DEEP DIVE INTO:

AI

ISSUE:

2 Limits of AI

A monthly in-depth report on AI, written by Byron Reese.

SUBSCRIBE NOW

GIGAOM

Introduction

GIGAOM

As one decade closes and a new one begins, the media marks the occasion by finding futurists of all stripes to make predictions about the next ten years. When it comes to AI, there is a wide range of opinions about what the near future holds.

Most of those in the AI field are, for obvious reasons, excited by the prospects that the technology holds. I say "for obvious reasons" because presumably if they weren't excited about it, they would be in another field that they were excited about. But even among those enthusiastic about the technology, there is a wide range of beliefs about what the technology will actually be able to deliver in the foreseeable future. That's why I wanted to write this, the second issue of Deep Dive into AI, on the topic of the limits of artificial intelligence.

I am incredibly bullish about the technology. When I hear it described as "the new electricity", by folks like Kevin Kelly and Andrew Ng, I find myself nodding in agreement. When others talk about how it will animate everything, and everything around us will become smart, I nod even more.

However, I am keenly aware that the technology has very real limits; that you can't just say "AI will do that" about anything. The techniques that we use for AI are highly specialized, and thus can only do a narrow range of things. Those things, luckily, have broad applications, ***so AI really will impact our lives a hundred times a day, a thousand even, the way that computers already do. Can you imagine a world without computers at this point? They run everything around you. I would bet dollars to donuts that wherever you happen to be right now, there are ten computers within thirty feet of you. AI will be like that, but even more so.***

Often, on my podcast, [Voices in AI](), the topic turns to whether AI is overhyped. Going back through those episodes, I was struck by the perspectives of my guests. After all, these are people in the AI world. Let me share just a few of the comments my guests have offered. If this isn't particularly interesting to you, skip down to the end of the indented quotes and our narrative continues there.

Episode 47 of Voices in AI **Podcast:**

Ira Cohen

Co-founder and Chief Data Scientist at Anodot

People think, because of the hype, that it does a lot more than it really does. We know that it's really good at classification tasks, it's not yet very good at anything that's not classification, unsupervised tasks, it's not being able to learn new concepts all by itself, you really have to code it, and it's really hard. You need a lot of good people that know the art of applying neural nets to different problems. It doesn't happen just magically, the way people think.

Episode 88 of Voices in AI **Podcast:**

Ron Green

CTO of KUNGFU.AI

Every now and then we will speak with somebody who will say, "[AI is] all hype, there's nothing really there." And I get that. Tech cycles come and go, and things can get overhyped. But I've been doing this a long time, I've been in this space long enough to have lived through the AI winter of the '90s. And what we can do now compared to 20 years ago, it's just breathtaking.

Episode 73 of Voices in AI **Podcast:**

Konstantinos Krachalios

Managing Director at the IEEE Standards Association

I'm not a supporter of the hype around artificial intelligence, also I'm not even supporting the term itself. I find it obfuscates more than it reveals, and so I think we need to re-frame this dialogue, and it also takes away from human agency.

Episode 60 of Voices in AI **Podcast:**

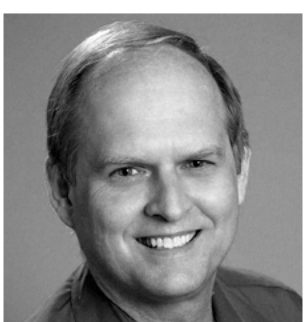

Robin Hanson

Author, Research Associate at the Future of Humanity Institute of Oxford University, Chief Scientist at Consensus Point, and Associate professor of Economics at George Mason University.

I don't see anything recently changing relative to the long-term trends. Now that doesn't mean something couldn't change, I mean there's no guarantee that past trends will continue. We should be ready for the possibility of change and think of how to prepare for that. But that's different from saying we're seeing a dramatic change. We're seeing another burst of interest and attention like the ones we saw before. So I was caught up in a previous burst you see, this is personal for me. I was a graduate student in 1983, and I read about cool things happening in AI, and so I left my graduate program to go out to Silicon Valley to get involved in AI, because I bought the hype at the time.

[Episode 53 of Voices in AI](#) **Podcast:**

Nova Spivak

Leading Technology Futurist, Serial Entrepreneur and Angel Investor.

First of all, it's great to see AI in vogue again. I lived through the first AI winter and the second sort of unacknowledged AI winter around the birth and death of the semantic web, and now here we are in the neural network machine learning renaissance. It's wonderful to see this happening. However, I think that the level of hype that we see is probably not calibrated with reality and that inevitably there's going to be a period of disillusionment as some of the promises that have been made don't pan out.

[Episode 50 of Voices in AI](#) **Podcast:**

Steve Pratt

Chief Executive Officer at Noodle AI

There's a tremendous amount of hype and B.S. right now out there about AI. People anthropomorphize AI. You see robots with scary eyes, or you see crystal balls, or you see things that – it's all magic. So, we're trying to be explainers in chief, and to kind of demystify this, and basically say it's just data and math, and supercomputers, and business expertise. It's all of those four things, coming together.

[Episode 35 of Voices in AI](#) **Podcast:**

Lorien Pratt

Chief Scientist and Co-founder of Quantellia

I think we'll have an AI fall, but it won't be a winter and here's why–we're seeing a level of substantive use cases for AI being deployed, especially in the enterprise, you know, widespread large businesses, at a level that never happened before. I was just talking to a guy earlier about the last AI hype cycle in the '80s, where VLSI computer design by AI was this giant thing and the "fifth generation," and the Japanese and people were putting tens, hundreds of millions of dollars into these companies, and there was never any substance. There was no "there" there, right? Nobody ever had deployed systems. AI and law, same thing, there's been this AI and law effort for years and years and years, and it really never produced any commercial systems, for like a decade, and now we're starting to see some commercial solidity there.

[Episode 19 of Voices in AI](#) **Podcast:**

Manoj Saxena

Investor & Educator in Trusted AI. Chairman CognitiveScale and AI Global. First GM of IBM Watson. UT Austin Faculty.

I think there is a lot of hype, and there is a lot of misperception about AI right now. I like saying that, "AI today is both: AI equals 'artificially inflated,' and AI equals 'amazing innovations.'"

Kinds of AI

So what is real and what is hype? What can technology do? To answer those questions systematically we have to take a long path to the answer. We have to understand what intelligence is and what computers can do, then we can have a solid basis for understanding just what is and isn't possible. To do that, we need to work down through each rectangle in this image.

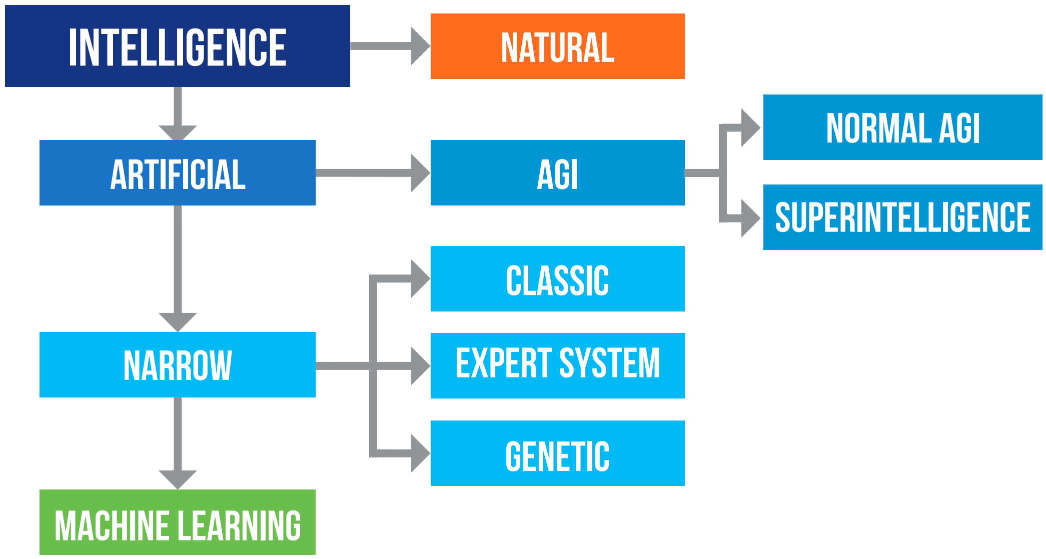

Let's start with **intelligence**. Our journey gets off to a rocky start because there is no agreed-upon definition of what intelligence is. In fact, that is overstating it a bit. There's no definition at all worth anything really. There's not even a "Here, are the five definitions people can choose between." The best research suggests there are hundreds of definitions, none of which has anything like a major following. So, you can either say it's a meaningless term or it's a term we don't understand, but that's our starting point. It is worth pointing out that intelligence is not alone in this ambiguity. We don't have a consensus definition for life or death either, but they clearly exist.

But whatever it is, there are two kinds of it. There is **natural** intelligence, which is what we have and is biologically based, analog in nature, and largely driven by chemistry and there is **artificial** intelligence, which at present and in the foreseeable future mechanical, binary, and electrical.

In terms of artificial intelligence, there are two kinds. One is **AGI** (artificial general intelligence). An AGI is the kind of AI you see in the movies. It is an AI as versatile as a human. It is Commander Data, it is C-3PO. No one knows how to build an AGI and the estimates for when we will have one vary from five to five hundred years. And yet, 95 percent of the guests on my podcast believe we will build one. How can this be if we don't know how? Simple: Those who believe we can build it readily admit that we don't know how, but base their confidence on the simple fact that people are machines, and we have AGI. So, the reasoning goes, since everything in the human body is governed by physics, then whatever we can do can someday be duplicated my machines obeying those same laws of physics. They reject any argument against this as a kind of "magical thinking."

The important thing is that the kind of AI we know how to do today may not, and I believe is almost certainly not, the "junior" version of an AGI. It isn't as though we have these techniques and are 5 percent the way to an AGI. Rather, what we know how to do today is a trick that can't lead to AGI because of the inherent limits of that trick.

Although this issue isn't about AGI, the two boxes off to the right of it in our diagram deserve passing mention. There are, in theory, two kinds of AGI. One is a **Normal AGI**, just a really smart computer, like the fictional robots I mentioned above. But there is another kind of AGI that is theorized, a **Superintelligence**. The idea here is that most humans have IQs in the range of say 80 to 120. Don't suppose that an AGI would be like 160. If computers double in power on a regular basis, someday the AGI will have an IQ of 1000, then 10,000, then a million. And something with an IQ of a million would have no more regard for us, one way or the other than we do ants or even bacteria. As Eliezer Yudkowsky succinctly said, "The AI does not hate you, nor does it love you, but you are made out of atoms which it can use for something else."

Those who are worried about AI are worried about AGI's and the possibility of Superintelligence. No one is worried about what we know how to build now, that is, narrow AI except to the extent it may impact jobs. But in terms of the big worriers out there, Elon Musk, Bill Gates and the late Stephen Hawking, they worry about an AGI, a technology that we don't know how to build and may actually not be possible, although this is a minority view.

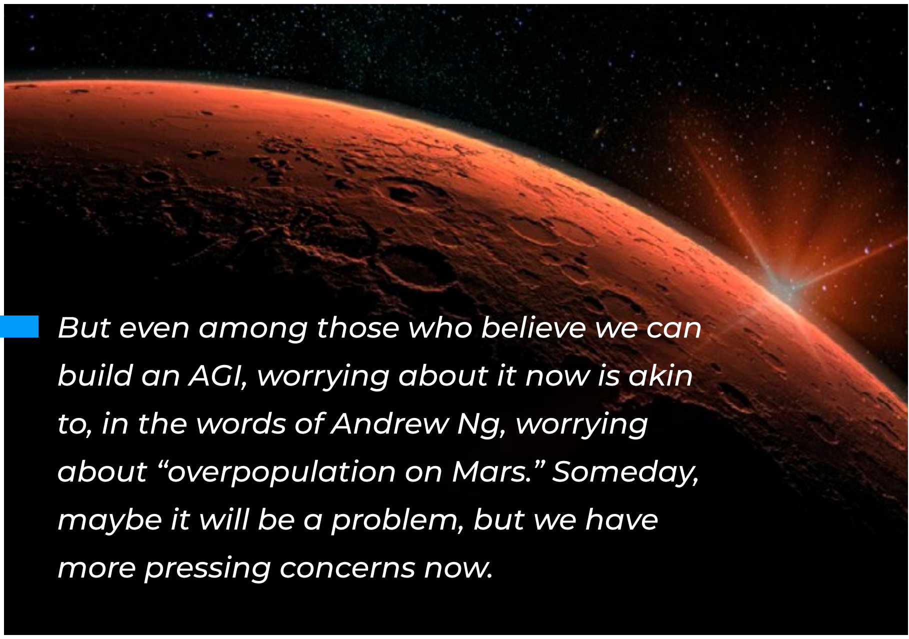

Musk and company disagree, and Musk has termed it an "existential threat" that he believes we will have in less than five years. He goes so far as to say that, "I hope humans are more than just the boot loader for the AGI, but it turns out that's probably what's going to happen." Time will tell, but for the record, I don't agree. In fact, I emphatically don't agree.

But enough about AGI. I only mention it here because we are snaking down the path to get to what we can do and what the limits of AI really are. And that brings us to **Narrow AI**, also called weak AI, this is an artificial intelligence that can do one thing very well. Narrow AI is what we know how to code now. It is what routes you through traffic on your way to work or filters SPAM out of your email. And there are 3 flavors of that, two of which we're not really going to discuss, so they go off to the side, one of those is **Classic AI**. Classic AI is what they did in the '50s and '60s, and it's still with us today. It's an old idea, hundreds of years old, that posits that since, "all of math can be reduced to symbols," perhaps you can reduce all of science to symbols as well? Can you make the words that I'm writing right now into some kind of logical, mathematical script? If so, you can apply rules of logic to it and arrive at Truth, or at least truth. And that was the hope of classic AI, that the world could be represented meaningfully by manipulating logical symbolism. That hope is what fueled the early optimism about AI. By "early" I mean the '50s and '60s. And we still use this all the time for little things. For instance, if you were going to build an AI that played tic-tac-toe, you wouldn't build a machine-learning algorithm that studied millions of games of tic-tac-toe, you would represent the game itself and every possibility of how to play it into logic. But life is complex, and classic AI caps out pretty early.

The second one of these narrow ones is an **Expert System**. An expert system is a rejection of this classical AI. It says, "Let's not try to figure out the system in question like in classic AI. Instead, let's just ask all the experts to tell us how they do something and then we'll instantiate all the rules in a hierarchy. For instance, we might take a group of doctors and ask, "How do you tell the difference between the cold and the flu?" and one doctor might say, "Well is there a fever? If there's a fever, then we …"

And you can usually spot expert systems because they are decision trees, so, when you call into a place and you push the button, you're navigating through an expert system. Expert systems are still very much with us since they work well for certain kinds of applications, such as in factories where a complex set of rules can be applied to how items flow through the factory. But you can't make a self-driving car this way.

Another kind of AI is listed on our chart as Genetic, but this is a catchall term for a variety of evolutionary algorithms. These techniques apply Darwinian "survival of the fittest" to the creation of AIs. They are often used in specialized cases, usually when you don't have training data to work with.

Say you wanted to design the most aerodynamic car possible. What data would you use to do that? There's not really any traditional data you could train that on, instead, you would say, "Here are the basic principles of physics, I want you to start evolving the perfect shape," and so it starts going through different possibilities, and when one is a little better, it "selects" for that attribute and starts trying variations on that theme.

Likewise, there is something known as Generative Adversarial Networks (GANs) that involves having two AI's, one trying to "trick" the other one. To do this, one iterates on an algorithm then tests the quality of its output by running it through the other AI.

These techniques are CPU intensive and not really the sorts of methods that the typical enterprise will work with, but they deserve a mention for two reasons. First, the innovative solutions they come up with and second, their lack of reliance on labeled data.

Now within Narrow AI, we have finally arrived at **machine learning**, and that's the thing that gets everybody excited, today. Machine learning is a really simple idea. It says, "Let's study past to find patterns with which to predict the future. Let's take all the data about the past that we have, let's pull all the patterns out of it that we can, and let's build a model of future behavior. At the root, that is all that machine learning is.

Imagine for a moment that somebody gives you 1,000 photographs of cats, and 1,000 photographs of dogs and they say, "Figure out an algorithm in your head that can tell a cat from a dog." It's actually a hard problem, isn't it? Because you might say "Cats have pointy ears." But then you find those flat eared cats and find dogs with pointy ears, then you might say "Well dogs have long snouts," and then you find those dogs with pushed-in faces, and then you might say "Cats have whiskers," or whatever. It is really hard to articulate the rules, yet people are really good at saying, "That's a cat, that's a dog." Humans, children even, can even handle special cases where there wasn't any data provided. A child can see a Manx cat, one of the tailless cats, and say, "Look! There is a cat without a tail," even if they have never been told that such a thing exists. No one knows how we do this, but AI researchers would really like to find out.

With machine learning the process is different. What computers do is take each of the photos and divide them up into, say, 2x2 pixel images, and label each of those as cats and dogs. Then take the same photos and divide them into 3x3, 4x4, 5x5 images, and so forth. Then it might take those photos and figure out all the lines, all the vectors, in each one of them and label all those as cats and dogs. Then all the color in the images, and anything else it can tease out. Now, instead of 1,000 photos, you have billions and billions of pieces of data, all of which have these attributes of catness and dogness. Then when a new photo is fed into the model, the computer starts looking for matches in that huge database. And it finds a certain number of matches in the cat database and a certain number in the dog database, and it uses all of that to make a catness score.

So unlike expert systems and classic AI, where the reasons that the model gets to the conclusions that it does, machine learning never tries to understand anything. The computer is not explicitly taught what makes a cat a cat.

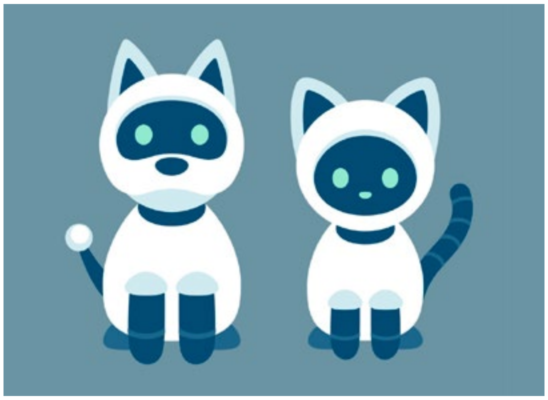

This is why explainability is such a challenge in a machine learning world because there is no "why" that would make sense to a human. Why does the computer think that is a cat? Not because it has pointy ears or whiskers or anything like that. The explanation of what really is going on doesn't map well to what we think of as human explanations.

Goals and Approaches

So, that is what machine learning is. Let's now take a minute to look at the goals and techniques we use, because these speak to what ML can do.

Let's start with goals. What is ML trying to do? There are only really three of them, Classification, Regression, and Clustering.

Classification is where you say, "There's a finite number of things this could be, and I want the computer to figure out what it is." Think of OCR. That's a classification problem. The computer has to decide between 26 letters, 10 digits, and a handful of punctuation marks. Is this a photo of a cat or a dog? That's a classification problem with two choices. There is nothing in between them – there are no half-cat half-dog hybrids.

Regression is a statistical modeling technique where the AI is asked to estimate a value along a continuum. So if you have a smart thermostat, it is trying to figure out where along the range of temperatures you likely want it to be set. If you are trying to estimate someone's income or likelihood to die next year or anything like that, you are looking at a regression problem. This is different than classification in that you are asking the machine to calculate a value, not pick between choices.

Clustering is the final goal. Say you looking for credit card fraud. You have all this data, maybe thousands of sources. But for the purpose of being able to visualize this, let's just say you have three pieces of data. You have the amount of a transaction, the distance the transaction occurs from the person's billing zip code, and the time of the transaction. Imagine a three-dimensional space where those three data sources serve as the axes. So you start entering transactions. Place a dot in your mind at $27, 41 miles, and 4:51 PM, for instance. Place a thousand white dots in this space. Then imagine in those 1000 transactions, 30 were fraudulent. Turn those dots red in your head. Now, with clustering, you are asking the computer to find clusters of red dots that contain as few white dots as possible. Once it solves this problem, you have an algorithm that determines if a transaction is fraudulent. If a new transaction's dot shows up in one of those clusters, it is likely fraud. Explainability is hard with clustering. Why is this? Well, imagine that instead of three-dimensional space, you are dealing with a hundred-dimensional space. Then, when a transaction is flagged as fraud and someone asks why the only answer is "It was in the cluster of other transactions that turned out to be fraudulent." There is no "why" beyond that because there is no attempt to understand fraud, just to predict it.

Those are the three goals AIs try to accomplish. And, broadly speaking, there are three techniques or approaches we use to achieve those goals:

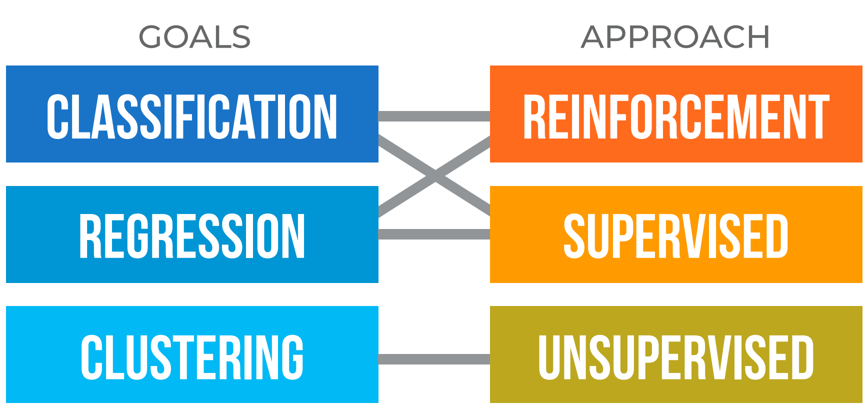

Supervised learning is the most common. That is where you take a bunch of well-labeled data and use it to teach the AI something. If you have a thousand photos of cats and a thousand photos of dogs, and they are labeled as such, that is supervised learning. This is the most common because most bread and butter AI problems are things like, "Given the sales data from this store last year, what inventory should I order for next year?" or "Given a list of plane tickets sold last year, what airplanes should the airline buy for next year?" We are usually dealing with structured data. Don't get me wrong the data may still be a mess. The vast majority of the effort in most AI projects is usually getting the data in a civilized format. But supervised learning is the goal.

Reinforcement learning is a bit different. With supervised learning, we label the data before it goes into the model. With reinforcement learning, we label the data as it comes out. Say you have one cat photo and one dog photo and you tell the computer to use those two photos to try to figure out if a third photo is a cat or dog. It spits out an answer, "Dog" and you say, "Wrong, computer." Then you give it another photo, it says "Dog" again, which happens to be the correct answer, so you say, "Right, computer." Rinse and repeat this a million times, and you will have a great algorithm. We use this approach when we don't have the data on hand. Imagine you were making a self-driving car. You might start by making your best effort at programming the computer, then put it in a car with a human driver. The AI isn't actually driving the car, the person is. But the AI is supposed to say to itself, "I would brake now" and if the human driver does in fact brake, then that's regarded as a "Correct." Likewise, if the car says, "I would change lanes now" and the driver doesn't, then that is an "Incorrect." Over time, the AI learns from the reinforcement it is given.

***Unsupervised learning** is where you can't do the other two techniques and thus rely on the computer to figure out the correct answer on its own. Clustering is generally unsupervised since you are basically throwing everything you have at the AI and saying, "You figure it all out."*

The holy grail of AI is the creation of a generalized unsupervised learner, something you can just point to the Internet and say, "Go figure out how the world works." Then, when it is finished, you can ask it "How can we extend human lifespans?" "How can we lower crime?" "How can we generate abundant clean energy?" and so on. Of course, it won't be that easy, but given that our world is full of messy, unlabeled data – which is what most of the Internet is made out of – getting data to train AIs is pretty hard. But an unsupervised learner would just sort it all out. In the second Avenger's movie, however, they plugged Ultron into the Internet for like fifteen minutes and then he decided to destroy the world.

Can we build a generalized unsupervised learner? There is a great book on this topic written by one of the giants in the field of AI, Pedro Domingos. It's called The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World. He believes there is a master algorithm to be found and suggests it might be really simple, the kind of thing that someone in Des Moines comes up with one boring winter's evening.

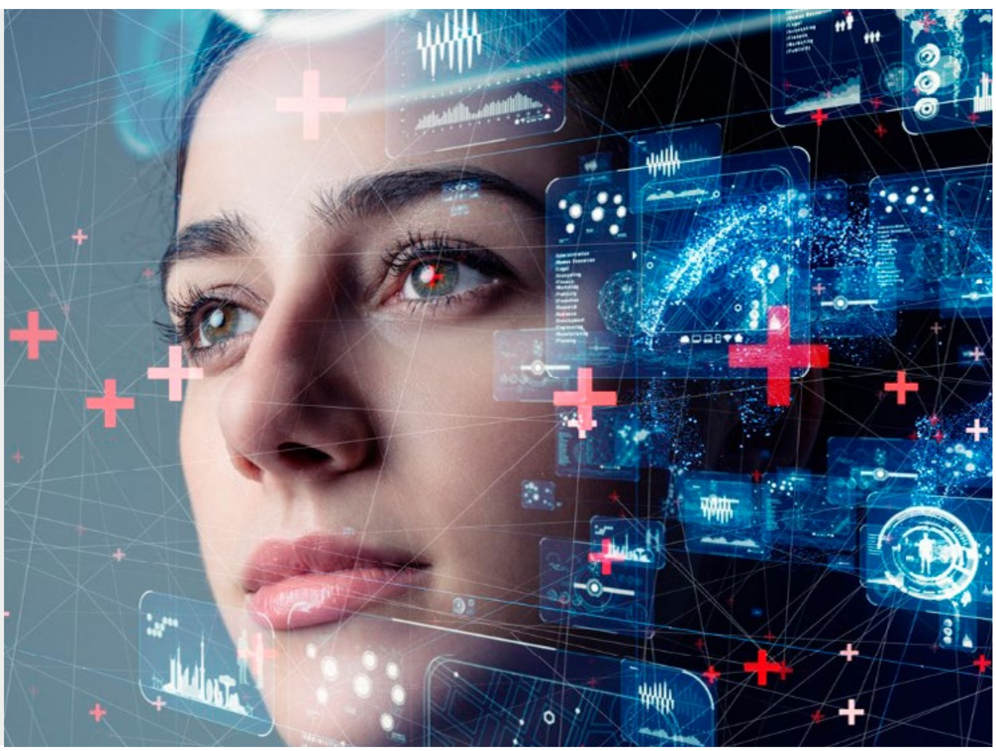

Examples of Goals and Approaches

- **Ecommerce.** The kinds of recommendation engines that you see on Amazon and other sites are classification systems. They have a finite number of products that they can suggest to you, so the question is, "Which of this set of products should I recommend to you based on you ordering this particular item?" And then, how do those companies hone those algorithms? It's reinforcement learning. Did you buy the item they suggested? If so, that is a "Correct" and if not, an "Incorrect." They get better and better every day, and the knowledge contained in those becomes an almost insurmountable barrier to competitors.

- **Matching symptoms to illnesses.** That is a classification system as well. There are a finite number of symptoms and diseases and systems like these are generally built using supervised learning. It would be helpful, by the way, to build a reinforcement learning system along these lines. For any entrepreneurs out there, here is an idea for you: Most of us, when we have a health issue, say, a twitching eye, go to good old Dr. Google and type in our symptoms. Invariably, what comes back is something like, "It is either benign" or "You have Lou Gehrig's Disease." Of course, there is a 99.9 percent chance that it is the former. What we need is a system where a month after you do this search, you are sent an email that asks, "What did it turn out to be?" Then that data, the reinforcement learning, is rolled back into the system so that when other's do that search, it says, "In 99.9 percent of the cases, this is…" This, and many more examples, are great places where reinforcement learning can, over time, build a really sophisticated system with little deliberate effort.

- **Recidivism rates among criminals.** That's a regression analysis where the machine tries to compute a likelihood that a given criminal will commit another crime. Usually, this is accomplished by supervised learning, where we have a list of criminals, what they did, and whether they later committed another crime. Hopefully, there is an aspect of reinforcement learning as well, that over time, there is a feedback loop into the system where new data can easily be integrated.

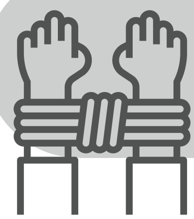

- **Face Recognition** is another classification system. When a person picks up your phone and it has to decide if you are you, that is a classification problem where the machine has two choices: It is you or it is not you. Likewise, a digital camera that automatically focuses on faces has to classify all the objects in a frame as either faces or not faces. The digital camera has to be trained with supervised learning. There is no opportunity for reinforcement learning. The camera never knows if it got it right, so your specific camera can't gradually get better. Sure, you could hypothetically have a camera where if the focus is not on a face, you hit the focus button again, basically telling the camera it is wrong. This isn't all that practical, though. The thing about AI models is that they are CPU intensive to build, but once built are simple to run. Cramming that kind of computational power and energy usage into the camera is probably not the best way forward.

- **Language Translation**. Similarly, language translation is another classification problem. It has recently gotten a good deal better, pretty quickly. Why is this? There are people who argue, and I am not one of them, that all of the advances in AI that we see today are because of Moore's Law, that's it's just computers getting faster. But there is a kernel of truth in there. We are able to do so much more because we have such powerful machines. Years ago, you may have tried to train a system with a few thousand pieces of data. But now we can handle billions and billions and billions of pieces of data. This is why translation got so good. Before the machine might say, "I have this word, and it means 'good' sometimes, and it means 'cool' sometimes." And so back in the day, you would say, "Well, we'll just call it 'good'." Years later the AI could say, "Well, when it's used before this word and after this word, it usually means good, however, if it's used before this word and after this word, but there's the word 'this' in the sentence, it means, cool, unless of course there's this word in the sentence, and it's used here, but not with this, unless this, on a Tuesday after a full moon, in which case it means good." So it can get as complicated as you want, and so the more data you have to train it, the more it gets these really rich nuanced examples.

Likewise, you might remember a few years back when AlphaGo stunned the world by defeating Lee Sedol four games out of five. This also was achieved through a great expense of hardware. That's not all, of course, it also used evolutionary algorithms to play against itself to train it. But in Game 3 there was this move, Move 37, that was so unexpected when AlphaGo made it that at first it looked like a bug. Sedol stared at it for something like twenty minutes before realizing it was what would later be referred to as a "divine move." The folks operating AlphaGo scrambled to see what was going on. They figured out that the odds a human player would have made that move was 1 in 10,000. And yet, this move won the game for AlphaGo. So the question is, where did that move come from? People said it was "creative," that "the machine is thinking for itself" but it isn't really. It is just that with enough data and with enough cycles, and with enough training, the AI can transcend human ability. And that is, of course, the goal.

Implementing AI in Your Enterprise

How do you implement AI at your enterprise? At GigaOm, we are asked this question frequently. It really is pretty straightforward, and can be broken down into five steps:

The first is to brainstorm. Sit down with your team and just think. Don't worry about technology. Instead just ask yourself, "What areas of our business could be improved using data?" Think of any different thing in your business that you could, in theory, transform into something bigger, again not that you're going to, nothing, just think of all your business challenges, and ask yourself, which of these could data inform.

Sort all of your ideas into three buckets. The first is task automation, how do you take something you already do and use a machine to do it automatically. Second is insight generation. That's, "How do I price better? How do I put better products in different locations? How do I stock my stores better?" You're asking the AI for knowledge to help you make choices on the various business problems you face. They might be completely new things, like scoring leads, that you may not be able to do right now or simply aren't doing. In that case, you might go to your sales team and ask, "How can you tell a good lead from a bad lead?" And you might take that as the basis to build a new system to score leads. And last, there are conversational technologies, such as chatbots, which can be used for support, customer service, presales and more. Look for areas where you're interfacing with other people outside of the company.

Step 1. Think of everything you would want to automate in your business that data might inform, insights that might be able to be generated through areas, and any business processes that are fundamentally conversational. Often times these are going to be places where you're already using analytics, or that are manual processes someone is doing in Excel.

Computers are really good at playing games because, games have defined rule sets, constrained universes, point systems of some kind, and clear winners and losers. There are a lot of things in life that don't behave like games. If you have a conversation with somebody, did you win that or did you lose it? But there are other things that look a lot like games. For instance, imagine that your company has a bunch of job candidate resumes. Additionally, you also have existing and former employees and their performance reviews. Can that be made into something like looks game like? Maybe. Perhaps you could find the resumes in the applicant pool that most closely match successful employees and use that to sort incoming resumes. This particular example has all kinds of problems, of course, not the least of which that resumes aren't structured consistently, maybe poor predictors of success, and you probably don't have enough current and former employees. But I include this as an example to illustrate the type of task that AI might be able to inform in your company.

Another example. Not all sales leads are created equal, right? So how do you score sales leads better? It's something that companies are probably doing informally in a way that looks like an expert system by looking at the source of the lead, the person's title, their company size, their industry, and so on. So you can go to your sales folks and ask them how they knew a certain lead was good or not, and try to automate that. Or you could take a machine learning approach and design it like a game. Leads that didn't purchase go in one bucket, leads that did purchase go in another, and the machine crunches all of that.

Step 2. Once you have your list of all of the places where AI might help your company, what next? The second step is to look at all of those tasks and ask yourself, "What data is needed to do this? What data assets are necessary?" And you should look at two kinds of data assets - internal and external, that is, data you have in house or data that you can buy. You have to ask questions around the accuracy of the data, the volume of it, and whether it really has enough of both successes and failures to train an AI. You can't teach a machine to tell the difference between a cat and a dog with 999 cat photos and 1 dog photo, but perhaps you could with 800 and 200.

Step 3. Ask yourself, "What individual services would this effort require? What technology would be required?" Maybe you need something that's really good at natural language processing, or that is math-heavy and does great regression or something that does different kinds of deep learning that's big on insight generation.

Step 4. You know the full range of business problems you might be able to solve, you know what data is available, and what sort of abilities you need in a solution. Next is to start shopping for products. This is pretty easy since there are many fine solutions out there. This wasn't the case not so long ago. AI has made much of its recent progress because the toolkits are so much better and, shall we say, more civilized. There are good solutions from companies big and small, and so the challenge isn't finding a good one, but finding one that best fits your needs.

Step 5. Spin up a number of projects. Projects should probably have timelines of weeks, not months, you should probably do multiples. And what you are likely looking for is not proving some technical ability to do a task. That's pretty certain that it can be done by the time you get to this step. What is important is whether it moves the needle. Does it create compelling value? Is there a business case for scaling it?

Companies that I have found to be the most successful do AI incrementally. They do a small project and it works, and then they make it bigger and then it works better. It is all pretty straightforward stuff, but many companies never take that first step because the subject matter, AI, has this aura of complexity about it that isn't usually warranted. You don't have to have a data scientist to do these sorts of projects that target low-hanging fruit.

For smaller enterprises, the best solutions are probably ones where you buy someone's "AI product" and use it on your own data. There are innumerable solutions out there, and while "AI" is often just a marketing label on a product, it isn't always.

Two Examples of Implementing AI in Your Enterprise

Examples of AI working in the world are too numerous to count. But often enterprises view their successes as proprietary information, for obvious reasons. They often don't want to share specifics of how they did things, and what did and didn't work, and so on. But even given that, it is easy to find real successes out there that you can learn from. On the **Episode 50 of Voices in AI**, I chatted with Stephen Pratt, the CEO at **Noodle AI** . They will themselves as a company that helps enterprises use "Artificial Intelligence to create a world without waste." I asked him to share one of their success stories:

Steve Pratt - Episode 50 of Voices in AI

Sure. I can take the work we're doing at XOJET, which is the largest private aviation company in the U.S. If you want to charter a jet, XOJET is the leading company to do that. The way they were doing pricing before we got there was basically old, static rules that they had developed several years earlier. That's how they were doing pricing. What we did is we worked with them to take into account where all of their jets currently were, where all of their competitors' jets are, what the demand was going to be, based on a lot of internal and external data; like what events were happening in what locations, what was the weather forecast, what [were] the economic conditions, what were historic prices and results? And then, basically came up with all of the different pricing options they could come up with, and then basically made a recommendation on what the price should be. As soon as they put in our application, which was in Q4 of 2016, the EBITDA of the company, which is basically the net margin – not quite, but – went up 5 percent, in the company.

The next thing we did for them was to develop an application that looked at the balance in their fleet, which is: "Do you have the right jets in the right place, at the right time?" This takes into account having to look at the next day. Where is the demand going to be the next day? So, you make sure you don't have too many jets in low demand locations or not enough jets in high demand locations. We actually adjusted the prices, to create an economic incentive to drive the jets to the right place at the right time.

We also, again, looked at competitive position, which is through Federal Aviation Administration data. You can track the tail numbers of all of their jets, and all of the competitor jets, so you could calculate competitive position. Then, based on that algorithm, the length of haul, which is the amount of hours flown per jet, went up 11 percent.

This was really dramatic, and dramatically reduced the number of "deadheads" they were flying, which is the amount of empty jets they were flying to reposition their jets. I think that's a great success story. There's tremendous leadership at that company, very innovative, and I think that's really transformed their business.

A second example I would like to relay might seem a bit whimsical, but it has an important lesson. It is about a former embedded systems designer from the Japanese automobile industry named Makoto Koike. Makato's parents are cucumber farmers. And, according to a Google post about this:

Watch the Cucumber Sorter video here.

Makoto's father is very proud of his thorny cucumber, for instance, having dedicated his life to delivering fresh and crispy cucumbers, with many prickles still on them. Straight and thick cucumbers with vivid color and lots of prickles are considered premium grade and command much higher prices on the market.

But Makoto learned very quickly that sorting cucumbers is as hard and tricky as actually growing them. "Each cucumber has different color, shape, quality, and freshness," Makoto says.

Sorting cucumbers is a tedious task, and according to Makato:

The sorting work is not an easy task to learn. You have to look at not only the size and thickness, but also the color, texture, small scratches, whether or not they are crooked and whether they have prickles. It takes months to learn the system and you can't just hire part-time workers during the busiest period. I myself only recently learned to sort cucumbers well.

So Makato built an AI system to do it. Of course he did. The main controller is a Raspberry Pi 3 which is attached to a camera that photographs the cucumber. Then, he uses TensorFlow to make sure it is a cucumber, and if it is, the image is forwarded to a more robust Linux box also running TensorFlow to be classified according to the various cucumber attributes. He built the mechanics of the system, the conveyor belts and such, using inexpensive off-the-shelf technologies. And it works.

I tell this story to point out that if cucumber sorting is a potential use case for AI, then just imagine how many more there are in a typical enterprise. I don't know the ROI on the sorter, and perhaps it was a passion project, but the principle is clear: "Applications of AI are everywhere."

Four Ways to Fail

01

Unclear Ownership - If you remember the '90s, people had "web departments" because the world wide web was new and only a few people had this new knowledge, and so they were all siloed. Of course, you wouldn't ever have that today because the Internet is in every department. But I see a lot of successful companies having an AI team or strike force or some other action-oriented group. The bottom line is if it's unclear where this lives on your org chart, whether it's in IT or whether the CIO owns it or what, it will be hard to act with the kind of focus this technology merits.

02

Lack of Talent - People fail by picking the wrong problem to solve. You have that big list of all your business challenges, you have all that data, you have all that technology, but which projects actually get the green light? There is no easy answer here, and perhaps this a place to seek some outside counsel.

There are 4 ways to fail with AI. Well, I guess technically there's an infinite number of ways to fail. But there are four pretty common ones.

03

Align to Existing Skill Set - Strive to find projects at the right skill level. The projects you pick need to realistically align not just with the data at hand, but with the technical talent at your disposal. You don't need a data scientist to do most AI projects, which is good news, because they're in short supply, but, some projects are much harder than others, and they require different skill sets. So try to make sure that you're picking things that you can execute on. Make sure that you have a good fit between the problems you're trying to solve and the skills that you have on your team.

04

Poor UI - The first three may be obvious, but the fourth one is less so. UI matters a whole lot. The worst thing you can do is build a system that works, but which no one uses due to poor UI. Technologists in the AI world tend to be laser-focused on measurable results and less so on implementation.

Machine Learning's Three Virtues

Machine learning, as a technique, has three great virtues.

Scales - First, it scales very well. The more data you feed it, the better it gets. Because of this, it can be used to solve really big problems, humanity-scale problems, where we have enough data. When we finally have self-driving cars, and billions of miles are being logged a day, the cars will get better and better.

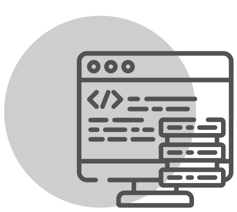

Low Barriers to Entry - The second benefit is that generally speaking, machine learning has low barriers to entry. Sure, the AlphaGo team had something like $25 million in hardware, but with a Linux box, the cloud, and some open-source libraries, you can go to town.

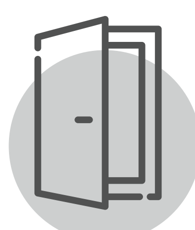

Bias Free - Finally, it has the benefit of being bias-free. It isn't inherently bias-free, because data selection is sometimes an editorial process, and thus prone to human bias. But in the vast majority of cases, it is bias-free. If your enterprise ties point-of-sale data into an AI prediction system, you should get unbiased results.

If your AI system is collecting rainfall data to predict future droughts, there isn't necessarily going to be bias inherent in the data. It is good at removing anecdotal evidence and vague opinions and replacing them with findings based on data. Again, humans can be biased, but the machines are not. Humans can feed in biased data, and end up with a biased AI, but that's not the computer's fault.

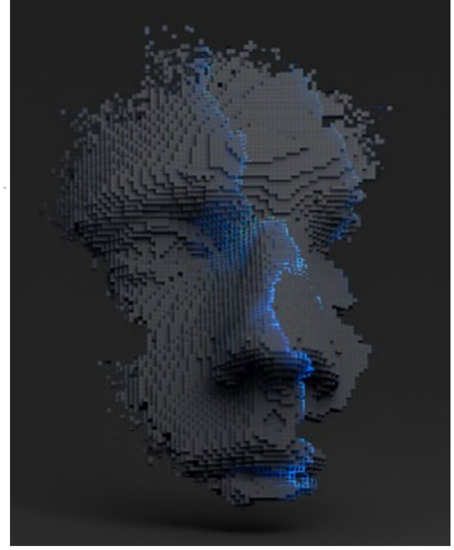

Machine Learning's Five Faults

But, in addition to these three virtues, machine learning has five faults.

- **The first fault is that ML is very superstitious.** This ties back to the biased remark above. You know how sports teams often get these superstitions, where, maybe a guy stubs his toe one morning, and that afternoon, they win the big game. The next day, he stubs his toe again and they win the game. All of a sudden, everybody on the team is stubbing their toe. Because many things in life and business actually have few data points, we sometimes connect things that we shouldn't. We can build AI's that are superstitious by using too little data. If two times out of two the stubbed toe leads to a win, then that's probably superstition. But if you stub your toe 100,000 times and you win 100,000 times, and you didn't stub it 100,000 times and you lost 100,000 times, well, you've just discovered a truth about the universe. You can't call that superstition.

- **The second fault is that transfer learning is not solved.** This is the biggest problem by far in all of artificial intelligence, in my option. Transfer learning is where we take knowledge from one discipline and apply it elsewhere. Humans do it naturally, from early childhood. If I drop a nickel in the water and it sinks, what do you think will happen if I drop a quarter in? You know that, but a machine doesn't. If you learn how to look for deer hiding in the woods, and then someone tells you to keep an eye open for bears, you are already three quarters the way there. It is so natural to us, we don't think about it. But to a computer, well, it has to learn everything from a blank slate. There are a couple of places we are making progress, largely in language and dialects, but the trick behind it is elusive. It has something to do with how humans build internal models of reality and know what elements transfer to other things. Drop a feather into water, and you know it will float. Drop a dime in the air and it will hit the floor, but so will a feather. How do you know that? You probably haven't tried it.

- **The third fault, machine learning can't solve hierarchical problems very well, like, natural language processing.** For example if you wanted to teach a computer to have a conversation, you can't just feed it the full text of the Internet. It is good at figuring out what you might type next in an email, but you can't give it every novel ever written and hope it will write a bestseller. I can't even find a chatbot that can answer the question, "What's bigger, a nickel or the moon?"

- **The fourth fault, it cannot distinguish correlation from causality,** a classic problem that humans deal with as well.

- **The fifth fault relies on the future being just like the past.** If the future isn't like the past, all it does is just train on the past worked, and all it does is it figures out how it could have made a million bucks in the past. But, as the mutual fund ads say, past success is no guarantee of future performance.

5 FAULTS

SUPERSTITIOUS

CAN'T TRANSFER LEARN

CAN'T SOLVE HIERARCHICAL PROBLEMS

CAN'T TELL CAUSATION FROM CAUSALITY

DATA DEPENDENT

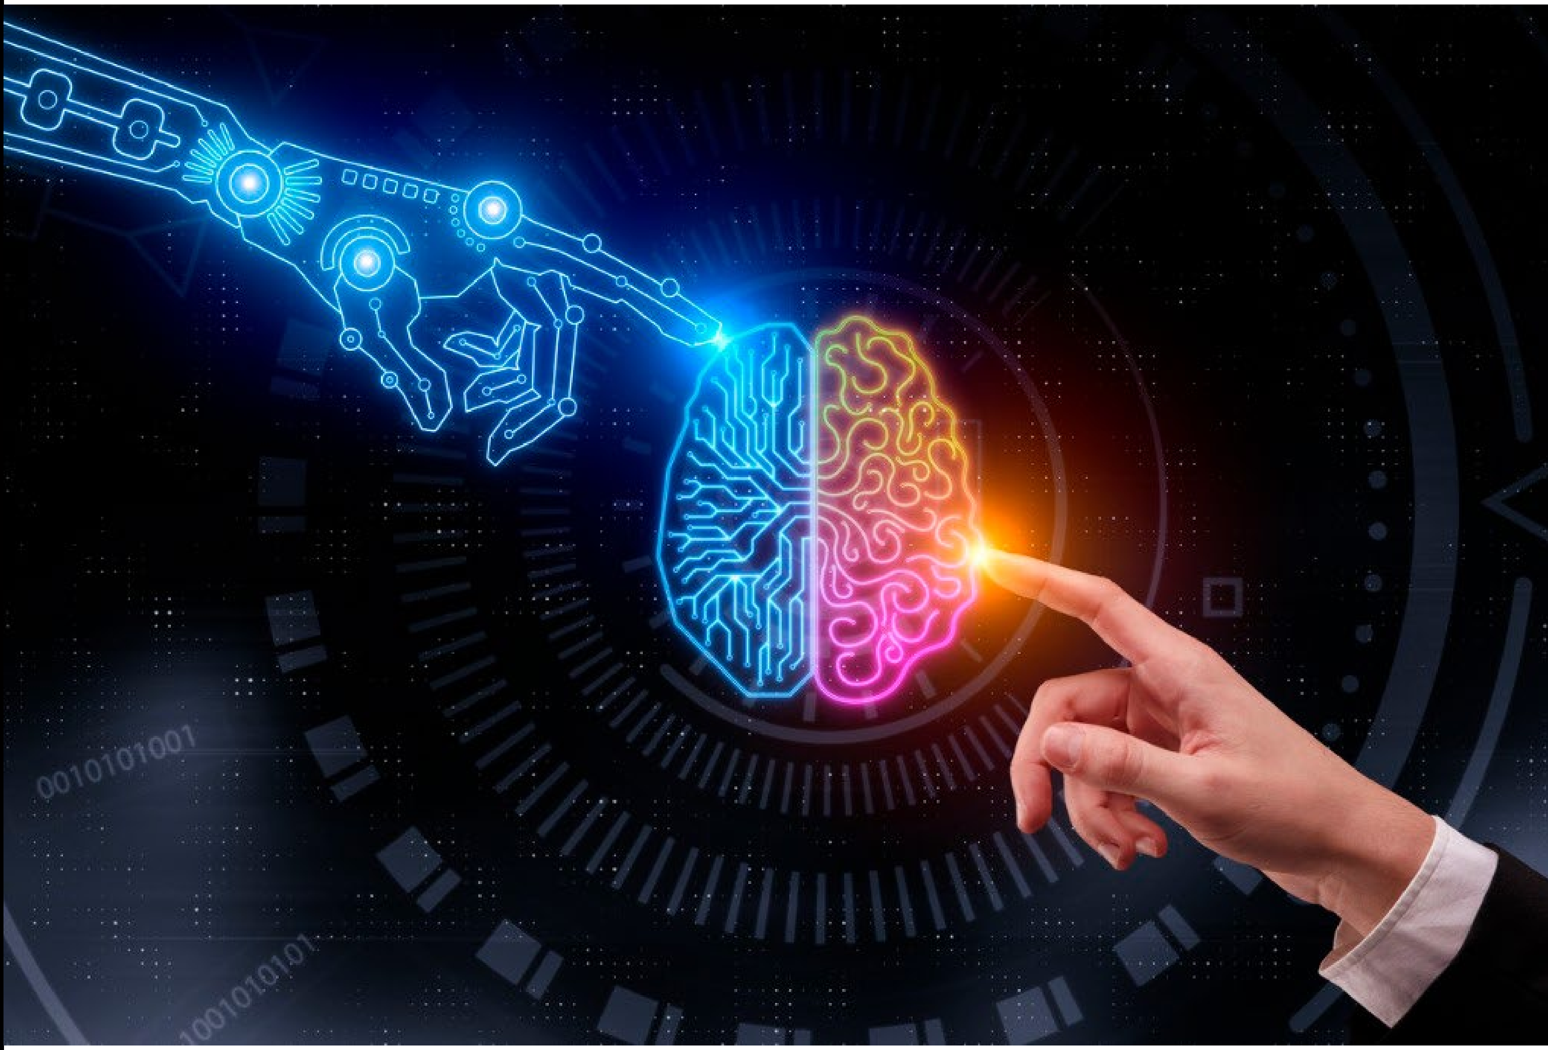

Conclusion

We've taken a long path to define the limits of AI. I hope at this point, that they are a little clearer. Since AI is driven by data from the past, it works on problems where the future and the past are similar. It works in places where enough data exists and is, or can be, labeled to train the machine. It works best in domain spaces that resemble games in the sense that there are defined rules and clear objectives.

This broad outline is both expansive and constrained. It is expansive in that there exist an abundance of things that fit it. We live highly unoptimized lives where we follow our gut, or at best, anecdotal evidence. The number of places where data could improve our lives and our businesses boggles the mind. But it is constrained in the sense that there are more things in life and business that AI can't help on than there are that it can. Maybe this will change at some point, but right, don't ask it to write a sonnet by feeding it Shakespeare. JK Rowling and Lin Manual Miranda aren't going to be out of work any time soon. It can't tell you what to get your spouse for their birthday, unless you have data on their last few thousand birthday gifts and what they liked and didn't like. It won't help you find purpose in your life or business. But all of that is fine. That's all human work. But there is plenty of stuff out there that the machines can do better than we can, and the sooner we put them on that job, the better.

DEEPDIVE INTO AI

About Byron Reese

As the CEO and Publisher of GigaOm, Byron leads the company in their mission to help business leaders understand the implications of emerging technologies and their impacts on business, media, and society. Byron produces and hosts GigaOm's podcasts Voices in AI and The AI Minute and has published two books: "Infinite Progress: How Technology and the Internet Will End Ignorance, Disease, Hunger, Poverty, and War" and "The Fourth Age: Smart Robots, Conscious Computers, and the Future of Humanity."

Byron possesses a diverse body of patented work, and enjoys exploring the intersection of technology, history and the future with both technical and non-technical audiences around the world.

LEARN MORE:

bryonreese.com
@byronreese
voicesinai.com
@voicesinai
gigaom.com/shows/

GIGAOM

To learn more about how we help transform enterprises in AI-enriched data-driven world, visit gigaom.com

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization. To apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.