



UNITARIAN UNIVERSALIST

CONGREGATION OF FREDERICK
Spirituality · Community · Justice

**“No Fate But What We Make”:
Darwin Day, Transformational Artificial Intelligence & Human Values**

The Rev. Dr. J. Carl Gregg

February 12, 2023

frederickuu.org

Charles Darwin was born 214 years ago on February 12, 1809. His birthday has come to be known as [International Darwin Day](#), an annual opportunity to celebrate the principles that guided his life: **“perpetual curiosity, scientific thinking, and hunger for truth.”** These values resonate with our UU [Fourth Principle](#) of “A free and responsible search for truth and meaning” as well as our UU [Fifth Source](#) of “reason and the results of science.”

One reason it is significant to celebrate Darwin Day in UU congregations is that both sides of Darwin’s family were [in large part Unitarian](#). Some of our Unitarian and Universalist forebears were also among the earliest religious leaders to embrace the paradigm-shifting implications of Darwin’s discoveries—that **we humans are *not* a little lower than the angels; we are “a little higher than the apes” with whom we share a common ancestor.**

In the wake of the Human Genome Project, we now know that at the DNA level, there is only a [1.23 percent difference between humans and chimpanzees](#). In other words, humans and chimpanzees are 98.77% *the same* on the DNA level. We humans are not uniquely-superior creations; we are merely *one evolved species, among many*, deeply interconnected with all the ecosystems and other varied forms of life on this planet. As our UU [Seventh Principle](#) affirms, our invitation is to have “respect for the interdependent web of all existence of which we are a part.”

Denying our place within the animal kingdom and the larger natural world has contributed to humans exploiting this incredible planet on which we find ourselves. Darwin Day is an annual reminder urging we humans to recalibrate our place in the grand scheme of things — how we think of ourselves — from a more scientific, ecological, and evolutionary point of view.

This year for Darwin Day, I want to draw our attention to the many ways that Darwin’s guiding principles — “perpetual curiosity, scientific thinking, and hunger for truth” — are now stimulating the evolution of *artificial* intelligence.

In the past, I preached about AI in a sermon titled “**Immigrants Aren't Coming for Your Job, Robots Are.**” So, why am I revisiting this topic now, five years later?

Each summer, I create a draft of our Sunday Service themes for the coming year. In the middle of that process last June, [The Washington Post](#) reported that an engineer at Google was raising the alarm that Google’s AI chatbot had already become *sentient*.

How many of you have seen the 1968 film *2001: A Space Odyssey*? It’s a classic example of science fiction writers speculating that AI will one day become conscious, self-aware. The most famous scene involves an astronaut named Dave and a newly-*sentient* AI named HAL:

Dave Bowman: Open the pod bay doors, HAL.

HAL: I'm sorry, Dave. I'm afraid I can't do that....

Dave Bowman: I don't know what you're talking about, HAL.

HAL: I know that you...were planning to disconnect me....

Dave Bowman: HAL, I won't argue with you anymore! Open the doors!

HAL: Dave, this conversation can serve no purpose anymore. Goodbye.

That’s a future most of us want to avoid. We don’t want hostile computers taking over. We don’t want Skynet coming online and launching a nuclear attack, as depicted in the *Terminator* movies. I could keep going with many more pop culture examples from *Ex Machina* to *Her* to the early seasons of *Westworld*.

The question raised last summer was whether an AI had already become sentient, not in some imagined sci-fi future, but in our own present reality. Google

denied the allegation, and the engineer was fired the next month. In the meantime, the Internet was set ablaze with interest about the future of AI.

That flashpoint last summer persuaded me to schedule today's Darwin Day sermon on the subject of AI. Then, last fall, the plot thickened. (But wait, there's more!) Just a few months ago, in November, the research lab OpenAI released a free public version of ChatGPT, which is their version of Google's LaMDA chatbot, the one the engineer claimed had become sentient. Headlines began exploding again about the future implications of AI.

ChatGPT is the website where I typed in "Write a sermon about Artificial Intelligence for a Unitarian Universalist congregation." Within seconds, it generated the fairly impressive 350 words that you heard for our Spoken Meditation. How many of you have played around with ChatGPT or a similar AI chatbot? To be honest, 350 words is about 2000 words short of a full sermon, since the currently available *free* version of ChatGPT has built-in limitations that reduce the processing load on the AI system.

Also, remember the Internet adage that, "**If you're not paying for the product, then you are the product**" (The Social Dilemma). The free version of ChatGPT is constantly improving, evolving — *learning* — from all the ways we humans are using it. A future *paid* version of ChatGPT is coming soon which will likely be more capable of producing fairly convincing full-length sermons, college-level papers, and much more. If your job involves writing, perhaps you can think of ways these sorts of AI power-writing tools could impact your work — in both positive and negative ways — in the coming years.

If I were asked to grade the AI-generated spoken meditation we heard earlier, I wouldn't give it an "A." It's not the most elegantly written piece of prose; but if I didn't know it was written by an AI chatbot, I might give it a "B." It is *clearly* written with some perceptive insights.

How are these AI feats possible? At the risk of oversimplifying, these AI chatbots have been trained on a massive amount of data reflecting how we humans communicate, data widely available on the Internet and in various digital archives. The AI chatbot uses all that data to help it predict the next most-likely word in a sentence,

the next after that, and so on. The result is new, original texts that can feel deceptively humanlike — indeed, sufficiently humanlike to lead at least one hapless Google engineer to believe that the Google chatbot had passed the tipping point into sentience.

I want us to reflect on the implications of these technological advances, but first there is one more important piece of information to get out on the table. We've been focusing so far on *text*-based AI, but this fall there was still more breaking news about the power of AI to create *visual art* in response to short written prompts. (But wait, there's more!) Similar to the chatbots, these AI image generators are trained on the huge number of digital images out there on the Internet and in various archives. The AI image generator continuously refines pixels until it “matches” a particular text description to a visual image.

Some of the most well known AI image generators are MidJourney and OpenAI's DALL-E2, which is named after both the surrealist artist Salvador Dalí and the animated robot WALL-E from the Pixar film. How many of you have already played around with AI image software?

You can think of the most random things, and the AI system can create one or more similar or related visual representations incredibly quickly. A famous example that OpenAI used to demonstrate their new DALL-E2 technology was typed into the system as: “Teddy bears working on new AI research underwater with 1990s technology.” Within seconds they got [this image](#).

Pretty good, right? Looks photo-realistic to me, but it's completely computer generated within seconds.

Let me give you a few examples I created. Many of you will be familiar with the famous painting, *The Scream*, by Edvard Munch. The painting's agonized face is an iconic representation of human existential anxiety. Since I'm a meditation teacher, and since the first Buddhist Noble Truth is, “Life is suffering,” I wondered what the painting would have looked like if Munch had been a Buddhist.

So, I typed into DALLE2, “Buddha Statue in the style of Edvard Munch.” Within seconds it produced these images:



At least to me, these renderings of the Buddha feel quite profound and worth meditating on. Also, note that DALL-E2 often generates not one, but *four* options. You can then choose which of these you prefer, and then generate four more variations. And then keep going to refine and co-create an image that best represents what feels right to you.


I’m also a big fan of the Mexican painter, Frida Kahlo. What if we tried “Buddha Statue in the style of Frieda Kahlo?”



These Buddha-Kahlo images also feel quite well worth meditating on.

Whereas ‘*The Scream Buddha*’ tears my heart open in compassion for the existential reality in which all sentient beings find ourselves, my response to the ‘Frieda Kahlo Buddha’ is a softening into a beautiful serenity.

Or what if van Gogh, famous for his painting *Starry Night*, had been a Buddhist? “DALL-E2, show me “Buddha Statue in the style of van Gogh.”



I love these images too. And I could envision building a daylong meditation retreat around just these three sets of images that we've explored so far. Ok, enough Buddhism.

Since the surrealist painter Salvador Dalí is the partial namesake of the AI tool generating these images, what if Dalí had been born in *India* instead of Spain? "DALL-E2, show me, "Hindu images in the style of Salvador Dalí."




Or what if *Rembrandt* had converted to Hinduism and started painting the goddess Kali?



I love these too.

Or since I'll be leading a pilgrimage to Israel and Palestine this summer, what if the French impressionist painter Claude Monet had moved to the Holy Land to paint landscapes?

Show me, “The Al-Aqsa Mosque in the Old City of Jerusalem in the style of Monet.”



Or what about Georgie O’Keefe, known for her moving paintings of colossal flowers and New Mexico landscapes? What if her muse had instead been the Al-Aqsa Mosque?



Or Pablo Picasso, co-founder of the Cubist movement. What if he had been a Christian?



Cubist Christianity. I’m here for it!

Ok, just two more from me. Since I lived in Louisiana for many years, I wonder how many of you know the remarkable work of Clementine Hunter, a self-taught Black folk artist. What if Clementine Hunter had been practitioner of the traditional African religion of Yoruba?

Show me, “Yoruba religious images in the style of Clementine Hunter.”



That’s a somewhat obscure request. But DALL-E2 knows all about Clementine Hunter and Yoruba, and within seconds can create an original image combining the two!

Or I was on Zoom the other day with my meditation teachers’ cohort, and we were playing around with the AI art generator Midjourney using the share screen feature. We wondered what it would look like if an opossum experienced Buddhist enlightenment.



A friend who saw this picture made the savvy observation that this bodhisattva opossum could be the cover of a bestselling children's book!

That was a lot of images, but that’s part of the point. **In total, it took me only a matter of minutes to generate *all* of those images.** Even the most talented and

efficient human artist would require vastly more time to research and create anything like what we just saw. Such created images can also be a *beginning* to greater creativity, not an end, not the final product.

The images I've been showing you so are also all very quick first drafts. But if you keep working with such an AI systems, you can infinitely refine and enhance your own imagination, skills, and art.

As a final example, last fall *The Washington Post* ran an [article](#) on AI-generated art that won first place at the Colorado State Fair arts competition. It is “so finely detailed the judges couldn't tell” that a *computer* created it instead of a human.

The current state of AI innovation is already impressive, and it's only going to get better and faster from here on. How quickly will it improve? To give you one point of reference, a senior analyst who professionally studies A.I. risk currently estimates that there is a **35 percent chance of “transformational A.I.’ — defined as A.I. that is good enough to usher in large-scale economic and societal changes, such as eliminating most white-collar knowledge jobs — emerging by 2036”** ([The New York Times](#)). That's a mere 13 years from now. Don't miss that there is a *66% chance* it will take longer to reach transformational A.I. Regardless, evidence is steadily emerging that transformational AI could well happen within the lifespan of humans alive today.

Programs like ChatGPT are already being used to “write screenplays, compose marketing emails and develop video games” ([The New York Times](#)):

- What about **finance professionals**? Since 2017, JPMorgan Chase has been using AI to almost instantaneously review certain types of financial contracts that previously took more than three hundred thousand hours annually for humans to review. That's a lot of people's jobs that are now being done by computers.
- What about **doctors**? In 2018, AI was developed that, “diagnosed brain cancer and other diseases faster and more accurately than a team of fifteen top doctors” and “identified malignant tumors on a CT scan with an error rate *twenty times lower* than human radiologists.”
- What about **lawyers**? An AI system trained to spot potential legal issues in contracts has been shown to have a 94 percent accuracy rate as compared with an 85 percent average for human lawyers. The humans needed 92 minutes on average to complete

the task, whereas the AI needed *twenty-six seconds*. When you are being billed by the hour, online AI legal assistance starts looking like an appealing option. (Roose 29-31)

- What about **therapists**? It might sound terrible, but what if that AI application was programed with perfect recall of all the best therapy training in the world; was willing to patiently listen to you for as long as you needed; and was available twenty-four hours a day year-round?

I could keep going with other professions, including my own; what if they trained an AI on every UU sermon ever preached? But you get the idea.

Futurists are likely correct that, **“Artificial intelligence, cloning, genetic engineering, virtual reality, robots, nanotechnology, bio-hacking, space colonization, and autonomous machines are all likely coming, one way or another”** (Douglas Rushkoff, Team Human, 211). There’s a lot to say about all that, and we’ll keep exploring these forthcoming changes in future Sunday Services. But regardless of what technological innovations the future brings, our challenge in the present in light of our UU values is to **demand that *human rights and environmental sustainability be an integral part of any future.***

We also need to be honest that there are fortunate people like me — and a likely number of you as well — who love your jobs, or most parts of your your jobs. But for a lot of other humans on this planet, jobs aren’t great: the hours are too long and the work isn’t life-giving; many are literally backbreaking, hazardous, degrading, unethical, boring. In such cases, maybe it’s good news that the robots are coming for at least those jobs (120-121). **What’s so great about jobs?!**

Many **psychologists have long assured us that what truly makes life worth living is not merely work; it’s *meaning and purpose.*** If AI and robots take over many or most of our jobs, many other options for meaning include refocusing our lives on creating community, cultivating hobbies, enjoying this beautiful world in which we find ourselves, spending time with family and friends, and so much more. That doesn’t sound like a bad future. To be honest, such a future world could also include a lot of

people playing video games all day, but that's all part of the challenges of co-creating a better world — together.

In the words of one philosopher who studies future existential threats to humanity: **“There is a long-distance race between humanity’s technological capability, which is like a stallion galloping across the fields, and humanity’s wisdom, which is more like a foal on unsteady legs”** (Brian Christian, [The Alignment Problem: Machine Learning and Human Values](#), 313). We need to up our wisdom game — and fast!

In the face of coming advancements in AI, one major tool for maintaining human rights may be a **Universal Basic Income** ([Patheos](#)). How would we pay for that? At least in part by *taxing* some of the income being made by AI systems. If the robots take our jobs, **we are fools if we allow tech billionaires to amass the vast majority of AI wealth for themselves when we could instead mandate that wealth be shared**, so that all human beings have a chance to live dignified lives.

We don't have to reject technology, but we must demand that all sentient beings — all stakeholders on this beautiful planet — be part of the processes of guiding that emerging technology. In the words of Kevin Roose from his book [Futureproof: 9 Rules for Humans in the Age of Automation](#), **“There is no looming machine takeover, no army of malevolent robots plotting to rise up and enslave us. It’s just people, deciding what kind of society we want”** (Roose xxxvi).

Remember the Declaration of Independence: governments derive “their just powers from the *consent of the governed*.” Our choice is whether “We the people” will passively let technology — and tech billionaires — shape and *use us*, or whether we will instead choose to *use technology* skillfully, *wisely*, and compassionately.

“Remember...’The future is not set. There is no fate, but what we make for ourselves” ([Terminator 2](#)). Whatever will be is up to *us*, both individually and collectively. And I am grateful to be on this journey with all of you.