MFS transcript 6

Continuity: How To Grow A Human: My Frankenstein Summer with Dr. Philip Ball, episode six, Living With Our Creator.

Dr Philip Ball: I really enjoyed the cafes in Boston,

the city, once renowned for its trading tea, peppermint, dirty chai. 1%. These days, Boston is better known for its coffee, but you have to wonder how much longer it'll be before those baristas are all. By an interactive terminal that takes your order and makes it just how you want it.

Computer speaker: Hello Dr Ball, can I interest you in a cappucino today?

Dr Philip Ball: Might it be necessary to automate even more tasks like this to reduce human contact? The robotic barista could be pretty simple for today's AI, but these technologies are already finding roles in far more complicated tasks. Like the healthcare diagnoses being trialed by IBM's Watson health at the labs I visited in the last.

Of course sometimes what we want from doctors and nurses. It's the human interaction that after all is why a good barrister can still add to the atmosphere of a cafe, but the COVID-19 pandemic has forced us to face the possibility that human interaction might in the future be precisely what we'll need to avoid.

Can AI ever be a real substitute for. That's the issue I'm going to investigate in this final episode, the interactions between people and machines and how society might change as these become increasingly a part of our lives. For many people, the kinds of fears raised by Mary Shelley's Frankenstein of technology run a mock are today invoked by developments like AI, which in the popular imagination become very much like the creature that will turn on its creator.

Passage reader: But it's true. I am a rich murdered, the lovely and helpless. I have diverted my creator to that irredeemable ruin. He lies white and cold and death.

I look on the hands, which executed the deed. Which do you mark your notion of it was conceived them long for a moment when these hands will be tomorrow or I'll be marching, motion reform, more theater or more,

Dr Philip Ball: but is that how it'll really be, or will we forge a more amenable, more productive relationship?

I knew just the person to ask. His name is . And until recently he was working at the Massachusetts Institute of technology. MIT. Iyad is very much a rising star. He's Syrian by birth. And I first came across him in 2012, when he was working in Abu Dhabi at the Mazda Institute of science and technology at MIT, he's been looking at the intersection between the computer sciences and the social science.

In particular, the interface of humans and machines issues such as collective intelligence, how people exchange information by machines, how they interact with machines and the general social aspects of artificial intelligence.

Iyad Rahwen: Okay, well maybe, maybe. Talk about the, sort of how I got to here first, quickly. And I did my PhD in AI, but not the cool type, you know, the cool type. It turned out to be machine learning and deep learning. I think the other type, which is a symbolic reasoning and that sort of thing in multi-agency systems, um, and then sort of slowly drifted away from, uh, interesting, interested in AI and more interested in people.

So I was interested in stuff. Social media and how it facilitates new forms of cooperation or mobilization. And this is how I got really interested in things like the use of Twitter and social media to mobilize millions of people, to search for things or to search for individuals in remote cities. And you've covered kindly one of our

stories

Dr Philip Ball: in 2011 Iyad took part in a competition set by the us state department called the tag challenge. We came to use social media to locate people in remote cities using only their mugshots pictures of them released to the public.

Iyad Rahwen: I meant to hear you. We were meant to find people using their mugshots in remote cities. Can you use social media to do this? And I remember, um, and we won this, uh, this challenge from Dubai where we were. Uh, using these incentive structures to incentivize people, to look for the information, but also tend to recruit their friends to the cause.

And we wrote a paper and it was sort of still, you know, in, um, we're about to submit it. And then there was the Boston marathon bombing here.

Dr Philip Ball: But as it was unfolding, the bombing at the Boston marathon took place at a challenge. Like this became very real. As people tried to identify just from TV and camera footage, the perpetrators of that crime

Iyad Rahwen: and people started using social media to search for the perpetrators. This is in the two days when they were on the run.

Um, actually. You know, they ended up being on the campus and they actually killed one of the MIT police. Um, we released the paper at the time because we thought, you know, it can be. You know, you can find people through the crowd, but they're in this case, they're specifically hiding. So I was interested in these sort of processes, like what, what makes social networks able to, uh, buy online social networks, specifically able to mobilize and to do this, you need to, it's not really a computer science problem as much of a social science problem or a network science problem, or a problem of understanding incentives that change that mobilized people, or, um, why people do these things.

Um, how fast do they respond to each other's messages? So this is all behavioral. So it became so interested in the collective behavior

Dr Philip Ball: online with this background in socially mediated information technologies, I had found himself wondering how our own behaviors are likely to be affected when we hand over some of the decision-making to machines and AI and how we will interact with them as they become increasingly like entities that.

Iyad Rahwen: I felt like, well, I have a background in AI and I've been studying humans on social media, and now these things are going to merge. So cars are going to drive themselves on the road where other car, where people are driving and people are crossing the street. Um, there will be bots online sharing information, filtering information, manipulating information, potentially together with inference with people.

So again, there's more and more. Sort of systems where the combination of human behavior and AI behavior is going to interact. And I became fascinated by the source of questions that you can ask because that's a new kind of. And we've never seen this before. And I felt like, well, science also has perspective

on these things, will people who operate with machines, if you know, if they know their machines, will people, what people expect from machines?

Do they expect the same thing from a, from an AI car as they expect from a human. So we need to, I think, understand the human side. And how it interacts with these entities to be able to thrive with

Dr Philip Ball: them in the future. Very recently, I had became involved in a project looking into the ethics of self-driving vehicles and the question of how we're going to program them to determine what they will do in the case of an emergency, a situation that might cause an accident.

I asked him how he ended up pondering issues like this

Iyad Rahwen: philosopher. Uh, say we already figured all of these problems out since Thomas equines, right. Told us that we just gave you the answer done.

Dr Philip Ball: Thomas Aquinas. Well, it wasn't entirely what I expected, but this 13th century Italian priest was an astute, moral philosopher who laid down a number of ethical rules for judging morality and complex situations.

Even today, moral philosophers still pay heed to what is. Moral philosophy crops up. When you think about how to program driverless cars to deal with collisions, because it's related to a famous thought experiment called the trolley problem. The scenario is that there's a runaway trolley car and a track headed towards a group of people who will be killed by the collision, but you're standing by a lever that could divert the trolley onto another track where it would only hit one person.

The question is, do you pull the lever? Do you let the accident happen to several people or intentionally kill one. This used to be just an abstract philosophical problem, but driverless vehicles make it a very real dilemma.

Iyad Rahwen: And of course that leaves out the question of how we would implement. So Thomas Aquinas in the car, but, you know, as in terms of ethical principle, it's been figured out, but that's a normative principles that comes from a particular culture.

And guess what, not everybody subscribes to that culture, right? There are many other cultures in the world. The machinist has new capabilities that necessitate new answers. So for example, human may not know the outcome with such

precision in which case anything sort of is defensible, right. But maybe a machine can figure out outcomes with more certainty.

So if that is the. Can we still apply the same role or is a new role necessary if you're, if you're applying a rule in one situation, but then 10 million cars get the same rule instantly, what sort of outcome does that create? That's different from human society, where rules are sort of applied and implemented to a different degree by different people because you have different.

The different ethical standards or motivations

Dr Philip Ball: or so, and it sounds like there's nothing kind of preordained about that. I mean, maybe the temptation, I think you said was to imagine that we have moral frameworks and we just have to fit the machines into those. Whereas in fact, it seems clear that the new possibility you solve social media and machines generally.

Bring up new questions that we haven't considered before or alter our behavior in ways that we haven't considered before. So it's not as though we already have more answers and we know how we just need to know how to implement that.

Iyad Rahwen: I think it's a combination of an analytical exercise with also. Maybe a psychological exercise to see if people would trust car programmed, according to that standards, because they may say, let me say, no, I am okay with my doctor doing, you know, um, either sacrificing the mother or the child in, in a complicated birth.

Right? It's this, I think some of the cases it's always like when it's looked at, but I'm not comfortable with the machine doing that. Right. Making that call. Maybe I need a human being, uh, to be involved and so on. Right. So. I think we still don't know. You know, there's so many questions.

Dr Philip Ball: I mean, I guess that also comes down to the kind of reasoning that we can.

Machines to have, um, because confronted with a situation like that, and certainly with a doctor confronted with a situation like that, it's not as though he or she looks up in a table, you know what to do. And the answer is there. There's a balancing of so many things. Some of them not even articulated, but that are just normative in our society.

And we kind of accept that, that, you know, a doctor will do the best that they can under those situations. Whereas if you're deciding if you're actually going to program into a machine and algorithm to make that very, it's not obvious that it's the same process,

Iyad Rahwen: is it? I guess? Correct. And I think, I think also finding the right metaphors or finding, what is it, what is this problem that we're facing with the machine now?

How does it, which older problem does it map to. Sometimes kind of surprisingly, uh, difficult to anticipate ahead of time. And I think this is where science also helps. So to give you an example, um, there was around 20 13, 20 14, all these articles about how you should program a car, you know, to it, to behave in an imminent accident where there's unavoidable heart, lots of normative opinions about.

So we started by saying, well, uh, let's ask people and see what they think. Um, and maybe we can just provide a public of public opinion as a lens, you know, and as an additional input to this process, to this discussion. But then we asked people which we just thought we might ask people two questions, let's ask them, what do they think the car should do?

But also which car would they buy? And we instantly, you know, found this in hindsight, obvious. Conclusion, which is they want a car that minimizes harm, but they will never buy it. They want the car that only protects them keep. So I think this is a very simple experiment, right? And, but it taught us that AI ethics this in this case, it's not a question of what of normative finding the normative rule.

It's a question of cooperation. It's enforcing the outcome that people think. Normatively. Correct. And then we did the subsequent study, which asks people to decide what the car should do in an imminent accident. And then we S we had 40 million decisions from people worldwide, which allowed us to then not only find out that there's a disconnect between consumer preferences and citizen preferences, but also that there are cross-cultural differences in those preferences.

And, and again, I think that's another dimension that we, people who are interested in. Regulating AI and building AI need to take into account. So this is

how I see my role as using science to bring in new things that that should be taken into account to inform this discussion

Dr Philip Ball: know there are self-driving cars, but there are a long way.

I mean further, I think, than people often realize from being something that you can go out buy. Um, but it feels almost as though this is still a kind of, uh, an, uh, an experimental test case for the, kind of, for much more general reasoning. I mean, it almost sounds like, you know, what you say about people's responses to what they would choose.

Um, you know, as opposed to what I think is kind of right, January. It's kind of like a, uh, uh, an analogy, a metaphor for society. Generally, the people tend to say, yes, I want public services. Yes. I want good spending on libraries and schools, but I don't want to pay lots of taxes. Um, so we've lived with this kind of dilemma for, for a long time, but it's interesting that now.

It's technology's really forcing us to, and

Iyad Rahwen: then, you know, how we, I guess how we deal with it is also brings up new challenges, I suppose, because in the old days, you know, we, we, we dealt with those. The way we dealt with those dilemmas has had to do a lot with. Our understanding of human nature, but now with machines making decisions, you know, whether it's, you know, um, uh, decisions about policing or, or, uh, decisions related to selling products or price differentiation, all sorts of things, it's machines that are now learning those new behaviors.

And we, I guess we don't necessarily have a mental model of the machine, the way that we did with, with the human. As humans are kind of like us. We've been living with them for a long time. You know, we're living with each other, uh, for, for millennia and we've evolved a set of norms and expectations, but machines don't have those constraints.

So if two, given a certain, and I'm not talking about paperclip, you know, colonizing the universe to make paperclips scenarios, you know, it's a much more benign version of that scenario is you simply ask the, the machine to. Sell more insurance policies or sell more IVF services, right. And then machines could learn all kinds of sophisticated, unexpected, potentially seriously unethical strategies to achieve those goals.

Like let's encourage people to delay, you know, who do want to have kids to delay, uh, wanting to have children because maybe they'll need IDs in the future. So how can we, maybe we should push more. Advertisements of, uh, holiday opportunities, you know, give them deals on holidays. It's not far-fetched for an algorithm, reinforcement learning algorithm tasked with, with the goal to learn the strategy.

I mean, it has to be very. Humans to think of something like this, but a machine could learn this and people who run these companies don't even need to know. They wouldn't even know because the machines are learning their own representations of causal relationships in the world and exploiting those two presentations.

So unless we build a, a science of the behavior of the machine, then we're not going to be able to. Keep this, these things under control and thrive with them, right. Which is what we

Dr Philip Ball: want. And at the moment, the machines that do this kind of thing, if they're based on machine learning and what they're doing is, is really just looking for correlations.

That's all that, oh, if this happens, it tends to make this happen and they're exploiting that. So, um, so obviously there's, there's, there's no, uh, moral reasoning in there, but it's also that they will. Spot unexpected correlations that we don't even think of looking for. I mean, I guess that was what would exactly what came out of the, the AlphaGo algorithm that, you know, defeated the best.

Human, uh, go prayers. And it came up with strategies that no one had ever thought of. Let me explain this a little more. Google deep mind. An AI project run by Alphabet Inc. Has developed a program called alpha go, which was able to play the game of go a Chinese game. That's one of the oldest and most complex on the planet in 2015, the algorithm beat the renowned Korean professional player.

Lisa. AlphaGo is what's called a deep learning algorithm. As we heard in the previous episode, this looks at connections correlations between some set of input data and outcomes based on that data to find effective strategies for achieving desired outcome. It learns strategies using a set of training data where you already know the outcome, and then it's ready to go in situations where the outcome is still open.

It's simple minded in principle, but astonishingly effective in practice. At least for some times of problem, we have good reason to think that machines just built on this really quite simple minded approach. Isn't it just looking for correlations that they will find things, route strategies that no one.

The w we wouldn't imagine. So we could certainly couldn't anticipate building into the machine some way of preventing them from acting in these ways. We think

Iyad Rahwen: we're on that. Exactly. Because we can anticipate. Yes. So I think Africa is a perfect example because, you know, you give the machine one objective and a very powerful technology for achieving that objective, which is, you know, the.

Deep reinforcement learning and the objective is to win the game. And then it discovered this completely foreign alien way of achieving the goal. And it took, you know, experts in go a while, you know, and I think they're still digging into those games to try and understand how this machine is thinking and how, you know, what is it about those strategies that that's so effective.

Right? And so you can imagine, and this is in the closed world, closed full information world. Of course. So imagine if you have a similar machine learning to do something else, like sell things, uh, you know, and this machine works for a conglomerate that sells a whole bunch of products. You know, what could this machine machine figure out in terms of.

And it's reinforcement learning. It's actually exploiting causal relationships, not just correlation. Right? What could it learn about what sort of things you get children into when they are young, so that when they get older, they become your customer. I mean, that is not any possible thing for an algorithm to learn.

If it's, I think if it's learnable and if you have the data. They will learn before we discover it. So they are going to discover new ways of sort of capitalists strategization and maybe consumer expectation potentially, you know, and it sounds, it's great when it's empowering us and it's giving me personalizing products for us and that's, that's perfect and it helps us live better lives, but it can also have the opposite effect.

And we wouldn't know. So I think. The only other strategy that we can do is the strategy with that has been effective so far with another sophisticated intelligence enabled human intelligence. We live in this world, uh, reasonably

peacefully with other humans, despite not understanding the machine. And it's because we have good models of them, of their reasoning and behavior, and thanks to behavioral science and, uh, improving, you know, over the last few decades, we have increasingly, uh, a better understanding of this human behavior, individual and collective behavior.

So, um, sort of group. Collected, uh, you know, I've invited a collection of computer scientists as well as behavioral scientists from political science and economics and anthropology and biology. And we recently wrote a kind of manifesto. Uh, I published in nature a couple of months ago, uh, titled machine behavior and as basically a call to arms that, to be able to live in the world with, with these increasingly sophisticated AIS, we need a behavioral science of machines, as much as an engineering science.

So it's a machine being racist. Is it discriminating in, in the, when it's source CVS? Is it discriminating when it is, when it's recommends, who should go on parole or who shouldn't or when it allocates bed space in a hospital and all and so on and so forth, these are all behavioral questions and presumably, um, pivotal scientists have a lot to say about those questions.

They're the experts in these things.

Dr Philip Ball: This is a radical but timely idea to build a science of machine behavior, which we would pursue just as we do in studying human or animal behavior. It accepts that even though we're the architects of the machines, we don't know for sure what they'll do, and we'll just have to find it.

I was just in time to speak to out about these things in Boston, because he was literally on the Eve of departing for Berlin, where he is going to set up a new center for studying machine behavior. It's interesting that you, um, the way you talk about. The machines and understanding that the way they think, I mean, we don't even have a language to do it that isn't anthropomorphizing the machine.

And, you know, we say the machine is being racist. So of course the machine is it's completely insane, exactly. In a way that we recognize as, as racist. And I mean, Um, you know, there, there seems already to be this sort of, um, perception perhaps at the sort of popular level that if we make AI sophisticated enough, either it's going to at some stage become like us, or it's going to become like the robot overlord that will take us over some, for some reason it will be malicious.

Um, but. The way it's working at the moment with machine learning. Doesn't obviously it's not obviously like the way we think. I

Iyad Rahwen: think you're absolutely right. And maybe AlphaGo example is destructive in that, you know, it didn't win using a conventional strategy. Um, so if we were only looking for conventional strategies, we wouldn't have.

So I think we need to be open. And this is, I think why we need a broader scientific approach to it. Because if we just take the lay person approach, we're going to accessibly, anthropomorphize the machine and maybe make MIS what it's actually doing. But if we have a systematic way of studying its behavior, um, that is grounded in the scientific method and behaviors, good behavioral science, uh, practices, then hopefully we will come up.

You know, as this Cambrian explosion of AI is taking place, we will develop classification methodologies for all of the AI species and the different kinds of AI intelligences that are emerging and the kinds of pathologies that they could develop if you like. And, and, and, and I think that's a precursor precondition for having everything under control, right?

Like a very small example, uh, that. Uh, worked on published last year, is this idea of cooperating with machines? Uh, so there's a lot of work on, uh, how humans cooperate, even though they have the temptation to be selfish. What we did is we, we picked an AI reinforcement learning, uh, program, and we paired it with people for, uh, for the first time.

And we started studying the dynamics of. Uh, interaction. And, you know, one of the interesting things we found is, uh, there was this particular algorithm that can signal, you know, so it can do things like threats and promises and suggestions, um, and it can cooperate with people. But then when you look at the behavioral, um, level of how it actually achieved that, so it was able to achieve cooperation on par with human, human cooperation.

But it was doing it using different strategies. So for example, it was less positive. It was more threatening. Wow. You know, so that's one example and it is a very early example. Right. So, but an interesting followup question would be suppose that I have to interact with humans as well as algorithms. If machines are more vindictive, let's say that.

I may have to develop new ways of reacting, you know, new norms as myself as a human would this then impact the way I interact with other humans, you

know, is there some kind of behavioral contagion between our attraction, with machines and interaction with other human. We have no clue, right?

Dr Philip Ball: Presumably we're already seeing an aspect of that when we're interacting with one another through machines, it doesn't mean that behavior changes, you know, as online trolling and all the rest of it.

And you know, the anonymity of that. Ready concerning, I guess.

Iyad Rahwen: And I think, I think some people have also been, uh, saying that, uh, children who interact with chat bots, like Alexa starts using more imperative, you know, uh, language rather than asking other people, you know, with other children, you know? So they, they, they start ordering other children other than asking them politely.

So, you know, I still think we're, we're at an early stage in the science, but I think we really need to figure it out.

Dr Philip Ball: So there's a lot for, I had to explore in Berlin. It seems the challenge here in anticipating the interactions of human and machine is partly, but the machines and algorithms we use at the moment create no guarantee of moral probity.

We can't assume that they will observe the moral rules. We recognize and try to cultivate in society. They're not immoral, but amoral. Now you might hope that we could simply enforce morality by programming it in at the outset. But remember Isaac Asimov seminar, 1950 book, I robot, it showed how attempts to do something like that are likely to fail because they'll never anticipate all the circumstances in which the rules will be applied.

If we're going to increasingly depend on AI based decision-making, then we're going to have to grapple with this. Cold logic might never be enough to ensure that the machine behaves the way we'd like it to empathy. Isn't obviously something you can reduce to a computer code at the same time. What I had implied is that whatever morality or lack of it, the machine embodies seems likely to impact on us.

And this brings me back to my first interviewee in Boston, George Church of Harvard medical school. I asked him if we should be concerned that the technologies we're developing and perhaps using to alter and augment ourselves are changing us as human beings. Seems to me so much of the discussion about

AI in particular kind of converges on some assumption that if we continue down this path, we'll end up with something that's a bit like us, and that will communicate a bit like us.

And there's no obvious reason why that should be. So, or even if you're talking about recently, not so, so far. And in fact, we're not like us. Uh, w w we're becoming increasingly like robots. Uh, we treat each other more and more like robots. We increasingly expect people to respond to their text messages as, uh, as text messages where the nuance of body language and eye contact, all that stuff is lost in the text messages.

Okay. Um, so I think we are not us. And in fact, we're embracing. Atypical is more than we ever have before. Um, they're, they're getting more and more rights and privileges that high, high functioning ex high functioning, OCD, autistic dyslexic, narcoleptic, all these things are being, they're seeing the advantages.

We're getting better drugs that allow you to be conditional. Uh, high functioning, meaning that, that, uh, you take this drug and you'll, you'll, you'll be a little bit better in one category. Uh, you know, even though you're normally dysfunctional, you can become functional or you take this other drug and you'll be microdosing and you'll be, you'll be different.

The point is that humanity is embracing, um, a broader range and there's a lot more hybrids too. And hybrids re create a whole new set of problems. Like some people have photographic. They're allowed into a courtroom, but a camera's not allowed to port room what's going on there, you know, is that going to, are we going to be able to maintain that, that duality?

Um, or are we going to start planning people, um, that have, uh, photographic memories? Um, what if you need a camera in order to be able to. You're visually impaired. You have a camera with optical implants. What if that camera records things you have because you have short-term memory problems, um, et cetera.

Are you going to ban those? Where do you draw the line between disabled and someone who just wants to be augmented? We, we will be hybrid. And it's a hybrid, this, that we should fear more than, than, than any particular. So, so is there some set of people that are obsessed with germline manipulation of intelligence and another set of people that are obsessed with, um, like people downloading the brains and the silicone?

And I think really, probably neither of those is that, um, uh, amazingly problematic what's problematic are the, are the hybrids where we have huge corporations of people with huge machines. Working in concert so that neither, neither the people nor machines alone are threatening, but the whole ensemble is, but that's the thing that we're least prepared to do something about.

I mean, we can pass laws, you know, banning germline editing, and we can, uh, you know, we can have all kinds of, uh, efforts to constrain artificial intelligence, but w we have trouble developing the will to construct. Uh, companies that do exactly what we want them to do, that they're addictive. Uh, they're providing us with what.

Our urges want. Um, but they, but in a way, satisfying our urges is one of the most dangerous things you can do. I start worrying about new technologies, the incident, I think of it. And I encourage other people as well. I mean, some scientists or some scientists are perceived as being reassuring. Some of them actually feel that they're being restrained.

I don't feel that that's one of my job description says I don't need to reassure people. I want them to have a realistic, possibly pessimistic view so that, so that we don't, uh, so we avoid. Um, disasters. Or even minor inconveniences. Uh, if we have a false positive where I worry people about something unnecessarily, that's not so bad, you know, it makes for nice movies and, you know, and, uh, books, you know, that are sort of on the dark side or the pessimistic side.

In fact, most of, I think most of science fiction literature has a, has a pretty good. Uh, tone to it. So, so that's, that's, there's, that's almost, that's harmless to create a false positives of false, uh, concerns. Um, but false negatives are very dangerous because it means that when it arrives and you weren't ready for it, then you go into full reactive mode rather than proactive.

And for reactive mode, you tend to make more mistakes and you create problems. Um, your solutions are, are bigger, probably. Then the problem. So I think it's much more, much very important to, to talk about these things way in advance, which is what I try to do. George might not consider it his job to reassure us, but I think it is vital that we're thinking about these things now that unlike Victor Frankenstein, we are curious about how our creations will behave even while we're building them.

And in particular, how they will interact with us and affect our own behavior. That's what struck me most in several of these conversations I had in Boston

and thinking about how our artificial being might behave. The question comes back to our own behavior because that in my view is what Mary Shelley's book was really all about what responsibilities we have to our creations, to each other, and to ourselves.

How do we make sure that we don't ourselves become the monster?

It was a huge pleasure to spend my summer in Boston, but the scientists who gave their time for these conversations for that, I'm deeply grateful to George Church, Robert Langer, Alan Jackson, Thoma Allman David Cox. And I had Rowan. If you've not yet had a chance to listen to all six episodes, I hope you will.

Oh, and for. About my second brain. As I said, at the outset, it is no more. Once it was grown, it was sliced up for inspection under the microscope, but that's not necessarily the end of that little bit of me. I still have some of my stem cells frozen in a lab in LA. And who knows what lives they may lead.

Continuity: How To Grow A Human: My Frankenstein Summer is written and presented by Dr Philip Ball and directed and edited by Keith English. This show is brought to you by Aurra Studios. Listen to the full series on Apple podcasts or wherever you get your podcasts. .