

Dear readers: I interviewed Charlie Beckett for a report I've just finished writing for INMA about generative AI. For those of you who want a deep dive before the report releases this month, enjoy!

Ariane Bernard 00:00

You have a unique perspective because you're a media specialist, but you're not trying to sell technology and you're not a publisher. So from your perspective, what is urgent for a publisher to get on when it comes to generative AI?

Charlie Beckett (LSE, Polis - JournalismAI) 03:23

I can talk about this without a personal stake. Everything I say is repeating what other people have said to me. So I've been listening rather. And I think there's a lot of really interesting speculation at the moment. Much of it is kind of fictional. A little bit is kind of ideologically or imaginatively-driven. And it's kind of racing ahead of the technology. And that's not a bad thing.

But I would say the first thing that people should do, drawing upon in the last five years of us looking at this, is just get across the basics. Before you start thinking about ChatGPT or MidJourney, or anything like that, just wind right back and make sure that you've got any sense of what we're really talking about. I don't mean that in a patronizing way because I don't know. I'm not a technologist. I spent the last five years trying to get across and understanding the concepts and so on. But leaping off the top diving board is not the way to proceed. And it's not just about you personally, doesn't matter whether you're the CEO or whether you're a foot soldier journalist. Everyone just needs to get across some basics about it because artificial intelligence does not exist. It's a terrible term.

And I'm cursed, I think, to have a project called JournalismAI because there is no such thing as intelligent, artificial computer programming. It's all much more mundane, and much more various as well. It's everything from you know, kind of machine learning to automation to personalization. We use it every day. When we do search for example, we've been doing that for about a decade. So this isn't something completely new. And at some level, even with the new generative AI, it's just a form of autocomplete. All it does, is predicts the next word. It doesn't know what it's talking about. It has no let alone any intelligence. It doesn't know any kind of concept, let alone emotions, or feelings or insights at all, is purely programmed by humans. And then, as you know very well because you know more about the technology, it's become exponentially good at doing this, learn from the data and being able to give the appearance of being able to do quite extraordinary things and that's what's different. I think about the generative AI. I am very excited about it. I think that its creative potential augmenting journalists is amazing. But first of all, before we get there, make sure that you and ideally your colleagues are across the basics and you know, I'm

sure you've been linked to the basic training courses that we do. So you got to walk before you can run.

And the second thing I would say is about the sort of generative AI in the chapter up and so on is don't do anything with it. Or at least don't do anything that is direct to publishing with it. Don't do anything that risks the quality of your editorial processes. And that could be from news gathering right through to publication at the moment is much too volatile. It's much too risky to do anything that goes to print if you like. Of course, I would absolutely recommend that you experiment with it. That you start playing with it. I would suggest you do not write checks to people to buy in tools and programs yet, until you've been through a process of thinking through what problem does it solve? Does it do it any better than what you had before? And do you understand how it works? If you can't say yes to all those things, then you know that why are you doing it?

So I mean, we said this about AI generally. And to be honest, I think it applies to many of the technologies and tools and processes frankly, you know, you wouldn't hire a fantastically clever intern, age 19, and make them editor of your publication just because of their potential. And because they sound brilliant. You'd want to know what they could really achieve and how they fit into your processes, and also how your processes might have to change. You know, I think this is the implications of AI, generally, or normal sort of generative AI are probably even bigger.

The last bit I'd say is that journalism is different. I think this set of technologies has got risks and opportunities and other fields which are much greater, frankly, at one level, you know, that old adage about technological change, we always exaggerate the short term and underestimate the long term I think is really really true of this one in terms of journalism, that journalists have gone around saying, Wow, this thing is inaccurate. I gave it a really clever, weird prompt and it got it wrong.

And this is going to take over all our jobs and it's going to end the world literally, you know, they're talking about the end of the world. No, certainly not overnight, but longer term. I think it may well reshape the nature of news information and the industry and the way that we work just as it's going to have interesting implications for everything from health, you know, to retail to security, etc.

Ariane Bernard 09:41

This is absolutely great. I use the adolescent metaphor too. ...I feel like you kind of touched on this, but when you know, hear some of the sort of more negative headlines, what gives you pause?

Charlie Beckett (LSE, Polis - JournalismAI) 11:13

I think one of the problems with these debates is that they are dystopian or utopian. And we know very well that technological change is much more mixed and there are unexpected consequences. You know, I think some of the kind of imagery and metaphors about you know, robots taking over is incredibly unhelpful. The whole Matrix thing.

But I think that one of the reasons why we should know a bit more about AI, all journalists should know a bit more about it: It's not just it's going to change their jobs, but it's going to change the world that they report upon. You know, it's really important that you get across this a bit, doesn't matter what beat you're on. AI is going to be part of it. And you know, we can see that we should be alert and critical, because there are all sorts of potential consequences. Most of them I would argue, to do with the choices that we as humans made.

So on one level, we face the same risks that we've always faced in terms of the detail about editorial strategy. When the internet came along, you had a choice: You could be a kind of clickbait factory producing terrible journalism. Or you could go you could use all these wonderful digital tools to do the amazing kind of data journalism, And provide a really fantastic service to your subscribers. So we face similar choices about it. I think there are fear is definitely when people say "it's gonna take your job". Yeah, it might, actually. Yeah – If your job is producing pretty crappy automatable journalism, or journalism, that's just you know, very automatable like the weather, then you may well... that job may well go.

And indeed, much journalistic labor is going to be replaced, just as mobile phones and search replaced a lot of journalistic labor in the past and we welcome that, generally speaking. The risks for individual journalists is that they have to learn new skills, and also that they may have to work harder. It's interesting that yes, it can save you enormous amounts of time. But generally speaking, these technologies often create work in the sense that we're creating new products. And you can't just let them loose. You know, it's not just it's not like the robots that cut your lawn. They ended up in the swimming pool, don't they? And this similar thing can happen with this stuff, that you have to edit the technology you have to program it, you have to feed it data, and you have to oversee what it does, and adjust it and change over time. And so new roles and new tasks will arrive.

So in the current constraint and journalism where it's so competitive and so revenue constrained the idea that we're all suddenly going to be able to get our deck chairs out and take it easy with cocktails all afternoon: it is nonsense. This is not some sort of panacea of journalism's ills. Journalism has enormous problems around you know, the desertion of advertising, political pressure, persecution, the competition from other sources: AI is not going to solve that. But from what I've seen, if you're not using these and other associated technologies, then you are putting yourself at a disadvantage. So I don't think this is some sort of dream scenario. Conversely, I'd say

that there's a kind of positive template which is that in the same way that you know, I would love to be young journalists now you know, the amazing tools that you have the amazing range of sources and platforms and ways of reaching your audience are quite extraordinary and AI offers lots of interesting ways of doing things. So anybody with any kind of creativity would be looking at these and thinking -- this is interesting. A bigger point, which is a bit more kind of, I think, applies to AI in general. Not just to journalism, which is the usual stuff about who made this, who controls it? And who regulates it, and those are huge societal questions.

Ariane Bernard 17:16

amazing transition from my next question. Thank you so much. I wanted to ask you about the kind of like the meta environment right for the development of generative AI, whether that's the legal environment, whether that's fair use copyrights, both in terms of how we inherit copyrights that may or may not actually shake out to be actual copyrights.

And, you know, what's interesting is there's like a question, what question for you that's very obvious, which is like what you think is the need for regulation like what are the things that you feel whether it will happen, is a different thing, but if you were a legislator, where would you wedge yourself and say “this cannot just be left to private interests and enterprise”.

You could treat it as a separate question if that's easier, but if you are a publisher and of course they some of them range for the super risk-averse to the kind of like, you know, “let's do it until we're told that this is no longer legal”, how should you approach using generative AI?

Charlie Beckett (LSE, Polis - JournalismAI) 19:49

well, on the first part, I'm both sort of fatalistic, and perhaps instinctively, sort of anti-regulation. I don't mean I'm anti-regulation in general: I'm a complete believer in regulation. We regulate cars, we regulate food, regulate medicine, all good things. But what worries me about regulation in this area is it doesn't matter where you are. That it's one thing it say in the UK to say let's limit the potential harm of these things. But elsewhere in the world, the problem is lack of freedom, and authoritarian regimes of all types using any excuse to clamp down on the internet, you know, we'll see not just China, not just Russia, look at India, you know, look across Africa, we can see that both media freedom and freedom in general is declining. And technology, which we hoped and still can open up societies and give people voice are being constrained. So I'm actually more fatalistic. I sort of say: let's accept some harm, some risk. I'm more concerned about accountability, I think, rather than regulation. So accountability in terms of what these companies are doing, and transparency about what these companies are telling us about what they're doing. I mean, OpenAI is relatively open.

Ariane Bernard 21:32

How do you enforce that though?

Charlie Beckett (LSE, Polis - JournalismAI) 21:35

Yeah, I know. It's incredibly difficult. But you can say, for example, that like existing social media platforms, you can say that if you want to be in our territory, you've got to give access to your APIs to academics, for example, so they can do research. That's a condition of being in our territory, or Europe who is being quite progressive on this, you know, quite activist on this. Nothing that's relatively healthy.

Ariane Bernard 22:08

That's an argument for auditability replacing legislation.

Charlie Beckett (LSE, Polis - JournalismAI) 22:11

Yeah. But what I don't think is, you know, telling, which is happening everywhere else, which is, like India -- is telling people what you can have on those platforms. You know, that's what worries me from a news media perspective. From an industry perspective, from a publisher's perspective, again, I would say to them: Be careful what you wish for. I think some of the, you know, deals about link tax and so on are basically just kind of ... extortion rackets. But it's very interesting talking to publishers of all ilks, I don't just talk to nice, cozy not-for-profits. I talk to highly competitive, powerful media organizations as well, and who are not friendly to the tech companies.

Ariane Bernard 23:03

They're fresh from their success with Google...

Charlie Beckett (LSE, Polis - JournalismAI) 23:06

Indeed. And they see potential for more competition, actually, because there's not just one company creating generative AI and they see opportunities going back to something I spoke about, like 10-15 years ago, when I talked about fake news or misinformation, which I described as an opportunity for at least good news organizations. That if there's a whole bunch of bad content out there, than if least you're better than that, and of course, journalism has always been a source of bias and misinformation itself at times, but at least if you're better than what's out there, then you stand a better chance of earning people's confidence. Those people that want confidence and if people want rubbish and nasty stuff, they'll go and get it. But if you want to get people, you know, let's call it quality news, right? Then, you've got to make yourself distinctive, you'll stand out. You've got to tell them what you do and that works. It is actually working. Journalism may not feel it, if you're a badly paid, precariously employed person doing it. But it's actually a success story: There's more journalism consumed there than ever before; there's more good journalism out there than ever before. So in that sense, you know, we shouldn't be fearful

and the people I've talked to do say, for example, that there may be [...] people who are producing generative AI tools want to have licenses as it were, arrangements with news organizations to say: 'Look, we're not the crazy generative AI people. Our ChatGPT is going to declare its sources of where it got its data, and we're going to reward the people who gave us that data, or helped us train the data to demonstrate to you. You run us, rather than that dodgy ChatGPT by somebody else, which is just crazy and hallucinates and doesn't link back to anything. ... So again, you know, I may or may be being over optimistic, but Gresham's Law may in some sense, apply to this.

Ariane Bernard 25:25

It's great that you're drawing this parallel with the sort of like, how one fights misinformation or spam or low quality content on the internet, which is more of an organic sort of battle where the good basically occupies space against the bad rather than legislating out spam, which, good luck.

Charlie Beckett (LSE, Polis - JournalismAI) 25:55

Look at America, you know, the home of the free. Where, you know, right wingers are trying to ban various books and stop teachers saying various things... These are humans who are restricting speech. And you know, some people think they're right to do so. But that's a real problem. You know, I think we have to accept, I know sounds complacent ... But in some senses, we have to accept the existence of misinformation and it's much more visible and it can be generated faster and spread quicker. But you know, I've looked at this passionately for the last 10 years. I did a whole commission at the LSE on misinformation back in 2018. And we did come up with ways that you can structure the environment to make it healthier, you can improve media literacy, you can support better journalism, but [...] you just have to accept that there is this mixed economy, which includes some terrible stuff. And generally, the research shows that perhaps we've exaggerated the downsides, the moral panic, and I suspect we're probably the downside, but the caveat with that is that I also don't want to say, "Don't worry". I think we shouldn't be worrying massively. And I think that the rise of you know, the AI ethics sector is incredibly welcome. You know, there's a tendency to, like with all of these sectors, to spend all your time finding fault and finding danger. And that sometimes means we could decontextualize what's happening. Because you're so busy looking for terrible things that, I think, you know, having a more critical approach has been fantastic. I think that's one of the one of the healthiest developments over the last few years, not just in academia, I mean, in society, you know, organizations like yourselves, thinking critically about the sort of downsides and risks.

Ariane Bernard 28:44

I mean, that's the thing that I think sometimes interrogates which is this assumption that AI is going to enable bad actors to do things that couldn't do before. It's like, no, they could already do all that. They just would do it by hand. And if you're Russia, you have cash to pay as many human trolls on the internet to do whatever it is that you're trying to do that month. So not that I'm

thrilled that Russia gets access to cheaper ways to do whatever it is doing. But fundamentally, and maybe more, but I mean, fundamentally, it's not necessarily moving the needle that much on.

Charlie Beckett (LSE, Polis - JournalismAI) 29:20

Yeah, I think there's a bigger question around I think there's a sort of shift. For example, on the trust debate. [...]. Now, I think this kind of exponential effect of generative AI where you are going to have increasing difficulty identifying authenticity, for example. Definitely something in some ways, that's beautiful because it means that anyone can create any kind of image and that's so wonderful and creative and you don't know quite who's created it, where it came from and stuff. In a way, that's beautiful. Now, my point with that is not so much that there's going to be the direct deep fakes like the Pope one. I think it just proves the opposite effect: that it had no impact directly. No one really thought it was true. And if they did... [...]

When there's a gray area a human has to then look at the tech. [We have a] false expectation that the technology can fact-check everything. [But] as we all know, as journalists, there's facts, and then there's "fact-ey" things. There's things in context and there's grades of factual-ness and so on. And that won't go away. So I suppose it's not like a worry but in some ways, I think this makes us more intelligent actually, thinking that things aren't straightforward. So when people say, oh, there's this decline in trust, for example, in media, I'm less worried about that. To be honest, you know, I'm kind of worried if people don't want to get towards some kind of truth and they don't use journalism or journalism doesn't help. But all the evidence is the opposite, that people do still use journalism to try at least to get to something that is helpful for them as a version of reality. That's a bit philosophical but yeah. Whether there is an AI or not, it actually has more to do with news literacy. Yeah, that is a early childhood to adolescence type of building block of building This technology makes it more blurry. No doubt about it.

[..] But on the other hand, somebody was saying "this is going to be a real boon for Wikipedia" for example. Wikipedia generally speaking is a pretty straight, consensual by its very nature, it tends to drift.

Ariane Bernard 34:45

It's regression to the mean...

Charlie Beckett (LSE, Polis - JournalismAI) 34:46

A lot of ChatGPT regresses to the mean, it's quite broad, it's quite safe. Yeah, because that's what works best. So in that sense, you know, I think that there's a countervailing thing to the fake news and extremism is that this the algorithms will tend towards giving you quite a safe conservative version of reality and it may be slightly what's the word Panglossian is that I kind of feel with all this technology.

But this technology, it will put a premium, in a financial sense, as well as ethical sense, a premium on human creativity. If the machine can do all this stuff, and I think there is a big warning sign here to journalism, to the kind of routine commodity journalism.

If these machines can do it, then why go to a news organization to tell me what time of day it is, or what the weather is, and what's just happened in New York? If it can do all that, if it can answer so many questions for people, you know, 'is my flight on time?' 'What's happening with this company?' If this machine can do it, then why do we need journalism? Well, we need journalism for the same reasons as we always did: which is about judgment, it's about reporting. It's going out into the real world, it is asking better questions. It's coming up with clever ways to engage people and get people to think about things. It's about having expertise and experience of things to say, "Oh, well, this happened before. Perhaps this is different or not". So all those things I become your Unique Selling Proposition as a journalist. So you're going to spend more time doing that. If you're going to stand out, you know, as far as being as nice things like reliable and accurate and informed and things like that. So I think increasingly the journalists will be really there's another reason why we've got to get across all this stuff. You need to do that basic work. And oftentimes when I wasn't sure that you do that basic work, and that was the end product for your day. Do that now so you come up with something else?

Ariane Bernard 37:10

But the one thing and I guess this is where these things are narrow AI rather than general AI, but where [...] we use emotions which are of course, unknown to AI and the way we filter our lived experience through values which are also unknown to AI. The AI might see our values without realizing that that's what they are because it gets text and content from us in which our values are encoded, but those things are the reason why we anticipate things a certain way. And that might be really difficult and why journalists therefore think 'this is the story and this isn't' and 'this is an angle, and this isn't' in a way [...]. But these things are not statistically observable. They're not in our text, but they drive our decisions. [...] I was thinking well, no one got into journalism right to write 10 different versions of a headline in order to be tested. No one. But we all wanted in some ways, even if we wanted to contribute to the business of news, to think that we're helping shape society around some of the unknowns of being human in the world. And those aren't things where the computer is anywhere near [competence]

Charlie Beckett (LSE, Polis - JournalismAI) 39:59

Yeah, and I think most people working with this. I think that's one of the kind of false goals. I don't know if I don't know enough about the researchers goals. But if the goal is to do what you said so eloquently, you know, to understand values to understand...

Ariane Bernard 40:19

That's not the goal. They don't have that goal.

Charlie Beckett (LSE, Polis - JournalismAI) 40:21

Yeah, well, then that's great. And I think that one of the mistakes over the last couple of months, or you know, it's been fun, but journalists saying 'hey, it doesn't have any morality', 'hey, I can turn up racist'. Of course it can, that's not the point of this bloody thing. I suspect that certainly in the short term anyway, that the real use of this will be more along the Microsoft line. It's basically is a really clever version of Mr. CLIPPY you know,

There's going to be this thing where suddenly it's integrated into all our workflows, quite seamlessly. Nick Diakopoulos did a nice blog post about possible use cases and it was things like, for brainstorming and so on. ... And other people i've talked to, obviously, they're finding out things they can't do, or doesn't do very well. But I talked to a video editor and 'they say, oh, yeah, it saved me 10 minutes throughout our edit, because it can do this thing where it just puts the text straight onto the right place'. You know? So it's often those little things and you talk to people you already integrating it in their workflow and often quite mundane ways. And of course, somebody will be super clever just as people use AI for you know, major data investigations. Each bank leaks stories and things. Which couldn't have been done without it. And certainly on the visual, I think that's incredible. I spent a lot of time earlier this week talking to audio people. And that's amazing.

And when they get it, they get it you know, I think the most exciting things is obviously the coding bit this the way they can, but also the idea that both with audio and images that it can start to create its own metadata. Obviously, there are worries around that, you know about how discriminatory or biased it is or inaccurate. But just generally that it's able to do it is, as you know, is a huge stepping stone. [...]