Spiritual AF Sundays with The Mystic Geek

Season 1, Episode 1

The Art of Artificial Intelligence

Transcript

Hi everyone.

We're going to be talking about artificial intelligence. What are we talking about when it comes to AI? Why are we even having this conversation?

You likely have seen this app out there called Lensa. I jumped in on the hype train with it.

And what Lensa is: For a small fee and uploading up to 20 images of yourself, it will mash together all these different styles of artwork to create - one hundred, two hundred, however many you pay for - profile images of yourself. Some of these are cool, and some are weird.

I jumped on the hype train. I was curious. Then, I started seeing social media and other areas: how the data for the artwork was derived. That's why I wanted to have that topic.

To recap what we'll be talking about today. Over the past few years, artificial intelligence has made leaps and bounds in its abilities as it redefines how we interact with technology. Things like robotics, machine learning, natural language processing, and big data are a few ways that AI has transformed virtually every industry out there. If your industry deals with any type of data - any kind of trying to predict what people want - it more than likely uses AI.

As for common concerns when it comes to this emerging technology, it's usually apprehension around a couple of different things:

- Job security or the ability for people to continue making an income
- Privacy issues when it comes to that data
- Copyrights
- Accuracy when it comes to these data models making decisions or predictions

And these are all valid issues to consider!

But the most thought-provoking question that AI raises today is how it changes our understanding of consciousness. So we're going beyond the practical concerns, starting to go philosophical.

- Will AI eventually have any type of autonomy?
- Will it have free will?
- Can we teach it to make moral decisions?

And then flip to focus, what does AI's development say when it comes to traits that we initially thought were either uniquely human or unique to living beings, like creativity, innovation, complex decision-making, and empathy?

Consider these things today as we go into the topic of artificial intelligence.

When we look at artificial intelligence, we are all in different places regarding understanding. First, start by discussing concepts and terms to ensure that we are all on the same page. I might say one thing, and you might think I mean another. I want to handle that right off the bat.

The next bit is going to be the concerns concerning AI. Before we jump into the spiritual implications, we need to look at real life, now and today, and the situations that are going on concerning the use of

artificial intelligence. Last but not least, now that we've built that foundation, we can look at the spiritual implications of AI.

What is Artificial Intelligence?

First, let's talk about concepts and terms: what is artificial intelligence? In the fifties, experts in computing wanted to develop a way to differentiate machine-based intelligence from what we see in humans and animals.

During that timeframe, mathematician and logician Alan Turing came up with this philosophical question, "Can machines think?".

Let's think about this for a moment. We can't observe whether or not someone else is thinking. There's no way of measuring this construct that we call the mind. So, he tweaked his question. Rather than asking, "Can machines think?" he changed it to, "can we observe them engaging in intelligent behavior?".

In 1995, Stuart Russell and Peter Norvig published “Artificial Intelligence: A Modern Approach.” This paper lists four different potential goals for AI.

- Systems that think like humans

- Systems that act (make decisions) like humans
- Systems that reason—thus taking emotion out of the picture
- Systems that act rationally

When you look at various uses of AI or ideas around AI, you can put them in one of these four buckets.

Narrow AI

There are different types or flavors of artificial intelligence that are out there. The one that most of us are familiar with is what's called Narrow AI. The idea behind this is it's very focused, and the machines need a way to go beyond what developers train them to do.

For instance, if you're using Netflix or Twitter, the platform recommends things for you to watch. Even though this is narrow, it is still reasonably complex in what it is doing. Narrow AI means the machine can't learn to act or do things beyond the specific models provided to it.

General AI / Super AI

That leads us to the concept of artificial general intelligence and superintelligence. These are more in the realm of science fiction.

Suppose you're looking at science fiction books or movies about AI, for instance, "Westworld" or "I, Robot"; these stories are about machines being to that point of being human, being self-aware, making decisions, making moral decisions even. We are still far from that. That's not to say this is impossible because technology moves so quickly, but these are more of the theoretical realms of artificial intelligence versus the practical everyday topics we explore.

Let's go into terms of artificial intelligence.

Machine Learning

In machine learning, we don’t know what algorithms to use to get the desired outcome. There are hundreds upon thousands of data points. We tell the machine, "Here's the data. Here's the outcome we're looking for. You determine the algorithm".

What machine learning does is much like what we would do with the scientific method. It comes up with a hypothesis. It runs the scenarios. It matches the generated outcomes up against the result that we're

looking for. And it says, okay, is this it or not? And then it'll keep running experiments.

Machine learning focuses on optimization and identifying a process for decision-making. In contrast, we might need to figure out what to do, but we have the time and energy to figure it out. We might have limited experience and perspective. This scenario is where we benefit from working with machine learning.

Deep Learning

Deep learning is a subset of machine learning, which is a subset of artificial intelligence. This system uses layers of algorithms, called neural networks, to make conclusions.

One example of the use of deep learning is image recognition. You have a layer that assesses shapes. Another layer identifies the eyes and studies their placement. Another layer determines body gestures or facial features.

What Are the Concerns Surrounding the Use of AI?

Let’s discuss the elephants in the room: concerns regarding artificial intelligence.

Data, Copyright, and Privacy

First and foremost, there are concerns surrounding data collection and usage. When we're utilizing machine learning or deep learning, there has to be a foundation of data for those machines to understand the problem and formulate their algorithms to make decisions.

We're asking about the data. Individuals are the initial owners of their data and creative works. If a machine uses that data or information, do they have a right to it? Did the owners give informed consent?

We're not just talking about them clicking "Yes, I agree" on a Terms of Service (we'll talk more about them in a bit). Do the people who agree to the TOS fully understand the implications of what's happening?

Some companies are scraping data from internet sites like Pinterest and DeviantArt, where users might have yet to give full permission to access that data. What are the implications of that?

We're looking at some more significant problems with data mining or sourcing. Large Scale Artificial Intelligence, Open Network, LAION, is

a nonprofit that releases free data sets for AI research and development. Take some time to look into how it gathers this data and dodges copyright laws or ideas of who owns or should get attribution.

Data scientists use LAION's material to train image generator models like Stable Diffusion, the model that Lensa and other image generation platforms use. Many of us didn't realize that this dataset includes people's medical records, which LAION likely shared without consent.

It's essential to see how resource aggregators pull in data for use. If you're curious whether AI models are using your photos or images, haveibeentrained.com is the name of that website to check out.

Skewed Data Models

So this then leads us over to the question of machine learning. Who is setting the direction? So when we talk about machine learning, we come up with a starting point. We come up with the endpoint. The machine does all the testing to come up with the models to be able to optimize what it is that it's setting out to do.

Now here's the issue - how someone uses machine learning or engages with it impacts future users' experiences. Suppose a machine learning tool is utilized by people who may be bad actors,

bringing up issues or ideas or adding problematic data. In that case, that will impact future users' experience, and it can cause problems down the road.

Let's talk about machine learning and Lensa and how this all comes together. So Lensa uses a mix of Stable Diffusion AI, which we talked about a moment ago, plus user-submitted photos to generate profile images. The source artwork plus end-user feedback can skew the model that produces these images. One of the things that came up was that whatever gender you choose determines how sexualized the pictures of you are in that final output. Some of the resulting images can be NSFW if you indicate that you identify as female. You're paying for these generative profile photos, and suddenly you're being provided with nudes without your consent.

Biased Data, Biased Data Models

When we look over the use of machine learning and data, another problem that could come up is skewed input equals skewed output. Anyone who does any type of analytics is probably familiar with that garbage and garbage out. The same thing can happen here regarding machine learning and its impact on equality.

Without oversight, machine learning applications can reinforce cultural biases. We look at healthcare, for instance. The concepts of "sick" and "healthy" are not concrete; they delineate who meets a certain standard and who does not. How do developers define these terms? Then there’s the issue around audience sampling and demographics. Was it equal across the board of various ethnicities or backgrounds, or were people of specific backgrounds included as a majority? There’s nuance when it comes to other issues in life they don't get handled.

In profit-based systems, machine learning can adversely affect how members of marginalized groups receive care because it's just looking at the surface info. It does not realize what's going on or the whys.

The Terms of Service Don't Equal Transparency

And then more than transparency is needed. You can put the fricking "Mein Kampf" in a Terms of Service; people will agree to it because they're not reading it. Also, the end user is not meant to be the expert; we are not experts on all these different things, nor should we be.

Most of us don't have the time, energy, or expertise to understand the legalese in a Terms of Service fully, nor do we have that understanding to see the consequences of our actions when it comes to things like machine learning until it is too late.

Is AI "Stealing Jobs"?

As Artificial Intelligence continues to revolutionize the workforce, it's forcing us to reevaluate our view on economic prosperity. What would happen if we adopted a system that ensured fundamental human rights for all - like secure housing, healthcare, and nutrition? We could navigate through this paradigm shift with greater ease!

Instead, access to those things often depends on the value of your labor in society, which determines your income and what you can afford. The problem isn’t necessarily that artificial intelligence is taking jobs. It’s that automation is taking away jobs within a society where someone may need that work and income to maintain a fundamental standard of living.

We're dealing with the context issue here. Not necessarily AI being the bad guy, machine learning, or the apps, but us living in a pretty

effed-up world is the issue here. So we have all this stuff that is going on.

Blurred Accountability

We are starting to identify the increased need for societal and legislative oversight. Societies that leverage artificial intelligence need better laws surrounding copyrights and privacy. We also require supervision to ensure that artificial intelligence and machine learning models are generated and used ethically, particularly if they impact business and social policies.

The Spiritual Impact of AI

Now we can talk about the spiritual implications. We needed to be on the same page regarding the terms. We also needed to make sure that we addressed the practical concerns when it came to AI. So let's start getting into some of its more philosophical and spiritual aspects.

WTF Is Spirituality?

I'm going to lay out another term for you because some of you're probably thinking, "What the heck is spirituality?"

Spirituality looks at connection and meaning. It focuses on our relationship with the world around us as well as ourselves. So understanding what the meaning of life is, what is our purpose? What is our place in the world? Those are essential aspects that fall into the realm of spirituality. We come up with our values and view of our life purpose through all of those. And through that, we can see whether we will live with integrity with those values through our actions.

AI Pros and Cons on the Spiritual Aspects

Some people believe that artificial intelligence will help us spiritually because it helps free our time and focus. If we're using AI to handle the menial, repetitive tasks, it could free us up to focus our time and effort on the things that we value in life.

So, for instance, if we have a lot of automation through artificial intelligence and other tools, and we have a world that uses universal basic income that gives us the ability to put our time, effort, and energy into things that could benefit the social good. That's the plus side.

Now here's the shadow side.

If we don't have oversight, or if people are bad actors, it can encourage social engineering when it comes to using AI.

In some cases, it has.

If you have used social media since the 2016 U.S. Presidential election, you probably noticed how much things have changed.

Free platforms like social media thrive on our engagement, so they adjusted their algorithms to keep us on their apps longer. They did this by curating content that they believed would maintain our attention. Since these apps could identify our preferences, we saw content aligned with our beliefs.

This process created engagement bubbles, limiting our exposure to different views and further polarizing us. It hindered empathy and skewed our worldviews.

Redefining Humanity

Artificial intelligence has the potential to challenge what it means to be human. Are there characteristics that we consider to be solely human traits? Are there traits exclusive to living beings, such as complex

problem-solving or empathy? As we've observed in other animals, we know those traits aren’t exclusively human. Heaven forbid you upset a crow or raven; everyone else in their little flock will find out about it.

What happens when we observe machines exhibiting traits - empathy, adaptability, or the ability to make moral decisions - that we initially thought were exclusive to humans or organic beings?

Tangent: AI and Ethical Decision-Making in Science Fiction

Let's dig deep to discuss artificial intelligence and ethical decision-making.

There are two examples from entertainment that I’d like to share.

The first is a movie I previously referenced, "I, Robot." In this movie, Will Smith's character, Detective Spooner, was in a terrible car accident. Spooner was aware of another person in danger,12-year-old Sarah Lloyd. A rescue robot intervened but could only save one of them.

If two people were facing death and you could only save one, who would you choose: the grown man or the child? According to Detective Spooner, he believed that Sarah’s life should have taken priority and that any human being would have come to the same conclusion.

However, the robot scanned the vital signs of both him and Sarah and determined that the girl only had an 11% chance of surviving with intervention, whereas Spooner had a 45% chance and thus saved the detective.

The robot didn't make this decision in a vacuum. Someone had to program the robot to focus purely on the survival factor and not consider cultural programming, such as believing that women and children take priority over the lives of men. There's so much to unpack from that moment when it comes to ethics and the impact of culture.

Let's switch to the other example.

I am a big fan of the Mass Effect video game series. The primary antagonist throughout the series is the Reapers, a robotic race that targets sentient life within the galaxy. We don’t learn about their origins until a DLC for Mass 3 shares a cautionary tale.

The Leviathans, an aquatic race whose existence precedes the Reapers, saw themselves as superior to the other races. The other sentient races died off once they reached a certain level of technological development, and the Leviathans intervened by tasking their machine intelligence to preserve life at any cost.

Those machines do precisely that, but not in the way that their makers intended. This intelligence determined that the mutual development of organic and synthetic (AI) life led to inevitable mutual destruction. Rather than focusing purely on survival, the machines defined preservation as maintaining records. The conclusion was to harvest those races and preserve their genetic memory in machine form while allowing the other evolving races to develop, where the Reapers would collect them at a later cycle. In a sense, the machine intelligence, the Reapers, followed their mandate to preserve life at any cost. Were they acting morally? That depends on your values and frame of perspective.

Okay, enough tangents.

AI and Art, Creativity, and Innovation

That brings us to spirituality and art. I recommend you check out the slides in the show notes because I included screenshots of Twitter conversations on this topic.

The question, “what is art?” has been debated for who knows how long. One fascinating way I've seen this brought up is the concept of art as a conversation between the artist and the audience. Artists may

not talk one-on-one with you, but the artist has a vision and expresses it through their work. The audience seeks to understand what the artist is “saying” through their work.

AI-generated creations are more akin to conversing with the machine or yourself. What ends up happening is you're inferring the meaning. The audience is trying to piece together the essence of what they see versus taking in the artist's intention.

With human-generated art, you can connect with the artist's vision or just freaking ask them. Whereas with a machine, not so much. It's a black box where you cannot fully see what's going on that generates that outcome.

That brings us over to the whole concept of imagination. What is imagination that fuels creative works? Who is the artist? Is it the machine or the user who provides the prompts?

Another question emerges when you try to simplify the artistic process. Human artists observe the world, possibly studying existing works to determine what "good” art is. They create based on an inner drive to relay a particular emotion or experience.

Meanwhile, machines gather information, and with prompting, they create.

What is the difference between AI-generated work and humans creating art? Both draw upon the outside world. Do we say AI-generated work is not art because it requires external direction?

If you can think back to art class, your teacher was giving you prompts of, "This is what I want you to do for your art assignment," and you create it, another human being is essentially influencing you to make something. Is artificial intelligence - is the human prompt in that - is that a similar type of process?

So even if it requires human influence, such as a prompt, can an AI's process still be considered imagination?

We don't know.

How We Impact the World Through Our Choices

A considerable component of spirituality is reflecting on the impact of our actions. If we're looking at things from a grand scheme, an individual's actions may seem small, but they are still important because of how we can influence others and impact one another.

Supporting Artists

We can influence the future on an individual and a grand scale by being more aware of how we spend our time, money, and focus. We can be more intentional about the AI platforms with which we interact.

Talking primarily about the United States in this point - since we live in a society where income determines a person’s access to housing, food, and health care, we can support individual creators through patronage or promotion. So we can either pay for their work or promote them. We must provide attribution when we share another person’s artwork online (assuming the artist permitted us to do so). While we may not patronize a specific artist, we can connect them with someone who will. Since art sales rely on a connection (marketing and networking), we can support creators who may not have access to such things.

Supporting Watchdog Organizations

We need to support organizations trying to ensure that companies ethically develop artificial intelligence tools. The Algorithmic Justice League is just one example out there. They're focusing on removing

bias in AI algorithms, and you can find out more about them from ajlunited.org.

Legislation

For those of us in the United States, it's a democracy. We are participants in things that go on. We shouldn’t discourage ourselves from acting because of skewed systems and power dynamics. Small actions can have a significant impact, especially if we organize. There are ways that we can work with our representatives in Congress to create laws to impede bad actors when it comes to artificial intelligence or machine learning, data usage, and those sorts of things.

In February 2022, the Senate introduced The Algorithmic Accountability Act. S 3572. Congress referred this bill to the Committee on Commerce, Science, and Transportation. We can keep track of this bill and the committee’s findings and work with our representatives to move this bill forward. Go to congress.gov and then search for artificial intelligence. Read up on the existing bills and then reach out to your representatives in the House of Representatives or the Senate to support further action.

Keep in touch, and focus your time and energy on understanding how our government creates change that can handle things on a broader level. It can impact corporations, impact the direction of society. So this is more than just simply voting people. We need to be active participants in all of this.

Concluding Thoughts

I hope this provided you with additional insight into the issues surrounding the current uses of artificial intelligence.

And again, the big focus here is to be mindful of what you're consuming and using. Try to understand a bit more. While you don’t need to be an expert, staying informed is a good idea. We want to support the people doing good work who can understand what's happening. We need laws and policies to protect those of us whose privacy and livelihoods may be affected by further developments in artificial intelligence. Even if it doesn’t negatively impact you, these issues can negatively affect others. As good global citizens and spiritual people who see our connection with others, it's crucial for us to take action and be more aware of what's happening.