

Hi, and welcome to the introduction.

Today, we'll be talking about AI, in particular the business context of AI, and especially AI that you are already interacting with that you may not even be aware of.

Everyday AI

Depending on your level of exposure to artificial intelligence, you may feel that this is a foreign concept to you that you're gonna be starting from scratch. But the reality is that's not the case. You have been using it AI-powered resources for frankly, a number of years. And we're going to talk about a few of those examples so that you can understand, "Oh, that's AI".

Alexa

If you have an Alexa device, you know that it's a machine that takes voice, responds in voice, processes it, and can take actions on your behalf, giving you information. If you have a smart home, interact with the smart home, that sort of thing. Alexa is AI-powered. As a matter of fact, a lot of Amazon's growth has been AI-powered.

Siri

If you have an Apple iPhone or even an Apple computer, you can use Siri to do the same thing. Roughly the same thing that Alexa does, maybe not as complex with the smart home stuff. But you have the opportunity to interact with the machine using voice, can process that voice message, and provide you with an output, hopefully, that's accurate.

Amazon

Another example that we discussed in the written lesson would be Amazon again with their recommendation engine. You go to buy, in this case, a digital camera. When you get ready to check out, before you get there, Amazon's going to say, "Individuals who bought this camera, also got these things." That small tweak with AI power has resulted in Amazon and most categories of seeing a 30% or higher increase in sales prior to when they had their AI-powered recommendation engine deliver specific additional purchases for the individual. These are all tools that we've all been using for quite a while now. It was AI.

Three terms that you're going to hear as far as AI goes that relate to the capabilities of artificial intelligence are **Artificial Natural Intelligence**, **Artificial General Intelligence**, and **Artificial Super Intelligence**.

We dove into them in the written content, particularly focusing on **Artificial Narrow Intelligence**, which is artificial intelligence that is built to perform a narrow set of tasks, such as those that we just described.

“Siri, what's the weather today?” It can give you that. *“Siri, can you create a new symphony for me?”* It cannot do that. It serves a very narrow function, but is AI-powered.

Artificial General Intelligence is the level at which, collectively, we believe that when AI hits AGI level, it will be comparable to human capabilities. At that point, it will be at parity with human intellect.

Then, **Artificial Super Intelligence** is the one that the media loves to talk about. That's, we suspect, years away. Artificial Superintelligence is the point at which the artificial intelligence itself has capabilities beyond what humans thought was possible or are capable of as a human brain.

A funny example I saw that kind of brought it home, Artificial Narrow Intelligence. If you'll see the little narrow effect that he's got in his hand. Artificial general intelligence – it can think like you, it can determine the next best course of action rather than deliver a result prompted by input, which you'll learn all about prompting in just a moment.

Artificial General Intelligence is the level at which AI can make decisions based on future projections and determine the next course of action. And as it references here, take knowledge from one domain and transfer it into another domain.

And then **Artificial Superintelligence**, as you can see here, an order of magnitude smarter than humans.

All of this course is going to be focused on applications of Artificial Narrow Intelligence in business.

So we just introduced you to those other two topics so that you were at least familiar with the distinction between where we are now and where perhaps the media may suggest that it becomes dangerous.

Machine Learning

In artificial intelligence, there's a few key concepts that we're gonna talk about. Machine learning is simply programming an algorithm that is able to take in data and make and create output without human participation in the creation of that output. The machine has been fed data. It's got a data set. Now, based on the prompt or the input that you give it, it now provides output without you participating in the process of creating the output.

Deep Learning

Deep learning is a form of machine learning that is modeled on literally the human brain. If you were ever to have a microscope and be able to see the activity that occurs in a brain while an individual is learning something, watching something, participating in an activity, there are neural nets that are engaging in the brain and actually being created real time as new experiences are occurring by the human. Deep learning is machine learning structured on the neural net model, very similar to the human brain, allowing it to be much more robust in its development of processes to deliver the output based on your input.

Large Language Models (LLM)

LLMs are what drives ChatGPT and BARD. They are huge data sets of information that allow the interaction with the human to provide input. The Large Language Model, because of its broad expanse of data that it's drawing from, with the machine learning and the deep learning to be able to give you output that is creative. Output that is an unexpected. Output that is hallucinatory or absurd.

So, those are the three terms that you'll hear a lot: **Machine Learning, Deep Learning** and **Large Language Model**.

Generative AI will talk about in just a moment, which is the focus of the application of Generative AI is the focus of the rest of this program.



One of the things to remember about machine learning and deep learning and LLMs is that the larger the data set, the higher the accuracy.

In an LLMs case, if we gave it a data set of 10 books on carpentry and our proprietary LLM, that was carpentry focused, would be able to draw from the content of those 10 books. And give us an answer based on those 10 books.

But if that same LLM was trained on educational studies of carpentry, and carpentry is art, and the profession of carpentry, and A to Z carpentry, and carpentry 1500 years ago versus today. If the data set was that robust and I prompted it for output related to a carpentry topic, it would have a much broader understanding of the subject matter. We call it an understanding, but it would have a larger data set in order to be able to provide a more factually accurate, detailed, actionable output than that smaller data set. So this is a simple concept, but the more data in, the more precise the result is based on your input.

Machine Learning Examples

So again, these are algorithms that can learn from data to make predictions without you telling it what to do.

Netflix Recommendation Engine

Netflix has said that the viewership, since they have introduced the recommendation engine, has increased. It keeps users on much longer. I can't remember the numbers but it was astounding. Without that, Netflix would be a great company used by many. But because of that particular application of machine learning, it now knows that whatever associations are related to this movie that you just watched, all of these other movies are tagged similar.

So, it's giving you accurate recommendations based on the feedback that you've given it. The larger data set that it's created from its exposure to you saying, I like that, or it notices when you don't finish something, or when you watch something multiple times. It uses that data to more accurately give you a recommendation that has a high accuracy of being something that you like.



Google Maps

"4:59, time to go home. Let's check Google Maps to see which way is the fastest way home."

This is a machine learning environment as well. Google, which a lot of their growth has been powered by their AI initiatives over the past 20 years, the machine learning algo for Google maps takes into account: *"Is it a holiday? What is the traffic typically been like on this particular route at this particular time? What's the weather look like? Is it daytime? Nighttime?"*

All of those types of inputs from Google allow it to predict the most expedited way to, from point a to point B. Again, machine learning driven.

Grammarly

Finally, one that we've used a lot in the production of this course is Grammarly. Grammarly is an editorial assistant, I guess you call it. It allows you to take a copy or a body of text and Grammarly, using its machine learning algorithm, will represent that text to you in five or six different ways of saying something.

So in our case, if we were preparing something, we wanted to get a point across, but we didn't feel like the way that we worded it was the right way, highlight, hit Grammarly, and it would give us the five or six best alternative ways to get that concept across using similar language. All of those recommendations were generated by a machine learning algorithm.

Deep Learning

Deep learning again, as we talked about, this is based on neural network. An example of use of that, we've all seen the commercials of the Gecko and Flo and Liberty telling you that there's savings waiting for you should you go to their website and go through their process.

That process is a deep learning powered process that takes into account information such as: *Where you park your car? Or where your home is? What's the likelihood of fires in that? What's the flood zone in that area? What's the theft rate in that area? What's the violent crime in that area? What's your social security number?.*



We can pull up all of the records that we have access to. *Does this person pay their bills on time? Is this person under 25 years old?* All of these variables allow for the deep learning algorithm to predict what would be an appropriate premium for you to pay for that insurance.

Insurance companies, your premium is based on the likelihood of risk. That they will have to pay up insurance payout because of actions that you took or related to where your property is located. That deep learning algorithm can pretty accurately mitigate risk for the insurance company by adjusting the premium for the individual.

The smarter the algo, the more precise that premium. That's how they're able to say that "*We can save you \$750 a year over any other insurance company.*" It's because their algorithm they believe is precise enough.

That it can accurately predict. And that precision allows for usually a lower price because the insurance companies want to air in their favor. If they're not sure how it's going to work out, at some point we can charge them enough to where it'll work out for the insurance company. These algos make sure that the risk is managed by the insurance company while also giving us better rates.

UPS

If you were to see the process of UPS, I mean, their business is physical delivery in an expedited fashion from a specific point to a specific point. The whole process of getting that package from your grandma to your house is powered by deep learning.

Once UPS receives that product, it's scanned. And now, the system takes over. It knows that this truck goes near that place. And if we put this package on here as the third package in and we structure this package in this package, by the time that truck is loaded, the driver can look at their phone. There's not a lot of discretion involved. They follow the route that has been mapped by AI and is real time updated from the deep learning algorithm based on driving conditions and input on weather and things like that to help ensure that the timeline of expectation is met for the delivery of that package. It's not as simple as somebody looking at a map and digging through packages in the back.



Another example, **John Deere**. We think of them as really agricultural equipment company. John Deere has very sophisticated deep learning algorithms that take in a lot of information, not related necessarily to the sales of their products, but to build relationships with their clients.

For instance, John Deere has very accurate harvest predictions for their clients. This type of information doesn't necessarily on the surface support John Deere's growth. However, if that service is provided when you buy John Deere's equipment, they know that it's more likely that they will maintain a long relationship of re-purchases from those clients because those clients are now reliant on the information that's being given to them by John Deere's deep learning algorithms in order for them to have higher crop yield, harvest before the big storm, whatever the case might be.

Large Language Models

Some examples of Large Language Models, we've mentioned ChatGPT and Bing and Bard and Claude but some that you may not be familiar with.

Stability AI is a large language model that's image driven. Stability is the LLM behind tools such as Midjourney. And if you've seen any of the amazing AI generated art that's popping up on social, most of it is coming from Midjourney or tools that are run on a Stability AI LLM.

Another example that you might be familiar with is **Spotify**. When you set a station on Spotify, let's say it's Billie Holiday. That large language model, based on tags and associations with Billie Holiday, understands that all of these other songs have very similar profiles with tagging and structure that Billie Holiday was.

It's not going to play Green Day next up in your queue. These large language models allow, again, for that accuracy to make sure that if you click on Billie Holiday, you're probably going to get some jazz that is not fast-paced. It's going to be relaxing jazz as compared to Green Day or Metallica popping up on the next song after Billie Holiday. The LLM that Spotify has built is the tool that allows it to have that level of accuracy.



Another example is **HubSpot**. Commonly used business tool for email, customer relationship management, that sort of things. HubSpot uses its large language model that has been fed with billions of emails. And it understands the structure of different types of emails and that if this word is in there, it probably is this structure and that sort of thing.

Similar to how Spotify and Stability can generate output based on the data set that has been fed into that large language mode, HubSpot does the same. It can auto complete sentences or even entire emails based on the initial subject matter that you type into it, powered by an LLM.

Everyday AI. These are all services or tools that we have encountered and may not have understood the contribution that AI had into their operations.

As I mentioned at the beginning of this introduction video, the bulk of the tools that you'll be learning how to apply are under the category of Generative AI.

Generative AI is a relatively new category, and it simply means that it's AI that generates output based on your query or your prompt. Again, we'll talk about prompts in later modules.

Generative AI is how people are creating those images. And how people are having AI create a SEO-optimized content piece. Or how marketing teams are understanding which advertisement is going to work the best for their audience. And having that info, having that input generates a very accurate advertisement for that particular product.

It's blown up. The landscape has blown up since the launch of ChatGPT in November of 2022. As it says here, it's a work in progress. Depending on the narrow application that you have for a tool, there's going to be a handful of them. And some of them are going to be pretty good. Some of them are going to still be a little buggy. However, the number of options that you have when it comes to specific tools for artificial narrow intelligence and a generative AI application are growing daily.

One of the things that we do with our community is we make sure that we are consistently and constantly testing the new tools that are coming out. Comparing them to the tools that are filling that niche for us now, whether it's a workflow management tool, or it's a text generation tool, or it's a scheduling predictor or whatever the case might be.

So just to give you an example of how much is coming out, this was from a few months ago, but this gentleman tried 200 AI tools. And this was four months ago, according to when that video was shot. Today there may be hundreds coming out per month that are specific to voice to text. Taking your calls and converting them into transcripts and summarizing those transcripts into the relevant points and identifying who in the conversation was tasked with doing certain things. All the way to the creative side taking a handful of words and generating a symphony or a rap or a children's rhyme.

Generative AI is almost limitless in its application. The limitations are you. How do you think? How do you communicate with the machine? Do you understand what's possible? Once those things happen and once you developed a level of AI literacy, AI intelligence, and understand your business well enough to know that, "there's an AI for that," the benefit that you'll receive from AI and in particular Generative AI in your business is exciting and incalculable almost.

It's limited only by how well you understand your business and how creatively you can elicit the information to be generated from the LLM.

In the next modules, you're going to be learning all about the tactical application of the tools. You're gonna be learning what tools exist. How do I use them? And how do I use them in business?

And in later modules once you've learned the tools, we're going to dive into: how do I create an AI business strategy? Now that I know the tools, what process do I go through to fully deploy the benefits possible from the use of AI in my business?



We're not focused on the technical stuff. If you're a data scientist, if you're a coder, fantastic. We're happy that you're here. However, the bulk of what you'll be learning in this is specific to application in business workflows, business processes. Whether that's generating more content for your social media engagement to again, the advertising example. Having more precise advertising that gets better results from every dollar that we spend in marketing.

So with that, I'm going to close up the introduction. And in the future modules, you're going to be introduced to the rest of the executive team here. The next one you'll be working with Dan Linden, who is the Chief AI Officer of Chief AI Officer.

See you in future modules.