

AI and Data Ethics: Closing the gap between priority and practice

Community Roundtable: June 2024

**Data
Leadership
Collaborative**

A circular logo with the text "SPONSORED BY TABLEAU" around the perimeter.

Table of Contents

Introduction and Panelists	3–4
01. The Murky Ethics of GenAI	5–6
02. Toward a More Ethical AI Model	7–8
03. Legislating Ethical AI	9–10
04. Ethical AI in the Development Lifecycle	11–12
05. Data Leaders and the Ethical AI Conversation	13–15



Data Ethics in the GenAI Era

Nearly every organization says that the ethical use of data is a priority, but is data ethics a practical possibility in the GenAI era? LLMs already have been trained on copyrighted data, and removing this data from model training leaves a negligible amount of public domain information. At the same time, content creators and artists deserve recognition for their efforts, and everyone should have access to data interactions free from bias.

DLC gathered four top data leaders to discuss issues like these: how we got where we are today with GenAI, where and when ethics should be built into AI models, and prioritizing the value of transparency in digital ecosystems.

Panelists



Caroline Carruthers
Chief Executive,
Carruthers and Jackson



Adam Mico
Principal, Data and Analytics
Strategist, Moderna



Maya Mikhailov
Founder and CEO,
Savvi AI



Kevin O'Callaghan
Head of Data Analytics,
Teamwork.com

The Murky Ethics of GenAI

Data Leadership Collaborative (DLC): Let's kick this off with the ethics of GenAI. Why do they seem so hard to pin down?

Adam Mico: I see several factors. First, these technologies often blur the lines between original creation and intellectual property, which leads to complex copyright and ownership issues. Many LLMs rely on massive datasets that very likely include personal or proprietary information, which raises significant concerns about data privacy and consent. The potential for AI to perpetuate biases present in training data also adds to the ethical complexity, as does the impact of AI on employment and the broader economy.

Furthermore you've got issues of transparency and accountability that complicate determinations of responsibility for AI-generated outcomes, while the potential for misuse heightens safety and security concerns.

Kevin O'Callaghan: I would sum it up in one word Adam used just now: transparency. Transparency is where the ethical risk lies, but that's also where the win is. Your colleagues and customers want to know: How does this AI work? What are the risks and challenges? How does it play into my job? What does it mean to me? Imagine being in a hospital and the doctor or the nurse comes in and says, "I'm sorry, we're stopping your treatment because the algorithm said that there's not much hope for you." The first thing you're going to say is, "I'd like to talk to a person and understand what the meaning is behind that decision." You rightly demand transparency.

We always should be clear what's going into our AI and what's making the rules and how that affects it. Even if it's about as interesting to you as reading the terms and conditions of a software contract, you should have the ability to be aware of what that transparency means, what that algorithm could suggest, what the outcomes may be. No one likes being a victim of a decision that doesn't go their way or is not explained to them in a way that makes sense to them.

01.

01. The Murky Ethics of GenAI (cont'd)

Maya Mikhailov: For me AI ethics starts at the root of how the AI was trained and what we consider to be acceptable and fair usage of training materials. If you start with the premise that everything's fair game and no one should be compensated for work that makes your model better, is that ethical? I heard a leader in GenAI interviewed recently and he said, If we couldn't use these materials, we wouldn't really have anything.

I was shocked at the candor of that reply because we have lived through this before. Napster's business model only worked if music was freely distributed without compensating the artist. Here is a GenAI leader saying that their business model only works without compensating or respecting the creators. We all know how Napster ended, and this can be equally as problematic.

Caroline Carruthers: I think this could be a little bit of a circular argument, because even before we really started having these conversations we used to talk about similar ethical issues for writers. There are only so many stories to tell, so virtually everything written as a story could be seen as an iteration of something that had come before. If that's true, AI is just the machines doing the iteration rather than us in this particular process. Every one of us has probably drawn inspiration from a story we've heard, or an artist will have drawn inspiration from something they've seen already. So is the AI just doing what we have been doing but doing it with a lot more inputs? And then why would that be ethically different from us doing it?

DLC: Some people argue that the horse already is out of the barn on AI ethics. Would you all agree?

Caroline Carruthers: I understand the point that people make about the horse leaving the barn already, but that only works as long as you believe that ethics is a fixed static situation. And I don't believe it is. I can give many examples on how different societies change at different rates. As a society our ethics have changed because ethics are no more than a reflection of what we deem acceptable at a particular time. So if you buy into that view, then ethics isn't a line that you cross or an "i" that you dot. It's a continuous process that you have to self-check yourself against. And my argument would be that it's never too late to start on that journey.

I'm also a big advocate for the human always being in the loop somewhere. So even if it's just checking the consequences and the outcome of the actions that we're taking and having the ability to go back if we need to, I think that the moment we've given up, we've lost.

Every one of us has probably drawn inspiration from a story we've heard... So is the AI just doing what we have been doing but doing it with a lot more inputs? And then why would that be ethically different from us doing it?

– Caroline Carruthers
Chief Executive, Carruthers and Jackson



Toward a More Ethical AI Model

DLC: What are some of the considerations that come to your mind around ethical AI models. And is one even possible?

Adam Mico: While a completely unbiased or purely ethical AI may be an ideal rather than a practical thing, we can make significant progress towards reducing biases and enhancing ethical considerations. This involves a multi-faceted approach that includes diversifying data sources to better reflect a variety of perspectives and actively identifying and correcting biases in AI systems.

Continuous monitoring and updating of AI models as societal norms evolve are also crucial. Engaging diverse groups in the development and review of AI systems can help ensure that multiple ethical viewpoints are considered, leading to more responsible and equitable AI solutions. So while perfection may be unattainable, I believe that striving for improvement in ethics and fairness in AI is both necessary and possible.

Caroline Carruthers: I absolutely think an ethical AI model is possible. There's certain things that I can see GenAI in different models being really useful for and other things I wouldn't touch with a barge pole. The simple example I would give is anything that I have to prove my decision-making ability on. So if you are feeding into a regulator, for instance, I would not let GenAI touch that because I need to be able to demonstrate every single step of the process.

A consideration for me is that the tool we now have access to is incredibly dynamic, very powerful, and changing at such a high pace. So we can do a lot with it. We also have a lot of people engaging with GenAI that maybe aren't the traditional people that would've been engaging around data and things like that before as well. So it's almost like we've got a new audience as well to deal with and we shouldn't forget that.

02

02. Toward a More Ethical AI Model (cont'd)

Maya Mikhailov: Not to mention that training these LLMs is so darn expensive. Billions of dollars have been and will continue to be spent on the race for the best model. That means that those that can afford to train will be the voices that control. Again, we are back at the search engine wars where the gatekeepers of information were a handful of companies, like AOL or Yahoo, who controlled access to information because they had enough money. This spawned many ethical debates about walled gardens and who should be the arbiter of content and narrative.

There are other questions, too, like Is AI going to exacerbate isolation and loneliness? Will it drive a deglobalization if nation or large enough affinity group, that can afford it. If these players start saying, You know what, your LLM does not speak for my culture, my people, my values and I need my own lens on the world for my people.

It concerns me that just as the internet brought us together, GenAI has the possibility of ripping us apart where we're in our own little AI bubbles being fed information that is hyper-tuned to what someone wants us to think and feel. It bears unpacking on so many levels besides the fact that for all intents and purposes it is, in fact in many ways still a black box.

Kevin O'Callaghan: I would be of the hopeful type here, though we might be looking at a lot more safeguards in the future. Equally, the risks of an AI that's exclusive to a region or country and one that gives you a certain view of the world or of a position or a politic or anything is always going to be a bit of a worry. And then how do you factor that into decisions?

DLC: Should different AI models be used for high-risk versus low-risk decisions?

Adam Mico: I would advise that as a way to appropriately manage potential consequences.

If you think about a high-risk decision — something in healthcare, law enforcement, or financial services — there's a significant potential impact to human lives. That would require AI models with enhanced safety, reliability, and transparency. These models should undergo rigorous testing, validation, and adhere to strict regulatory standards to ensure they meet the highest ethical and operational benchmarks and only be used as an additional assistant but not to actually make decisions on solely.

Then conversely you've got low-risk applications, like content generation for entertainment or automated customer support. These AI applications can utilize less stringent models, allowing for greater innovation and faster development. However, ethical development needs to be considered at every step. This approach not only mitigates risks but also optimizes resource allocation, ensuring that the most critical AI applications are the most robust and secure, thereby balancing innovation with responsibility.

Legislating Ethical AI

DLC: The European Parliament has taken a big step forward with the first draft of its [EU AI Act](#). How are people reacting to it from where you sit?

Caroline Carruthers: I haven't seen a lot of impact coming out of it at the moment. Before it came out I saw quite a bit of stagnation going around companies. It was almost like they were waiting to find out what was going to happen. I think the devil is in the detail with a lot of this and we haven't really seen the impact yet, nor will we for a while. The first few test cases will be really interesting to see where we're going to be focusing.

Kevin O'Callaghan: I think if you're in the industry, you're watching it and you're paying attention to what it means, and what are these large models, and how does this impact the decisions? I think a lot of folks still are not quite tuned into it. If you are in this industry or business or you're working in legal, you're on it. I think you're keeping an eye because it's no different to GDPR being implemented and the risks and fines there. I assume positive intent that this AI act is going to be the same, but for most folks, they're not quite up to speed on it yet. I think once we see the act being ratified and really being brought into place and companies and organizations start adhering to it, that's when the rubber will really hit the road in terms of European focus.

03.

Engaging diverse groups in the development and review of AI systems can help ensure that multiple ethical viewpoints are considered, leading to more responsible and equitable AI solutions.

– Adam Mico
Principal, Data and Analytics Strategist, Moderna

DLC: And how about America's attempt at AI legislation, otherwise known as the [American Privacy Rights Act](#)?

Adam Mico: APRA represents a significant step toward establishing national privacy standards in the U.S. because it addresses a longstanding gap filled by state-specific regulations. Whether these efforts are too little, too late hinges on the law's enforcement and adaptability to rapid technological advances, such as those in AI and big data. It could be argued that these measures are overdue given the frequency and scale of data breaches, but the APRA is still a vital foundation for safeguarding consumer privacy. I like that it sets a crucial baseline for privacy, providing a framework that can be adapted and built upon as digital technologies evolve. So I see it as a necessary action at a critical time rather than a delayed response.

But since all these technologies are evolving so quickly, I see APRA as more of a guideline. Just think of how many resources and how much knowledge it would take to investigate and enforce it. Ultimately, much of the application of APRA would require private self-interested parties to develop AI being mindful not to violate its guidelines.

Ethical AI in the Development Lifecycle

DLC: To the extent that we can build ethics into an AI model, where should that happen?

Caroline Carruthers: It's a bit like the question we always ask ourselves: When is the best time to do something? The best time is ten years ago, the next second best time is now. So if we haven't already put ethical considerations into our AI, we need to start immediately. It should be built in right at the very beginning. We should be constantly thinking about what we are prepared to accept and what we're not. I tend to tag ethics with the governance side of what we do because we're trying to put guardrails in. Those guardrails should be just far enough from the road that we can't drive off the edge of the cliff but not so far in that we can't still drive and use the road. At the moment there's still an awful lot of nervousness around ethics in AI models because it's new. Have we got it right? Are we doing the right thing? You can almost see on a daily basis the guardrails being picked up and moved subtly around to see if we found the right place for them. So I think there's still a lot of work to be done in this area.

Adam Mico: Agreed, ethics should be integrated into every stage of the AI development lifecycle to ensure responsible creation and application. This begins with planning and design, where ethical risk assessments identify potential impacts on privacy, fairness, and human rights. That continues with data collection and preparation, emphasizing the responsible acquisition and use of data to avoid biases. During the model development phase, ethical principles should guide algorithm selection and training to ensure transparency and fairness.

Testing and validation also should rigorously evaluate AI systems for ethical integrity. This includes assessing them against diverse scenarios to ensure they meet set ethical guidelines. Once deployed, continuous monitoring of an AI is necessary to maintain ethical compliance, with real-time checks for



04. Ethical AI in the Development Lifecycle (cont'd)

discriminatory outcomes or other ethical breaches. Lastly, post-deployment involves regular audits and adjustments based on feedback, ensuring the AI system remains compliant with evolving ethical standards and societal expectations. That way you're fostering trust and sustainability in AI technologies.

Kevin O'Callaghan: I'll chime in as well and say, yes, from the get-go. If you were to do the right thing in developing an AI, you would ensure that ethics and fairness are in there right from the start. You'd know exactly what you're getting yourself in for, what is fair and what isn't. Checks and balances along the way are actually good in terms of making sure that that everything has been approved by a human, not an AI.

You also need to ensure that the experience and outcome of your model matches the experience that you wish to have. I think that is going to become part of a quality process that's no different to quality processes we've seen in software development, in applications, and so forth. You need someone in there to actually say, Hang on a second, what about X, or is Y in the right place? And that might be a new role or it might be an existing role. But I think having someone to ask the questions and asking the awkward questions that align to the values and own ethics of an organization or individual should be embedded throughout.

Maya Mikhailov: And I think the more voices at the table, especially if they're voices of non-developers — stakeholders who the AI is effecting are critical. We've read what happened when AI is deployed as a crime prevention technique. It's been applied to communities where the developers had no knowledge of the community dynamics. Nor did they incorporate feedback into the model if the suggested policies were detrimental or simply failed to lead to the outcomes they were looking for. Developers and Data Scientists should not be forced to be ethicists or civic leaders. It's an unfair burden unless it's one they are entering into willingly. Not all of us go to work every day with someone's life, home, freedom in our hands. That's a very big flag to bear. Instead, community voices should have a seat at the table in these higher risk modeling situations. You could have the very best of intentions and still have an adverse outcome.



Data Leaders and the Ethical AI Conversation

DLC: Clearly there are a lot of challenges around AI ethics. What are some of the things that data leaders can do here?

Adam Mico: Data leaders can take several proactive steps. Establishing clear guidelines and standards for AI development and usage is crucial, as is promoting transparency through the adoption of explainable AI techniques. Engaging in open dialogues among technologists, ethicists, policymakers, and the public can help forge consensus on ethical AI practices. Additionally, implementing regular auditing and reporting mechanisms can ensure compliance with ethical standards and assess the societal impact of AI technologies.

Finally, investing in education to increase awareness of AI capabilities and limitations, as well as advocating for diverse and inclusive development teams, can mitigate biases and lead to more equitable and responsible AI solutions.

Kevin O'Callaghan: I think it's staying true to what we see as fairness and openness, and I'm going to say transparency again because it really is that important for me in terms of what people are going to get out of AI. Data leaders should focus in terms of explaining and understanding the challenges and the problems that they are trying to solve, while also explaining that back out to an organization; e.g., "This is what our AIs are looking to solve, these are the outcomes they're looking to drive, and this is how we come to those decisions."

I don't see this any different to when you're making a decision on purchasing something. If you're going to purchase a software product, you're probably going to test it, trial it, make sure it works how you want it to work, and that there's no hidden surprises. I think a similar method that we'll probably use with AI as well in terms of asking ourselves, What are we involving here? Awareness, clarity, visibility and explaining that to

05.

05. Data Leaders and the Ethical AI Conversation (cont'd)

others about these risks and even just the way they're architected at the moment is so important for data leaders. And it's not just the LLMs, it's the day-to-day stuff. It's your machine learning. It's what constitutes that model, what brings it in. And I think that's where a mature discussion needs to be had, really doing your diligence on what is it you're implementing here or what are you buying here or what are you integrating into your tech stack or in your processes? And does it pass the smell test, so to speak?

Maya Mikhailov: Honestly, I just have more questions. And I think that from a private enterprise point of view, the key is to go back to basics of your rules, your guidelines, and start using that as a framework because AI doesn't supersede it. It needs to behave like it. So if you have guidelines and policies in place and safety procedures and ways of ethically doing business, this is what you have to use as the guardrails for any AI. In a weird way, you don't have to reinvent the wheel. As Caroline said, you just have to add guardrails. And just as you wouldn't want a customer service representative to behave negatively towards your end consumers, you wouldn't want your chatbots to behave that way either.

Kevin O'Callaghan: I think that's a fantastic point. And if businesses show that they care, be it a chatbot or a person on the phone, that they really care and AI isn't just something put in to save money or to be cheaper, customers will stay true. But even internally, when they're looking at what they're integrating, building or buying, I think particularly thanks to privacy and all of the checks and balances that you need to do there, you're almost applying that diligence also to these systems of what's the give and take and how much is the business willing to take and how much is it willing to give in order to get certain functionality out of it? Because I think a lot of businesses today feel under pressure because they're seeing AI everywhere. Smaller companies and medium enterprises and startups and all are kind of going, Oh, I have to do something with AI because everyone's doing AI. But at the same time it's looking at the AI arms race and saying, Wait, how does this stay true to what we want to do? How does it get us to where we want to go?

Caroline Carruthers: I think every data professional out there has an obligation or responsibility to get our heads together on this one. So I think this problem should surpass individual company loyalties. This is a bigger issue that we need to deal with. So it's a general call to arms. But I think it'll take a little bit of groundswell to get it because it's not going to be an easy thing. I think it is one of those things we need to start with.

DLC: These are great insights. In closing, do you have other thoughts about how an ethical approach to AI becomes an opportunity and not a burden?

Kevin O'Callaghan: I think trust in terms of doing the right thing, but also Does it match the values that I want to have because not all AIs are going to match my values and not all companies are going to match my values? That's the luxury of choice that I have. So ask yourself as a data leader, Does this organization, does this algorithm match my values?

Maya Mikhailov: The flip side of this is, Are these values we want to convey or is this how we want to behave? Is this how we want to be known to behave? Simply having a reactive ethical Whac-A-Mole strategy is really stressful for your compliance, legal and comms teams. But knowing and living your values toward your team and your customers and just architect backwards from there. And if particular technology fits, great; if it doesn't, that is okay also.

Caroline Carruthers: The starting point is to understand what is the outcome that you're trying to drive to. So I talk about purpose a lot, but if you really understand the purpose of what you're trying to get to, then it should be relatively straightforward work back from that. And when you talk about such a big question, it's more about the process and the framework that you're putting in place rather than the specific examples at that stage, isn't it? But if you understand the purpose, you understand the people involved with ensuring that purpose is what you want it to be and checking that it takes place and then having a process to backtrack if you need to. That's the sort of building blocks to put in place.

Adam Mico: Another opportunity for data leaders is that as AI legislation evolves, they can proactively engage by educating themselves and their teams about new laws, actively participating in policymaking, and implementing ethical best practices ahead of regulation. They can influence legislation by providing expert insights and perspectives through direct dialogue, public consultations, or industry groups. So it's a great opportunity for a data leader to build their brand.

Internally, fostering an organizational culture that prioritizes ethical AI development is crucial, involving setting guidelines that exceed legal requirements, conducting fairness audits, ensuring privacy protection, and establishing clear accountability for AI outcomes.

You also need to ensure that the experience and outcome of your model matches the experience that you wish to offer. I think that is going to become part of a quality process that's no different to quality processes we've seen in software development.

– Kevin O'Callaghan
Head of Data Analytics, Teamwork.com

Additionally, data leaders should facilitate open dialogues within their organizations and with external stakeholders, including regulators and the public, to collaboratively address the challenges and impacts of AI technologies. By taking these steps, data leaders not only ensure compliance but also position their organizations as ethical leaders in the AI space, driving responsible innovation.

Maya Mikhailov: I guess my big takeaway is that we're all going to muddle through the ethical implications of this new technology for a while. And what companies have to just reconcile with is what sort of legal or reputational harm are they willing to endure to muddle through it? And if the answer is very little, then they need guardrails, transparency, and a firm framework in place before you just willy-nilly implement the technology in production. And if the answer is, I'm a small start up and no one cares about what I do over here, then maybe you have less protections and guardrails. But for large institutional players and especially players that are in regulated, risk-averse lines of industry, they're going to be very cautious. And instead of calling them laggards, let's understand that their rules are different.

DLC: Well said. Thank you all for a great conversation.

What's your opinion on AI and Data Ethics? We want to hear from you!

Follow the DLC  

[Data Leadership Collaborative](#) 

**Data
Leadership
Collaborative** 