HR Tech & AI

Beyond the Algorithm – Exploring the Pitfalls and Transformative Possibilities of AI

Interview with Dr. Joe Perez

Dr. Joe Perez speaks at numerous conferences all around the world each year, he's a technology and professional development expert and an Amazon bestselling author with multiple new releases. He currently serves as senior systems specialist and team leader at the North Carolina Department of Health and Human Services.

What are some of the greatest pitfalls you have seen with regards to AI implementation?

One of the greatest pitfalls of AI implementation that I see is the erosion of human creativity. When organizations begin to rely too heavily on AI. I think there's a real danger of stifling innovation. C. A. I. Algorithms, while powerful indeed operate within the confines of their programming right and the data that they have been fed and this can lead to what I would like to call a homogenization of content and ideas where the unique spark of human imagination is extinguished by, I don't know, formulaic and uninspired outcomes.

Imagine a future where, say, artists, musicians, and writers simply become operators of the AI engine, their creative expression reduced to the mere pushing a button. Buttons that generate predictive results, the vibrant flame of human creativity, which has always driven progress for centuries, risks being smothered by the uniformity of this algorithmic output.

That's what you get when you trust in generative AI alone. Another significant pitfall that I see, is the loss of objectivity, when AI is trusted without any human oversight. AI systems, even though they're efficient at processing vast amounts of data are only as objective as the data that they've been trained on.

Now, when biases are embedded into that data, AI can perpetuate and even exacerbate these biases, and that leads to discrimination outcomes and this becomes especially concerning to me in areas like criminal justice, hiring practices, and financial practices, where the consequences of bias are particularly severe. Without the warmth of human empathy, and the ability to understand context and nuance, AI can turn justice into a cold mechanical process, one that strips away the humanity that's necessary for justice. For fair and compassionate decision making in that realm.

Lastly, I think there's the peril of over reliance on AI. We depend on it too much and some people see it as a replacement for human judgment and expertise rather than an enhancement thereof, that is supposed to be a tool to make us better. When AI is dependent upon too heavily, the essential skills that make us human: critical thinking, problem solving, adaptability, that kind of thing, that begins to atrophy. And this over reliance creates a fragile system where the failure of AI can lead to catastrophic consequences. Like the ones that we saw in the tragic example of the Space Shuttle Challenger back in 1986, where a single misjudgment based on a misunderstood tool, the O Ring, led to disaster.

So yeah, sure, AI's potential for good? Immense! But it demands respect and understanding, and to avoid all these pitfalls that I've talked about, we need to remember that AI is only a tool and not a master. And it's our responsibility, we, us humans, it's our

responsibility to wield that tool wisely, enhancing our capabilities without surrendering our humanity.

What foundational groundwork is needed/needs to be double-checked with regards to data quality and governance in order to leverage AI effectively?

To leverage AI effectively, I think that foundational groundwork of data quality and data governance is absolutely crucial. I would say that the first step would be ensuring that the data being fed into the AI systems is both accurate and comprehensive. Poor quality data leads to poor quality insights. We already know that and when decisions are made based on this incomplete or incorrect information, the consequences can be severe, and this isn't just related to having a large data set. No, I'm talking about having the right data: Data that's representative, relevant, up to date, before any AI implementation. It's essential to thoroughly clean and validate the data, identifying and eliminating any errors or any biases or redundancies, this process ensures that the AI system is working with the most reliable information available, leading to outcomes that are both trustworthy and effective. Another critical aspect I see is establishing strong data governance frameworks. These frameworks dictate how data is collected and stored, processed and shared out within an organization. Effective governance ensures that data is handled in a way that's both ethical and compliant with any regulatory standards. This would include things like setting clear policies on data privacy, data security, access control. As well as creating protocols for auditing and monitoring the whole data cycle. Good data governance also involves making sure that all stakeholders from the data engineers all the way to end users and everybody in between, that they all understand and stick to these policies. When data governance is strong, it not only protects the organization from ethical and legal risks, but it also enhances the integrity and reliability of the AI systems that rely on this data. Finally, creating a culture of responsibility around data management, that's essential too. And that means that everybody involved, everybody in the process from data scientists to decision makers need to take responsibility for the quality and ethical use of the data. It's not enough just to set up the systems and the frameworks I was talking about. There also needs to be an ongoing commitment to education and diligence, regularly revisiting and revitalizing and revising the strategies, especially as new data sources and technologies emerge. That is the key to effective maintenance of AI.

So by prioritizing data quality and governance, organizations can ensure that AI becomes this powerful tool we're talking about. For both innovation and decision making rather than, say, a source of risk and uncertainty.

What kind of framework is needed to critically evaluate AI's impact on creativity, objectivity, liability, and dependability?

I believe that you have to have some robust framework, a multi-dimensional framework to accomplish this goal. It needs to begin with purpose-driven assessment criteria.

For creativity: the focus is going to be on whether AI is enabling innovation or merely producing repetitive content. And this involves analyzing not only the output, but also the process from one step to the other, whether AI is being used as a catalyst to stimulate new ideas. and break the creative block that we sometimes have, or if it's only dampening originality by enforcing some stiff norm. The framework needs to include both qualitative and quantitative metrics, and this will gauge both the diversity, and the uniqueness of the ideas generated with AI's assistance, not them taking over. It also needs to assess whether AI is empowering creative professionals to explore new frontiers, or if it's simply reducing them to just mere operators of a pre-programmed tool.

Objectivity: I think the framework needs to include a rigorous examination of the data sources along with the algorithms that AI relies on. This involves scrutinizing the data for those biases we talked about and ensuring that the algorithms are designed to minimize them rather than perpetuate them. The evaluation process needs to include regular audits of the decisions that AI makes, especially in sensitive areas like hiring and finance impartiality is critical. A key aspect of the framework needs to be measuring how well AI enhances objectivity by filtering out the human biases versus how often it falls prey to the very biases that are embedded within the data it processes. The goal is to create a level playing field where AI decisions are transparent, explainable and free from human prejudices.

When it comes to liability and dependability the framework needs to integrate shared accountability mechanisms and reality checks too.

Liability: It's essential to establish clear guidelines. Who is responsible? Whenever AI systems might fail or produce harmful outcomes, the hallucinations people talk about, this includes not just the developers and the operators of the systems, but also the organizations that deploy them, everyone's responsible. The framework needs to evaluate whether these whole systems have built-in safeguards and whether there's a clear path to recourse when things go wrong.

Dependability: It needs to include stress tests, maybe some scenario planning, to assess how AI systems perform under unexpected conditions. Regular monitoring, maintenance protocols, that's crucial too. And that ensures that AI remains a reliable partner rather than a fragile crutch that falters in critical moments. So by implementing this comprehensive framework that I've outlined, organizations can harness the power of AI while at the same time, safeguarding against its potential pitfalls, ensuring that those four things, creativity, objectivity, liability, and dependability are upheld at every step. and revising the strategies, especially as new data sources and technologies emerge. That is the key to effective maintenance of AI.

How can companies address potential biases in AI algorithms?

Companies first need to acknowledge that bias is not just a technical issue. It's also a profound human issue as well. The biases in AI are not in and of themselves. They're often nothing more than a mirror of the biases in the data that we ourselves feed into it, the decisions that we ourselves make during its development. Companies need to start by committing to a deep introspective audit of the data sources all the way back. **Going beyond surface level checks, but diving into the origins of the data**, the demographic that it represents, the historical context, all these things are factors that might skew it. Companies need to ask these hard questions:

Whose voices are missing?

What perspectives might have been historically misrepresented or not represented?

By actively seeking these things out and integrating diverse viewpoints, companies can begin to mitigate the risk of these biases and they do it at the root. But addressing bias does not stop just in the collection part of it, the design and the development process itself further upstream needs to be rethought too, to process and include inclusivity. Bringing diverse teams to the table, not just talking about it in terms of the technical skills. Diverse skills in terms of cultural backgrounds, even experiences, perspectives. What happens when algorithmic decisions are made by some homogeneous team, a team that's not diverse? Well, they are more likely to reflect a narrower worldview. But by assuring that those who design, build and test the AI system come from varied works of life, companies can inject a much broader spectrum of thought into that development process, and that's going to lead to algorithms that perhaps a more balanced and representative. This diverse input needs to be combined with rigorous testing environments, things that simulate real-world scenarios. Particularly those that involve marginalized or vulnerable groups to identify and correct biases before they start manifesting themselves out in the wild, so to speak.

Lastly, companies need to embrace accountability as a core principle. It's not enough just to aim for unbiased algorithms. There needs to be a true, transparent and ongoing commitment to identifying, reporting and correcting any biases that are seen, whenever they appear. Now, this means implementing some sort of continuous monitoring, creating channels for feedback from users who might experience or notice biases themselves in the output of the AI. Companies ought to be willing to take personal responsibility to own up to the shortcomings of their algorithms and to take swift action to correct them once they figure them out. And in the end, the question goes far beyond eliminating bias. It should also include whether companies are willing to challenge their own assumptions, listening to those people who are effective and evolve, right? The real mic drop moment comes, when companies realize that the fight against AI bias is not just a one-time fix, but rather a continuous evolving commitment to both fairness and equity. The moment they stop pretending that perfection is achievable and start focusing on accountability instead, accountability and adaptability, that's when true progress begins.

What are some actionable strategies to leverage AI responsibly and ethically?

I'd say that it starts with the **active integration of human oversight** that we were talking about earlier, and we're talking about, throughout the entire AI lifecycle. Let me break that down. That means that at every stage from development on the one side, all the way to deployment on the other, all the way through the spectrum there needs to be a human in the loop. Or more than one, humans in the loop, to make sure that AI output is not blindly trusted, but critically evaluated and not just an oversight that's micromanaging every intricate minutia of technology. No, I'm talking about maintaining balance between machine efficiency and human wisdom blended together.

For instance, when AI is used in a decision-making process, it should serve as a tool to augment human judgment, augment not replace it. I believe that organizations need to create protocols that require human review of any AI generated decision, especially in high stakes situations.

Healthcare diagnosis, legal judgments, financial approvals that we talked about. By embedding this distinct layer of human scrutiny, companies can prevent the unchecked spread of algorithmic errors. And well, again, those biases that we've been hitting on off and on.

Let's see another key strategy: **the establishment of ethical guidelines and accountability frameworks**. Responsible AI use requires clear and forcible standards that people know ahead of time. And it dictates how AI should be designed, deployed and monitored. These guidelines need to be rooted in ethical principles like fairness, transparency, respect for individual privacy, that kind of thing. And to be effective they need to be more than just abstract, ethereal ideas. They must be translated into concrete actions that people can follow. Regular ethics audits, transparent reporting on AI's impact, mechanism for redress in case AI systems cause harm when something goes wrong. Companies also need to create that culture where ethical considerations aren't just an afterthought, but it needs to be the central focus. It needs to be part of our default response and it can be achieved by setting up ethics committees and they'll work alongside AI development teams. This makes sure that whatever decision gets weighted against the potential ethical implications.

Lastly companies should embrace the concept of AI being an enhancer. The most ethical and responsible use of AI comes when it's viewed as a partner. A partner that complements human strengths and not some sort of substitute for human ingenuity. Our laziness gets us to do that. This approach encourages use for AI for tasks that it excels in like analyzing vast, ginormous data sets, automating routine processes. It's much better at that than we are, but leaving complex, nuanced decision making, the gray areas, in human hands. Framing AI as a lever to elevate human potential rather than a crutch that diminishes it when we over depend on it. Companies can harvest AI's power without losing the very qualities that make us humans: Creativity, empathy, critical thinking. The things that AI cannot possibly do, at least not yet, and not for a long time.

Responsible AI use, it's so much more than just preventing harm, it's ensuring that technology serves to enhance our humanity rather than overshadow it. And the future of AI? It's more than about what it can do, but how we choose to wield its immense power.

How do you feel about the latest developments with AI training AI?

AI training AI is like the blind leading the blind. I can see how if AI already has the algorithms that it perpetuates upon itself and it can be a force multiplier to achieve this training so much more quickly. But. I still think there needs to be a human in the loop to oversee it.

I heard and I don't know if it's true, I didn't get a chance to verify that it's true that Meta AI or whatever the name of the engine is, was working on developing some new strategy, came up with a language that humans couldn't even understand. And they shut that thing down. Because you don't want to give the keys of the kingdom to some ethereal entity that you don't know what the outcome is going to be. It's like letting a password generator generate a password that you don't know and generate a backdoor that you don't know about either. I can see the benefit in involving the actual AI engine in training another AI engine, but I don't see it happening without human intervention. I think we're asking for trouble when we let that happen without human intervention.

What is the role of individuals and organizations in shaping the future of AI and how do skillset, mindset and company culture needs to shift to fulfill it?

I'd say that it starts with understanding that AI is not just a tool, but a transformative force that reflects the values and intentions of the people who wield it. To truly fulfill this role, individuals and organizations both need to shift both their skill set and their mindset. This is necessary to keep pace with the AIs evolving capabilities. It's evolving even faster than we think. At the executive level for instance, it means embracing continuous learning. Everybody: executives, frontline workers, everybody in between all need to become fluent in the language of AI. And I don't mean the technical proficiency necessarily. Not everybody can be a programmer. That's not necessarily what's needed. It's about developing a critical understanding of AIs potential and its limitations at the same time. Individuals need to cultivate a mindset that sees AI not as an adversary, not as a replacement, but as an enhancer. A partner that amplifies its abilities while still requiring their own unique judgment and creativity. The individual. I've said it several times in different ways, but I don't think it can be overemphasized.

Now on an organizational level, organizations, company culture needs to undergo a significant transformation. The traditional top-down approach to decision making needs to be tempered. Yes, I believe that things rise and fall on leadership, but at the same time, it needs to give rise to more of a collaborative cross-disciplinary environment. Ethical considerations and diverse perspectives get embedded into every AI project. Organizations need to create spaces where technical experts, maybe ethicists and end users too, can collaborate in the same room to ensure that AI systems get designed with the holistic understanding of their impact. This is a cultural shift and it's going to involve creating an atmosphere of transparency and accountability where mistakes are openly acknowledged and there's a commitment to continuous improvement, not just punishing people when they mess up. The lesson from the space shuttle disaster I mentioned earlier. It was not just a technical failure; it was a failure of communication and a failure of culture. I'm talking about the engineers who warned about those risks were ignored. Why? Because the pressure to launch outweighed the caution to wait. I'm not trying to blame anybody. I'm just saying, there was no AI back then, so this has nothing to do with AI, but we are talking about the same attitude, the same mindset. Today, AI stakes are high. Organizations need to cultivate that culture where caution, ethics, and diverse voices are valued just as much as the speed and innovation that you get from it.

So to fulfill this, both individuals and organizations, both of them, need to adopt a mindset of shared responsibility. I touched on it before. AI's future depends on it. It isn't just solely in the hands of the technologist or the technology or the developer. No, it's in the hands of every single person who interacts with it and is affected by it. And this type of thinking, this type of mindset requires organizations to empower their people with not just the technical skills that they need to work alongside AI, but also the ethical literacy to navigate its complexities responsibly.

Like the Wright brothers, you remember these guys, right? They didn't just build a plane, they mastered the thing. Every single detail, every aspect of flight, they mastered it. Relentless learning, adaptation, and a deep understanding of all the nuances involved. With the same mindset as the Wright brothers, you and me today, more than a century later, we need to approach AI with that same level of dedication, the mindset, the skillset. On the skillset, the skillset embraces AI. On the mindset, the mindset that critically evaluates its role. And then the culture prioritizes ethics and collaboration. All those things are going to work together to ensure that AI becomes a tool for progress rather than a force for division and conflict.

You know, the future of AI is not predetermined. I think it's shaped by the choices we make today.

You mentioned the importance of education. What is your take on the responsibility of companies to offer AI training, especially while they might still be exploring and just in the piloting stage.

I think it comes down to the mission, vision and goals of the organization. What are your policies as an organization to further the mission and vision and goals that you have set for yourselves as an organization? Are you utilizing the technology? Are you just exploring it? Are you in the process of implementing some sort of strategy? Implementing a strategy without educating people on how to use it is like putting gas in a car without an engine under the hood, or like giving a 16 year old kid a driver's license in the US without showing him or her, how to park the car, how to merge on the traffic, how hard you press the brake pedal before you either put somebody's head in the dashboard or how not to slam on the accelerator and make a mess. It is an individual's responsibility to be ethical and be careful and be cognizant and mindful of what's going on. But I think it's also an organization's responsibility if they're going to adopt a new technology, you cannot adopt unless your people are adept.

If they're going to adopt the technology, then the individuals that are tasked with that implementation need to understand the nuances, need to understand how to use that technology, how to leverage the technology, how to keep themselves from using it incorrectly. Organizations are only hurting themselves if they don't show people how to use things correctly, because invariably they're going to start using it incorrectly. Or else they'll finally figure it out on their own, but they'll do so in a very inefficient manner and that wastes even more time. Time and resources and energy and money and so forth.

Getting back to the mission, vision and goals, if they're not attuned to developing the strategy, then it doesn't make any sense to have that huge expenditure, if they're only in the exploratory phase and they have not yet figured out how to do it, what direction they're going or the manner in which they're going to implement it, that would be a little bit premature. Yes, I'm all for education, but if I own a landscaping company, I'm not going to show my employees how to use a jackhammer, because they're not going to use a jackhammer in the course of their job. Right tool, right job, right place, right environment. There's something to be said about jumping the gun and wait until you've got your strategy figured out. Once you have your path outlined, once you have whatever it is that you're going to do, whatever direction that you're taking, aligned with this mission, vision and goals, that would be the time to start bringing people in.

And then you bring them in from the get-go. Involve representatives from different levels of your organization in figuring out: this is the idea that we have. This is the direction we want to go. I want to hear your ideas on your individual perspective on ways that we can implement this in the most strategic and efficient way possible. And so once you get to that point, jump in, move forward, but before you get to that point take care, no reason to dive before you figure out how to swim.

The full video recording of the interview is available **here**