Episode 2: Infrastructural Justice: developing safer technologies

[MUSIC PLAYING] KSENIA BAKINA: Welcome to The Cyber Armor, a podcast which champions voices and safety of women and girls in the digital world. This podcast is brought to you by the Center for Protecting Women Online and the Open University. I'm Dr. Ksenia Bakina, your host for the podcast and the research fellow in law and policy at the center.
In today's episode, we will be discussing responsible software engineering and the concept of infrastructural justice. We'll address how this concept could be incorporated into responsible software engineering in practice, and why it's important to do so. We will also discuss what difference this could possibly make to women's online safety.
And in order to address these issues properly, I have two brilliant guests here with me today, Professor Arosha Bandara and Dr. Sarah Robinson. Professor Bandara is a Professor of Software Engineering at the Open University. He also leads The Future of Responsible Techstream at the Center for Protecting Women Online.
Professor Bandara's research and teaching focuses on software engineering for adaptive systems. He has a particular interest in techniques for building adaptive security and privacy, mechanisms for internet of things systems.
He is also a member of the steering group for the Open University's Center for Policing, Research and Learning, and is currently Associate Dean and Director of STEM research. Dr. Sarah Robinson is an Applied Psychologist at the People and Technology Group, School of Applied Psychology UCC and Ireland, and Lero Science Foundation Ireland Research Center for software.
Her work contributes to understandings of responsibility, justice, and care across a number of contexts, including software engineering. Arosha and Sarah, welcome to this podcast. It's a great pleasure to have you here with me today.
AROSHA BANDARA: Thanks, Ksenia. It's great to be here.
SARAH ROBINSON: Yeah, thanks, Ksenia.
[MUSIC PLAYING]
KSENIA BAKINA: So, I want to kick off this discussion, this concept of responsible software engineering, and explain to our listeners what this actually means and how is responsible software engineering might be different to ordinary software engineering. And, Arosha, are you able to give us some insights on that?
AROSHA BANDARA: Sure, yes. So let's start with software engineering. And so this is a term we use for the scientific approach to designing, building, deploying and maintaining software-based systems. Its goal is to try and make the software we build deliver on our needs, the needs of organizations or individuals who might be the end users of that software or paying for its production.
So for software engineers, this means using a collection of processes, methods, and tools in order to effectively design, build, deploy, and maintain these systems. In that context, the concept of responsible software engineering comes about because there's the recognition now that software-based systems have become pervasive in our lives. And that means that they have both beneficial and potentially harmful impacts on people individually and at society-- on society at large.
That highlights the necessity for software engineers to really consider the ethical implications of the systems that they're involved in building, the effects it has on individuals and society. And so thinking about ethics as a branch of philosophy that's concerned with how to live a good life, responsible software engineering, we could say, is about addressing the challenge for software engineers to understand and address the question of, how do I develop software that is beneficial and not harmful to individuals and society more generally?
In social media platforms, the content delivery algorithms were designed to increase engagement, and therefore they have a benefit to the content creators and the advertisers who are trying to generate revenue and maintain their businesses on the back of those platforms.

But they're clearly also have-- there's increasing evidence on how they harm users' well-being by encouraging addictive behaviors like what we call doom scrolling. And also sometimes really tragic outcomes like the case of Molly Russell, another example that's highlighted in Sarah's work about how the pushing of harmful content to a vulnerable individual led to very tragic outcome of suicide in that case.

This is what responsible software engineering is trying to address is to try and maximize the benefits and really minimize the potential harms, and actively consider them as part of the software engineering process. And I think that's primarily because of the way software-based systems, software technologies that embed software in them have become such an integral part to the way we live in the modern world.

When software was first being developed in the 1960s and '70s-- well, even earlier, actually, my history of technology isn't quite up to speed as I would like-- but early software very much focused on the functionality that was being developed. So if we think about the kind of first, what would be considered widely used application of computing in the mainstream, was something like spreadsheets, for example. So VisiCalc and Lotus 1 2, 3.

These were products that really embedded themselves in business, but they had very limited scope and focus. And the focus of the developers was just on providing the functionality that those users and companies wanted. But we've moved on so far from that kind of software to software that determines how decisions about our finances are made, how our health is provided, that we use for our entertainment and education, and so on.

So all of these things mean we can't treat software as simply a tool that is completely in the hands and control of its users. The development of it has to consider all of these wider societal and individual implications as well.

KSENIA BAKINA: Yes. Thank you so much for emphasizing that. And I think it's really interesting to really go back and highlight where it started because it really sets the tone and understanding that we've moved on from spreadsheets to smartphones, to apps to having technology and software in our pockets.

We have so much control about what we do with it and how we use it, whether we use it for harm through notify apps, or whether we use it for good. And with that in mind, I want to come to you, Sarah. And I know in your research with your colleagues, you've identified this concept of infrastructural justice that can help improve and develop further this idea of responsible software engineering. But before we get to talking about that, could you just unpack a little bit, what do you mean by infrastructural justice?

SARAH ROBINSON: Yeah, thanks, Ksenia. And thanks, Arosha, for that really, really good introduction. I guess maybe just before I answer that, just to maybe respond to some of the things Arosha has been saying and then to position this idea of infrastructural justice within that kind of dialogue, one of the things that I am interested in-- I work in an applied psychology division and work very closely with colleagues who are software engineers and in software engineering disciplines.

And I guess I'm interested in the institutional arrangements of how software is developed. And I think one of the things that we've seen over the last 50 years is that that has really dramatically changed from software engineers-- academic software engineers, but also people in industry working either with small companies, maybe as an independent software engineer supporting them, or with some of the big tech organizations.

And we're really seeing that that has been, kind of, distributed now. And so in our paper, we talk about the open source community and the way in which software development, to a certain extent, has been de-professionalized and opened up to other people, to people who are tinkering at home with software. So that's kind of an emerging area that I think is important for us to consider when we think about responsibility.

And then there's also this kind of growing body of people that work freelance software developers across the globe, often in contractual relationships that might include maybe a client in the global north with software developers in the global south, et cetera. So that's just to give you a sense of the distributed institutional contexts, I suppose, that are changing all the time in terms of how software is developed and has implications for responsibility.

And I guess our idea of infrastructural justice stemmed from looking back at literature, over the period of time that computing and software has been rapidly developing to think, what are some of those historic ideas around ethics and practice in institutional contexts? What might help us for this kind of contemporary changing landscape? And one of the ideas that really inspired us was this idea of infrastructure.

So when we think about infrastructure, we often think of critical infrastructures, maybe our transport, our rail networks, et cetera. But the idea that we draw on thinks of infrastructure as not just those substrates or those things, but actually as something that's deeply relational. That infrastructure only becomes infrastructure to us through our use.
So, for example, if I'm a wheelchair user, a stairs is not infrastructure to me. If I have the ability to walk, it is, et cetera. So infrastructure both includes and excludes. And when we think about software systems as infrastructure and the institutional context of software development as infrastructure, it also includes and excludes. So we found that concept very illuminating. And I'll explain a little bit more about that as we go along.
But we were also interested in the fact that there are, I guess, the responsibility landscape. There's over 180, probably over 200 now, different guidelines about responsible AI, for example. And as Arosha mentioned, responsible software engineering as a kind of term really only came into discourse in the last five years.
We've had 50 years of computing ethics, so it's a kind of crowded, discursive terrain. And some of the principles that are dominant in there are things like fairness, explainability and transparency, but one of the principles that we were surprised not to see, given the societal implications of some software applications, was justice.
And so we wanted to bring these two ideas of infrastructure that both can include and exclude in justice, which we took from a political theorist called Iris Marion Young. And she was interested not solely in distributional justice, this idea that if we redistribute resources in society, we will have a just society. We wanted to go a step further.
And she has this idea of structural justice, which is that actually societal structures which are technologically supported can, in fact, create domination and oppression. And in order to overcome that, we need to have what Iris Marion Young would say, structural justice and what we would say, infrastructural justice.
And what I mean by infrastructural justice is that all of us who are technology and software users are in some way connected and responsible for ensuring those systems and infrastructures that they support are inclusive and do not lead to domination and oppression. And, indeed, if they do, we are all in some way responsible, not to blame, for trying to make that structure, that infrastructure free of domination and oppression.
KSENIA BAKINA: Are there any key principles that we can tease out from this theory that we can latch on to and try to embed further into responsible software engineering?
SARAH ROBINSON: Yeah. So maybe if I could start with that. So I guess one of the key differences of this way of considering responsibility is that often when we think about responsibility, we think about blame and we think about liability. So we think, for example, in the case of Molly Russell, the coroner indicated that the companies are in some way to blame for that situation.
And I think for Molly Russell's family, that's very important, isn't it? And I think that does bring a certain amount of justice. But what we would say is that in tandem with that and differently from that is that responsibility, when we only consider it in a liability way, can isolate responsibility. It can socially sanction and it can say, OK, that's over now. That's addressed. But, actually, what we know is that, really, that situation is not fully addressed. And there are still harms and injustices that are created by those particular social media systems, for example. And in our approach, what we're saying is we want to spread responsibility to everybody who is in some way connected to this issue.
And by spreading responsibility, we're saying that we want to reflexively learn from it. And I guess what we're trying to do with this idea of responsibility is to say, yes, liability is important and it is there and it feeds into things like requirements. But if we only have liability, we know we also create these cultures that aren't reflexive, that aren't open to understanding failure as generative.
So to go back to your question, I think a key principle then is reflexivity. It's about creating these dialogical spaces where the public, where people like me who work in psychology and social sciences are in dialogue with our engineering colleagues about improving systems and processes. So I think reflexivity is something that's really important and maybe not emphasized enough currently in the literature.
The second thing that I would say, Ksenia, is that we often-- in tandem with liability, we isolate responsibility, have blamed someone or a corporation and there's a social sanction in some way. But, actually, what we want to do is to say-- is to look at norms, is to look at the taken for granted, morally accepted practices that both software engineers, but also those of us who

use software systems do together, and how we unwittingly often contribute to harms and injustices.
In our paper we talk-- we give many examples, but an example that some of my students are working at the moment is the fact that there are six billion smartphone users in the world. All of those smartphones are connected to a cloud. And the fact that that cloud is not tangible to those users means they don't fully understand the material impact in terms of sustainability, which-- we're interested in women-- disproportionately affects women, and often marginalized women in the global south. So to get back to a principle, what the principle would be is making those background conditions that are often intangible to us and immaterial visible to us.
KSENIA BAKINA: I was just reflecting on what you've said from my own legal background, and I remember that when, first, the issue of image-based sexual abuse surfaced back in 2013, it was viewed like revenge pornography. So it was viewed that only something that ex-partners did to each other in terms of posting their nude images.
And as a result of that perception that only the person who posted your image is liable, we had a law which just dealt with the initial perpetrator. It didn't address at all the fact that the image may have gone viral and the responsibility potentially of software platform, the tech platforms for facilitating creation of these images, facilitating the spread of these images.
And it's only now, in 2025-- almost 12 years later-- we're finally talking about, actually, platforms such as Meta, such as Google also have responsibility and can we develop some guidelines or codes of practice for them. So that idea of spreading that responsibility beyond the initial course and effect perpetrator is definitely really, really important. What problems might continue to exist in software engineering if we don't incorporate this concept of infrastructural justice?
AROSHA BANDARA: If I may comment first, if that's OK. So, yes, I think that's a really important valid point and challenge about what's the implication of not doing this and not looking at this. And I think one of the things that Sarah highlighted very importantly is that our current way of dealing with these challenges has this isolating effect that each incident or event is treated as happened, let's figure out what went wrong and who to blame and then move on.
That makes the arc of justice move very slowly as a result of that, I would say, because it means you get that piecemeal response in the infrastructures of justice like legislation resulting from these individual incidents. So that flipping of the narrative and approach, the reflexive approach that Sarah was highlighting where in an environment where there is scope to have these open conversations and learning about where the challenges will arise or there's anticipation, so being able to try and anticipate where the potential harms might arise and understanding collectively how to deal with them and how to minimize those harms, can change them-- the actual thing we build-- to be safer and better for everyone.
And, actually, reduce the need for that liability-based approach because there needs to be a shared understanding that-- if there is a shared understanding that collectively software developers, engineers, technology companies, regulators and legislators, the public at large, have collectively made a joint effort to do their best to minimize injustice, then I think when things go wrong-- which they will still likely go wrong because the world changes-- I think there's space to accept that we did our best, something went wrong, but we can now learn and try and improve.
Whereas in the current model, that doesn't happen. And the liability model shuts down those conversations because no one wants to share what they're doing in case that makes them liable for financial penalty or sanction of some kind.
Another potential critique of this approach is diffusion of responsibility means no one will be responsible. But I think my response to that is to say it's not about diffusing the responsibility to avoid it, but it is to clearly assign it, to understand the problem and clearly assign it to the right people, because at the moment-- and the right processes and methods-- because at the moment, that's not happening. And this kind of post hoc figuring out who might be liable and penalizing them is clearly piecemeal at best, ineffective at worst.
SARAH ROBINSON: I think I completely agree, Arosha, with everything that you said. And we haven't asked the question of who is responsible, so who is responsible when these harms asked? And, I guess, in our infrastructural justice paper, we're trying to say we're not all to blame, but we are all responsible in this way of reflexively contributing to a more just society and a more just software systems.

And Iris Marion Young, who we draw on very heavily in the paper, her theory talks about the importance of shared responsibility, which is different from collective responsibility. So collective responsibility might happen when a country, for example, is responsible for something, or a corporation, but individuals within that company or that country may not feel responsible or be allocated responsibility. And so it isn't necessarily distributed.

With shared responsibility it's different. So shared responsibility is personal responsibility that you identify as being responsible for the outcomes or the risks of harmful outcomes that are produced in a group. So, for example, the example of 6 billion smartphone users-- which I'm pretty sure all three of us are part of that-- we Are all partially responsible for the environmental impact of smartphone-- of our daily digital practices.

We're not fully responsible for that, but there is a partial responsibility, and it's very difficult for us to quantify what that is. And in those situations, shared responsibility becomes more important. So we can't quite tangibly delineate it, but we can know it and we can act on it.

KSENIA BAKINA: So we've discussed a lot about the concept of infrastructural justice and its meaning. And we've talked about these quite broad concepts of shared responsibilities, but these broad concepts also disrupt ordinary ways of way of thinking. So it seems that in order to really achieve infrastructural justice, we've got to challenge the status quo in the way we go about our daily lives-- the amount of photos we store on our phones, the fact that we update a new iPhone every year without thinking-- and that seems to be quite a challenge.

So I wanted to challenge you a little bit more and start thinking, well, how can we address this concept of infrastructural justice? How can we incorporate this into our responsible software engineering in practice? What would we need to do?

AROSHA BANDARA: The key challenge here in your question is that in practice part. I mean, I guess partly we've talked a lot about the background theory aspects of this already. So some of the things we've already highlighted, Sarah highlighted about making information knowledge more explicit or surfacing it in ways that allow individuals and groups to be reflexive, be deliberative in the choices that they make as a result. I think that that's a key aspect of this.

But also I think there are things that we can look at to create those spaces for dialogue and discussion and building consensus around decision making. For example, software, and if we take the area of mobile apps, for example, there are already some public fora in where some issues are discussed.

And I would point to app stores and the reviews and the dialogue that happens between developers and users of the apps through that review mechanism of the infrastructure, I think, is a very tentative step in the right direction of opening up that space for that interaction to happen. But we need to go much further and be able to not just silo those conversations in each app's page on an App Store, but also to be able to bring them together more collectively around categories of technology like generative AI or image generation technology.

I mean, I think this is the challenge, that the level of abstraction you choose can also be inclusive, but also then exclusive in that it loses the nuance that you need to progress. So I think what we need to do is to create those spaces, enable those spaces for conversation to happen, and to have mechanisms to take the output of that conversation and make decisions that drive how we build and deploy the technologies.

SARAH ROBINSON: Yeah, I completely agree, Arosha. I think dialogue is the key term that I was hearing and dialogicality. And in my own research, we've also been consulting the public about some of the ethical concerns that they have with software and what they are doing to address it.

And what often comes up is that people feel powerless and that they have contacted maybe software providers, they have made complaints but the complaints are not addressed, or there isn't a kind of satisfactory response, or maybe they're stuck in a loop talking to a bot. And then eventually they get worn down, then they don't dialogue anymore.

So I think that as a society is something that we really need to change. And that's linked to my second point. So your question was about practice, and, obviously, engineering practice is central to this, but I think it's not just engineering practice. I think software engineers and engineering is responsible, but it is responsible in tandem with the corporate world, for example, states and semi-state organizations that it engages with, and the public as well.

So all of us are in some way responsible. And I think at the moment we're quite siloed in that rather than collectively coming together. There are obvious moments and examples where that does happen, as Arosha has mentioned with the App Store, but kind of creating more

spaces where people across different silos can come together and consider these issues I think is really important.
And, Ksenia, you mentioned image-based abuse and initially how the idea was about perpetrator and victim, and how, of course, that kind of de-responsiblized the platforms, for example. So that is an example, I think, of where that kind of siloed way of thinking about this can be part of the problem, that we miss the bigger picture.
So I think an infrastructure lens makes us see this across an infrastructure. So it makes us look at the kind of ordinary normative background condition practices, and how they're facilitated and how they would unwittingly create harms.
So it helps us think about how the standards and conventions of each of our disciplines, each of our professions might impact how the material conditions, et cetera. And one of the other things that we're working on, led by Arosha and my joint colleague, Dr. Aaron Ralph, who has been looking at freelance software development-- and I mentioned earlier that kind of disparity sometimes with clients in the global north working with clients in the global south in more precarious situations and contexts.
And I think it's shedding a light on those kind of challenges, I suppose, for responsible software development. And that's sometimes the harms can actually be to the software engineer as well, or the software developer. And so it's kind of shifting our gaze in lots of different ways and breaking down silos, I think.
AROSHA BANDARA: Yeah. And I completely agree. And I think one of the other reasons why this is really important, I think, is because some of the challenge we face comes from assumptions about responsibility and motivation and intent of different stakeholders in this space.
An example I can draw on from research we've done here at the Open University with colleagues in software engineering, looking at cybersecurity practices in software development teams, in professional software development teams, which stemmed from this assumption and a kind of external assumption from the team's perspective that software engineers are not sufficiently caring of cybersecurity, and so they're producing products that are inherently insecure.
And that was an assumption being made by the external perspectives on the software. But when we went and actually did some ethnographic work to understand software engineering practice in these professional contexts, what was clear is that software engineers actually care very deeply about security and want to embed practices that increase security of software.
But they're operating in a context that has lots of other pressures. And there are organizational pressures, market pressures, things that are preventing them having the space and time to do those things effectively. And what we see then is insecure shadow practices that come about because there are these external drivers that are pushing things in that direction.
So I think that this importance of dialogue is key in order to avoid assumptions, incorrect assumptions being made, and, therefore, responsibility being ascribed to the wrong place, essentially. And that a lot of-- many of the problems we see and the kind of reactive responses have this element of the responsibility being put in a wrong, at best, a suboptimal place. And if we had a more infrastructural justice approach, we would understand better where responsibility lies in this shared responsibility model that we're trying to.
SARAH ROBINSON: In the infrastructural justice model of a shared responsibility, it says that, yes, we are all partially responsible, but we're not all equally responsible. I think that's really important. So those with power, and often the power to define a system or to create those decision making, make the decisions and influence the decision making architecture are more responsible than those that don't have that power.
Similarly, those who are privileged from a system have more responsibility. And many of us in the global north are incredibly privileged, actually, and don't want to give up our privileges, when we think about climate action, for example. So it's a really difficult thing, but it's a core tension.
And then there are those with interest. So, again, we're not talking about blame. This is not victim blaming, which sometimes the critique and can suggest. It's actually saying that those with lived experience of harms, of injustices, are often the most qualified to suggest what needs to change. And keeping them out of the equation is a problem.
So engaging with people who are directly harmed is very important, and others who have an interest, maybe allies or people connected to those people with lived experience. And then,

lastly, those with ability. So not all of us are able to address the injustices or harms, perhaps because we don't have the technical capacity or we don't have the resources.
And in our model, it would be about looking at those across these systems to understand where responsibility lies. And acknowledging that, really, shared responsibility, there needs to be a shared response. So a collective response, that brings these kind of different layers together to address it.
KSENIA BAKINA: And as you were talking, I am sitting here, and I'm trying to imagine this space. And I'm trying to imagine how this would work where this dialogue could happen between, whether it's users, governments, platforms.
And I'm just-- so I wanted to ask you, in terms of-- so two questions that were in my mind were, first of all, what needs to be done for all of these players in the idea of infrastructural justice and shared responsibility could come together? Do you think-- should government issue a law that once a year we need to have this discussion? Or is it going to be led by the big players, the big platforms? Or should it be more those freelance engineers that are potentially working from home? Or should it be led by the users?
And so how do we create that space in the first place? And then, secondly, how do we safeguard it from, A, this idea that this could turn into shifting blame where everyone's blaming each other? For instance, such as in the case of Molly Russell, that this platform say, no, no, no, it's the users who are to blame because they're the ones posting images that glorify suicide or other harmful practices.
Or it's the government that's to blame because they don't regulate, or they make it unprofitable for us to do-- to safeguard our users. So how do we safeguard it from passing the blame and also safeguard from those who are with power potentially not being so open? Because they might not see it as they're financially beneficial for them to engage in that open dialogue, to say, oh, we're working on this and this and this, and this is how we see it impacting the world, because they might afraid that somebody's going to steal their idea, for instance. So, yeah. So I think what can we visualize the space? What needs to happen for it to exist? And how do we safeguard it from misinterpretation of this concept of infrastructural justice?
SARAH ROBINSON: That's a really hard question, Ksenia, and it's one that keeps me awake at night and I think about a lot. I guess, maybe to break the question down in terms of what is the space and how do we define the space, when we talk about infrastructural justice, we're talking about all sorts of things, aren't we?
And what we've taken, I suppose, from the seminal work in computer supported cooperative work-- which I don't know whether the listeners would kind fully understand that term, but essentially it's a kind of multidisciplinary space that looks at work and practice that are supported by computer systems, often software enabled systems.
And from that work is this idea of infrastructure as relational, as something that becomes useful through practice. And what we learn from that work is that every context is different. So when we talk about infrastructural justice, we're talking about different context specific things and ways, and different ways in which domination and oppression are created.
So I think the first thing is about how we define the multiple problems within that very big overarching problem space. Arosha's example earlier of the research which understood a kind of problem space or an assumption of software engineers don't care about cybersecurity, which wasn't at all the case, actually, when that more reflexive ethnographic work showed that there were much more what I would call, actually, infrastructural normative pressures that created unwittingly harms.
So I think it's something about how we define the problem space. And at the moment, the problem space is saturated by we need guidelines, we need principles. And, actually, we have almost 200, if not more than them, just for responsible AI, and that doesn't really seem to help. That seems to actually hyper focus us in that space and distract us from practice.
So I think there's something really important about how we define the situation and the contextualization of that. And as a qualitative researcher, primarily, I think there are a lot of practices from reflexive thematic analysis, for example, from ethnography, from discourse analysis that can help us understand and maybe reframe problem spaces. So that's the first thing that I think needs to happen, and is happening in all sorts of places.
And the second thing that I think needs to happen is education. So, for example, in my university we have an undergrad degree that's psychology and computing. And in that degree, students are exposed to social science ideas, pluralistic methods, as well as software development and computing. And by no means perfect but I think that kind of model of

learning about society and societal ways of understanding society, and understanding software as part of society is really, I think, important.
AROSHA BANDARA: I think, yeah, it's a really challenging space. And one of the issues, I think, that we need to make explicit and address is this challenge about addressing the broad problem versus approaching it in a more focused way, in particular contexts. And I agree with you, Sarah, about what you said about the importance of context and looking at the specific nature of particular social technical systems in order to make meaningful progress in this space.
Because in one of the potential misinterpretations or misapplications of this goal of infrastructural justice is that it becomes a kind of tool to prevent progress-- or that's a critique of it, I guess, that could be made that the emphasis or the need for all of this dialogue will take too much time and then that will inhibit technological progress, and so on.
And that's not the intent or goal, but I think it's also raising the mirror to the fact that technological progress is not a kind of inherent good as well. So we need to recognize that there are goods in technological progress, but there are also harms, and, therefore, this is about saying let's make sure that we're maximizing the goods.
And so it needs us to take that focused approach in practice, but then take those learnings to upper level of abstraction that allows us to learn more broadly, address the norms and practices and cultural aspects that Sarah has been highlighting in what she's been saying. Because if we try to address the abstract problem, that's too big a problem and we won't-- we will get stuck. And all of these criticisms of this is hindering progress will come true and be valid.
But I think in context, we can make progress, but the importance is that those contextual investigations and solutions need to take what would be termed a transdisciplinary approach, an approach that understands that there are these different perspectives that need to be considered.
Including the perspective of experts in lived experience of a domain that the people who have been affected by are users and have interest in the domain of study, but also the expertise from different academic perspectives and practice perspectives that are important.
And I think that's what excites me about the Center for Protecting Women Online is that we're bringing together those perspectives in a place to explore and try and address it, the specific problem of online harms to women and girls. But even that, I would argue, is too broad a problem, and we will need to break it down into specific contexts to look at.
KSENIA BAKINA: I wanted to move on to talk a little bit about online harms affecting women, as Arosha, you mentioned. And we've talked today-- we've mentioned the case of Molly Russell and what that meant for-- in terms of infrastructural justice. But I wanted to ask more specifically, how do you think this infrastructural justice can help to address online harms against women generally? Or perhaps you might give some examples surrounding specific harmful practices affecting women today.
SARAH ROBINSON: Well, maybe, Ksenia, I could just give you, firstly, the example that we have in our paper about infrastructural justice which relates to deepfake technology and the development of deepfake technology. And so for listeners who aren't familiar with deepfake technology, it's essentially when photographic or video evidence is adapted to impose someone's face or body on video or photographic images that they did not participate in, in the events that it depicts.
And, unfortunately, most deepfake technology is pornographic, so the vast majority of it. And a lot of it and the action orientation of it is harmful to women. So we see particularly celebrity women where there are a lot of images that can be taken and imposed on maybe pornographic video images to depict celebrity women in particular ways, but it's also been used in terms of revenge porn situations, bullying situations, et cetera.
And I guess what I want to illustrate to the way of thinking about infrastructural justice is that we could think about deepfake technology in the victim perpetrator lens, but it would only get us so far. But if we bring an infrastructural justice lens, we're looking at who's powerful within the infrastructural arrangements that enable the development of that software. And who gains privilege from it?
And when we think about that, we can see the landscape of practice. So deepfake technologies is primarily developed in open source communities that are supported by platforms and repositories like GitHub, which we know is managed by Microsoft, a very big technology company.

Often people who develop deepfake technology communicate through Reddit and other platforms. And there is what we call an installed base, which is the material and social and cultural conditions which enable the development of this technology.
So, for example, there are these very large open source databases that have been used to develop what people call generative AI, but things like ChatGPT, but also to develop deepfake. And those particular, those amalgam of different platforms and repositories come together to create these conditions where deepfake development can flourish.
And so if we break that down then and think about where responsibility lies and where change can come, we can think about things like platform moderation on GitHub, we can think about things like addressing the no-safe-for-work content that is readily available in open source data repositories. We can think about open source itself, which often is considered a very benign, actually, pro-social space where people can come together collectively to develop software, and it's not commercial.
But the idea of open source licensing is that we don't think about downstream harms. So we don't control for that. And responsibility is distributed in the sense that there can be a lot of anonymous contributions to software. And what we're seeing in this instance is, actually, that can be really harmful for women.
We're also seeing that there is a norm of toxic-- what the literature calls toxic geek masculinity, so men that maybe feel that they have not lived up to the kind of norms of masculinity that society proposes and have come into these spaces to enact masculinity, unfortunately, in quite a toxic way in this instance.
AROSHA BANDARA: Yeah. And I think part of this is also not trying to say that consideration and-- yes, consideration of these potential harms necessarily means that the technology-- the work of developing a particular technology should be prevented or stopped dead in its tracks, so to speak. But it's more about, again, making this knowledge and information explicit so that it can be responded to appropriately.
So, I mean, the example of deepfakes, there are-- the deep part of it comes from the fact that it's using these deep neural net technologies, the generative AI technologies that enable the creation of these images and audio and video.
I think that underlying technology is also what is driving a lot of the generative AI that we're seeing being used, and sometimes used very beneficially to reduce burdens of some types of work in particular contexts. Obviously that has other implications for the nature of work and so on, but this is all, I think, what infrastructural justice is trying to address is to make sure that those things are being not just being ignored, but are being actively considered.
But it also means that rather than stopping a technology from being developed or deployed as a broad response, it means appropriate steps can be taken, for example, to make sure that the uses of that technology are actually tracked and-- so that interventions can be put in place downstream if needed to prevent harms that were not fully anticipated when a technology was developed.
SARAH ROBINSON: I think in the case of deepfakes, we saw, for example, Reddit have born-- have banned non-consensual porn from its site, but yet it still acts as a repository for people who want to make it in the sense of people can still congregate there and find other ways to share this information. So even when we have something that's designed to make a space safer, it can sometimes have the opposite effect.
I think there's something important when we think of software development in that particular context of de-professionalized open source communities. I've seen that whole landscape and how everything contributes and connects, and seeing the leverage points across that system.
KSENIA BAKINA: Thank you so much for highlighting these examples. And I think these examples really help to bring the idea of infrastructural justice to life. And as you were talking, I was thinking about potential issues beyond software engineering, and how infrastructural justice can help to-- the idea of shared responsibility can help to highlight those as well as, for instance, with the example of deepfakes that you gave that responsibility goes beyond just the initial creator of deepfakes, but also to the apps, to Reddit, to Microsoft.
And even to those users who might just want to learn about deepfake technology, but having to engage with the sexualized, non-consensual content. And that kind of also made me think about what broader responsibility and thinking about, well, what about pornography companies? What about the pornographic industry that sexualizes non-consent? Because that's also a responsibility outside of software engineering.

But it's an important-- I think, your point about toxic masculinity, Sarah, really brought that to life for me because I started thinking, well, if non-consent was not so sexualized, would we want to create these non-consensual deepfakes in the first place?
So I think that really-- that example really helps to bring to light the shared responsibility across different layers of society within software engineering, but also beyond that as well.
And so I want to wrap up this episode by asking this question-- particularly to you, Arosha-- in terms of how can we best use this idea of infrastructural justice and shared responsibility at the Center for Protecting Women Online and further our projects, potentially, in light of what we talked about today?
AROSHA BANDARA: Yeah. I mean, I think there are quite a few things to learn and take into what we're doing, but also reflect on the context that the center provides that can actually be a real enabler for moving forward this idea of infrastructural justice as well.
And, primarily, I think I would say it's about that space for dialogue, for interaction between the different disciplinary perspectives that the center brings together. So we have experts in law, policy and regulation, we have human behavior represented, software engineering, artificial intelligence and policing.
So those five-- let me put in quotes-- "academic" perspectives on this problem space are important, but even more than that, are the broader non-academic partners that are involved in the center who have interest and motivation to try and address online safety for women and girls, and make online spaces safer and more inclusive. And that means that we have some of the-- well, an environment in which these dialogues can happen, which I think is an important first step in that it's not always there, it's not always easy.
The other aspect of it is just through those conversations, one of the things I think we can contribute to infrastructural justice is the shared language of having a shared language to be able to progress what responsibility means, how it could be shared amongst those who have power and interest, and all of these other facets that the infrastructural justice framework helps us unpick and understand better.
So, yeah, I mean, I think it has real potential in that way. I mean, the challenge for us is-- as is reflective of the challenges we've highlighted through the conversation, I think, is what are the mechanisms and tools by which this dialogue can be both conducted but also captured and understood so that we can learn and put into practice appropriate interventions, appropriate solutions and tools to help improve online safety for women and girls.
KSENIA BAKINA: Thank you very much, Arosha and Sarah. It's been a pleasure to have you both on today's episode. And thank you for engaging in such a really interesting and insightful discussion. In today's episode of Cyber Armor, we discussed issues in relation to responsible software engineering and how this can be combined with the concept of infrastructural justice. We also addressed why this concept is important for women's online safety, and how it can be put into practice.
I hope you enjoyed listening to today's episode, and you stay tuned for the next installment. In the next episode, I will be talking with my colleagues from the Center for Protecting Women Online about deepfakes. And we'll dive deeper into this issue from both legal, psychological, and technological perspectives. You can also, in the meantime, follow the Center for Protecting Women Online on our LinkedIn page.
[MUSIC PLAYING]