



Take on Board Podcast – Episode 319

Transcript – Candice DeVille to answer some of the best questions ever - her words.

Helga Svendsen: Welcome to this event with Candice Deville around ai. I know there's been lots of interest in it, so I would like to start as always by acknowledging the traditional custodians of the land on which we meet. For me, I am on the very chilly lands of Wurundjeri people of the Kulin nation, and I pay my respects to elders past and present.

I acknowledge their continuing connection to land. Waters, skies, culture and country. And I'd also like to pay my respects to any First Nations people that we may have here today. As I'm sure many of you have heard me say, I support voice treaty and truth for Aboriginal and Torres Strait Islanders peoples in Australia.

And I encourage others in the Take on Board community to do the same. Candice, hello, hello, hello. I can't actually remember who introduced me to you, Candice, but I know when we met, I'm just like, oh, you are awesome. You can explain stuff to me so that I understand, and that is a good sign and that is what most people in the boardroom need to understand.

So of course, I got her on the podcast to talk about all things AI and AI readiness, and then during that conversation I said, oh. I know what's gonna happen here. I'm gonna do this miniseries on ai. People are gonna have loads of questions, so Candice come along to this event. So that wanted Candice to come along and share all her gorgeous wisdom.

I reckon that's all I need to say. Candice, I'm gonna pass to you to share with us your beautiful wisdom. So Candice, over to you.

Candice DeVille: Thank you so much Helga, and thank you everyone for joining me this morning. I didn't wanna take up too much time doing a big presentation, but what I did wanna do was share with you a little of the training that we do with different boards and what the starting point for that is.

So I am gonna go through that just for a few minutes with you and then that will sort of help kick us off for. Questions a little bit more understanding, but take us right from the beginning. So what I wanted to do today was really go through this from the very beginning where we get asked as board members to do the ai.

So this particular presentation today, we're only gonna cover the first part, which is translating board directives into actionable strategy. If anyone is interested, I'll give you some more details on those. Other parts, and you can join us for that at a different time. But the main thing that we are doing today is looking at how do we translate those instructions into something we can do something with.

So we are not just sitting there going, okay, now what? I don't know if anyone's ever seen this show. This is one of my favorites. Tea crowd, but this is what I feel like when I hear from boards. We need some ai. And then I say to them, what do you mean by that? And they're like, ai, everybody's doing it. We can't fall behind.

We need some of this. It's incredibly vague. And non-specific, and it makes it extremely difficult for anyone else on the board, let alone the executive team to understand, well, what is the intention here? What does we need some AI actually mean, and how can you help lead the board in this particular project?

So what happens when we find that the board comes to you and says, where do we need ai? And we've experienced this directly with clients who've been in this position. The CEO then comes to us and says, well, we need this, but can you figure this out? We don't have a plan. And what happens is that when there is no really clear directive and no clear strategy, we see that teams interpret what this means very, very differently.

And they. Overlay their own bias, priorities, ideas, anything that they've brought with them. A lot of, you know, media bias, particularly as to what this actually means and interpret it really, really differently. And then what happens is we often see business units go off and start doing their own experiments.

So there's no sort of cohesive plan. No really. Oversight and then six months down the track when the board's saying, well, what have we done? We don't actually have any consistent progress. We have no way to prove value, or ROI, we have outcomes that don't line up with what the original intent was. And so if the board says, you know, bring me an elephant.

Well, which elephant do we want? Do we need the pink one? Do we need the gray one? The Indian, the Asian? We don't actually know. Until we have a better understanding of what that landscape looks like, and we have a coordinated approach. So today I'm just gonna give you the really simple steps to understand this conversation when it first comes up with you for the board, although it's probably already come up now, to be honest.

So the first thing that we wanna do is get real clarity on what. Is meant in this conversation, what does it mean we need to get the ai? So understanding the intent behind the directive. So asking those deeper questions as to the why and what they're actually thinking. Getting the rest of the board to actually articulate that what is it, the, the feeling that they have.

Is it one of fomo? Is it one of getting ahead? Is it finding competitive edge? What are we actually looking for here? And then the big part is, well, what business problem is this solving? Or is it actually just because we feel like we're behind and we have to jump on this trend? And then lastly, what is the success in this project look like?

And I mean, in hard numbers, not in feelings. So I actually had a really interesting and amusing conversation with the CEO the other day and I said, okay, so what are we taking back to the board? What KPIs are we gonna pull out of this project to take back? And they said, that's completely up to you, because the board only asks me, are you doing the AI things?

And I thought that is the least quantifiable metric that we could possibly have for success. So in hearing that, that's what's actually happening at these levels. We need to really start asking some questions. So I'm gonna give you four key questions to ask in this. And the first one is, what is the intent behind the directive?

So what are you really trying to achieve? Are we looking at improved speed, more savings? More innovation for the company or even compliance, like what is actually behind that directive. Then when we talk about the specific business problem we're looking at, at actually breaking that down into certain areas of the business that that can be applied to.

So in the same way as we think of medicine, what are we actually trying to heal here? What is the actual solution? To that specific problem. So we are talking about things like customer churn, manual overhead, risk exposure, anything else. Maybe it's profitability, maybe it is expanding to a new region.

There could be all sorts of business problems, but we really need those clarified. Then when that comes to measuring success and how we do that in numbers, we want to understand what are the metrics that we're going to use to prove that success. So whether that's things like time saved, cost reduced, customer satisfaction, and so on, these are all things that can be measured independently.

All our different levers for pushing the business forward and some of those things you may want and some you may not want more. Business is one of the goals that we're looking at. We change the process to be able to do things better. And then last but not least, we want to look at the delivery. When is this actually going to come through?

What is our timeline for making this. Happen and timelines tend to expand according to, to what we give them. And if we don't give them and say, how long is a piece of string? Then the project will just keep going and going. It is important to understand that with AI, because things change so, so very quickly, planning a project that is going to last for a year, two years, or three years without a final date to say, now we start measuring that outcome is a really bad idea because what tends to happen is that that delivery date gets.

Pushed out beyond where the technology is up to. And then what we find is we're measuring a project that is now not best in class anymore. It's actually already been overtaken. So this is a whole new way that businesses actually need to get their head around delivery time on a lot of things where, you know, a traditional software project might take 12 months, two years to roll out.

You'd be looking at a much, much more condensed timeframe to actually have this rollout with the technology being current and then being able to see what those metrics are. So our next thing, once we've actually gone through and asked those four questions, is alignment. So this one is very, very important and I don't know how much of a problem it is for you with the boards that you're working with, but we see this constantly.

We are constantly looking at the AI question and wanting to map that to the organisation's strategic objective. So if it's not tied to revenue, cost reduction, risk mitigation, customer experience, those are the things that you can clearly define. Then it's just a distraction. It's just being used as a toy and it's really never going to provide the value that you're really looking for.

So again, we're gonna have four questions that we wanna ask to make sure that these projects and these requests are in alignment. And the first one is, which strategic objective does this support? We need to be able to clearly articulate what this AI project is tied to. Now, we don't need to define in 20 different projects in this.

Even if we just look at that overarching alignment, we wanna see can we at least get two strategic objectives that this actually meets. Then we wanna see how this AI initiative is going to move the

needle. So if we cannot link this to a measurable business driver, this is not a strategy. It's again, a bit of play for the business.

And jumping on that bandwagon. Here's another really important one. Would we fund this if AI was not involved? And we all know that every VC out there right now is clamoring for anything that has AI attached to it. And you could pretty much see that any SAS coming out right now or just about any business is saying, we are AI enabled.

We're doing this, but what does that mean? And would you actually work with this or fund this if AI was not involved? Do you have a solid business case for this application? And lastly, what happens if we do nothing at all? What if we do not adopt AI? What does that mean for our business? And if the answer is, well, not really, very much is going to change, then it's really not worth doing.

And so our jobs as board members is to ask these really hard questions of everybody there. And you don't have to know AI in order to answer these questions, but you do have to understand how to be a lot more strategic. Our third one is prioritization. Something that we all struggle with. I know I can't do everything at once, but somehow I think I've got 20 hands and I can, but we need to really have a simple way of prioritizing what we wanna do.

And I'm gonna give you something very easy to take away because we want to have a matrix that we can actually put all of these projects or ideas in and very quickly and easily see. Should we be starting with this? Should we even give this time to look at? So our feasibility matrix, which we'll look at on the next slide, is your high impact, high feasibility, and we're gonna start there.

Then we can review those projects for what their risk profile is, things that should be done at board level before they even start going into discussions with the executive. So a quick feasibility matrix looks like this. We have our impact, so our, uh, value to organisation on one axis. And then we have our ability to achieve this.

Like how hard is it going to be to actually put this in place? And you can see in the colors here, we've got like strategic bets, quick wins, avoid and defer, and a low hanging fruit. And so we've got some example projects in here, things like demand forecasting, predictive maintenance, and so on. And we've just put those into each of those quadrants to say, alright, some of these are going to be much more difficult to achieve than others.

We can do them really quickly. For example, like doing your content or your marketing copy, the impact on that is not as high, but it's easier and faster to achieve. So we start to see what quadrants these go into. Once we have that, we can review the risks for these projects. So we'll, we should have, say your top three that you want to be able to look at and then apply this risk review.

So first of all, our data risks. Privacy breaches. Who owns the data? Where does it go? Is it going to be used beyond your control? Some of these projects will have data risks. Others will have really none at all. So if you are talking about using AI to write blog posts, you are not going to have a data risk.

So you can look at that there. Our operational risks, poor integration, downtime over reliance. How do we work this into legacy systems? How risky is that going to be? And our model risks models. It's never a one size fits all, and you don't need to know all of those things, but you do need to ask the question, what are we using here?

Are we using models that are inherently biased, that have inaccuracy or even that lack of explainability? We know that to some degree AI is a black box, but we still want to have a little bit more explainability and transparency over the process of how that's used. Our regulatory risks are very big, and these are tightening globally.

They're often things that the business case manager does not necessarily even understand these because they're coming up so quickly and those regulations are changing regularly. So if there are no regulations in the space that you are looking at, then the best thing to do is see what is the EU doing right now?

And plan accordingly. Use them as a bit of a litmus test for where we might be going, and use that as a safety system to say, if we can at least meet those standards, we are doing well. And our reputational risk. When people don't know that we are using AI in the organisation, it can actually really damage that public trust and it can damage the brand.

So those, we've seen them with bad AI outputs, that there's not been a human in the loop and what that's actually done to organisations. It could be the simplest thing ever. You may have seen the story about Secret Lab. Just last week, somebody used AI to. Essentially search for a phone number. They were given a phone number that was not correct.

It was for a different business called Secret Lab. And then the whole thing ends up being a big problem with people not following the governance and the AI applications policies inside of organisations. So it says even tiny little things that can cause problems to the brand reputation. When we review the risks, again, we can use our quadrant type system and.

Again, look at, you know, how risky are these sorts of things. And you'll see in here we've got our low priority and low impact versus things that are really critical risks. So regulatory, non-compliance, customer data breach, and so on. So you'll know this about your organisation better than anyone else.

And in terms of actually determining how likely is this to happen if you are working with a good AI consultant across these projects, this is the part where we actually help you work through that, then it really does. Or it could be something that we break up that project so only a tiny little piece of that is causing issues.

And just lastly, I wanted to share this one with you. This is a risk pyramid that is used in the AI industry. Just to sort of say really clearly what is minimal to unacceptable risk. So at a board level, if you're trying to say, okay, is this something that we should even think about without having to dig into the minutiae of the risk, you can very, very quickly see where your projects are going to fit in this risk pyramid.

So lastly, the big thing I want to tell you is the point is not to create fear. What we really want to do is just to be able to enable people to give them more ability to make good decisions and to make success something that is a clear and measurable goal. And just lastly, if anybody at the end of today would like some more information, I've got a copy of my book for you, which is about navigating AI.

So how to be AI ready. And this is the part that comes before you even think about putting AI into an organisation. And it will really help you work through a very, very clear step-by-step checklist on what you should be looking at in the organisation before you can say. Okay, we're ready to start considering AI projects, so that will be there for you as well as a free resource that you can take away and work on.

So that's my very, very quick whip through the beginning of how we start the conversation with our boards about ai.

Helga Svendsen: Awesome. Thank you so much, Candice. We had a few questions. All right, so first up, Melanie Eagle, you are first up the one about micro courses. So Melanie, can I get you to introduce yourself and to ask your question?

Melanie Eagle: Thanks a lot. Uh, Helga and Candice so any recommended courses, I'm particularly interested at the governance level so that we can get up to speed. I'm increasingly realizing it's gotta span the fundamentals along as strategic approach, or at least for me, the fundamentals heard Shiva make the comment about the University of Sydney course.

Any other recommended courses?

Candice DeVille: Yeah. Yeah. It's actually a really tricky space because AI governance is changing a lot, and what I've found is that some of the, the shorter courses that are coming out from university do tend to be a bit behind what's actually happening in the space. So one of the things that we've done as a company is we started putting together our own short courses about two years ago, where we were just, we have them online, we have free fundamental courses for anyone who wants to learn like AI 101.

What does it all mean? What are the names? What are, what does all the jargon mean? And then we've set up a Mastering AI for Professionals course, which is an online program that you can join. It gives you all the fundamentals, and then starts to move into other areas that are more. Specific to what you want, and you can kind of pick and choose from there, whether you're looking at things like governance, how to write a policy, how to actually apply it to your work, anything like that.

And we have that available as like an ongoing monthly program that anyone can just jump in and learn any time that suits you. So that's all there with the resources and available through our website, but I'll make sure that I send direct links to the information with that for you all. And we're actually about to up update the whole board and governance thing after last week and spending two days doing, uh, the tech and government summit in Canberra.

So there's a whole lot of new stuff coming out from that that will be added into there in the coming weeks.

Helga Svendsen: Thanks, Candice. Thanks, Melanie. For what it's worth, I reckon. Just find stuff and do it, but do it consistently. I've, you know, I am far from an AI or a tech expert, but I've found just playing with lots of different little things, each time I learn something new and it's just that constant build of stuff.

Beautiful. Thanks, Melanie. Next up, Julie. Julie, can you introduce yourself and ask your question?

Julie Lyford: Sure. Hi everyone. Julie Lyford. I'm chair of the Women's Environmental Leadership Australia Board and also on the Green Institute Board Wrestling so much with the AI issue, but from a climate perspective, climate and water perspective with all of the NGOs working in the environment space.

This is a really tricky issue because of the data centers and what's happening with climate. Big issue with Britain running out of water for data centers. So it's almost like the big elephant in the room.

And we are wrestling with what to do with ai, with Weller, and one of my major concerns is that we all jump in and make the climate issue even worse.

Candice DeVille: So it is a bit of a tricky one, this one because there's a lot of misinformation out there about the impact on climate, uh, of AI and what that actually does in terms of water. So there is more information that I can send you in terms of the actual carbon footprint in terms of all of this. And honestly, what we're doing right now, our zoom call with this many people is using way more water and energy than.

3,500 chat GPT inquiries, what we're doing now. So I'll send you the data on all of that so that you can sort of see the comparisons. And so that's one thing, yes, it does have an impact, but we need to be really clear on, in terms of what, in that ecosystem of what we're already doing. You could use chat GPT for a week and you would not use as much energy as you use.

Sitting on one 60 minute zoom call. So it's a very, very big difference. The other side to that is what we're doing about the problem and how are we solving this? And there is, um, a lot of local entrepreneurs who are really, really focused. On the sustainability angle, there's one in particular that we're working with in Gippsland who are building a data centre where the servers are actually water cooled as well, and they're using the water for aquaponics and using that to actually, with the, the fish filter, the water, grow the food, and so on.

So actually making this complete cycle of being able to produce more food. By using the, uh, the machines in the data centre there in Gippsland. So again, that's another one I'll send you some more info on if you're really interested.

Helga Svendsen: That is interesting to know Candice, and I think it is. So first up, there's an episode tomorrow on the day Take on Board podcast about this exact question about some of the ethical considerations and climate considerations.

It would be awesome to have some of the data. Candice. I think that would help enormously because there are many organisations I know that are like Julie torn on this, so getting some of the data would be incredibly useful. Thank you. Excellent. Thank you. Thanks for the question, Julie. It's on all of our minds.

Alright, next up. There's no name against this one, so I. We'll ask it, and my punt is, it'll be one that we need to pop over to the Facebook group, to be honest, but Candice, are there funding opportunities for not-for-profits to build digital infrastructure and workforce AI capability? Any thoughts on that?

Candice DeVille: That's a really tricky question. At the moment, I haven't seen any funding for. Anybody to do this other than there are some tax credits for digital spend in certain areas that I think run out this year. And then there's also R and D tax credits that can be achieved for certain things that you may want to build.

But as far as actually hiring those skills, I haven't seen anything on the funding side. The really tricky part there is. Finding people with the skills to start with because it's a very, very small pool of people who do have that ability and those skills to do it. And then also the, the salary levels for those tend to be quite prohibitive for a lot of not-for-profit.

This is one of the reasons that of, like with our business, we have created some pro bono work that we do where we actually go in and work with not-for-profits pro bono because we know that they

just, it's simply not an option for a lot of them. And being able to train internally. Is usually a much, much more effective option for them in the longer term than trying to hire just certain specialists from outside.

Helga Svendsen: For whoever asked that question, can I suggest that you pop it over either in the Take on Board Facebook group or there's a Facebook group called it's Run by the Institute of Community Directors, I think it's called Not-for-Profit Happy Hour. I think that would be a good question in either of those to to check out in the world to see whether there's any, but I think, I think it's.

It's gonna be tricky, quite frankly. Alrighty, next up Tracy, your question about Secret Lab AI misstep. If you can introduce yourself.

Tracy Bongiorno: Hi, I'm Tracy Bongiorno. I'm CTO mentor looking for my first board role. Yeah, I was just really interested in the story you shared about Secret Labs and someone accessing, um, a phone number through their ai and I just, I did a quick Google search and couldn't find it, so if you could share.

Candice DeVille: The article or or reference that you had for that? Yes, I will. It was in Smart Business about 48 hours ago, so I'll send it to you. Yeah. Okay, great. Thank you very much. Awesome. We'll share it in the follow-up email that comes out as well.

Helga Svendsen: Beautiful. Thanks, Tracy. Mel, uh, Melanie McGrath, the other Melanie that's here who's about to do Kickstarter with Candice, but I should say you're both in the program and I think Vicki is here as well.

But Melanie, can you introduce yourself and can you ask your question?

Melanie McGrath: Yeah, absolutely. I'm Melanie and my question comes out of, of work that I was doing for the CSRO for the last few years on specifically the role of trust in bringing AI into organisations. And what we found in investigating a lot of that, and I wondered if it's something you've come across, is even if someone has the idea of, of the metric for success.

Do you often find that they're not actually measuring that accurately or well or effectively at the baseline level, and so you really go in and you have a large piece of work to do around. Measuring that metric of interest before you've even introduced the ai.

Yeah, yeah. Sorry. The metric of interest in terms of do people want ai, are you just talking generally?

How do you measure those? No, just like if they're, if they have actually established, we wanna bring in AI because of productivity gains or, or efficiency or, or any of those potential gains that you highlighted. Do you find that they're often not actually measuring those or capturing those at baseline, and so you don't really have a point of comparison to work out if your project has worked or not?

Candice DeVille: Yeah. I love this question. Mel, absolutely love it because I never get asked this, and it is so, so important, and I can tell you without a shadow of a doubt, I have not yet met a business who is currently measuring their KPIs correctly. And so this is a big part of like where we come into the process where we're working with boards because we have a very deep background in business intelligence.

So we look at, okay, this is what they think the KPIs are, but then we go the extra mile of how is this actually being measured? How is this being tracked? Looking at the technology behind how they're tracking it, and is there something that is reliable numbers or are we relying on say, salespeople to click a button over here or fill in a note or.

Well, whatever that might be, that is really unreliable. How do we do that? How do we get the most reliable data first? If we don't have enough of that, then one of the things that my team does is we'll use what's there to actually model out good KPIs so that we can potentially synthesize some of that data, make sure that we do have good foundational KPIs, and then actually build from there so that we make sure that the whole process has been validated right from the very beginning, and we know that what we're working with is.

Good data. If we don't have good data, then we find other ways to get that out of the business, even if it's anecdotal. Pull all of that together. Do some data synthesis and. Then extrapolate from there, knowing that our baseline has come from that synthetic data.

Helga Svendsen: Awesome. Thank you. Thanks Melanie.

Thanks Candice. Alright, Mica, I'm gonna come to you next.

Mica Huynh: The deck that you presented earlier, Candice, it felt like it was to help the executive team respond to board ask. Is there a similar deck that helps the board better articulate their ask?

Candice DeVille: The deck that I have here should really be used by the board before they have those questions answered.

They shouldn't really be going to the executive. So even if it's KPIs, if they don't know the KPIs, that's. Fine, but these are the questions that you should be asking to the executives. So before any decisions are made, what that will help you do is as you will work through these, it will help you really clarify where do we not have this information?

Where do we need to get this impact from, this information from the board before we can go ahead. But I do feel like these are the things that as board members, we should be having this discussion. We should be talking about this together, because aside from that. The question of, well do we get the AI is really no question at all.

It doesn't actually have anything very actionable in it, so that would be the first part. Then from your question, I would actually have a separate set for the executive to be able to answer those questions. So where they should be looking in the business for these answers. Who should they go to? What department heads, what teams, what metrics should they be looking at?

That would be the executive instructions or deck.

Helga Svendsen: Awesome. Thank you, Alice. Your one is next, so if you can introduce yourself and ask your question. Alice, over to you.

Alice Chiew: I'm Alice Chiew. I sit on the board of a not-for-profit that's completely student founded and student run called Regional Education Support Network.

Our mission is to make Australia a place where geography has no impact on your education. My question is around sort of a clear AI policy as a starting point. I'm a little bit concerned about the, we

are all dabbling and playing, which is great, but sometimes I wonder if we fully understand how the engines we are using are taught what data they use and whether we are accidentally sharing data, they would never share in a heartbeat.

Candice DeVille: For the most part, people who are dabbling do not know the answers to those sorts of things, and it is really, really important to get that fundamental understanding. In fact, that is the, the linchpin of everything that we do in the business where, you know, whether we're doing a lecture or doing a project with a company or anything.

Always gotta be that fundamental understanding because otherwise you just don't know what kind of animal you're dealing with. For the type of group that you're talking about, Alice, um, we actually created a whole lot of free tutorials that we can deliver online. We do have the option to deliver them in just about any language that we need to, thanks to the power of ai.

Which is absolutely wonderful, but we have all those available for free. I'll make sure that you get the links to those. We currently have just over 24,000 students across the world in these courses, so they're really nice and simple, bite-sized. Just an hour kind of thing. And I've broken them down into smaller lessons so that it's accessible not just by geography or language, but also how much time you have.

So, you know, if you're busy cooking dinner and you just wanna learn something for five or 10 minutes and watch it on your mobile, you can do that. I did that because that's how I have to learn a lot of the time while I'm busy, juggling kids or food or whatever it is.

Helga Svendsen: So Candice, I don't know if it was in the conversation with you to be honest in the podcast series or somebody else, but one thing that I did post doing some of these interviews, and it was around this entirely because I've been playing, is I now don't.

Use, I've only got a paid account. That was one of the things, is that right? Is that if you are using a free account just to dabble, woo-hoo, there goes all your data wherever you go, if you've got a paid account, it is more protected. Not a hundred percent protected, but much more protected. Is that the case?

Candice DeVille: Yeah. Yeah. So exactly like Alice is dropping in the chat, if you don't pay, you are the product. We know this from social, we know that for everything, it is the biggest warning sign ever. You don't pay your other product. And we've seen this in a lot of things where you can, you know, they'll say, don't train on my data, et cetera, but the, the level of protection is not the same, the paid version.

You do have different options in terms of privacy where things can be turned off or deleted. There is still some controversy about that where OpenAI said they were deleting things and they didn't delete things. But there's more secure versions as you kind of go up the ladder, so to speak. So if we're talking about chat, GPT, the paid version gives you a lot more privacy to what you can do.

The actual OpenAI models that are accessed via API. Is different again, and you can put your own layers of security and things onto that. But for dabbling, the safest thing to do and the best thing to do is just pay for the simple Pro subscription. I would still recommend doing that with chat GPT just because of the flexibility of the model and what you can learn with that.

But for 20 US dollars a month. It's a small price to pay for that peace of mind.

Helga Svendsen: So Laura, I'm gonna skip down to yours. If you can introduce yourself and ask your question, that would be great, Laura?

Laura Haines: Absolutely. As I'm Laura Haines, based in Sydney, I'm Chief Technology Officer for CES, which is the technology behind your, return your container and get 10 cents across Australia.

And I've previously sat on a different board, um, and currently I'm on an advisory board and looking for kind of my next step. But my question is there's lots of different groups that are needing support in terms of upskilling in AI. There's the boards themselves, there's the business leaders and the people in the business.

But then there's also the people, like people in this room who are trying to take their knowledge. Learn AI and upscale so they can help boards or individuals. So I just wondered what you're kind of seeing Candice in terms of the biggest areas of impact. Because there's so many people to help, we can't.

Focus on everything.

Candice DeVille: Yeah. Yeah, it, it is really tricky because there are so many people to help and educational learning is one of my biggest passions. So helping people reach their potential, and when AI sort of first hit the headlines and people wanted to learn about it, everybody was. Racing to catch up from behind how do we actually do this?

And there was really two kinds of learning methods on the market. One is the, the academic road. You know, going to the university doing this course, it does tend to be a lot more technical. So that's more for like your engineers, your technicians, and so on. But also the university courses don't often keep up with where things are today because it changes so quickly their syllabuses cannot move fast enough.

And then. On the other hand, we have a lot of like little training courses, things that are going on from, you know, solopreneurs, people that are just putting things out. Local government might be hosting things in your area, but it's still not addressing that business case. Use, and this is why over the last few years as we've worked with clients, we've realized there are very, very clear segments that we need to address in a lot of this training and being able to make things for them.

So that's why we've built that. We haven't yet actually found anybody else that's. Doing this in the same way because there's a lot of opportunities for, you know, come in, do a workshop, that's all great. Yes, we do that too, but we wanna make things available to all of those levels. And that does start with that fundamental understanding and then branching off into what your specialty is there.

So whether that's learning more about AI for HR. For example, what are the risks there? What should you be doing? What should we be looking at with your team? There's a whole section on things like building your AI use policy in the organisation. How do you do that? And is this something that should be a top-down directive or should it really be something that involves everyone in the organisation and come from bottom up, which is where we find that.

The policy is adhered to a lot better if it actually comes from the bottom up. And we start talking to people about, well, what are you really doing? What's called shadow AI use, which is people in the organisation that you don't know have a free chat GPT account and are just sort of putting things in there to do that and.

It's one of the biggest problems in organisations right now. So rather than from the top down saying, we have to have these policies and you are going to do this, what we wanna do is talk to people who are using it, find out how this is helping them, why they're going to these places, and then understand what they love about that, and we can make policies from there.

So. To your question about how do we train all of this? The only solution that I can give you right now is the one that we're continually building all of the time because we're doing this from real live questions. And even today, everything you guys are asking me will help inform the next version of this that comes out to be able to answer these questions for people, because we're just finding that it's such a desperate need and we want to be able to deliver this in the most accessible way possible.

People who are not technologists who don't need to understand all of this. They just have to have that fundamental understanding of how it works and then be able to communicate that to other people in the team in a way that they can understand with no technology background and it's entirely doable and something that sounds a lot scarier than it is when it comes to learning ai.

But honestly. Even with one or two hours, you can get those fundamentals of how this works. What is a transformer, you know, how does AI take my question and then come up with this answer. What is the process for that? Does that help answer that question? It's a, it's a very evolved one. I think so. So the biggest impact right now is helping non technologists understand the value of ai.

Yeah, I think so. And because we're already seeing this split between the people who know how to use it and those that don't, and that gap will widen significantly and much faster than we've ever seen before because of the speed of development of this technology, and then the speed of those who are actually applying that, so that that gap is happening quite quickly.

So you don't have to be a technologist to understand it or even to use it, but to know where your sweet spot is. And what is right for you and your organisations? That's a really big piece at the moment. The other second piece, and we can probably deal with this on another question another time, but looking at how you hire for ai, and that's a really, really big problem right now.

We're seeing organisations who are trying to hire AI experts, but have absolutely no idea what they really need, what that skillset looks like, how they validate any of those skills. Do they need an engineer or somebody completely different? So there's a whole lot to unpack in that space as well.

Helga Svendsen: Kate, can I get you to introduce yourself and ask your question?

Kate Wilkinson: Thanks, Helga, you and Candice. Uh, I'm Kate Wilkinson. I am a consultant for the health sector and I also work with an Aboriginal and Torres Strait Islander Wellbeing Research Centre at ANU. I'm seeing a lot of different new AI platforms promising to do certain things for certain types of organisations.

Just wondering how do organisations assess what's actually going to, you know. Be worthwhile investing in, or should they be building their capability in-house to use AI to build their own tools to achieve those things?

Candice DeVille: So many good questions. I love those questions. I would say that these are the best questions I've ever been asked in one of these sessions.

Helga Svendsen: So the Take on Board community is smart. We love it.

Candice DeVille: Okay, so this is a really, really great question because there are more AI tools coming out every single second at the moment, and a lot of them are being made obsolete within months, and a lot of them are also coming from smaller startups that don't necessarily have a huge support element to it.

So there is a lot to consider even in terms of the security of these third party applications. So generally when working with a company, we would say this is an area that you do need to consult with an expert on. If you are trying to evaluate third party tools versus building it inside, it's really going to depend on the exact project and the application of that, because you may want a bit of both.

It's going to. It'll depend on the size of your organisation, your risk appetite, all these sorts of things. So the short answer is consult with someone who really knows. Bring in an expert to talk about this with you, because they'll be able to identify very clearly what you should and shouldn't be using for those use cases.

And then if you want to build internally, which. Absolutely. We do that all the time. Build internal solutions for clients. It could be because of privacy and compliance, but it could also be that you can build a much more elegant, cost effective solution that you can manage internally if you have the resources to do it.

So, short answer, it depends. Get expert advice.

Helga Svendsen: So the next question is, and you talked about shadow ai, which I hadn't heard the term, so that was very useful for me. So the next question is, what tips do you have for piloting AI brackets, co-piloting Microsoft to address shadow ai?

Candice DeVille: Very, very difficult because.

The only way you can address shadow AI is to know that it's happening, and there's only really two ways that you can know that it's happening. Either you ask people and they're honest, or you invest in a tool that actually finds this for you. So far, I've only found one that does it through Cisco. You need to have that kind of enterprise budget in order to do it.

So for most organisations, that's not really an option. You actually have to know, and one of the good ways to pull that out of people is instead of asking them, are you doing something that is kind of against our policy or shadow AI and making that feel like a bad thing to encourage them to actually work with you on developing your AI policy.

Because when you bring them into that conversation and say, what have you learned through all of this, you can get alignment from everybody. They'll tell you what they've learned, what they've been using. They can contribute to the building of that policy. You get a lot more openness because you're asking for help and feedback of what they've already learned rather than just.

Closing the gate on them. Once you've done that and you've got a good policy in place, you're not necessarily saying to them, Hey, you can't use this thing that you love. It's, you can use it within these safe parameters, and that's the best way to find it and address it without having to invest in very expensive monitoring software.

Helga Svendsen: You know, it, it feels like, you know, telling teenagers they can't do something just means they're gonna do it. But if you actually have the conversation about how it can be done, well

exactly. It's gentle parenting to find out about shadow AI use. Makes perfect sense. Alright, the next question again, I think you've touched on this a little bit, but how do organisations assess which AI platforms are worth investing in?

Or should organisations be learning how to use AI to build their own tools?

Candice DeVille: The short answer to that is don't start investing in building your own tools until you really understand ai, because it's not an investment, it's a risk if you don't know. There are definitely going to be risks in using any third party tool, but in terms of having a sandbox to play in, at least when you're doing that, you can then look at, well, what is their security documentation?

How do they deal with data? Learning to ask some of those questions while you're working with a third party tool and getting your AI fundamentals will be the first step before even considering building that in-house. It's hard enough to know where to start with ai, let alone how to build your own. I, that's my view anyway.

But that's 'cause I don't know much. But yes, look, honestly, one of the things that we've seen come out, and you will see this in headlines and over the last six months, is the whole vibe coding thing where you've got a lot of people who are vibe coding, hang on, vibe, coding, what is that? So Vibe coding is essentially being able to work with AI to code and build something when you've never done it before.

You just go with your gut. You go with what the AI tells you. You work through that, and that can be super, super helpful, particularly for people who are already technologists. But if you don't understand the process, the infrastructure, the documentation, the security, all of those other layers around the coating as well, you're also walking into a trap.

Helga Svendsen: Alrighty. Our next question is from Melanie McGrath. So bound to be a good one 'cause she's asked some awesome questions this morning. She asks, how can we improve our ability to identify the unintended consequences of our AI implementation plan? Ooh, really good question. The one thing I love to do is ask the why question multiple times.

Candice DeVille: So when we first start out and we say, okay, we wanna put AI in the organisation, that's our first question. Why every answer that comes through keep asking why, why, why? Even if it's, we are choosing this model for the project, why are we choosing this model? Then we can start to see things emerge from that, that answer this question.

Things like, we're choosing it because the data training set is different. It doesn't have this bias or it has this ability and it's the best one for coding. So asking those why questions the whole way through the process is the best way.

Helga Svendsen: How do boards manage the ethical risks around data sovereignty and safety?

For example, in research, in health, and for young people. And I think that's probably connected to another question that is here about what are the key considerations regarding risk, human rights, and privacy? Should we keep in mind for data, especially for those who are disenfranchised, it feels to me like they're connected.

What are your thoughts around that?

Candice DeVille: So they are kind of connected. And the first one about the, you know, how do we keep privacy in mind for, for those that are disenfranchised, that should always be. Front of mind on any of these projects, and as I was talking about in the the risk section, if you look at a project and say, okay, this is high risk because we are dealing with health data, that doesn't necessarily mean that we throw the project out.

It means now we need to talk to experts about where that PPI actually sits in the project and how we can manage that. In the best and safest possible way. So there's all kinds of layers of security that can be looked at, but the first things is making sure that you know it's GDPR compliant, SOC two, and so on.

So all of those standards, we want to be able to meet them with our project, but particularly to just dig into that a little bit. The other part of the question was around data sovereignty. And this is a very, very interesting question. I was actually, um. Speaking with Google about this at length the other day in the Google Cloud team, having an understanding of what data sovereignty actually means.

Is really important because data sovereignty and data localization are two different things. So just because your data is housed in Australia doesn't necessarily mean you have data sovereignty. By the same token, it could be housed in Singapore and you could still have data sovereignty because of the control.

So there's a lot of like. Nuanced questions to ask around that. And again, it comes back to the why. Why do we need data sovereignty? How, what does our exposure look like? And when we are working on these projects where a board has said, okay, we have a policy with. All of our clients that of data sovereignty, that their data cannot leave the shores.

That's data localization and different again. So we need to know what that actually means. Then we can work with them on how do we create the exact right solution. That really addresses the question, not just ticks a box on a list that says data sovereignty. Oh, this is so awesome and so break my brain all the time.

Helga Svendsen: Every time I think I've got a tiny little grasp on things, there's just more, more, more. Awesome. Thank you. The other part of that question, the key considerations around risk, human rights and privacy. I'm just wondering particularly on the human rights side of things in data, is there particular considerations there that people should be thinking about?

Candice DeVille: That is a really, really tricky one because there's no clear legislation on any of that. It's really down to the organisations themselves and how they choose to app actually apply this. So it's certainly an emerging field where people are looking at this and how that data is. Collected how it is used, how it can then be used to manipulate people.

It really is an emerging field and something that you need to pay close attention to because we know that it has not been used well in the past. And not to name names, but we can look at meta as a great example of that. How people are manipulated because of their data and they are more marginalized or they're hidden, or they don't see the groups that they need to.

There's an awful lot in there, and it's not just a question that has come up now. Since people have started using chat GPT, although it's now front of mind, but this has been a big problem for many years and anyone who's used Facebook will know that. So this person asks, while working in

organisations collecting clients' data, apart from data security, what other risks are involved in using ai?

Yes, it depends on how you use that data. So one of the big risks is hallucinations. Now you've all heard the term hallucination. That means essentially that the GPT is giving you back something that isn't true, but it looks really true these days. It's considered a security flaw, although it's actually part of how AI is built, right?

So it is just like that all the time. When you give it a whole lot of data, let's say you've given it a spreadsheet and you ask it to do things. Inside A GPT, there is a high likelihood that some of that data will come back and it will not be correct, it will be hallucinated. So there are other ways to work with that data where you can prevent that happening because of how the solution is built.

But. For a lot of individual users and dabblers who are just starting to get into it. If you do that, then what will happen is chances are you will not get back the correct data, and unless you're going through, you know, 70 lines of things, you may not actually see where that problem is, or you may not notice because it might be a change of phone number, a change of spelling, an email that does doesn't exist.

Helga Svendsen: We're down to our final question. So this person asks, my board recommends management use ai, but it seems to fall on deaf ears. How do we convince or influence management?

Candice DeVille: Okay. Another really, really, really big question. First of all, management needs to understand AI before they can even consider it. So they need to have some foundational fundamental training.

Once they have that, then. They ideally should be using some AI in their daily life to have an experience with it. And when we look at change management of bringing AI in, whether that's for, you know, the, the sales rep that needs to use it, or for management, they need to have a win. They need to have that experience of going, wow, this actually changed something really meaningful for me.

And without that experience, they don't have an understanding of what it could do. And so we like to, when we're actually working with management or training people, it. Give them a little tiny task or a tiny project to do where they can have that win. And it might be even something as simple as go home and menu plan for your entire family with all of the things that they hate, and no Brussels sprouts and no gluten and all of it, and how hard that is to shop for.

Tell your ai, get it to help you. And then when they can actually see the impact, you'll start to see that the walls will come down and they're more interested in doing that because at the end of the day when you use AI and you get back that time to go to your kids' play or knock off early on a Friday, you're gonna want to use it again.

Helga Svendsen: And the only thing I would add to that is I think if, if boards are asking management to pick it up, it might be a capability gap and a playing with it gap. But it also could be boards are not explaining. Clearly enough. What the ask is and what the why is as you say. So it might be almost back to where we started this conversation about, you know, we've gotta use ai.

I'm not saying that's what's happening here. 'cause I know the name of the person and I know that there would be some good conversations about it. But I think if boards are, are pushing on it and not

getting buy-in, then for some reason management's not hearing the message. The board is pushing out there.

And it might be that boards are not clear enough in their message. Maybe it's look back on. Look back on yourself first. It's like, why are they not hearing the message? What is it about the message that we are sending that we are not explaining as well as we could? So it might be about boards engaging in that conversation better.

Candice DeVille: Yeah. Yeah. The other missing piece is making sure that the board facilitates training for the management. Not just go do the thing that you don't know what it is. Even just give them, you know, a half a day course book in a session. Get us to come and sit with them, work with through it. Like very important.

Helga Svendsen: And do it together. Do it with the board to work through some of these things. Yeah. Awesome. So again, thank you so much for taking the time to do the original podcast. Thank you so much for taking the time to turn up today. And I answer, everybody's gonna say they're my stupid questions, but everybody else's fantastic questions and for continuing going to share your wisdom with the Take on Board community.

Candice DeVille: So helpful. Thanks for being a superstar. You are so welcome. I look forward to seeing you again soon.

Transcript by Descript