. That's been our big driver for the longest time. This is just things like image recognition, and machine vision, and speech to text and text to speech. Even to some extent, robotics. The sense I have to be able to sense the world around me much of the same way that the human senses, whether I'm reading a newspaper, or trying to understand a stop sign in a self-driving car.

Then the second piece to it is the sense once I have all of this information, well, what do I do with it? How do I turn that into something that is something that allows the piece of software to make decisions? And that typically comes in a broad bucket called machine learning. A lot of different elements of machine learning. But for that, the best way to think about it is I am a machine learning piece of software is one that can learn and come to understand the world around it without being explicitly programmed.

So traditionally, software—if I take a piece of software, and I'll give you a good example, you take something like DL Wills, DL Wills is a form of artificial intelligence. Specifically, it's a form of augmented intelligence. So it augments human decision making. It does this by taking lots and lots of information that we know about how, say, to write a will, and it breaks into very, very sophisticated flowcharts that then prompts the user, the lawyer, to work through it in a certain cadence, and come out with an enhanced result.

But there are other types of learning in machine learning that are—I'm sorry. We call that an expert system. That's an expert system because it takes expertise and needs to be programmed by experts. The other piece to it is machine learning, often a component of that is something called deep learning. Where rather than having to explicitly program a piece of software, I can simply take data that has been curated and determined to be relevant, provide that software with that data, and the data under a host of different conditions can actually—the software under a host and different conditions can use that data to actually learn much the same way that humans learn.

You think about how your child may learn, or you learn a new thing. You learn through experimentation. You learn through reading. You learn about finding what works and what doesn't work. And basically with machine learning, and sometimes referred to as deep learning systems, we've built technology that learns exactly the same way.

Maj Rick Hanrahan:
So deep learning, Sir I think we've talked about this before offline. I think you'd mentioned that this is really where the revolution is occurring?

Col Frank Coppersmith:
That's correct. So deep learning has been around for good 30, 40 years. Deep learning is often referred to sometimes as learning via a neural network. Inside your brain right now are neurons. Those neurons all talk to each other. It's how you set down memories. It's how your child learns or children learn. Neuro networks are made of inputs, outputs, and what are described as activations or sometimes activation functions.

So the neurons are all connected to each other. They get inputs. They based on criteria that they learn over time, send out outputs that's going to other neurons. Put together in an enormously complex web of information and interconnectivity, they're able to make decisions. What has changed for us in artificial intelligence, or those are as building software powered artificial intelligence, is that we're now able to build neural networks where before we could not.

We couldn't build neural networks that were successful before because we didn't have the hardware that could make the calculations fast enough. And we didn't have the amount of data needed to actually teach neural networks. Today with the advent of smartphones and modern microprocessors, those two barriers have fallen. And now neural networks are becoming the preferred way of building decision making software.

Maj Rick Hanrahan:
And so Sir maybe to circle back a little bit from the beginning when we were talking about artificial intelligence, just the way that I think about it, I think of maybe human involvement, like you mentioned a toaster, where that would be a supervised involvement and then you have what would be, I guess, semi involved human. And then on the furthest spectrum, which you're talking about here, deep learning, is that more or less unsupervised by the human?

Col Frank Coppersmith:
It depends. So deep learning, you can take a deep learning neural network and you can either provide it supervision, or non-supervision. A supervised deep learning network is one where I want to build something that in an automated fashion recognizes an outcome. Say, recognizes a stop sign. So it's a self-driving car. While a user driving the car can note that there is a stop sign.

And if the AI does or does not recognize that the AI can know, I should have seen that as a stop sign, and use that to learn much faster. That is how most neural networks are being built because for the most part, we know the right answer. We are simply trying to build a piece of software that either comes to that answer faster, or comes that answer without the intervention of a human, for example, a self-driving car.

Unsupervised deep learning is probably where the great breakthroughs are going to ultimately come from. Supervised learning what are we trying to do? We're just replacing a person. Unsupervised learning we're trying to figure out answers and use the AI to find answers when the massive information is far greater, and far more dense than a human being could ever see. It's sometimes called finding the signal inside the noise.

So for example, rather than try to teach a car or teach a piece of software to drive a car, and so I have to teach it what a stop sign is. Instead, I might take an aircraft and fly it over untold millions of miles of hostile territory. And then maybe the AI is able to identify military trends, motion, something that we just couldn't see because there's too much information.

You sometimes hear the sense of using things like Twitter and social media to determine sentiment analysis, or to determine the risk of a potential terrorist attack. You could never review the untold billions of tweets that are tweeted every hour, but a machine sure could. You could never review the untold millions and millions of postings that happen on social media, but a machine could. In fact, they could review it in real time.

And so when it's going to have to do is it's going to have the ability of making our decision makers smarter, better because it'll find things in this enormous digital noise and surface them for us to evaluate. So that's really the difference between supervised and unsupervised learning.

Maj Rick Hanrahan:
So maybe switching gears and then talking about how this is applicable to our modern legal practice, maybe you could give a little bit of a historical context to just how the legal practice has evolved, and where we are

today in 2019, and where you expect AI and how AI will work with the legal practice of the future?

Col Frank Coppersmith:
Yeah, that's a good question. So I would say the first thing is we have to recognize that outside of government practice, the practice of law is ripe for disruption. The lawyer technology has, as a general rule, served as a breaker of intermediaries, whether it's buying direct from amazon.com and wiping out mid-level retailers, or other places where you can just buy your airplane flight direct and don't have to go through a travel agent anymore.

The role of an intermediary is rapidly being replaced by technology. That's what we're going to see happen in the law because the lawyers serve as intermediaries. The law is opaque. It's complex. It's hard to understand. It's difficult to even know what section law applies to your legal problem. But the problem we've run into when you say, "Well, let's just all go get a lawyer," so the cost of that actually going to get a lawyer, it's incredibly expensive.

Even at $350 an hour, a very low rate for general practitioners, the vast majority of people cannot afford that. In fact, we see anywhere between 50 and 80% of all the people who have a legal problem in the middle class never ever have a lawyer get engaged in it. And then you realize that at the top end of that market, some lawyers are charging as much as 50, 100 dollars an hour. That is a market screaming for disruption.

We have millions of people who have legal problems and they can't afford it, and the price point per hour is more than $300 an hour. That's an exciting. And that is bringing technology into our, I mean, into our world. And that's where a lot of this is coming from. It's really to help satisfy these unmet needs. The reason it works, the reason why it's going to work so quickly in the law is that the law is set up perfectly for AI.

When you think about how the law is structured, the law is strategic subject, predicate object. It's structured almost like a code. I have a slide and a presentation I do that has a flow chart for admitting a piece of evidence, and next to it it has some Python code. When you put the two next to each other you realize they read exactly the same, and that an engineer could look at the law and go, "Well, that makes sense," and a lawyer could look at the code and go, "Well, I get how this works."

It feels the same way. Such technology, I think, is really set up to come in and change them. The last thing I think I would say is that we're already seeing a lot of people do this. Some of the examples, and I think where it really hits is that routine legal practice is going to go away. If you work at a big law firm, you think about the associate work at these big firms it's a lot of document review. It's a lot of document drafting.

We know machines are already better at that than lawyers are. We're already seeing that. There's a technology called **Curia Systems** that's used for by law firms to do due diligence. This system basically goes through untold millions of legal documents bearing say merger and acquisition due diligence process, and is able to surface important findings, important information, places where the documents don't seem to agree.

And importantly, the lawyers can go in and actually set the criteria. Actually, adjust what it's looking for much the same way you might give your associate different instructions after they present you with some documents they found. But here's the thing, it might take your associate another week of research to get back to you on an answer. It takes Curia Systems a couple of seconds. You simply give Curia Systems another natural language question. It gives you back another series of answers, and surfaces some documents.

And you can tell it, "Hey, I really want to see more of this document." Or, "I need to see documents like this. Or, "I want to see documents that are different from this one. This one's good. Are there any NDAs that are different than this one?" And it can surface all that information to you almost instantly. And it's not just theoretical. This

is in work today, and has been adopted by some of the biggest law firms in the country.

Maj Rick Hanrahan:
I think you and I have talked about that and that to me was fascinating to hear some of those actual concrete examples of where it's happening today, right? This is not just theoretical. It's actually happening today. You gave an example of some AI that was in the UK, I believe. And you also talked about some examples with law firms as well.

Col Frank Coppersmith:
Yeah. So there are two I think that are really interesting to me. So we talk about Curia Systems, which is a great research tool. But a lot of lawyers might look at that and say, "Well, that's okay. That doesn't really affect my practice." So let's take one that's more near and dear to our hearts. Let's thinking about the special victim, special victim counsel.

Obviously, the Air Force and DoD suffer from challenges as it relates to sexual assault, and our solution to that was to create a special counsel to deal with that specifically. That's not how some Cambridge students decide to address the exact same problems. Four Cambridge students got together about three years ago. One was 21 years old. And they create a piece of software called **LawBot**.

What LawBot does is it helps users understand and learn about complex legal problems. It addresses a whole host of different types of legal offenses, specifically around sexual assault. They leveraged using a therapist to come in and make sure that the language, that the LawBot was appropriate for potential victims.

What it was, what LawBot really was this software that allowed users to put in basically their facts, their experience, and LawBot would prompt them with questions, and then eventually give them advice, tell them whether they might be a victim of a crime, what crimes they might be a victim of, what should their next steps be.

But here's what's important about that. It was free. It was available on your smartphone. They didn't collect any data from the users. Why does it do all that? Because that's how the rising generations want to interact with technology, and that's how they want to get their problems solved. You see, we've created this massive bureaucracy in the Air Force around the special victim counsel, again, for good reasons and using really good people to solve a very important problem.

But you see those students in Cambridge, they saw that same problem with a little piece of technology but sometimes referred to as a chat bot, and they did it in six weeks. They did it in six weeks. And they created a tool that is actually what people want to, how people want to get help. I think going to talk to a special victim counsel, if you were potentially a victim of a sexual assault would be traumatizing, be very uncomfortable. Downloading an app on your phone and getting that same advice without having to go through that experience without maybe having to retell that story again publicly, that seems like a really powerful tool.

And then there's one other, I think, that is really, really getting to the point of how it puts some pressure on actual legal practitioners. Again, everything we've talked about before could be seen as something that just helps lawyers. If there's technology coming down the pike, it's going to replace lawyers. One of the example is a software called **Do Not Pay**. Do Not Pay was built in the UK. It was a 19-year-old Stanford student who spent the summer in the UK. He didn't understand the law very well, and he got a lot of parking tickets. I mean, thousands of dollars of parking tickets. And so he decided to fight them, fight all of them. And he did. He fought them and he won. He beat be a lot of them.

Maj Rick Hanrahan:
Like any reasonable 19-year-old would?

Col Frank Coppersmith:
Like any reasonable 19-year-old would, exactly right, and who attends Stanford. And so when he got done,

he realized he'd learned all these amazing things. And so he put it into a piece of software. This piece of software prompts users about their parking ticket problem, many, many questions. If you answer all of these questions, it automatically generates a letter to print. It prints the letter that says, "This is how you should respond." It actually creates the document you should send into the parking authorities and tells you where to send it.

Now, here's the thing. There's two really important facts out of this. The first, more than a year ago, it had handled already 250,000 cases. You think, "Okay, does it win?" Well, its win rate is about 67%, which is about the same rate you get if you hire a barrister to represent you. But then here's the part that should make all lawyers nervous. He gave it away for free. He gave it away for free.

There is an entire bar in London that focuses entirely on managing people's parking tickets and speeding tickets, and their driver's license. Because obviously, big city, lots of people, big problem, lots of things to deal with, and he replaced it with an app that he built in the summer. And so that's the power of this technology. It is wildly disruptive to what we do.

And the last piece of it is as we think about, I'm going to talk about this in a minute, as we think about where we're going to be in a few years, we're going to be in a place where our potential legal opponents certainly are other lawyers offering in our space are going to be using this technology. And if we're not using this technology, we're going to find ourselves falling quickly behind.

Maj Rick Hanrahan:
Wow. That's fascinating stuff, sir. You had also mentioned using a corporate context, which I think we will discuss. We're going to be moving into the military context here, but I wanted to start a little broad first and then we'll focus strictly on how this applies to the military. But you had mentioned in the corporate world how AI is already starting to disrupt how settlement negotiations are conducted, and how litigation is conducted.

Col Frank Coppersmith:
Yeah, I think more on the litigation side. So one of the powerful pieces is that there is now AI that will—and there's two good examples of this. The first there is one that will draft NDAs for you. And two parties can say, "We need an NDA." And you can both work with this piece of software. Again, it's also free. And the NDA will work back and forth with the two parties until they're able to draft NDA that meets both of their needs. No lawyers involved.

But then the one that's coming down the pike that everyone is starting to look at hard is the sense of, can I get something predictive? Can I get something predictive about my case? And that is really powerful because if I know all of the lawyers involved, and I know the judges involved. And I'm able to look at untold millions of legal cases that have already existed, and I'm able to pull all the data from my case in a central location and review it, can the AI give me some prediction of what my settlement likelihood should be?

You see, I've worked many years as a corporate counsel and big corporations with some rare exceptions do not want to litigate. What they just want to do is get to an answer, discover how much they have to pay or how much they are going to get paid, and then move on because most companies typically are not in the business of litigation.

There are a few high-profile ones like Apple and Microsoft that use litigation differently. The vast majority of your mid cap companies litigation's just an enormous distraction. They just want to know how quickly can you close it, and how little can you spend? What that means is they pay, right now, lawyers a lot of money to make those estimates. Shockingly, large amounts of money.

I guarantee that the CFO, given the opportunity to just drive and drop all the case files into an Oracle somewhere for an answer to that question, would love that answer. Because you see if that Oracle also has, say, upper and lower bounds of what two partners are

willing to settle for, then that's just what mediation does. You think about mediation or even, at some extent, arbitration, it's just about seeing if the upper and lower balance of someone's settlement authority and comfort level are overlap.

And if they do, you get a result. Well, if I can do that in 15 minutes with two chief financial officers on a piece of software on a computer somewhere, and not spend tens of thousands if not hundred thousand dollars in lawyers doing preparation, is that interesting to me? Do I like that? Am I excited about that if I'm a corporate executive? You are absolutely certain. I'm thrilled to do that. And so, that is the next step for us. And that is going to have enormous, enormous impact because it won't just apply in civil matters. I think we'll see it move over into criminal matters as well.

Maj Rick Hanrahan:
Thank you for listening. We hope you've learned a thing or two about artificial intelligence. Please check out episode two, which is a continuation of this interview with Colonel Coppersmith, where we dive into AI's role in the military in JAG Corps at large. Last, feel free to check out Colonel Coppersmith's article entitled, **Autonomous Weapons Need Autonomous Lawyers** as published under Digital Reporter on 10 April, 2019. Thanks again, and see you on the next episode.

Announcer:
Thank you for listening to another episode of the Air Force Judge Advocate General's Reporter Podcast. You can find this episode's transcription and show notes along with others at reporter.dodlive.mil (now **JAGreporter.af.mil**). We welcome your feedback. Please subscribe to our show on **iTunes** or **Stitcher** and leave a review. This helps us grow, innovate, and develop an even better JAG Corps. Until next time.

Disclaimer:
Nothing from this show or any others should be construed as legal advice. Please consult an attorney for any legal issue. Nothing from this show is endorsed by the Federal Government, Air Force, or any of its components. All content and opinions are those of our guests and host. Thank you.