"Introduction to AI in Business - The 5th Industrial Revolution" is a 3-part series presented by Certified Information Security's founding Principal and content author, Allen Keele:

Session 1 (This session) - Embark on a journey through the basics of AI in simple terms. Dive into the 2025 AI trends and its integration into business operations, all supported by well-recognized and reputable sources.

Explore the follow-on sessions in this 3-part series:

Session 2: Delve into the reasons behind organizations incorporating AI and its practical applications across various industries like healthcare, retail, services, manufacturing, quality management, and supply-chain management.

Session 3: Understand the critical aspects of AI accountability, safety, and risk management in business settings. Discover key frameworks such as ISO 42001 AI Management Systems, ISO 23894 AI Risk Management, and NIST AI Risk Management Framework 1.0. Uncover the compliance mandates of the EU AI Act of 2024.

Transcript

Welcome to this short course on artificial intelligence in business, the fifth Industrial Revolution. I'm Allen Keele. Thanks for joining me. We'll begin the course by getting a basic understanding of artificial intelligence in the business context. We'll then expand upon that by going into AI adoption and trends globally in 2024. We'll then look at how to benefit from and leverage AI in an organization. We'll then get more specific examples of how organizations are benefiting from integrating AI in various industry sectors. We'll then move on to managing AI risk, AI, accountability, safety and risk. We'll then learn more about how we can leverage existing newly emerging standardized frameworks to plan, integrate, manage and improve AI in business. And then finally, we'll go into late-breaking newly emerging. AI legislation and regulation that governs how AI can be used in business.

So what is artificial intelligence and how can we use it in business? Let's take a look at some of the basics of AI in business. Microsoft Co founder Bill Gates once said the power of artificial intelligence is so incredible that it will change society in some very deep. So let's look at a very basic definition of what AI is. Artificial intelligence is a machine's ability to mimic human action and to imitate logical thinking, learning, planning and creativity. AI systems can perceive and analyze their environment to solve problems independently. So a unique feature of AI systems is that they learn from past situations. They use this information. To adapt their future actions. So how does AI? How does it benefit us in business and otherwise?

Well, the desired objective for artificial intelligence. Is to enable computers to perform complex tasks beyond simple automation by using human-like faculties such as data

processing, pattern recognition, and even decision making. So how is this then delivered? Well, we need to go beyond simple automation to have technology actually learn on its own. So machine learning is a key aspect. Of AI that differentiates it from simple automation. AI uses algorithms and learning processes to adapt to new data and improve performance overtime. In other words, it incorporates lessons learned. So what does all of this need to work? Well, we need some serious computing power. Computing power is crucial for the functioning of AI systems. It involves memory capacity and rapid data processing to enable. An AI system. To perform complex tasks efficiently. So as we look at this data science, artificial intelligence, aren't they the same thing? How do they differ? Let's take a look at that. So when we look at data science, for example, in the background, we look at the data science involving. Collecting massive data sets for analytics and visualization. Whereas artificial intelligence relies on predictive models to do more than just analyze and collect data, it actually uses these predictive models. To forecast future events. For future activities and decisions. The data science provide statistics, whereas artificial intelligence applies algorithms to statistics to derive new information. Again, for new decisions when it comes to techniques, data science. Has data analytics techniques, whereas artificial intelligence is taking it a step beyond to actually attempt to not just analyze information that it's already collected that is already there. It's actually trying to develop new information. It's incorporating what we'll learn about later, deep learning and machine learning. Learning the focus of data science is to focus again on patterns within data sets. Whereas artificial intelligence focuses on going beyond that to actually mimic and imitate human intelligence. And when it comes to problem solving, data science is using specific if then programs as opposed to AI. Using these programs to then actually solve problems. The goal of data science is using data to make good business decisions. But their decisions made by. Good people using data science to make their decisions as opposed to artificial intelligence is where we are teaching machines and computers to actually act intelligently, to think on their own. And mimic human thinking as it considers data to make good business decisions on its own. With tools, we look at data science with various existing technical tools. Sass, SPSS, Keras, R, Python And so on with artificial intelligence. Also have tools available. The purpose of data science being to help companies make informed decisions based on information that can be made available to them. Whereas artificial intelligence goes beyond to try to improve lives and improve business by developing systems to better. Support work and everyday life. Now you're going to find in trying to understand what AI is. Well, that depends. We have different maturity levels of AI, but we also have different applications of AI that could differentiate one AI from another. So let's take a look at this. We have traditional neural networks. Machine. This is an umbrella term that covers a lot of ground for algorithms that learn from data to recognize patterns. And make predictions. But then we have a complex neural network that we term Dee learning. Is a secialized subfield of machine learning macro to micro. It uses multilayer artificial neural networks. To learn complex patterns in large data sets. To then be able to mimic human like thought processes. So let's see this graphically. So we look at artificial intelligence and understand that within artificial intelligence that allows us to mimic human intelligence such as a self driving car that's making decisions. On how to react to changing circumstance as the car goes down the road, well within that it turns out that it needs to be able to solve tasks as it

continues to drive that car. Now that means that it's not really prepared to do so until it's been given enough information and taught. How to consider that information to make decisions on driving a car so it needs machine learning? Now in another example of machine learning would be for example facial recognition that you may have to unlock your smartphone. It's actually analyzing the data that's given, namely the dynamics of your face that it's looking at. In the. In the graphic representation of the image, and then it's actually analyzing that to determine. Is this close enough to what I understand to really feel that I can authenticate and authorize this person to use the phone? Well, deep learning is where it goes beyond just analyzing what it's already been given. According to the algorithms that it's supposed to use. This is where with deep learning, it's the machine's ability to actually train. Itself overtime to become smarter, if you will, through neural networks and big data. An example of this is chatgbt. So we've learned that a key difference between AI and simple automation is the ability to learn and become more capable over time. O let's go ahead and take a look at that. We have reinforcement. This is a an approach for AI to learn. Where the model is trained through interactions with the environment that it's operating within. It understands when it gets things right, but it also finds out when it gets things wrong. And punishments. So for example, you can have AI learning. In a game to learn a game through the same method that a human learns learning when it fails by going through the wrong door or when it fails by a particular action or particular decision that was made, that reinforces. What it needs to do to try to attempt to learn the right decision and when it makes the right decision and see success, that success is a reward that reinforces making decisions along that line. So again, just like a human, it's learning from its own rewards and punishments of learning what was right versus what was wrong. To then become better and more accurate at it as it goes. It is learning as it goes. So we then have something called unsupervised. Unsupervised learning is where the model is trained using unlabeled data. We'll talk about labeled data in a moment. And it is allowed to then use this data to then discover patterns and insights itself, without the human interaction. We then have supervised. This is model is being trained using labeled data. This is where for an example of this would be recognizing spam emails based on previously identified labeled data with spammed emails. OK. Again, you'll see that there is an apparent hyperlink here. If you're actually using this course from a in our learning management system, it's certified information security. You'll be able to click on that link and use that to then go learn more about labeled data if you wish. Not trying to stretch this course into a full data science course. It's beyond the scope. But for, for those of you that have thirsty minds and you want to go learn more about this, why there's a nice convenient link for you as we move on. Semi supervised. This is a model that is trained using both labeled as well as unlabeled data. An example of this would be recognizing handwriting using a small number of labeled samples. So we didn't figure out AI all at. There's been an evolution, a learning evolution, that we've gone through to try to determine how to go beyond machines just operating by rules to the point where they can actually learn. So we begin with reactive machines. This is classic weak AI that can only perform a task for which it was specifically programmed. It's been taught one way, one way. It doesn't learn from what it's doing, it's simply performs essentially, if then automation, we then go on to limited memory. Is a common type of weak AI that is still used today. Collects, analyzes data and applies it to current events. We have theory of mind.

This is theoretical and powerful. AI is able to perceive, understand, and respond to human emotions. We are still working our way there to really pull the. Off we then have self-awareness and again this is where in theory we have AI that can reach or even exceed human level consciousness. Sentient AI again, for better or worse, and it could be good. Might be terrible. This is kind of where things are evolving too. So what can AI systems do? Well, they can do a lot of different things. We have text to speech, speech to text that is facilitated by a level of AI. We have expert systems that will attempt to make decisions based upon. Well, the expertise that's been fed to them. That they didn't necessarily learn on their own. We have robotics where we can have AI that helps to make automation. Performing repetitive and predictable tasks, we can actually use AI again for image recognition and machine vision. With assisting with vision, we have again machine learning with both deep learning and predictive analytics that I talked about. We have, which is what most people now recognize in the public as AI generative AI, where we actually use AI to generate new information such as sounds, music or even. Graphics videos. We then have AI that is used for planning, optimization, logistics and then we have natural language processing for classification, translation and data extraction. Going back into making data science smarter. So let's take a snapshot look at three. Oh, oh, oh, oh. Foot view of some of the benefits and the risk of AI. Within AI, we can find better automation of routine tasks. Not only do we have simple automation, but perhaps it will evolve. And improve from lessons learned over time, we'll have improved decision making, personalized services, progress and medical research. We have increased efficiency in production and logistics. We have more creativity and innovation through recognizing patterns and problems. However, we also have to be concerned with data protection and privacy, because as AI learns, it tends not to be. So circumspect as to where it gets its information, it gets it where it can find it, and that might put us all at risk of inadvertently. Breaching data privacy or losing data. Privacy. We also have ethics and accountability concerns, of course. We'll be talking about some of this later. We have potential fairness and discrimination problems. In other words, bias that can actually be inherent. In artificial intelligence, and if you think about it, if AI is learning from us and if we have bias. What are the chances that AI learns that bias? Yes it can. We also have to be concerned about security and opportunity for misuse, fairness and discrimination. We have to be concerned about. Transparency and explainability of AI systems. This is a key concept, especially with AI. You'll learn more about this in ISO 42, double 01 for the AI management systems. Transparency and explainability goes to. The AI systems themselves. And how trustworthy they are, how much we can have confidence in the AI and that in large part is determined by how much do we know about it and where did it get its decisions from. Let's face it, if I came along and said hey. You don't know me, but trust the advice I'm about to give you. You might want to ask where I got my ideas right. Well, what if I wouldn't be transparent about telling you where I based my opinions from? What if I what if I would share what I used to base my opinions on, but I couldn't explain it to you in a way that you. Understand. Either way, you wouldn't be able to trust my advice. OK, well, I just put it in layman's terms. We. The same concerns with AI. Where did? Get this information to make its decisions. And how did it go about making decisions from that information anyway? Transparency and. Key things to support trustworthiness of AI. And then finally, we also have misinformation that can also taint

AI as well. OK, that little presentation explanation that I gave you about AI transparency and AI explainability and and why we need those within an AI system. Also shows you. That there may be good reason to be skeptical of AI. And the dilemma is one of our objectives of a good AI management system is to achieve confidence and achieve trustworthy AI, OK. What are some of the reasons that people are skeptical of artificial intelligence? The first is they're concerned that AI destroys jobs. OS. Well, with any industrial revolution we have activities and jobs that are no longer necessary. As we've evolved with. Trains and planes and automobiles, but also jobs that will lose that. AI can replace, but on the other hand we'll have even greater need for people who know how to leverage AI for even greater effect. So again, will some jobs be? Yes. Will there be new jobs, of course. But still there is a reason for skepticism with many people. Because of this, OK. Also recognize that just as a human can make mistakes, so can AI. All we make mistakes based upon poor consideration of information. So can AI. We make mistakes, perhaps because we didn't think through something very thoroughly. So can AI. So we not only consider the wrong information or we don't really consider the information very deeply. AI can make mistakes, just like we do. We also have concerns with AI potentially disrupting personal privacy. Not only is it. Information. It's constantly exploring for new information and using that, keeping it, learning from it, but also you know, it turns out that you can access information. By deducing it from other information. If the money was in the room before you came, but then it was gone. After you left, I could deduce maybe that you took it. If no one else had access to the room, you don't know that you didn't see me take it. I know, but it was there before you were there and it's gone after. So in other words. There could be personal private information that is deduced. By a system that can think. And AI is designed to think 0. Not only could it potentially access. Personal private information. It could also potentially deduce private information without access. AI could be used for criminal. Perhaps we can use AI to make cyber attacks better and smarter and more effective. Better. Well, worse, I. Deending, uon uon. How you're looking at it. Right O it can make cybercriminals more effective. Can make them faster, it can make. Another words, if it helps good people do good things better. It can help bad people do bad things better, more effectively. So AI could be used for criminal. We're already seeing that now with deep fakes and masquerading and things of this nature. We'll talk more about that later. We also recognize that we're skeptical of AI because we're worried that AI is going to take control. How? Well, we fear power, I. And we're concerned that we can't trust AI to well be the right person in charge if it's in charge. So we're skeptical of that. And then finally, we're also as humans. Concerned that we won't be the head humans anymore, that we have AI that will dehumanize the world. And I think that subjectively my own opinion is yes, I've experienced where. Things have lost some personalization, some human touch, if you will. But then again, maybe we just start to the point yet where AI is good enough to be human touch. Dangerous conversation. So we have great hopes for AI, but what are some of its limitations? Well, we have lack of initiative. We have dependency on data and what data the AI can get to. We have a lack of critical thinking skills that we're trying to learn to emulate with AI. A lack of emotional understanding. That gives humans a perspective that is very difficult to really imitate in AI. AI doesn't then understand bias, and it can't necessarily catch itself at applying. And that's a problem where as a person, can we have a potential lack of creativity, a

lack of common sense when you think about it, what a humorous term for our discussion here? Common sense, what is common sense? Again, that's like trying to define integrity. Doing the right thing even when you won't get caught for doing the wrong thing. Well, that requires a sense of what's right and what's wrong. Does AI have that sense? Could it then practice that with common sense? Food for thought, right? So we have also moral and ethical limits that may not be boundaries, that AI understands, AI doesn't. It doesn't have feelings at this point. Again, it goes back to that lack of emotional understanding and finally need for regular maintenance. I don't know if that's necessarily a problem because people have a need for regular maintenance as well. So we've looked at evolution through the industrial revolutions. Now let's go ahead and take a look at some of the AI adoption and trends in 2025. Innovation why do we need it? Henry Ford once said if I had asked people what they wanted, they would have said faster horses. So in other words. Staying in traditional thought and not innovating because that's what everybody seems to want. Is not necessarily how you make progress. So earlier you heard me mention industrial revolution when I was talking about, well, some jobs leaving and some jobs coming, OK. When we look back through time, not so far back, we can look to our first industrial revolution. Where we were looking at the end of the 18th century, late 1700s using steam power to enable mechanical production facilities. And thereby replacing some of our previous manual agrarian tasks. So we then moved on to the second industrial revolution where we had electrical energy enabling assembly line production that then progressed into our third industrial revolution where it and computer technology. Then radically changed how we live and how we work. This enabled further production automation beginning in the 1970s. This progressed into our fourth industrial revolution. Do note that they're coming faster and faster where the Internet of Things enables network production. And then finally, we are now in our fifth industrial revolution. Where further development of technologies from the fourth industrial revolution are driven by AI and learning with virtual reality as well as mixed reality. So what is the future of AI? Well, the next 10 years will be pivotal. Matter of fact, it's expected that we will make more progress in AI in the next 10 years than we have in the past 50. However. More breakthroughs are needed before machine consciousness can actually be considered feasible. Feasible. Interestingly enough, the recent winner of the Nobel Prize for Physics for his work in artificial neural networks, Doctor Jeffrey Hinton, said. I think we should think of. AI as the intellectual equivalent of a backhoe. It will be much better than us at doing a lot of things. So that is the insight from Doctor Jeffrey Hinton. So as things are evolving today, we look at AI expertise becoming the foundation of corporate development. We know that AI will increase efficiency that will potentially facilitate even a four day work week. We understand that processing allows human senses as input and this will continue to progress. We have AI generated music and films and this is becoming increasingly prominent we have. A GPU shortage that is driving smaller. To then end, Nvidia's dominance in the AI we have data quality becoming more important than data quantity when it comes to AI training. So let's look at some of the data quality trends for AI in 2025. We now see that AI has gone mainstream. AI tools that can be used directly by end users are becoming increasingly popular. Microsoft Co. Examle. But now we see AI being integrated into online web applications into almost all software has now become AI enabled. AI facilitated OK. Also, recognize that open source

models are now starting to compete with our closed source. Models and AI transparancy, along with Explainability, is becoming more relevant for companies, especially when it's required by law. Europe's AI act is now emerging as the benchmark of safety AI. Be. More about Europe's AI act later. We also have the United States Presidential Executive Order 14110 that recently was released to require AI risk management and implementation of the NIST AI Risk management framework throughout all federal executive departments and agencies. We'll be talking more about that later as well. We also have ethical questions about AI systems that are now gaining momentum. Let's look at how this is impacting the way that we do business. Is it impacting our organizations and companies? Well, the very first and most obvious one is the competitive environment. Early adopters of AI. Organizations that integrate AI to their benefit to increase their service offerings and to increase the quality of service and to cut costs increase productivity. They are naturally going to become more. Competitive for it. And they will actually be able to get ahead of competition that lags in adopting and integrating AI. So AI technology will not just lead to strained relations between companies, but also between governments in different countries. And towards that end, we're going to see a greater need and a drive towards international collaboration as countries strive to become AI leaders. The companies within those countries. We'll need to align themselves with national strategies. For AI, so as countries enact legislation, you'll have companies that operate within those countries that now fall under different laws and different requirements perhaps than in other countries. We need to straighten this out. Needs to again. Improve through collaboration through international cooperation. We also see with talent acquisition that AI expertise is becoming one of the most sought after skills. So again, if you're concerned about losing your job, become AI savvy and you will become then in demand, companies will need to develop strategies for acquiring. And retaining AI savvy talent. Also, we are concerned with supply chain problems, bottlenecks in compute resource in the semiconductor industry will force companies to diversify so they can better mitigate their risk within the supply chain. So I talked about how AI is affecting business and the trends of business as we can expect in the future with AI. But what about the business of AI itself? Wow, take a look at this $196 billion worth of market value in 2023 alone. Forecasted revenue for the AI market in 2031.812 trillion. We see that 35% of employees surveyed in Europe. Regularly use AI in their daily work and where is all of this coming from? This is all coming from Mackenzie, the state of AI. In 2023, Generative Ai's breakout year. We also see Grandview research, artificial intelligence, market size, share and trends analysis report. Released in 2023. So Speaking of countries, adoption trends in AI, which also will contribute to their success moving forward or perhaps their failure, their challenges. Who are our market leaders in AI adoption and AI experimentation? We begin with #1 Singapore, followed by United Arab Emirates, followed by South Korea. Well, if those are the top three that are leading. Who are the bottom three that are lagging? Well, the prize goes to USA and France. Tie for lagging the most. And then next U from there is Italy. No, Allen, it can't be true. United States can't be lagging. Okay don't. Don't kill the messenger. This was actually reported by IBM with IBM Global AI adoption Index Enterprise report. In 2023. But Allen, our organization is small, but we don't have the kind of investment necessary. Don't have the. We don't have the resources to develop and integrate AI into our organization. Won't be able to compete. Ah, never fear. There is an emerging

boutique cottage service industry. AI is a service. So with traditional AI development, we were concerned with needing to, well set up and pay for and equip large teams consisting of data engineers, infrastructure architects and developers. But today, with AI as a service, we can have smaller teams only requiring a data scientist to interface as well as perhaps an app developer to develop an AI solution. Traditionally, we had to worry about upfront costs, very large investment both in people and in equipment and in data access, very large investment before a is value to the business could ever become recognized. But with AI as a service, the return is much quicker. Includes the work of data scientists, data engineers. Infrastructure architects and developers, through a subscription access pay as you go. Well, with traditional AI developing it and realizing return from it could take weeks or even years. This required significant time for all of that investment for set up experimentation, customization, training, and ultimately. For the software development itself. If you ever join me for the AI risk management framework training that I do, you're going to see the actual NIST AI development life cycle. And how the AI risk management framework that NIST created maps to that? That is a beast. Well, if that's too much for you, turns out that that can be part of your outsourced effort. Weeks and years can be reduced, potentially to hours and weeks, where your AI as a service helps you automatically find the appropriate model. Training it at a faster speed with less effort. We also were concerned previously with traditional AI with poor scalability. This could take several years to develop. What it is that you could possibly use only to find out that it's obsolete by the time you've created it. So now we recognize that instead of taking several years before AAL models can be converted into actual software for production, we can speed that up with scalability from a service where the model could be ready for implementation. And be more easily scaled moving forward. Earlier we talked about our skepticism of AI because while it would, it would destroy jobs. Well, like I said. Now it'll destroy some jobs as we move forward in this industrial revolution, but it will create new ones with this revolution. So we have some jobs that will be replaced by AI, but we'll now have new jobs created by AI for those people who are AI. Heavy and are needed to be able to operate the AI that replace the other jobs. So we'll have new positions available and new need for process automation specialists, machine learning specialist AI specialists, big data specialists and again. Where does this information come from? I didn't make it up. It turns out that this was published in the World Economic Forum. Future of Jobs Report 2023.