JARED: The reality is technology is fundamentally neutral. If you have, say, a nuclear bomb that can be used to hurt people, but it could also be used to...

Hello, and welcome to Dead Code. I'm Jared Norman. Today, we are talking ethics and open source with Coraline Ada Ehmke. For most of you, she probably needs no introduction, so I'm not going to give her one. Let's get into the episode.

Hey, Coraline, welcome to the podcast.

CORALINE: Hi, Jared. It's really nice to be here.

JARED: So, for our listeners that maybe haven't spent as much time as I have arguing about which codes of conduct to adopt for their open-source projects...

CORALINE: [laughs]

JARED: Could you tell our listeners a little bit about who you are?

CORALINE: Sure. My name is Coraline Ada Ehmke. I'm a veteran, we will say, emerita, software developer emerita. I spent about 25 years in corporate tech, and I spent the last 10 years of my career focusing on bringing diversity and equity to the realm of open source; probably best known as the creator of Contributor Covenant, which is the first and by far most popular code of conduct for open source and other digital communities.

In 2019, I wrote the Hippocratic License, the first ethical open-source license, and co-founded the Organization for Ethical Source, where I'm currently the executive director.

JARED: Very cool. So, one of the reasons I'm excited to have you on the podcast is that when I attended my first conference in 2014, that was RailsConf, I actually got to see you speak and give a talk called Artisans and Apprentices, where you did a really great job of comparing the sort of career progression of 12th-century apprenticeships to modern bootcamps and made a very impassioned plea on the community to improve how we onboard and bring people into the tech community. It was definitely a different time back then for the tech industry. Before we dive into, you know, the main topic today, do you think we've improved how we onboard people into tech in the last 10 years? Has it gotten any better?

CORALINE: I really think that that talk is just as relevant today as it was back [chuckles] in 2014. Because, in my experience, onboarding has always been...it continues to be a very difficult problem. Kind of what I've seen in my last tech job was getting people to [02:24] their first PR against the company's codebase. And I'm not sure that that that's the best contribution right off the bat.

I think maybe something I mentioned in the talk was, when you onboard a new developer, that's a chance to see your own codebase for the first time all over again. I think we should be

listening to folks that we onboard a lot more than we're talking to them, getting their impressions of, you know, the way we're doing things and what makes sense and what doesn't, those fresh eyes. I don't think we do a good job of that.

JARED: Yeah, no, no, I definitely agree. And it's interesting, I just came from Rails World, where there's been a lot of buzz around technological solutions for onboarding people onto that framework. But there was also a keynote discussing how, which I think is certainly the thing we should be focusing on, which is the culture and, you know, the people side of these communities and how we bring people into them. So, fortunately, there was also a talk about how, you know, tech doesn't solve people problems, which hopefully, will spur more conversation on that topic.

CORALINE: And hence, it's all people problems [laughs].

JARED: It's all people problems indeed. So, I had the pleasure of seeing you speak at Madison Ruby earlier this year, and you gave a talk that, unfortunately, not recorded, so we can't link to it, but we will be discussing it here. So, you gave a talk called Four Reasons Not to Care About Ethics in Open Source, which it's a great title for a talk because I thought I was supposed to care.

CORALINE: [laughs]

JARED: So, what was this talk about?

CORALINE: Well, it was intended to be a provocative title.

So, I do a lot of work at the intersection of ethics and open-source and ethics and tech in general. And I came across a paper called Limits and Possibilities For "Ethical Al" in Open Source: A Study of Deep Fakes. And this was an academic paper put out by David Gray Widder, Dawn Nafus, Laura Dabbish, and James Herbsleb. And David Gray Widder is someone who's very familiar to me. His work really influences a lot of what we do at the Organization for Ethical Source at OES.

So, this paper was a study of a particular open-source deepfake community. What they were trying to do, what the paper's authors were studying was, how did the members of this project justify the work that they were doing from an ethical perspective? Because, obviously, deep fakes have more negative and abusive use cases than they do positive ones, whether it be misinformation or all the way down to revenge porn, or non-consensual porn.

So, they're very interested in the topic of what are these developers saying to themselves and saying to each other to kind of justify working on something that is so obviously rife with abuse? And they boiled it down to four essential reasons. And then, in the talk, I explored what those reasons or justifications were and talked about why they weren't satisfactory or why they were, in some cases, even fallacies.

JARED: Right. And that's something that's really interesting to me because, you know, there are people working on all kinds of projects, both, you know, commercial and open source. And I've talked to people who have chosen to, you know, abdicate their responsibility for the consequences of those tools for a variety of different reasons. What's the first reason that, you know, people gave that this is okay?

CORALINE: The first reason is something that I've run into a lot in my work with ethical licensure, and that is the freedom zero argument. So, this derives from Richard Stallman's work, The Four Essential Software Freedoms. And freedom zero is basically the freedom to use software for any purpose whatsoever without restriction. And this is a value that is very essential to free and open-source traditionalists. It's considered the only ethical stance when it comes to software is the software be free to use for any purpose.

And, of course, this is a carte blanche invitation to not care about consequences whatsoever. If free and open is the end all be all of your ethical responsibilities, then your license satisfies all of your ethical requirements, which frankly is a lot of bullshit.

JARED: Right. Right. So, do you think that argument, in particular, these freedoms that they're trying to allow, you know, freedom to do anything, was this a response to restrictions that existed, particular restrictions? Or was this just like, hey, anyone should be allowed to do anything?

CORALINE: I think it stemmed from a frustration that Stallman had, and let me be clear, Stallman is an awful person, but he's a very important historical figure, so I'm going to be talking about him without commentary. Stallman had a frustration with a printer driver, with an HP printer driver, and he wanted to make a modification to it so that it would report paper jams back to the kind of resource management system. And he was very frustrated that he couldn't do that because the software was proprietary, and HP wouldn't create access to the source code for modification.

And this kind of went against the nascent hacker ethic that came up in the '70s and '80s at places like the MIT AI LAB. It was a matter that, of course, if you wanted to make an improvement to a piece of software, of course, you could do it. Of course, you had access to the source code because that's how things worked in academia. And that's not how things worked in the private industry, in the emerging software industry.

So, it arose out of his frustration at not being able to make a modification to software that he was using. And he took that on as an ethical stance. And there's, you know, there's some justification for that. We do prefer software that is open. We do prefer, as developers, to be able to make changes to the software to improve it or make it better suited to our use case. But then, taking it a step further, Stallman actually treats freedom zero as a fundamental human right, and I think that's taking things a step too far. And we need to recognize that freedom zero is great in principle, but it's not the only ethical consideration when it comes to how software is used.

JARED: Right. It was sort of a reactionary response to the proprietary software causing problems for him personally, and he chose to not just adopt it as an ethical stance but sort of take the freedoms that he was looking for being infringed and say, "No, we need 100% freedom here."

CORALINE: Right. And I think one of the other things to consider here is that, in terms of the limitations of software freedom, and that's who are we extending freedom to? Essentially, Freedom Zero is extending freedom to other software developers. It doesn't take into consideration end users, and it certainly doesn't take into consideration what we call collateral users, which are people that the software is used upon or forced upon without their consent.

A freedom that only extends as far as software developers is insufficient in an age where we see software being used and abused and having significant and sometimes even fatal influences over people's lives. Being able to change the source code for an autonomous drone that carries out extrajudicial killings, who cares if that software is open, right?

JARED: Right, right, absolutely. So, that, I think, transitions us pretty well into the second argument that came up as part of your talk, the open argument. So, what's the open argument?

CORALINE: It kind of relates to what I was saying there at the end about freedom zero. The open argument is, as long as it's licensed with an open-source license, which is defined by a license that's approved by the Open Source Initiative, as long as access to the source code is free, that's a remedy to any power imbalances that may exist, whether power imbalances between big software companies, big tech companies, or imbalances in how the software is used or applied.

So, as long as it's open source, as long as everyone has access to the source code, can read it and review it, we're fine. That's as far as we need to go in terms of thinking about the ethics of the software that we're creating. And the open argument builds on freedom zero, but it extends the effect into the licensing realm and extends it to access to the source code as being a remedy, which it can be, but in a very fixed set of circumstances.

JARED: Right, right. You know, it not only makes the OSI, you know, a moral authority in this situation, but it also allows the person making this argument to sort of ignore the rest of the world's, you know, like, these...it's as if this openness exists in a vacuum and says, "Okay, well, anyone can use it, so it's fair."

CORALINE: Yeah, and I think the other thing that's a little bit more subtle, maybe is that it moves ethics and morality into the realm of intellectual property law. I mean, intellectual property law was a great hack that really enables the entire open-source ecosystem to function, but it has limitations, too. Intellectual property is not the greatest framing when we're talking about societal harm, when we're talking about the opportunity for social good. These are things that are beyond, you know, legal rights and licenses. But the OSI, their only ethical consideration is, is the license open or not? And that's as far as it goes.

And I think it's important to note, too, that in terms of the paper, in terms of the research paper that the talk was based on, the developers in the deepfake project took both freedom zero and the open argument as default positions. They said, "You know, we can't do anything about it because you have to be able to be free to use software for any reason." They didn't question that. They said, "We couldn't do anything about it because it's an open-source license." And they are accepting these things as kind of default parameters for the software that they were working on, which maybe excused them from thinking through some of the ethical implications.

They thought that as long as they were doing the default of an open-source license, that covered their responsibility, that was the limit of their responsibility. And, in fact, they said...one of the developers said, "You know, we can't really put protections in place in the software, like watermarks or anything like that, because the source code is open, and someone could just fork it and pull that code out," not realizing that the agency that they had they yielded when they adopted an open license.

JARED: Yeah. And that argument almost seems like they found themselves in a trolley problem, and they're like, well, I'm not going to pull the lever because somebody else could just pull it back [laughs]. Not a very compelling argument when stated like that.

CORALINE: Right, yeah.

JARED: One argument that I've absolutely seen wielded in a variety of settings, not just when it comes to software, is that technology is fundamentally a neutral tool, that it has no fundamental ethical stance. Truly, we heard this argument from the maintainers of this project.

CORALINE: Yeah. And this is what I call the hammer argument, and it derives from an often-repeated quote by Noam Chomsky, philosopher Noam Chomsky, where he said that technology is like a hammer. It can be used to build a house, or it can be used by a torturer. And what Noam Chomsky was talking about was general technology at a very high level. He was talking about the cell phone. He was talking about the internet. He was talking about the personal computer. He wasn't talking about drone navigation software. He wasn't talking about software that's used to determine sentencing for people who are convicted of a crime. These are not the technologies that he's talking about.

And, in fact, when we look at specific technologies, of course, their use is dictated by the ethical decisions of their makers. They are designed for a purpose. There are hammers that are designed to build houses, and there are hammers that are designed for torturers. And to pretend that they're the same thing is ignoring the entire field [laughs] of design and [inaudible 14:47] because we build technologies for very specific problems.

And if we're not thinking about how to solve those problems, if we're just thinking, well, it doesn't matter because it could be used for good or bad, and it's on balance; it's fine, then we're not

paying attention to those negative use cases. And most importantly, we're not designing for those harmful use cases because we've abdicated responsibility for even thinking about them.

JARED: Right. It's almost as if the argument looks at tools as being some sort of just sort of universal thing. Tools can do anything. I don't know. Who knows?

CORALINE: It's a very simplistic view, and it's increasingly inaccurate as our tools get more and more and more specific. And we know that that's the direction of things. That's the direction of the software that we build, the frameworks that we use, the languages that we use. All is to get down to a level of specificity that I think we conveniently don't apply when we're thinking about the overall, you know, is this technology going to be pro-social? Is it going to make improvements to the world? And are those improvements on balance more impactful than any negative impact, or any misuse, or any abuse of the software?

But if we're not planning for abuse and misuse, then we're enabling it. Think about social media platforms at their height, where moderation was really key, where community and safety features were really key, where we were anticipating ways that technology could be abused and trying to build tools to address them. You don't get to a place where you're building community management, or anti-harassment, or pro-privacy, or anti-spam technologies by assuming good intent on the part of developers.

JARED: Now, there's one last argument that you could say it was inevitable that was going to come up in this list.

CORALINE: [laughs] Yes, the inevitability argument.

JARED: Right. What's that?

CORALINE: One of the participants in the study basically said, "Well, if I don't work on this project, someone else will." So, it's the notion that technology is inevitable, that a particular development in technology is also inevitable, and that it doesn't matter who does it. I might as well get paid for it. And, in the talk, I gave the example of the parable of the locksmith, which was basically some mysterious stranger comes to a locksmith with a suspicious job, and the locksmith sort of says, "Well if I don't open the safe for this mysterious stranger, they'll just go to the next best locksmith."

And what that is, this is actually a logical fallacy because it's saying, as long as I can imagine someone taking this job, as long as I can imagine someone building a deep fake technology if I can imagine one person for whom there's not an ethical problem, then it's fine for me, too. So, we're actually abdicating ethical responsibility to the least ethical person that we can imagine and claiming that that's carte blanche to do whatever we want with a given technology.

And when you frame it that way, it's kind of ridiculous, right? You're saying that, you know, you only have to behave as well as the worst person in the world. And as long as you behave no

worse than them, then you're fine. You're in the clear. That argument, I think, really bothers me the most, and I think it's the most dangerous as well, in part because it does make that assumption that all technological developments are inevitable. And this is a symptomatic of, you know, Western thinking, White Western thinking that civilization is always a curve that goes up into the right, you know, we're not doing anything to set ourselves back significantly. And that's just...that's a myth.

JARED: Right. And it's...I do, you know, well, I totally agree, but I also see where people might fall for this one in that when we look at these tools, we are supposed to be looking at their potential for harm and what people might do with them. So, you are supposed, you know, I should be thinking about, oh, well, could someone take this and make these changes and do harm with that? That's, you know, I think we'd both agree that that's on me to consider. But, obviously, you know, as you say, I shouldn't abdicate responsibility to the point that I'm making those changes that I'm, you know, putting that out into the world just because somebody else could hypothetically also do it.

CORALINE: Yeah. It kind of feels like, well, I was just ordered to do it. Advocating responsibility, except in this case to an imaginary person, which is...that's quite an argument [laughs] to make. I think we can go...we can do one better than simply not being the person who does that technology. Well, first of all, to understand the things that we're not going to build, understand the things that...where our lines are. And I always encourage people to do what I call a negative road map. And you can do this for yourself, or you can do this at your company. Make a list of stuff that you refuse to build.

If a product manager comes to you and asks you to put spyware in the learning management system, you are within your rights to say, "No." And that might get you fired. It might cause conflict. It might cause a big stink that impacts things long term, but it's worth it if you're saving the human rights and the fundamental privacy rights of students that are in high school.

And how about the inevitability argument there? If I don't stand up against being asked to build this feature, no one else will, right? So, I think knowing what your lines are, and knowing what your boundaries are before you're placed in a situation where you have to make that on-the-fly decision at stand up or what have you, I think, is very valuable to kind of know what your boundaries are going in.

JARED: Absolutely. And it's a piece of advice I certainly wish that I'd had myself before. I've found myself in situations where I chose to step back from contracts or other things where I was not comfortable with the ethical implications of moving forward. And having laid down those lines for myself without any money on the line would help avoid, you know, the situation where you're trying to make a more perhaps pragmatic call in that situation rather than the correct ethical one.

CORALINE: And think about, you know, we opened the conversation talking about onboarding. Think about the junior developers on your team, the people straight out of bootcamp, the people

who don't have the standing or the reputation to be able to say, "No," necessarily, with less consequences. So, do you want to be the person who, in front of those newcomers to our industry, do you want to be the person that builds that spyware? Or do you want to be the person who says, "No" to that spyware? What are you inspiring? And what are you saying to the next generation of developers that's coming up right behind you? We have to think about that as well.

It's not just our decision. We're deciding on the part of our teams. We're deciding on the part of our industry. We're deciding on the part of, you know, what we owe society. And if we're not thinking about these things, then we're guaranteed to do harm.

JARED: Yeah. Looking at what we can do to sort of change the culture, you know, you mentioned earlier these positions were viewed as default positions, that they're viewed as default positions because they assumed that place in the open-source development culture. And if the culture were different, you know, as more people push back on these, certainly, we would see fewer people viewing them as default positions.

CORALINE: And I think, you know, the default position of open, the default position of software should be free; these are valuable, and they're historical. And they are a huge part of how we got to where we are with open source having eaten the world. But these were also values that were written down in 1998, where, you know, the biggest threat to humanity in terms of what software developers saw was the market domination of Microsoft with, you know, 98% of the browser market.

And we had scrappy underdogs like Mozilla that came along, open-sourced their code, and were thus able to compete with the behemoth that was Microsoft. And, eventually, for a period, at least, we won the web. We lost it to JavaScript frameworks [laughs] in the meantime. But for a while there, we won the web, and it required that confrontational stance. It required that almost zealotry behind free and open in order to rescue us from that moment in time.

But we have to go back and revisit these things because the technological context has changed, and the social context is radically different. So, it's important to revisit those decisions and maybe our motivations or our guardrails that were effective and appropriate in 1998. Maybe the world's changed a little bit, and we might want to revisit those. And that's essentially the work that we're doing at the Organization for Ethical Source. We're saying open is great; free is great. What's next? And what do we have to do to adapt, to continue being that revolutionary force for broad social good? And we're not going to do it by doing the same things we did in 1998.

JARED: So, what do you see as sort of next steps for open-source technology when it comes to navigating this sort of ethics of what's being put out into the world?

CORALINE: There are lots of touch points along the way in the evolution of a technology or, a software project, or an open-source community. There are touch points along the way where we need to be checking in.

Really, the first place to start is deciding if you're going to build it in the first place, and this comes into the negative roadmap idea a little bit. Someone made a decision to make deepfake technology. Someone made a decision to start that project. What if, at that moment in time, before that first commit message, what if someone had come along and said, "Hey, there's a set of questions that you should ask yourself before embarking on a new technology project. And we're going to walk you through it, and we're going to tell you why it's important to ask these questions"?

And those questions include things like, have you thought about the potential for harm? And does the potential for good outweigh the potential for harm? Are you designing things to be fundamentally supportive and to promote fundamental human rights, including emerging digital rights frameworks like privacy and data autonomy? Are you working against those interests? We're developing at OES a complete set of what we're calling a pre-flight checklist of questions that you ask yourself at the very beginning when you're imagining the solution, when you're imagining this technology you want to bring in the world. What are the considerations at that stage in time?

And then, moving on from there, you have your open-source project. You're choosing your license. Open licenses are great, but some software, in some cases, you may want to consider using ethical license, which the Hippocratic License, which OES put out, is tied to the United Nations Universal Declaration of Human Rights. Human rights should not be controversial, but people have very strong feelings about licensing. So, that's an option you can consider as an Ethical Source License correct for this project. Or am I fine? Am I comfortable with an open license? You're making a decision about how you manage your community by the adoption of the code of conduct and other governance documents.

And I think kind of the last step in the process is, when your software is actually in the world, when the technology you've created is actually in the world, you're not done. You need to go and look to see, you know, how is it being used? How is it impacting people? Is there harm that's being caused? And is there something I can do to mitigate that harm—to prevent it, to heal it, to repair it? There has to be a cycle.

All of these decisions and all of these checkpoints are cyclical. There are questions we should be asking ourselves at every step along the way to make sure that what we're putting into the world is truly good. And I think, in most cases, our intentions are good, but that doesn't mean that it's simple, or easy or that we can just accept the default settings and get what we expect to get out of the other end.

JARED: Well, Coraline, thank you so much for being on the podcast. Where can people follow, you know, the work that you're doing online?

CORALINE: Well, if you're interested in the work that we're doing with the Organization for Ethical Source, you can go to ethicalsource.dev. You can find me on Mastodon. I'm Coraline Ada on Ruby Social. I abandoned X and Twitter a while ago.

And you can also keep an eye out for my upcoming book. I'm happy to say that this is the first public announcement, I think, besides my Mastodon post. I'm writing a book tentatively called We Just Build Hammers, stories from the past, present, and future of responsible tech. It's going to be published by Apress. And you should look for that coming out in March.

We go into a lot more depth not only on the sort of ramifications of technology decisions we make but also connecting with the history and connecting with different visions of the future by speculative fiction authors, so an interesting blend of both history and future coming out in the book. I'm super excited about that.

JARED: Awesome. I'm very much looking forward to that. Thanks so much for coming on the podcast.

CORALINE: Thanks so much for the opportunity, Jared.

JARED: The nature of responsibility when it comes to, you know, our interactions with the world around us is something that, you know, I've talked about on this podcast, and, you know, I care very much about. I think that the reality is deciding you do or do not have responsibility for the world around you, in general, is, you know, a fundamental ethical position.

If you have power over something, in some sense, you are responsible for it. And whether you use one of these arguments that Coraline talked about or something that you've come up with, you know, some other argument to abdicate that responsibility, doesn't change your role in these systems. You are able to effect change, and you are choosing not to.

It's not always the case that, you know, those roles and those decisions are ethical ones, but sometimes they are. And it's worth evaluating, you know, your role in these systems and deciding, you know, where your ethical boundaries are and where your boundaries in general are, what you are and are not willing to do, and what impact do you want to have on the systems that you're a part of.

This episode has been produced and edited by Mandy Moore.

Now go delete some...