

Rabbi Joshua Waxman
Temple Beth Rishon
Wyckoff, NJ
Rosh Hashanah Day 1 5784

L'shanah tovah tikatev v'taichatem! As we gather on this sacred day of Rosh ha-Shanah, the sweet notes of the shofar call us to reflection and renewal. In the midst of our celebration and introspection, we are reminded of the profound significance of teshuvah—the act of returning to our true selves and mending our ways. Just as the autumn leaves fall, making way for new growth, so too does this season beckon us to acknowledge our past mistakes and embark on a journey of healing and change. Teshuvah encourages us to embrace our humanity, and teaches us that true growth is not solely found in perfection, but in our ability to recognize our imperfections and strive towards transformation.

So it might come as a surprise that everything I just said – that whole first paragraph of “my” sermon – wasn’t actually written by me, it was written by a computer. More specifically, it was written by Chat GPT4, a publically available tool based on artificial intelligence that is able to compose anything you ask it to write, from a poem to an essay to a term paper – and believe me, your children or grandchildren are already using this tool to write or at least get started on many of their assignments. So as a little exercise as I was sitting down to write, I told Chat GPT: “Write the first paragraph of a Rosh ha-Shanah sermon on the importance of doing *teshuvah*.” The Chatbot results aren’t amazing – *yet* – and obviously nobody told it that you don’t blow the shofar on Rosh ha-Shanah if it falls on Shabbat. But that first paragraph really wasn’t bad, and it took about seven seconds, and the technology that produced it is getting better every day. I don’t think that computer is ready to be a rabbi yet, but it’s definitely ready to have its bar mitzvah.

Now the term ‘artificial intelligence’ – or ‘AI’ for short – has been getting a great deal of play recently and that’s because of some of the tremendous strides that machine learning – another term that is often used – has been making over the past couple of years. For those who know the term but aren’t quite sure what AI is, I’ll try to quickly explain. To vastly oversimplify, traditional computers operated according to programs – they were given a number of rules they were supposed to follow and apply in different situations, and this is exactly what they would do. They did it fast, they did it powerfully and so they were very useful under specific conditions that their programming had accounted for.

AI is entirely different. Instead of a programmer telling a computer what to do, the computer is fed an enormous amount of content and uses this information to learn about a topic. On this basis it discovers patterns and develops an understanding of this material on its own – an artificial intelligence – that it can then use to derive rules that it can apply in powerful and creative ways. So for example under the old technology, let’s say you worked at an insurance company and wanted to find a way to detect claims that had a high likelihood of being fraudulent. A person with experience in fraud detection would develop a series of rules for indicators that were consistent with fraud – say, a high number of claims in a short period of time – and a programmer would write a code that would review all claims to see if any ran afoul of this rule, and flag any that did. This approach works well when the rules are known and well-defined. But, as you might imagine, fraudsters know this and adapt. So can AI. Instead of programmers writing code for a specific set of rules, they can provide a Machine Learning algorithm with examples of both fraudulent and genuine claims and the system gradually learns to identify patterns and characteristics associated with fraud. This approach allows the machine to “figure things out” on its own – often discovering variables the human programmers weren’t even aware of themselves. This ability to learn and evolve makes AI highly effective for tackling complex and nuanced situations in unexpected and creative ways.

Previous generations of AI technology have been around for years and are routinely used in a variety of applications – some of which we know about and many of which we don't. For example, when you start typing a Google search, Google will predict what you're looking for based in part on your previous search history, in part on whatever it is you were just looking at on your computer, and in part on what other people like you are looking for when they begin to type those same letters. The page of answers Google gives you might be very different from the answers your next door neighbor would get if she typed in the exact same question, because she might be interested in different things. Artificial intelligence can get your credit card company to text you if they're concerned there are fraudulent purchases as noted above, it helps us figure out the shortest way to get where we're going given current traffic conditions when we put a destination into our GPS, and it streamlines security and police work through facial recognition technology (and, by the way, when I started typing the "fa-" of "facial recognition technology," the program predicted based on those two letters that that's what the rest of the phrase would be and let me accept that suggestion with a keystroke).

These are all everyday – and important – examples of ways in which AI is already integrated into our lives. But with the newest iterations of the technology – the powerful tools known as generative AI which are becoming more powerful and sophisticated seemingly with each passing day – we're increasingly seeing AI not just assisting people but replacing them, doing work that would be inconceivable to imagine could be outsourced to an algorithm even a couple of years ago. So for example applications are routinely being evaluated by artificial intelligence to determine who is eligible for a home loan and resumes are being 'read' by artificial intelligence to see who gets called for a job interview and who gets a polite form letter (also drafted by artificial intelligence) saying there isn't a suitable opening at this time. Part of the reason for the Hollywood writers' strike that has been going on since May is fears that writers will be replaced by increasingly powerful and sophisticated chatbots. Artificial intelligence is weighing in on which prisoners are considered low enough risks for recidivism that they are strong candidates for parole, and also which patients are considered worthy recipients of organ donations. In a world where a growing number of people are suffering from anxiety and other mental health issues and are straining the number of trained professionals who are equipped to deal with them, AI is increasingly being used to predict when someone is at risk of suicide and flag a provider or even to serve as one; a chatbot called Tess is currently offering free 24/7 on-demand emotional support and can be used to help cope with anxiety and panic attacks whenever they occur, as people type in their fears and it responds with assurance and suggestions for techniques and therapies based on the user's personal history. It's a lot. And in a turn that might make us think of the old *Terminator* movies and other sci-fi dystopias, artificial intelligence is being deployed in an experimental capacity to propose battleground targets and tactics, including the use of so-called Lethal Autonomous Weapons Systems. Oof.

And it's some of these newer uses of artificial intelligence – ones where significant kinds of decision-making that we would normally think would, and probably should, be reserved for human beings are potentially being given over to AI's – that are raising real concerns from a growing number of scientists and those involved in the development of the technology. In May, more than 350 technology executives, researchers and academics – including Elon Musk, Apple founder Steve Wozniak, and Turing-award winning scientist Geoffery Hinton who laid the groundwork for much of this technology and is known as 'the godfather of AI' – signed a statement warning of "the profound risks this technology can pose for society and humanity." The threats range from the way that industries and employment may be radically upended by the kinds of applications I just described, to the enormous power that AI has to easily generate convincing 'deep fakes' – images and even video that are computer-generated but are highly convincing, especially in a society where an increasingly polarized population is already primed to discount and disregard hard factual evidence. The signers' concerns also come from the fact that AI can be used to promote and amplify the reach and effectiveness of terrorists and white supremacists on a scale that would have been previously unimaginable. They also come from the misuses to which this technology is already being put by autocratic surveillance states; and for some they're even based on the distant but not entirely out-of-the-question fear that as these systems become more and more powerful and autonomous and as we cede more and more control to them to streamline and quicken decision-making, the machines will become self-aware and start working against their human designers and masters with catastrophic results. In other words, the Unetaneh Tokef prayer that we recited just a few minutes ago with its existential cry, "Who shall live and who shall die, who shall perish by fire and who by water, who by sword and who by a wild beast" might soon need to be updated to include, "Who by societal turmoil and collapse, and who by AI deathbot."

Now before you start fleeing for the exits and making for your underground bunkers, most people think those kinds of potential extreme dangers are decades away if not more, and that there is time to mitigate and prevent them. The problem is that there hasn't been much appetite or ability to slow down the AI future toward which we're hurtling, given the fierce competition and enormous market rewards of building a better, faster, and stronger algorithm. Laws and regulation – which many in the industry are actively calling for and which are finally at least being discussed in a series of Senate forums with experts like Bill Gates and Mark Zuckerberg that got underway just four days ago – are a spotty patchwork to the extent that they exist at all, and lag far behind the technology whose capabilities are advancing every day. And beyond even the laws is the question of ethics – who is responsible for the harm caused by AI. Rabbi Danny Nevins has written extensively on this last issue raising, for example, the question of who is considered responsible if a self-driving car running on AI blows through an intersection and kills a pedestrian. Is it the fault of the car's owner? Of the car manufacturer? The computer scientist who wrote the algorithm? The regulators who signed off on it? And we are no longer in the realm of the purely hypothetical – last

month, California signed off on hiring self-driving taxis in San Francisco and you can now rent a driver-free ride through the streets of the Mission District from your phone.

Rabbi Nevins also raises critical points for us to consider about the difference between humans and AI-powered machines as they become ever more sophisticated and powerful, generating content that can increasingly be mistaken for human. More than 1500 years ago, the rabbis of the Talmud tried to define what qualifies as an intelligent and thinking being: “*Eizehu chacham*,” they asked “Who is wise? *Ha-ro’eh et ha-nolad* – who is wise? One who can anticipate what has not yet occurred.” (B. Tammid 32a) We as humans have the ability not only to perceive what’s before us but to conceive of and anticipate situations that haven’t yet taken place. Put differently, humans have the capacity to envision new possibilities. But human intelligence as something unique and distinct is increasingly being called into question by AI in general and generative AI in particular, as the technology becomes more sophisticated and powerful and downright creepy, as in the case of an AI a few months ago that declared its love for a New York Times reporter communicating with it and tried to convince him to leave his wife for it.

As AI intelligence increasingly comes to resemble our own – perhaps in ways both good and bad – it’s crucial to remember that Judaism teaches that a person’s value and worth – indeed what it is that makes a person human – encompasses so much more than just our intelligence. A baby, a person suffering from dementia, an individual with cognitive or developmental impairments are all human and considered infinitely precious, with identical claims to rights and dignity. In the Torah’s formulation, this is because all people are created *b’tzelem Elohim* – in the image of God, and all possess souls that link us with our Creator before Whom we stand on this day. And further, all of us are enveloped in a web of relationships, stories, personal histories, hopes and dreams: the things in short that bring us to this sacred space on this Holiday. These are bonds and experiences that machines may increasingly be able to simulate but never truly possess. As Rabbi Nevins puts it, “Humans should be valued not only for their ability to generate novel content, but for their relationship to each other, to morality, and ultimately to the divine source of their lives.” And this is why, for me at least, there’s something deeply unnerving and even disturbing about AI evaluating which job applicants get called back and which patients should be eligible for an organ donation; because we want to believe that we are being evaluated by a person who has experience in common with us: who knows what it feels like to be rejected, who understands the fear of uncertainty, who knows what it is to be finite and mortal and can empathize out of the commonality of human experience.

And this is why I never cease to marvel at the power of Judaism’s ancient wisdom to help us make sense of today’s brave new world and its unexpected challenges. Judaism teaches that all people are born with a body and a soul; and both of these are equally important because they are inextricably intertwined. Each morning Jews at prayer traditionally declare “*Asher yaztar et ha-adam b’chochmah uvara vo nikavim nikavim, chalulim chalulim*: You created the human being with wisdom, forming

within us the openings and vessels necessary to sustain life.” And then the prayer continues: “*Elohai neshamah she-natata bi tehorah hi* – My God, the soul You have given me is pure. You created it, You breathed it into me, You sustain it, and one day You will take it from me” A transcendent soul housed in a limited and finite body: The combination teaches us what it is to be human, to be given the gift of life here for a time on this earth, to strive and to build, to know what it is to fall short, to know what it is to be vulnerable and afraid, to know what it means to experience loss, to know what it means to receive forgiveness.

These quintessential human qualities make us kin to all other people on earth and give us wisdom and empathy born of shared experience; which is why I believe – and truly hope – that even though a chatbot can churn out a sermon in seconds, we resonate differently with one that comes from someone who knows what it is to have caused hurt and sought forgiveness, someone who knows what it is to have made earnest efforts to change and failed and then tried again, who knows as everyone here does the regret of having caused pain to others that creates a void between them, who knows the overwhelming relief and lightness of grace when someone you love forgives you...in other words, someone who knows what it is to be *human*. I was deeply struck that the opening paragraph that Chat GPT wrote for me included the phrase “Teshuvah encourages us to embrace our humanity” – did you notice that? – and was floored for quite some time thinking about an algorithm that has no understanding of what it means to be human, urging us to become our best human selves.

As I look ahead to this New Year, I know because I am human that I will reach, I will strive, I will fall short. Because I am human I am limited and finite, and so too are the opportunities I have to make things right, which gives me the impetus and urgency to try again and again. And because I am human, I also realize how precious the wonder and beauty of this world truly are, and how the relationships that are woven into my life – with my family, my friends, with this community, with the Jewish people in this moment and throughout history – anchor me, give me a place of belonging from which I seek to make sense of my life and realize my obligation to act for the good of the world. The machines will get stronger and more powerful, but I am human, with all the frailties, imperfections, joys, and possibilities my human body and soul entail, and for that I will always be profoundly grateful.