Transcript for ELC 092
Beyond Content Generation: A More Ambitious Vision for AI
A Conversation with Phillipa Hardman
The Elearning Coach Podcast

Hello, learning people. Welcome to Episode 92 of The Elearning Coach Podcast.

How can AI make you better at what you do, not just faster? In this episode, I speak with Phil Hardman, PhD, an AI and learning specialist with more than 20 years of experience designing and researching learning experiences across higher education, corporate learning, and large-scale digital platforms. She has developed an AI-powered learning design system grounded in learning science, leads a popular AI & Learning Design Bootcamp, and writes an excellent newsletter covering AI and learning.

We begin by discussing how to use AI for research and how to evaluate the quality of studies. Then we explore learner-facing AI, including tutors, coaching, and feedback, before turning to practical ways AI can help learning professionals innovate. We wrap up by discussing the skills that will matter most as AI continues to transform our profession.

Connie: Hello, Phil. Welcome to the eLearning Coach Podcast.

Phil: Thank you, Connie. It's so great to be here.

Connie: I'd like to talk about research first. So, when you come across a new study, how do you decide whether it's solid enough to share with your audience?

Phil: This, I guess, is one of the benefits of my having an academic background is that we spent a lot of time trying to answer this question.

It's really complicated. But in a very practical way, I suppose what I do is I, I first and foremost look at where it's come from, and I try to avoid sources that are effectively trying to sell something because--always biased. And so, what I find, therefore, is that I focus in more on academic research or, or broad industry studies. The longer period of time, the better.

I think it's really important to acknowledge that some research that comes out is, by its definition, just an observation or it's just a hypothesis. Actually, there were no findings. It's not trying to claim that we know X, Y, Z. It's saying, "This looks interesting. We think there should be more research in this space, but actually, you know...

So I also look at the limitations of the research and I think, you know, the, the sum of all of that is that I tend to prioritize, as I say, academic research, which tends to be more broad brush and by definition to take into lots of different considerations and, and like a big backstory because, you know, the literature review is a key part of that.

And I also find myself looking really for the principles at the heart of the research. What is this actually telling us? It might be about AI, but actually what's it telling us about some of the learning principles that we know are interesting and validated? So that might be feedback or something like that.

So yeah, I think just being aware of who's written it for sure and what its purpose is, whether it's just a briefing and a hypothesis or actually a set of findings with controlled studies. And then from there, just make a judgment as to like what's just interesting and worth keeping an eye on and what should we share out to the world.

I think it's challenging, as well, because I think a lot of the stuff that I read that is emerging is the most interesting stuff to share. So sometimes I share stuff. There's been a lot of stuff recently about, for example, cognitive offloading. So, this idea that, you know, if our learners learn with AI, then actually they learn less.

Is AI helping them to learn or actually de-skilling them? And it's a very emerging debate and there's lots of mixed messages in the research. But still, it's interesting to report on that too. So sometimes I also gather together bits of research from different pockets like that and just comment as openly as I can and as transparently as I can about what it tells us and what it doesn't tell us yet.

So, that's a slightly different angle as well rather than saying, "This is what's happening." It's saying, "Oh, this is what's emerging." And it's, you know, but that is often where a lot of the, the juicy bits are.

Connie: That's true. How do you deal with the fact that so much of the research in AI comes from the world of education or a good amount of it does? So how do you translate that? How do you know whether you should translate it? Like I know when I read something, if it's about first and second graders, it might be kind of interesting, but I'm not going to say, "Yeah, this is true for adults." But then there's a certain point, high school, college, you know, it's a different environment, but I never know quite when to translate it to L&D.

So how do you make that switch?

Phil: Yeah. And it is, it is tricky. There's so much interesting research like you say about like K-12 that I read. It's like, ooh, how transferable is this to higher education and/or the workplace in particular, which feels even further away?

The way that I do it is to look for, again, the principles, I suppose.

So, I'm trying to think of a concrete example, but one concrete example I think was some research that has been done in the university setting, for example, about intelligent tutoring systems like AI coaches effectively. And if you look into some of the detail, there are things there which are definitely repeatable and applicable to workplace training.

So, things, for example, about how to design that coach well versus not that depend not on the university setting, but on like how humans engage with AI. A tutor that just keeps asking questions back and refuses to give you an answer leads to lower engagement because that's frustrating for human beings.

And that's true in both settings. So, there's just that sort of trying to get it down to what, what's the crux of the point here? Where are the things which are transferable, which as you say, aren't dependent on context, but actually tell us something about the value or maybe the danger, the risk of, AI, for example, in the workplace as well.

But I think it is a real gap, and I do wonder sometimes why there isn't more research done specifically on the corporate setting. I think it's because, you know, most academic research tends to be more self-interested, and they have more access to academics in schools. It's easier, it's faster, it's more relevant.

I think it’s’ an interesting opportunity, and this is what I try to do with corporates, and I do try to be quite systematic in some of the, the testing that I do. Sometimes transferring what we've learned from education over and seeing if it does transfer in order to, yeah, try to merge those two worlds together.

But it can be challenging, can't it, to have an evidence-based approach when a lot of that evidence is based in a totally different world from the one we might be working in if it's corporate L&D.

Connie: I find it challenging. I mean, a lot of it is common sense. How to design a tutor, it's how to design a tutor, most likely.

And most likely, adults are going to get frustrated. We know that from elearning where they may not ever give you the answer and the user has to click retry, retry, retry. After a while, they're just gonna shut it down. So yeah, there's a

certain amount of subjective human experience in there that helps us to make the transfer.

Phil: I think so. And then there's, I guess the flip side as well is that there are some things that we know to be transferrable because they've been tested and proven to transfer to different contexts. So, one for example, which I always talk about, and I'll probably talk about multiple times in this chat is feedback.

Often people say to me, "What's the one superpower? What's the one thing that I should change about my whatever design or delivery, whatever, to improve outcomes?" And the answer is always feedback. Right. and so that we know that there's been enough research into that as, as something which seems to drive humans of different ages and in different contexts to both engage more and to learn more as like that's a good example of, yes, I would trust that remains true in all contexts because, it's been proven to, I guess, be portable across different contexts and times.

So, I totally agree. Like a lot of it is stuff that we already know to be true. And then some of what the research tells us is powerful. Like there's enough weight of evidence to know that it will work.

Connie: Those are good guidelines. I like that. You know, a lot of us are using AI research tools, Consensus, Undermine, Illicit, all of that to maybe find research, summarize research.

Of course, we always have to check. We know that. But what do you think the AI research tools do well and where do they fall short?

Phil: I mean, I use these tools a lot. I'm also often asked, like, how has AI changed your work, my work most? And I think it's this. I guess I've always had the privilege of being a researcher.

I've always had my, my library card from a university so I can get at the research, but of course it's challenging. It's time-consuming. It's, it's like a whole job in itself. And so, I love both for me and for learning designers more generally that we now have more access than ever to the research. It feels a little bit like the paywall's been taken down, the physical walls have been taken down, and it feels like an important moment really that more of us have access to ways of validating that actually what we do is right.

I think often as learning designers, and there's so many different directions that people come from, and I think a lot of people have imposter syndrome. A lot of people feel like they design from their gut or that it's somehow creative but not

necessarily robust. And I see a lot of people coming through the bootcamp, once they start to work with these research tools, they're really starting to feel like, "Ah, I do know what I'm doing. I am more confident." And also, I see their relationships with their stakeholders changing for the better. And so, it becomes like, "Well, Phil says this, but the professor says this, so we're going with what the professor says." And it becomes more of a conversation about, "But the research says this, that if we do X, we'll get this.

And if we do Y, we'll get this. " And it's helping people to become more influential, I suppose. And so, headline, I guess, is that I find this a very positive move shift both for us, like, as an industry and as individuals. There are, however, shortcomings, as ever. one thing I often say on the bootcamp is that, like, wandering into an AI research tool is like wandering into a library.

So, you know, usually, unless that's, you know, some crazy library like the British Library or Cambridge University, and even then, basically there's not a copy of everything. And everything that there is a copy of has come through a very stringent and biased process of being selected or not being published or not, being identified as important or not.

And so, one thing I love to do is to ask these different research tools exactly the same research question and identify the, the bias of all of them because they all have biases. Of course, the information within them is as reliable as information can be. They are published, they're peer reviewed, whatever.

It is a selection. And what you find is that actually, yeah, like consensus will give you some overlap. It's a bit of a ven diagram situation. So, like, there's a lot of overlap, particularly if there's been, like, a lot of very well-cited research. You get overlap between, say, two, three different research tools, but there's also always things on the edges which other research tools miss.

And as I say, one of my favorite pastimes is to basically open up an LLM. So, let's say I open up Claude, and I'll give it one paper from, say, consensus and say, "This is what my research assistant said was blah, blah, blah." And then I'll give it a second one and I'll ask it to debate, like, which is right and what, why are they different?

And, and what you soon realize is that, yeah, they are a sample. So, I think one of the risks is that we, I guess, first of all, don't acknowledge that proven knowledge that is published subjective. You know, I know that everybody is very, very concerned and rightly so about the biases of AI. So, you know, it's sexist, it's racist, it's all of those things because we are and we built it.

But so too is written literature and I don't think this is a new problem. So, I think making sure that people understand that the research itself has biases and then realizing that there are gaps in the different collections and learning how to triangulate them. I guess they're the methods I use myself and the things I teach on the bootcamp to make sure that we're getting value, but as you say, also remaining the very critical, human in the loop who's doing that validation and checking and cross-checking across different tools.

Connie: Do you have some favorite tools that perhaps I haven't mentioned? Maybe I don't know about?

Phil: The one that I recommend that people use on the bootcamp is always Perplexity in Academic mode, because I just find that has an amount of simplicity and power that's really good if it's the first time that you're using it.

But in terms of triangulation, I tend to, to go to that plus Consensus. And then I'm not sure if you mentioned this, Connie, maybe you did, but I also have a look at a tool called Elicit. And they're my, like, three go-to's because I've found that they each have a slightly different perspective on what research is.

And so, I find if I triangulate those three, I get some, decent results. There are some other interesting tools out there. I think actually finding research papers, there's, you know, Google Scholars, really good semantic scholar, but sometimes it just gives you the reference, and you have to then go and find the paper.

And if you don't have the login, then you're stuck. What's great about those other tools is that they give you direct access to the material. I think there's a really interesting tool called Storm, which was built by academics at Stanford. The idea being that you can write a personal Wikipedia page on anything just using very reliable sources.

And it is pretty powerful, but it's still, it's an emerging tool. But it has some interesting features like have it generate you as a research summary or generate you a debate as if like experts sitting around a table and debating a topic.

Connie: Good one to know.

Phil: Yeah. And there's a growing number out there and it's really interesting that they are, so popular. But yeah, all with the caveat, of course, they are incomplete and there's a lot still behind paywalls. You know, journals haven't stopped charging. mostly the research that we're getting through these tools is the stuff that's open source. But yeah, certainly better than zero. And zero is where most of us started.

So, I think this is, great. And I would say to anybody listening, you know, it's probably one of the most powerful things you can do aside from feedback. Just go to the research and think about what are the things that I do a lot. Like I, I write a lot of learning objectives. I have to decide a lot on whether something is a course or a job aid or whatever.

You know, I have to decide often what mode of delivery I'm going to use. For these big decisions that we make often, you can very quickly get together a beautiful little cheat sheet, based in the research. and it really does bolster your confidence and also ultimately your impact. So, I think that's, you know, it's a bit of a power move for people who are just getting into research tools to just, yeah, just to find a few of the big decisions that you make in line with the research.

And you'll also be surprised how much of what you're already doing is best practice.

Connie: That's true. And I love that as a power move, especially when people insist on things being a certain way who have no background or true understanding of how people learn. So that's great.

Phil: I do believe that AI and it's having a very mixed effect overall, I think, on our, on our jobs and it's going in lots of different directions. But certainly, in some ways it is helping us to professionalize. And I think this is one of those ways. I think. And it's happened before in other industries there are times when you don't have access to data about what great practice looks like.

It really holds back an industry. But now the concept that we do and, and that more people than ever can access that information is, I think, game changing.

Connie: And just one little addition, I think it's so important to research in user experience design, product design, and many other fields, because our field is just one giant Venn diagram. And the other fields contribute so much that I just think it's important to research everywhere. Anything that's slightly related.

Phil: And it makes you realize, as you say, how many jobs we actually have. The answer is definitely too many. I'm not sure what the number is.

Connie: We better move on to learner-facing AI. I wanted to ask you, what are some of the most effective uses of learner-facing AI for the workplace that you've seen?

Phil: Let me just share very quickly an experiment that I do. Cause I think actually we're not experimenting much. So, most of the experimentation that is

happening at the moment on the ground, especially in corporate M&B, is very much focused on almost like the backend.

So, we're using AI to make the process of designing faster, but not as much innovating the other side of things, which is the experience itself. So just picture a world where you have more e-learning courses. the experience for the user is at best the same and maybe even worse.

But for us, it feels great. Like, oh, I'm so much faster. So, I've been doing a little bit of experimentation with this on my bootcamp, which I run every month. And as you know, like, it means I have access to learner reactions to exactly this. And so, I experiment with a couple of things. So, I give my cohort, they have access to me, and they have access to each other, and they have access to materials.

But within this mix, I very intentionally, I tell them at the beginning. They also have access to AI-generated content. I made it, but AI made it. You know, I might use NotebookLM and give them a slide deck, an infographic, flashcards, an audio summary. And, and part of my experiment here is like, what's in it for them?

Like in it for me is, "Oh, I'm so generous. It's great." And they get to choose and we often sell this idea of powering learners through choice. But actually, on the receiving end of all of this, I look at the stats and barely any of them look at any of this stuff, which in itself is interesting when we talk about this.

You know, partly it's because they have a choice. And there's also a lot of materials that I shared that are marked with a little DNA thing, which is like, well, this was written by human Phil. And they much prefer that even if it's just text to an infographic that looks sort of familiar and they know it's me plus AI, and they're a bit like noise.

And that's what I hear a lot is like it's too much information. And so, the learner experience is becoming a little bit overwhelming, I think, for the learners. But the other thing that I do on the bootcamp, not just content, but I give them also AI support. What I mean by that is they have like an AI version of me finish.

And so, if I'm asleep, whatever, I'm not available. They can at any point in time ask me AI me a question and get feedback on a project that they're working on or whatever it is. And I put this data out recently on a post. I think it's 98% close to of the students who come on my bootcamp engage with that.

They really like having a safe space. That's the words they often use. It's like, I'm not going to embarrass myself in front of others or real Phil. I'm going to ask what

I consider to be stupid questions to this AI, and I know they seem to be reassured by the fact that I built it. They said they would feel very different if it was like a generic off-the-shelf AI, but because they know who built it and because they have a real trust for that human, they actually really value both that safety as they see it.

And another element, which is just endless amounts of feedback. So, I've taught this little AI to give really good feedback and they love it. Like more is more. And sometimes at the end of the bootcamp, it's the favorite thing ever.

Connie: Is it a custom GPT that you're talking about?

Phil: I use Claude Project. And I also sometimes use NotebookLM for people who can't access that.

So, they don't like the freedom of all the content. It's not freedom to them. That's like, "Well, noise. I don't wanna have to be your editor. Thank you very much. And also, I'm not that interested in watching your AI avatar. I have an AI avatar as well. Tell me about the course. I'd rather just read the bullet points and move on.

And so, and I think this is important because I think it tells us that actually there's a lot of the stuff that we're doing at the moment with AI is maybe great for us and certainly great for the leadership who are now paying, I don't know what the numbers are, but like one pound for a video, not 1,000 pounds making that up, but similar.

But for the learners, they're like, "No, please not another AI avatar. Please not another infographic." But what they do love, and I think this is the key example—is a well-built tutor study mate type setup where they can get safe and reliable support and feedback in a way that just was not possible when it was just the humans.

That's the most exciting thing that I've learned so far is that it seems to be much less about the content and more about the support that AI is giving.

Connie: Yeah, that makes sense.

Phil: Things which have more traction than say an AI video or whatever the content. Things which make the learning experience more active.

So, people playing with ChatGPT now and Claude and Copilot can help you to build branching scenarios. They're not new, but they were very expensive and

slow to build. So, you know, they were quite unusual historically. And now we're seeing more active learning and, and there's, there's some traction there.

But I think above and beyond everything is this concept of an AI tutor with an identifiable maker. Somebody's brain. Exactly. Somebody's brain. Yeah, that's perfect. And I think that's interesting because also there's loads of stuff we don't know there. Like how would you feel if you never got access to me, the human me, you only got the AI tutor I built different.

Something about having the choice and the power matters. But maybe it wouldn't be different if they weren't paying. Maybe if you could have free access to the tutor and never ever meet real human tutor, that would be okay. So, there's still a lot that we don't know, but generally it's much less about content.

The experience humans seem to like is access to an expert to give me the teaching, the support, whatever I need. And this is backed up as well, Connie, by data at OpenAI. So, I work with them quite closely. But they've published it widely now. The number one use case of ChatGPT is people trying to learn stuff.

Like this is fascinating. You know, people love to learn. And it seems that they actually quite like learning by just having a bit of a back and forth. And if that back and forth is with a trusted source, then great, even better. So, I think there's lots of data to suggest that this is where the learning experience should go.

More emphasis, I think, on AI coaching, let's say, versus AI content. But also, there's just so much that we don't know. Most of the, the effort that's going into AI use at the moment, as I say, is in, in the backend to make the production process faster rather than thinking about what's the experience of the learner.

But yeah, my data overwhelmingly every month is nobody's looking at your AI content, Phil. nobody cares that they are very much chatting to your AI teaching bot thing.

Connie: I think that you've got like 80% of it. You're right. People want support. They want feedback. In a workplace setting, you may not be able to have it be based on an individual's expertise.

It may just be very general four SMEs in the room building it.

Phil: Yeah.

Connie: But I also think two things I've learned from having a community is one, that people are insanely busy. And so, whatever's going to get them the result they want the quickest is where they're going to go. The other thing is that

Notebook LLM, their infographics, unless you go round and round and round with it, are just so overwhelming.

There are so many visual cues. They're frequently impossible to understand. And so that's a problem. And to get it as simple as you want, you know, you have to pull your hair out.

Phil: It's such a great point. And I spent a good couple of days doing a bit of research on this because, because learning designers seem to love NotebookLM.

And so, the hypothesis I had was that maybe with the right input, I might be able to get it to follow Maya's multimedia principles and actually so that the content that we were producing and giving to learners was actually at least in some ways optimized for learning. Absolutely not possible. Like the system prompt is so powerful I could not break it.

However hard I tried, it defaults to its defaults, not to mine. And so, I think, yeah, there is still, as you say, a lot of frustration and a lot of not time saved. AI is always sold to us as time saved, you know, put your feet up, have a coffee. In fact, it's just time, repurposed. And in some cases, time expanded because, yeah, the job is not the production of the stuff, but it is the editing of the stuff, be that text or infographic or whatever.

So yeah, I think there's a lot of hair pulling out going on.

Connie: Eventually, I just shut it down and go to Illustrator and do it.

Phil: Old school.

Connie: Unfortunately, yes. So, from what you've seen, what do you think is the best way for someone to design and deploy a learner-facing tutor, coach, customized chatbot, or something similar?

Phil: So, you start with a hypothesis of, first of all, the KPI is not get as much AI out as possible. It's more like increased value with AI. So, in the learning context, that might be we want more people to be selling more things like, you know, behaviorally we want them to be doing something different by this amount, by this time.

And the former hypothesis about if and how AI can help. What I'm seeing in, in more innovative, corporates are, is, is the creation usually just in a sort of research and development setting, just to see, but of like little pods of people

who own problems. So, they might, the business problem might be we're not selling enough.

And the pod of people will include a learning expert. It will also include a, you know, the business partner or the subject matter expert, an engineer, a data person. And together they form hypotheses about if and how AI might be able to help solve this problem and they build a prototype. But it does, doesn't go anywhere near the world until it's proven to be, to work.

And so, the testing happens in controlled environments. So, you know, whatever it might be like. We have a hypothesis that if we give everybody in learning design access to NotebookLM, X, Y, Z will happen. Test it with a sample. And then there's a decision point, which is, are we gona shelve this or are we gonna scale this?

And then it goes out. And that much more sort of agile testing scrappy approach is the version of deployment that's working best. And it's driven by KPIs about proven impact and data and the shelve and scale decision rather than get it in everyone's hands, get people adopting it as quickly as possible.

And if you're interested, I think LinkedIn is the, like, the archetypal example of the pod type structure. And they experimented with that early. And I mean, it's not new. I think it's very familiar to people in engineering and, and product, but it's interesting how it's also impacting different teams, including learning teams.

And it's definitely, in terms of the methods of deployment it's much more successful than the more typical approach of buy it, give everyone access to it, see what happens. I'm not sure I know what you mean. What is LinkedIn doing? They have this vision that everybody that works at LinkedIn, regardless of team, is a, what they refer to as like a full stack builder.

And they work either individually or in pods but empowered by AI. For example, as a learning designer, AI does give you superpowers at every stage of our workflow. So empowered by AI, you can do deep analysis on data and work out what's the problem we're trying to solve. You know, what actually is the root cause here of this problem and therefore what's the best solution?

And then with the help of AI, you can prototype that thing. And then with the help of AI, you can actually test it with real learners. Yeah. And then you decide whether it's a shelf or scale it. So, they've changed their, they've actually dr - quite dramatically changed the structure of the organization away from specialized teams towards these full stack builder roles.

And that CPO has a very ambitious vision, which is that we should delegate everything to AI aside from judgment. So, the human should always be making the decisions, but AI does all of the data gathering, the presenting of the stuff that, you know, it's, it's like having a, everyone has a, a sidekick or an, an assistant to help them to do the work.

And it's, yeah, it's changed the structure quite a lot. Even the concept of team and department is totally switched now.

Connie: I'll have to read more about that. I wanted to turn to the topic that most people are using AI for, which is how to make our work more effective. I think root cause analysis, which you mentioned previously, should just be one of the top tools in a learning designer's toolbox because we are so often asked to produce training for which training is not the appropriate solution.

What are some of the best ways to use AI when conducting a root cause analysis?

Phil: I agree. I think this is where any project either succeeds or fails to some extent, or at least it can fail at the very beginning. Yeah. Just assume the problem is X and the solution is usually course, then the whole thing might be a waste of time, but if the problem isn't well defined.

And so, it feels to me very powerful. And, and if you think about why we've not done it in the past, I think there's lots of reasons we didn't have access to data. We didn't have it's yet another one of the examples of the 500 jobs that we do. And yeah, actually again, to go back to the bootcamp, a lot of people in their first week when they do their analysis project design AI tools that help them to do really great needs analysis problem definition in ways that they just couldn't do before, I think primarily because it just took too long.

There are still challenges, and this is always the case with AI, that data, you need the data to be able to make the decision. But just to give you a couple of concrete examples, even a simple GPT, somebody, last month built a simple GPT to give to people who made requests for training in a corporate environment, which walks them through the five whys.

And so, what they get, as in the, the learning designer gets, is not a vague statement about, you know, we need to sell more things. It's a much clearer problem statement for them to then more clearly make some decisions and take some actions.

Similarly, people have done that with the Mager-Pipe (the Mager and Pipe Model). It's essentially like a decision tree. So, you know, if the problem is X flavor, then it's likely an information problem and therefore not a learning solution. It's an, it's a, it's a PDF or it's an email. And so, yeah, I've seen lots of examples of people creating either prompts for themselves to turn vague requests into well-formed problem statements and therefore like solution suggestions.

Or as I say, some people build GPTs which go out and, you know, can be shared with the subject matter expert or whoever in order to get to something more valuable as an, as an input. I often talk on the bootcamp about these three buckets. So, there's one bucket that's just like, do what you already do faster.

Yeah, I don't know. For needs analysis, I suppose it would be just taking the request and quickly turning that into a brief, not questioning it, but just quickly making a document. For the better angle it is actually working with AI to apply some of these methods like the five whys to actually get to something more substantial.

But then at the transformative end, it's possible to do things like build agents, which constantly monitor things like our Slack channels or Teams channels. and I've actually set this up as an experiment and it does work quite well. so that instead of having somebody request a training, the AI helps me to decipher actually what's causing problems on the ground.

What are people asking? What's stopping people from doing their jobs? So, it's not the opinion of somebody or, you know, something that's come off the back of a new priority. It's actually what does the data tell us about where people need help to do their job better? And you can set these channels up for courses as well so it could be like, "Oh, people are struggling with this or that.

But the most powerful versions are, you know, actual readings of the real work that tell us about where people are struggling to take out this whole back and forth about like, is this the thing we should be designing? The data speaks for itself, and we can then, you know, hit the ground running and, and design from that.

So there, there's some concrete examples of the sorts of ways in which we can just become both faster and better at gathering information about the problem and then of course with the, with the aim of then turning that into the right solution.

Connie: I did want to ask you how an organization can set up some type of continuous feedback. Can you get a little more concrete about that? For example, how in your example are they getting that data from the ground?

Phil: Really interesting. So, there's different levels of ambition. In my very basic prototype, you know, some of the clients I work with, it's literally a reading of a Slack channel or a Teams channel, which is set up to capture what people are saying about what they're finding difficult in their job.

So, there might be a channel for sales and people are encouraged to just share like what's tripping in, what, what's this, what's this? And then it gathers that data themes it, and then that becomes the input for my prioritization. So, it's pretty simple. There are then much more advanced complex ways of doing this.

And I mean, they sound, I guess, quite transformative and quite futuristic, but actually they are happening. So, a couple of examples, one is like in a sales organization, imagine a world in which you can use AI to track sales calls and monitor from that data where the closing of that call missed X, Y, Z, which means the likelihood of them actually selling is lower.

And again, effectively what we would then get as a learning designer is a data report on like the most powerful thing we could do for the sales team is to help everybody to be better at closing a sales call the last two minutes. And then they're like, great. I'll find the subject matter expert to do that.

So, there's some monitoring of like actual performance on the ground. I see it across lots of different places. It's like the same vision is emerging quite consistently and it's not new either, but it's this concept of learning in the flow of work. So, no more stop and learn courses. And I mean, basically in this vision, AI is like Clippie, the paperclip that people famously hate.

So, you know, it takes us back to that, but how does everyone feel about this question? But we get support to complete a task in the flow of the task. So, imagine a shift from a course, stop and learn course to a, basically a personalized coach who's like, yeah, Phil's never been good at decks. Let's help her do X, Y, Z.

But also, in the course of that interaction with AI and the flow of work, we're also being stretched. So, it's deliberately building in a little bit of friction and a little bit of stretch so that next time we're better. And again, that's just all monitoring. It's like there's then a record somewhere behind the scenes of how well, you know, where I started on writing decks and how well I do now and what kind of feedback the AI gives me in this kind of thing.

And one thing I notice about this is like it takes a lot of data. Again, what's in it for the learner, for the employee? I'm thinking quite deep paranoia about the amount of data that you have and the amount of monitoring is what it feels like - Right. That is required to achieve this. And so again, I think it's really important to get very excited about what's in this for us.

Yeah, what's the human machine interaction at the other end? And what's the overall impact on what it is that we're trying to do, which is to engage these people and help them grow? And yeah. So, there are interesting things going on, but I think still a lot to learn particularly about that. How does it feel?

Connie: That was one of the things that came to mind was all the monitoring and how people will feel about that. If there is some way to anonymize the data like you do with healthcare, I think that would make people feel a lot better.

Phil: Yeah. So, we'll see where that goes, but there is this big converging vision, and I see experiments happening on part of them. We'll see if this happens. But as I say, the vision isn't new. It's just we now have the tools to do it. And maybe actually we need to get a little bit more ambitious in our vision, before we start building new systems.

Connie: The last thing I wanted to ask you, and I know you only have a minute left is people are concerned about losing their jobs. What do you recommend for people to future proof their careers in terms of AI?

Phil: Yeah, and they definitely are fearful. I think the, the key skills are not complicated. They are about working with AI really well. If you can work with AI in a way that makes you not just faster, but also get better results that you have evidence, then you are a purple squirrel by which I mean, there's lots of people looking for you out in the world.

And what that means in practice is, you know, there's these two basic things that I teach on the bootcamp if you look at the kernel of it, the heart of it. One is, first of all, how to decide what to build and what not to build. That's like actually a new discipline because when you can build anything, it's like don't, don't though. Don't build at all. So, like problem definition and prioritization, key skills.

And then the other things which are much more technical, is like how to write a really good system instruction and how to build a brilliant knowledge base so that instead of just defaulting to AI's default behaviors, which are biased and lazy and repetitive and all of these bad things, you know how to actually turn it into something that's quite magical and quite expert.

And all of that requires very deep expertise of what you do, how you do it. Effectively AI needs you in order to be powerful. So, I think go deep on that practice. Go, go to the research tools. Understand if you were training AI like an apprentice, what instructions would you give it? What resources? And that's what I'm seeing set people apart.

It's not the ability to, you know, use Synthesia. It's the ability to shape the quality of what AI produces.

Connie: Love that. Get better results. Okay. Thank you so much. It's wonderful to speak with you.

Phil: You too. I feel like we could like do this every week. There's so much to talk about.

Connie: The one idea that stood out to me, it's that AI has the potential to do much more than help us create content faster. Throughout our conversation, Phil challenged us to think about ways to innovate with the creative superpowers that AI provides. That's a more ambitious vision than simply generating courses more quickly.