Please cite the Published Version

Interlinked Computing in 2040: Safety, Truth, Ownership, and Accountability

Charles Weir [ID] **and Anna Dyson** [ID] **,** Lancaster University

Olamide Jogunola [ID] **,** Manchester Metropolitan University

Louise Dennis [ID] **,** Manchester University

Katie Paxton-Fear, Manchester Metropolitan University

Computer systems are increasingly interconnected, magnifying benefits and risks, especially with AI integration. Using a Delphi-based method, we interviewed technology futurists about potential trends towards 2040 and their societal impacts. Our findings highlight five key forecasts related to artificial intelligence and system complexity, and suggest six interventions to mitigate negative impacts.

What is going to happen in the world of digital technology? How might things develop in the next 15 years, as computer systems become more and more advanced? In particular, what issues will arise as digital systems become more *interlinked* and integrate *novel technologies*, with software systems owned and controlled by different organizations interacting directly and indirectly with each other? These interlinking and novel technologies include direct web application programming interfaces, smart grids, blockchains, user-programmed online bots, and many others.

INTRODUCTION

While the pace of technology seems to surge ahead quickly each year with large leaps, organizations must react with caution, thinking about the technological maturity before investing. If a technology is adopted too soon, it can lead to issues in overall system stability or security, and if it is not adopted soon enough, organizations may find themselves lagging behind their competitors.

What should we be concerned about, and what actions should we plan in the face of these challenges? These two questions are pivotal for those designing and planning interconnected computer systems and, given how much digital technologies now impact our lives, vital for almost everyone.

To address that question, we undertook a form of Delphi study, a well-known technique for forecasting. We interviewed a range of respected futurists about how they see different aspects of new digital technologies and their interactions with interlinked computing affecting our world by 2040. From those interviews, we produced a series of predictions; then, to build a more complete picture, we went back to the interviewees and asked them for reactions and comments on the initial forecasts. In this article we explore five of the arising forecasts and six of the recommended interventions.

The aim of this article, therefore, is to help policy makers and technology professionals make strategic decisions around developing and deploying novel interlinked computing technologies using the information in these five forecasts.

The rest of the article is as follows. The "Background" section explores the art of futures forecasting. The section "Interlinked Study Method" looks at approaches to forecasting the future and describes the approach we used. The section "Interlinked Study Results" describes the interview participants and explores five specific forecasts in detail. The section "Suggested Solutions" discusses some possible interventions suggested by the participants to address the problems involved in the forecasts, and finally, a summary is provided.

BACKGROUND

This section explores the methods for technology forecasting and its past usage and effectiveness.

Technology forecasting methods

Future studies is now a well-established discipline, with its own conferences and experts.[1] Technology forecasting is a major component of this.

It is, of course, rarely possible to predict the future reliably. The science of looking forward to the future relies therefore on *forecasting*: characterizing each of several outcomes, with some idea of the likelihood of each outcome and of what events and developments might lead to each coming about. The methods for predicting technological change or the pace of maturity have differed over time and include macrohistory, field anomaly relaxation (FAR), Delphi studies, scenario planning, and futures wheels.

Macrohistory is the sociologically informed analysis of long-term patterns of political, economic, and social change.[2] Macrohistory researchers identify trends in past history as a basis for extrapolation into the future. It enables a rigorous approach to analyzing social trends but is not commonly used as a basis for technological prediction. Macrohistory is often used as a source of data when combined with other methods of analysis.

In FAR, researchers identify a set of different topic areas, called "sectors."[3] They then use interviews with experts, literature surveys, and a range of other techniques to identify many different outcomes ("factors") over time in each of the sectors. They then use *relaxation*, identifying incompatible combinations of factors, to focus in on a small number of total outcomes ("scenarios"), which are often expressed as timelines. FAR can be remarkably effective but is heavyweight and time consuming.[4]

The Delphi method involves several rounds of interviews with experts until a consensus has been reached.[5] This is labor intensive, requiring multiple experts to spend their time speaking with researchers, but it provides some measure of what experts think. Delphi studies are widely used in situations where there is expertise but little concrete information available.[6] Such situations include policy and strategy creation in a wide range of settings, as well as forecasting. The Delphi method has two advantages over more direct methods of gathering expert opinion. It is anonymous, which encourages participants to be more honest and open in their feedback, and it is iterative, allowing for feedback among participants so the results are refined with each iteration.

Scenario planning takes a different approach. Scenario planning takes place within an organization that has a measure of influence on the outcomes. A team of researchers starts by identifying scenarios. Specifically, they identify "internal scenarios" that can be influenced by the organization, "external scenarios" that cannot, and "system scenarios" that combine elements of both.[7] The team then identifies key factors within the scenarios and explores "projections": how

combinations of such key factors may play out. The team then clusters such projections to identify new scenarios, each of which they explore for impact on the organization.

The futures wheel, by contrast, typically involves a single workshop with expert participants. It starts with an event, trend, or idea, and participants are asked to think about its consequences, intentional or otherwise.[8] Participants then repeat the process with the set of consequences, creating a list of consequences of the consequences, and so forth. These consequences are represented as a wheel infographic, with different sectors for different kinds of consequence and with first-, second-, and third-order consequences at increasing distances from the center. This provides an effective overview of the prediction space for discussion or planning.

Past forecasting success

A seminal prediction was published by Kahn and Wiener in 1967, *From The Year 2000: A Framework for Speculation on the Next Thirty-Three Years*.[9] Kahn and Wiener use a combination of macrohistory and an approach similar to FAR to make many forecasts about society in the 2000s. There are some laughably wrong predictions, such as the use of nuclear weapons for mining and an artificial moon to light up areas at night. However, many others have proved correct, especially those involving in software and communications.

Indeed, a 2002 review of Kahn and Wiener's predictions[10] judged that 81% of those within the theme of communications and computers had occurred by 2000. This theme's forecasts includes data processing, computers, networks, video, and additive manufacturing technologies such as 3D printing. For other themes the predictions were less accurate, with defense, materials, biotech/agriculture, and environment achieving 50%–40% accuracy. The least successful theme was aerospace, which primarily predicted interplanetary travel and habitation and achieved only 18% accuracy.

It is therefore reasonable to use forecasting methods to inform strategy around the development of digital technologies.

INTERLINKED STUDY METHOD

In this project we used a Delphi-style study to leverage expert advice in a
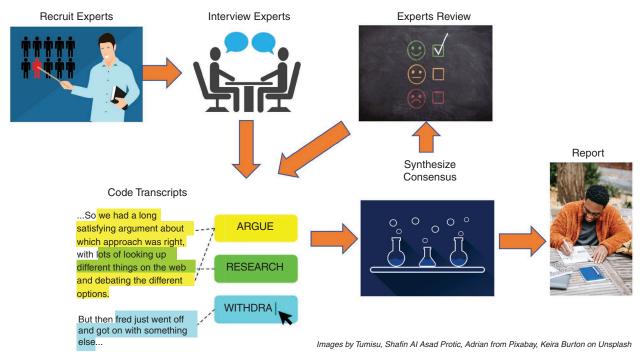

Images by Tumisu, Shafin Al Asad Protic, Adrian from Pixabay, Keira Burton on Unsplash

FIGURE 1. The Delphi process.

relatively short time. Figure 1 shows an overview of this approach.[6] The recruited experts give their best predictions by answering open questions. The researchers then use qualitative analysis techniques to identify themes in those predictions and synthesize a summary from the opinions. The experts then provide feedback on the summary; this process may continue for as long as there is convergence toward consensus.

Research questions

To provide focus for the study, we started with four research questions (RQs) on the theme of the interaction of novel digital technologies and interlinked computing.

RQ1. What novel computing and interlinking technologies are organizations anticipating incorporating in digital systems by 2040?

RQ2. What issues, particularly in terms of stability, security, privacy, safety, and environmental management, might be magnified, mitigated, or changed by the use of such technologies?

RQ3. What problems with responsibility and accountability (legal, moral, and energy-related) might arise as these technologies are incorporated?

RQ4. What approaches might offer future ways of addressing these issues?

Study design

We used a two-round Delphi study, as follows. We devised a set of open-ended questions around these four RQs, aiming to generate roughly an hour of discussion.

Specifically, we asked each expert which technologies they saw as the biggest game-changers being implemented now, what effect and impact each suggested technology might have, and how combinations of such technologies might play together. We then addressed interlinked systems directly, asking about types of linking and their implications in terms of responsibility, accounting, and control; the problems that might arise; and where one might look for solutions.

We piloted the questions in interviews with two people working in similar areas to the topic, to establish timing, comprehensibility, and completeness. Two researchers separately used thematic analysis,[11] "open coding" the pilot interview transcripts, and identified problems in the question set's effectiveness to answer the RQs. We amended the set of questions and the coding approaches accordingly.

In parallel with the above, we recruited a range of experts, such as futurologists, software industry journalists, and technology leaders who specialize in understanding future trends. We used a "snowballing" approach, starting with futurist personal contacts of the lead author. We carried out interviews of around one hour with each expert, recording and automatically transcribing each discussion.

Two researchers separately again "open coded" each interview transcript according to the agreed coding approach. Both coders then met and discussed different findings, to synthesize findings addressing the RQs. We then wrote up the results in a short editable document structured as a questionnaire around the finding statements.

We circulated the document to the experts and solicited feedback, in writing or as a further interview (according to preference). A researcher coded each response document or transcript, looking for changes to incorporate, areas of agreement, and areas of disagreement. The researchers selected five predictions to give a cross-section of the outcomes and explored the comments around them in more detail.

All interviews were carried out using Microsoft Teams, using that service's recording and automated transcription features. Coding was done using the NVivo tool (supported by the online transcript conversion tool Teams2Nvivo). Further analysis used Excel and Python on Jupyter notebooks.

INTERLINKED STUDY RESULTS

The process of recruiting interviewees took several months. We interviewed nine people in the first round, coded the results, extracted forecasts, and circulated the result as a questionnaire. Technical issues with the questionnaire led to us doing second face-to-face interviews with three of the nine. Three further experts expressed a subsequent interest in participating; we interviewed two of them face to face, and one completed a questionnaire.

Table 1 summarizes the expert interviewees. P1 and P2, the pilot study participants, are omitted. All but P3 and P6 have at least 30 years' experience in research or industry. The columns "Round 1" and "Round 2" indicate who participated in rounds 1 and 2, respectively; F indicates a face-to-face interview; and R indicates a response form.

The set of 36 forecasts from the round 1 analysis is available online.[12]

Areas of innovation

The interviewees identified a range of key areas of innovation: automation,

artificial intelligence (AI), the metaverse, quantum computing, and the Internet of Things. The research coincided with much media interest around generative AI, and all interviewees discussed implications of AI as the leading concern.

Several interviewees mentioned blockchain, but most dismissed it as a source of major change in future: "Blockchain has now proved its irrelevance (P11)." Indeed, no interviewees made forecasts specifically around blockchain technology.

Forecasts

Although the interview questions focused on the impact of interconnecting computers, in practice, many of the trends and issues identified were not focused on such interconnection.

From the 36 forecasts that the survey explored, we identified five that we believed were of general interest to digital technology practitioners and users. Each had at least some agreement from the 12 experts we consulted:

F1. In 2040, competition, both among states such as the United States and China and among big tech companies, will have led to corners being cut in the development of safe AI.

F2. Quantum computing will have limited impact by 2040.

F3. In 2040, there will be ownership of public web assets, and it will be identified and traded using technology such as tokenization.

F4. In 2040, it will be more difficult to distinguish truth from fiction because widely accessible AI can mass-generate doubtful content. AI will be a threat to objective truth and verification.

F5. In 2040, there will be less ability to distinguish accidents from criminal incidents due to the decentralized nature and complexity of systems.

We note that only F5 relates primarily to the topic of interlinked systems; of the rest, all but F2 anticipate system interlinking as a cause of their impact.

The following sections explore each forecast. Summaries and statements are backed up where appropriate by quotations from the interviews; these are shown in quotation marks, identifying the speaker (for example, P1).

F1: Cutting corners in safe AI

Unsurprisingly, given the media coverage of concerns about AI safety and related issues, this was a frequent forecast. Indeed, all the experts agreed on it. However, we noted that often the driver cited for corner cutting was competition among geopolitical superpowers rather than competition among multinational IT companies.

"The weaponization of code is what's going to drive economic opportunities ... for major nation states. Groups of states partnering together as well (P12). Legislation on this is bound to be retrospective, and in some cases, non-existent (P7)."

The possible scale of the problems anticipated by our interviewees is eye-watering. Most of the experts predicted exponential growth in AI over the next 15 years, leading to a large possibility of systemic problems.

ID	Description	Round 1	Round 2
TABLE 1. Expert interviewees.			
P3	Researcher into future of smart grids	F	R
P4	Futurist	F	R
P5	Marketing futurist and pioneer	F	R
P6	Smart building specialist	F	R
P7	Consulting futurist	F	R
P8	Consulting futurist in brand and retail	F	R
P9	Journalist specializing in smartphones	F	F
P10	Consultant futurist on financial retail	F	F
P11	CTO of large independent software vendor (ISV)	F	F
P12	CTO of security ISV		F
P13	CEO tech startup		F
P14	University professor and public strategist		R

F: face-to-face interview (blue); R: response form (amber).

Two experts (P4 and P10) considered a million-death incident attributable to the irresponsible deployment of AI systems reasonable in the 2040 time frame. Others disagreed, although several suggested many smaller, perhaps ten-thousand-death, incidents.

A frequently cited contributor to this trend was "regulatory capture," the difficulty of regulating politically powerful entities.

> "There's little evidence that the upcoming regulations have enough 'teeth' to tackle regulatory capture (P3)."

F2: Limited impact from quantum technology

One frequently mentioned promising item of new technology was rated as unlikely to have much impact in the 2040 time frame: quantum computing. It was seen generally as a long-term prospect; some saw it unlikely to have much impact in the near time frame and possibly in the longer term.

> "If [quantum computing is a game changer in 2040] it will be only just. I don't think it's predictable at this point. The quantum stuff feels to me very similar to the transputer and other new computing models. The problem was that for everyone who worked on them, by the time they got to commercial reality the existing models were past where they were at. And I suspect that's what's going to happen with quantum computing (P13)."

Several experts anticipated that early uses of quantum technology would be in supporting AI models. Some disagreed with the forecast's pessimism, expecting progress to be faster.

F3: Ownership of public web assets

What the media hype over blockchain and the success of blockchain-based currencies *have* achieved is a public understanding of the potential for cryptography "leading to decentralized finance and tokenization and things that go with that (P10)." The resulting forecast was for a new trend of digital assets, such as stakes in companies, being identified and traded using technology such as tokenization.

> "Everything that's data can be copied, but tokens can't. I don't send you a copy of a Bitcoin; I transfer the ownership of the Bitcoin from my wallet to your wallet. …

Larry Fink at BlackRock, the world's biggest asset manager, says tokens are going to be the new markets (P10)."

The introduction of such trading will involve considerable changes to social and legal frameworks:

> "There are going to be some big legal fights about [trading assets using tokens], but I expect the establishment to prevail (P7)."

A token, as an easily identifiable reference to an underlying publicly visible entity, can be traded fast and efficiently, whether or not blockchain is used. This is an attractive prospect for further financial automation.[13]

F4: Difficulty distinguishing truth from fiction

The ability of AI as a tool to mass-produce misinformation was a regular theme throughout the interviews. "This is already happening, such as Chinese information warfare in Taiwan, and I can't see how it will be solved by 2040 (P11)." Several experts mentioned the problem this poses for the Western democratic process:

> "We're not going to be living in a George Orwell world…. We're going to be living in a Philip K. Dick world [where] nobody knows what's true (P10)."

Some anticipated a future battle of truth and fiction among AIs: "…a sort of Cold War between AI trying to fake things and AI trying to detect them (P10)." Others shared concerns over conspiracy theories and their proliferation through AI:

> "There's a good chance that our Internet and our shared

> A TOKEN, AS AN EASILY IDENTIFIABLE REFERENCE TO AN UNDERLYING PUBLICLY VISIBLE ENTITY, CAN BE TRADED FAST AND EFFICIENTLY, WHETHER OR NOT BLOCKCHAIN IS USED.

knowledge base gets poisoned...we're prone to believe in conspiracy theories (P7)."

A related issue is that AI can be used as a topic to leverage political gain. The notion of using conspiracy theories about AI to gain political support from marginalized sectors of society is one example:

"[A populist politician] who will go and pick up the last 15% and try to rile them against the system...whatever crazy conspiracy theory he will throw them and run with it. Quite often it's about technology (P5)."

Thus, AI was seen as a threat to objective truth and verification, holding significant implications for democracy when considering its possible impacts on perception and judgement. Examples of issues include subversion of democratic decision making, poor individual health decisions in the COVID pandemic, and subversion of the academic review process itself.[14]

Some looked toward AI itself for possible solutions: "Improvements in AI itself might solve it (P4)." Others looked to sociotechnical solutions: "The way we determine what's true might be an interesting question in the future (P13)."

F5: Inability to distinguish accidents from criminal incidents

Systems are becoming more complex, and many will be beyond the capability of humans to understand. Potentially, that complexity may lead to more accidents; probably, it will make cybersecurity more difficult. This forecast, though, derives from the increasing complexity of systems leading to difficulty of coordination:

"Technical reliability between systems is difficult. At scale, it's very difficult. As soon as you bring in [multiple organizations] it becomes even more difficult, so the chances of having a cascading failure get much higher... [Finding a] root cause is a lot more difficult because there isn't the same degree of openness and information sharing (P11)."

In addition, the human skills required to analyze incidents are not widely available:

"It's a big problem in critical infrastructure. It's just people don't know it and root cause analysis is something that is quite early in maturity (P3)."

The result of this problem of determining the root cause will be difficulty distinguishing between accidents and the effects of a malicious actor,

TABLE 2. Agreement with statements.

	Forecast	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	P13	P14
F1	Cutting corners in safe AI	blue	blue	blue	blue	blue	blue	blue	blue	blue	blue	blue	blue
	... leading to a megadeath incident		blue	amber		amber	amber	blue	blue	amber	amber	blue	blue
	... due to regulatory capture	blue	blue	blue	blue	blue	blue	blue	blue	blue	blue		blue
	... which has developed exponentially	amber	blue	blue	blue	blue	blue	blue	blue	blue	blue		blue
F2	Quantum processing only just being used		blue	blue		blue	blue		blue	blue	blue	blue	amber
F3	Difficulty distinguishing truth from fiction	blue		blue	blue	blue	blue	blue	blue	blue	blue	blue	blue
F3	Tokenization for ownership of web assets	blue	blue	blue	blue	blue	blue	blue	blue	blue	blue		blue
F5	Cannot distinguish accidents from incidents	blue	blue	blue			blue	blue	amber		amber	blue	blue
	... systems beyond human understanding	blue		blue	amber	blue	blue	blue	blue		blue	blue	
	... with more accidents due to complexity	blue	blue	blue	amber	blue	blue	amber	blue		blue		blue

Blue: agree; white: no opinion or unsure; amber: disagree.

making it hard to learn from incidents and to develop mitigating strategies.

Disagreement on forecasts

As with any group of experts, our interviewees were by no means in agreement about the forecasts. Table 2 shows the statements from round 1 related to the above five forecasts, with the level of agreement of each expert. It highlights the range of opinions. We identified several sources of disagreement:

> time scale disagreements, where the expert agreed on the forecast but predicted a different time scale
> impact disagreements, where there was agreement on the forecast but not on whether it would have a major impact ("game changer")
> competing forecasts, where the expert expected a different, and incompatible, outcome.

For instance, while F2, quantum technology, shows a poor level of agreement, those who disagreed were either predicting an earlier time frame or suggesting low levels of usage.

SUGGESTED SOLUTIONS

In the interviews we asked the experts what solutions they could identify for the problems their forecasts outlined. Surprisingly, there were fewer suggestions than negative forecasts. Perhaps this was because the main role of a futurist is to forecast the future, not to change it. Perhaps it was because the scopes of most of the forecasts were too wide to address with single interventions, or perhaps the experts felt no need to mention what they considered obvious interventions.

Several experts suggested solutions for the two forecast problems related to AI (F1 and F3) and for the problem related to interconnected systems (F5). These suggestions were ambient accountability (S1), government policy (S2), new forms of education (S3), and the involvement of social science expertise (S4):

S1. Ambient accountability is a technical approach: the inclusion of checks such as safety audits and outcome validation in the code of the systems themselves. This includes the design of systems that lack the ability to support wrong outcomes—effectively, code that checks itself. Two experts, though, doubted the effectiveness of such an approach in an AI

world: "We've gone well beyond real time governance; that horse has truly bolted (P12)."

S2. Government policy can address the AI concerns in two ways. First is the establishment of government AI purchasing safety principles; second is legislation around AI safety. This carrot and stick approach appealed to many of the experts, although others thought it unlikely to have much impact. They pointed out that different countries will have different interests and competence, so the effect on a global market will be small. They also predicted "regulatory capture," the power of the major technology companies to influence regulation in their favor.

S3. New forms of education were suggested to address the complexity of the kinds of legislation recommended in S2. There was a particular emphasis on the provision of university courses combining technical skills and legislation.

S4. As a more general approach to addressing the problems, there was agreement that these

	Solution	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	P13	P14
TABLE 3. Proposed solutions.													
S1	Ambient accountability					🟦	🟧	🟦	🟦	🟦	🟧	🟦	
S2	AI safety legislation and government purchasing	🟦	🟦	🟦	🟦	🟦	🟦	🟧	🟦	🟦	🟦	🟦	🟦
S3	Courses with technical and legislation	🟦	🟦	🟦			🟦	🟦	🟦	🟦	🟦	🟦	🟦
S4	Expertise from social sciences	🟦	🟦	🟦	🟦		🟦	🟦	🟦	🟦	🟦	🟦	🟦

Blue: agree; white: no opinion or unsure; amber: disagree.

issues are only partially technical ones and that addressing them requires social science expertise.

Table 3 summarizes agreement around those approaches. To them, we can add two further recommendations implied directly from the forecasts:

S5. investment, both by governments and companies, in the development of responsible development and deployment methodologies for AI systems and other novel computing systems with an emphasis on existing best practice in software engineering

S6. investment by government and media in the development of a healthy ecosystem, representing a wide range of viewpoints, of information sources with high standards of fact checking.

Based on Delphi study interviews and surveys of 12 experts, we addressed four questions relating to the future of interlinked systems up to 2040, specifically as follows:

RQ1. What novel computing and interlinking technologies will be adopted?

RQ2. What issues will those technologies impact?

RQ3. What problems will arise related to responsibility and accountability?

RQ4. What approaches might offer ways of addressing these issues?

From the experts' responses, we extracted forecasts of five important trends toward 2040. These were as follows, where the proportion of experts expressing an opinion who agreed is shown in brackets:

F1. Cutting corners in AI safety will have led to major disasters (12/12).

F2. Quantum processing will only be starting to have an impact (3/7).

F3. Cryptographic tokens representing ownership of web assets will be traded widely (10/10).

F4. It will become difficult to distinguish truth from fiction (8/10).

F5. Organizations will become unable to distinguish accidents from adversarial incidents (5/7).

We observed that only F5 is caused directly by systems interlinking, although all relate to it as a topic.

To help address the human problems involved with predictions F1, F4, and F5, we identified six recommendations, four from the experts and two implied directly:

S1. use of "ambient accountability," where systems verify themselves (5/7)

S2. AI safety legislation and government purchasing rules (11/12)

S3. degree courses combining technical learning with legislation (10/10)

S4. expertise from social sciences applied to these technical problems (11/11)

S5. investment to create responsible development and deployment methods for AI systems

S6. government and media investment in a fact-checking ecosystem.

Investment and action on these six points is likely to yield strong benefits for software users and society over the next 15 years. ∎

REFERENCES

1. R. A. Slaughter, "Futures studies as an intellectual and applied discipline," *Amer. Behav. Scientist*, vol. 42, no. 3, pp. 372–385, 1998, doi: 10.1177/0002764298042003008.

2. R. Collins, *Macrohistory: Essays in Sociology of the Long Run*. Stanford, CA, USA: Stanford Univ. Press, 1999.

3. G. Coyle, "The nature and value of futures studies or do futures have a future?" *Futures*, vol. 29, no. 1, pp. 77–93, Feb. 1997, doi: 10.1016/S0016-3287(96)00067-5.

4. R. Rhyne, "Whole-pattern futures projection, using field anomaly relaxation," *Technol. Forecasting Social Change*, vol. 19, no. 4, pp. 331–360, 1981, doi: 10.1016/0040-1625(81)90005-6.

ABOUT THE AUTHORS

CHARLES WEIR is a research fellow at Lancaster University, LA1 4YW Lancaster, U.K. His research interests include developer-centered security, risk assessments, and leading-edge software development. Weir received a Ph.D. from Lancaster University. He is a member of the U.K. IET. Contact him at charles.weir@lancaster.ac.uk.

ANNA DYSON is a research associate at Lancaster University, LA1 4YW Lancaster, U.K. Her research interests include disruptive technologies, unmanned aerial vehicles, and contemporary conflict. Dyson is currently completing a Ph.D. in international relations from Lancaster University. Contact her at a.dyson@lancaster.ac.uk.

OLAMIDE JOGUNOLA is a lecturer in cybersecurity and smart grids at Manchester Metropolitan University, M1 5GF Manchester, U.K. Her research interests include energy transitions, local energy markets, and the cybersecurity of smart grids. Jogunola received a Ph.D. in electrical engineering from Manchester Metropolitan University. She is a Member of IEEE. Contact her at o.jogunola@mmu.ac.uk.

LOUISE DENNIS is a reader at the University of Manchester, M13 9PL Manchester, U.K. Her research interests include autonomous systems, AI transparency and ethics, and verification. Dennis received a Ph.D. from the University of Edinburgh. Contact her at louise.dennis@manchester.ac.uk.

KATIE PAXTON-FEAR is a lecturer in cybersecurity at Manchester Metropolitan University, M1 5GF Manchester, U.K. Her research interests include data science, AI, and cybersecurity. Paxton-Fear received a Ph.D. from Cranfield University. Contact her at k.paxton-fear@mmu.ac.uk.

5. D. Beiderbeck, N. Evans, N. Frevel, and S. L. Schmidt, "The impact of technology on the future of football – A global Delphi study," *Technol. Forecasting Social Change*, vol. 187, Feb. 2023, Art. no. 122186, doi: 10.1016/J.TECHFORE.2022.122186.

6. D. Beiderbeck, N. Frevel, H. A. von der Gracht, S. L. Schmidt, and V. M. Schweitzer, "Preparing, conducting, and analyzing Delphi surveys: Cross-disciplinary practices, new directions, and advancements," *MethodsX*, vol. 8, May 2021, Art. no. 101401, doi: 10.1016/J.MEX.2021.101401.

7. T. Witt, M. Dumeier, and J. Geldermann, "Combining scenario planning, energy system analysis, and multicriteria analysis to develop and evaluate energy scenarios," *J. Cleaner Prod.*, vol. 242, Jan. 2020, Art. no. 118414. doi: 10.1016/j.jclepro.2019.118414.

8. P. Daffara, "Applying the futures wheel and macrohistory to the COVID19 global pandemic," *J. Futures Stud.*, vol. 25, no. 2, pp. 35–48, 2020, doi: 10.6531/JFS.202012_25(2).0006.

9. H. Kahn and A. J. Wiener, "The next thirty-three years: A framework for speculation," *Daedalus*, vol. 96, no. 3, pp. 705–732, Summer 1967. [Online]. Available: https://www.jstor.org/stable/20027066

10. R. E. Albright, "What can past technology forecasts tell us about the future?" *Technol. Forecasting Social Change*, vol. 69, no. 5, pp. 443–464, Jun. 2002, doi: 10.1016/S0040-1625(02)00186-5.

11. V. Braun and V. Clarke, "Thematic analysis," in *APA Handbook of Research Methods in Psychology*, H. Cooper et al., Eds. Washington, DC, USA: American Psychological Association, 2012, pp. 57–71.

12. C. Weir, A. Dyson, L. Dennis, O. Jogunola, and K. Paxton-Fear, "Interlinked computing: Initial forecasts," 2023. [Online]. Available: https://doi.org/10.5281/zenodo.8187042

13. P. Laurent, T. Chollet, M. Burke, and T. Seers, "The tokenization of assets is disrupting the financial industry. Are you ready?" *Inside Mag.*, vol. 19, pp. 62–67, Oct. 2018. [Online]. Available: https://www2.deloitte.com/content/dam/Deloitte/lu/Documents/about-deloitte/Inside/lu-inside19-full.pdf

14. J. Bartlett, *The People vs Tech: How the Internet is Killing Democracy (and How We Save It)*. New York, NY, USA: Random House, 2018.