

Transcript: Exploring Digital Life, Technology Change and Attitudes to AI through National Datasets



[0:00:00]

Athina Vlachantoni: Hi. I'm Athina Vlachantoni, based in the University of Southampton, and my presentation is the result of a project in the Centre for Population Change – Connecting Generations Programme funded by the ESRC and the Centre for Research on Ageing, both at the University of Southampton.

The work is the result of a team which also includes Jane Falkingham, Maria Evandrou and Min Qin.

The background to the work is research that we already know is showing that family ties are quite strong in England, and we've seen that in terms of reciprocity of support across the life course. We know that using analysis, using cohort data, has shown that most midlife adults had received some support from their parents earlier in life, and then they themselves support their elderly parents later in life. Within this context, the pandemic, the COVID pandemic, has had direct and indirect impacts in a number of areas, including health, for example, producing severe morbidity and greater mortality among the older age groups.

The pandemic also disrupted family interactions, specifically for those individuals in groups who were encouraged by the governments to shield, to stay indoors and not have contact with other individuals. And in this context, digital solidarity, which is often in the literature described as an example of associational solidarity, it provided a way of using technology to maintain contact between generations with the potential of improving health and wellbeing. And that digital solidarity was particularly valuable for older

persons in the earlier parts of the pandemic when traditional ways of staying in contact with your family could not be used.

So previous research has shown that increasing age reduces the likelihood of using the internet and using digital means of communication, although we know that at the same time it's not one's age generally, but at a specific point in time that matters, and that is defined by cohort membership.

We also know that there are gender differences in the use of the internet. So, we also know that persons with higher educational qualifications are more likely to use the internet and we know that people who have poorer health are less likely to use technologies in general.

Our aim in this project was to try to understand whether these relationships held in the context of the pandemic.

So overall, the objective of the research was to examine patterns of intergenerational digital contact before and during the COVID-19 pandemic in England, and specifically, we were interested in three research questions. Firstly, whether sociodemographic factors and internet experiences in the past were associated with intergenerational digital communication before the pandemic, whether that changed during the pandemic, which groups shifted to using more communication, which groups shifted to using less. And finally, we were interested in understanding whether inequalities in the use of digital communications had been reinforced or mitigated during the pandemic.

The data came from the English Longitudinal Study of Ageing, Wave 9, so 2018 to 2019, and the first wave of the specific COVID sub-study, which was collected in 2020. And the analysis included just over 4,000 adults aged 65 and above who had at least one surviving immediate family member living in another household, and that could include an adult child, a parent or a grandchild.

Digital communication was defined as emailing, texting or video calling and intergenerational digital communication was that contact with children or other family members outside the household of the older person.

[0:05:00]

We defined frequent or less frequent by using the threshold of at least once a week and we applied multivariate binary logistic regression to understand the determinants of frequent intergenerational digital communication.

Some findings. This is the first figure that shows you the percentage of all the respondents who engaged in frequent intergenerational digital communication before the pandemic, that's in blue, so about 54% of the respondents reported such frequent contact, and about 25% in yellow reported shifting towards such contact during the pandemic, while 4% reported shifting away from such contact during the pandemic. So that's the total population.

What happens when we look at gender differences? Well, about 46% of men compared to 61% of women reported engaging in frequent communication before the pandemic, and during the pandemic, 30% of men and 21% of women shifted towards such communication, compared to 5% of men and 3.5% of women who shifted away from such communication.

And finally, did the experience of having used digital means of communication in the past make a difference? The answer is yes. So we found that those who had ever used the internet before the pandemic were more likely to engage in frequent intergenerational digital contact compared to those who had never used the internet. And the pandemic actually had a polarising effect in that slightly more of those who had ever used the internet shifted towards such contact than those who hadn't. And on the other hand, a greater proportion of those who hadn't used the internet shifted away from such contact than those who had.

In this table, we summarise the direction of the changes during the pandemic. These are the net changes, a shift towards digital communication subtracted by the shift away. And we find that women had a higher proportion of frequent communication pre-pandemic than men, but the gender disparity narrowed during the pandemic as more men shifted towards more frequent digital communication. A much higher proportion of older adults who reported having no close family members shifted towards digital communication and that narrowed the digital disparity in terms of family closeness. In contrast, the pre-pandemic gaps between internet users and non-users in intergenerational communication, these widened during the pandemic.

We also explored whether that frequent digital communication between the generations during the pandemic was linked to health and social care usage and loneliness during the pandemic. These results show that older people with frequent intergenerational communication, they had a lower level of unmet health and social care needs, and those are measured with hospital operations or treatment cancelled or not accessing needed community health and social care services. And they also had less loneliness, they reported less loneliness, and that's measured through variables reporting feeling lonely, feeling left out, isolated from others, and that compared with their counterparts.

So overall, over half of the respondents aged 50 and over reported frequent intergenerational communication, and it was women, younger cohorts, those living with someone or having close family members, those who were healthier, and those who had been internet users in the past who were more likely to have frequent intergenerational digital contact.

During the pandemic, we also saw that many people shifted towards frequent intergenerational digital contact and those shifts were specifically observed across all groups with different demographic socioeconomic characteristics and internet skills and previous use.

The digital gender gap and also disparities in the use of digital communication related to family closeness in the family, those narrowed during the period, but inequalities in terms of having used internet in the past or not, those widened during the pandemic. And overall, we found that intergenerational digital communication was important, had a positive role to play in terms of meeting older people's health and social care needs during this time.

[0:10:35]

So overall, the context of this research is that many older people experienced greater associational solidarity with their families during the first phase of the pandemic as a result of digital intergenerational communication, but at the same time, there remained a subgroup who were, we call digitally excluded, who were left behind during this time. What can be done in the future in our paper that has been published, we have suggested reducing the digital divide through education programmes for digital skills that highlight the benefits of using such technology. That's important where in the context of about more than half of the persons who lack basic internet skills being over the age of 65. So that's an important way forward for policymakers.

The paper has been published in the Journal of Applied Gerontology. Thank you for your attention, and I look forward to the discussion.

Roshni Modhvadia: Hi, everybody. I'm Roshni, a senior researcher at the Ada Lovelace Institute, and today we're going to be talking about our two major surveys that we've conducted capturing public attitudes to AI which were carried out in partnership with the Alan Turing Institute.

But as a bit of context, Ada is an independent research institute with a mission of making data and AI work for people and society. And public attitudes research plays an important role for us to be able to do this. We use it to help build our understanding of public hopes, expectations and

fears as well surrounding these technologies. And we often use a range of methods from deliberative dialogues, qualitative interviews through to polling and surveying work.

Today, I'll start by giving some context into why we decided to do this nationally representative surveying in the first place, and then I'll describe the two datasets that we have as a result of this work, both their commonalities as well as some important differences, and reflect on what we've learned methodologically in this process.

So starting with context, our first survey went into field in 2022, and it was published in 2023. And quite notably, fieldwork predated the launch of ChatGPT.

We were, in 2022, and I think still now really, we were kind of operating in a landscape where AI felt like quite an abstract term, that on its own doesn't really hold a clear meaning. It's often associated with things like highly sentient robots, things of sci-fi, things of the future. And even now when you search this term, I think the images are quite unhelpful. It comes up with things like complex circuit boards that appear to mimic brain functioning. You see these images associated with sci-fi, robotics, advanced robotics. And all of this speaks to potentially some sort of future problem or future facing technology, when in reality, these technologies are in the here and now, impacting people in society and they're often in the here and now in quite mundane and everyday ways.

And we found that at the same time, there was very little on the horizon in terms of the governance and regulation of these tools. Often this fell into other regulatory domains and there was definitely nothing comprehensive in terms of how we were thinking about AI and its regulation.

And so this prompted us to look at the evidence base in this field and we saw some of these problems and challenges playing out here as well. So focusing specifically on surveys, we found that surveys would often ask

about AI as quite a general concept and often as well in terms of dichotomies. So firstly asking, generally, how do you feel about AI without really providing much of an explanation of what this is, what kind of systems it might relate to, how it might be operationalised in the here and now. And then in terms of measuring attitudes, ranking things from being quite helpful to harmful, feeling optimistic or pessimistic, but not really an opportunity to express holding both of these values at the same time.

[0:15:23]

And while there was some research on specific use cases or applications of AI, a lot of this was concentrated in specific spaces, so things like self-driving cars and to a lesser extent facial recognition more generally or automation in general.

And just again, we felt that these framings were quite limited. It didn't really speak to some of those more mundane applications of AI that might be affecting people in society now, be that applications on phones or probabilistic based kind of decision-making. And it didn't really give space for the public to express both their hopes and concerns around these tools in a more detailed way.

And so we carried out this work, and in some of those principles you can see through both waves, and they felt really important design choices to make which is why we've held them quite consistent.

So firstly, we focused on specific uses of AI, so presenting people with very clear examples of how these technologies might be used or embedded. That might be something like facial recognition used to unlock a mobile phone, it might be a more behind the scenes application of AI, and I'll talk to some of the specifics that we surveyed in a moment. But as I've said, it was really important for us to contextually ground this work.

We asked about people's awareness of these tools and their personal experience with them, the level of benefit and concern that they associated

with each application in our survey, as well as the key advantages and disadvantages they saw from these uses. And then we also asked the public about their expectations around the governance and regulation of these tools, as well as their overarching trust in different institutions as well and different stakeholders within this ecosystem.

It was important for us that this piece of work was nationally representative using strong research methods and so we made a decision to opt for random probability sampling to help ensure that. And just a small detail, the first wave of the survey focused on a Great British public, whereas the second wave focuses on a UK public. And moving forward, if we were to do further iterations, it would use that UK framing.

As I've mentioned, there are two iterations of this work and the first wave went into field in 2022.

So as I've mentioned, it was really important for us to ask about a wide variety of use cases ranging from the highly impactful to the quite mundane to provide some of that comparison, contrast context in terms of how the types of attitudes different applications of AI elicit. We had 17 applications of AI in this first wave of the survey, and we used a split survey approach to manage this. So this meant that everybody answered questions relating to the use of facial recognition for unlocking a phone, and then they answered about half of the other applications in the survey. So for each other application, half of the sample answered on it.

The applications fell into six broad buckets. So as I've mentioned, facial recognition as well as risk and eligibility technologies, so things and systems that might determine how eligible you were for social welfare or for job applications. We had both political and consumer targeted advertisements. We asked about virtual assistants in the form of smart speakers as well as healthcare chatbots. Advanced robotics, so this included autonomous weapons, robotic care assistants, driverless cars, as well as again that more

mundane robotic vacuum cleaners and also applications of AI in more research or education contexts.

[0:19:51]

And we found that our approach of multiple use cases to be quite informative. There were stark differences across technologies in terms of attitudes. So some technologies were met with quite overwhelming positivity, others with equal levels of concern and optimism, and then for again some quite high levels of negativity. So for instance, for the most part, the majority of the public thought that a diagnostic application of AI in the context of cancer detection to be quite beneficial overall, whilst very few felt the same about targeted political advertising, for instance, or even driverless cars. And we flattened the picture here a little bit. So we've subtracted overall levels of concern from overall levels of benefit, so the extent to which people rated these applications on those two scales, just to demonstrate some of this nuance. So with those applications that have net positive values, it indicates where benefit might outweigh concern and then vice versa. Yeah, even across these applications with both high and low levels of public support, we found a lot of nuance. So being able to see the overall benefit of AI in cancer detection, for instance, did it mean that the public don't recognise potential drawbacks of this use, for instance when it related to over-reliance on technology over professional judgements, or what this might mean for accountability if mistakes were made?

And then when we asked the public what would make them more comfortable with the use of AI, 62% of people said laws and regulation and 59% said clear procedures in place to appeal AI decisions. And in terms of governance and who should be doing this regulation, an independent regulator was the most popular option selected. And so together these findings suggest a desire for avenues to appeal decisions made by AI and broadly for the regulation of these tools.

And so stepping back, while the survey really offered a lot of new insight, I think for the landscape in terms of public attitudes to AI as an organisation, as a team, it also made us reflect a lot on what was potentially missing from the evidence or what we weren't able to do. A large limitation, I think, because we found that some voices it felt like were getting lost within a nationally representative framing. A lot of our institute's work uses these more qualitative or deliberative methods, often targeting specific populations or thinking about who might be most impacted by technologies, and then centring their voices and I think within a nationally representative context, it becomes harder to do that. And this challenge was heightened by the split survey design that we used. So we were trading off breadth of use cases, potentially for more detailed insights relating to specific publics.

And we also found ourselves wanting to explore additional aspects of governance and regulation that we weren't able to in the confines of this first wave of the survey because in some ways it required us to do the thing that we didn't want to do, which was flatten AI into quite an abstract concept again asking about the regulation of AI quite generally or preferences for governance quite generally. So these were very live challenges that we had as we were reflecting on this piece of work.

And so when it came to the most recent iteration of our survey, so to the second wave, we knew there were some things that we wanted to retain in the design, some things that worked really well for us, as well as some things we wanted to explore a bit further and potentially modify.

And so we still used random probability sampling as a methodological basis for this piece of work, but we added in some demographic boosting to complement this nationally representative picture that we were presenting in our findings. And this was quite an iterative process. It was informed by both the literature as well as just what was feasible, both within the survey landscape we were working in and our own budgetary restrictions. And so ultimately we decided to oversample based on ethnicity, specifically Asian/Asian British participants and Black/Black British participants.

Income, we wanted to over-represent people who might be on lower levels of income which included both kind of like an equivalised income measure as well as self-reported measures of how well people felt or how comfortable people felt financially. And we also looked at digital skills and those who might have lower levels of digital skills or digital confidence as another group to focus on. And again, we were thinking about who might be disproportionately affected by AI and whose voices tend to be suppressed in nationally representative reporting.

[0:25:27]

And to ensure our sample sizes didn't really dwindle for subgroup analysis, we also narrowed our AI use cases examples down from 17 to eight, which all participants answered questions about. So this included repeats from our first wave, so facial recognition and policing and surveillance, and some risk and eligibility applications, as well as some uses of advanced robotics. And we also added in two new use cases based on developments in the field. So firstly, large language models in general, like ChatGPT, and then secondly, an example of a commercially available mental health chatbot built on a foundation model that was able to respond seemingly naturally to text input from a user and offer mental health support and guidance.

And this approach allowed us to explore subgroup differences and it demonstrated potential disproportionate negative impacts of AI, or certain applications of AI for some groups of society in a way that we weren't able to really discuss or explore in our first wave. So for instance, while the national picture of attitudes to facial recognition in policing seems quite positive at face value, when we dig deeper, we do see differences in attitudes with Black and Asian people feeling more concerned than the national average about the use of these tools. And at the same time, we also found that those on lower income were more concerned overall about nearly all of the technologies we surveyed. And overall, again, we found that people do see the benefits of many use of AI, but that concern level seems to be rising. So

while reported levels of benefit is quite similar wave on wave, there is a slight increase in the level of concern that people are reporting.

And so for instance, if we use the example of the use of AI to determine welfare eligibility, 59% of the public reported feeling concerned about this tool in our second wave, compared to 44% in our first iteration.

And then looking at some more specific measures of governance or regulations surrounding AI, we found that firstly, the public self-reports high levels of exposure to potentially AI-generated harms. So we asked about things like whether people had encountered deepfakes online or potential misinformation or fraudulent activities that could have been AI-generated. And at the same time, they also expect the government to be equipped in relation to AI safety. So we asked about preference for the government, not just private companies having specific powers around the use of AI, kind of strong support for those powers not to just be concentrated with private companies.

And to close, I'll offer some overarching reflections now having conducted two waves of this work.

So this project has made us and continues to make us think about how to appropriately leverage quantitative approaches when it comes to quite complex attitudes, sometimes very challenging subject matter. We feel like our data does show that the public understands often intuitively the risks and benefits posed by AI and often has quite consistent views on how some of these risks in particular can be managed. So for instance, things like independent regulation or removing concentrations of power from private companies. They crop up elsewhere in not just our research but research within the wider landscape.

[0:29:35]

But again, digging into the specifics of things like safety and governance mechanism is still a challenge that we're thinking about on how to do this

effectively using quantitative approaches in ways that mean that results aren't necessarily artefacts of survey design, question wording and things like that.

And by thinking about representation over sampling, we also uncovered some challenges in the wider survey ecosystem. So for instance, adding boost to a project that already uses random probability sampling is incredibly costly and sometimes these panels are not always set up for it in the first place. So, finding agreed ways of thinking about things like low income and digital exclusion that align with what is currently available in panel structures can be quite difficult. And even still, there are issues of representation. So for instance, we found that Black British participants are quite poorly represented in many panels to begin with making oversampling approaches really difficult.

And the funding structures we were operating in also added time pressures to deliver and it would have been great to have been able to pilot in more depth questions around governance regulation. Those questions where we really needed to, it would have been great to unpick some of the complexity of these questions, understandability and things like that.

And then narrowing down our use cases in the second wave of the survey also meant sacrificing some breadth of our coverage. And so we're considering ways of prioritising use cases going forward that ensure that we have enough variety in the survey to provide evidence of attitudes across a range of sectors, both private and public, commercial applications versus more personal use compared to more institutional use of these tools and systems, and doing that in a way that doesn't lead to participant fatigue or too much repetition in the survey.

And lastly, we're thinking about the utility of the tracking function of this public attitudes work. And this is where I think we'd really value any feedback that you have in terms of how you have or might use these datasets. So we now have over two, we have two survey waves of data that

really highlight nuances in public attitudes, just both demonstrating that these remain quite contextually bound and therefore showing that asking about AI is a singular concept can be quite unhelpful.

And we're thinking about how this tracking work could be accompanied by things like wider social listening which might explain some differences in attitudes we see wave on wave, thinking about the social narratives that surround AI between survey iterations, as well as thinking more broadly about what is useful to continuously track in the ecosystem and where we have flexibility to respond to changes in AI development and the policy landscape more broadly and how to balance this work alongside our strong portfolio, both qualitative and deliberative inquiries that often complement, but sometimes also raise more questions from this data.

And so finally, finally, I just want to say thank you for listening. Apologies for not being able to be there with you live to discuss this work further.

And I'll just leave you with some key resources. So here you'll find our microsite which hosts the findings from both iterations of the survey, has detailed their methodology and obviously more detail around how we did this work. And here you'll find our raw datasets, our data tables, code books, any kind of accompanying materials you might need for both waves of the survey. And my email address is there as well if anybody wants to reach out.

And just to conclude, acknowledging the various partners that are involved in both of these surveys. So our first iteration was a collaboration between the Ada Lovelace Institute, the Alan Turing Institute, as well as LSE's Methodology Department, and it received funding from the Alan Turing Institute and the Arts and Humanities Research Council. And the second iteration was a partnership between the Ada Lovelace Institute, Alan Turing Institute within the Public Voices in AI Research Programme and supported in part by BRAID as well. Thank you.

[0:34:32]

Mark Elliot: Hello, and welcome to this talk on the Public Perceptions of AI. So this is based on a study carried out by myself and my colleagues, Filip Bialy and Robert Meckin.

I'm going to cover the project, give you an outline of what we did, and then a summary of the findings, and then I'll talk a little bit about data sources.

So, the project was funded by the North West Partnership for Security and Trust, and ran for five months between November 24 and March 25.

So, we had quite a few different research questions, but they're all framing around what constitutes the components of public perceptions of AI and how is this mediated by demographic, contextual, and technical factors? And then what's the consequences for that? And also thinking about mediating factors such as institutional trust, transparency and human oversight in the acceptance of AI technologies.

So, the methods that we employed were a systematic review. As you'll see, there's a lot of research going on in this area, a review of available data sources, which I'll report on at the end, and then secondary analysis of some survey data.

So, for the systematic review, these were the search terms that we used. And you see that there's perception, opinions, attitudes and there's different ways of framing this essential concept. We weren't in this particular study focused on knowledge, which obviously is one of the key things you could think about how much we know about AI, but we looked at those other kind of concepts. And there's also this one here of trust and that's an important component, as you'll see later.

Interestingly, that didn't pick up a lot of extra hits because trust has imbued all of the other three types of terms in terms of papers what quite frequently,

very frequently in fact, we're referring to trust as part of their understanding of perception.

So, key findings. So, one of the things to say is that the research on public perception of AI is growing rapidly. Most of the research that we captured happened since 2021. We had a little bit of 1,000, more than 1,000. We whittled that down to 150 through a scoring process. We wanted to get it really focused on stuff that actually was about that rather than just mentioning it in passing, and then snowballed that to an eventual sample of 251 papers for thorough review.

So, here we have the graph which shows the growth in papers on this topic. And as you can see, it's an exponential growth. Unsurprising, perhaps, given the fact that AI has erupted in public consciousness, so this would be a topic which would be now more thoroughly researched. But nevertheless, it was quite an eye-opener for us, the volume that's been produced in the last few years.

A second key finding was the contextualisation of public perception by country, culture, political narratives. And there was a real difference, for example, in respondents from Western countries and from countries from the South East and East Asia, and with China particularly noted as being techno optimistic. And this has come up in a variety of contexts. But the key thing here is that there is not a singular global reaction to AI, and that's going to affect political narratives going forward.

[0:38:40]

So demographics and education had a complex relationship with public perception of AI. So in general, younger, higher educated male respondents who are familiar with the technology tend to be more positive towards AI. Nevertheless, some studies found that AI literacy also correlates with an increased awareness of risks and concerns about what will happen with AI and the problems that might arise. So it's a double-edged sword. And we

found that in our own work, when we were doing the survey study, it was very much very important to focus on both concerns and benefits. And Roshni has already referred to that study, the data which we used, as having that particular feature.

Despite generative AI being the most salient in public discourse and dominating media coverage, research has focused on domain specific applications rather than particular types of technology for AI and machine learning. So, it's very much on how it's going to be used, how it's going to be deployed that has been the focus of most of the research that has happened.

Trust is transferred from institutions to technology. This is a really critical element of the findings that we had. So, essentially, rather than thinking about trust in AI or perceptions of AI, it's the perceptions of the organisations deploying the AI which is really critical in framing how the public view the technology. So this trust transfer mechanism suggests that when institutions are well regarded, their stewardship of AI systems is more readily accepted. So, that really overrides the specific use case, and that's really important, a very important finding. And as you will see, trust comes through as really, really vital.

So technical features about applications did affect public trust. AI systems designed to appear over the human life actually met with scepticism. That's a general finding. And there's a preference for functionality, human control and transparency over anthropomorphism. This is interesting because some of the narratives in the media would suggest otherwise, and that in order to trust a robot, it has to be reasonably human-like rather than in some sense alien. And this pattern is observed across domains and use cases.

Now, returning to this issue of trust, so this is a piece of software we ran over the entire corpus and picked out key terms that were relevant to our understanding of opinions and attitudes. And you can see here that trust

absolutely dominates in terms of the intentional verbs used to capture that relationship.

Now, you'll note, of course, that one of our search terms did include trust, so we ran this again removing those papers that were only captured by that mechanism, and the graph was very, very similar, in fact. So we're very happy that this is a really dominant understanding of what public perception of AI means, that that trust relationship is really important.

And just to emphasise this further, this is the word cloud for the entire corpus, and you can see here again that trust dominates it. We've excluded all minor words like envy and so on and a few academically focused words like paper and so on which tend to be artificially captured obviously in anything that's written by academics, so to get at the core concept of the substantive area, and this is the word cloud, and you can see the essential notions. Things that we thought might be important, like privacy, are on this table somewhere, they're very small. Trust dominates everything and that kind of fits in with our general understanding of the corpus of papers.

Okay, to go on to the datasets. This was Public Attitudes Towards Data and AI, an ONS survey, PADAI. It covers perceptions, governance, AI literacy and knowledge and there's four waves so far. It's been run since 2021, there might even be a fifth wave now. And it's really good for looking at perhaps changes over time, certainly in recent years.

So this is a sub-dataset of the Opinions and Lifestyle Survey and it looks at awareness and attitudes towards AI. The latest release for that was December 24, so relatively recent. And again, it's available via ONS.

[0:44:03]

So if you want to do cross-country comparison in Europe, there are a few surveys, and I'll mention a few more in a moment. So this particular one is the PAICE survey, 4,000 European citizens, and includes a UK sample. And

it's focused on literacy, trust in AI and perceived AI impact. And that's available via arXiv.

So, this is the one that Roshni has already introduced, so I'm not going to go into any detail in terms of the overall spec of this, but it has a lot of value. We use the 2022 survey. There is one, as Roshni said, from 2024 now, and it's published by the Ada Lovelace Institute and Turing.

So three other European datasets that might be relevant to you if you're going to do cross-country work, there's the AI and the Future of Work, European Citizens' Knowledge and Attitudes Towards Science and Technology, which includes AI, and Attitudes Towards the Impact of Digitisation and Automation in Daily Life. So, those three surveys are available via the europa.eu website. And again, give you that opportunity of doing cross-national comparisons.

So, we used the Alan Turing and Ada Lovelace dataset, as mentioned earlier, and we used this to complement the systematic review with some secondary analysis. A lot of the results here confirm what we found from the systematic review, and we're going to just focus on one particular thread here, for what the dataset enables you to do.

So first of all, as Roshni indicated, capture of both benefits and concerns about AI uses. And this particular one's driverless cars. And it was interesting because it comes out as one of the most negative uses of AI in terms of its perceptions. And as indicated, I think in the headline results, male, young people and indeed people with degree level qualifications tend to see the benefits of AI more readily. But we can also see from this across all sectors of the population, there is more concern about the use of driverless cars than there is perceived benefits. And it's particularly acute in that higher educated group which is an interesting tension that exists in these data. And this is one of the most extreme versions, and Roshni

pointed this out in one of her slides, that actually is perceived more negatively if you subtract the number concerned what from the number seeing benefits than even automated weapons systems, which certainly surprised me.

The dataset also gives rich information on particular details of concerns and indeed benefits, and that allows you to do a sort of deeper analysis. And this is the list of concerns that were raised for driverless cars. And there's several different things going on here. Some are thinking about the reliability of the technology. That comes through quite strongly. But there's lots about how humans and the technology will interact, both at the societal level and at the individual level. So I think that it's interesting that it is so present for people, and that might be because of the ubiquity of cars, it might be just because, in Western culture, most people now drive or can drive and have cars as part of their daily existence. And maybe it's that embeddedness of the impact which brings out all of these kind of quite nuanced thoughts about what their concerns might be. So that's an interesting thing and I think that's worthy of further study and it'd be interesting to see how other technologies, as they become more salient, how the concerns get raised and benefits are perceived going forward in time. And obviously, you can look at this again in the 2024 dataset and examine where has this changed, has patterns changed here?

Okay, so there's a lot more detail on this in our full paper. The full paper itself is a distillation of the report that we produced for NWPST which ran to 300 pages. The paper is about 75 pages, so there's a lot of information in there on the review and the dataset and so on. So it's available on open access. The link is here. As I say, I think the slides are going to be shared with you so you should have that available to you if you want to look into that in more detail. Thank you very much.

[End of Transcript]