Caitlin McDonald:

All right. Welcome, everyone. For today's podcast I'm joined by Tim Gordon, who is the founding partner at Best Practice in AI, a research firm that helps organizations use AI to build competitive advantage. Tim has over 20 years international leadership experience in growing businesses and digital transformation in the media, financial services, consumer, not-for-profit campaigning, local search, and lead generation industries. Welcome, Tim.

Tim Gordon:

Great to be here. Thanks for having me.

Caitlin McDonald:

We're really lucky to have you today. I wanted to start. Given your work with ethics oversight for the board and the C-suite. And I wanted to start with some questions around the evolution of AI ethics from an emergence, but unformed issue on the senior leaders' minds into something that's growing into a more robust governance process. And why is it so important for senior leaders to get to grips with this subject?

Tim Gordon:

I think if you look at some of the headlines we're seeing in the press at the moment, so in the UK we've recently had a bit of a crisis with our AI algorithm marking. Now, that's not AI specifically, but it probably begins to indicate some of the challenges which senior leadership are going to face when these types of automated decisions start being made. And coming down the track, there's a series of regulatory shifts that are happening. So, for example, GDPR has already brought in the first signs of regulatory control on this issue. And we're going to see this increase and exploding as digital first business models expand. This is going to be an increased important topic by the boards.

Caitlin McDonald:

Excellent. And are there some things you want to talk about in terms of the ways? You've already mentioned GDPR. Are there other nascent, emergent governance models or frameworks that you see on the horizon that are coming towards that boards are going to have to be concerned?

Tim Gordon:

I think what's interesting is if you take a step back, what AI ethics actually are. Effectively, we have three big buckets of challenge. The first is essentially the technology isn't working very well. So, if you think about some of the stuff you read about in the press where, for example, image recognition systems come up with terrible classifications, particularly people of color. Or where you see automated cars crashing to people when they shouldn't. What you're seeing there is basically the technology is not yet really good to go in terms of being able to play at the full production level. So, that's one of the big areas of control and controversy around so-called AI ethics. That's interesting, but to some extent, it's not really ethical questions. It's kind of clear what the answer is. The answer is don't deploy until it's ready.

Tim Gordon:

The second big bucket I think, which is more interesting is where AI obviously relies on data and the data that it's processing is historical data. So, where you have tools being used for recruitment. And essentially it turns out that it's massively biasing towards white or male programmers, and basically

downgrading people of color or women or people with other characteristics. What you're seeing there is this is based on 10 years worth of training data, of 10 years worth of historical decisions that have been taken by the firm. I think what's interesting about this is basically putting a mirror to our current management processes, our management approach, our management decisions, and then say, "Look, mathematically, there's clearly something going on here." Now, you might historically been able to dismiss some decisions as being bad apples, or individual anecdotal stories. It's taking those anecdotes. There's pretty numbers around those. And it's saying, "Look, fundamentally we have some challenges." So, we're seeing that recruitment has been sexist. That legal decisions in the States have occasionally been racist, who knew? And these are questions which are really the AI is bringing to the fore.

Tim Gordon:

The third, I think most interesting set of questions around AI ethics. And these are real ethical questions are the sort of questions that are coming out to the new business models that AI enables. So, where you have a system where you could have massive personalization, you end up with that on one level a massive consumer win. But the downside is you probably end up with echo chambers in society through social media. Where you have questions around, for example, increased use of surveillance to generate all sorts of potential societal goods. You also have the potential downsides from that. And so, AI is effectively enabling business models that raise quite deep and quite broad ethical questions. Some of which are new, some which basically is all the society.

Caitlin McDonald:

Yeah. It's a really great point actually. And I think actually from my perspective, that's one of the great opportunities that AI offers is to question some of these existing biases that were already in play. That procedurally were already in play. And what we're seeing now is that as these decisions are becoming increasingly automated, we actually have to go and question not the automation part, but the decisions that were being made originally.

Tim Gordon:

The better the AI gets, the closer we're getting to really deep questions about our society and how actually we make decisions in our country. So, go back to the AI algorithm. That fundamentally was a question around what you prioritize. Are you prioritizing minimizing grade inflation? Are you prioritizing better social outcomes? Are you... What are you trying to prioritize in terms of a mix of questions, and AI with that ability to really nail things down makes these explicit decisions often as opposed to often implicit decisions.

Tim Gordon:

But I think the thing that's interesting with AI and we talked about a couple of minutes earlier is the way in which actually how you turtle some sort of quite notional force. So things that really matter at the board level. And so, we stopped counting when we passed 120, 130 international frameworks around how you should think about ethical AI. And those come from everyone, from governments, to pressure groups, to charities, to people at the World Economic Forum.

Tim Gordon:

You look at those actually broadly, there are six things that emerge in all of those which are, it's got to be inclusive and fair. It's got to explainable and transparent. It's got to basically operate safely. It's got to deliver some sort of social benefit, which obviously is debatable what that is, but that's a key

component to most of these frameworks. It's got to be responsible in terms of how it's been managed, and people got to be held accountable for this, which is where it very quickly comes back to the board because at the end of the day, who's fundamentally responsible for this? At the end of the day it's the board.

Tim Gordon:

When you go through what the board's responsibilities are, AI ethics is a big national topic. It comes down to four things. Legal questions, are you getting it sued? Regulatory, are you going to be in trouble with your regulators? Operational, is something going to go wrong? Are you going to make bad judgments and bad decisions? And reputational, which is you made a bad call. How do you justify that when the public coming knocking?

Caitlin McDonald:

Yeah, that's a great point. And that takes me nicely onto the Empowering AI leadership toolkit that you helped with the World Economic Forum to develop. And I noticed that the model that explains all the different modules for the toolkit has ethics positioned at the very heart of everything.

Tim Gordon:

Yeah. That I think was a... And there was a fair amount of debate about that, but I think that was very much trying to articulate to people from a board perspective how people have to be thinking about this. Not purely as a way to make more money in the short term, but also how you think about the longterm impact of what you're doing. And I think the thing we tried to do with that was ethics was at the heart of it. But we had chapters of each of the various elements. Whether it's the impact of your brand, and your customers on sales, on operations, and so on, but really try to just keep reminding people, you need to keep those sort of ethical questions at the heart of what you do.

Tim Gordon:

But the thing we tried to do, particularly building on since then is really bring out what the checklist, the board needs to think about that. So what are the questions that board members need to ask executive team members to really understand what's going on here in order to get to the heart of it. And what we come out very clearly is there are four main areas you need to start thinking about. Firstly, what's the governance process. How are you making sure that decisions are being covered off? Secondly, what are your operational management approaches? How are you making sure that the data you're getting is being done in the right way, you're running the right procurement processes, and so on.

Tim Gordon:

Third is I think just really articulating the human role in these decisions. There's a whole endless debate about, do you want humans in the loop, on the loop, over the loop? And also how the longterm impacts of that, but really making sure you've got the humans still involved, which by the way, is usually the best way to operationalize this stuff. And finally, stakeholder communications, which is how do you reach out to your stakeholders and communicate with them what's going on? And a lot of the conversation about AI is so driven by this image of the silver metal man reaching out to change the world, as opposed to seeing this as being fundamentally a spreadsheet with attitude. And I think that trying to get the stakeholders to understand what that means and why this isn't necessarily therefore all about taking a job or all about a whole series of other scary stories associated with AI is a huge part of the board responsibility.

Caitlin McDonald:

Yeah. And to that point about it not coming for your job. So there was a great article in the Guardian this week, which had purportedly been written by AI, a machine learning program that works with texts. And then of course, when you read the fine print, you discover there were eight or nine interventions that had been made by editors, and by authors to try and actually pull that together. So it's very much... It goes very way back before even AI there's a wonderful program in the 1960s called Eliza, which was essentially a fake... Not a fake, but a text generated therapist, which basically just repeated back to you what you said, and what was interestingly discovered by the researcher whose name I'm forgetting at the moment. But what was discovered was that people helped it out. So, there would be slight miscommunications or it would say something that was a slight non-sequitur, but then just as in real conversation a person would interject and make the conversation work. So, you'd work around it. So, there's actually a lot of human intervention when we're dealing with AI. We work with it. We're generous with it.

Tim Gordon:

I think you're entirely right. So, at Best Practice, www.bestpractice.ai, we put out our research into AI. So, we basically created, we've articulated 750 or so use cases for AI and 1,200 case studies of people using AI. The thing that really strikes me every time I go through AI case studies is how rare it is that you actually see job reductions coming through. So, we've tried to write these case studies very much in the way that a CFO would understand them rather than CTO, which is what are the business impact. What's the KPI impact of some of these roll-outs? And it's remarkable how rare it is that you actually find jobs disappearing. What you have, obviously, is jobs transforming. And I think we're just at the beginning of people understanding how you begin to shift the business model, and use AI to basically scale up things massively.

Tim Gordon:

So, chat bots are massive. They're great if you need to massively increase your call center capacity, but they're really rarely going to take away from those really knotty problems that actually really where the value added comes and people spend most of their time. And my favorite graph thinking about this historically is the graph of ATM expansion in the United States in the 1950s, '60s, and so on. From '60s, '70s, I don't remember when it was they came in. But what you see at the same time is the number of cashiers practically doing the same job dealing with customs, handing out money increases pretty much at the same rate in the same time period. So, I think the evidence that AI is coming for your job, at least this iteration of technology is pretty limited still.

Caitlin McDonald:

Yeah, that's a great point. And there was something you raised a little bit in your prior answer around the operational view versus the board view. And I wonder if you could delve a little bit into that idea between the difference in terms of what operationally as you are a worker who's working with these things or making decisions about them versus the board level strategic view and different responsibilities or different concerns that you might have in those roles.

Tim Gordon:

That's a great question. We just started with the World Economic Forum now preparing the C level version of the board documents. So, I believe... I think the board's job is basically to ask hard questions, supportive, but hard questions. The operational team, obviously, they're trying to deliver things which is

a totally different approach. I think what's interesting in the operational side though is the way in which a lot of the issues and the challenges, for example, around explainability, is this thing a black box and so on. Increasingly we're seeing software tools evolving to try and match and meet some of these requirements.

Tim Gordon:

And so, I think there is an increasing role for startups and actually bigger companies like Google who are trying to sell their cloud base... Not cloud base, their cloud real estate to produce tools and increase a use for managers. And obviously I think increasing it's important for executive leaders to understand what they can and can't do in the space. Board can operate at a level above that. They don't really need to understand technology. They need to understand the principles and understand what questions they're asking. And that's a different level of insights and education that AI is playing through.

Caitlin McDonald:

Yeah. And that takes me nicely on to a question about levels of maturity around AI ethics. So different companies of course have different operational processes, and different levels of understanding of what they need to be responsible for. What stages do you typically see a board going through as they go from just starting to get this issue on their radar to really becoming leading practitioners in the space?

Tim Gordon:

I think it's broadly true of AI there's a lot of conversation, and it's still at pretty early stages in terms of actual deployment and roll out. Even if you talk to some of them we work with or large financial services institutions say in the UK, they may have teams of hundreds of people working in this space, and they may be doing lots and lots of different tools, and trying different things. It's kind of still pretty rare they have stuff in production. And that's partly obviously because they're building on historical legacy programs, getting in that place is tough.

Tim Gordon:

Second thing is that a lot of these exercises are very distributed. So, you've got people in different places who will say, "Oh, let's try it. Let's see what we can do with our data. And in some places that's the IT department driving this. In many other places, it's the data or the analytics department. In some places, it's business heads. In some places it's chief executive's private project.

Tim Gordon:

So, you could see different ways in which this thing is running through. So, you have a sort of often a disparity across the piece. And part of what's often lacking is a conversation that brings things together. And we find a lot of the people we're talking to are the heads of AI inside organizations, but maybe two rungs down from the board, and they're struggling to have that full board conversation. And sometimes they have it because the board kind of buys this season of the future has been to WEF or listen to a podcast. But all too often, actually, these are conversations, which aren't always grounded in the hard reality, which is getting your plumbing to work and getting your IT systems to operate in such a way the data can flow is a really massive challenge.

Tim Gordon:

So, the compensations are often dislocated from the reality of what needs to happen. But then counter to that, you have the big US tech firms, the Microsofts, the Googles of this world, but are so far down

the track, they're now demanding regulatory intervention because they've actually looked... For example, Microsoft has looked at issues around facial recognition and they're going, "Well, we don't actually want a blunder into that. We want to be part of that conversation, but we fundamentally understand this needs to be government level intervention."

Tim Gordon:

And I think what you're seeing is in Europe, especially boards are mostly in that first place where it's very bitty. It's slightly all over the place, and from different conversations. In places like the US, possibly China, you have more sophisticated boards, and maybe a few UK companies who get this, but are still pretty rare. And most of the actual conversation is still at the sort of, well, what is this thing, and where's this going level?

Caitlin McDonald:

So, do you see the trajectory being developing those governance mechanisms internally then more industry level, and then ultimately the move is towards legislation. Is that what I'm hearing you say?

Tim Gordon:

Well, it's good. Look, one of the debates around AI ethics is what's driving wholly ethics debate. This is basically as capital takes it from labor, labor is basically the academics who used to drive this entire industry. And then also taking a step back, are they driving ethics up the agenda to try and make sure they still have a role? Is this ethics something which has been driven by the big tech corporations? You'd rather talk about ethics and talk about regulation. There are all these interesting questions, but I think regulation is coming. The European Union is ahead of the curve on this. Article 22 in GDPR basically is already put in place, that core foundational message around you need human involvement or your human explanation of the ultimate decision making taking place.

Tim Gordon:

And so, we're going to see that regulation, but we're still somewhere away from that regulation really impacting on the majority of corporates. And I think until you see some of these new far more often Chinese business models and financial, for example, driving AI first business models, and that starts to really impact on financial services, and so on. You're going to see some delays in this really coming up the agenda. I think the other challenge we've got at the moment, there's a lot of regulators are, it's probably fair to say, not yet quite as fully staffed up in terms of AI expertise as they might be. So, [inaudible 00:16:30] regulation is still several years away. Until then, we're going to have a bit of a Wild West that you referred to earlier.

Caitlin McDonald:

Yeah. Well, and also that raises an interesting question about international cooperation on legislation as well and international regulations because which I think we're already starting to see in terms of the way that firms, social media in particular are now being questioned in various ways, depending on where they are and the regional differences in how Facebook is regulatorily treated. For example, in the US versus in the EU. You might start to see similar battles. You can't see me because this is a podcast, but I'm making a movement with my hand, which is like a battle between two opposing forces, basically.

Tim Gordon:

Well, I think what's interesting is regulation and geopolitical conflict are almost becoming very closely aligned. So the US-China standoff, Europe's rise to try and become almost its competitive branch in a way. We're not having these big tech companies, but because we have the world defining regulatory standards around GDPR, for example, increasingly will define the rules of the game globally. And I think this is a huge area that is playing out in all sorts of level. I imagine if you're a big US tech firm, you probably want the world to be regulated in a similar fashion. And the European to some extent when you're successful, you like the European element because it has cost to all your competitors. But is the US going to accept that, clearly not. So, I think we're in a place where this game is playing out in different ways. And do we have the rise of the splint in it? A Chinese model, a US model, a European model? Probably, we're moving though to a more differentiated set of regulatory boundaries across the world due to broader geopolitical forces.

Caitlin McDonald:

Yes. And actually we've just released... The LEF has recently released a report on China, which you can find at the LEF website, which covers many of these issues in terms of rising cooperation versus competition across different kinds of industries in specifically China, and the US, but also looking at other kinds of regions as well. So, that is a follow-up for everyone. That takes me nicely on, and we've talked about this a little bit already in terms of power shifts in capital. You mentioned the idea that ethics being a rear guard action from academics who are trying to keep their relevance. But I wonder as well, there might be some other... The relationship between AI ethics and risk management, I think is quite clearly established, but there are other aspects of ethics that I think are more relevant to ways of identifying opportunity. Where do you think that side of ethics sits within board level priorities?

Tim Gordon:

Well, I think you've touched on what in my mind is the real opportunity in the space, which is trust is going to be absolutely key to succeeding in a data/AI driven world. And historically trust is a slightly nebulous thing that maybe gets measured on a few brand research items and so on. But if you think about the drivers of economic success in a data-driven world, your ability to gather data cheaply is going to be a core competitive advantage. So, if consumers trust you, they're therefore going to give you... They're going to give you more data, and they will require less economic incentive to do so. And so, you've moved to a world very quickly where the more trusted you are, the more data you're going to get more cheaply effectively. And obviously there are the other factors at play, but it will be an increasingly important factor.

Tim Gordon:

And you begin to think about the mix of stakeholders involved, that trust is going to be critical. So you think about in a world where employees can be mobile and increasingly they could be distantly mobile as well with sort of the post COVID working practices. They're going to want to work for companies that they like. And so, you've seen already walkouts and responses at a place like Google where the feeling is being that they've overset various ethical boundaries. Simply partners and suppliers who trust you will share information where they wouldn't otherwise. And so, you can see a whole series of ways which each of your stakeholders, there is a real clear, actual economic advantage you can put to the bottom line on increased trust.

Tim Gordon:

I think what's... It's a bit like the mirror element of AI is putting marks to some of these ethical questions, and therefore making the very real for particularly, leadership teams who maybe have always been mathematically driven. I think that's the real opportunity that emerges from the space. And so, you're going to have companies who really are going to focus on building that trust. And if you've got trust, you will then basically have a successful economic model.

Caitlin McDonald:

Yeah. That's a really interesting point. The mathematization of trust as what AI fundamentally is I think is an interesting way of looking at it.

Tim Gordon:

I think AI is that. AI is the trunk of maths over words. And so, huge amounts of the 20th century has been built on words, legal systems, linguistic practices, and so on. AI has a fundamental trend for maths over words. And so, as we're taking all those things. and we're putting them into digitalized numbers.

Caitlin McDonald:

Yeah. And I think speaking as a humanities background person, perhaps the ethics piece is you try to get a dig in for the words people again, potentially.

Tim Gordon:

Yeah. But I do think also given that what AI enables us to do is to put precision and numbers on ethical decisions that have always been... It begins to reduce the gray area, if you like, in terms of decision-making priorities. So, we talk about the trolley problem, which is the question of, is it... When put into the context of AI is an automated car that's about to crash. Does it kill the child or does it kill two old people? And then these hard decisions.

Tim Gordon:

Now, it's not really often really a practical question because the answer in almost every case told to engineers is you hit the brakes as hard as you can, and hope things work out. But you take a step back, and you think what this is about. This is about actually beginning to take these big ethical questions, which are always gray and trying to put them down into numbers. And to yes or no, one, zero type digital questions, and that's causing debates and discussions in ways that we hadn't had before. So something's bringing business ethics to the fore, casting them as AI ethics, but forcing those conversations in a way that has real numbers attached, and that probably is interesting at least

Caitlin McDonald:

I agree with you there. And I've gone, and sneakily looked up the name of the researcher. It was Joseph Weizenbaum who created the chat bot Eliza. And of course his original consideration was he was actually horrified to find that people engaged with Eliza as though it were a real person because his feeling was a machine is a machine and you shouldn't get emotionally, essentially caught up with what it can do or be confused by what it can do. So, he had a real kind of horror of what he had created. And I think that there's an interesting question as we approach some of these newer considerations around AI ethics. That question around what is it that AI enables us to do in terms of actually having an objective look at some of these ethical considerations I think is a new way of looking at it.

Tim Gordon:

And what's interesting is I've just talked about how it's all mathematical. I think about bringing my kids up and they're interacting with Siri and other chatbots. I'm going to tell them how they talk to Siri matters to me. And are they polite with Siri? Now, clearly it's a box. It doesn't care. But to some extent, increasingly their interactions with the world will be first through these computers and then with human beings. And as these computer get more and more human like, are you trying to get your kids to show good manners with these boxes? Probably yes, but it's this debate I have with myself on a regular occasion. [crosstalk 00:23:50].

Caitlin McDonald:

It's such a fascinating question because on the one hand you don't want them being confused about it not having agency because you don't want people to treat AI systems as though they do have agency because actually as you pointed out, the responsibility really lies with the humans that are in charge of the system. Whether that system is a company or whether that system is a piece of AI or a product that the company is running. And yet on the other hand, you're absolutely right. Our, our interactions with people are always rehearsals for other interactions with people. So, if you're not teaching your kids to be polite to the robot then can you expect them to be polite to other people in their lives? I think is a really interesting and very, very real and human question about this.

Tim Gordon:

Yeah. I think the broad of that is, we're all going to be living inside algorithms, and increasingly our behavior is going to be modified by algorithms. And we take a step back and think through what human... If you think about the way this technology is going to work economically, it's going to work by fundamentally changing how humans behave. So you think about cars. When cars are rolled out humans massively shifted our behavior to allow the economic success of cars to work. So, we basically accepted we wouldn't walk in front of them. We accepted that they would travel at huge, huge speeds. We accepted they'd kill thousands of our fellow citizens every year because we saw the bigger economic impact to them.

Tim Gordon:

And I think when you go to the ethical challenges that are going to come out of successful AI models, it would be the same sort of question. How do we change ourselves to make these models work? And so, all these questions about how we get AI to work for humans. The harsh reality is it's going to work the other way. We're going to basically ask ourselves, there's a bigger win here from this algorithm working. How far are we prepared to change our behavior to basically meet with the algorithm?

Caitlin McDonald:

That's a great point. And we could have a whole podcast just on that question. So, with that, any final thoughts, key takeaway messages for our listeners today?

Tim Gordon:

I just think the point I'd say to any listener is this is the debate or the questions that come out of these issues are going to define our lifetimes, our children's lifetimes. And it's something we cannot just allow philosophers. We cannot just allow the expert community. We cannot just allow the people who are working in the industry to drive this discussion. This is something which is going to such huge impact on all of our lives. It already is having a huge impact in our lives. And it's something which fundamentally is

going to become to my mind, one of the most pressing political challenges in terms of how we regulate and how we run our societies, what choices are we prepared to make?

Tim Gordon:

And fundamentally, these come down to these big political questions, which is, do we put the individual first? Do we put bigger society first? And so, these are constant trade offs that have to be made. And I would just encourage everyone to be part of that conversation because the more we're all part of it, the better chance we have of getting results that makes most of us happier.

Caitlin McDonald:

Excellent. Thank you so much, Tim. We really appreciate it.

Tim Gordon:

Lovely talking to you. Thanks a lot.

Caitlin McDonald:

Thanks for listening to the Growing Digital Ethics in Practice Podcast. You can find out more about Tim's work at www.bestpractice.ai, and you can find out more about the Leading Edge Forum perspective on digital ethics just by Googling the phrase, Stemming Sinister Tides, where the first link should be our 2019 position paper Stemming Sinister Tides: Sustainable Digital Ethics Through Evolution.